UNIVERSITY OF OXFORD
WADHAM COLLEGE

DOCTORAL THESIS

Choice and Auction Design in the Allocation of Food to Food Banks

Author:
Samuel Mitchell ALTMANN

Supervisor:
Professor Ian CRAWFORD

A thesis submitted in fulfillment of the requirements for the degree of Doctor of Philosophy

in the

Department of Economics

Approximate Length: *98,000 words*
September 10, 2024

Declaration of Authorship

I, Samuel Mitchell ALTMANN, declare that this thesis titled, "Choice and Auction Design in the Allocation of Food to Food Banks" and the work presented in it are my own. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Signed: Samuel Mitchell Altmann

Date: September 10, 2024

"Mostly Void, Partially Stars"

Cecil Gershwin Palmer

UNIVERSITY OF OXFORD

Abstract

Department of Economics

Doctor of Philosophy

Choice and Auction Design in the Allocation of Food to Food Banks

by Samuel Mitchell ALTMANN

Feeding America, an organisation responsible for feeding 130,000 Americans every day, distributes donated food among a network of participating food banks. Feeding America's allocation mechanism, the 'Choice System', uses repeated rounds of simultaneous first-price auctions to allow food banks to signal which types of food they need from Feeding America. This provides food banks a large degree of choice over the types of food they receive. This thesis examines the welfare and distributional consequences of enabling this choice.

I develop an empirical model of bidding in repeated rounds of simultaneous first-price auctions. I prove non-parametric identification of primitives in this dynamic multi-object auction model, and introduce a computationally feasible procedure to estimate this type of game. The difficulty arises because, when players cannot place package bids, bids do not uniquely identify bidders' valuations.

I then apply this model to Choice System bidding data, estimating the distribution of food banks' heterogeneous and time-varying needs. The central challenge is that I cannot observe food banks' inventories — the key determinant of bidding behaviour. Nonetheless, I prove that observations of food banks' winnings, which are just observed changes in their unobserved stocks, are sufficient for nonparametric identification. I propose a Bayesian estimator to estimate food banks' needs in the presence of the latent inventories. I then use these estimates to compare the Choice System to the previous allocation mechanism employed by Feeding America which gave food banks very limited choice. I estimate that the Choice System increased welfare by the equivalent of a 17.1% increase in the quantity of food being allocated, and that on average 85% of food banks are strictly better off from this change. I find that the majority of this welfare gain arises because the Choice System allocates food in batches, rather than sequentially.

Acknowledgements

I will forever be grateful for the supervision of Ian Crawford, for keeping my research on the right track and for generally enlightening me in the ways of the economist. I hope he can forgive me for the number of parametric assumptions I make in this thesis. Likewise, both Howard Smith and Alex Teytelboym deserve special mention for giving me so much of their time, and such valuable direction and support. Without either of them my work would look nothing like it does today.

I benefited from the guidance of the majority of Oxford's economics department, though I am particularly grateful for the support of Frank DiTraglia, Paul Klemperer, Meg Meyer, and Ludvig Sinander. Likewise, my fellow graduate students and friends who have made my time in Oxford so enjoyable and helped to alleviate the various stresses that come with doing a PhD. Particular thanks are due to Aidan Smith, for answering all my questions about mathematics and economic theory, and Luke Milsom, for answering all my questions about everything else.

The administrative teams at both Oxford and Wadham college have made my life as a PhD student that much easier, and I will always be grateful for their tireless work. Special thanks are also due to Alan Beggs, Martin Cook, Marie Milofsky, Sarah O'Keeffe, Judith Shapiro, Alex Voorhoeve, and Verity Watson, for their enormous contributions to my development as an economist.

My friends and family have been invaluable in supporting me through this journey; always keeping me grounded and reminding me that there are (some) things more important than economics. The biggest thanks go to Ellen Lees, my partner in everything, for embracing my life choices and reading more drafts of poorly written papers than anyone should ever be exposed to.

I have benefited from many conversations about Feeding America and food banking with Canice Prendergast, Mike Loeffl, Sarah Pennel, Dhanu Sherpa, and Ann Sheppard. The work in this thesis has substantially improved with the suggestions of Nikhil Agarwal, Robert Miller, Anh Nguyen, Neil Thakral, and Daniel Waldinger.

I am also grateful for the academic and financial support of the Economic and Social Research Council, The George Webb Medley fund, and the family of David Richard, for funding my graduate studies. Likewise, the many hours of support and computation time I received from the University of Oxford Advanced Research Computing (ARC) facility. Finally, a major thank you to Feeding America for giving me access to their data.

Contents

List of Figures

This page is intentionally left blank

Chapter 1

Introduction

1.1 The Question

Organisations are regularly faced with the problem of allocating scarce resources as efficiently and as equitably as possible. Governments must decide how to allocate contracts to contractors, local authorities must allocate school places to students, and hospital boards must allocate kidneys to transplant patients. Likewise, food relief organisations must decide how to allocate truckloads of donated food among their networks of regional food banks.

A fundamental question for these organisations concerns how much 'choice' to give agents. That is, to what extent agents should be able to signal their preferences to the central planner, and to what extent they should have control over their allocations. In this thesis I study the food allocation problem faced by food relief organisations. I focus on Feeding America, but other such organisations include Food bank Australia, Fareshare (U.K.), Banque Alimentaire, Tafel Deutschland, as well as all the other national members of The European Federation of Food Banks. Every day these organisations receive many donated truckloads, of various types of food, and must allocate these scarce and heterogenous loads among their food banks. Over the next one hundred pages I investigate the importance of giving food banks choice over the types of food they receive.

In the food bank setting, the importance of choice depends on the extent of heterogeneity in the 'match values' between food banks and truckloads of food. In turn, these match values depend on the substitutability of different types of food, as well as the heterogeneity in the types of food wanted by different food banks at different times. If food banks all needed the same types of food, so there was no heterogeneity, then a first-best allocation would be to just allocate food at random. In practice, allowing food banks to signal their heterogeneous needs might be welfare improving. Therefore, in order to assess the importance of choice and perform any sort of welfare analysis, it is important that we are able to accurately estimate the extent of this heterogeneity.

1.1.1 The Rise of Food Banking

One in ten Americans faces food insecurity, living without consistent access to enough food to lead a healthy and active lifestyle. Almost paradoxically, around 30% of all the food produced in the United States gets wasted.[1] This is a common pattern seen across the developed world: Around 6% of British people face food insecurity and 15% of the food produced is wasted, with similar numbers for continental Europe.[2] This disconnect highlights a missing market, which food banks attempt to fill. Food banks are charitable organisations that rescue food that would otherwise be wasted, and ensure it gets to those who need it.[3]

Food banking, as we know it today, began in the U.S. in the late 1960s, spreading to continental Europe in the early 1990s, and the U.K. in the early 2000s.[4] The rise in food banking coincided with a societal rediscovery of the existence of food poverty and insecurity, seeking to fill a gap that many had assumed would be filled by government welfare programmes (Riches, 2002). The increase in food banking is driven by both demand and supply factors. For example, the recent increase in the number of food banks in the U.K. has been attributed to the rise in food poverty caused by the Global Financial Crisis and a decade of government austerity (Loopstra et al., 2015). Meanwhile, tax incentives for corporations to donate surplus food to food banks in the North America lead to a similar increase in food banks, as there is now more food available for them to meet the demand (Leib et al., 2017, Kinach, Parizeau,

[1] These figures come from Feeding America's annual "Map The Meal Gap" report, available here.

[2] U.K. figures come from the Food Foundation, available here, while figures for Continental Europe are available through Eurostat.

[3] This food is never past use by date, and is often simply surplus food that a producer projected they were unlikely to sell in the first place. Food banks have a mandate to turn down food that is past its best. As well as food rescue and distribution, food banks perform numerous other important tasks including advocacy, offering training programmes, and access to funding schemes.

[4] See, for example, Wikipedia's discussion of the history of food banking or Williams and May (2022) on the recent rise of food banking in the U.K.

and Fraser, 2020). This global increase in food banking has both positives and negatives. On the positive side, people in need of food are now able to access it. But on the other hand, food banks only act as a sticking plaster and are not well placed to tackle the root causes of food insecurity. The need for food banks is arguably symptomatic of deeper societal problems; many have argued that the existence of food banks risks leading to political inaction on these root causes, leading to society becoming reliant on this charitable sector.[5]

There are a number of different models of food banking used around the world. Possibly the most common model is the 'warehouse model', where the food relief organisation operates a large central warehouse in each city or region. These central warehouses serve as the main point of food donations and storage, literally acting as a bank for food. These warehouses then distribute food among nearby charities, food pantries, and soup kitchens. This double hub and spoke system is effective as it makes use of scale efficiencies where storage is concerned. Other alternatives would require the partner charities to store food for extended lengths of time on their premises, which comes with additional costs. To the best of my knowledge, most systems have moved towards the warehouse model, owing to the benefits from centralisation.[6]

Food banks receive food donations from nearby farms, factories, grocery shops, and charity food drives, which I refer to as local donations. In many cases the food relief organisation, the central charity that coordinates the food bank network, also receives donations — this is the case for Feeding America, Food Bank Australia, Banque Alimentaire, among others. These are typically large donations - often truckloads of food coming from large corporate donors or farms. The food relief organisation then has to decide which, of their network of food banks, to send this food out to.[7] The choice of which food bank to send this food to is therefore of the utmost

[5]See, for example, discussions in Caplan (2016), Loopstra et al. (2015), Williams and May (2022), and Byrne and Just (2022).

[6]Other models of food banking involve the food pantries and soup kitchens themselves procuring donations in a more decentralised fashion. OzHarvest, and their related "XHarvest" organisations (such as UKHarvest), operate an "online" (in the Operations Research sense) model. Their fleet of vans are constantly roaming a city, picking up and dropping off food as it is needed, removing the need for a warehouse to store food in the first place.

[7]Many food relief organisations, particularly those that do not operate on the warehouse model, do not procure central donations in this way. Instead, they use their links with corporations and farms to increase the number of donations going directly to their member food banks. This can be thought of as similar to always sending these central donations to the food bank closest to the original donor.

importance.

1.1.2 Feeding America

Feeding America directly distributes around 100,000 tons of food each year across a nationwide network of 200 food banks. This is enough food to provide meals for around 130,000 people each day. These food banks, in turn, distribute around 500,000 tons of food each year, and it has been estimated that one in seven Americans make regular use of a Feeding America affiliated food bank (Gundersen et al., 2017). As a consequence, achieving an efficient and equitable allocation of food is a priority for Feeding America, to ensure that food banks can keep up with the ever-increasing demand for their services.

Like many food bank networks around the world, Feeding America previously allocated food using an allocation mechanism that allowed food banks very little choice in the food they received. Under this mechanism, referred to as the 'Old System', food banks would queue until they were offered a truckload of food, and return to the back of the queue regardless of whether they accepted or rejected this load. This mechanism was unpopular among food banks as they were rarely offered the types of food they needed. Efficient central planning is difficult because of unobserved heterogeneity in food banks' needs: Different food banks need different types of food at different times.[8] This heterogeneity arises because food banks in different parts of the country have access to different types of food from their local donors, and these types of food are liable to change over time. Feeding America's current allocation mechanism, the 'Choice System', consists of an auction market in which food banks are given an amount of virtual currency to bid on loads of donated food (Prendergast, 2017). This gives food banks a strong degree of control, and choice, over the food they receive.

Therefore, in this thesis I will investigate the magnitude of the welfare gain from Feeding America's transition from the Old System to their Choice System. We know from economic theory, including the welfare theorems, that choice will be beneficial.

[8] I use the term 'needs' to capture both what a food bank has a preference for, on behalf of their clients, as well as what they have room for in their warehouse. In this way, the term is intended to capture the determinants of a food bank's demand function, or their revealed preference from observed bids - a food bank with a warehouse full of cornflakes may still have positive marginal utility of additional cornflakes, but due to capacity constraints will not bid on additional cornflakes.

Likewise, Feeding America has found that the Choice System is extremely popular, relative to the Old System (reported in Prendergast (2017) and NPR (2015)). However, this alone cannot say anything about the size of these welfare benefits, nor whether the magnitude is large enough to justify investing in one of these auction platforms. Similarly, theory alone cannot tell us what is driving this welfare change — what it is about the Choice System that works so well. Therefore I will also investigate which features of the Choice System are driving these results. This will help highlight the key features of the mechanism we might take to other food bank networks or other settings. I will also investigate how these welfare benefits are distributed, as a common fear about such 'market mechanisms' is that they will increase inequality or inequity.

1.2 The Road Ahead

The structure of this thesis is as follows:

1.2.1 Chapter 2

Chapter 2 primarily serves to motivate my research question and methods. It serves as an introduction to the setting, the data, and various features of the data that are important for modelling purposes.

In Section 2.1 I give additional institutional details on Feeding America and their food allocation problem. In Section 2.2 I review the related literature on the food allocation problem, market design more broadly, and the econometric of auctions. In Section 2.3 I introduce the data sets used in this thesis and *briefly* explain how I clean and combined them. I then then go into detail describing the data in Section 2.4. Next, in Section 2.5 I present suggestive evidence on the extent of heterogeneity in food banks needs - both across food banks and within food banks over time. This is an important finding because the importance of choice directly depends on the extent of this heterogeneity — if food banks all had the same needs, which were invariant over time, then allocating food at random would be optimal.

In Section 2.6 I go on to discuss key features of the data, investigating important determinants of bidding behaviour. In particular, I highlight the existence of strategic bidding behaviour, that food banks treat food as a storable good (considering not just what to consume, but when to consume it), as well as non-additivities across different truck loads of food. These finding motivate my use of a structural auction model to analyse this problem. In Section 2.7 I discuss the positives and negatives of using a structural auction approach, as well as a number of other possible approaches. I settle on using a structural auction model despite the need to develop additional econometric tools to analyse this problem.

1.2.2 Chapter 3

Chapter 3 is primarily methodological, laying out the econometric framework that will be used to analyse the problem. In Section 3.2 I introduce an empirical model of bidding in repeated rounds of simultaneous first-price auctions. I then discuss identification and estimation of this dynamic multi-object auction model. The difficulties arise because simultaneous first-price auctions are not a direct revelation mechanism, so that bids do not uniquely determine a bidder's type. This means we cannot straightforwardly apply an inversion type argument to identify the model (Guerre, Perrigne, and Vuong, 2000), nor to estimate the model (Jofre-Bonet and Pesendorfer, 2003).

To identify the model we need some other type of observed variation. In Section 3.3 I show that observed variation in bidders' state variables is sufficient for identification. The idea is that variation in a bidders' state, such as variation in stocks, will cause variation in observed bidding behaviour. Under mild conditions, this reduced form relationship is sufficient to pin down the payoff function. Consequently, just like Jofre-Bonet and Pesendorfer (2003), I show that non-parametric identification of the model primitives follows from identification of the equilibrium bid distribution, conditional on state variables, and the state transition process.

In Section 3.4 I consider the estimation problem, which involves extending the procedure of Jofre-Bonet and Pesendorfer (2003) to the multi-object setting. In the single object setting, their procedure involves writing the value function as a function of the observed distribution of equilibrium bids *only*, a condition which fails in a multi-object environment (due to the lack of a Direct Revelation Mechanism). Instead I demonstrate that the value function can be written as a function of both the bid distributions and a correction term, which corrects for the non-additivities between lots allocated simultaneously. I then show how this correction term can be recovered from the estimates of a misspecified static model, allowing us to consistently estimate the model primitives. I then present an simulation study demonstrating the estimator's efficacy.

1.2.3 Chapter 4

Chapter 4 returns to Feeding America's food allocation problem, armed with the tools developed in the previous chapter. I employ this empirical framework to estimate food banks' needs from Choice System bidding data, and evaluate the welfare benefits of Feeding America's transition to the Choice System.

The framework presented in chapter 3 requires state variables are observed. Because I cannot observe food banks' inventories I extend the identification argument and estimation procedure to allow for unobserved states. To prove identification I expand on the framework presented in Hu and Shum (2012) to allow for multivariate latent states. I prove that the model is non-parametrically identified using observed variation in food banks' winnings. The intuition behind this argument is simple - winnings give observed changes to the unobserved states. When a food bank wins a truckload of a particular type of food, the change in their propensity to bid in the following periods allows us to pin down their payoff function. Likewise, the length of time until they return to their previous bidding behaviour allows us to pin down the distribution of food they receive from local donors. The idea is that the more of a particular type of food they receive from local donors, the longer it will take before they bid on that type of food after a win.

Extending the estimation procedure to allow for the unobserved stocks is relatively easy, and we can estimate the distribution of these local donations in the misspecified static estimation step. To overcome that the unobserved stocks will be correlated over time I employ a Gibbs Sampler to perform this estimation step - essentially iterating between drawing the unobserved stocks from their conditional posterior distribution, and estimating the misspecified static model, given a sample of stocks. I make parametric assumptions to facilitate estimation, assuming that demand curves are linear and that the unobserved donations are normally distributed.

Having estimated the model I use counterfactual simulations to compare the distribution of equilibrium allocations under the Choice System to allocations under the Old System. I find that on average welfare is 17.1% higher under the Choice System, and that on average 85% of food banks are strictly better off. I use additional simulations to tease out the most important features of the Choice System. I find that this welfare change is driven by Feeding America switching from a sequential allocation mechanism, to a simultaneous mechanism that allocates food in daily batches.

Chapter 2

Feeding America and the Food Allocation problem

2.1 Chapter Introduction

In this chapter I introduce the reader to Feeding America and detail their food allocation problem, paying particular attention to the role of 'choice'. I then introduce the data that I will use to analyse the problem, and present various descriptive statistics to give the reader a flavour of what the data looks like. I then present suggestive analysis, looking for evidence of heterogeneity in food banks' needs, suggestive of the importance of choice. Finally, I present additional exploratory analysis, investigating the key drivers of bidding behaviour, highlighting which features of the data are relevant for building an empirical model of the Choice System. I also use this chapter to review the relevant market design and auction literatures.

2.1.1 Feeding America

Feeding America, formerly America's Second Harvest, began in 1976 as a collection of food banks that would solicit donations from local grocery stores and farms. As additional food banks joined their network it became necessary to co-ordinate resource sharing. In 2001 the organisation merged with Foodchain, the then largest food rescue organisation, making it responsible for over 200 food banks.

Many of these food banks operate as standard food pantries — directly giving out food to those in need. However, the majority act as food distributors; themselves responsible for storing and sending out food to *hundreds* of local food pantries.[1] The food banks typically have fleets of trucks at their disposal, which they send out to pick up and drop off food as and where it is needed.

Figure 2.1 describes Feeding America's ecosystem. Corporate partners donate truckloads of food directly to Feeding America, who then decide which food bank to send this food out to. Food pantries, often churches or community centres, then come and take this food from their local food bank, before distributing it in their communities.[2] In this way, the food bank essentially acts as an intermediary, storing food before it is collected by a food pantry. In addition to receiving food from Feeding America, food banks also receive food from their local donors — nearby farms, factories, and grocery shops. In this way, food banks substitute food from Feeding America for food they do not have access to from their local donors. The average food bank receives around 75% of its food from local donors, however there is a lot of variation across food banks, with some relying on food from Feeding America much more than others.

2.1.2 The Old System

Under the Old System any truckload of food donated to Feeding America was offered to the head of a queue. The potential recipient had a few hours to accept or decline the load, before it was offered to the next food bank. This meant that each load could only be offered to around ten food banks before being returned to the donor. To discourage rejections, food banks would return to the back of the queue regardless of whether they accepted the loads. A food bank's relative position in the queue was determined jointly by whether they had recently been offered food, and their 'Goal Factor': A measure of the poverty in their local area relative to the national average. A higher Goal Factor implies more mouths to feed, so these food banks should be offered more food. Transportation costs were paid by the food banks, many of whom have fleets of trucks and lorries for this purpose.

[1] One point of clarification: Often, particularly in the U.K., the term "Food Bank" is used to describe what is known as a food pantry in most other countries.

[2] Different food banks give food out in different ways, however for the most part representatives from a food pantry go to the food bank, choose what they want, and load this into a van to take back to their pantry. Some food banks allocate slots for pantries to come collect food, others use on-line ordering systems. Studying some of these additional allocation mechanisms could be another fruitful area of research in future.

FIGURE 2.1: The Feeding America Ecosystem

The type of food offered in each truckload was essentially random, so that on average food banks received the same quantities of food per mouth. This would have been optimal if food banks all had the same preferences and capacities. In reality, different food banks needed different types of food at different times. Food banks use food from Feeding America to substitute for food they do not receive from their local donors. A food bank surrounded by farms is likely to have a weaker preference for fresh produce than a food bank in a city. Feeding America wanted to improve welfare by taking account of differing needs. They decided to use a market mechanism to give food banks control over the allocation they receive.

2.1.3 The Choice System

The Choice System consists of simultaneous first-price sealed-bid auctions. Two rounds of auctions occur each day, five days a week, with around 30 lots auctioned each day. Bidders observe the previous winning bids for a particular type of food, making it easier for food banks to know how to bid. Outcomes of auctions that occur simultaneously are independent, and bidders cannot place combination bids. Winners generally pay to transport their winnings.

Food banks bid with a virtual currency called 'shares'. Other than storage and transportation costs, the only opportunity cost a food bank faces when bidding is that they will have fewer shares to bid on other lots. Feeding America can control which food banks have the most shares, ensuring that food banks with larger Goal Factors are allocated more shares and, consequently, receive more food (in the spirit of the Second Welfare Theorem). All spent shares are redistributed each night.[3] Food banks can save shares from one day to the next. Those with less than the median allocation of shares have access to interest-free credit, so that food banks can smooth their consumption over time. The money supply is set to ensure that prices remain constant (on average) over time, reacting to changes in the supply of food.

Food banks can bid negative amounts, down to a reserve price of -2000 shares. This incentivises food banks to accept undesirable loads, helping Feeding America maintain good relations with their donors by ensuring that every lot is graciously accepted. This had been a problem under the Old System - donors whose donations are refused are less likely to donate in future.[4] As all lots are eventually sold, donors now feel like their donations are always graciously accepted, and so they continue to donate. On average 21% of lots are sold at strictly negative prices, and 10% are sold at the reservation price.[5] Negative prices occur because food banks are capacity constrained. The marginal value of an additional load of food to a food bank with an already full warehouse is negative. The extra load will likely spoil and have to be thrown away, which creates a bad image.

The introduction of a market mechanism had the potential to disadvantage smaller food banks. Credit use, joint bidding and fairness committee mechanisms were introduced to alleviate this risk. Smaller food banks often choose to bid jointly, because they might not need a whole truckload of a food. A small number of food banks bid jointly for more than half their winning bids. Otherwise, food banks rarely place joint bids. For this reason I generally ignore the decision to bid jointly. Discussion

[3]The redistribution creates a small positive externality. For every share spent, an individual food bank will receive around 1/210 of that share, which is negligible.

[4]This aspect of the Choice System contributed to the supply of donations to Feeding America increasing drastically since the introduction of the mechanism. Feeding America itself would often turn down donations under the Old System, fearing that no food bank would accept the load.

[5]While 22% of lots are not sold right away, most are sold the following day. Lots not sold right away are predominantly either multiple loads of fresh produce or large bottles of water. The numbers are skewed by 130 loads of 8 litre bottles of water that were sold over several months.

of how I consider joint bidding is given in Appendix A.1.3. Feeding America also employ a fairness committee to enable food banks to raise any problems they have with the Choice System. So far there have been no complaints, and food banks have universally reported great satisfaction with the Choice System.

When multiple homogenous lots are auctioned simultaneously lots are allocated to the top bidders (who pay their bid) until the lots have been exhausted. These auctions resemble discriminatory first-price auctions, rather than simultaneous auctions. 7% of auctions fall into this category. The main text ignores these auctions, while estimation and analysis does not. Details of how the model and estimation procedure are extended to account for these auctions is given in Appendix H.

Feeding America allows food banks to sell the food they receive from local donors, making up 4.5% of lots. There are several distortions in this submarket: For equity reasons Feeding America taxes and redistributes shares earned, reducing incentives for foodbanks to sell. Similarly, foodbanks have always happily shared excess food with one another for free and selling one's excess is looked down on by the foodbanks.[6] In this paper I generally ignore food banks' decisions to sell food. Selling food is rare, particularly for the food banks most reliant on the Choice System. Incorporating the decision to sell adds too much complexity to the analysis.

2.2 The Literature

In this section I discuss the related literature. I begin by discussing previous work on the food allocation problem, on which there is very little prior literature. In 2.2.2 I discuss the empirical literature on centralised allocation and market design, focusing on settings when the planner cannot use (real) money to clear the market. Finally, given the direction this thesis will take over the few chapters, in 2.2.3 I review the empirical auction literature.

[6] Reported in Planet Money, (NPR, 2015).

2.2.1 The Food Allocation Problem

Prendergast (2017) and Prendergast (2022) also study Feeding America's transition to the Choice System. Prendergast (2017) is predominantly descriptive in multiple ways. After outlining the details of Feeding America's Choice System (much of which is directly repeated in this thesis), Prendergast then presents descriptive analysis of the Choice System data. The key 'causal' contribution of this paper is to recognise that under the Old System food banks would, more or less, all receive the same food. Whereas under the Choice System, food banks could choose their allocations. Therefore, the degree to which food banks sort into consuming different types of food under the Choice System conveys information about the relative inefficiency of the Old System. If, under the Choice System, food banks still all consume the same food, then the Old System was not particularly inefficient. Prendergast (2017) then reports that food banks broadly sort into those who predominantly consume high quality food (such as peanut butter), and those who consume lower quality food (such as produce) in large quantities.

Prendergast (2022) goes into additional detail describing the data. They present additional descriptive statistics demonstrating heterogeneity in winning and bidding patterns across food banks, all of which emphasise that welfare is likely to be greater under the Choice System than the Old System. They also considers a simple welfare exercise, using a sufficient statistics approach to compute welfare under the Choice System compared to under the Old System. Assuming pure price taking and aggregating consumption at two month intervals across 21 categories of food, they calibrate a quadratic utility function with the results from the prior literature on food consumption. Under this model, the observed elasticity of demand can be used to compute a sufficient statistic for food bank welfare. They find that the observed dispersion in consumption across food types is associated with a 21% increase in the value of distributed food.

There are several important differences between this thesis and the previous work by Prendergast. First, they use data from 2005-2011, right after the Choice System was introduced, whereas my data (2014-2017), comes from a time when one might expect the market to be in a steady state. Likewise, my structural approach is complementary to their descriptive and sufficient statistic approaches, allowing me to perform detailed welfare analysis and consider additional counterfactuals of interest. This thesis employs richer data that is disaggregated at the auction level and includes information on losing bids. Instead of aggregating consumption over time, I can study the exact timing of food banks' consumption as well as their losing bids.[7] Consequently, I gain a more detailed understanding of how food banks make inter-temporal substitutions. This allows me to simulate alternative dynamic allocation mechanisms, for example investigating the benefits of batch versus sequential allocation. Because my approach does not require consumption aggregation over time I do not make assumptions about the correct period for aggregation, which is complicated by how drastically food banks differ in both the scale and scope of their operations.

While there is some literature studying food banking more broadly, there is little work examining the food allocation problem.[8] As noted by Byrne and Just (2022), research on food banking is still in a nascient stage, particularly in countries where widespread food bank use is still a recent phenomenon. Given the rise in food bank use worldwide, we can expect increasing interest in this area in future. There is increasing interest in fair-division and "many-to-one" matching markets (which characterises the food problem) from the Operations Research Literature, and several algorithms have been proposed for solving such problems. The only study which explicitly handles the food allocation problem is Walsh (2015), who tackle the problem faced by Food Bank Local in Australia. They develop an online-algorithm (the "Like" mechanism) for sequentially allocating food among food banks. This is one of the allocation mechanisms I consider as a counterfactual.

2.2.2 Market Design

Empirical market design is a growing literature analysing preferences and allocations in centralised assignment markets, often employing techniques from Empirical

[7] Their attempted consumption, if you will.

[8] There is an extensive literature seeking to understand the causes of the rise in food banking (Riches, 2002), the geography of food banking (Cloke, May, and Williams, 2017), who works in food banks (Williams et al., 2016), and an extensive literature on the effectiveness of food banks at providing nutrition to their clients (Bazerghi, McKay, and Dunn, 2016). The "How" of food banking remains an under-researched topic.

Industrial Organisation. Unlike the much of the theoretical market design literature, the empirical literature typically examines settings in which the planner's choice of mechanism depends on features of agents preferences. For example, in the food bank setting, the importance of choice depends directly on the existence and the extent of heterogeneity.

There is an extensive literature empirically analysing one-to-one matching markets, studying, for example, the allocation of doctors to hospitals (Agarwal, 2015), the allocation of students to schools (Abdulkadiroğlu, Agarwal, and Pathak, 2017), and the allocation of teachers to schools (Combe, Tercieux, and Terrier, 2022). See Agarwal and Budish (2021) for a broad overview of the empirical market design literature, Agarwal and Somaini (2020) for a detailed overview of the empirical school choice literature, and Chiappori and Salanié (2016) for a detailed overview of the econometrics of matching markets. However this literature typically employs static models, since students are only assigned a school once. The food allocation problem is both multi-object (many food banks are allocated many loads of food) and dynamic (food must be allocated repeatedly over time).

On the multi-object side Budish and Cantillon (2012) study the course allocation problem, which also uses a system of virtual currency to allocate MBA courses to students. Similarly, Fox and Bajari (2013) use methods from the stable matching literature to measure the efficiency of the 1994 US spectrum auction. Preference heterogeneity is an important theme in these papers, even though their data is primarily cross-sectional.

The empirical dynamic assignment literature often consists of evaluating waiting list design. Agarwal, Hodgson, and Somaini (2020) and Agarwal et al. (2021) study the mechanisms used to offer deceased donor kidneys to transplant patients. Likewise Waldinger (2021) studies the allocation of public housing. Similar to this paper, they assess the value of giving agents choice over their allocations, considering the trade-offs between efficiency and other concerns of policy makers. Other work analysing dynamic multi-object allocation problems include Verdier and Reeling (2022) on hunting licenses, Gandhi (2019) on nursing homes, Liu, Wan, and Yang (2019) on peer-to-peer ride sharing, and Robinson Cortés (2020) on foster-care placement. This literature highlights the importance of heterogeneity in preferences and match values, but typically do not consider the role of heterogeneity over time, which is an important factor in the food allocation problem.

Dynamic Games

This literature regularly borrows from the literature on the estimation of dynamic discrete choice models and dynamic games. Consequently, this thesis does the same, albeit most often through the an empirical auction lens, which I discuss shortly.

The Estimation Procedure I introduce in Chapter 3 is a generalisation of Jofre-Bonet and Pesendorfer (2003)'s procedure for estimating dynamic auction games, but in a discrete choice setting, collapses down to the Conditional Choice Probability methods of Hotz and Miller (1993) and Pesendorfer and Schmidt-Dengler (2008). The difference only arises when the agent is faced with a choice over lotteries.

This thesis also relates to the literature on the identification of dynamic models in the presence of unobserved states, drawing from Kasahara and Shimotsu (2009) and Connault (2014) among others. My identification argument expands on Hu and Shum (2012), introducing an additional (potentially endogenous) variable that acts as an observed shifter, or instrument, for the unobserved state. This has an advantage over their framework in reducing the restrictiveness of the assumptions needed for identification, as well as yielding this very intuitive identification argument. I also expand their framework to allow for multivariate latent states, an important contribution in itself. My non-parametric identification proof relies heavily on their framework, as I use a similar argument based on the spectral decomposition of linear operators as has also been used extensively in the measurement error literature, including in Hu and Schennach (2008) and Carroll, Chen, and Hu (2010). My identification argument is also similar to that of Berry and Compiani (2020), using an instrument (observed winnings) to identify changes in the unobserved state (food banks' stocks). Whereas they require the unobserved state is independent of their instrument (an exclusion restriction), I require dependence between my instrument and the unobserved state as instrument relevance.

Unobserved stocks are also a key feature of the inventory models of Hendel and Nevo (2006) and Erdem, Imai, and Keane (2003) among others. The key distinction between my model and these is that agents also receive food from an external source

- local donors. I do not need to assume that agents solve the optimal control problem for their consumption of these stocks, and can instead treat this as an identifiable process determined by their client food pantries and local donors.

2.2.3 Empirical Auctions

To the best of my knowledge, I am the first to consider the identification or estimation of dynamic multi-object auction games. However, I build strongly on a number of literatures, most notably the literature on the estimation and identification of dynamic auction games, and multi-object auction games.

The state-of-the-art approach to using structural econometrics to empirically analyse auctions comes from Guerre, Perrigne, and Vuong (2000), who established how to invert the bidders' first-order-conditions to write the underlying values as a function of the observed distribution of bids. This approach has been used to analyse, for example, timber auctions (Athey, Levin, and Seira, 2011), government procurement (Li and Zheng, 2009), allocating mineral right (Li, Perrigne, and Vuong, 2000), among other things. It has also been extended to various other settings, including the affiliated private value framework (Li, Perrigne, and Vuong, 2002), the interdependent value framework (Somaini, 2020), to allow for unobserved auction level heterogeneity (Krasnokutskaya, 2011), the multi-unit setting (Hortaçsu and McAdams, 2010), and a setting with dynamic linkages (Jofre-Bonet and Pesendorfer, 2003). See Athey and Haile (2007) for a comprehensive review.

Multi-Unit and Multi-Object Auctions

There is a large literature empirically modelling multi-unit auctions. The distinction between Multi-Unit and the Multi-Object auctions studied in this thesis concerns whether lots are divisible. Hortaçsu and McAdams (2010) was the first to extend Guerre, Perrigne, and Vuong (2000)'s identification and estimation approach to multi-unit auctions. The central idea here is to recognise how rather than simply inverting an individual bidding function, we can invert the whole demand schedule to recover an underlying cost function. This has since been applied to analyse, among other things, financial securities auctions (Kang and Puller, 2008) and electricity auctions (Reguant, 2014). See Hortaçsu and McAdams (2018) for a comprehensive review.[9]

On the Multi-Object side, Cantillon and Pesendorfer (2007) was the first paper to consider identification and estimation of simultaneous first-price auctions. They focus on first-price auctions in which bidders submit combination bids over London bus routes. Similarly, Kim, Olivares, and Weintraub (2014) study combinatorial first-price auctions in the allocation of contracts for Chilean school meals. In both cases, observations of combination bids are necessary to identify complementarities between lots. In contrast, Gentry, Komarova, and Schiraldi (2023)'s key contribution is to consider simultaneous first-price auctions *without* combination bidding. They showed that variation in the characteristics of rival bidders could be used similarly to an instrumental variable to identify the bidders' combination values. However, as I will demonstrate, these exclusion restrictions fail in a dynamic environment. This is because bidders' forward looking behaviour ensures that their continuation values, and hence bidding behaviour, are generally affected by every state variable. Instead, as I will show, these exclusion restrictions are not necessary for identification. As described above, Fox and Bajari (2013) also study an auction environment without combination bidding. In their application the equilibrium allocation is shown to be 'stable', enabling the identification of preferences and partly enabling them to get around the dimensionality problem. However, this stability property cannot be applied in general.

Dynamic Auctions

Jofre-Bonet and Pesendorfer (2003) was the first paper to empirically analyse dynamic auction games, focusing on repeated first-price auctions. They analyse highway contracts, finding backlog effects to be determinants of future bidding behaviour. A number of papers have built on this framework, for example Jeziorski and Krasnokutskaya (2016) who study dynamic auctions with subcontracting. Groeger (2014)

[9] The identification framework and estimation procedure presented in Sections 3.3 and 3.4 respectively can be extended with relative ease to Multi-unit auctions. This is left for future work.

study participation in first-price auctions in a dynamic setting. They assume (justified by the data) that while entry decisions are forward looking, bidding behaviour is myopic. Balat (2013) generalise participation in the dynamic setting, additionally allowing for unobserved heterogeneity in lot quality. Finally, Raisingh (2021) studies the effect of pre-announcements of auctions in the MDOT data, using a dynamic model of bidding and entry behaviour. These papers generally study data in which multiple auctions are held simultaneously - assuming that pay-offs are additively separable across auctions. I build on this literature by relaxing the assumption that pay-offs are additively separable.

Kong (2021) study an intermediate case in which gas and oil leases for a given tract are auctioned sequentially, finding evidence of both complementarities and affiliaiton in the values of sequentially auctioned leases. There is also an extensive literature on dynamic second price auctions, often studying on-line market places such as eBay. Sailer (2006) consider a model of identical objects, presenting various non-parametric identification results. Backus and Lewis (2016) extend this framework to allow objects to be heterogeneous and allow for a 'random coefficient demand' style model. Their framework has since been applied on a number of occasions, for example in Hendricks and Sorensen (2015) and Bodoh-Creed, Boehnke, and Hickman (2021) who study the efficiency of on-line marketplaces. Finally, Balat et al. (2015) use a dynamic multi-object auction model to analyse how dynamic strategising affects bidder's behaviour in electricity markets. The model they present is used to derive several hypotheses, which they test using reduced form analyses.

2.3 The Data

Three sources of data are used in this thesis. The main data is the Choice System dataset, which is not publicly available and was received directly from Feeding America. I also make use of an auxiliary data-set enabling the identification of the locations of 85% of the food banks. Lastly, I use information from Feeding America's on-line food poverty tracker tool to estimate Food banks' goal factors.[10] Detailed discussion of how I clean and categorise the data are included in Appendix A.

[10]This tool is accessible `https://map.feedingamerica.org`

2.3.1 Bidding Data

The Choice System dataset contains information on $26,617$ individual auctions run over the course of 44 months from January 2014, covering 165 food banks. The data included both winning and losing bids from each food bank, as well as information on the food composition and location of each lot.[11]

The sheer volume of types of food being auctioned makes categorisation necessary. I split food into 15 categories, largely the same categories used in Prendergast (2017). To capture different types food being imperfectly substitutable I further split food into 164 subcategories.[12] To capture storage costs I categorise food into five storage types: Dried, Tinned/Bottled, Refrigerated, Fresh, and Non-Food. Fresh food includes produce and baked goods, that generally have limited shelf-life. Refrigerated includes anything that needs to be stored in a fridge or freezer, such as meat and dairy. Tinned and Bottled food includes anything with a long shelf-life that is tinned or bottled, ranging from baked beans to bottled water. Dry food captures long shelf-life food such as cereal, pasta, or cookies. Non-Food includes anything not considered food, including cleaning and beauty products. Many loads contain multiple types of food. I allow lots to contain up to four different items which I assume evenly make up the load, unless explicitly stated otherwise. See Figures 2.3 and 2.4 below for the approximate composition of allocated food by Category and Subcategory respectively.

2.3.2 Auxiliary Data

Food banks in the main Choice System data were anonymised. Using data from Feeding America's Food Bank Locator tool[13] I identified the locations for 85% of food banks, who together consumed just over 98% of all food on the Choice System.

[11]Importantly, I do not observe whether any given auction happened in the morning or afternoon. I assume that all auctions in a day happen at the same time. This presents a potential weakness of this analysis, however anecdotal evidence suggests that most food banks bid in only one auction round each day. This was suggested by Canice Prendergast, one of the original designers of the Choice System. If food banks are optimally choosing not to bid on any auction in a given round then the inaccuracy of my results will be minor.

[12]See Appendix A for additional discussion of how food was categorised.

[13]Accessible at `https://www.feedingamerica.org/find-your-local-foodbank`

Figure 2.6 shows the approximate locations of foodbanks (black spots) and the origins of lots coming to auction, by storage type. I imputed the distance between each food bank and each lot using the distance "as the crow flies".

I did not receive access to recent Goal Factor figures. However, this data can be constructed using the locations of food banks, formulae given in (Prendergast, 2022), and information on local poverty and food insecurity rates from Feeding America's 'Hunger in America' on-line resource. Under the Old System a food bank's Goal Factor was given by:

$$GF_i^{OS} = \frac{\text{Population}_i}{\text{Population}_{US}} + \frac{\text{Poverty}_i}{\text{Poverty}_{US}} \tag{2.1}$$

Where Population_i refers to the number of people living in food bank i's catchment area, and Poverty_i refers to the number of people living below the poverty line in food bank i's catchment area. Food bank catchment areas, defined at the County level, are given in Feeding America's 'Hunger in America' on-line resource. Populations figures were then taken from the 2015 US census Small Area Income and Poverty Estimates (SAIPE). Poverty rates, by county, are given in an additional dataset received from Prendergast, in turn received from Feeding America. Presumably these were the figures used to construct the Goal Factors to begin with.

Under the Choice System, the Goal Factor formula was updated to reflect that even individuals above the poverty line often use food banks. The new formula includes $\text{Poverty}_i'$, the number of people between the poverty line and 185% of the poverty line, as well as $\text{Population}_i'$, the number of people above 185% of the poverty line. These figures were included in the dataset I received. These figures are weighted according to empirical usage weights. The updated formula is given by:

$$GF_i^{CS} = \frac{0.73\,\text{Poverty}_i + 0.22\,\text{Poverty}_i' + 0.05\,\text{Population}_i'}{0.73\,\text{Poverty}_{US} + 0.22\,\text{Poverty}_{US}' + 0.05\,\text{Population}_{US}'} \tag{2.2}$$

2.4 Descriptive Analysis

2.4.1 The Food

Figure 2.2 presents descriptive statistics on the lots being allocated, split by storage method. Several things are evident. First, that many lots are allocated simultaneously. On average around 30 truckloads of food are allocated each day across the two rounds of auctions. Fresh produce is by far the most abundant type of food, however part way through my sample the supply of produce drops to near zero. This is because Feeding America started allocated produce off platform instead.[14] Other than this, there are no other changes to supply occurring over the sample period. The composition of food by category are demonstrated in Figure 2.3. Likewise, the composition of food by subcategory is shown in Figure 2.4, displaying more common subcategories in proportionally larger font size.

A second key feature of the data is that lots come in very variable sizes. Produce is most often shipped in 20 ton truckloads, while other loads have averages between 10 and 15 tons, with standard deviations of 5 tons. This is important because it emphasises that even just descriptively and ignoring their different compositions, lots appear to be highly heterogenous.

The final feature that is clear from this figure is that only a small number of bidders bid on any given lot, with an average of around 3 bidders per lot, with a standard deviation of 3. Likewise, we see that a large proportion of lots sell for negative prices - particularly Fresh produce and high volume/low quality beverages (included in the Tinned storage type). This suggests low demand for these types of food.[15]

[14]This was most likely because these auctions were extremely uncompetitive (attracting on average only one bid per allocated lot, with only 71% of lots being allocated), with extremely deterministic demand. That is, a small set of food banks would consistently buy the bulk of the produce. When demand is so easily predictable, even the small costs associated with listing the food on the platform can make off-platform allocation worthwhile.

[15]Figure 2.2 shows that very few bidders are observed bidding in any given auction. This may suggest auctions are uncompetitive. In practice it is unlikely that food banks collude, given how this harms non-colluding food banks and that most food bank managers are extremely prosocial.

FIGURE 2.2: Descriptive Statistics, across lots

	Dried	Tinned	Fridge	Fresh	Non-food	Mixed	Total
Daily lots							
(mean)	9.19	5.3	4.32	10.56	2.52	4.6	29.36
(std)	5.95	3.81	2.92	5.79	2.06	3.22	13.32
Pounds per lot							
(mean, 000s)	22.5	34.3	28.3	40.1	20.4	27	28.8
(std, 000s)	9.7	8.3	10.1	3.9	12.2	10.6	11.3
Winning bid							
(mean)	2106	1085	2704	211	2967	2481	1802
(std)	5329	6414	6331	779	6436	5176	5375
No. bidders							
(mean)	2.95	2.7	2.54	1.22	3	2.78	2.59
(std)	3.14	3.5	3.17	0.64	3.26	3.06	3.04
% Allocated	93	83	91	71	91	96	88
% Negative prices	35	47	28	19	28	27	32

Note: Excludes multiple homogeneous loads. Mixed loads are presented as a separate type for this figure only. Winning bids includes the reservation price when no bids are received. 'Allocated immediately' refers to the percentage of lots that receive at least one bid above the reservation price. Negative prices include loads allocated for 0 shares.

2.4.2 The Food Banks

Figure 2.5 summarises relevant demographic information and bidding behaviour of food banks. The statistics are calculated by food bank, then quantiles are evaluated across food banks. There are three main takeaways from this figure. First, examining the population and poverty figures, we see that the scale of food banks' operations differs vastly. The 75th percentile food bank caters to around 5 times as many people as the 25th percentile food bank. This is partly driven by rural vs urban differences, as well as variation in population and poverty density across the country.

Similarly, food banks also differ vastly in terms of their broad behaviour on the platform, with some food banks bidding more often and winning much more often than others. A factor not clear from the marginal distributions is that there is little correlation between how often a food bank bids, and how much they bid. Some food banks bid regularly, bid low, yet win often. These food banks bid on the types of food that are less popular to other food banks, likely in areas with a lot of poverty, but few local donors. Other food banks bid rarely, but bid aggressively when they do. These food banks are likely those who receive a lot of food from local donors, so do not need to rely on Feeding America for their staples, and can instead focus on

FIGURE 2.3: Composition of food allocated, by Category

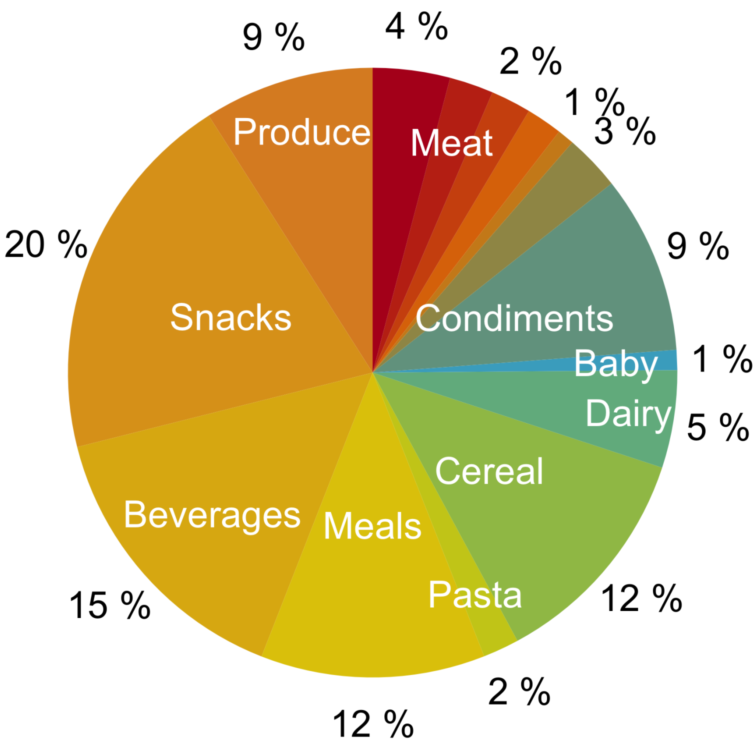

FIGURE 2.4: Composition of food allocated, by Subategory

trying to win just the highest quality food.

The final thing to note is that these statistics present a large rightward skew, with a small number of food banks winning the bulk of the food, and with a small number of food banks having the bulk of the mouths to feed (these are not necessarily the same food banks). Therefore, much of my analysis over the coming chapters will be driven by these smaller number of food banks for whom Feeding America is perhaps the main source of their food. This will be something to bare in mind in the next chapters.

Figure 2.6 plots the locations of lots allocated by Feeding America, as well as the approximate locations of food banks. The majority of food banks are in the eastern half of the country, as are most of the lots, with a reasonably large number of food banks and lots located in Texas and California. Food banks do not only bid on lots close to them, even though there is a strong positive correlation. The average lot still travels 500 miles to reach a food bank.

FIGURE 2.5: Descriptive Statistics, across food banks

	Mean	p10	p25	p50	p75	p90
Population (000s)	1913	384	676	1270	2543	4385
Poverty (000s)	284	64	99	191	373	645
Goal Factor	1	0.16	0.36	0.62	1.19	2.46
Bids Placed	380	13	44	166	442	844
Average Bid	3601	546	1202	2509	4067	6903
Lots Won	159	6	26	70	177	351
Average Payment	3803	485	1060	2476	4507	7414

Note: Statistics are calculated by food bank, then quantiles are evaluated across food banks. The mean Goal Factor is normalised to 1. Population and Poverty figures refer to the number of people in a food bank's catchment area.

2.5 Suggestive Evidence

Under the Old System every food bank was, ex ante, offered the same allocation. If food banks have heterogeneous needs, unknown to the social planner, welfare might be increased by allowing food banks greater choice in their allocations. For example, suppose food bank A is in a city and does not receive much fresh produce from its local donors, while food bank B in the countryside receives a lot of produce from nearby farms. Allowing B to choose less produce so that A can choose more

FIGURE 2.6: Locations of lots and food banks

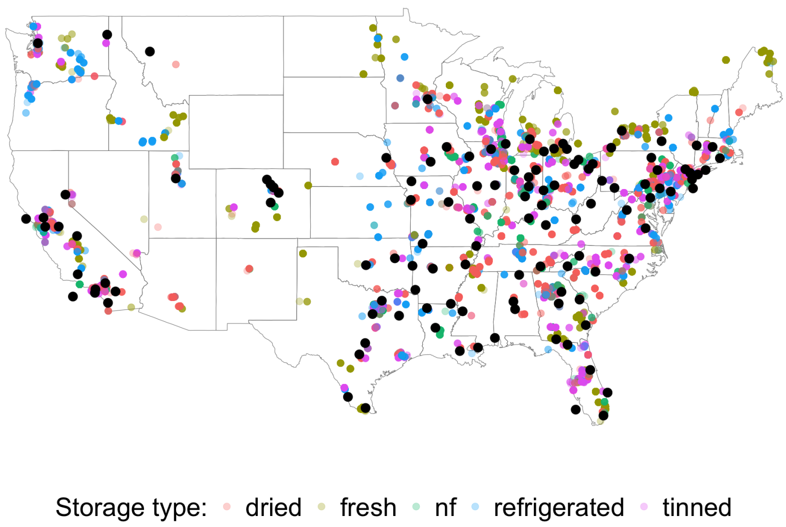

Note: Black spots give the approximate locations of food banks, jittered by an average of 200 miles to maintain anonymity. Excludes loads originating in Canada (2% of loads).

may present a welfare gain for both food banks. Therefore heterogeneity is a key determinant of the value of choice. I now investigate, in a reduced form manner, whether there is evidence of heterogeneity in the data.

2.5.1 Heterogeneity Across Food

There is a large degree of heterogeneity across different types of lots. Lots attract significantly different bids depending on their subcategory. Lots also sell for significantly different prices depending on the category. Goods such as cereal and pasta sell for much higher bids than fresh produce and beverages. These differences cannot only be explained by differences in supply: Both cereal and ready meals are in abundance, and sell for relatively high prices. Meanwhile, Health/Beauty and Baked goods are rare, and sell for relatively low prices. This suggests both demand and supply factors at work in determining the prices.

Figure 2.7 shows average winning bids across subcategories, controlling for the censoring caused by the reservation price. These averages are generally statistically different from one another, and a Likelihood Ratio test that subcategory coefficients are equal within a category is rejected at 1% significance level for all but the Pasta category. However there is still much variation in winning bids within subcategories: Variation in subcategories accounts for just 30% of the variation in winning bids. It is clear there is a great deal of heterogeneity between lots, and that these lots cannot be substituted one for one.

FIGURE 2.7: Heterogeneity in Lots

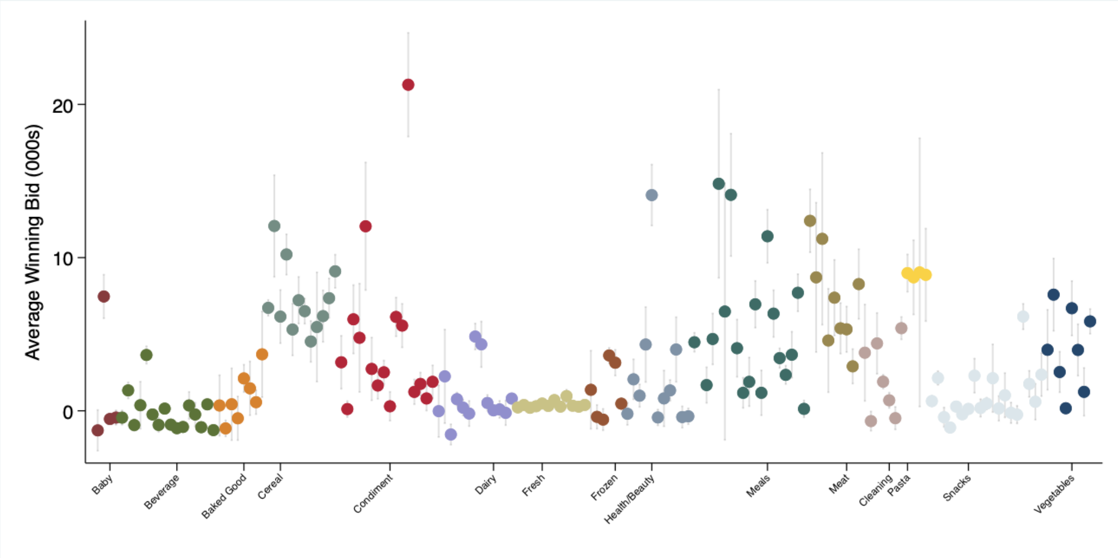

Note: Plots mean winning bids, and 95% confidence intervals, across subcategories, controlling for censoring and lot composition. Coefficients are ordered and coloured according to category.

2.5.2 Heterogeneity Across Food Banks

Food banks differ vastly in terms of their total consumption: Five food banks receive the same amount of food as 122 food banks who receive the least food from Feeding America. However, these food banks are also choosing very different types of food. These 122 food banks, in total, spend 4 times as much as the five high consumption food banks. Therefore these five food banks are choosing to receive much cheaper food. This is likely because they rely on Feeding America for their staples, having fewer local donors than the other 122 food banks. I now investigate how bidding behaviour on different types of food varies across food banks using additional simple reduced form analysis.

Empirical Specification

To gain insight into the degree of heterogeneity across food banks I run a simple Tobit model. The dependent variables are food bank i's bid at time t on lot l, where lot l comes from storage type g and subcategory h:

$$b_{itl} = \alpha_{ig} + \beta_i \, distance_{itl} + \varepsilon_{itl} \qquad b^*_{itl} = \begin{cases} b_{itl} & \text{if} \quad b_{itl} \geq R_h \\ R_h & \text{if} \quad \text{Otheriwse} \end{cases} \qquad \varepsilon_{itl} \sim N(0, \sigma_{ih}) \tag{2.3}$$

Where α_{ig} are food bank $\times$ storage type fixed effects, representing food bank i's average willingness to bid for a lot of type g.

Results

Figure 2.8 plots average bids and 95% confidence intervals across food banks and across different types of food. I focus on different types of food according to how they are stored. I use a Tobit specification to account for censoring caused by food banks only bidding on a small proportion of lots. Food banks are sorted according to their average bid for Dried food. That is, as one moves from left to right, the average 'Dried' bids increase monotonically. Several observations are clear. First, that there is evidence of systematic heterogeneity in average bids across food banks. Second, that there is evidence of systematic heterogeneity in average bids within food banks, across types of food. Third, that these two types of heterogeneity are not perfectly correlated: for some food banks average bids on Fresh food are higher than average bids on Dried food, but for other food banks this relationship is reversed. This demonstrates that bidding behaviour differs systematically across food banks.

As in Prendergast (2022) this permanent heterogeneity across food banks is likely being driven by different food banks having different access to different types of local donors. Food banks who have do not receive one type of food locally will try to get hold of that type of food from Feeding America instead. A food bank that receives many local donations of produce will rarely bid on produce auctioned through the Choice System. Some of this heterogeneity may also come from food banks' clients demanding different types of food, perhaps due to differences in the type of poverty. Vertical heterogeneity will arise from differences in scale, with smaller food banks (in terms of catchment areas, or physical warehouse size) bidding less often and less aggressively. Differences in storage capacities may also be driving some of the heterogeneity across food types - food banks with lots of refrigerated storage probably have greater demand for food that needs to be refrigerated.

FIGURE 2.8: Heterogeneity Across Food Banks

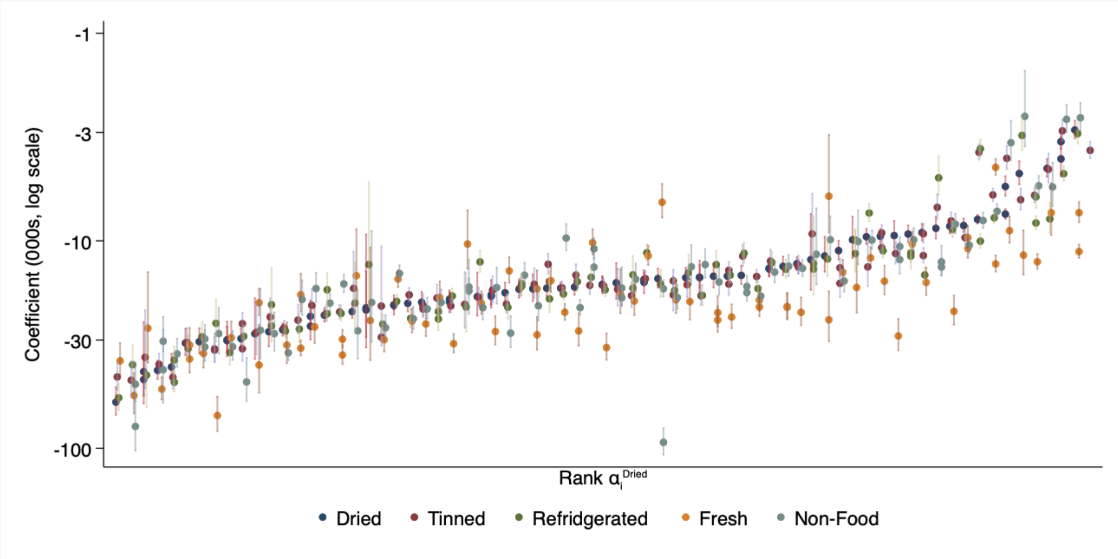

Note: The figure plots coefficients and 95% confidence intervals from a regression of Food bank $\times$ food storage type on bids, controlling for distance and censoring (non-bidding). Coefficients are ordered by the 'Dried' coefficients. Only includes results for food banks who placed at least 50 bids on a food type. Each food bank $\times$ type cell averages $3,450$ observations.

2.5.3 Heterogeneity Across Time

I investigate temporal variation in bidding behaviour by relaxing the Tobit specification given in equation 2.3 above. I investigate how each food bank i's bid on food of type g varies across months m, writing α_{igm} for these average bids, essentially allowing that α_{ig} parameters in equation 2.3 to vary from one month to the next. I estimate the model only on food banks who win at least 100 lots over the period, dropping the first and last months due to incomplete data. I also control for the distance between the food bank and the lot. Each food bank $\times$ type $\times$ month cell averages around 80 observations. I also estimate a restricted model with average bids α_{ig} fixed over time. The hypothesis test of interest is whether $\alpha_{igm} = \alpha_{ig}$ for all m.

This hypothesis test may be underpowered to reject a null hypothesis of constant average bids. It does not take into account that average bids likely don't shift neatly at the beginning of each month. It also does not take into account that large variation in bids within a month (which may cause failure to reject the null) are also indicative of variation in food banks' needs. If the within month variation is on a similar scale to the across month variation in average bids this reduces my power to reject the null hypothesis.

However, the test may be over powered if variation in factors other than food banks' needs is mistaken for variation in needs. For example, if the quality of food varies unobservably over time, this may cause systematic variation in bidding behaviour that should not be attributed to variation in food banks' needs. To account for this possibility I estimate a second restricted specification with food bank specific month fixed effects. These fixed effects will capture variation in bidding behaviour that is common across food types. Under this specification a rejection of the null is evidence of systematic variation over time in bidding behaviour *on specific types of food*. This specification is almost certainly underpowered. If food banks need more food of all types in certain months the fixed effects will also soak up this variation.

Figure 2.9 plots the likelihood ratio test statistic across food banks. The dotted lines gives the χ^2 critical values for tests at the 5% significance level. The red points give the baseline specification, while the blue points give the specification including month fixed effects.

I can reject the null of constant average bids over time, at 5% significance level, for 96% of food banks in my baseline specification, and 70% of food banks for the month fixed effects specification. So, I have strong evidence that food banks' bidding behaviour, and hence their needs, vary significantly over time.[16]

This heterogeneity over time is not driven by variation in supply as this is already controlled for. It is also unlikely to be entirely driven by seasonality, as seasonality should be captured by the time fixed effects included in the second specification. Instead, this variation is likely to be driven by idiosyncrasies in the local donations food banks receive from one month to the next. There may also be variation in demand from food pantries from month to month.

[16]While we see significant evidence of heterogeneity over time, as I demonstrate in Appendix F we do not see evidence of trends or systematic breaks in this behaviour. This is important because a key assumption I will make for structural estimation is that the data, and the equilibrium of the underlying auction game, are stationary.

FIGURE 2.9: Heterogeneity Across Time

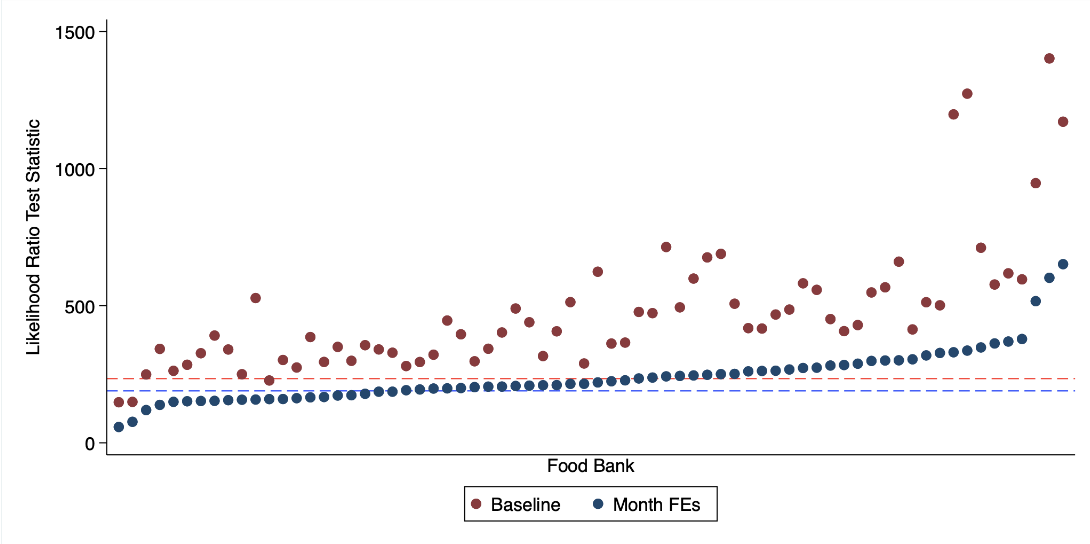

Note: This figure plots likelihood ratio test statistics for the hypothesis test that average bids for each type of food are constant over time, against the alternate hypothesis that bids vary by month. The estimated model controls for censoring, distance, and lot composition. The blue results also include month fixed effects. Under this null hypothesis the test statistic takes a χ^2 distribution with 200 (red) or 160 (blue) degrees of freedom. Critical values for tests at the 5% significance levels are plotted as horizontal lines.

2.6 Exploratory Analysis

I now investigate several stylised facts which point towards key determinants of bidding behaviour, motivating my model's key features. I have emphasised the role of heterogeneity, and established the existence of several types of heterogeneity that will become features of my model. In addition, I point to the importance of negative bidding, strategic behaviour, as well as both dynamic and static complementarities across lots.

2.6.1 Negative Bidding

Negative bidding is common: 27% of bids are negative. Furthermore, with a negative reservation price non-entry only happens when food banks have negative marginal valuations, when food banks must be paid to accept certain loads. This occurs in 98% of bidder $\times$ lot combinations. Negative valuations likely occur because of limited storage capacity, as emphasised in Prendergast, 2017. They cannot throw away excess (non-expired) food as this sends a bad signal to donors. Therefore, I require a model that incorporates these storage costs and negative marginal valuations.

2.6.2 Strategic Behaviour

Figure 2.10 shows that 13% of bids are at the reservation price. If food banks truthfully bid their values this implies point mass at the reservation price, and that the distribution of underlying values is non-smooth. Point mass at −2000 requires there is something special about this number that causes it to have this point mass. A simpler assumption is that food banks are strategic — that food banks with values just above the reserve price shade their bids down to this reserve. The small reduction in likelihood of winning is more than compensated by the increased surplus in the case of winning.

FIGURE 2.10: Strategic Bidding

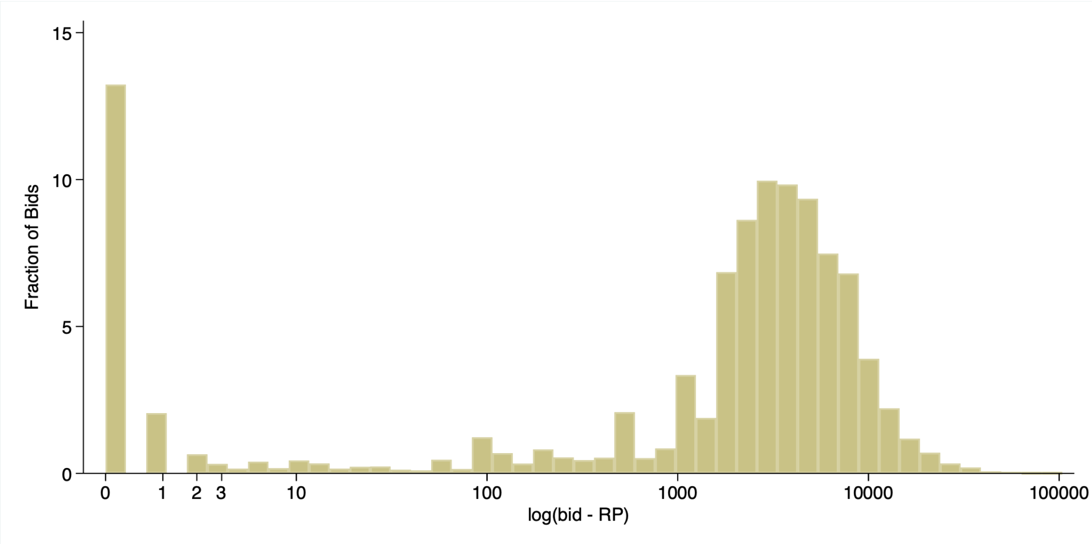

Note: This figure plots a histogram of the log distance between a bid and the reservation price for that lot. The figure shows large mass of bids at the reservation price, and smaller but still significant mass of bids just above the reservation price. If food banks bid their true values, we should not see this point mass.

Another possibility is that food banks are so unstrategic that they just bid the reserve price because it is a particularly salient number. However, figure 2.10 also shows excess mass of bids just above the reserve price, between one and 4 points above. This again requires unpalatable assumptions about the distribution of true valuations having point mass at these apparently special numbers. More likely is that bidders are aware of the mass of bids at the reserve price and are aware that, by bidding the reserve price, they risk tying for the object. Therefore, for some of these food banks it is worth bidding just above the reserve price to mitigate this risk.[17]

[17] I cannot rule out the possibility that only a subset of food banks are strategic, likely those who receive the most fake money and who bid most regularly on the system. However, strategic behaviour

2.6.3 Storable Goods

Figure 2.11 panel (A) demonstrates that, conditional on winning food of a particular type at time 0, the probability of bidding on lots of the same type falls by around 25% (1.5 pp) on subsequent days. This marks a fall from an average probability of bidding of 0.05 down to 0.035.[18] As food banks win more of a particular type of food, the less they are willing to pay for an additional lot from that type. This finding matches anecdotal stories of food banks treating truckloads of food from Feeding America as a Storable Good, in the sense of Hendel and Nevo (2006).

FIGURE 2.11: Evidence of Storable Goods

(A) (B)

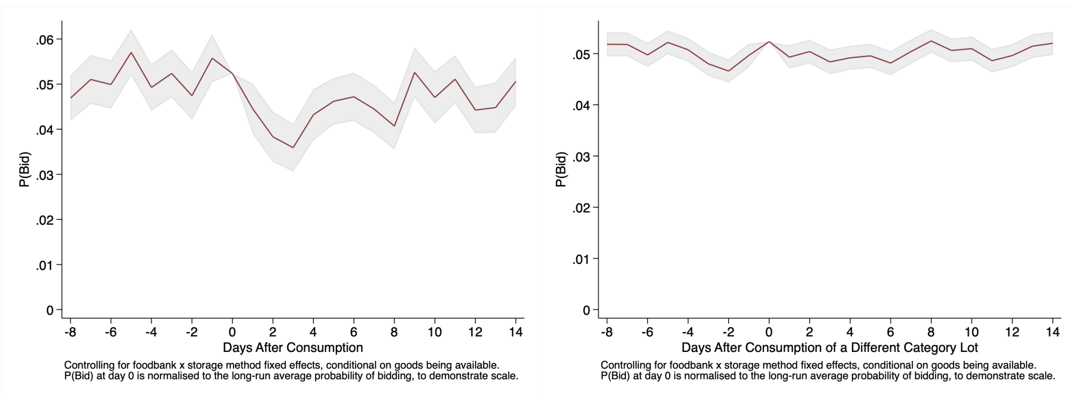

Figure 2.11 panel (B) demonstrates that conditional on winning a lot from one category, the probability of bidding on a different storage type does not fall by a meaningful amount. This finding demonstrates the importance of modelling heterogeneity across types of food. As food banks win more of a particular type, the less they are willing to pay for an additional lot from that type as shown in panel (A). This could be due to storage costs, but could also be because they are exhausting their budget constraints and cannot afford to bid on subsequent days. However, food banks remain willing to pay for lots from different categories, making it less likely that binding budget constraints are the key factor at play.

is most important at the high end of the bid distribution, where bidders shade their values by the largest margins. Whereas in the middle of the distribution we expect relatively little shading, due to the large number of bidders. So long as it is the strategic bidders who are bidding at the high end of the distribution (who tend to have the highest values) this is unlikely to lead to much inaccuracy.

[18] This is higher than the standard 2% probability of bidding as this analysis is performed using food bank × storage method combinations with more than 50 observations of bidding.

Given that food banks are almost certainly forward looking, this finding highlights the need to model dynamics. Food banks treat these large loads of food like storable goods, working through their current stocks before returning to bidding on the Choice System.[19]

2.6.4 Substitutes

Figure 2.12 panel (A) demonstrates that, for a particular type of food, as the number of lots auctioned on a given day increases, food banks bid on a smaller proportion of lots. If pay-offs were additively separable we would see a horizontal line. This suggests that lots exhibit a negative complementarity (substitutes) within a storage type - they do not want to win more food than they can afford to store. I cannot treat auction pay-offs as additively separable, and must instead take a multi-object approach, accounting for the simultaneous auction environment.

FIGURE 2.12: Evidence of Static Complementarities

(A) (B)

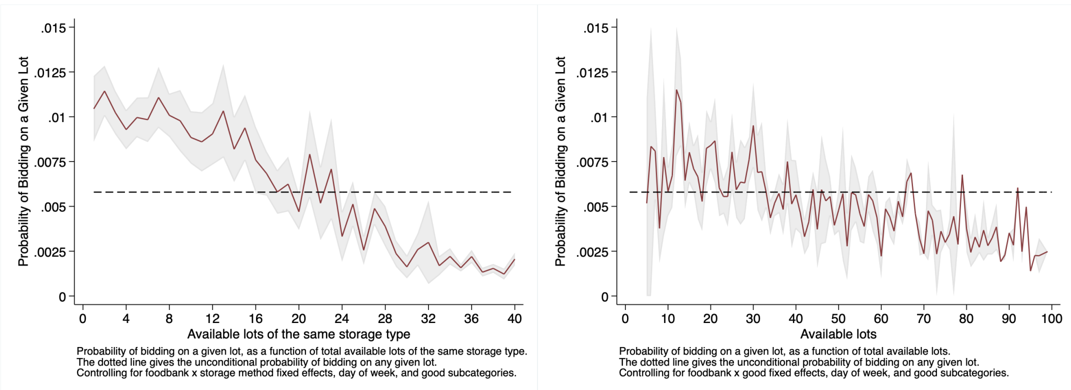

As above, this relationship could also be due to a binding budget constraint - food banks bid less when there is more food available because they cannot afford to bid on so many loads. This is unlikely to be the case because food banks can always

[19]This alone does not necessarily require a dynamic model of forward looking agents. A dynamic model is only strictly required if the counterfactuals of interest sufficiently change the strategic environment, sufficiently altering the agents' continuation values. The counterfactual mechanisms I consider in section 4.5 are strategically very different from the Choice System, as food banks are generally less able to access the food they need than under the Choice System. However, food banks are almost certainly forward looking. Auctions happen very frequently. It is unlikely that food banks do not recognise that if they do not win a lot today a similar lot is likely to be auctioned again in the next few days. This means the marginal benefit of winning a lot today, in a repeated game, is very different from a one-shot game.

bid zero. It may also be due to the transportation costs associated with winning food. However panel (B) demonstrates that we see a much weaker relationship as the total number of lots increases, which we would not expect if this was being driven by either budget constraints or transportation costs. This also further highlights the importance of treating different types of food as imperfect substitutes.

2.7 The Way Forward

There are various ways we could analyse this data and this problem. I take a structural auction approach. This approach has several downsides, predominantly methodological, and numerous upsides. I now discuss the pros and cons of this approach, as well as several alternate approaches I might have taken.

2.7.1 Structural Auction Approach

First, on the structural approach. There are downsides to analysing the problem using structural econometrics, requiring us to fully specify an empirical model and (as will become necessary) make strong parametric restrictions to make estimation feasible. This does have the upside that it forces us to make explicit all the necessary modelling assumptions. The other key upside, and the central reason I choose the structural approach, is that it facilitates welfare analysis. It enables me to consider detailed counterfactual allocation mechanisms, counterfactuals which are rarely identified with a reduced form approach.

An additional difficulty with using a structural model is that, in order to capture the extent of the heterogeneity seen in Section 2.5, we will need a structural model that can allow for large scale heterogeneity. In order to capture the welfare benefits of the Choice System It will be necessary to allow different types of food to be imperfectly substitutable, allow different food banks to want different types of food from Feeding America, and allow these different wants to change over time. Fortunately, if we restrict this heterogeneity at any point, particularly heterogeneity across food banks, this is expected to bias our results in favour of the Old System over the Choice System, making any welfare estimates conservative.

Next, on using an auction model. The key (predominantly methodological) downside, which arises from the results presented in Section 2.6, is that as well as using a strategic model we must also account for both dynamics and the multi-object auction environment. While the econometrics of dynamic auction models have been around for some time, as per Jofre-Bonet and Pesendorfer (2003), and researchers have recently begun working on multi-object auction models, as per Gentry, Komarova, and Schiraldi (2023), there has been no work on the estimation of dynamic multi-object auction models.

If we ignore dynamics, we ignore how under the Choice System food banks choose not just what to consume, but when to consume. Choice allows food banks to plan better. For example, they do not need to hold as much food in case of a large local demand shock (incurring precautionary storage costs), because they know they will always be able to get this type of food from Feeding America when needed. In other words, a food bank's continuation value under the Choice System is likely to be very different to under the Old System. By ignoring dynamics we ignore a key potential benefit of choice.

If we ignore the multi-object environment, treating lots as all additively separable, we ignore how food banks choose the combination of lots they bid on. We cannot rationalise a food bank choosing to only bid on a small number of similar lots because they do not want to win more than they can store. Instead we can only rationalise this as food banks receiving low idiosyncratic draws. A food bank choosing to bid on any given lot with only 5% probability is very different from a food bank *always* bidding on two out of forty lots auctioned each day. It is unclear exactly how this will bias my results on the importance of choice, except for making it more difficult to predict bidding behaviour.

A final key difficulty, that I will not tackle until the final chapter, concerns heterogeneity over time. This heterogeneity most likely arises due to variation in food banks' stocks (Prendergast, 2017), which are unobserved by the econometrician. Important steps have been made recently on the identification and estimation of structural models with unobserved states, including Kasahara and Shimotsu (2009) for the time-invariant case, Hu and Shum (2012) on the one dimensional unobserved state, and Arcidiacono and Miller (2011) for a convenient estimation procedure.

However none of these identification frameworks extends to my setting. In particular it is important to allow food banks' stocks to be multidimensional. This is key for capturing how food banks trade off consuming different types of food, and how different food banks receive different types of food from their local donors (i.e. heterogeneity across types of food). It is important that we allow for heterogeneity over time as it is a (or possibly *the*) key margin of heterogeneity by which food banks benefit from being given choice.

2.7.2 Other Approaches

There are other possible approaches we might take to studying this problem, all of which have various benefits and drawbacks. I now broadly summarise some of these approaches, and explain why I opted not to take them. I focus on approaches that place sufficient 'structure' on the data to allow us to evaluate welfare. If we dropped the requirement of making welfare comparisons then, as in Prendergast (2017), even the simple reduced form evidence of heterogeneity presented in Section 2.5 is sufficient to show that the Choice System performs better than the Old System. However, part of the point of this thesis, as distinct from work by Prendergast, is to make welfare statements.

Demand Estimation

One possible approach, that would be easier than the full structural auction approach, would be to structurally estimate food banks demand functions, either using a discrete or continuous choice approach. On the discrete side, one might estimate the pure characteristics model along the lines of Berry and Pakes (2007), as has become the workhorse of demand estimation. On the continuous side, we might estimate a demand system, such as the Almost Ideal Demand System of Deaton and Muellbauer (1980), either with demand defined at the daily level, or with consumption aggregated over multiple days. I discuss other aggregation based approaches in a separate section.

Both discrete and continuous demand estimation involves a simplifying assumption on the purchasing process, assuming a posted price mechanism instead of an auction mechanism. If the number of bidders was infinite, equilibrium allocations and payments would be the same under both sets of mechanisms. This is clearly not the case in practice, but may pose a reasonable simplification. One major downside of the posted price assumption, however, is that we only make use of data on consumption and payments. We would ignore data on losing bids (and non-bids), which provide important information about food banks' willingness to pay.

The upsides of these approaches are large. Both approaches allow me to easily model dynamics, either using dynamic discrete choice methods, or an Euler Equation approach for the continuous case. Both also enable me to somewhat more easily model the unobserved stocks. With a discrete approach I could estimating a storable goods model, similar to Hendel and Nevo (2006) or Erdem, Imai, and Keane (2003). Whereas with a continuous demand model, allowing for unobserved stocks is as easy as allowing for temporal autocorrelation in the demand residuals.

The downside of the discrete choice approach is that it makes it difficult to model food banks' choice over bundles. While there has been some recent headway on estimating demand models when consumers have preferences over bundles (including Iaria and Wang, 2020, Allen and Rehbeck, 2022, and Fox and Lazzati, 2017), most of these require the number of objects to be small, or are computationally burdensome. Furthermore, there is not yet literature on estimating dynamic multi-object demand, meaning the econometric requirements are still not there.[20]

The continuous demand approach circumvents these difficulties by using implicit optimisation and necessary first-order conditions to get over the dimensionality problems. However this comes at a cost of totally ignoring the discreteness of lots, incorrectly treating them as perfectly divisible. If there were an extremely large number of lots auctioned each day, and food banks all win a large number of lots, then assuming a continuous approximation might be reasonable. Whether the average of 30 lots allocated each day is large enough or not is debatable. More problematic is that each food bank only wins a small number of lots each day, so that the continuous approximation is almost certainly poor. A related issue is the large number of zeros in the data: The fact that most food banks frequently win no food of

[20] As we will see, the multi-object problem is distinctly easier to solve in the auction setting. It recasts the discrete options into a continuous action space (bids), so that estimation is based on a limited number First Order Conditions, instead of an exponential number of inequality conditions.

certain types. The selection and censoring based approaches typically used to correct for zeros make allowing for the unobserved stocks significantly more difficult.

Aggregation Based Approaches

Other approaches involve aggregating consumption over time, reducing the day-to-day noise in consumption, and overcoming the observed discreteness of lots. These aggregation approaches also require the fixed-price and price taking assumptions required for the demand estimation approach.

Building on the continuous demand system approach above, we could overcome the discreteness problems by aggregating consumption at, say, the monthly level. This type of aggregation is already used implicitly in most studies estimating continuous demand.

Rather than the fully structural approach, another possibility is to take a sufficient statistics approach, as in Prendergast (2022). We build a tractable model of consumption that yields some measure of heterogeneity as a sufficient statistic for the welfare under the Choice System, relative to the Old System. As an example, Prendergast aggregates consumption at the two month level, assumes a quadratic form for food banks' payoff function, and roughly shows that the variance of consumption across food banks can be used, in conjunction with the elasticity of demand, as this sufficient statistic. The computational ease of this approach is a major benefit, though questions remain about how to go about estimating these key statistics (Prendergast takes a figure for the elasticity of demand for different types of food from the literature).

A related possibility is to employ a Welfare Index. Using some sufficiently general model, we may be able to come up with an index number, likely a quantity index, that allows us to compare welfare across two allocations. For example, using a Tornqvist quantity index for comparison of allocations under the Choice System and the Old System would yield a second order approximation to any set of utility functions.[21] Likewise, we might take a non-parametric revealed preference approach in

[21] This approach has other difficulties, in that while I arguably observe allocations under the Old System (as every food bank got the same food in expectation), I do not observe prices under the Old System. Unless we are willing to assume homotheticity of payoffs, both sets of supporting prices would be needed to evaluate this index

the spirit of Samuelson (1947) and Afriat (1972). We can use the revealed preference axioms to compare allocations under the Choice System and the Old System, and observe what proportion of Old System bundles lie within their budget sets under the Choice System. This allows us to find a lower bound on the welfare change.[22]

The key problem with these aggregation approaches, aside from the fixed price assumption, is that they require us to make a statement about the correct period of aggregation. We must assume that agents have some accounting horizon over which they make their consumption decisions, and over which market fundamentals remain constant. It is true that food banks likely do look ahead, and may well account over a finite intervals (and we can test robustness to different interval lengths). However, it is difficult to argue that food banks operating on very different scales use the same accounting intervals. Large, and particularly busy, food banks likely account over a much shorter interval than those who only send out food maybe once or twice a day. In addition, it is difficult to justify the boundaries of these intervals - in the last few days of a month, food banks are likely already looking ahead to the next month, rather than only thinking about hitting their monthly targets from the current month. In practice, food banks make their consumption decisions every day, reacting to new information about local donations, making it difficult to justify this type of aggregation approach.

2.7.3 Next Steps

Having settled on this structural auction approach, I will now build a model of forward looking bidding in these repeated rounds of simultaneous first-price auctions. This will involve unifying the models of Jofre-Bonet and Pesendorfer (2003) and Gentry, Komarova, and Schiraldi (2023). I will begin by focusing on the simplest possible version of a dynamic multi-object auction model, abstracting away from reserve prices, entry decisions, inter-temporal budget constraints, and the unobserved states. With the exception of the unobserved stocks, which I will hold off from analysing until chapter 4 (as it requires a more substantial alteration to the baseline model), these other extensions will all be analysed in Appendices.

[22]Note however that both this revealed preference approach, and a welfare index approach, do not allow us to account for unobserved stocks, which essentially appear as though preferences change from one month to the next.

Chapter 3

A Model of Dynamic Multi-Object Auctions

3.1 Chapter Introduction

In this chapter I develop an empirical model of bidding in repeated rounds of simultaneous first-price auctions, and study identification and estimation in this framework. I begin with the simplest possible setting, assuming entry is exogeneous, reservation prices do not bind, and all state variables are observed. I relax these restrictions in Appendix E.

Analysing this problem has broader appeal than the specific Feeding America setting that I am interested in. First-price auctions, which are regularly used to allocate government procurement contracts, rarely take place in isolation. Multiple lots (contracts) are often auctioned simultaneously, and auctions are repeated whenever new contracts become available. In real world environments bidders' values may be non-additive across different lots. For example, bidders may face capacity constraints, facing higher costs the larger their current backlog. Or, they may benefit from economies of scale, facing lower costs when working on many of the same type of contract at once. The structure of these non-additive values is highly relevant for auction design - should similar contracts be auctioned simultaneously, or spaced out over time? When capacity constraints are the dominant factor, auctioning a large number of contracts simultaneously may create inefficiencies by depressing competition. However, if firms are able to exploit economies of scale it may be worth auctioning similar contracts simultaneously, or even bundling the lots together.

Previous research has either studied forward looking bidders and assumed auctions are single-object, or studied auctions of multiple objects and assumed bidders are myopic. For example, both Jofre-Bonet and Pesendorfer (2003) and Gentry, Komarova, and Schiraldi (2023) study synergies in bidding behaviour in repeated simultaneous first-price auctions for highway maintenance contracts. Jofre-Bonet and Pesendorfer (2003), who estimate a dynamic single object model, find significant negative effects of capacity constraints on bids. Meanwhile Gentry, Komarova, and Schiraldi (2023), who study simultaneous first-price auctions in a static setting, find similar capacity constraint effects; however, they also find evidence of positive synergies among similar contracts that allow firms to exploit economies of scale. The implication is that neither paper accurately models the effect of non-additive values on bidding decisions. To the best of the author's knowledge this study is the first to consider both a dynamic and multi-object approach to an empirical auction model.

My model is fundamentally a blend of the models presented in Jofre-Bonet and Pesendorfer (2003) and Gentry, Komarova, and Schiraldi (2023), henceforth referred to as JP and GKS respectively. Bidder pay-offs are represented as the sum of privately known and potentially correlated *lot specific* values, a *combination specific* flow payoff, and a *combination specific* continuation value. Following GKS, the combination specific flow payoff is treated as a deterministic function of state variables.[1] This is a natural framework that reflects known capacity constraints or economies of scale. The model primitives consist of the distribution of lot specific values and the combination specific flow payoff function.[2]

Building on this framework I make three key contributions to the empirical auction literature. First, I extend GKS' identification framework to make use of variation in a bidder's individual state variables, such as their backlog of contracts, to non-parametrically identify their combination specific value function *without* the need for exclusion restrictions. Intuitively, identification arises because variation in the state

[1] In Appendix E I consider several extensions when the combination value is stochastic. Most notably, I consider the case when the combination value is a function of low-dimensional unobservables. For example, if we do not observe firm backlogs. I show that model primitives remain identified in this case, so long as the function has a strictly monotonic first derivative.

[2] Like GKS this assumption allows me to separately identify complementarities and affiliations, the central problem studied by Kong, 2021. Affiliation across lots comes through correlation in the lot specific pay-offs, while the synergies remain deterministic. However, like both papers I assume the lot-specific pay-offs are independent across players.

causes variation in bidders' combination values, which in turn causes variation in their bidding behaviour. If lots are substitutes we expect to observe more aggressive bidding when backlogs are low. Following the approach presented in GKS I translate the inverse bidding system, conditional on a given state, into a system of linear equations in the unknown combination values. I show that the combination values are identified by combining these systems of linear equations across state variables, proving that the combined system has a unique solution. This result is important because it ensures the combination value is identified without the need for exclusion restrictions that prohibit identification of forward looking behaviour.[3]

Second, I extend this identification framework to a dynamic setting. I prove that identification of the joint distribution of equilibrium bids (conditional on the state), as well as identification of the state transition process, are sufficient for identification. Building on JP I show that the continuation value can be written as a recursive function of the observed distribution of bids and the sum of the immediate combination value and the discounted continuation value. This recursive formulation can be solved for the continuation value as a function of the observed distribution of bids and the immediate combination value. Substituting this identity into the system of first order conditions, combined across state variables, yields a system of equations in the unknown combination values. I prove that, under mild conditions, this system has a unique solution.

Third, I outline a three step procedure for estimating the model. The estimation procedure generalises that of JP for estimating dynamic auctions, overcoming the problem that the continuation value cannot be written as a function of the observed distribution of bids only. In the first step one estimates the probability that each bidder wins any given lot - the first step in most empirical auction studies. In the second step one estimates the distribution of underlying lot-specific values, and the sum of the immediate combination value function and the discounted continuation value function. I refer to this sum as the 'pseudo-static' pay-off function. It is essentially the object one would estimate if we incorrectly estimated the model as if it

[3]Without additional restrictions, excluded variable (such as common state variables or the states of other bidders) enter each bidder's continuation value, directly affecting their bidding behaviour and violating the exclusion restriction, rendering the model non-identified in GKS's framework.

was a static model.[4] In the third step one disentangles this sum, separately estimating both the immediate combination value and the continuation value. This exploits the fact that the continuation value can be written as a recursive function of the distribution of observed bids and the pseudo-static pay-off function.[5] This estimation procedure is shown to be computationally feasible, particularly in an environment where the interior fixed point algorithm is extremely slow. I then use a Monte Carlo study to evaluate the performance of this estimator.

The structure of this chapter is as follows: In section 3.2 I introduce the auction game and associated auction model. Section 3.3 introduces the identification framework and proves that model primitives are point identified. Section 3.4 outlines the proposed three step estimation procedure, and Section 3.4.3 presents a simulation study evaluating the performance of this estimator. Several additional results are presented in the Appendices. Appendices B - D present technical proofs. Appendix E presents several extensions to the identification and estimation framework.[6]

3.2 The Model

3.2.1 Environment

I now present the dynamic auction model that is the focus of this paper. I build on the models of Gentry, Komarova, and Schiraldi (2023) and Jofre-Bonet and Pesendorfer (2003). Assumptions necessary for identification are discussed in section 3.3.

Setup: Each period t, over an infinite horizon, n risk-neutral players i compete in a series of first-price Sealed Bid auctions. Player i wins lot l in period t if $b_{itl} \geq \max_{i' \neq i} \{b_{i'tl}\}$. Sealed bids are placed simultaneously, then winners are announced simultaneously. Winners pay their bids, and every player observes the bids and identities of winners. Define the $L \times 1$ vector $\mathbf{w}_t$ as the outcome at time t, where

[4] This step is similar to the second step in the estimation procedure presented by GKS. As in GKS, if one is non-parametrically estimating the model then this estimation step will actually consist of two substeps: First, estimating the deterministic combinatorial valuations, then estimating the distribution of stochastic lot specific valuations.

[5] This step is similar to the second step in JP's procedure, wherein the continuation value is written as a function of the distribution of observed bids, allowing us to back out the cost function.

[6] These include extensions for second-price auctions, reservation prices, endogenous entry, intertemporal budget constraints, and stochastic combination values.

$w_{tl} = i$ if i won lot l at time t. Ex-ante hypothetical outcomes are denoted by $\mathbf{w}_t^a$. The index l is a label and need not reflect anything substantive.

Reservation Prices and Ties: I assume reservation prices do not bind, that auction entry is exogenous, and that ties occur with probability zero.[7]

Lots and Lot Characteristics: Each period a set of lots $\mathbb{L}_t$ comes to auction. L possible lots can be auctioned in a given period, where L is finite. Each lot l is characterised by a row-vector of characteristics $\mathbf{x}_{tl}$, writing $\mathbf{X}_t$ for the stacked characteristics of all lots available in period t.[8] The set of characteristics, $\mathbb{X}$, is assumed finite. If $l \notin \mathbb{L}_t$ then $\mathbf{x}_{tl} = \emptyset$. Finally stack the lot characteristics, set of available lots, and other common state variables into $\mathbf{s}_{0t} \in \mathbb{S}_0$.

3.2.2 Primitives

Individual States: Player i begins the period in state $\mathbf{s}_{it}$. This may represent a player's existing stock of the good, or backlog of contracts. The set of possible individual states, $\mathbb{S}_i$, is assumed to be finite.[9] If the outcome at t is $\mathbf{w}_t^a$ then player i ends the period in state $\mathbf{s}_{it}^a$, referred to as the ex-post state. $\mathbf{s}_{it} = \mathbf{s}_{it}^a$ if and only if the player does not win a single lot. For notational convenience, define the set $\mathbb{S}_i^a(\mathbf{s}_i, \mathbf{s}_0)$ as the set of possible individual ex-post states $\mathbf{s}_i^a$ having started in state $\mathbf{s}_i$, given the available lots, lot characteristics, and other common state variables $\mathbf{s}_0$.

Total States: Stack the individual states $\{\mathbf{s}_{it}\}_{i \in \mathbb{I}}$, and $\mathbf{s}_{0t}$, into the total state variable $\mathbf{s}_t \in \mathbb{S}$, where $|\mathbb{S}| = S$ is finite. In section 3.3.6 I set out sufficient conditions on the set $\mathbb{S}$ which ensure identification. Similarly, Stack the ex-post states $\{\mathbf{s}_{it}^a\}_{i \in \mathbb{I}}$ and $\mathbf{s}_{0t}$ into the total ex-post state $\mathbf{s}_t^a \in \mathbb{S}$.

Transition Functions: At the beginning of each period, the state $\mathbf{s}_t$ is drawn stochastically from $T_{\mathbf{s}}(.|\mathbf{s}_{t-1}^a)$. Because $|\mathbb{S}|$ is finite, the transition probabilities can be described by a standard transition matrix T, such that $P(\mathbf{s}_{it} = \mathbf{s}_m | \mathbf{s}_{t-1}^a = \mathbf{s}_n) = T_{mn}$.

[7] However these assumptions are relaxed in appendices E.2, and E.3.

[8] Characteristics may include details of lots being auctioned. Such as the lot's location.

[9] This is predominantly for mathematical convenience, but is likely to hold in practice. For example, Highway maintenance companies likely have a maximum number of contracts they can feasible hold at any given time, and the measurement of their backlog of contracts can be arbitrarily discretised into days of work remaining.

Actions: Each player plays an L dimensional vector of bids each period, denoted $\mathbf{b}_{it}$. The set of possible bids conditional on available lots $\mathbb{B}(\mathbb{L})$ is assumed to be convex and compact, so that $b_{itl} \in [\underline{b}, \bar{b}]$ for l such that $l \in \mathbb{L}_t$.

Lot Specific Values: I focus on an independent private value framework. Player i's lot specific valuation, v_{itl}, conditional on $l \in \mathbb{L}_t$ is a mean zero random variable observed before they make their entry decisions.[10] If $l \notin \mathbb{L}_t$ then v_{itl} is normalised to zero. Stacking these values $\boldsymbol{v}_{it}$, a $L \times 1$ vector, is drawn from cumulative density function $F_i(.|\mathbf{s}_t)$ with support $[\underline{\boldsymbol{v}}_i, \bar{\boldsymbol{v}}_i]$.[11]

Combination Value: The combination value is given by $J_i(\mathbf{s}_t)$, a $2^L \times 1$ vector. Each row $J_{ia}(\mathbf{s}_t)$ gives the mean flow pay-off corresponding to a different outcome $\mathbf{w}_t^a$, ending the period in state $\mathbf{s}_{it}^a$.[12] Entries corresponding to winning lots that are not available are normalised to zero. J_i is assumed a deterministic function of $\mathbf{s}$, and assumed to be finite. A player's type is characterised by the tuple $(\boldsymbol{v}_i, J_i)$. I assume that F_i and J_i are both common knowledge.

3.2.3 The Agent's Problem

Strategies: A (pure) strategy consists of a mapping from a player's type $(\boldsymbol{v}_i, J_i)$ and the state of the world $\mathbf{s}$ onto a series of bids $\mathbf{b}_{it}$. Ex-ante a player's strategy admits a distribution of bids σ_i according to F_i as well as J_i and $\mathbf{s}$.

Marginal Win Probabilities: Denote $G_{i'l}(.;\sigma_{i'})$ and $g_{i'l}(.;\sigma_{i'})$ respectively the marginal cdf and pdf of individual i''s bid on lot l according to their strategy $\sigma_{i'}$. Denote $\Gamma_i(\mathbf{b};\sigma_{-i})$ the $L \times 1$ vector where row l contains the probability that i wins lot l, given their bid and entry decision, taking as given the strategies of other players. Because ties occur with zero probability we can write:

$$\Gamma_{il}(b_{ilt};\sigma_{-i}) = \prod_{i' \neq i} G_{i'l}(b_{ilt};\sigma_{i'})$$

[10] The mean zero assumption is without loss of generality. The mean of the variable will be absorbed into the combination value.

[11] In the current paper I assume these values are, conditional on the state, independent over time. This is predominantly for comprehensibility. The identification results carry over to the case of dependence over time, however depending on the structure of this dependence estimation may become more involved.

[12] Even though $J_i(\mathbf{s}_t)$ varies with $\mathbf{s}$ because $\mathbf{s}_{0t}$ dictates the type of lots available and hence the possible ex-post states. However I will assume that the values of $J_{ia}(\mathbf{s}_t)$ only depends on $\mathbf{s}_{it}^a$.

Combination Win Probabilities: Denote $P_i(\mathbf{b};\sigma_{-i})$ the $2^L \times 1$ vector where row a contains the probability, conditional on their bid and the strategies of other players, that player i's ex-post state will be $\mathbf{s}^a_{it}$.

Overall Combination Probabilities Denote $Q_i(\mathbf{b};\sigma_{-i})$ the $n^L \times 1$ vector where row a contains the probability, conditional on their bid and the strategies of other players, that the outcome from period t is $\mathbf{w}^a_t$, and so the overall ex-post state is $\mathbf{s}^a_t$. This object is extremely similar to the combination win probabilities presented previously, except this object also takes into account exactly which player $j \neq i$ wins each lot. Importantly, $P_a = \sum_{a' \text{ s.t. } \mathbf{s}^{a'}=\mathbf{s}^a} Q_{a'}$. That is, summing $Q_{a'}$ over all the ex-post outcomes for which player i's state is the same ($\mathbf{s}^{a'} = \mathbf{s}^a$) gives P_a. For both the marginal and combination win probabilities entries of both vectors that correspond to winning unavailable or un-entered lots are normalised to zero.

Discounting: Players are assumed to have temporally additively separable preferences, and make forward looking decisions with discount parameter $\beta \in (0,1)$. β is assumed known by players and the econometrician.

Expected Flow Pay-off: Payoffs are assumed quasi-linear in payments.[13] Consider player i with a realisation of $\boldsymbol{v} = \boldsymbol{v}_{it}$ who places bid $\mathbf{b}$ against players bidding according to strategies σ_{-i}. For notational convenience I also suppress subscripts.

$$\Pi(\mathbf{b}|\boldsymbol{v}_i,\mathbf{s};\sigma_{-i}) = \Gamma_i(\mathbf{b};\sigma_{-i})^T(\boldsymbol{v}_i - \mathbf{b}) + P_i(\mathbf{b};\sigma_{-i})^T J_i(\mathbf{s}) \tag{3.1}$$

Value Function: The Value Function for player i is then given as:

$$W_i(\boldsymbol{v}_{it},\mathbf{s}_t;\sigma_{-i}) = \max_{\mathbf{b}} \left\{ \Gamma_i(\mathbf{b};\sigma_{-i})^T(\boldsymbol{v}_i - \mathbf{b}) + P_i(\mathbf{b};\sigma_{-i})^T J_i(\mathbf{s}_t) + \beta Q_i(\mathbf{b};\sigma_{-i})^T V_i(\mathbf{s}_t;\sigma_{-i}) \right\} \tag{3.2}$$

Were $V_i(\mathbf{s}_t;\sigma_{-i})$ is the continuation value, to be discussed shortly.

Ex-Ante Value Function: Define the (scalar) ex-ante value function by taking an expectation over private information $\boldsymbol{v}_{it}$: $V^E_i(\mathbf{s}_t;\sigma_{-i}) = \int_{\boldsymbol{v}_{it}} W_i(\boldsymbol{v}_i,\mathbf{s}_t;\sigma_{-i}) dF(\boldsymbol{v}_i|\mathbf{s}_t)$

Continuation Value: The combination continuation value is given by $V_i(\mathbf{s}_t;\sigma_{-i})$, a $n^L \times 1$ vector. Each element a of this vector contains the continuation value $V_{ia}(\mathbf{s}_t;\sigma_{-i})$

[13] However, in appendix E.4 I extend the model to allow for an inter-temporal budget constraint. I prove that quasi-linearity is observationally equivalent to the inter-temporal budget constraint model in a stationary setting, when players have constant marginal utility of wealth.

corresponding to a different allocation of lots and ending the period in a different state $\mathbf{s}_t^a$. This relates to the ex-ante value function by taking an expectation over the next period state $\mathbf{s}_{t+1}$, given $\mathbf{s}_t^a$: $V_{ia}(\mathbf{s}_t;\sigma_{-i}) = \int_{\mathbf{s}_{t+1}} V_i^E(\mathbf{s};\sigma_{-i})dT(\mathbf{s}|\mathbf{s}_t^a)$.

3.2.4 Equilibrium

I now discuss the dynamic equilibrium, and the assumptions required for existence of an equilibrium. A full and general proof of equilibrium existence is beyond the scope of this paper.[14] In place of a complete proof I present a proof of equilibrium existence under the conjecture that a pure-strategy Bayesian Nash Equilibrium exists in the static game without entry. The lack of full existence proof should be considered only a theoretical issue, rather than a practical problem.

Markov Perfect Equilibrium I focus on symmetric markov perfect equilibria consisting of a set of strategies $\boldsymbol{\sigma}^*$ such that for any $(\boldsymbol{v}, J, \mathbf{s})$:

$$\mathbf{b}_i^{\sigma^*} = arg \max_{\mathbf{b}} \left\{ \Gamma_i(\mathbf{b};\sigma_{-i}^*)^T(\boldsymbol{v}_i - \mathbf{b}) + P_i(\mathbf{b};\sigma_{-i}^*)^T J_i(\mathbf{s}_i) + \beta Q_i(\mathbf{b};\sigma_{-i}^*)V_i(\mathbf{s};\sigma_{-i}^*) \right\} \tag{3.3}$$

Equilibrium Existence

To prove equilibrium existence, I rely on the following conjecture:

Conjecture 1. Existence and Uniqueness of a continuous Static Equilibrium

There exists a unique symmetric (non co-operative) Pure Strategy Bayesian Nash Equilibrium of the (myopic) stage game, such that for all i and l the expected pay-off is continuous in $\boldsymbol{v}_i$ and J_i.

This conjecture takes essentially the same form as the assumption that a continuous and unique equilibrium exists in GKS.

Proposition 1. *Under the assumptions of the game, and under Conjecture 1, a Symmetric Markov Perfect Equilibrium exists.*

[14] To my knowledge, no complete proof of equilibrium existence exists even for the static game without entry. This paper joins the line papers studying sufficiently complex auction games in which neither existence, nor uniqueness of equilibrium can be guaranteed. For example, Gentry, Komarova, and Schiraldi (2023) on simultaneous first-price auctions, Fox and Bajari (2013) on simultaneous ascending auctions, or Jofre-Bonet and Pesendorfer (2003) on dynamic single-object first-price auctions. If the bid space were discrete, then static equilibrium existence follows from Milgrom and Weber (1985).

Proof of Proposition 1 is relegated to Appendix B, as existence is not my main focus. The proof involves showing that the equilibrium payoff in the stage game is consistent with the continuation value, employing Kakutani's fixed point theorem.

An additional consequence of Conjecture 1 — in conjunction with the assumptions of a compact action space, a compact payoff space, and a finite state space — is that we can apply Theorem 3.2 from Bhattacharya and Majumdar (1989) to ensure that the value function exists and is unique.

3.3 Identification

I now demonstrate that the distribution of lot specific values F, and the combination value J, are non-parametrically point identified. The intuition behind this argument is that variation in $\mathbf{s}$ causes variation in payoffs which, in turn, cause variation in bidding behaviour. I use the observed bidding behaviour, with information about bidders' equilibrium beliefs, to 'back out' the distribution of values. For example, we see how players bid when their stocks are empty, and compare behaviour to when they have some food in stock, or a lot of food in stock. If their payoff function is concave we expect bidding to be less and less aggressive the more food in stock. This essentially allows us to trace out their payoff function.

A brief discussion of what I mean by non-parametric identification is helpful: A model is point identified if, given the implications of equilibrium behaviour, the joint distribution of bidder's pay-offs, $\{F_i, J_i\}_{i\in\mathbb{I}}$, are uniquely determined by the joint distribution of observables (Athey and Haile, 2002). Using the terminology of Lewbel (2019) we say that a model is non-parametrically identified if the identified objects are functions. This is in the sense that we do not assume a functional form, but identify the entire function $J_i(\mathbf{s})$ for every $\mathbf{s} \in \mathbb{S}$ and $F_i(\boldsymbol{v}|\mathbf{s})$ for every pair $(\boldsymbol{v}, \mathbf{s})$.

I begin by introducing the assumptions necessary for identification in subsection 3.3.1. In subsection 3.3.2 I consider the agent's optimisation problem, and establish necessary First Order Conditions. Next, I prove that the model primitives are identified in 4 steps. First, in subsection 3.3.3, I show that conditional on identification of J and V, F is non-parametrically identified from the inverse bidding system. This argument is only a minor extension to the argument presented in GKS.

Second, in subsection 3.3.4 I show that conditional on identification of J, V is non-parametrically identified. Third, in subsection 3.3.5 I show that identification of J collapses down to a rank condition. Finally, in subsection 3.3.6 I demonstrate that only very mild restrictions on the state space are sufficient for this rank condition to hold. In subsection 3.3.7 I then consider identification under several extensions of the model presented.

3.3.1 Assumptions

Define the objects $G_i(.|\mathbf{s}), \Gamma_i(.|\mathbf{s}), P_i(.|\mathbf{s})$, and $Q_i(.|\mathbf{s})$ as the empirical counterparts to the objects presented previously.

Assumption 1. *For each t, the econometrician has a set of observations as follows:*

$$\mathbb{O}_t = \left\{ \mathbf{w}_t, \mathbf{s}_t, \{\mathbf{b}_{it}\}_{i \in \{1,2,\ldots,n\}} \right\} \tag{3.4}$$

I assume the econometrician observes all bids and entry decisions, not just the winning bid. Under this assumption G, Γ, P, Q, and T are all non-parametrically identified. Therefore, in the remainder of this section I treat these objects as known.

Assumption 2. *The data $\{\mathbb{O}_t\}_{t=1\ldots T}$ are generated by strategy profile σ^* which is a symmetric Markov perfect equilibrium of the dynamic auction game.*

This assumption requires that the same equilibrium is played throughout the observed period, ensuring that strategies can be written as a function of the state. As a result, the continuation value can be written as a function of the state. We can then express the continuation value in vector form as $\mathbf{V}$, with elements corresponding to the expectation from ending a period in any particular ex-post state. It is then useful to define the relationship between the n^L vector $V(\mathbf{s})$ defined previously and $\mathbf{V}$:

$$\begin{pmatrix} V(\mathbf{s}_1) \\ \vdots \\ V(\mathbf{s}_S) \end{pmatrix} = A\mathbf{V}$$

Where A is an $Sn^L \times S$ dimensional matrix. I often also use the relationship $V(\mathbf{s}) = A_{\mathbf{s}}\mathbf{V}$ for the $n^L \times S$ submatrix $A_{\mathbf{s}}$. This matrix contains a 1 in entry am if the potential outcome $\mathbf{w}^a$ yields ex-post state $\mathbf{s}^a = \mathbf{s}_m$. That is, this matrix just selects the relevant continuation values that correspond to possible ex-post states.

Assumption 3. *For all i and $l \in \mathbb{L}$ $\Gamma_{il}(b_{il}|\mathbf{s};\sigma^*_{-i})$ is strictly increasing and differentiable in b_{il}. Similarly, for all i, j, and $l \in \mathbb{L}$ the object $Prob(j \text{ wins } l|b_{ilt};\sigma^*_{-i})$ is continuous and differentiable in b_{ilt}.*

This assumption ensures that the marginal, combination, and over-all combination win probabilities are continuous and differentiable in $\mathbf{b}$, enabling us to take first order conditions. As in GKS, when this assumption fails we lose point-identification, though the model primitives generally remain partially identified.

Assumption 4.

i) *Element a of $J_i(\mathbf{s})$ can be written as: $J_{ia}(\mathbf{s}) = j_i(\mathbf{s}_i^a)$ for some $j_i : \mathbb{S}_i \to \mathbb{R}$.*

ii) $E[\boldsymbol{v}|\mathbf{s}] = 0$.

Part i) of this assumption is relatively weak. I require the immediate combinatorial pay-off from ending the period in state $\mathbf{s}^a$ depends only on this final state.[15] Part ii) of the assumption just ensures that the mean of $\boldsymbol{v}$ is absorbed into J.

By stacking J over $\mathbf{s}$ and j over $\mathbf{s}_i$ I define a mapping between $J(\mathbf{s})$ and $j(\mathbf{s}_i)$:

$$\underbrace{\mathbf{J}}_{S2^L\times 1} = \begin{pmatrix} J(\mathbf{s}_1) \\ \vdots \\ J(\mathbf{s}_S) \end{pmatrix} \qquad \underbrace{\mathbf{j}}_{S_i\times 1} = \begin{pmatrix} j(\mathbf{s}_{i1}) \\ \vdots \\ j(\mathbf{s}_{iS}) \end{pmatrix} \qquad \mathbf{J} = B\mathbf{j} \tag{3.5}$$

Where B is an $S2^L \times S_i$ transformation matrix with rank S_i. I also write $J(\mathbf{s}) = B_{\mathbf{s}}\mathbf{j}$ using just the $2^L \times S_i$ sub-matrix $B_{\mathbf{s}}$. This matrix selects elements of $\mathbf{j}$ according to the possible ex-post states for player i, given they started the period in state $\mathbf{s}$. We can define the relationship $P(\mathbf{b}|\mathbf{s})^T B_{\mathbf{s}} = Q(\mathbf{b}|\mathbf{s})^T A_{\mathbf{s}} C$ for the $S \times S_i$ matrix C. entry mn of C is equal to 1 if $\mathbf{s}_i^m = \mathbf{s}_i^n$, and zero otherwise. Therefore, each row of C contains a single non-zero entry, while column n contains a 1 in rows for which $\mathbf{s}_i = \mathbf{s}_i^n$. This relationship holds because C collapses Q over states with the same $\mathbf{s}_i$.

[15]This differs from GKS' approach in which J_{ia} is able to depend on both $\mathbf{s}^a$, $\mathbf{s}$, and potentially even other $\mathbf{s}^{a'}$. For example, their approach permits identification of reference dependent preferences. However, in their empirical example they do impose this restriction. Furthermore, while I impose that j is independent of $\mathbf{s}_{-i}$, this restriction is not necessary for identification or estimation.

Based on these assumptions, I will prove the following proposition:[16]

Proposition 2. *Under assumptions 1 - 4, the model primitives F and $\mathbf{j}$ are non-parametrically identified up to β and $j(\mathbf{s}_{i1})$.*

As is standard, I take β as given. Meanwhile, we must normalise $j(\mathbf{s}_{i1})$ because the level of pay-offs is not-identified - only marginal pay-offs are identified. Just as, in the single-object case, we normalise the pay-off from losing to zero.

3.3.2 Optimising Behaviour

The agent's problem is to maximise their expected discounted pay-off, and so in each period the agent maximises the following object, with respect to $\mathbf{b}$:

$$\begin{aligned}\tilde{\Pi}(\mathbf{b}|v;\mathbf{s}) &= \Gamma(\mathbf{b}|\mathbf{s})^T(v-\mathbf{b}) + P(\mathbf{b}|\mathbf{s})^T J(\mathbf{s}) + \beta Q(\mathbf{b}|\mathbf{s})^T V(\mathbf{s}) \\ &= \Gamma(\mathbf{b}|\mathbf{s})^T(v-\mathbf{b}) + P(\mathbf{b}|\mathbf{s})^T B_{\mathbf{s}}\mathbf{j} + \beta Q(\mathbf{b}|\mathbf{s})^T A_{\mathbf{s}}\mathbf{V}\end{aligned}$$

Assumption 9 ensures that $P(\mathbf{b}|\mathbf{s})$, $Q(\mathbf{b}|\mathbf{s})$, and $\Gamma(\mathbf{b}|\mathbf{s})$ are continuously differentiable in $\mathbf{b}$. Necessary First Order Conditions of optimal bidding are then given as:

$$\underbrace{\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})}_{L\times L}\underbrace{(v-\mathbf{b}^*)}_{L\times 1} = \underbrace{\Gamma(\mathbf{b}^*|\mathbf{s})}_{L\times 1} - \underbrace{\nabla_{\mathbf{b}}P(\mathbf{b}^*|\mathbf{s})}_{L\times 2^L}\underbrace{B_{\mathbf{s}}\mathbf{j}}_{2^L\times 1} - \beta\underbrace{\nabla_{\mathbf{b}}Q(\mathbf{b}^*|\mathbf{s})}_{L\times n^L}\underbrace{A_{\mathbf{s}}\mathbf{V}}_{n^L\times 1} \tag{3.6}$$

As above, under the assumption of zero probability ties (or exogenous tie-breaking), $\Gamma_{il}(\mathbf{b}|\mathbf{s}) = \prod_{j\neq i} G_{jl}(b_{il}|\mathbf{s})$. Therefore $\nabla\Gamma$ must be a diagonal matrix with entry ll equal to $\sum_{j\neq i} g_{jl}(b_{il}|\mathbf{s})\prod_{k\neq j,i} G_{kl}(b_{il}|\mathbf{s})$, and so $\nabla\Gamma$ must be invertible for most $\mathbf{b}$.

3.3.3 The Inverse Bid System

I now consider how the first order conditions can be inverted to give the inverse bidding system. This enables me to show that F is non-parametrically identified,

[16] Importantly, this proposition does not collapse down to GKS' identification result, even when the game is static or bidders are myopic (e.g. $\beta = 0$). This proposition differs from GKS's key proposition in the source of variation, and exclusion restriction, used to identify J. They prove identification from variation in characteristics of non-bid on lots and characteristics of bidder i's competitors, which cause 'exogenous' variation in Γ and P. I focus on variation in the state variable, which creates variation in Γ and P but also directly creates variation in $J + \beta V$, all of which then cause variation in bids.

conditional on J and βV. This argument is almost precisely the same as that presented by GKS, which is a simple multi-object extension of Guerre, Perrigne, and Vuong (2000)'s identification result from inverting the first order conditions.

Proposition 2′. *Under assumptions 1 - 3, and conditional on J and βV being known, the cdf F is non-parametrically identified.*

Suppose both J, and βV are already identified The First Order Conditions can then be inverted to obtain the inverse bid function:[17]

$$\xi(\mathbf{b}^*|J,\beta V;\mathbf{s}) = \underbrace{\mathbf{b}^*}_{\text{observed}} + \underbrace{\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1}}_{Identified}[\Gamma(\mathbf{b}^*|\mathbf{s}) - \underbrace{\nabla_{\mathbf{b}}P(\mathbf{b}^*|\mathbf{s})}_{Identified}\, B_{\mathbf{s}}\mathbf{j} - \underbrace{\nabla_{\mathbf{b}}Q(\mathbf{b}^*|\mathbf{s})}_{Identified}\, A_{\mathbf{s}}\beta\mathbf{V}] \tag{3.7}$$

This system of equations is a natural extension of the standard inverse bid function. At the optimum the lot specific value is equal to bids $\mathbf{b}^*$ plus a lot specific markup $\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1}\Gamma(\mathbf{b}^*|\mathbf{s})$, minus a combination markup $\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1}\nabla_{\mathbf{b}}P(\mathbf{b}^*|\mathbf{s})B_{\mathbf{s}}\mathbf{j}$, minus the standard dynamic markup which depends on precisely who won each combination of lots $\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1}\nabla_{\mathbf{b}}Q(\mathbf{b}^*|\mathbf{s})A_{\mathbf{s}}\beta\mathbf{V}$.

We can evaluate this inverse bid function at the observed bids, which holds for a particular candidate $(J,\beta V)$. If this candidate $(J,\beta V)$ is correct, that is, if we have already identified $(J,\beta V)$, then $\boldsymbol{\xi}(\mathbf{b}^*|J,\beta V;\mathbf{s}) = \boldsymbol{v}$ trivially. From here it is simple to non-parametrically identify $F(.)$.

3.3.4 Identification of V

I now demonstrate that conditional on $\mathbf{j}$ being known, we can write $\mathbf{V}$ as a function of the distribution of bids and $\mathbf{j}$ only. This essentially extends Proposition 1 from JBP to the multi-object case.

Proposition 3. *Under assumptions 1 - 4, the expected stage pay-off is given by:*

$$\begin{aligned}\tilde{\Pi}(\mathbf{b}^*|\boldsymbol{v};\mathbf{s}) = {} & \Gamma(\mathbf{b}^*|\mathbf{s})^T\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1}\Gamma(\mathbf{b}^*|\mathbf{s}) \\ & + [P(\mathbf{b}^*|\mathbf{s})^T - \Gamma(\mathbf{b}^*|\mathbf{s})^T\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1}\nabla_{\mathbf{b}}P(\mathbf{b}^*|\mathbf{s})]B_{\mathbf{s}}\mathbf{j} \\ & + [Q(\mathbf{b}^*|\mathbf{s})^T - \Gamma(\mathbf{b}^*|\mathbf{s})^T\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1}\nabla_{\mathbf{b}}Q(\mathbf{b}^*|\mathbf{s})]A_{\mathbf{s}}\beta\mathbf{V}\end{aligned} \tag{3.8}$$

[17] For rows l such that $l \notin \mathbb{L}$ and hence $b_l = \emptyset$, we will also have $\xi_l = \emptyset$

Proof of this proposition is given in Appendix C. This relation is an extension of Proposition 1 presented in JP. The first term on the right hand side of the equation can be written as $\sum_l \frac{\Pi_{j\neq i} G_{jl}(b_{il})}{\sum_{j\neq i} g_{jl}(b_{il})}$ — the same term in JP's proposition. Unlike in the single unit case there is an additional adjustment for the non-additivity.

From Proposition 3, employing the identity $P(\mathbf{b}|\mathbf{s})^T B_{\mathbf{s}} = Q(\mathbf{b}|\mathbf{s})^T A_{\mathbf{s}} C$, and taking an expectation of the observed bids, we can write the ex-ante value function as:

$$V^e(\mathbf{s}) = \Phi(\mathbf{s}) + \Omega(\mathbf{s})[C\mathbf{j} + \beta\mathbf{V}]$$

$$\text{Where} \qquad \Phi(\mathbf{s}) = E_{\mathbf{b}}[\Gamma(\mathbf{b}^*|\mathbf{s})^T \nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1}\Gamma(\mathbf{b}^*|\mathbf{s})|\mathbf{s}]$$

$$\Omega(\mathbf{s}) = E_{\mathbf{b}}[Q(\mathbf{b}^*|\mathbf{s})^T - \Gamma(\mathbf{b}^*|\mathbf{s})^T \nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1}\nabla_{\mathbf{b}} Q(\mathbf{b}^*|\mathbf{s})|\mathbf{s}]A_{\mathbf{s}}$$

Stacking over $\mathbf{s}$ write the continuation value as $\mathbf{V} = T\mathbf{V}^e = T\Phi + T\Omega[C\mathbf{j} + \beta\mathbf{V}]$ Which we can invert for: $\mathbf{V} = (I_S - \beta T\Omega)^{-1}[T\Phi + T\Omega C\mathbf{j}]$. This yields a stationary solution for the continuation value. Non-singularity of $(I_S - \beta T\Omega)$ is a condition for stationarity, though this matrix is guaranteed to be non-singular anyway.[18] This ensures that, conditional on $\mathbf{j}$ being known, the continuation value is point identified.

3.3.5 Identification of j

I now prove that identification of $\mathbf{j}$ collapses to a rank condition. Impose the mean zero property of $\boldsymbol{v}$ for:

$$\begin{aligned}
0 = E_v[\boldsymbol{v}|\mathbf{s}] &= E_{\mathbf{b}^*}[\boldsymbol{\xi}(\mathbf{b}^*; \mathbf{s}, (\mathbf{j}, \mathbf{V}))|\mathbf{s}] \\
&= E_{\mathbf{b}^*}[\mathbf{b}^* + \nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1}\Gamma(\mathbf{b}^*|\mathbf{s})|\mathbf{s}] - E_{\mathbf{b}^*}[\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1}\nabla_{\mathbf{b}} Q(\mathbf{b}^*|\mathbf{s})|\mathbf{s}]A_{\mathbf{s}}[C\mathbf{j} + \beta\mathbf{V}] \\
&= \Upsilon(\mathbf{s}) - \Psi(\mathbf{s})[C\mathbf{j} + \beta\mathbf{V}] \qquad (3.9)
\end{aligned}$$

[18] Non-singularity follows from the matrix being strictly diagonally dominant. This result is known as the Levy-Desplanques Theorem. Strict diagonal dominance arises because every element of $T\Omega$ is weakly positive, and rows sum to 1. Therefore, off diagonals of $I - \beta T\Omega$ lie in the interval $(-\beta, 0]$, while diagonals are strictly positive, and rows sum to $1 - \beta$.

Stacking over $\mathbf{s}$, and substituting in the expression for $\mathbf{V}$, we get:

$$
\begin{aligned}
0 = \mathrm{Y} - \Psi[C\mathbf{j} + \beta\mathbf{V}] \\
= \mathrm{Y} - \beta\Psi(I_S - \beta T\Omega)^{-1}T\Phi - [\Psi C + \beta\Psi(I_S - \beta T\Omega)^{-1}T\Omega C]\mathbf{j} \\
= \mathrm{Y} - \beta\Psi(I_S - \beta T\Omega)^{-1}T\Phi - \Psi(I_S - \beta T\Omega)^{-1}C\mathbf{j}
\end{aligned}
\tag{3.10}
$$

This system of LS equations in $S_i - 1$ unknowns overcomes the standard order condition discussed in GKS. There exists a unique solution to this system ($\mathbf{j}$ is point identified) if and only if the $LS \times S_i$ matrix $\Psi(I_S - \beta T\Omega)^{-1}C$ has rank $S_i - 1$.

3.3.6 Rank Condition

This rank condition requires that observations of bidding behaviour, across all S states, produces sufficient information about $\mathbf{j}$ to uniquely pin down all $S_i - 1$ elements. We gain information about $j(\mathbf{s}_i)$ from how bidding behaviour changes when $\mathbf{s}_i$ is a possible outcome from the round of auctions. Stacking the moment conditions in equation 3.10 stitches together information about $\mathbf{j}$ across different state observations. As well as information as $\mathbf{s}_i$ varies, we gain information as $\mathbf{s}_{-i}$ varies - even if this is excluded from the function j — providing additional identifying variation. One additional assumption is sufficient for this rank condition to hold:

Assumption 5. The set $\mathbb{S}_i$ is partially ordered according to the strict partial ordering $\succeq$, such that if $\mathbf{s}_i' \in \mathbb{S}_i^a(\mathbf{s}_i, \mathbf{s}_0)$ then $\mathbf{s}_i' \succeq \mathbf{s}_i$. In addition, the maximal elements of $\mathbb{S}_i$ do not outnumber the non-maximal elements.

The partial ordering assumption is very mild, really only imposing the transitivity of partially ordered sets. A requirement for these partial orderings is that winning an auction is monotonic: one cannot gain an object from winning one auction and give it away by winning a different auction. I limit the number of maximal elements because observations of bidding from maximal elements are not informative.[19]

Proposition 4. *Under assumption 1 - 5 $\Psi(I_S - \beta T\Omega)^{-1}C$ has rank $S_i - 1$*

[19] An element $\mathbf{s}_i$ is defined as maximal if there does not exist an $\mathbf{s}_i' \in \mathbb{S}_i$ such that $\mathbf{s}_i' \succ \mathbf{s}_i$. The interpretation of how an element can be considered maximal is left ambiguous. One interpretation is that these maximal elements are the largest (in partial ordering terms) states that are observed as possible ex-post outcomes, but are never observed as ex-ante outcomes. In this way, we want to try to identify j for these states, but do not get to use observations beginning in these states.

Proof of this proposition is given in Appendix D. It is omitted from the main text for brevity. This rank condition is not trivial, since Ψ is certainly rank deficient. Likewise, it is not ex-ante obvious whether stacking $\Psi(\mathbf{s})$ across states yields information about every element of $\mathbf{j}$. The bulk of the proof requires establishing the rank of Ψ and finding it's null space. As we stitch together observations of bidding from each state, stacking $\Psi(\mathbf{s})$ across $\mathbf{s}$, the rank increases by at least two each time. I then consider the image of $(I_S - \beta T \Omega)^{-1} C$, proving that the only element in the intersection of this image and the null space of Ψ is the constant vector.[20]

While the algebraic proof is long, it has a relatively intuitive geometric argument behind it, focusing on the rank of Ψ. Consider a simple setting with $L = 2$ and $|\mathbb{S}_i| = 13$, and $\mathbb{S} = \mathbb{S}_i$ (rival states and common states do not vary). The partial ordering of these states is shown in Figure 3.1, demonstrating that $\mathbf{s}_0$ is the unique minimal state, and $\mathbf{s}_6 - \mathbf{s}_{12}$ are maximal states. Beginning in state $\mathbf{s}^0$, if the bidder wins nothing their ex-post state is $\mathbf{s}^0$. If they win one lot then $\mathbf{s}^a = \mathbf{s}^1$ or $\mathbf{s}^2$ depending on which lot. If they win both lots then $\mathbf{s}^a = \mathbf{s}^4$. From each state, possible ex-post states are shown by the dotted arrows (ignoring self-loops).

We want to identify the 12 marginal payoffs of moving from $\mathbf{s}^0$ to the other states. Each observation, beginning in a given state, yields two pieces of information from the two bids. However, there are always three marginal payoffs to identify.

These marginal payoffs are identified if, by adding two edges from each state (shown in red), we can connect the graph, so the adjacency matrix (and hence Ψ) has rank 12. It turns out that we can always do this, so long as the maximal states (for which we have no useful data) do not outnumber the non-maximal states. The fact that the set of possible ex-post states overlap is key to this argument — $\mathbf{s}^4$ is a possible ex-post state for $\mathbf{s}^0$, $\mathbf{s}^1$, and $\mathbf{s}^2$. Likewise, each non-maximal state has a unique maximal element of its ex-post set — in that only $\mathbf{s}^0$ has $\mathbf{s}^4$ as the maximal element of its ex-post set. This ensures that the two pieces of information we receive from each state are all informative.

[20]This proof only holds for the setting when the state space is finite. However the underlying idea extends to the case with infinite states: Even though the rank of an infinitely large matrix is undefined, it is clear how the logic of combining observations across states yields identification.

FIGURE 3.1: The Geometric Identification Argument

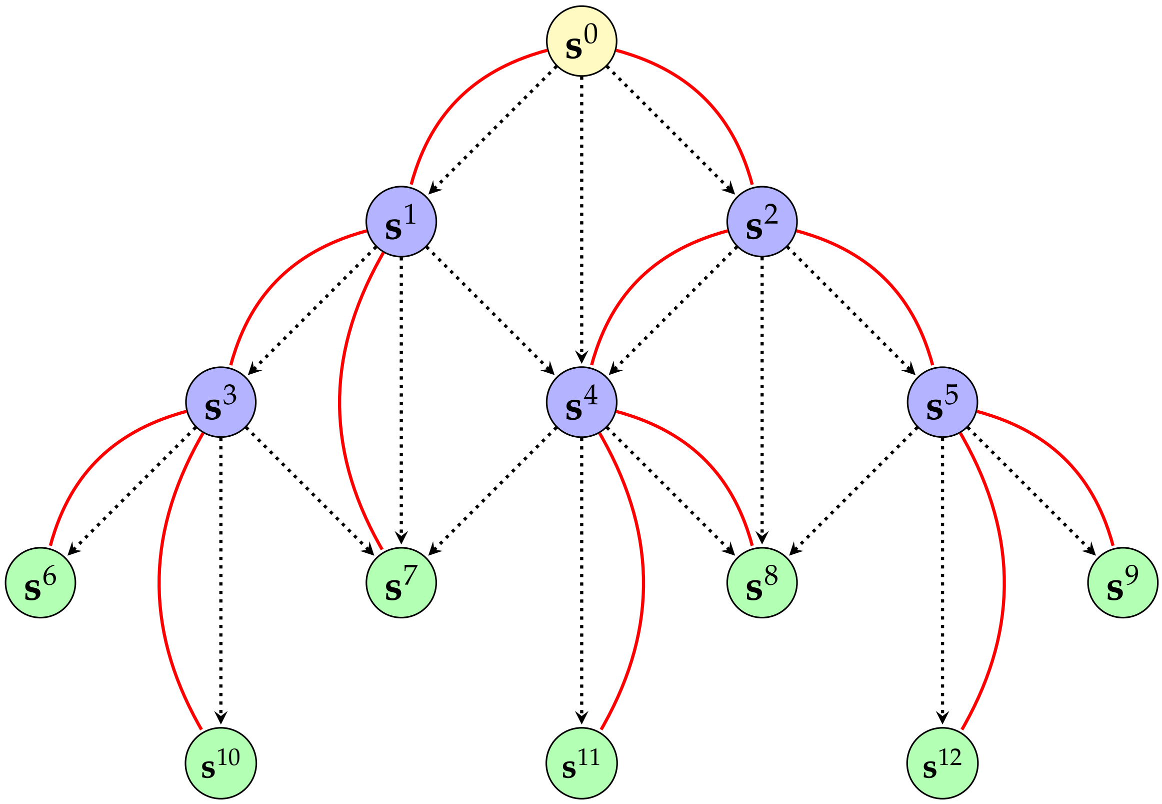

3.3.7 Extensions

I now consider several extensions to this identification framework. Some of these extensions are extremely mild, often relaxing some of the previously stated assumptions. Other extensions significantly alter the empirical and model framework.

Second-price auctions: In Appendix E.1 I demonstrate that identification extends, almost trivially, to simultaneous second-price auctions. Holding constant all but one bid, optimal bidding requires bidding the expected marginal value of winning that additional lot, given the other bids. This yields a set of necessary first order conditions similar to the first-price case. These can then be inverted for an inverse bid system that is identical to the first-price case, except that it is missing the markup term $\nabla_{\mathbf{b}}\Gamma(\mathbf{b})^{-1}\Gamma(\mathbf{b})$. Proposition 3 is then extended in a similar fashion. From here the identification and estimation arguments extend intuitively.

Binding reservation prices: In Appendix E.2 I consider how the presence of binding reservation prices impact identification. The binding reservation prices essentially cause censoring in the data. We immediately lose point identification of both F and j. However, F remains partially identified, using a similar argument as presented in subsection 3.3.3. We can no longer use moment conditions to identify j, as we did in subsection 3.3.5, and instead make used of quantile conditions. As a result, some elements of j can only be bounded. While reservation prices are a mathematical nuisance, they do not have a meaningful impact on identification, particularly when econometricians make parametric assumptions on F.

Endogenous Entry: In Appendix E.3 I consider an additional stage in-which the bidder choose a subset of auctions to enter, where entering each subset has an associated cost. This creates a minor change to the representation of V as a function of j. The identification of j and F then follow from previous arguments. The distribution of the entry cost distribution then follows from standard results.

Inter-temporal Budget Constraint: In Appendix E.4 I relax the assumption that pay-offs are quasi-linear in wealth and instead allow for an inter-temporal budget constraint, equivalent to the standard no-ponzi condition. I prove that in a stationary environment the quasi-linear model is observationally equivalent to the budget constraint model. This arises because the two models are observationally equivalent when when the marginal value of wealth is constant over time, and I demonstrate that under mild conditions stationarity requires that the opportunity cost of bidding in any given period must be time-invariant.

Stochastic Combination Value: In Appendix E.5 I allow the combination value to be stochastic. I consider two cases. First, when the combination value depends on the number of lots won, not which lots are won. Second, when the combination value is a function of low dimensional ($< L$) unobservables, such as unobserved states or unknown parameters. The only necessary restriction is that this function is strictly monotonic in the unobservables. In both cases I prove identification by inverting the first order conditions to identify the unobservables.

3.4 Estimation

Having established the non-parametric identification of the dynamic game, I now describe a computationally feasible procedure to estimate F and j.[21] I begin with

[21] Standard procedures are either infeasible or inapplicable. JP's method is inapplicable since we cannot write the maximised expected payoff as a function of bids only. Interior fixed point methods are

a general description of the novel procedure, outlining the key intuition. I then describe in detail the three estimation steps.

3.4.1 Premise

We cannot employ the estimation procedure of JP because we cannot write the value function, and hence the continuation value, as a function of the observed distribution of equilibrium bids and transition process *only*. However, as we saw with Proposition 3, we can write the continuation value as a function of the bid distribution, the transition process, *and* a correction term. This correction term, which simply corrects for the non-additivities across lots, is a known function of the bid distribution and the sum of the deterministic flow payoff function j and the discounted continuation value βV. This relationship is exemplified in Equation 3.11:

$$V(\mathbf{s}') = \int_{\mathbf{s}} \int_{\mathbf{b}} \Pi(\mathbf{b}|\mathbf{s}; K) dG(\mathbf{b}|\mathbf{s}) dT(\mathbf{s}|\mathbf{s}') \tag{3.11}$$

Where $\Pi(\mathbf{b}|\mathbf{s}; K) =$

$$\underbrace{\Gamma(\mathbf{b}|\mathbf{s})^T \nabla_{\mathbf{b}} \Gamma(\mathbf{b}|\mathbf{s})^{-1} \Gamma(\mathbf{b}|\mathbf{s})}_{\text{as in JP}} + \underbrace{[Q(\mathbf{b}|\mathbf{s})^T - \Gamma(\mathbf{b}|\mathbf{s})^T \nabla_{\mathbf{b}} \Gamma(\mathbf{b}|\mathbf{s})^{-1} \nabla_{\mathbf{b}} Q(\mathbf{b}|\mathbf{s})] K(\mathbf{s})}_{\text{correction term}}$$

$$\text{and} \quad [K(\mathbf{s})]_a = k(\mathbf{s}^a) = j(\mathbf{s}_i^a) + \beta V(\mathbf{s}^a)$$

The central premise of the procedure exploits the fact that if we 'knew' this object $k = j + \beta V$, we could plug this into the formula for the value function, and so continue as per JP's procedure. I refer to this sum as the 'pseudo-static' pay-off, as it is essentially the object we estimate if we incorrectly assume myopic bidding. Because both G and T are easily estimated using standard methods, if we had a consistent estimate for the function $k : \mathbb{S} \to \mathbb{R}$, then we would have a consistent estimate for V,

infeasible due to the computation difficulty of repeatedly numerically maximising expected payoffs. Finally, CCP style methods are applicable, but are more computationally expensive than the approach presented below. In the most abstract case presented in the previous section, fully non-parametric with a finite set of observed states, mine and the CCP approach are equivalent.

and hence j $(= k - \beta V)$. Like CCPs or the distribution of equilibrium bids, this function $k(.)$ is not a model primitive but a function of primitives. The central estimation problem then concerns estimating k.

This procedure is closely related to the estimation of a static model. If players are myopic, so that $\beta = 0$, then $k = j$. Therefore, the procedure is almost like estimating the model *as if* it were static. While we may expect that j is independent of $\mathbf{s}_{-i}$, k in general is not. Therefore, the procedure involves estimating the model as if it were a generalised static model, in which pay-offs are allowed to depend on every element of the state space that enters the continuation value.[22]

In part, the procedure is a generalisation of JP, since we write the Continuation Value as a function of the distribution bids and this additional combinatorial term. When k is additively separable $\Pi(\mathbf{b}|\mathbf{s};K) = \Gamma(\mathbf{b}|\mathbf{s})^T \nabla_{\mathbf{b}} \Gamma(\mathbf{b}|\mathbf{s})^{-1} \Gamma(\mathbf{b}|\mathbf{s})$, just as in JP. The procedure is also related to CCP methods. Both CCP methods and JP's procedure involve writing the continuation value as a function of the distribution of observed actions. The pseudo-static pay-off function, if it is identified from standard choice data, will itself be a function of the distribution of observed actions.[23] Therefore the continuation value is still a function of the distribution of observed actions, however it requires the additional intermediate step of estimating k.

Finally, in a discrete choice setting, this estimation procedure is numerically equivalent to conditional choice probability (CCP) estimation when the researcher fits a parametric functional form to the conditional choice probabilities. For example, in a binary choice setting one might fit a logit link function, so that $p(d_t = 1|\mathbf{s}_t) = \frac{\exp(\delta \mathbf{s}_t^1)}{\exp(\delta \mathbf{s}_t^0) + \exp(\delta \mathbf{s}_t^1)}$. This is equivalent to specifying the choice specific value function $U(d_t = a|\varepsilon_t^a, \mathbf{s}_t) = \varepsilon_t^a + j(\mathbf{s}_t^a) + \beta V(\mathbf{s}_t^a)$, where ε_t^a is Type 1 extreme value, and $j(\mathbf{s}_t^a) + \beta V(\mathbf{s}_t^a) = k(\mathbf{s}_t^a) = \delta \mathbf{s}_t^a$. Then we have $V(\mathbf{s}_t) = E[\ln(\exp(k(\mathbf{s}_{t+1}^1)) + \exp(k(\mathbf{s}_{t+1}^0)))|\mathbf{s}_t] + \gamma$, where γ is Euler's constant. The difference is that we give the reduced form coefficients δ a structural interpretation. This procedure only differs from a CCP estimator in an environment with choice over lotteries with multiple possible outcomes.

[22] This permits a simple test of forward looking behaviour and a way to identify β following the 'exclusion restriction' logic of Magnac and Thesmar (2002): If the model is correctly specified and $\mathbf{s}_{-i}$ is excluded from j, then observing that k varies with $\mathbf{s}_{-i}$ is sufficient to reject myopia.

[23] In this application, identification of k just uses the assumption that $E[v|\mathbf{s}] = 0$. Point Identification is then ensured by proposition 4, ensuring the first line of equation 3.10 is invertible, as the matrix Ψ is invertible up to normalisations.

I now outline the three key estimation steps. The discussion is kept general within the dynamic multi-object auction framework. The procedure can be written succinctly as the following three step procedure:

1. Estimate equilibrium win probabilities Γ and Q, and transition functions T.
2. Given $\hat{\Gamma}$ and $\hat{Q}$, estimate F and k - the primitives of the pseudo-static model.
3. Given $\hat{\Gamma}$, $\hat{Q}$, $\hat{T}$, $\hat{F}$, and $\hat{k}$, evaluate $\hat{V}$ then back out $\hat{j}$

3.4.2 3 Steps

Step 1.

The First Stage represents the standard first stage in the empirical auction literature. There are several possible approaches the research might take. The researcher could estimate the conditional cumulative distribution of bids G_{il}, then form $\Gamma_{il}(b)$, $P_{il}(b)$, and $Q_{il}(b)$ respectively. This is the approach taken in both GKS and JP. Otherwise the researcher may directly estimate Γ by estimating the probability that i wins any given lot given their bid, estimating $P(\max_{i' \neq i} \{b_{i'}\} \leq b|b)$, essentially estimating the distribution of highest bids other than bids belonging to i. This is the approach taken in Cantillon and Pesendorfer (2007) and Raisingh (2021).

The researcher can easily take either a parametric or non-parametric approach, so long as the ensuing estimated object is continuously differentiable. That is, the researcher must continue to impose the assumptions required for identification. The same goes for estimating the transition distribution $T_{\mathbf{s}}(.|\mathbf{s}^a_{t-1})$.

Step 2.

In the second step we estimate the sum of the immediate combination value and the discounted continuation value; the pseudo-static pay-off function $k(\mathbf{s})$. We can broadly follows the second stage in the estimation procedure presented in GKS. This step will be done by exploiting the moment conditions used in the identification argument: $E[\boldsymbol{v}|\mathbf{s}] = 0$. Set $\hat{k}$ such that, for all l and all $\mathbf{s}$, $E[\xi_l(\mathbf{b}^*|k;\mathbf{s})|\mathbf{s}] = 0$.

However, the identifying conditions gives rise to an additional set of moments:

$$E[v_l\mathbf{h}(\mathbf{s})] = E[E[v_l\mathbf{h}(\mathbf{s})|\mathbf{s}]] = E[E[v_l|\mathbf{s}]\mathbf{h}(\mathbf{s})] = 0$$

Where $\mathbf{h}(\mathbf{s})$ is a known vector valued function of $\mathbf{s}$, such as $\mathbf{h}(\mathbf{s}) = \mathbf{s}$ or $\mathbf{h}(\mathbf{s}) = \mathbb{I}[\mathbf{s} = \bar{\mathbf{s}}]$ (dummy variables for each possible state). That is, we can also just set $\hat{k}$ to ensure that $\xi_l(\mathbf{b}^*|k;\mathbf{s})$ is mean independent of $\mathbf{s}$. These moment conditions help us interpret the estimation problem as an instrumental variables procedure with $\mathbf{h}(\mathbf{s})$ as the instruments. Rewrite the FOCs as a regression equation:

$$\underbrace{b_{lt} + \frac{\Gamma_l(b_{lt}|\mathbf{s}_t)}{\nabla_{b_l}\Gamma_l(b_{lt}|\mathbf{s}_t)}}_{y_t} = \underbrace{-[\frac{\nabla_{\mathbf{b}}Q(\mathbf{b}_t|\mathbf{s}_t)}{\nabla_{b_l}\Gamma_l(b_{lt}|\mathbf{s}_t)}]_{l.}K(\mathbf{s}_t)}_{\mathbf{x}_t\beta} + v_{lt}$$

We could estimate $K(\mathbf{s}_t)$ using a least squares procedure; set k to minimise the sum squared residuals $\sum_t \sum_l v_{lt}^2$. But, in general $E[v_{lt}[\frac{\nabla_{\mathbf{b}}Q(\mathbf{b}_t|\mathbf{s}_t)}{\nabla_{b_l}\Gamma_l(b_{lt}|\mathbf{s}_t)}]_{l.}] \neq 0$ because $E[v_{lt}b_{l't}] \neq 0$. Therefore we have a standard endogeneity problem. Fortunately, we have a set of candidate instruments, $\mathbf{h}(\mathbf{s}_t)$, so we can write the first stage as:

$$\underbrace{-[\frac{\nabla_{\mathbf{b}}Q(\mathbf{b}_t|\mathbf{s}_t)}{\nabla_{b_l}\Gamma_l(b_{lt}|\mathbf{s}_t)}]_{l.}}_{\mathbf{x}_t} = \pi_l \underbrace{\mathbf{h}(\mathbf{s}_t)}_{\mathbf{z}_t} + \varepsilon_{lt}$$

This interpretation is helpful as it enables the researcher to make use of a number of standard instrumental variable procedures, such as analysing the relevance and validity of our instruments.[24]

The econometrician is able to estimate k either parametrically or non-parametrically. In the Monte-Carlo study presented below I consider both a non-parametric spline estimator, as well as a semi-parametric specification, assuming K takes a convenient linear-in-parameters form, with $k(\mathbf{s}) = \mathbf{h}(\mathbf{s})^T\theta^k$, and estimate θ^k using Generalised Method of Moments, imposing $E[v_l\mathbf{h}(\mathbf{s})] = 0$. Having estimated k we back out the distribution F using the empirical cdf of inverse bids ξ.

Step 3.

The Third Stage broadly corresponds to the second stage in JP's estimation procedure. Given the distribution of bids and the pseudo-static pay-off function K; use these objects to evaluate the continuation value.

[24]Depending on the form of $\mathbf{h}$ the first stage may not need to be estimated and π_l may be known.

Building on Proposition 3 we use equation 3.11 to evaluate the continuation value, evaluated using estimates of T, G, and K. First, we evaluate $\hat{\Pi}(\mathbf{b}_t|\mathbf{s};\hat{K})$ for each t. Next, they evaluate the ex-ante value function by either numerically integrating over the estimated distribution of bids or by taking a conditional expectation over observed bids. This latter approach ensures the distribution of bids never has to be explicitly estimated. The continuation value is then formed by taking a conditional expectation over the transition process, which can again be performed using numerical integration or by averaging over the observed distribution.

Finally, $\hat{j}$ can be backed out using the identity $\hat{j} = \hat{k} - \beta\hat{V}$ for some fixed value of β.[25] At this stage we must average our $\hat{j}$s over $\mathbf{s}_{-i}$. With a correctly specified model and an infinite amount of data there should be no variation. The choice of weighting procedure is important, since it will be necessary to place less weight on estimates for $\mathbf{s}_{-i}$ with fewer observations, so are more imprecisely estimated.[26]

Inference

Inference in this estimation procedure is complicated by the multiple stage nature. Inference on the first two stages is standard, making use of two step variance estimates for the second stage. Inference in the second stage enables hypothesis testing on model specification. For example, testing for the presence of non-additivity across lots, justifying the need for the dynamic multi-object approach.

In the third stage we use the delta-method, taking as given the first two stages. Bootstrap and Monte-Carlo methods are also possible, though may be time intensive. In some settings it may be appropriate to employ the same linear-in-parameters specification for both the second and third stages, as I do in the Monte-Carlo study below. This makes inference in the third stage extremely simple, as $j(\mathbf{s})$ inherits the same functional form as k and V. The variance of $\hat{j}$ is a function of the variances from the second and third stages, and their covariance.

[25]Note that at this stage one can vary β to see how $\hat{j}$ varies. In the spirit of Magnac and Thesmar (2002) β is identified from our exclusion restrictions on j, so in principle we *could* set β such that $\hat{j}$ is independent of $\mathbf{s}_{-i}$. Discussion of this possibility is left for future work.

[26]This weighting problem is related to the standard small sample problem for CCP methods, which is generally alleviated by reducing the weight on observations with poorly estimated continuation values. Similarly, like how these problems can be alleviated by taking a parametric approach to estimating conditional choice probabilities, a parametric approach to estimating k can alleviate these problems.

3.4.3 Monte Carlo Study

Next, I discuss the efficacy of the proposed estimator, presenting the results of a Monte-Carlo study. As discussed in GKS, the difficulty in simulating these types of games is that solving for equilibrium bidding strategies is intractable. Meanwhile, numerically finding equilibrium bidding strategies - by iterating over equilibrium beliefs and actions until a fixed point is found - is extremely computationally intensive. This is because, for each hypothesised set of beliefs, we must find the equilibrium continuation value through value function iteration, which requires numerically finding optimum bids a large number of times.

For simplicity I focus on the case where bidders are bidding against a parametric set of beliefs. That is, I essentially take the equilibrium as given. Furthermore, in the current simulations, I focus on an equilibrium in which equilibrium beliefs do not depend on each bidder's individual states $\{\mathbf{s}_{it}\}_{i \in \mathbb{N}}$. This is similar to many applications seen in practice, including GKS, Backus and Lewis (2016), Groeger (2014), Balat (2013).

Set up

Every period there are two auctions ($L = 2$) and two types of object, denoted x and y. Each lot contains one type of object, and one lot of each type of good comes to auction each period. However, some lots contain ten units of the good, while other lots contain only 5 units. The set of available lots is denoted (z^x, z^y), so lot 1 consists of z^x units of lot x, and lot 2 consists of z^y units of lot y. Each period the possible set of lots $\mathbb{L}_t \in \{(5,5),(10,5),(5,10)\}$. In this simplified setting the common state is just given by $\mathbb{L}_t$. For simplicity, this transitions stochastically such that each of these states occurs with equal probability, independent of previous states. This ensures that the continuation value does not depend on the common state.

Bidders' states consist of stocks of the two objects, both of which only come in integer values: $s^x_{it} \in \{0, 1, \ldots, 100\}$, likewise for good y. At the end of each period, bidders consume 3 units of good x with probability 0.4 and three units of good y with probability 0.3, until their stocks fall to 0 for either good. A bidder's combinatorial flow pay-off is given by:

$$j(s^x, s^y) = \theta_1 \log(s^x + 1) + \theta_2 \log(s^x + 1) \log(s^y + 1)$$

where (θ_1, θ_2) are parameters, which I set to 20 and 10 respectively. θ_1 ensures pay-offs are not additively separable over time, while $\theta_2 > 0$ ensures the two goods, and so the two lots, are complements within a period. I can check that bidding behaviour varies strongly with $\mathbf{s}_{it}$, ensuring this baseline instrument is strong. Meanwhile, the lot-specific pay-offs are drawn from:

$$\boldsymbol{v}_{it} \sim N \begin{pmatrix} 0 & & 900z_t^x & 100z_t^x z_t^y \\ 0 & , & 100z_t^x z_t^y & 400z_t^y \end{pmatrix}$$

I take as given the equilibrium distribution of the highest rival bids, which is assumed to follow a type 2 extreme value distribution. The mean of this distribution is given by the average (across states) marginal pay-off from winning each lot ($\approx (17.1z^x, 12.5z^y)$). The standard deviation is tuned to the standard deviation (across states and lot-specific pay-offs) of the marginal pay-offs from winning each lot. The shape parameter is set to 0.1.

I perform value function iteration to find the continuation value under this distribution of pay-offs and these equilibrium beliefs. Having found a continuation value, I can then simulate a dataset. Given the set-up the state space consists of $30,000$ unique elements. Focusing on a large number of elements is intended to simulate my real world application when the state space will be treated as continuous.

I simulate $1,000$ datasets of bids and states, each with $N \in \{ 1000, 10,000, 100,000 \}$ observations uniformly sampled from the state space. I then consider 3 estimators: $1)$ a parametric estimator using the same functional form as j, and $2)$ a quadratic polynomial estimator, and $3)$ a fully non-parametric cubic spline estimator. For the spline, I use uniformly spaced knots, setting the number of knots to ensure at least 30 observations per knot. For each estimator I consider estimates from using no instruments, the baseline "initial state" instruments, as well as using all the possible ex-post states as instrument.

Results

Results are presented in figure 3.2. Each estimator yields estimates of $\hat{j}(\mathbf{s}_i)$ for each $\mathbf{s}_i \in \mathbb{S}_i$. I then fit the correctly specified j across these states, extracting $\hat{\theta}_1$ and $\hat{\theta}_2$.

The non-parametric estimator (3) outperforms the two semi-parametric estimators, even in relatively small samples. However, it is very computationally intensive, with estimation taking almost 100 times longer than the semi-parametric estimators. Semi-parametric estimator (1), which fits the true functional form of j to both k and V, performs poorest. This is because we should not expect either k or V to inherit the functional form of j. Likewise, estimator (2), the flexible polynomial, performs reasonably well despite being misspecified. The choice of instruments is also shown to be important. For the most part, using no instruments $(\emptyset)$ out performs the initial state instrument. This arises for the combination of two reasons. First, as discussed previously, the initial state instruments may suffer from weak instrument problems, as variation in the initial state may not induce enough variation in bidding behaviour. Second, the degree of bias in the least squares estimation is expected to be small, depending on the correlation between $\Gamma_l(b_l)$ and $v_{l'}$. This correlation is relatively small because b_l varies much more with other variables, such as v_l and the state variables. Finally, using the ex-post states as instruments performs much better, but does not dominate (nor is dominated by) the no-instrument estimator.

3.5 The Way Forward

In this chapter I did three things: First, I set-up an dynamic model of bidding in repeated rounds of simultaneous first-price auctions. Second, I proved that the model primitives are identified from standard bidding data, using observed variation in state variables. Finally, I proposed a computationally feasible estimation procedure, overcoming the technical challenges of estimating model primitives in this setting.

While this chapter was primarily motivated by the study of Feeding America's allocation mechanism, it was also motivated by the prevalence of such repeated, multi-object auctions. Significant complementarities between auctioned objects have been found in both the dynamic single-object literature, and the static multi-object

FIGURE 3.2: Monte Carlo Study

Instrument =			$\emptyset$			$\mathbf{s}_t$			$\{\mathbf{s}_t^a\}$		
	θ	N	Mean	SD	rMSE	Mean	SD	rMSE	Mean	SD	rMSE
(1)	θ_1	1,000	5.79	5.11	15.1	4.46	6.11	16.7	4.74	5.64	16.3
		10,000	6.19	3	14.1	4.9	3.54	15.5	4.8	3.25	15.5
		100,000	6.57	2.63	13.7	5.56	2.95	14.7	5.5	2.75	14.8
	θ_2	1,000	5.76	0.596	4.28	6.71	0.775	3.38	6.33	0.733	3.75
		10,000	5.83	0.347	4.19	6.76	0.436	3.27	6.41	0.401	3.61
		100,000	5.98	0.274	4.03	6.9	0.374	3.12	6.53	0.34	3.49
(2)	θ_1	1,000	24.2	1.79	4.55	24.3	4.74	6.38	22.8	4.63	5.4
		10,000	24.2	0.602	4.27	24.7	1.52	4.89	22.9	1.49	3.23
		100,000	24.2	0.301	4.25	24.6	0.529	4.64	22.9	0.52	2.98
	θ_2	1,000	12.1	0.289	2.12	11.2	0.717	1.38	12.1	0.692	2.23
		10,000	12.1	0.0963	2.15	11.1	0.224	1.16	12.1	0.227	2.14
		100,000	12.2	0.0458	2.17	11.2	0.0803	1.16	12.1	0.0802	2.15
(3)	θ_1	1,000	20	2.04	2.04	20.7	14.9	14.9	19	3.73	3.85
		10,000	21.2	1.06	1.58	22.4	4.07	4.71	20.4	1.54	1.59
		100,000	21.3	0.413	1.32	22.5	1.19	2.74	20.4	0.48	0.653
	θ_2	1,000	10.5	0.306	0.558	10.5	2.17	2.22	10.7	0.587	0.928
		10,000	10	0.157	0.157	9.93	0.614	0.617	10.1	0.257	0.29
		100,000	10	0.0481	0.0486	9.92	0.172	0.189	10.1	0.0741	0.159

Note: The true values for θ_1 and θ_2 are 20 and 10 respectively. For each study I use 1,000 simulated datasets. The three instruments are: $\emptyset$ = no instrument (OLS), $\mathbf{s}_t$ = initial states, $\{\mathbf{s}_t^a\}$ = all the possible ex-post states, given the period began in $\mathbf{s}_t$. Estimator (1) is a semi-parametric estimator, using the true functional form of j to fit k and V. Estimator (2) fits a cubic polynomial, while Estimator (3) fits a cubic spline.

literature, most notably in JP and GKS. However, these two types of model had not, until this point, been unified in a single framework.

The Feeding America setting differs from the simplified model presented above in a number of ways. The inter-temporal budget constraint, the reservation prices, as well as endogenous entry are all factors that I analysed separately in Appendix E. I showed that none of these factors undermined the key identification arguments. However, moving forward, a key difficulty is that the identification strategy outlined above relies on observed variation in state variables. In the Feeding America setting, I do not observe bidders' state variables, as I do not have data on food banks' stocks. Therefore, one remaining methodological challenge is to extend these identification arguments and my estimation procedure to allow for these unobserved states.

Chapter 4

A Structural Analysis of the Food Allocation Problem

4.1 Chapter Introduction

In this chapter I use a rich model of food bank bidding behaviour to investigate welfare under the Choice System, compared to alternative mechanisms that allow food banks varying degrees of choice. The central challenge is that I do not observe food banks' inventories — a key determinant of bidding behaviour. I do, however, observe food banks' previous winnings, which act as observed shifters of the unobserved stocks, and I prove that this is sufficient to non-parametrically identify the model. I then develop a novel empirical strategy to estimate food banks' demand functions despite not observing their inventories, applying a dynamic auction model to detailed Choice System data. I exploit the panel dimension of the data to allow demand to vary across food banks and over time, as different food banks have different storage capacities, cater to different numbers of clients, and receive different types of food at different times from their local donors. I then use these estimates to evaluate equilibrium allocations under a number of alternative allocation mechanisms. Counterfactual simulations demonstrate that, relative to the Old System, the Choice System is extremely effective at achieving Feeding America's welfare and equity goals: Welfare is 17.1% higher under the Choice System than the Old System. This is roughly equivalent to an additional 50 tons of food allocated each day, enough to support an additional 22,300 people.

In order to investigate food banks' needs, and so evaluate welfare under various allocation mechanisms, I first develop a structural model of food banks bidding for food on the Choice System. The structural model follows the framework outlined in Chapter 3, unifying the models of Gentry, Komarova, and Schiraldi (2023) and Jofre-Bonet and Pesendorfer (2003). As we saw in Chapter 2, descriptive evidence demonstrates the need for both the dynamic and multi-object framework: First, when multiple similar loads are auctioned simultaneously food banks are less likely to bid on any given load. This suggests that similar loads are substitutable, and requires a multi-object model to account for the simultaneous auction environment. Second, food banks certainly act as forward looking bidders, given that auctions happen so frequently. Conditional on winning a load, food banks are less likely to bid on similar loads on subsequent days. This suggests food banks treat loads as storable goods subject to storage costs, emphasising the need for an empirical model that accounts for the dynamic environment.

As established by Prendergast (2017), the importance of choice depends on the degree of unobserved heterogeneity in food banks' preferences and storage costs, as well as the degree of substitutability of different types of food. In Section 2.5 I presented evidence of this heterogeneity, demonstrating large differences in bidding behaviour across different types of food, across different food banks, and within food banks over time. The model incorporates this heterogeneity in three key ways. First, food is classified by how it is stored (capturing storage capacity), and how it is used. It is further divided into 15 broad categories and 164 subcategories, each associated with distinct preference parameters. Second, the long panel (around 900 days) allows me to estimate distinct parameters for each food bank, allowing for permanent heterogeneity across food banks.[1] Finally, I allow for time-varying unobserved heterogeneity, which I attribute to the fact that I do not observe food banks' stocks of various types of food. This captures how food banks may irregularly receive donations from their local donors and irregularly give out food to clients.

The main challenge is that I do not observe food banks' stocks. Current stocks

[1] I lack identifying variation for certain food banks which rarely bid or rarely win. This is particularly a problem for estimating food bank specific storage costs. Therefore I build this heterogeneity into a Bayesian Hierarchical framework - assuming that individual parameters are themselves drawn from a parent distribution. This approach is natural as it allows food banks to differ in their ability to store various types of food, while pooling information for those food banks for whom identifying variation is scarce.

are a key determinant of demand — if a food bank suddenly stops bidding on a particular type of food it might be because, unobserved by the econometrician, they recently received this from a local donor. I make two important methodological contributions: First, I prove that variation in food banks' winnings are sufficient to non-parametrically identify the model primitives. Second, I develop a procedure to estimate bidders' values in a dynamic multi-object auction environment when individual state variables (stocks) are unobserved. I overcome this problem using a Gibbs Sampling procedure, employing a data-augmentation step to draw the unobserved stocks from their conditional posterior distribution. To the best of my knowledge, this is the first study to estimate a model of this type.

The model presented in Chapter 3 was identified using observed variation in individual state variables. In contrast, this model is identified through observed variation in food banks' winnings, which drive systematic variation in bidding behaviour. Winnings act as observed changes in the unobserved stocks, acting like an instrumental variable that shifts these stocks in a known way. Expanding on the framework used by Hu and Shum (2012) I prove that variation in food banks' winnings are sufficient to non-parametrically identify the primitives of the dynamic auction model, as well as the transition process of food banks' stocks. The idea behind this argument is that the change in the propensity to bid immediately after winning a lot identifies food banks' storage capacities: After a recent win, capacity constrained food banks will stop bidding on that type of food. Meanwhile, the length of time before food banks return to their average bidding propensity enables identification of the unobserved state transition process: If it takes them a long time to return to bidding on a particular type of food, this suggests they generally have access to that food from their local donors. This is essentially the same variation shown in Figure 2.11 in Chapter 2. This identification argument has broader applicability beyond the auction framework, for example, in the storable goods literature.

I employ the three step estimation procedure introduced in Chapter 3. In the first step I estimate equilibrium beliefs by estimating the conditional distribution of winning bids. I then invert food banks' first order conditions for optimal bidding, obtaining an inverse bidding system as in Guerre, Perrigne, and Vuong (2000) and

Gentry, Komarova, and Schiraldi (2023). In the second step, using the inverse bidding system, I estimate the distribution of food banks' 'Pseudo-Static' payoffs from winning combinations of lots. This means I estimate the sum of bidders' flow payoff and their discounted continuation value — essentially estimating the model as though food banks were myopic. During this step I also estimate the transition process for food banks' stocks. I employ the Gibbs Sampling procedure in this step to efficiently sample stocks from their conditional posterior distribution. This means I essentially iterate between 1) estimating food bank demand functions given a sample of observed stocks, and 2) sampling stocks from their posterior, given the food banks' demand functions, and so estimating the stock transition process. Finally, in the spirit of Jofre-Bonet and Pesendorfer (2003), the continuation value can be written as a function of observed bids, beliefs, and the pseudo-static payoff function. Therefore, in the third step I evaluate the estimated continuation value, before backing out the distribution of flow payoffs from the definition of the pseudo-static payoffs.

I find significant evidence of demand heterogeneity both across food banks and over time. I estimate large differences in access to local donors. For example, I find that food banks in urban areas generally have little access to fresh food, such as produce, so get much of it from Feeding America. This contrasts with more rural food banks which rarely need fresh food from Feeding America. This means it is important for food banks to be able to sort across types of food. Likewise, some food banks' local donations are estimated to be very variable over time, meaning it is important for them to be able to pick and choose different types of food as and when they are most needed. I estimate that day to day variation in stocks account for 45% of the daily variation in bidding behaviour. The model also suggests that food banks go through extended periods with high stocks, during which they very rarely place bids, and periods with low stocks, during which they bid very frequently. Therefore, over the long-run, I estimate that 72% of the variation in bidding behaviour can be attributed to unobserved variation in stocks.

Using the estimated model I consider equilibrium allocations under a number of alternative mechanisms that permit food banks varying degrees of choice. First, I consider the mechanism previously employed by Feeding America (The 'Old System'). This allows me to quantify and qualify the benefits of choice and the Choice System over lack of choice, building on the evidence presented in Prendergast (2017) and Prendergast (2022). I find that welfare is 17.1% higher under the Choice System than under the Old System. The majority of this welfare gain is due to food banks having more control over their stocks, better tailoring their allocations to fit their most pressing needs first. This is as opposed to accepting sub-optimal food when they face significant storage costs; food that may be used more effectively by another food bank at that point in time. As a result, around 85% of food banks are estimated to be better off under the Choice System.

Feeding America's allocation problem is faced by numerous other food bank networks around the world, such as the European Federation of Food Banks (FEBA), and Food Bank Australia. Therefore I also consider mechanisms employed (often implicitly) by some of these other food bank networks.[2] A mechanism that offers food only to the nearest food bank, aiming to minimise transportation costs but allowing food banks even less choice than the Old System, achieves only 65% of the welfare under the Old System. This result arises because even under the Old System each load was offered to multiple food banks. Many food bank networks, including the Trussell Trust in the U.K., implicitly use this mechanism by not allocating food centrally and instead linking food banks up with nearby donors. However, even if food is offered to every food bank (in order of distance), this is only marginally better than the Old System, but much worse in its distributional effects — food banks which happen to be well situated consume the most valuable food. This is because I estimate that transportation costs are not a large cost factor for most food banks.

The Choice System allocates food simultaneously in batches, rather than allocating food as donations arrive. Among other benefits, this 'batching' ensures food banks have information about all the food being allocated on a given day when making decisions, giving them more control over their allocations. The majority of other

[2]Other food bank networks often face somewhat different problems, in both scale and scope, to Feeding America. For example, transportation cost are known to be a larger factor in Australia. Likewise, FareShare (U.K.) face an allocation problem closer to an individual food bank allocating food to its associated food pantries. Therefore the results from this paper cannot be exactly applied to these other settings. That said, certain broader lessons are still valuable to these organisations. Future work, ideally using data from these other settings, is certainly needed.

mechanisms employed for food allocation are sequential in nature.[3] I simulate an 'efficient' sequential mechanism and find that welfare is still around 12% lower than under the Choice System. This is because, while food is always allocated to the food bank that values it most, food banks are not always allocated the type of food they need the most. Then, when a donation does come along that they really want, they no longer have capacity to store it. This is essentially the same effect driving the poor welfare results for the Old System.

The chapter proceeds as follows: Section 4.2 outlines the empirical model of food bank bidding behaviour and discusses non-parametric identification of the model primitives. Section 4.3 describes the estimation procedure and parametric assumptions employed for the structural model, while Section 4.4 details the estimation results. Section 4.5 details the counterfactual mechanism considered before presenting the simulation results.

4.2 A Model of Food Banks

4.2.1 Rules

I now present the empirical model of food banks bidding in the Choice System. Section 4.2.1 introduces the market environment and the model primitives. Section 4.2.3 introduces the food banks' dynamic optimisation problem. Section 4.2.4 discusses the Markov Perfect Equilibrium and stationarity in this dynamic context. Finally, section 4.2.5 discusses Identification. Assumptions necessary for identification and feasibility of estimation are introduced as and when they are needed.

Each period t, over an infinite horizon, N food banks compete in up to L First-Price Sealed-Bid auctions. Food banks are denoted by i and lots are denoted by l. a is used to denote the combination outcome from a round of auctions. That is, which combination of lots food bank i won.

[3]Theoretical results suggest this is suboptimal (Akbarpour, Li, and Gharan (2020), Baccara, Lee, and Yariv (2020)).

Actions

Players simultaneously choose which lots to enter and what to bid. Entry decisions consist of an L dimensional vector $\mathbf{d}_{it}$. Entry $d_{itl} = 1$ if they enter lot l, $d_{itl} = 0$ otherwise. Each player plays an L dimensional vector of bids each period, denoted $\mathbf{b}_{it}$, with $b_{itl} = \emptyset$ if $d_{itl} = 0$. Bids must weakly exceed the reservation price, so that $b_{itl} \geq R_{tl}$ if $d_{itl} = 1$. Auctions are costless to enter.

Outcomes

Winners are announced simultaneously. Winners pay their bids, and every player observes the identities and bids of winners. Define player i's individual outcome vector $\mathbf{w}_{it}$ as the $L \times 1$ vector such that $w_{ilt} = 1$ if food bank i won lot l at time t, and zero otherwise. Ex-ante hypothetical outcomes are denoted by $\mathbf{w}^a_{it}$.

Lots and lot characteristics

Each period up to L lots come to auction. Each available lot l is characterised by a row-vector of characteristics $\mathbf{c}_{tl}$, consisting of the the location, size, categories (c), subcategories (h), and storage method (g) of the lot. The number of pounds in each lot from each category/subcategory/storage method is denoted by $\{\mathbf{z}^c_{tl}, \mathbf{z}^h_{tl}, \mathbf{z}^g_{tl}\}$, so that if a food bank wins lot l their stock of food from each category increases by $\mathbf{z}^c_{tl}$. For notational convenience I absorb these variables into the common state variable $\mathbf{s}_{0t}$. I make the following assumption about the common state variables:

Assumption 6. $\mathbf{s}_{0t}$ *follows an exogenous Markov process, drawn from* $F^0(.|\mathbf{s}_{0t-1})$

This assumption ensures that supply and lot characteristics are exogenous. This requires that supply does not react to prices in the Choice System.

4.2.2 Primitives

States

Food bank i begins the period in state $\mathbf{s}_{it} \in \mathbb{S}$. This represent the food bank's current stock of food. I primarily focus on their stocks from each storage method, so that the individual state has 5 dimensions.[4] This captures the dynamic costs of storing durable goods. If the outcome from period t is $\mathbf{w}_t^a$ they end the period in state $\mathbf{s}_{it}^a$. $\mathbf{s}_{it} = \mathbf{s}_{it}^a$ if and only if the player does not win a single lot. Writing $\mathbf{w}_{it}^T\mathbf{z}_t^g$ as i's winnings from period t I make the following assumptions about how stocks transition:

Assumption 7. *(i)* $\mathbf{s}_{it}$ *transitions according to the following process:*

$$\mathbf{s}_{it} = \mathbf{s}_{it-1} + \mathbf{x}_{it} + \mathbf{w}_{it-1}^T\mathbf{z}_{t-1}^g$$

(ii) $\mathbf{x}_{it} \sim F_i^x = N(\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i)$ *is an exogenous innovation.*

I do not assume that individual states are observed. I assume that stocks are continuous.[5] Day-to-day variation in stocks is likely a major source of variation in bidding behaviour. Food banks supplement their stocks of one type of food they have not recently received from local donors with food from Feeding America. The random variable $\mathbf{x}_{it}$ is observed each morning before items are posted on the Choice System. It can be interpreted as the net daily change in food banks stocks - the food received from local donors, less the food given out to clients. Part (ii) of this assumption imposes that these changes are exogenous. Food banks don't turn down (or request additional) donations from their local donors.[6] Meanwhile, normality is reasonable for these large food distributors receiving many donations from many different sources, and sending out food to many different food pantries.[7]

[4] I also focus on their stock of each subcategory h in order to capture food banks' preferences over how the food is used. However, I will assume that pay-offs are affine in subcategory stock (not subject to diminishing returns - food banks always have people to feed), meaning that the level of the stock of each subcategory is neither identified nor welfare relevant (up to normalisation).

[5] Difficulties in estimating dynamic models with continuous state variables is well known. However, continuous states are common in models of dynamic auctions. A previous version of this model discretised stocks. However the state had to be very finely discretised to capture all the possible combinatorial outcomes from a day's auctions.

[6] This simplification is likely to bias my results in favour of the Old System over the Choice System. If net donations were endogenous this would allow food banks to use their winnings to influence future net donations. Choice would be even more valuable. For example, if local donations were negatively correlated with previous winnings food banks could focus on only winning food from the Choice System they know they cannot get from local donors.

[7] This transition process incorporates two additional assumptions. First, food received from Feeding America, and food from local donors, are perfect substitutes. This is a necessary normalisation. Second, stocks do not degrade over time. This assumption was motivated by discussions with food bank volunteers. Most of the donated food, even fresh produce, have long shelf lives, so that any daily decay parameter is close to 1.

Pay-offs

Following the model outlined in Section 3.2 and Gentry, Komarova, and Schiraldi (2023) I decompose the flow pay-off into a stochastic lot-specific component and a deterministic function of stocks:

Assumption 8. *(i) The flow pay-off from outcome a can be written as*

$$\mathbf{w}_{it}^{aT}\boldsymbol{v}_{it} + j(\mathbf{s}_{it}^{a})$$

(ii) The lot-specific pay-off $\boldsymbol{v}_{it}$ is a random variable with $v_{ilt} \sim F_i^v = N(\boldsymbol{\alpha}_i^T\mathbf{c}_{tl}, \sigma_l^2)$, known privately, observed before entry, and drawn independently across i and t.

(iii) The deterministic function $j : \mathbb{S}_i \to \mathbb{R}$ is finite, with $j(0)$ normalised to 0.

(iv) Pay-offs are quasi-linear in shares (virtual currency).

The flow payoff function j captures both the costs of storing food, and the utility from holding various types of food to be able to distribute them to clients. Part (ii) embeds two assumptions. Assuming the privately known $\boldsymbol{v}_{it}$ is conditionally independent across individuals imposes independent private values. Assuming conditional independence across time is a standard assumption in most dynamic models.[8] The assumption that j has finite range is predominantly for mathematical convenience, while the normalisation is required as only marginal pay-offs are identified. While I assume that pay-offs are quasi-linear in shares, as is standard in auction studies, I allow food banks to differ in their marginal value of wealth, given by $\lambda_i > 0$.[9]

I also assume players have temporally additively separable preferences, and make forward looking decisions with *annual* discount parameter $\beta = 0.99$, so that food banks are extremely patient. I assume F, j, $\mathbf{s}$, and β are common knowledge.

[8] Note that the fact I don't observe the state variables is a violation of the conditional independence assumption. Assumption 7 is the weaker assumption required instead.

[9] However, in Appendix E.4 I showed that in a stationary environment, as I assume shortly, the quasi-linear model is observationally equivalent to a model with an inter-temporal budget constraint. Therefore this assumption is not inaccurate.

4.2.3 The Agent's problem

A (pure) strategy consists of a mapping from a player's type and the state of the world onto entry decisions and bids $(\mathbf{d}_{it}, \mathbf{b}_{it})$. Ex-ante a player's strategy, Λ_i, admits a distribution of bids according to F_i, j_i and $\mathbf{s}$.

Beliefs

Denote $\Gamma_{il}(\mathbf{b}, \mathbf{d}; \Lambda_{-i})$ player i's belief about the marginal probability that they wins lot l, given their bid and entry decision, taking as given the strategies of other players. Denote $P_{ia}(\mathbf{b}, \mathbf{d}; \Lambda_{-i})$ player i's belief about the joint probability, conditional on $(\mathbf{b}, \mathbf{d}, \Lambda_{-i})$, that the outcome from the round of auctions is $\mathbf{w}_t^a$. These objects constitute food banks' beliefs about other players' equilibrium behaviour. In section 4.2.4 I make assumptions about these beliefs to make estimation feasible.[10]

Value Function

Assuming risk neutrality the bellman equation is given by:

$$W(\boldsymbol{v}, \mathbf{s}; j, \Lambda_{-i}) = \max_{\mathbf{b},\mathbf{d}} \{\Pi(\mathbf{b}, \mathbf{d}; \boldsymbol{v}, \mathbf{s}, j, \Lambda_{-i})\}$$

$$\text{Where} \quad \Pi(\mathbf{b}, \mathbf{d}; \boldsymbol{v}, \mathbf{s}, j, \Lambda_{-i}) =$$

$$\sum_l \underbrace{\Gamma_l(b_l, d_l; \Lambda_{-i})(v_l - \lambda_i b_l)}_{\text{lot specific}} + \sum_a \underbrace{P_a(\mathbf{b}, \mathbf{d}; \Lambda_{-i})[j(\mathbf{s}_i^a) + \beta \overbrace{\int_{\tilde{\mathbf{s}}} \int_{\tilde{\boldsymbol{v}}} W(\tilde{\boldsymbol{v}}, \tilde{\mathbf{s}}; j, \Lambda_{-i}) dF_i^{\boldsymbol{v}}(\tilde{\boldsymbol{v}}|\tilde{\mathbf{s}}) dF^{\mathbf{s}}(\tilde{\mathbf{s}}|\mathbf{s}^a)}^{\text{continuation value}}]}_{\text{combination specific}}$$

Continuation Value

The continuation value gives the expected pay-off from the start of the following period having ended the current period in state $\mathbf{s}^a$. This can be written as follows:

$$V(\mathbf{s}^a; \Lambda_{-i}) = \int_{\tilde{\mathbf{s}}} \int_{\tilde{\boldsymbol{v}}} W(\tilde{\boldsymbol{v}}, \tilde{\mathbf{s}}; j, \Lambda_{-i}) dF_i^{\boldsymbol{v}}(\tilde{\boldsymbol{v}}|\tilde{\mathbf{s}}) dF^{\mathbf{s}}(\tilde{\mathbf{s}}|\mathbf{s}^a)$$

A further important object is the sum of the deterministic flow pay-off function and the discounted continuation value, denoted by $k(\mathbf{s}^a; \Lambda_{-i}) = j(\mathbf{s}_i^a) + \beta V(\mathbf{s}^a; \Lambda_{-i})$ and

[10] Assumption 9, discussed shortly, ensures I only need to evaluate the 2^L probabilities of combinatorial outcomes for player i, not all N^L probabilities detailing which player won which lot.

referred to as the 'Pseudo-Static' pay-off function. This is essentially the object one would estimate if one were to incorrectly assume bidders are myopic. Estimating this equilibrium object is key to estimating primitives in a dynamic multi-object model. The importance of this object arises because the value function (and hence the continuation value) can be written as functions of this pseudo-static pay-off:

$$W(\boldsymbol{v}, \mathbf{s}; j, \Lambda_{-i}) = \max_{\mathbf{b},\mathbf{d}} \left\{ \sum_{l} \Gamma_l(b_l, d_l; \Lambda_{-i})(v_l - \lambda_i b_l) + \sum_{a} P_a(\mathbf{b}, \mathbf{d}; \Lambda_{-i}) k(\mathbf{s}^a; \Lambda_{-i}) \right\} \tag{4.1}$$

4.2.4 Equilibrium

I focus on symmetric Markov Perfect Equilibria (MPE), defined as follows:

Definition 4.2.1. : *An MPE consists of a set of strategies* $\boldsymbol{\Lambda}^*$ *and beliefs* $\Gamma(\boldsymbol{\Lambda}^*)$*, such that for any* $(\boldsymbol{v}, j, \mathbf{s})$*:*

Optimality: $(\mathbf{b}_i^{\Lambda^*}, \mathbf{d}_i^{\Lambda^*}) = arg\ \max \left\{ \Pi(\mathbf{b}, \mathbf{d}; \boldsymbol{v}, \mathbf{s}, j, \Lambda^*_{-i}) \right\}$

Consistency: $\Gamma_{il}(b_{il}, d_{il}; \Lambda^*_{-i}) = \mathbb{I}[d_{il} = 1] P(b_{il} > \max_{i' \neq i} \{b_{i'l}\} | \Lambda^*_{-i})$

The optimality condition requires that agents maximise the net present value of pay-offs. The consistency condition requires that bidders' beliefs are consistent with the observed distribution of winning bids.[11] Symmetry requires that bidders with the same 'type', and the same beliefs, place the same bids. This allows us to write the equilibrium strategies as a function of the state: $\boldsymbol{\Lambda}^* = \boldsymbol{\Lambda}(\mathbf{s})$.

In Section 3.2 I proved that, conditional on existence of a symmetric Pure Strategy Nash Equilibrium in the bidding game conditional on entry, such an equilibrium exists.[12] I make the following assumptions about equilibrium:

Assumption 9. *(i) The data are generated by strategy profile* $\boldsymbol{\Lambda}^*$*, a symmetric MPE of the dynamic auction game, with the same MPE played each period.*

[11] This also requires bidders' beliefs about P are consistent with observed joint probabilities.

[12] We also required that equilibrium pay-offs are continuous in $j + \beta V$. A full existence proof remains elusive. However this is not a practical problem. Numerous other papers studying sufficiently complex auction games are unable to guarantee neither existence nor uniqueness of equilibrium. This list includes, for example, Gentry, Komarova, and Schiraldi (2023) on simultaneous first-price auctions, Fox and Bajari (2013) on simultaneous ascending auctions, and Jofre-Bonet and Pesendorfer (2003) on dynamic single-object first-price auctions. The empirical strategy outlined in section 4.3 does not require existence of a MPE. Instead, it only requires that food banks have beliefs that are consistent with observed play.

(ii) $\forall$ *i, l, and* $b_{il} > R_l$, $\Gamma_{il}(b_{il}, 1|\mathbf{s})$ *is strictly increasing and differentiable in* b_{il}.

(iii) $\forall$ *i and* $\mathbf{s}_i$ *the Hessian of the pseudo-static pay-off function k has full rank*

(iv) $\forall$ *i and joint outcome a* $P_{ia}(\mathbf{b}, \mathbf{d}|\mathbf{s}) = \prod_l \Gamma_{il}(b_{il}, d_{il}|\mathbf{s})^{w^a_{il}}(1 - \Gamma_{il}(b_{il}, d_{il}|\mathbf{s}))^{1-w^a_{il}}$

(v) $\forall$ *i, l,* b_{il} *and* d_{il} $\Gamma_{il}(b_{il}, d_{il}|\mathbf{s}) = \Gamma_l(b_{il}, d_{il}|\boldsymbol{\vartheta}(\{\mathbf{s}_i\}_N), \mathbf{s}_0)$

Part (i) is reasonably standard in studies of dynamic games, ensuring that the observed data is stationary. However, it embeds the stronger assumption that food banks' states are stationary. We expect $\mu_i < 0$; without access to Feeding America stocks will drift downwards over time. However food banks use the Choice System to supplement their stocks. When stocks get low, the food bank begins bidding to keep stocks up. This requires that, in equilibrium, food banks have enough control over their winnings to make this possible.[13]

Part (ii) of this assumption is required to ensure that standard first order conditions are necessary for optimality, so that primitives are point identified. I allow for the possibility of ties at the reservation price, which imply non-differentiability of Γ at R. Likewise, part (iii) is necessary for identification, as conditional on the function k it allows the first order conditions to be inverted for $\mathbf{s}_i$.

Part (iv) requires that, in equilibrium, food banks believe winning one lot is conditionally independent of winning any other lot. This essentially assumes that winning bids are conditionally independent across auctions, simplifying estimation considerably. In Appendix M.1 I test and present support for this simplification.

Part (v) is necessary for estimation to be feasible. Without additional assumptions the continuation value for food bank i depends on the state of every food bank, creating an infeasibly large state-space. However, $\mathbf{s}_{-i}$ only enters the continuation value of player i through $\Gamma_{il}(.|\mathbf{s}_{t+1})$. As the number of bidders grows the probability of any individual and their state influencing prices falls to zero. This assumption ensures that equilibrium win probabilities Γ_i do not depend on the states of every player. Instead, they only depend on aggregate statistics of $\mathbf{s}$, using the aggregator

[13] Appendix F discusses the stationarity assumption in additional detail, demonstrating how we can test for stationarity. This assumption essentially requires that equilibrium winnings and net local donations are co-integrated, so that the *equilibrium* stock process is stationary. I demonstrate how we can test co-integration through equilibrium winnings without needing to observe stocks. Broadly, I find evidence of stationarity. Stationarity also requires that the distribution of net local donations is constant over my 3 year period. Feeding America's 'Hunger in America' resource shows that food bank usage and food insecurity remains stable over this period.

$\boldsymbol{\vartheta}$ with known functional form.[14] For notational convenience I absorb $\boldsymbol{\vartheta}(\mathbf{s})$ into the common state variable $\mathbf{s}_0$. This assumption also implies that we can write the value function, continuation value, and k, as functions of $\mathbf{s}_i$ and $\mathbf{s}_0$.

4.2.5 Identification

I now discuss the identification of this model. In Chapter 3 I proved non-parametric identification of the model when all state variables are observed. In particular, I proved that conditional on having identified (i) the equilibrium bid distribution, conditional on the state, and (ii) the state transition process, that the model primitives j and F^v were non-parametrically point identified. The question is then whether this conditional bid distribution and transition process are non-parametrically identified when we do not observe the state, but we do observe bidders' winnings and choice sets. The difficulty is that we are trying to identify the distribution of an observed variable, conditional on an unobserved variable, as well as the transition process of that unobserved variable. Key to the identification argument is using observed winnings almost like an instrument, shifting food banks unobserved stocks, and creating observed variation in bidding behaviour.

Proposition 5. *Under Assumptions 6 - 9 the distribution of equilibrium bids, conditional on both observed and unobserved states, and the state transition process, are non-parametrically identified.*

Proof of this Proposition is presented in Appendix I. I focus on proving a general proposition, requiring only assumptions on: 1) the causal structure of the model, 2) an 'instrument relevance' condition on winnings, 3) an invertibility assumption on the bid distribution, and 4) that food from Feeding America and food from local donors are perfect substitutes. The proof builds on the argument and tools presented in Hu and Shum (2012), restricting their framework to incorporate this observed shifter of the unobserved state, while broadening it to allow for multivariate

[14]This assumption presents a departure from both Jofre-Bonet and Pesendorfer (2003) and Gentry, Komarova, and Schiraldi (2023). It is similar to the large market Oblivious Equilibrium (Weintraub, Benkard, and Van Roy, 2008) and Moment-based Equilibrium (Ifrach and Weintraub, 2017), albeit in a game of incomplete information. Backus and Lewis, 2016 employ a similar assumption in their dynamic auction framework. They argue that because there are many competitors it is unlikely that bidders follow the identities of which other bidders are likely to bid at any given time, and their states. It is unlikely that any given food bank keeps track of competitors' states. This assumption is tested on the empirical equilibrium winning probabilities in Appendix M.1.

latent states. The full proof is omitted from the main text for brevity. I rely on spectral decompositions of linear operators, building on the methodology developed in Hu and Schennach (2008) and Carroll, Chen, and Hu (2010) for the identification of measurement error models.[15] The methodological contribution of this identification argument is broadly applicable well beyond the auction setting. This idea of an observed shifter of the unobserved state is common in the engineering and signal processing literature, known as an "input control variable". It is also useful in the storable goods setting, where the econometrician observes purchases of the storable goods, but may not observe their consumption.

One major benefit of this identification argument argument, using observed shifters of the unobserved state, is that identification is very intuitive, as I will now discuss.

4.2.6 Intuition Behind Identification

First, several objects are trivially identified in this model. For example, food banks' equilibrium beliefs are identified trivially from the observed distribution of winning bids, conditional on lot characteristics. Conditional on beliefs, the key question for identification is whether we can separately identify the deterministic flow payoff function j_i (predominantly storage costs), and the distribution of net donations F_i^x. As we saw in Section 4.3, identification of F_i^x and the pseudo-static payoff function k_i (equivalently, food banks' demand functions) is sufficient for identification of j_i.

Furthermore, I am easily able to identify permanent heterogeneity across bidders, identifying distinct payoff functions and net donation distributions for each food bank, using my long panels. If I wanted to, I could essentially estimate individual models for each food bank.

Identification of Net Donations and Demand Curves

Whenever we observe a food bank win a load of, say, cereal, this gives an observed change in their unobserved stock. We then examine how this impacts their bidding behaviour for cereal and other types of food in the following period. If they always

[15] Identification of this model does not require randomisation nor strict exogeneity. Instead, conditional weak exogeneity is sufficient. The bidder's winnings in periods $t' < t$, as well as the set of available lots at time t, will affect their bidding behaviour at time t. I assume, plausibly, that there is no contemporaneous reverse causation; their bidding behaviour at time t does not impact the set of lots available at t, nor previous winnings.

stop bidding on cereal after winning cereal this implies they have steep demand curves — increasing their stocks of cereal by one load causes a large reduction in their willingness to pay for cereal. This suggests they are relatively capacity constrained, facing large storage costs for cereal. Similarly, if it only causes a relatively smaller change in their bidding for other types of food, this implies that cereal and other types of food are not perfectly substitutable. In contrast, if they only stop bidding on cereal when they have won four or five loads of cereal this indicates relatively flat demand curves and so relatively smaller storage costs. This is essentially the same type of variation we see in Figure 2.11. After winning a truckload of food, we saw a 20% reduction in the propensity to bid on similar types of food.

Identification of the demand curves is also driven by observed variation in the size and composition of lots that are auctioned each day. While winnings acts as observed changes in the unobserved state, variation in choice sets acts as observed *hypothetical* variation in the unobserved state. This is a similar argument to one presented in Gentry, Komarova, and Schiraldi (2023). We examine how a food bank bids when there is one lot of cereal available, and we see that they almost always bid aggressively on this lot. Then we examine how they bid when there are, say, ten lots of cereal available. But now, we see that they still only place bids, or bid aggressively, on one or two loads of cereal. This signals that the food bank does not want to win too many loads of cereal, pinning down the slope of their demand curve. We can similarly identify cross elasticities by examining how bidding behaviour for cereal depends on the number of available loads of a different type of food. This is essentially the same relationship, and the same variation, seen in figure 2.12 previously.

We then separately identify the distribution of net donations by considering how bidding behaviour changes over the course of several days after an observed win. After the initial drop in their propensity to bid right after a win we might see that, on average, it takes around four days for them to return to their pre-win bidding behaviour. This means that, on average, it takes four days for the food bank to give out that load of food to their clients, and so return to their previous stock levels. This pins down the mean of their net donations. We can then use variation in this recovery time to identify the variance of their net donations. This is again presented in figure 2.11, where food banks gradually return to their pre-win bidding behaviour after several days.

Identification of F_i^v and λ_i

Given identification of j_i and F_i^x (and hence k_i), we identify the lot-specific pay-off distribution F^v using variation in lot characteristics, and how they are associated with differential bidding and entry decisions. For example, transportation costs are identified from the relationship between the distance between food bank and lot, and their bidding behaviour.

Finally, the marginal value of wealth parameters λ_i are identified by variation in the scale of bids across food banks, behaving in a similar manner to food bank specific variances. If a food bank has a high marginal value of wealth, and is very budget constrained, their bids will tend to be much more concentrated around zero. These food banks likely do not receive much food from local donors (relative to the demand for their services), and access most of their staples through the Choice System. This means they cannot afford to 'splash out', so tend to bid modestly. But for food banks with a low marginal value of wealth, typically those with access to relatively many local donors, shares (the fake money) are cheap. If they do not spend their shares today, there is rarely anything else they will want to spend them on in future. So, these food banks place bids with high variance. Most often they do not bid, already receiving staples food from local donors. But they do bid for the highest quality food; and when they bid, because there is little else for them to spend their shares on, they splash out and bid exorbitantly high.[16]

[16]In most empirical auction models λ_i are equalised across bidders, enabling inter-bidder welfare comparisons. Instead, I make the assumption that the variance of v_{ilt} is the same across food banks, equalising the unmodelled variation in lot specific attributes across food banks. For example, for food from a particular subcategory, such as a load of cookies (biscuits), this might contain high quality, wholewheat raisin cookies, or low quality sugary cookies. This assumption pegs the scale of the relative payoffs from these different types of cookie to be the same across food banks. This does not allow one food bank to treat all types of cookie as the same, while another only ever wants the healthy types of cookie. Allowing the marginal utility of wealth to vary across food banks is important for analysing the distributional consequences of choice. Meanwhile, restricting the variance of the lot specific value to be equal across food banks simply restricts the degree of heterogeneity, biasing my results in favour of the Old System.

4.3 Estimation

4.3.1 Empirical Strategy

This section describes the estimation procedure used to estimate the model. Section 4.3.2 outlines the three step procedure, noting the relationship to the procedure of Jofre-Bonet and Pesendorfer (2003). Section 4.3.3 discusses parametrisation and estimation of beliefs, which are estimated using a likelihood procedure. Section 4.3.4 discusses the second estimation step, in which I simultaneously estimate the state transition process, the distribution of lot-specific values, and the pseudo-static pay-off function. In section 4.3.5 I detail how I disentangle the combinatorial flow pay-off j and the discounted continuation value from the pseudo-static pay-off. Full details of the estimation procedure are given in Appendix K.

4.3.2 The 3-Step Procedure

The standard approach to estimating dynamic auction games, from Jofre-Bonet and Pesendorfer (2003), relies on the ability to write the continuation value as a function of the distribution of bids *only*. This is not possible in the multi-object context because of an order problem: Bids are L dimensional, while values, and continuation values are 2^L dimensional. Full solution methods of Rust (1987) are computationally intractable in this setting - recursively evaluating the value function requires numerically maximising bids for each $\mathbf{s}, \boldsymbol{v}_i$.

Instead, as we saw in Section 3.3, we can write the continuation value as a function of the distribution of bids *and* pseudo-static pay-offs. If we know the pseudo-static pay-offs we can find the continuation value, which then allows us to back out the flow pay-off j from the definition of pseudo-static pay-offs: $k = j + \beta V$. To estimate the pseudo-static pay-off function we estimate the model as if we were estimating a static model, but allow pay-offs to depend on $\mathbf{s}_0$.

4.3.3 Step 1: Beliefs

Assumption 9 ensures food banks form beliefs consistent with observed play. Therefore, we can estimate beliefs using the observed distribution of winning bids. Estimating beliefs in this way avoids the need to solve the model for equilibrium. This procedure is common in the empirical auction literature due to the extensive computational cost of finding equilibrium beliefs (Athey and Haile, 2007).

I make parametric assumptions about Γ to facilitate estimation. I assume winning bids follow a generalised extreme value distribution, censored at the reservation price:

$$\Gamma_{il}(.|\mathbf{s}) = GEV(.;\xi_c,\zeta_c,\mathbf{c}_{lt}^T\mu + d_{lt}) \qquad \text{where} \quad d_{lt} = \mathbf{s}_{0t}^T\vartheta \tag{4.2}$$

Where the shape and scale parameters ξ and ζ are category specific. $\mathbf{c}_{lt}$ gives a vector of lot specific location shifters, such as the subcategory composition.

d_{lt} describes how the distribution varies with the state of the world, forming an index to be estimated. The index is a linear function of the quantity of food, by usage type, auctioned at t and also the quantity over the previous 30 days, up to $t-1$. This is designed to capture competitive pressures on prices. If very little food has been auctioned over the previous month, one would expect a higher price. Estimating this demand index in the first stage allows us to use the estimated index in later estimation objects. In particular, when considering the transition process of common state variables, I can focus on just the transition process of d_{lt}.

The Fisher–Tippett–Gnedenko theorem establishes that the Generalised Extreme Value distribution is the limiting distribution of the maximum of independently distributed random variables. In an Independent Private Value framework, the winning bid is just the maximum of (conditionally) independent random variables. Therefore the GEV assumption is easily justified. Meanwhile the parametrisation is chosen to be suitably flexible, given the available data. Full details of how I estimate beliefs are included in Appendix K.1.[17]

4.3.4 Step 2: The 'Pseudo-Static' Payoff

I now describe the second part of the estimation procedure, in which I jointly estimate $F^{\mathbf{x}}, F^{v}, k$ for each food bank. I begin by discussing the problems that must be

[17]This details additional covariates included in estimation. Also, how I estimate the probability of tieing at the reservation price. If no food bank bid on a lot, then a food bank would have won if they bid the reservation price. But, if some other food bank won at the reservation price, then the food bank would have tied had they bid the reservation price. Importantly, allowing for ties rationalises food banks choosing to bid just above the reservation price.

overcome in this estimation step, before detailing the parametric assumptions made to enable estimation of each of the three sets of objects. The key functional form restriction is assuming that the pseudo-static pay-offs are quadratic in stocks, similar to the standard assumption of quadratic storage costs and, therefore, linear demand curves. This also has the benefit that the non-additivities in payoffs across lots auctioned simultaneously are pairwise. Estimating the second-stage then requires estimating a censored Linear Gaussian State Space model.

The central estimation difficulty concerns the unobserved state and the unobserved bids. Bids are unobserved when a food bank chooses not to enter a particular auction, which occurs frequently. Ignoring these bids introduces the standard problems of censoring in econometrics. I estimate the model using a Gibbs Sampling procedure. I use data augmentation to iteratively sample both unobserved bids and unobserved states from their conditional posterior distributions, before updating my parameter estimates given the augmented data. Full details of the estimation procedure, including assumptions on prior distributions, are given in Appendix K.2.[18][19]

Individual States

I estimate individual $\times$ storage type specific mean and variance parameters (μ_{ig}, Σ_{ig}) for the normally distributed net local donations $\mathbf{x}_{it}$. I make use of the prior information from Assumption 9, which requires the stock transition process is stationary. In Appendix F.4 I prove that stationarity requires:

$$\boldsymbol{\mu}_i = -E[\mathbf{w}_{it}^T \mathbf{z}_t^g] \qquad \boldsymbol{\Sigma}_i < 2Var[\mathbf{w}_{it}^T \mathbf{z}_t^g]$$

[18] Because heterogeneity is an important theme of the model I generally estimate separate parameters for each food bank. However, I do not always have enough identifying variation for each individual food bank. I use a Bayesian Hierarchical framework to flexibly introduce information pooling across bidders in my model. This approach is flexible enough to allow pooling for food banks that lack identifying variation, placing more weight on the hierarchical parameters, and allow separation for food banks that have a lot of identifying variation, placing more weight on the data.

[19] Due to computational requirements I focus my estimation on the 34 food banks that each won at least 150 lots (Type 1 food banks). These food banks consume 70% of the food from the Choice System. It is standard in empirical auction studies to estimate a main model and a model of 'fringe' bidders (For example Jofre-Bonet and Pesendorfer, 2003 and Gentry, Komarova, and Schiraldi, 2023). I estimate a simpler (myopic) model for the remaining 88 food banks who won at least 30 lots and whose locations are known (Type 2 food banks). Details of this model and estimation is included in Appendix K.4. All counterfactual analysis uses the models from both sets of food banks.

On average winnings must offset mean net donations and the variation in winnings reflects the variation in net donations. However I do not impose that either relationship holds exactly, and instead use them to build informative priors.[20] I have a standard initial state problem. The quadratic assumption I impose on k ensures the level of the state is not identified, so I normalise the initial state to zero.

Lot-Specific Pay-offs

I specify the mean of the lot specific pay-off v_{ilt} as $\alpha_i distance_{ilt}$, so that the mean lot specific pay-off depends linearly on the distance between food bank i and lot l. The variance σ_l^2 is category combination specific. Assumption 8 imposes that the lot specific pay-offs are uncorrelated across t and i. To simplify estimation I also assume these variables are conditionally uncorrelated across lots l.

Combinatorial Pay-offs

I fit a parametric form to the function $k(\mathbf{s}_i, \mathbf{s}_0)$. Within feasibility constraints, I choose a parametric function to reflect how food banks gain benefits from food according to how the food is used (according to it's subcategory) and how they face costs of storing the food (according to the storage method). I assume the following:

$$k(\mathbf{s}_i) = \Phi \mathbf{s}_i^h + \mathbf{s}_i^{gT} \Psi_i \mathbf{s}_i^g \tag{4.3}$$

Where Φ is an 1×164 row vector, and Ψ_i is a 5×5 dimensional matrix. The form of k as I have presented it above depends on both the stock of each storage type $\mathbf{s}_i^g$ and the stock of each subcategory $\mathbf{s}_i^h$. However, consider the marginal pseudo-static pay-off from winning lot l with characteristics $\mathbf{c}_l$:

$$k(\mathbf{s}_i + \mathbf{z}_l) - k(\mathbf{s}_i) = \Phi \mathbf{z}_l^h + \mathbf{z}_l^{gT} \Psi_i (\mathbf{z}_l^g + 2\mathbf{s}_i^g)$$

This does not depend on $\mathbf{s}_i^h$, because the hessian of the pseudo-static pay-off, with respect to $\mathbf{s}_i^h$, has rank 0. This means that $\mathbf{s}_i^h$ 'falls out' of the model, so that I can focus on $\mathbf{s}_i^g$ as the state variable. The quadratic specification ensures that non-additivities

[20]This is because the relationships only holds, averaged over time. Winnings $\mathbf{w}_{it}^T \mathbf{z}_t$ are both extremely 'lumpy' as well as auto-correlated, so that convergence to the true mean is slow.

in payoffs are pairwise: The change in marginal payoff from winning lot l as well as lot m is independent of whether they also win lot n. This drastically simplifies estimation as I no longer need to consider all 2^L possible combinations of lots, and instead only need to consider the $LC2$ pairs of lots.

In theory k should depend on $\mathbf{s}_0$, capturing how the continuation value depends on food banks' beliefs about future supply. If they believe many lots will be auctioned next period, prices are likely to be low in future, lowering the opportunity cost from *not* winning today. In practice I assume k is independent of $\mathbf{s}_0$ for two reasons. First, the supply of shares varies with supply to ensure prices remain approximately constant over time. Therefore we expect little variation in average prices over time. However, relative prices may vary with supply and this may impact the continuation value. As I show in Results section 4.4.1, the relationship between supply of different types of food and prices is not economically significant. Nonetheless, in Appendix M.2.1 I present results from an econometric specification that includes d_{ltg} (the demand index for food type g) as an input to k.

I impose that Φ is constant across i. Allowing it to vary introduces an unwieldy number of parameters to the model. This is a reasonable assumption as food banks likely gain the same benefit from different subcategories.

Three Equations

The model presented above leads to necessary optimality conditions for bidding which can be inverted for the Inverse Bid System, $\xi_{ilt}(\mathbf{b}, \mathbf{d}|\mathbf{s}_i, \mathbf{s}_0)$. This gives us the following three equation model, consisting of a 'Transition Equation', an 'Observation Equation', and a 'Censoring Equation':

$$\mathbf{s}^g_{it} = \mathbf{s}^g_{it-1} + \mathbf{x}_{it} + \mathbf{w}^T_{it-1}\mathbf{z}^g_{t-1} \qquad \rightarrow \textbf{Transition Eq.}$$

$$\lambda_i y_{ilt} = \Phi \mathbf{z}^h_{tl} + \mathbf{z}^{gT}_{tl}\Psi_i(\mathbf{z}^g_{tl} + 2\mathbf{s}^g_{it} + 2\textstyle\sum_{m\neq l}\Gamma_m(b_{itm})\mathbf{z}^g_{tm}) + v_{ilt} \qquad \rightarrow \textbf{Observation Eq.}$$

$$y^*_{itl} = \begin{cases} b_{itl} + \frac{\Gamma_l(b_{itl})}{\nabla_b\Gamma_l(b_{itl})} & if \;\; b_{itl} > R \\ R + \frac{\Gamma_l(R+1)}{\Gamma_l(R+1)-\Gamma_l(R)} & if \;\; d_{itl} = 1, b_{itl} = R \\ R & if \;\; d_{itl} = 0 \end{cases} \qquad \rightarrow \textbf{Censoring Eq.} \tag{4.4}$$

I derive this system of equations in Appendix G. The observation and censoring equations come from the inverse bid system, while the transition equation was defined in Section 4.2.2. Importantly, the Observation Equation is affine in the unobserved state $\mathbf{s}^g_{it}$, and so this is just a Censored Linear Gaussian State-Space model.[21]

Estimation procedure

Unlike the non-censored case, the likelihood of the Censored Linear Gaussian State Space model is intractable. Instead, estimation is performed using a Gibbs Sampler, which consists of the following steps:[22]

1. Draw beliefs Γ from their posterior distribution using Metropolis Hastings
2. Given Γ, the parameters of the pseudo-static model $\{k_i, F^v_i, F^{\mathbf{x}}_i\}_N$, and states $\{\mathbf{s}^g_{it}\}_{T,N}$, draw censored values of $\{y_{ilt}\}_{NTL}$ using the Censoring Equation
3. Given Γ, $\{k_i, F^v_i, F^{\mathbf{x}}_i\}_N$, and $\{y_{ilt}\}_{NTL}$, use the Carter-Kohn Algorithm to draw $\{\mathbf{s}^g_{it}\}_{T,N}$ using the Transition and Observation equations.

[21]This Observation Equation is endogenous - it contains a dependent variable on the right hand side, through b_{itm}. In general, b_{itm} may be correlated with v_{itl} — when v_{itl} is large, food bank i may prefer to win lot l instead of lot m (assuming negative complementarities), so lower their bid on lot m. In practice, however, simulations suggest the inconsistency caused by this endogeneity is very small, as $\Gamma_{im}(b_{itm})$ is generally very unresponsive to v_{itl}, depending much more on v_{itm}, $\mathbf{z}_{itm}$ and even $\mathbf{z}_{itl}$. In Appendix M.2.2 I employ the instrumental variable procedure from 3.4, using $\mathbf{z}^g_{tl} + 2\mathbf{s}^g_{it}$ as an instrument for $\mathbf{z}^g_{tl} + 2\mathbf{s}^g_{it} + 2\sum_{m\neq l}\Gamma_{im}(b_{itm})\mathbf{z}^g_{tm}$.

[22]Recognise how this procedure builds on the identification argument presented in 4.2.5. In step 3. I use variation in winnings and the effect on bidding behaviour to infer changes in stocks, pinning down the distribution of net donations. In step 4. I use variation in $\mathbf{z}_t$ as well as winnings (through the sampled states), and how these impact bidding, to pin down k.

4. Given Γ, $\{y_{ilt}\}_{NTL}$ and $\{\mathbf{s}_{it}^{g}\}_{T,N}$, draw $\{k_i, F_i^{v}, F_i^{\mathbf{x}}\}_N$ from their posterior distributions using the Observation Equation.

5. Repeat

Additional details of the estimation procedure are given in Appendix K.2.

4.3.5 Step 3: The Dynamic Decomposition

At this point we have draws of beliefs, $\{k_i, F_i^{v}, F_i^{\mathbf{x}}\}_N$, and $\{\mathbf{s}_{it}^{g}\}_{T,N}$ from the posterior distribution. I now describe how I evaluate the continuation value $V(\mathbf{s}_i^{g}, \mathbf{s}_0)$. I make use of the following proposition:

Proposition 6. *The ex-ante Value Function can be expressed as:*

$$E[W(\boldsymbol{v}_{it}, \mathbf{s}_i, \mathbf{s}_0)|\mathbf{s}_i, \mathbf{s}_0] = \frac{E[q_t(\mathbf{s}_i^{g})\pi(\mathbf{b}_{it}, \mathbf{d}_{it}|\mathbf{s}_i^{g}, \mathbf{s}_0)|\mathbf{s}_0]}{E[q_t(\mathbf{s}_i^{g})|\mathbf{s}_0]}$$

Where $q_t(\mathbf{s}_i^{g})$ gives the posterior probability that $\mathbf{s}_{it}^{g} = \mathbf{s}_i^{g}$ and

$$\pi(\mathbf{b}, \mathbf{d}|\mathbf{s}_i^{g}, \mathbf{s}_0) = \sum_l \lambda_i \frac{\Gamma_l(b_l, d_l; \mathbf{s}_0)^2}{\nabla_b \Gamma_l(b_l, d_l; \mathbf{s}_0)} - \sum_{m \neq l} \Gamma_l(b_l, d_l; \mathbf{s}_0)\mathbf{z}_l^{gT}\Psi_i \mathbf{z}_m^{g}\Gamma_m(b_m, d_m; \mathbf{s}_0) + \mathbf{s}_i^{gT}\Psi_i \mathbf{s}_i^{g}$$

This proposition is proven in Appendix J. The identity $\pi(\mathbf{b}, \mathbf{d}|\mathbf{s}_i^{g}, \mathbf{s}_0)$ arises from substituting the first order conditions back into the maximand, writing the ex-ante value function as a function of bids and the pseudo-static pay-off function. The main proof then extends the key result from Arcidiacono and Miller (2011) to the continuous choice case. The sample counter-part to this object is then easily found. Full details of this procedure are given in Appendix K.3. I evaluate the ex-ante value function across a grid of states. I use a 20^5 grid evaluated evenly across points from the posterior sampled states.[23]

Having evaluated the ex-ante value function for a parameter draw, I evaluate the continuation value using $V(\mathbf{s}_i, \mathbf{s}_0) = \int\int E[W(\boldsymbol{v}, \tilde{\mathbf{s}}_i, \tilde{\mathbf{s}}_0)|\tilde{\mathbf{s}}_i, \tilde{\mathbf{s}}_0] dF(\tilde{\mathbf{s}}_0|\mathbf{s}_0) dF(\tilde{\mathbf{s}}_i|\mathbf{s}_i)$. Finally I back out $j(\mathbf{s}_i) = k(\mathbf{s}_i, \mathbf{s}_0) - \beta V(\mathbf{s}_i, \mathbf{s}_0)$

[23] Such a large grid is feasible in this context as I only need to perform the procedure once. However, storing 34,000 grids, one for each food bank $\times$ parameter draw, is not. I use a quadratic approximation of the ex-ante value function. In Appendix M.3.1 I evaluate the fit of this approximation by considering the R^2 of the approximation regression. 100% of these R^2s lie between 0.99 and 1. The fit is strong due to the quadratic term that appears in the ex-ante value function.

4.4 Results

4.4.1 Step 1: Beliefs

This section presents the results from the three stages of estimation described in section 4.4.1. Only a subset of key results are reported in the text, focusing on the theme of heterogeneity and only presenting results for the 34 largest ('Type 1') food banks, who consume 70% of the food on the Choice System. Full results are reported in Appendix L, including Gelman-Rubin convergence statistics. When discussing statistical significance I focus on 95% credible intervals. I present several graphs plotting the individual parameters, and credible intervals, across food banks.[24]

The key parameters estimated in the first stage are the shape, scale, and location parameters that describe the generalised extreme value distribution. The Shape parameters lie significantly within the interval $(-0.1, 0.5)$, with none of the parameters significantly below zero. The scale parameters are all estimated to be between 2000 and 5000. The implied variance is much higher than the variance of winning bids. This variation is needed to rationalise the relatively high likelihood (around 0.3 on average) of winning at the reservation price.

The estimated subcategory fixed effects are precisely estimated, widely dispersed, and strongly correlated with the average winning bids across subcategories presented in Figure 2.7 ($R^2 = 0.74$). The standard deviation of posterior means across subcategories is 2400, while the mean posterior standard deviation is 800. This suggests much more variation across subcategories than uncertainty about subcategory posterior means. The previous 30 days supply of food is estimated to have a significant negative effect on prices for every type of good except Non-Food and Condiments. The coefficient on Meals is the largest, estimating that each additional increase in the previous 30 day supply by one thousand tons (approximately one hundred loads) decreases the winning bid by 350 shares. This magnitude, while statistically significant, is not economically significant (around 0.017 standard deviations),

[24] Appendix L.6 discusses model fit, both in and out of sample. Broadly, the model fits the data well, matching average patterns of consumption across food banks and food types, as well as average propensities to place bids across food banks, categories of food, and months. The simulated distribution of bids conditional on entry does not fit the data as well, failing to match the observed long right tail of bids and over-estimating the mean and standard deviation of bids by a magnitude of 50%. However this is not a major problem, since for my counterfactual exercises it is food banks' allocations that matter, rather than their signal of preferences.

relative to the variation seen across different types of food through the subcategory parameters. The present day's supply of food is not estimated to have a significant effect on prices, however these estimates are noisy.

4.4.2 Step 2: The 'Pseudo-Static' Payoff

Unobserved State

Figure 4.1 plots the estimated μ_i and $\sqrt{\Sigma_i}$ parameters for each of the Type 1 food banks. 95% credible intervals are also plotted. Estimates are sorted according to the estimates for the Dried food type.

There are two key takeaways from these results: First, the extent and significance of the heterogeneity. The variation across food banks in the distribution of net donations of for Fresh food is particularly stark. It is also clear that the variation is not purely vertical: Some food banks have higher estimates for Dried than Tinned, while other food banks exhibit the opposite relationship. The second key takeaway concerns the differences in the scale between the two sets of graphs - the standard deviations are generally larger than the means, so we expect the unobserved state process to be noisy.[25]

FIGURE 4.1: Estimated unobserved state parameters

(A) (B)

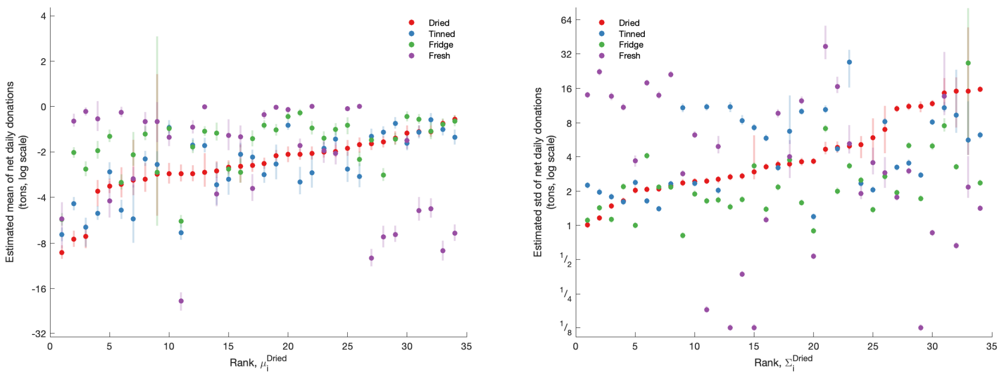

Note: The figure plots posterior means for the mean and standard deviations of net local donations, as well as 95% credible intervals. Results are sorted according to the estimates for the Dried storage type. The plot excludes Type 2 food banks, and the 'non-food' type, to improve graphability.

[25] I also correlate my estimates with observable characteristics of food banks, such as population density in their catchment area, and agricultural rents. This analysis is omitted as I do not find any particularly striking results. Correlations are reasonable and in the expected directions. For example, food banks in areas with higher population density or lower agricultural rents are estimated to have lower average net donations of fresh produce.

Lot specific pay-off

The key lot specific parameters are the coefficient on distance between food bank i and lot l, the constant marginal value of shares λ_i, and the standard deviation of the idiosyncratic lot specific shock v_{ilt}.

Distance coefficients vary across food banks from a cost of 5 to 98 shares per km, with an average of 23. λ_i for the Type 1 food bank with median consumption is normalised to 1. The posterior means then vary from 0.5 to 5. I observe a negative relationship between the shadow price of shares and goal factor - food banks with a higher goal factor receive more shares. However this relationship is very weak, stressing the importance of unobserved food wealth. Standard deviation parameters are estimated to be large, with a mean posterior mean of $32,000$. Such large standard deviations are needed to rationalise the small probability of bidding ($\approx 2\%$) with the relatively large variation in bids conditional on bidding.[26]

Combination pay-off

The estimated Φ parameters, associated with the marginal value of winning a pound of each subcategory, are strongly correlated with the first stage subcategory parameters ($R^2 = 0.82$). Panel A of Figure 4.2 plots food banks' willingness to pay for an additional $40,000$ pounds from each storage type, evaluated when stocks are zero.[27] I estimate significant variation across food banks, and within food banks across storage types. The willingness to pays are generally (significantly) negative, as expected. I find broadly low figures for Dried food, which is driven by food banks generally bidding on fewer fresh items at a time, and bidding on fewer items after winning.

4.4.3 Step 3: The Dynamic Decomposition

Figure 4.2 panel (B) plots posterior means of the marginal flow-payoffs from receiving $40,000$ pounds, evaluated when stocks are zero. Estimates are plotted for Type 1

[26] I decompose the residual variation into idiosyncratic variation in the lot specific value, and unobserved variation in the state. Condition on observing the previous period's state, 45% of this short-run variation is due to variation in the unobserved state, and 55% due to lot specific variation. When we consider long run variation, 72% of the unobserved variation is estimated to come from variation in the unobserved state.

[27] Due to the normalisation made in section 4.3.4, stocks on March 1st, 2014 were normalised to zero, as the level of stocks is not identified.

FIGURE 4.2: Estimates of Ψ_i and j_i

(A) (B)

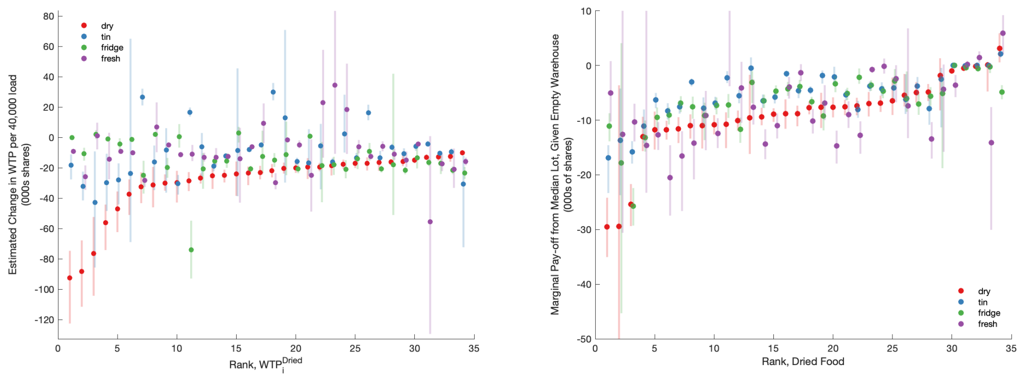

Note: Figures plot posterior mean equilibrium willingness to pay (A) and marginal flow pay-offs (B) for a 40,000 load for each storage type. Bars give the 95% credible intervals. Estimates are ordered according to the estimates for Dried loads. The plot excludes Type 2 food banks and estimates for non-food storage type. WTPs and marginal flow pay-offs are evaluated when stocks are zero.

food banks, sorted according to the estimate for the Dried storage type.

I estimate significant differences across food banks and across types of food. This suggests different food banks have different capacities for storing different types of food. Marginal flow pay-offs are generally negative, indicative of storage costs. Positive marginal pay-offs suggest food banks also benefit from not having an empty warehouse. The results are broadly similar to those plotted in figure 4.2 panel (A) as expected. Because stocks are persistent, the continuation value from storing food tomorrow should be similar to the cost of storing food today. However, the absolute magnitudes in panel (B) are generally smaller than in panel (A). This is because the pseudo-static pay-offs account, not just for present storage costs, but also expected future storage costs and the future *Opportunity Cost* of storage.

To summarise, I estimate that food banks differ systematically in their net local donations, storage costs, transportation costs and marginal value of shares. This suggests that giving food banks choice over their allocations is welfare improving.

4.5 Counterfactuals

Feeding America introduced the Choice System, replacing the Old System, to give food banks choice over the food they received. Feeding America also put significant resources into minimising downsides from the system. They were worried that it might lead to an inequitable distribution of food, allowing smaller food banks to 'fall through the cracks'. In this Section I investigate the welfare and distributional consequences of introducing the Choice System. I then consider a number of additional mechanisms used by other food bank networks around the world. For example, what would happen if Feeding America did not allocate food centrally, and instead linked food banks up with the closest donors.[28] Under this mechanism food is implicitly only offered to the nearest food bank.

In Section 4.5.1 I discuss welfare metrics, given fake money has no value outside the Choice System. Section 4.5.2 briefly explains how I simulate equilibrium allocations under the Old System, before presenting results describing the welfare and distributional consequences of the transition from the Old System to the Choice System. In Section 4.5.3 I use additional simulations to analyse several possible drivers of these results; in particular I consider the role of signalling versus simultaneous allocation. Section 4.5.4 introduces several additional mechanisms and presents results from these additional counterfactual exercises. Details of how I simulate these the mechanisms are given in Appendix N.

4.5.1 Welfare

My counterfactual simulations produce welfare measures in terms of consumer surplus, measured in shares. This has the benefit that consumer surplus is a cardinal measure, enabling inter-food bank comparisons. However the value of shares is difficult to interpret as they have no value outside the Choice System. Instead, similar to Agarwal et al. (2021), I report welfare as the equivalent increase in the supply of food that would have the same total value in shares.[29] This measure is valid under competitive equilibrium because the money supply adjusts to ensure prices are constant, given changes to the supply of food. If the supply of food doubles, the money supply adjusts so that expenditure doubles. Therefore, if consumer surplus under

[28]This approach is taken by the Trussel Trust (U.K.) and Second Bite (Australia), among others.

[29]This is similar to how consumer surplus is typically measured in dollars, except that here I am measuring it in terms of how much food those dollars could purchase. I could also use distance travelled as a numeraire, reporting the equivalent reduction in the total distance travelled. Prendergast (2022) then measures welfare in dollar terms using estimates of trucking costs from the literature. Given my results that different food banks face very different transportation costs, as shown in Appendix L.2, this approach is unattractive.

the Old System is double that under the Choice System, I liken this to double the nominal expenditure, which equates to double the supply of food.[30]

Importantly, the 'level' of welfare is not identified because the levels of both stocks and flow payoffs $j(\mathbf{s}_i)$ are not identified. I use a random allocation as a benchmark counterfactual. This is a relevant benchmark since it can be considered the baseline worst case allocation mechanism. Results are reported on a scale of zero (food is allocated no better than random) to 292 tons (the daily average amount allocated under the Choice System).

I report both utilitarian welfare and a weighted sum using Goal Factors as priority weights. I also report descriptive measures of welfare. For example, the total amount of food allocated. This is an important measure given the political cost to Feeding America of wasting food, or the indirect harm from donors being less likely to donate again in future.[31] The distance food must travel is another key metric. We expect food banks to sort on location as food banks choose nearby lots.

4.5.2 The Old System

Details of the Old System were given in Section 2.1.2. For the sake of clarity, I remind the reader of the core elements of this allocation mechanism. Under the Old System loads were sequentially offered to whichever food bank was at the head of a queue. The queue was determined by how recently they had been offered food, as well as their Goal Factor. Food banks were given an hour to decide whether to accept or reject a given load, and so loads could only be offered to a fairly limited number of food banks.[32] I assume loads could be offered to up to 10 food banks before being returned to the donor.

[30]In reality this figure is a lower bound. This is because, if supply and expenditure were to double, consumer surplus would less than double due to the concavity of payoffs (storage costs are convex). Therefore doubling consumer surplus would require more than doubling the food supply.

[31]The welfare measures reported in this system do not account for endogeneity in the supply of food, with respect to the allocation mechanism employed. This relationship is unfortunately not identified, given that I only observe data from the Choice System. However, given that my results generally show more food being accepted under the Choice System, this simply means my results can be interpreted as lower bounds on the value of choice.

[32]Originally food banks would not lose their place in the queue if they rejected a load of fresh produce. For tractability I assume they lose their place for rejecting such a load. This assumption makes the Old System relatively better, in incentivising food banks to accept these undesirable loads, but also relatively worse if food banks are made worse off by taking these loads. As we will see in Section 4.5.4 this is unlikely to make much difference.

I model the Old System in continuous time. This is realistic since food banks could receive a call from Feeding America at any time. I also assume that food is given out to clients and received from local donors at random times during the day. Continuity of time ensures that the probability of a call from Feeding America, or local donors, occur simultaneously with probability almost surely zero. I assume food banks do not observe offers made to, nor decisions of, other food banks. They do not know their place in the queue; only their own Goal Factor, and when they were last offered a load. I assume they form beliefs about the rate they receive calls from Feeding America, and also the probability of being offered a load with characteristics $\mathbf{c}_l$, conditional on receiving a call. I assume these objects are independent of the time since their previous offer. In practice, given the frequency and irregularity with which food banks are offered food (on average, around 5 times per day with a standard deviation of 8) this simplification is unlikely to cause significant inaccuracy.

I assume a Markov Perfect Equilibrium in symmetric strategies, as defined in Section 4.2.4. This requires food banks make optimal accept/reject decisions given their beliefs, and that beliefs are consistent with the observed realisation of offer rates. Appendix N.1 details how equilibrium beliefs and value functions are formed. Given beliefs I find each food banks' value function by numerically solving the Hamilton-Jacobi-Bellman differential equation. I then simulate the mechanism and update beliefs using observed offer rates, repeating until convergence.

Importance of Choice

Figure 4.3 presents the posterior mean and 95% credible intervals for various measures under the Choice System and Old System. The first column gives the unweighted sum of estimated welfare in equivalent tons of food. All welfare results are relative to the baseline random allocation. The mean welfare under the Choice System is mechanically equal to the average daily amount of food. While both Systems achieve significantly more welfare than a random allocation, the Choice System yields significantly higher welfare than the Old System. Welfare is on average 17.1% higher under the Choice System than the Old System, which is enough food to provide an additional 22,300 meals each day. When welfare is weighted according to Goal Factor, this figure increases to 22.9% higher under the Choice System. These results are extremely similar to those in Prendergast, 2022, who finds that welfare is roughly 21% higher under the Choice System.

The third column shows that, under the Choice System, food banks sort into consuming closer lots, with around 6,000 km less transportation required each day. This is in spite of the result from the fourth column that around 22 additional tons of food is accepted each day under the Choice System.

FIGURE 4.3: Counterfactual Results

Mechanism	Welfare (unweighted)	Welfare (weighted)	Distance (000 km per day)	Allocated (tons per day)
Choice System	292 (276, 309)	745 (672, 815)	16 (14.6, 17.4)	271 (253, 284)
Old System	242 (203, 276)	576 (415, 706)	22.3 (22, 22.6)	249 (248, 251)

Note: This table displays posterior means and 95% credible intervals for various measures of welfare. Welfare is measured in food equivalent terms relative to a purely random allocation, and mean welfare of the Choice System is normalised to 292. Therefore welfare should be interpreted as pegged to this scale of 0 (as good as random) to 292 (as good as the Choice System).

When I decompose these welfare differences into the stock dependant component $j(\mathbf{s}_{it})$ and the lot specific component v_{ilt} (which contains transportation costs) we see that 81.5% of the welfare gains come from the stock dependant component. Reduced transport costs account for 6.53% of the gain. This is because transport costs are only estimated to be major cost for a small number of food banks. The additional food that is accepted under the Choice System only explains 1.72% additional welfare, as this food is typically lower quality. The remainder is attributed to food banks sorting into food with higher unobserved idiosyncratic payoffs.

A likely driving force behind these results is that under the Old System food banks accept food that does not meet their most pressing needs at that particular point in time. They accept food that might be more useful to a different food bank, and may prevent them from accepting food they value more in the near future.[33] This is evident for three reasons. First, most food banks are still offered enough food to prevent their stocks from trending downwards. Therefore it is not that food banks

[33] Some food banks also accept food they would not have accepted under the Choice System just in case they need the food in future. Whereas under the Choice System they know they will be able to bid on this food when they need it, rather than having to wait and hope they will be offered it.

do not receive enough food. Second, stocks are more variable under the Choice System than the Old System. The average short run variance is around 25% higher under the Choice System, and the long variance is almost 100% higher. So, it is not about food banks smoothing their stocks. As evident from the main result, their stocks spend more time close to the optimum - the maximum of $j(\mathbf{s}_i)$. Therefore food banks are choosing to have more variable stocks, occasionally increasing stocks for food that is particularly valuable. Finally, under the Choice System food banks were free to bid zero on the types of food they were offered in my counterfactuals, but chose not to. This revealed preference implies that, often, they only want certain types of food at certain times. Under the Old System they just have to wait to see what they are offered. Their stocks remain at consistent levels, suggesting they accept different types of food instead, accepting food that is just good enough.

Distributional Consequences

Figure 4.4 presents welfare results by food bank, plotting the difference between food bank specific welfare under the Choice System and the Old System. On average 85% of food banks are better off under the Choice System.

FIGURE 4.4: Individual Welfare

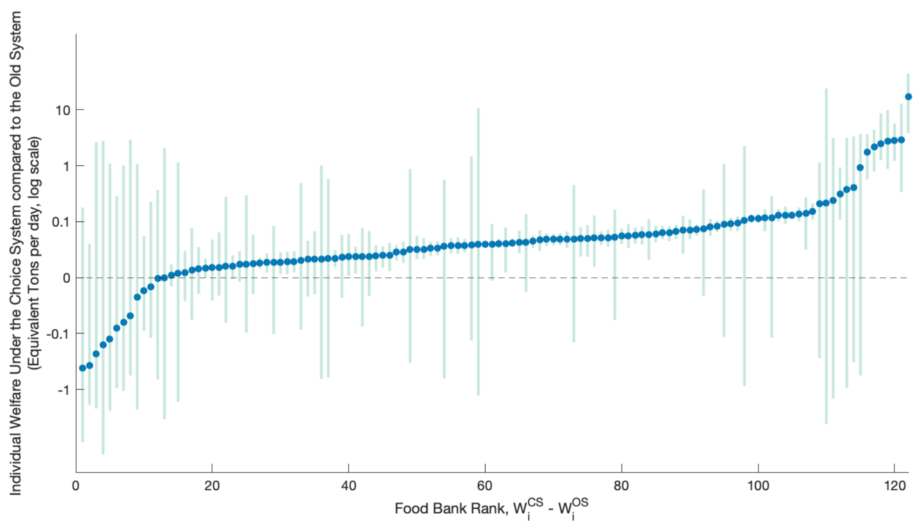

Note: Plots food bank specific welfare under the Choice System minus welfare under the Old System, ordered by the welfare difference, with 95% credible intervals across posterior draws. On average 85% of food banks are better off under the Choice System than the Old System.

Given the difference between the weighted and unweighted welfare estimates, it is unsurprising that there is a positive correlation between this welfare difference and Goal Factor ($\rho = 0.184$). I find negative correlation of -0.205 with estimated λ_is, the marginal value of wealth. This suggests that food banks who rely less on food from Feeding America benefit more from choice, from being able to be picky.[34]

Relative welfare is negatively correlated with food banks' mean net donations $\boldsymbol{\mu}_i$ ($\rho \approx -0.2$). This is driven by food banks in the tail of the distribution of net donations, particular those whose stocks trend downwards under the Old System. These are the food banks that regularly bid and win food at negative prices under the Choice System. Welfare is also positively correlated with sampled $\boldsymbol{\Sigma}_i$ parameters, the variance of net donations, but only for non-food, dried, refrigerated, and fresh stocks ($\rho \approx 0.26$). This is sensible — food banks with more uncertain net donations benefit from being able to choose the food they receive from Feeding America. However I do not see these correlations for Tinned/Bottled food.

Additional analysis of which factors are associated with the distribution of welfare. However, these correlations should be interpreted with caution, as they are very dependent on the assumptions that underpin the structural model.

4.5.3 Sequential versus Simultaneous Allocation

There are many differences between the Old System and the Choice System. I now investigate the relative importance of two key differences. This will help us understand which features of the Choice System are the most important, enabling us to potentially take these features to other food bank networks and other settings.

Simultaneity

First, Simultaneity. Under the Old System, food was allocated sequentially — every truckload of food was distributed before the next donation arrived. In contrast, under the Choice System multiple loads of food are allocated simultaneously.

[34]This relationship is small but worrying. Feeding America may not be setting budgets optimally - even a utilitarian social planner would equate marginal utility of wealth across food banks. I estimate a lot of variance in these parameters, shown in Figure L.4 in Appendix L.2. However the λ_is are not well identified in my model, resting on the strong assumption that lot-specific payoffs v_{ilt} have the same variances across food banks. This is certainly an area for future work.

Most other food bank networks allocate food sequentially, allocating them one at a time as they arrive.[35] This has a potential benefit over the simultaneous allocation in the Choice System as food banks do not risk winning too many or too few loads, alleviating the exposure problem highlighted in Gentry, Komarova, and Schiraldi (2023). However, as Akbarpour, Li, and Gharan (2020) and others have highlighted, waiting for a day then allocating all the donations simultaneously may yield better matches. In this setting, food banks may benefit from having information about everything being allocated that day. They do not risk accepting cornflakes when they really needed ready meals.

Auctions

The other key distinction is that under the Choice System, food is allocated using an auction format. This means food banks have a cardinal signal of their preferences, so can express the intensity of their values for a particular lot. In a standard one-shot, single-unit first price auction, this helps ensure that the lot always goes to the bidder that value it most. Likewise, in the multi-object setting, even though bidders are not able to fully signal their preferences (as without combination bids, simultaneous first-price auctions are not a direct revelation mechanism), this still helps the bidders ensure that they are more likely to win lots they value more.

Meanwhile, under the Old System, other than being able to turn down food they do not want, there is no relationship between the intensity of a food bank's preference for a particular lot and whether their receive it. To the best of my knowledge, Feeding America is the only food bank network that explicitly allows bidders to express a cardinal signal in this way.

Proposed Counterfactual

We have reason to believe that both allocating food simultaneously, as well as allocating food using an auction mechanism, are likely to be contributing factors to the

[35] Many food networks use modern technologies to make sequential allocation more feasible than it was for pre-Choice System Feeding America. Before 2005 Feeding America could only sequentially offer food to a very limited number of food banks, as it had to call each one and give them time to checks their stocks and decide whether to accept the food. Many organisations offer food using apps, and food banks make use of inventory management tools to quickly check the types of food they need.

welfare benefits of the Choice System. However, it is apriori unclear how much the welfare benefit depends on either of these factors over the other.

To investigate this question I consider two additional mechanisms. First, I consider a mechanism exactly the same as the Old System, but offering each truckload of food to every food bank until one of them accepts. This is an important because another key distinction between the Old and Choice System is that under the Choice System, every food bank can bid on every load. Next, I consider an efficient sequential mechanism. That is, every time Feeding America receives a donation, they always and immediately allocate it to whichever food bank values it the most *at that point in time*. Essentially, to whichever food bank has the highest marginal flow payoff as well as continuation value (given this mechanism). Importantly, this mechanism weakly dominates both sequential first and second price auctions, and yields the same outcomes in competitive equilibria.

By comparing the Old System to this "All Offer" System, we find the value of being able to offer every lot to (potentially) every food bank. Then, by comparing the All Offer System with this Efficient Sequential Mechanism we find the value of allocating food to the food bank with the highest value for that food, as opposed to offering it out essentially randomly. Essentially, this is informative of the value of allocation by auction. As we then compare this Efficient Sequential Mechanism to the Choice System, this is informative of the value of simultaneous allocation, since we are essentially comparing simultaneous and sequential auctions. And so, this additional counterfactual allows us to speak to the relative merits of using an auction mechanism and using a batch allocation mechanism.[36]

Details of how I numerically solve for equilibrium value functions and other aspects of the simulations are given in Appendix N.

Results

Welfare estimates from these two additional counterfactuals are displayed in Figure 4.5. Full details, including differences in the amount of food allocated, distance

[36] Ideally, I would also consider a simultaneous, non-auction mechanism. For example, where food banks choose a selection of lots they would like, then lots are randomly allocated among food banks who chose them. This would enable me to investigate whether auctions and simultaneous allocation are complements, substitutes, or neither. Unfortunately simulating a dynamic equilibrium under such a mechanism is computationally difficult.

FIGURE 4.5: Simultaneous vs Sequential Allocation

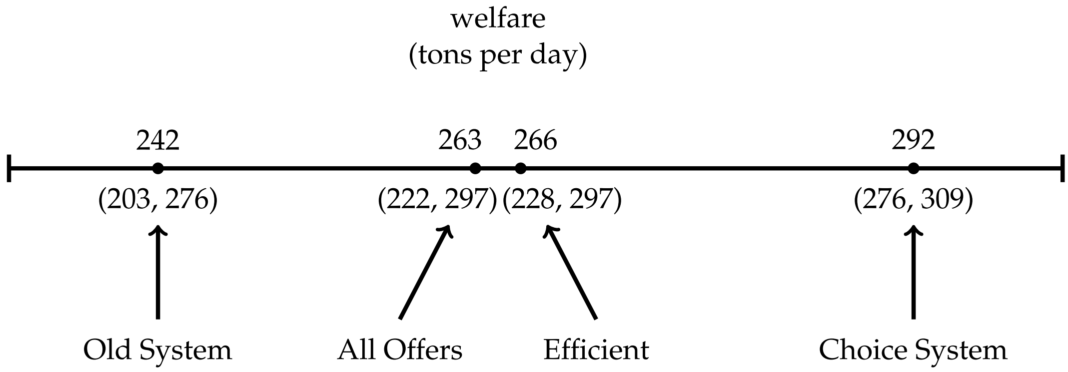

Note: This figure shows estimated welfare, in equivalent tons of food allocated each day, under the Old System, the Old System when food is sequentially offered to every food bank (not just the first 10), an Efficient Sequential mechanism (such as sequential first-price auctions), and the Choice System. 95% Credible Intervals are given in parenthesis. The (marginal) Credible Intervals do not indicate the degree of (positive) correlation between welfare estimates, and only demonstrate marginal distributions.

travelled, and distributional effects, are presented in Figure 4.6.

The results show that moving from the Old System to the All Offer System leads to an increase in welfare equivalent to increasing the supply of distributed food by 21 tons each day, with a credible interval of $(11, 24)$. Therefore, being able to offer food to every food bank explains around 40% of the welfare change as Feeding America moved from the Old System to the Choice System. However, moving from the All Offer mechanism to the Efficient Sequential Mechanism only leads to an increase in welfare equivalent to a 3 ton daily increase $(-4, 11)$. Finally, moving from the Efficient Sequential mechanism, to the Choice System, leads to an increase in welfare of 26 tons per day $(5, 31)$.

This is a particularly striking result, suggesting that the benefits of using an auction to allocate food are practically zero, so long as food is still being allocated sequentially. In harsher terms, it does not matter how sophisticated your allocation mechanism — if you are still allocating food sequentially, you may as well be offering it out at random. The reason for this result is that even though food is always allocated to the food bank who thinks they need it most in that moment. But it may not be the food most needed by the food bank at that time, and they may forfeit more needed food later in the day because storage is now more costly. This is very similar to the key reasons identified for the value of the Choice System over the Old

System, that food banks were accepting sub-optimal food.[37]

4.5.4 Alternative Mechanisms

The allocation problem faced by Feeding America is faced by numerous other food bank organisations around the world, such as the European Foodbank Federation (FEBA), FareShare (U.K.), and Food bank Australia.[38] To my knowledge no other network employs an allocation mechanism that gives food banks nearly as much choice. Therefore an important question concerns how their mechanisms compare to the Choice System, and to what extent they might be able to benefit by giving food banks more choice over the types of food they receive. For example, many organisations (including the U.K.'s Trussell Trust) do not allocate food centrally at all, but link up food banks with additional local donors. This could be an extremely inefficient use of resources, particularly if food banks who are slightly further away might be willing to pay the added transportation costs.

Therefore, I consider three additional counterfactual mechanisms. These are designed to emulate other mechanisms used in practice around the world. First, I consider the 'Closest' System. Under the Closest System, the lot is offered to the closest food bank, and no others. This mechanism is used implicitly by food networks who do not allocate food centrally, and instead link food banks up with additional local donors. Instinctively, if food banks have high transportation costs, so that distance is a (if not *the*) deciding factor, this mechanism may perform well, at least relative to the Old System. However, it is probably that there may exist a food bank slightly further away with a larger willingness to pay to transport the lot.

Second, in the interest of comparability, I the "Closest, All Offer" mechanism under which food is offered to food banks in ascending order of the distance to the lot. While not used, to my knowledge, by any food bank networks, it still provides an interesting basis for comparison. Finally, I consider the "Like" mechanism proposed

[37]It should be noted that in reality, the Choice System allocates using two rounds of simultaneous auctions each period, not just one as I am forced to assume. Therefore, these results should be interpreted with a grain of salt. That said, it is unclear why the bias in my results from this simplification would spurriously cause such striking results.

[38]It is worth recognising that these organisations often face different allocation problems to Feeding America. For example, transport costs are much more pertinent in Australia. Likewise, many of these organisations face a problem closer to the scale of individual food banks allocating food among food pantries. Nonetheless these results remain a useful starting point in analysing the efficacy of their mechanisms and proposed changes.

by Walsh (2015). Under this mechanism each lot is offered to every food bank simultaneously. Food banks can "Like" the lot, or not. The lot is randomly (weighted by Goal Factor) assigned to one of the food banks who liked it. While neither of these mechanisms are used explicitly by other food bank networks, they give a good overview of some of these types of mechanisms. For example, the mechanism used by the Foodiverse platform in Ireland first offers lots to the nearest food bank, and then offers it to all the other food banks simultaneously as per the Like mechanism.

Given the results presented in the previous section, an obvious downside of all these mechanism is that they are sequential in nature. However, given the mechanisms used in practice, results from these counterfactuals may provide useful results for policy purposes. A detailed discussion of the mechanisms and how I solve for the equilibrium value functions is given in Appendix N.

Results

FIGURE 4.6: Counterfactual Results (2)

Mechanism	Welfare (unweighted)	Welfare (weighted)	Distance (000 km per day)	Allocated (tons per day)	% Better Off (under CS)
Choice	292	745	16	271	0
System	(276, 309)	(672, 815)	(14.6, 17.4)	(253, 284)	(0, 0)
Old	242	576	22.3	249	0.85
System	(203, 276)	(415, 706)	(22, 22.6)	(248, 251)	(0.803, 0.893)
Old System	263	649	19.4	264	0.799
All offers	(222, 297)	(477, 785)	(19, 20)	(263, 266)	(0.738, 0.852)
Closest	134	311	0.305	58	0.945
	(75.7, 183)	(43.9, 530)	(0.298, 0.311)	(57.4, 58.6)	(0.918, 0.967)
Closest	258	632	14	265	0.738
All offers	(218, 294)	(460, 772)	(13.7, 14.3)	(264, 267)	(0.676, 0.795)
Like	264	653	19.4	264	0.819
	(226, 296)	(489, 784)	(19, 20)	(263, 266)	(0.762, 0.877)
Efficient	266	661	19.2	246	0.718
	(228, 297)	(510, 783)	(18.7, 19.8)	(244, 249)	(0.656, 0.779)

Note: This table displays posterior means and 95% credible intervals for various measures of welfare. The final column gives the percentage of food banks who are estimated to be (weakly) better off under the Choice System than each alternative mechanism. A higher number is worse.

Welfare estimates are presented in Figure 4.6. I present results for each of the 7 mechanism I have considered so far. I present results for welfare estimates (both weighted and unweighted), average daily distance travelled, average quantity of food allocated, as well as the proportion of food banks who are better off under the Choice System than under each mechanism I consider. This final variable is useful for understanding how the distribution of welfare would change if a food bank network using, say, the "Closest" mechanism, transitioned to allocating food using the Choice System. What proportion of food banks should we expect to benefit from this transition?

The key takeaway from this analysis is that the Closest mechanism performs very poorly. It achieves only 46.2% of the Welfare under the Choice System, and only around two thirds of the welfare under the Old System. This is indicative of the welfare benefits from centralised allocation. Food banks who are slightly further away from the donor are very often willing to pay the added transport costs. Even when food is offered to every food bank, in order of distance, this is still worse than both the "All Offer" Old System and the Like Mechanism. This is driven by the finding that food banks are not particularly sensitive to distance.[39]

Another, somewhat unsurprising takeaway, is that there are significant benefits from offering food to every food bank, as in the Like and both of the 'all offers' mechanisms, all achieving around a 5% increase on the welfare from the standard Old System. This is very intuitive, since it ensures food is not turned away when there is still a food bank that needs it. This highlights that, where possible, food bank networks should ensure that food is at least offered to every food bank before it is returned to the donor.

[39]It is true that transportation costs are likely a larger factor in other settings, meaning these results cannot immediately be applied to another food bank network. However it does highlight that there should be a serious discussion of how important transportation costs actually are in any given setting.

Chapter 5

Conclusion

In this thesis I examined the importance of giving food banks choice over the types of food they are allocated. I investigated this question by examining the welfare and distributional consequences of Feeding America's implementation of their 'Choice System'. The Choice System facilitated choice by allocating food using repeated rounds of simultaneous first-price auctions, allowing food banks to signal their preferences over various types of food and giving them control over their allocations. This allocation mechanism replaced the Old System which greatly restricted food bank choice, essentially only offering food out at random.

The importance of choice depends on the degree of heterogeneity in the types of food needed by different food banks. By examining variation in bidding behaviour I presented reduced form evidence of heterogeneity across different types of food, across different food banks, and within food banks over time. This is evidence that choice is valuable, and that food banks likely benefited from the transition to the Choice System. However, in order to perform welfare analysis considering the distributional effects of the transition, or to investigate the key mechanisms driving the welfare benefits, I needed to use a structural model.

To perform welfare analysis, given I only had access to data from the Choice System, I developed a dynamic multi-object auction model allowing me to model forward looking bidding behaviour in these simultaneous first-price auctions. My modelling framework unified the dynamic auction model of Jofre-Bonet and Pesendorfer (2003) with the multi-object auction model of Gentry, Komarova, and Schiraldi (2023). Because of the dynamic environment I was unable to rely on the identification arguments in Gentry, Komarova, and Schiraldi (2023) which leveraged exclusion restrictions. Therefore I proved non-parametric point identification of the model primitives by leveraging observed variation in state variables. Because of the multi-object setting (specifically, that simultaneous auctions are not a direct revelation mechanism), I was unable to apply the estimation procedure presented in Jofre-Bonet and Pesendorfer (2003). Therefore I proposed a novel estimation procedure, essentially generalising Jofre-Bonet and Pesendorfer (2003)'s method to the multi-object setting. The procedure is computationally convenient, and only involves performing one additional estimation step in which we estimate a misspecified static multi-object model. I then presented various extensions to this framework that capture other key elements of the Choice System, including endogenous entry, reservation prices, and the inter-temporal budget constraint.

In Chapter 4 I applied this model to Choice System data. To capture the possible benefits of choice I specified a model that allowed for a broad degree of heterogeneity across different types of food and across food banks. I modelled heterogeneity over time by recognising that food banks' stocks, particularly the food sent out to food pantries and their local donations, are likely to be unobserved and time varying. I then extended the framework presented in Chapter 3 to allow for these unobserved states. Rather than using observed stocks, I proved that this model is nonparametrically identified using variation in food banks' winnings, which are just observed changes in these unobserved stocks. I proposed a similar adjustment to the estimation procedure using a Gibbs Sampler, making use of data augmentation algorithms to sample the unobserved stocks from their conditional distribution. Having estimated the model, this allowed me to investigate the degree of heterogeneity across the types of food Feeding America allocates, as well as heterogeneity in food banks' needs - both across food banks and within food banks over time. This heterogeneity is important to understand how food banks' needs are determined by the types of food they have access to from their local donors.

I found that welfare was 17.1% higher under the Choice System than under the

Old System. These results are driven by this heterogeneity, particularly heterogeneity over time. Choice allows food banks to focus their allocations on their most pressing needs, whereas under the Old System they might be offered food that was more useful to a different food bank at that particular time. I estimate that 85% of food banks are better off under the Choice System. The largest benefits are seen by food banks with the fewest and most variable local donations, benefiting from the flexibility the Choice System permits. This study has important policy implications, both for Feeding America and other food bank networks around the world. I find that welfare under the Choice System significantly exceeds welfare under a number of alternative mechanisms. I found particularly poor welfare consequences of sending food only to the nearest food bank, and that mechanisms which allocate food sequentially as donations arrive are very limited in their efficacy. These finding highlight the importance of good market design.

Future work should consider the external validity of these results, and their applicability in other food bank settings. For example, applying the analysis to data from other food bank networks. Future work should also consider additional mechanisms, potentially building on the Choice System, for application in these other settings and even perhaps improving on the important work already done for Feeding America's allocation problem.

Bibliography

Abdulkadiroğlu, Atila, Nikhil Agarwal, and Parag A Pathak (2017). "The welfare effects of coordinated assignment: Evidence from the New York City high school match". In: *American Economic Review* 107.12, pp. 3635–89.

Afriat, Sidney N (1972). "Efficiency estimation of production functions". In: *International Economic Review*, pp. 568–598.

Agarwal, Nikhil (2015). "An empirical model of the medical match". In: *American Economic Review* 105.7, pp. 1939–78.

Agarwal, Nikhil and Eric Budish (2021). "Market design". In: *Handbook of Industrial Organization*. Vol. 5. 1. Elsevier, pp. 1–79.

Agarwal, Nikhil, Charles Hodgson, and Paulo Somaini (2020). *Choices and outcomes in assignment mechanisms: The allocation of deceased donor kidneys*. Tech. rep. National Bureau of Economic Research.

Agarwal, Nikhil and Paulo Somaini (2020). "Revealed preference analysis of school choice models". In: *Annual Review of Economics* 12, pp. 471–501.

Agarwal, Nikhil et al. (2021). "Equilibrium allocations under alternative waitlist designs: Evidence from deceased donor kidneys". In: *Econometrica* 89.1, pp. 37–76.

Akbarpour, Mohammad, Shengwu Li, and Shayan Oveis Gharan (2020). "Thickness and information in dynamic matching markets". In: *Journal of Political Economy* 128.3, pp. 783–815.

Allen, Roy and John Rehbeck (2022). "Latent complementarity in bundles models". In: *Journal of Econometrics* 228.2, pp. 322–341.

Arcidiacono, Peter and Robert A Miller (2011). "Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity". In: *Econometrica* 79.6, pp. 1823–1867.

Atchadé, Yves F and Jeffrey S Rosenthal (2005). "On adaptive markov chain monte carlo algorithms". In: *Bernoulli* 11.5, pp. 815–828.

Athey, Susan and Philip A Haile (2002). “Identification of standard auction models”. In: *Econometrica* 70.6, pp. 2107–2140.

— (2007). “Nonparametric approaches to auctions”. In: *Handbook of econometrics* 6, pp. 3847–3965.

Athey, Susan, Jonathan Levin, and Enrique Seira (2011). “Comparing open and sealed bid auctions: Evidence from timber auctions”. In: *The Quarterly Journal of Economics* 126.1, pp. 207–257.

Aumann, Robert J (1974). “Subjectivity and correlation in randomized strategies”. In: *Journal of mathematical Economics* 1.1, pp. 67–96.

Baccara, Mariagiovanna, SangMok Lee, and Leeat Yariv (2020). “Optimal dynamic matching”. In: *Theoretical Economics* 15.3, pp. 1221–1278.

Backus, Matthew and Gregory Lewis (2016). *Dynamic demand estimation in auction markets*. Tech. rep. National Bureau of Economic Research.

Balat, Jorge (2013). “Highway procurement and the stimulus package: Identification and estimation of dynamic auctions with unobserved heterogeneity”. In: *Johns Hopkins University Mimeo*.

Balat, Jorge et al. (2015). “Dynamic and strategic behavior in hydropower-dominated electricity markets: empirical evidence for Colombia”. In: *Borradores de Economía; No. 886*.

Bazerghi, Chantelle, Fiona H McKay, and Matthew Dunn (2016). “The role of food banks in addressing food insecurity: a systematic review”. In: *Journal of community health* 41, pp. 732–740.

Berry, Steven and Ariel Pakes (2007). “The pure characteristics demand model”. In: *International Economic Review* 48.4, pp. 1193–1225.

Berry, Steven T and Giovanni Compiani (2020). *An instrumental variable approach to dynamic models*. Tech. rep. National Bureau of Economic Research.

Bhattacharya, Rabi N and Mukul Majumdar (1989). “Controlled semi-markov models-the discounted case”. In: *Journal of Statistical Planning and Inference* 21.3, pp. 365–381.

Bodoh-Creed, Aaron L, Joern Boehnke, and Brent Hickman (2021). “How efficient are decentralized auction platforms?” In: *The Review of Economic Studies* 88.1, pp. 91–125.

Botev, Zdravko I (2017). "The normal law under linear restrictions: simulation and estimation via minimax tilting". In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 79.1, pp. 125–148.

Budish, Eric and Estelle Cantillon (2012). "The multi-unit assignment problem: Theory and evidence from course allocation at Harvard". In: *American Economic Review* 102.5, pp. 2237–71.

Byrne, Anne T and David R Just (2022). "Private food assistance in high income countries: A guide for practitioners, policymakers, and researchers". In: *Food Policy* 111, p. 102300.

Cantillon, Estelle and Martin Pesendorfer (2007). "Combination bidding in multi-unit auctions". In:

Caplan, Pat (2016). "Big society or broken society?: Food banks in the UK". In: *Anthropology Today* 32.1, pp. 5–9.

Carroll, Raymond J, Xiaohong Chen, and Yingyao Hu (2010). "Identification and estimation of nonlinear models using two samples with nonclassical measurement errors". In: *Journal of nonparametric statistics* 22.4, pp. 379–399.

Carter, Chris K and Robert Kohn (1994). "On Gibbs sampling for state space models". In: *Biometrika* 81.3, pp. 541–553.

Charalambos, D and Border Aliprantis (2013). *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer-Verlag Berlin and Heidelberg GmbH & Company KG.

Chiappori, Pierre-André and Bernard Salanié (2016). "The Econometrics of Matching Models". In: *Journal of Economic Literature* 54.3, pp. 832–61. DOI: 10.1257/jel.20140917. URL: https://www.aeaweb.org/articles?id=10.1257/jel.20140917.

Cloke, Paul, Jon May, and Andrew Williams (2017). "The geographies of food banks in the meantime". In: *Progress in Human Geography* 41.6, pp. 703–726.

Combe, Julien, Olivier Tercieux, and Camille Terrier (2022). "The design of teacher assignment: Theory and evidence". In: *The Review of Economic Studies* 89.6, pp. 3154–3222.

Connault, Benjamin (2014). "Hidden rust models". In: *Job Market Paper*.

Deaton, Angus and John Muellbauer (1980). "An almost ideal demand system". In: *The American economic review* 70.3, pp. 312–326.

Dunford, N and J Schwartz (1971). *Linear Operators*. Vol. 3. Wiley.

Engle, Robert F and Clive WJ Granger (1987). "Co-integration and error correction: representation, estimation, and testing". In: *Econometrica: journal of the Econometric Society*, pp. 251–276.

Erdem, Tülin, Susumu Imai, and Michael P Keane (2003). "Brand and quantity choice dynamics under price uncertainty". In: *Quantitative Marketing and economics* 1.1, pp. 5–64.

Fox, Jeremy T and Patrick Bajari (2013). "Measuring the efficiency of an FCC spectrum auction". In: *American Economic Journal: Microeconomics* 5.1, pp. 100–146.

Fox, Jeremy T and Natalia Lazzati (2017). "A note on identification of discrete choice models for bundles and binary games". In: *Quantitative Economics* 8.3, pp. 1021–1036.

Fudenberg, Drew and Eric Maskin (1991). "On the dispensability of public randomization in discounted repeated games". In: *Journal of Economic Theory* 53.2, pp. 428–438.

— (2009). "The folk theorem in repeated games with discounting or with incomplete information". In: *A Long-Run Collaboration On Long-Run Games*. World Scientific, pp. 209–230.

Gandhi, Ashvin (2019). "Picking your patients: Selective admissions in the nursing home industry". In: *Available at SSRN 3613950*.

Gelman, Andrew et al. (1995). *Bayesian data analysis*. Chapman and Hall/CRC.

Gentry, Matthew, Tatiana Komarova, and Pasquale Schiraldi (2023). "Preferences and performance in simultaneous first-price auctions: A structural analysis". In: *The Review of Economic Studies* 90.2, pp. 852–878.

Groeger, Joachim R (2014). "A study of participation in dynamic auctions". In: *International Economic Review* 55.4, pp. 1129–1154.

Guerre, Emmanuel, Isabelle Perrigne, and Quang Vuong (2000). "Optimal nonparametric estimation of first-price auctions". In: *Econometrica* 68.3, pp. 525–574.

Gundersen, C et al. (2017). *Map the Meal Gap 2017: A Report on County and Congressional District Food Insecurity and County Food Cost in the United States in 2015.*

Hendel, Igal and Aviv Nevo (2006). "Measuring the implications of sales and consumer inventory behavior". In: *Econometrica* 74.6, pp. 1637–1673.

Hendricks, Kenneth and Alan Sorensen (2015). "The role of intermediaries in dynamic auction markets". In: *Working Paper*.

Hortaçsu, Ali and David McAdams (2010). "Mechanism choice and strategic bidding in divisible good auctions: An empirical analysis of the turkish treasury auction market". In: *Journal of Political Economy* 118.5, pp. 833–865.

— (2018). "Empirical work on auctions of multiple objects". In: *Journal of Economic Literature* 56.1, pp. 157–84.

Hotz, V Joseph and Robert A Miller (1993). "Conditional choice probabilities and the estimation of dynamic models". In: *The Review of Economic Studies* 60.3, pp. 497–529.

Hu, Yingyao and Susanne M Schennach (2008). "Instrumental variable treatment of nonclassical measurement error models". In: *Econometrica* 76.1, pp. 195–216.

Hu, Yingyao and Matthew Shum (2012). "Nonparametric identification of dynamic models with unobserved state variables". In: *Journal of Econometrics* 171.1, pp. 32–44.

Iaria, Alessandro and Ao Wang (2020). "Identification and estimation of demand for bundles". In:

Ifrach, Bar and Gabriel Y Weintraub (2017). "A framework for dynamic oligopoly in concentrated industries". In: *The Review of Economic Studies* 84.3, pp. 1106–1150.

Jeziorski, Przemyslaw and Elena Krasnokutskaya (2016). "Dynamic auction environment with subcontracting". In: *The RAND Journal of Economics* 47.4, pp. 751–791.

Jofre-Bonet, Mireia and Martin Pesendorfer (2003). "Estimation of a dynamic auction game". In: *Econometrica* 71.5, pp. 1443–1489.

Kang, Boo-Sung and Steven L Puller (2008). "The effect of auction format on efficiency and revenue in divisible goods auctions: A test using Korean treasury auctions". In: *The Journal of Industrial Economics* 56.2, pp. 290–332.

Kasahara, Hiroyuki and Katsumi Shimotsu (2009). "Nonparametric identification of finite mixture models of dynamic discrete choices". In: *Econometrica* 77.1, pp. 135–175.

Kim, Sang Won, Marcelo Olivares, and Gabriel Y Weintraub (2014). "Measuring the performance of large-scale combinatorial auctions: A structural estimation approach". In: *Management Science* 60.5, pp. 1180–1201.

Kinach, Lesia, Kate Parizeau, and Evan DG Fraser (2020). "Do food donation tax credits for farmers address food loss/waste and food insecurity? A case study from Ontario". In: *Agriculture and Human Values* 37, pp. 383–396.

Kong, Yunmi (2021). "Sequential auctions with synergy and affiliation across auctions". In: *Journal of Political Economy* 129.1, pp. 148–181.

Krasnokutskaya, Elena (2011). "Identification and estimation of auction models with unobserved heterogeneity". In: *The Review of Economic Studies* 78.1, pp. 293–327.

Leib, EB et al. (2017). "Don't waste, donate: Enhancing food donations through federal policy". In: *NRDC*.

Lewbel, Arthur (2019). "The identification zoo: Meanings of identification in econometrics". In: *Journal of Economic Literature* 57.4, pp. 835–903.

Li, Tong, Isabelle Perrigne, and Quang Vuong (2000). "Conditionally independent private information in OCS wildcat auctions". In: *Journal of econometrics* 98.1, pp. 129–161.

— (2002). "Structural estimation of the affiliated private value auction model". In: *RAND Journal of Economics*, pp. 171–193.

Li, Tong and Xiaoyong Zheng (2009). "Entry and competition effects in first-price auctions: Theory and evidence from procurement auctions". In: *The Review of Economic Studies* 76.4, pp. 1397–1429.

Liu, Tracy Xiao, Zhixi Wan, and Chenyu Yang (2019). "The efficiency of a dynamic decentralized two-sided matching market". In: *Available at SSRN* 3339394.

Loopstra, Rachel et al. (2015). "Austerity, sanctions, and the rise of food banks in the UK". In: *Bmj* 350.

Magnac, Thierry and David Thesmar (2002). "Identifying dynamic discrete decision processes". In: *Econometrica* 70.2, pp. 801–816.

Milgrom, Paul R and Robert J Weber (1985). "Distributional strategies for games with incomplete information". In: *Mathematics of operations research* 10.4, pp. 619–632.

Newey, Whitney K and James L Powell (2003). "Instrumental variable estimation of nonparametric models". In: *Econometrica* 71.5, pp. 1565–1578.

NPR (2015). *Planet Money, Accessed: 2018-04-19*. https://www.npr.org/templates/transcript/transcript.php

Pesendorfer, Martin and Philipp Schmidt-Dengler (2008). "Asymptotic least squares estimators for dynamic games". In: *The Review of Economic Studies* 75.3, pp. 901–928.

Prendergast, Canice (2017). "How Food Banks Use Markets to Feed the Poor". In: *Journal of Economic Perspectives* 31.4, pp. 145–62.

— (2022). "The allocation of food to food banks". In: *Journal of Political Economy* 130.8, pp. 000–000.

Raisingh, Diwakar (2021). "The Effect of Pre-announcements on Participation and Bidding in Dynamic Auctions". In:

Reguant, Mar (2014). "Complementary bidding mechanisms and startup costs in electricity markets". In: *The Review of Economic Studies* 81.4, pp. 1708–1742.

Riches, Graham (2002). "Food banks and food security: welfare reform, human rights and social policy. Lessons from Canada?" In: *Social Policy and Administration* 36.6, pp. 648–663.

Robinson Cortés, Alejandro (2020). "Essays on Market Design and Industrial Organization". PhD thesis. California Institute of Technology.

Rust, John (1987). "Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher". In: *Econometrica: Journal of the Econometric Society*, pp. 999–1033.

Sailer, Katharina (2006). "Searching the eBay marketplace". In: *CESifo working paper*.

Samuelson, Paul A (1947). "Foundations of Economic Analysis. Atheneum, New York, 1965". In: *Google Scholar*.

Somaini, Paulo (2020). "Identification in auction models with interdependent costs". In: *Journal of Political Economy* 128.10, pp. 3820–3871.

Verdier, Valentin and Carson Reeling (2022). "Welfare effects of dynamic matching: An empirical analysis". In: *The Review of Economic Studies* 89.2, pp. 1008–1037.

Waldinger, Daniel (2021). "Targeting in-kind transfers through market design: A revealed preference analysis of public housing allocation". In: *American Economic Review* 111.8, pp. 2660–96.

Walsh, Toby (2015). "Challenges in Resource and Cost Allocation." In: *AAAI*, pp. 4073–4077.

Weintraub, Gabriel Y, C Lanier Benkard, and Benjamin Van Roy (2008). "Markov perfect industry dynamics with many firms". In: *Econometrica* 76.6, pp. 1375–1411.

Williams, Andrew and Jon May (2022). "A genealogy of the food bank: Historicising the rise of food charity in the UK". In: *Transactions of the Institute of British Geographers* 47.3, pp. 618–634.

Williams, Andrew et al. (2016). "Contested space: The contradictory political dynamics of food banking in the UK". In: *Environment and Planning A: Economy and Space* 48.11, pp. 2291–2316.

Appendix A

Data

In this Appendix I present additional details on how I constructed the dataset used in my analysis. Appendix A.1 focuses on the Choice System data, received from Feeding America, outlining how I cleaned and categorised the data. This Appendix also details how I identified joint bidding and which lots were sold by food banks ('Maroon Pounds'). In Appendix A.2 I detail the auxiliary datasets used in my analysis, used to locate food banks and construct their Goal Factors.

A.1 Choice System data

The variables included in the data were as follows:

1. A unique auction identifier
2. Date of the auction
3. Details on all the goods included
4. The number of Pounds in the lot
5. The number of identical lots being auctioned
6. Bids placed on the lot
7. anonymised Foodbank ID that placed each bid
8. The winning Bid(s)
9. An indicator stating whether the Auction/Bid was cancelled
10. The geographic location of the lot

Food banks were anonymised and indexed from 1 to 165. I did not observe whether a bid was placed jointly, nor whether a load was sold by a food bank. I also do not observe whether an auction occurred in the morning or in the afternoon. Finally, I only observe an auction on a particular date if at least one bid is received on that lot. If an auction that consists of two identical loads is observed with just one bid on day t, and another observation with just one bid on day $t + 2$ I assume that the auction also appeared on day $t + 1$ with only one load available. I must assume that it is not the case that auction appeared on day $t - 1$ but no one placed a bid.

A.1.1 data cleaning

The 1344 cancelled auctions and bids were removed from the data, with the assumption that bidding behaviour was not affected by cancellations.

There were various errors in the record data. Some errors could be corrected, such as misspelt names of products, while several had to be removed. Every load listed as being heavier than $97,000$ pounds (the maximum weight for a flat bed truck) was assumed to be a mistake, and fixed to $40,000$ pounds (the modal weight). Every load weighing less than 5000 pounds was also fixed to $40,000$. 6 auctions were removed from the data. These lots included items such as karaoke machines and a flat bed truck. These lots were removed under the assumption that the they fit outside the food banks' ordinary remits.

A.1.2 categorisation

Goods are classified into categories (mostly taken from Prendergast, 2022), subcategories,[1] Uses and Storage Method.

'Uses' are not used in the current model specifications. Uses includes Meals, Ingredients, Condiments, Snacks, and Non-Food. Meals are items that could be eaten on its own as part of a reasonably healthy diet for either breakfast, lunch, or dinner. Multiple Ingredients can be mixed together to form a meal. Condiments can be added to a meal to enhance it. Snacks can be eaten on their own, though not necessarily part of a meal. Snacks includes drinks. Non-food items are inedible items, such as cleaning products. This also includes formula and baby food.

Storage methods includes Shelf, Tinned, Refrigerated, Fresh, and Non-Food. The Non-food category is identical to the non-Food Use category. Shelf items can be

[1] This is performed to ensure at least 30 lots per subcategory. Subcategories are more granular the more observations there are. E.g. for cereal and beverages this includes brands, whereas all cheese is lumped together.

stored on a shelf, are generally dried goods, and have extremely long shelf lives. They are generally light but bulky. Tinned food, which includes jars and bottles, have long shelf-lives and are generally compact and heavy. Refrigerated food must be stored in a fridge, but still expire reasonably quickly. Fresh food is food one wouldn't generally store in a fridge, and generally has only a limited shelf life. This includes both fresh produce and freshly baked goods such as bread. Any item that was additionally listed as 'Shelf Stable', such as UHT milk was put in the tinned storage category.

Code	Category	Subcategory	Use	Storage	Description
1	Baby		Non-food	N F	unspecified baby
2	Baby	diaper	Non-food	N F	nappies
3	Baby	food	Non-food	N F	Baby food
4	Baby	formula	Non-food	N F	Baby formula
5	Beverage		Snack	Tin	unspecified Bev
6	Beverage	capri	Snack	Tin	capri-sun
7	Beverage	coffee	Snack	Shelf	ground/instant
8	Beverage	dry	Snack	Shelf	chocolate/milk powder
9	Beverage	fj	Snack	Tin	Orange/Apple/Grape juice - high quality, "pure"
10	Beverage	gator	Snack	Tin	Sports drink (gatorade)
11	Beverage	ic	Snack	Tin	Iced/Alternate coffee
12	Beverage	juice	Snack	Tin	juices, lower quality, mixed, e.g. tropical punch, fruit shoot
13	Beverage	ka	Snack	Shelf	Kool-Aid
14	Beverage	pop	Snack	Tin	fizzy drinks, e.g. coke
15	Beverage	propel	Snack	Tin	Propel brand water/sports water
16	Beverage	pshake	Snack	Shelf	Protein shake/powder
17	Beverage	shake	Snack	Tin	Milk shakes
18	Beverage	tea	Snack	Shelf	Tea/Tea bags
19	Beverage	vf	Snack	Tin	V8 juices
20	Beverage	water	Snack	Tin	Bottled water
21	Baked Good		Snack	Shelf	unspecified BP
22	Baked Good	bread	Ingredient	Fresh	Bread
23	Baked Good	cake	Meal	Fresh	cake, cupcakes, muffins
24	Baked Good	dough	Ingredient	Fridge	cookie dough, bread dough, etc
25	Baked Good	flour	Ingredient	Shelf	flour, cake mix, bread mix
26	Baked Good	other	Snack	Shelf	miscellaneous BP
27	Baked Good	pastry	Snack	Fresh	croissants, waffles, pancakes etc
28	Baked Good	stuffing	Condiment	Shelf	Stuffing mix
29	Cereal		Meal	Shelf	unspecified cereal
30	Cereal	bran	Meal	Shelf	healthy bran cereal (fibre)
31	Cereal	cheerio	Meal	Shelf	Cheerios
32	Cereal	flake	Meal	Shelf	un-sweatened flakes (spK, corn etc)
33	Cereal	gran	Meal	Shelf	granola
34	Cereal	Kashi	Meal	Shelf	unspecified Kashi
35	Cereal	Kellogg	Meal	Shelf	unspecified Kellogg
36	Cereal	ns	Meal	Shelf	non-sugared cereal (rice-krispies)
37	Cereal	other	Meal	Fridge	miscellaneous non-dry cereal
38	Cereal	oat	Meal	Shelf	oats/grits/porridge
39	Cereal	PL	Meal	Shelf	unspecified Private Label (e.g. Post)
40	Cereal	sugar	Meal	Shelf	fruit loops, apple jacks etc
41	Condiment		Condiment	Tin	unspecified condiments
42	Condiment	dressing	Condiment	Tin	Salad dressings, glazes
43	Condiment	fruit	Condiment	Tin	Fruit sauces, preserves
44	Condiment	gravy	Condiment	Shelf	Gravy granules
45	Condiment	jelly	Condiment	Tin	Jam
46	Condiment	ketchup	Condiment	Tin	Ketchup
47	Condiment	mayo	Condiment	Tin	Mayonnaise
48	Condiment	mustard	Condiment	Tin	Mustard
49	Condiment	other	Condiment	Tin	Miscellaneous cond (e.g. frosting)
50	Condiment	oil	Ingredient	Tin	Cooking oils
51	Condiment	pasta	Condiment	Tin	Pasta sauces
52	Condiment	PB	Condiment	Tin	Peanut Butter
53	Condiment	pickle	Condiment	Tin	Pickled Gherkins
54	Condiment	salsa	Condiment	Tin	Salsa/Guacamole/dips
55	Condiment	sauce	Condiment	Tin	BBQ sauce, etc
56	Condiment	stock	Ingredient	Shelf	Stock (assumed cube form)
57	Dairy		Ingredient	Fridge	unspecified Dairy
58	Dairy	butter	Condiment	Fridge	Butter/Margarine/Spread
59	Dairy	cc	Condiment	Fridge	coffee-creamer, coffee-mate (assumed liquid)
60	Dairy	cheese	Ingredient	Fridge	mostly cottage/cream cheese
61	Dairy	cream	Condiment	Fridge	Mostly sour cream
62	Dairy	dessert	Meal	Fridge	cheese cake etc
63	Dairy	egg	Ingredient	Fresh	eggs
64	Dairy	ll	Ingredient	Tin	evaporated/preserved milk
65	Dairy	milk	Condiment	Fridge	milk
66	Dairy	milk-alt	Condiment	Fridge	non-Dairy milk
67	Dairy	milk-flav	Snack	Fridge	flavoured (chocolate) milk
68	Dairy	pie	Meal	Fridge	Sweet pies, e.g. Apple/custard
69	Dairy	yog	Snack	Fridge	Yoghurt

70	Fresh		Ingredient	Fresh	unspecified produce
71	Fresh	apple	Snack	Fresh	Apples
72	Fresh	cabbage	Ingredient	Fresh	Cabbages
73	Fresh	carrot	Ingredient	Fresh	Carrots
74	Fresh	citrus	Snack	Fresh	Citrus fruits
75	Fresh	corn	Ingredient	Fresh	Corn (maize)
76	Fresh	fruit	Snack	Fresh	unspec/misc fruit
77	Fresh	melon	Snack	Fresh	melons
78	Fresh	other	Ingredient	Fresh	Miscellaneous veg
79	Fresh	onion	Ingredient	Fresh	Onions or garlic
80	Fresh	potato	Ingredient	Fresh	Potatoes
81	Fresh	squash	Ingredient	Fresh	Squash/Pumpkin,Yams
82	Frozen		Ingredient	Fridge	unspecified/misc frozen
83	Frozen	bp	Snack	Fridge	Frozen Baked Goods, e.g. bread rolls
84	Frozen	dairy	Ingredient	Fridge	Frozen milk, butter, eggs
85	Frozen	meal	Meal	Fridge	Frozen meals/pies/pizza
86	Frozen	meat	Ingredient	Fridge	Frozen chickens etc
87	Frozen	veg	Ingredient	Fridge	peas, carrots etc
88	Health/Beauty		Non-Food	Non-Food	unspecified HBC
89	Health/Beauty	body	Non-Food	Non-Food	body creams/moisturiser
90	Health/Beauty	dental	Non-Food	Non-Food	dental hygiene
91	Health/Beauty	deod	Non-Food	Non-Food	deodorant
92	Health/Beauty	detergent	Non-Food	Non-Food	detergent powder/tablets
93	Health/Beauty	drug	Non-Food	Non-Food	medicines/ointments
94	Health/Beauty	nutri	Non-Food	Non-Food	vitamins / unspecified nutritional items (e.g. protein powder)
95	Health/Beauty	other	Non-Food	Non-Food	miscellaneous (e.g. razors)
96	Health/Beauty	shampoo	Non-Food	Non-Food	shampoo/conditioner
97	Health/Beauty	soap	Non-Food	Non-Food	hand/body soap
98	Health/Beauty	sun	Non-Food	Non-Food	sun-cream/block
99	Meal		Meal	Fridge	unspecified Meals
100	Meal	bert	Meal	Fridge	Bertolli ready meals
101	Meal	breakfast	Meal	Fridge	breakfast meals
102	Meal	broth	Ingredient	Tin	Broth - assumed carton stock
103	Meal	cb	Meal	Tin	Chef Boyardee ready meals
104	Meal	chang	Meal	Fridge	P.F. Chang ready meals
105	Meal	chilli	Meal	Tin	Tinned Chilli / meat 'n' beans
106	Meal	healthy	Meal	Fridge	Healthy/Nutritious ready meals (e.g. weight-watchers, fish)
107	Meal	lunch	Meal	Fridge	Lunchables (ready packed lunches)
108	Meal	mc	Meal	Shelf	Marie Callender ready meals
109	Meal	meat	Meal	Fridge	Meat based ready meals
110	Meal	other	Meal	Fridge	miscellaneous ready meals
111	Meal	pasta	Meal	Shelf	Pasta ready meals, mac n' cheese etc
112	Meal	pie	Meal	Fridge/Shelf	Savoury pies / pastries (often shelf stable)
113	Meal	pizza	Meal	Fridge	pizzas
114	Meal	sand	Meal	Fridge	sandwiches
115	Meal	side	Snack	Fridge	ready meal sides
116	Meal	soup	Meal	Tin	tinned soups
117	Meal	veggie	Meal	Fridge	vegetarian/vegan meals
118	Meat		Ingredient	Fridge	unspecified meat
119	Meat	bacon	Ingredient	Fridge	Bacon
120	Meat	beef	Ingredient	Tin	Mostly tinned savoury mince
121	Meat	burger	Ingredient	Fridge	various burger patties
122	Meat	chicken	Ingredient	Fridge	Chicken
123	Meat	fish	Ingredient	Fridge	Fish
124	Meat	lunch	Ingredient	Fridge	Deli/luncheon meat
125	Meat	other	Ingredient	Fridge	miscellaneous meats (e.g. pork)
126	Meat	sausage	Ingredient	Fridge	Mostly hot dog sausages
127	Non Food		Non-Food	Non-Food	unspecified non-food
128	Non Food	battery	Non-food	Non-food	batteries
129	Non Food	bleach	Non-food	Non-food	Bleach/solvent cleaning products
130	Non Food	box	Non-food	Non-food	banana boxes/crates
131	Non Food	other	Non-food	Non-food	e.g. clothes, bags, window cleaner, wipes
132	Non Food	salt	Non-food	Non-food	non-food salt
133	Non Food	towel	Non-food	Non-food	paper towels
134	Pasta		Meal	Shelf	Dried pasta
135	Pasta	ben	Ingredient	Shelf	Uncle Ben's rice
136	Pasta	other	Ingredient	Shelf	Miscellaneous pasta product (lasagna sheets etc)
137	Pasta	rice	Ingredient	Shelf	dried rice
138	Snack		Snack	Shelf	unspecified snack
139	Snack	bar	Snack	Shelf	snack/granola bars
140	Snack	bp	Snack	Shelf	baked snacks, e.g. butterfinger
141	Snack	candy	Snack	Shelf	candy/chocolate
142	Snack	chips	Snack	Shelf	crisps
143	Snack	cookies	Snack	Shelf	biscuits
144	Snack	crackers	Snack	Shelf	crackers
145	Snack	fruit	Snack	Shelf	rollups/cups
146	Snack	jelly	Snack	Shelf	jello (pre/unmixed)
147	Snack	kellogg	Snack	Shelf	unspecified Kellogg brand snacks
148	Snack	nuts	Snack	Shelf	nuts/trailmix
149	Snack	other	Snack	Shelf	miscellaneous snacks
150	Snack	pbar	Snack	Shelf	protein bars
151	Snack	pc	Snack	Shelf	pop-corn (mostly popped)
152	Snack	pretzel	Snack	Shelf	pretzels
153	Snack	pt	Snack	Shelf	pop-tarts
154	Snack	pud	Snack	Tin	tinned pudding
155	Snack	seed	Snack	Shelf	sunflower seeds
156	Snack	sj	Snack	Shelf	slim Jims, jerky, biltong
157	Vegetables		Ingredient	Tinned	unspecified non-fresh
158	Vegetables	beans	Ingredient	Tinned	baked beans
159	Vegetables	fruit	Ingredient	Tinned	canned fruit (escaloped apples etc)
160	Vegetables	fry	Ingredient	Fridge	chips/potato wedges/ fries
161	Vegetables	gbean	Ingredient	Tinned	beans (non-baked, mostly green)
162	Vegetables	other	Ingredient	Tinned	miscellaneous veg
163	Vegetables	potato	Ingredient	Fridge	ready to cook potatoes
164	Vegetables	tomato	Ingredient	Tinned	tinned tomatoes

A.1.3 Joint bidding

I did not receive information of joint bidding. However, in some circumstances joint bidding can be inferred. For example, when only one load is auctioned but multiple foodbanks are listed as winning. Likewise, I observe the amount paid by winners (separate to their bid): If two food banks jointly bid 50 shares each I observe that the "bid paid" was 100. I use these cases to identify common bidding coalitions. I then assume that whenever one of these coalitions appears to place a bid, that they are placing a joint bid. By this method I identify around 30 coalitions, and infer that 4.5% of bids are joint bids. This is slightly lower than the true value of 5% reported in Prendergast, 2017. This is likely because I do not detect coalitions that never won together in the data.

I also risk classifying non-joint bids as joint bids when a coalition chooses not to bid jointly on occasions. This is only a problem if they did not win, and unlikely to lead to much inaccuracy if they do not win. A further problem is that I occasionally see multiple lots being auctioned, with more winners than lots (without a known coalition among these food banks). I am unable to infer which subset of bidders forms a coalition, and so am are forced to assume, incorrectly, that none of the bids are joint. This only happens a small fraction of the time, around 0.01%, so is unlikely to lead to much inaccuracy.

Joint bidding in the model

I do not consider the strategic considerations behind joint bidding. I do however consider how this impacts the inverse bid function and winnings. If a bid was joint between n people, I assume the pounds won are divided equally among the n bidders. As are the distance costs and the 'lot specific value'.[2] I also recognise how the food banks' beliefs about the probability they win given their joint bid is higher than either individual bid. Therefore I am able to recognise how the total (expected) surplus of the joint bid may exceed either individual surplus from placing a single bid equal to the joint bid - if storage costs are convex, sharing these load reduces the

[2]In principle I could split the lot according to the fraction of final expenditure, however joint bids in which one bidder bids an extremely small or zero amount, are not uncommon.

total cost incurred. When simulating the Choice System I am unable to simulate the joint bidding procedure.

A.1.4 Maroon Pounds

I do not observe which loads were sold by food banks ('Maroon Pounds'). However, after I had located food banks (discussed in Appendix A.2) considered whether any of the auction origin zipcodes matched the zipcodes of the food banks I had identified. Matched observations all had auction identifier codes that began with "ML" rather than "L" (followed by a string of numbers). Therefore, I focused on these auctions as Maroon Pounds, which make up 4.5% of unique auctions.

Maroon Pounds do not enter the current version of my model. Endogenising the decision to sell food adds too much complexity. However, the food banks responsible for consuming the most through the Choice System almost never sell food. As my results are predominantly driven by these food banks, ignoring Maroon Pounds is unlikely to lead to much inaccuracy. However, for the sake of posterity I will continue to describe how I match food banks to Maroon Loads.

One difficulty with matching food banks to maroon loads is that food banks move over time, often merging with other food bank organisations, so that the zipcode of a lot auctioned in 2014 may not match the zipcode I found for that food bank in 2019. Broadly speaking, I located food banks by finding the name of the city in which they are located, as well as their state. It is rare to have multiple food banks in the same. I therefore decided to match food banks to maroon loads under 3 conditions: First, if the zipcodes matched. Failing that, if they are located in the same city. Failing that, if they are within 20 miles of one another. I assume that the remaining Maroon Pounds are sold by the small food banks and food rescue organisations whose locations I cannot identify, or who are never observed bidding in my data.

A.2 Auxiliary data

I use five additional datasets in my analysis. Two data sets received from Prendergast (one of the original designers of the Choice System) containing losing bidders by auction for 2014, and also poverty figures by county. Third, food bank zipcode data from Feeding America's Food Bank Locator online tool.[3]. Fourth, Food bank catchment areas, defined at the County level, from Feeding America's 'Hunger in America' on-line resource.[4] Finally, Populations figures by county were then taken from the 2015 US census Small Area Income and Poverty Estimates (SAIPE).

These datasets were used to locate food banks and evaluate their Goal Factors.

A.2.1 Locating food banks

Prendergast kindly sent me a dataset containing data on losing bidders by auction for 2014. Importantly, this data contained the nearby towns of the bidding food banks. That is, food banks were identified by the town they were located in. I was able to cross-reference this data with the Choice System data for 2014, merging by date and the origin of each lot.

For each anonymised ID I found the town that appeared in the largest proportion matched auctions. For each town I found the anonymised ID that appeared in the largest proportion of matched auctions (these two proportions need not be equal). If the two sets of pairings were identical, I listed the ID/town combination as matched, removed it from the pool of remaining IDs and towns, and continued the process until I was unable to remove any more matched pairs. This process allowed me to infer the nearby towns of all food banks who placed a bid in 2014. This allowed me to infer approximate locations for 85% of food banks, who together consumed just over 98% of all food on the Choice System. It was clear that my food bank ID numbers had been listed in alphabetical order from 1 to 165 before anonymisation, validating my location matches.[5]

Given knowledge of nearby towns, I then used Feeding America's Food Bank Locator online tool to find zip codes for these food banks. I was unable to find three food bank's locations in this way, as they listed town names which were nowhere near any of Feeding America's food banks. I kept their locations as unknown.

[3] Accessible at https://www.feedingamerica.org/find-your-local-foodbank

[4] Available here: https://map.feedingamerica.org

[5] Based on this alphetical order, and knowledge of all Feeding America's associated food banks, I was able to match an additional 3 food banks by visual inspection

One of the most frequent bidders in the Choice System has a commonly occurring town name, with food banks listed in two of these towns. For these two candidate food banks I examined their annual financial statements from 2014 to find how much non-monetary donations they had received from Feeding America. One received an extremely large amount, while the other received a reasonably small amount. Because the food bank in question consumed an extremely large amount of food on the Choice System I reasoned it was most likely the food bank that received the larger non-monetary donations from Feeding America.

A.2.2 Distance

To find the distance between every lot $\times$ food bank combination I converted zip-codes into longitude/latitudes, then used the "distGeo" function from the R package "geosphere". This package finds the distance of the geodesic between any two points on the globe. In principle I could have found the shortest road distance using arcGIS software, as this would more accurately represent the transportation costs. However, this software is generally extremely computationally intensive. Given the large number of food bank $\times$ lot combinations ($\approx$ 3.6 million) this option was not feasible.

A.2.3 Calculating Goal Factors

I set the Goal Factors of food banks with unknown locations to the smallest known Goal Factor.[6] In this paper I only use the new Goal Factors.

For a small number of food banks their expenditure did not match up with their Goal Factors. They spent *significantly* more shares than the amount they received (implied by their Goal Factor). This is the case even when I take into account that food banks stop receiving new shares once they hit a 200,000 limit.

I calibrate Goal Factors and initial budgets to take this into account. I find the smallest absolute deviation from the Goal Factors implied by the formulae such that: 1. No food bank is ever in debt for longer than 30 days. 2. Food banks with above average Goal Factor are never in debt. 3. Food banks' budgets cannot exceed the

[6]Given that these food banks did not bid regularly enough for me to identify their location, presuming that these food banks also would not accept any loads they are offered in counter-factual simulations ensures that my estimates remain conservative. For some catchment areas Feeding America appears to have used slightly different cut-offs than the 100% and 185% lines.

FIGURE A.1: Distribution of Goal Factors and intial budgets

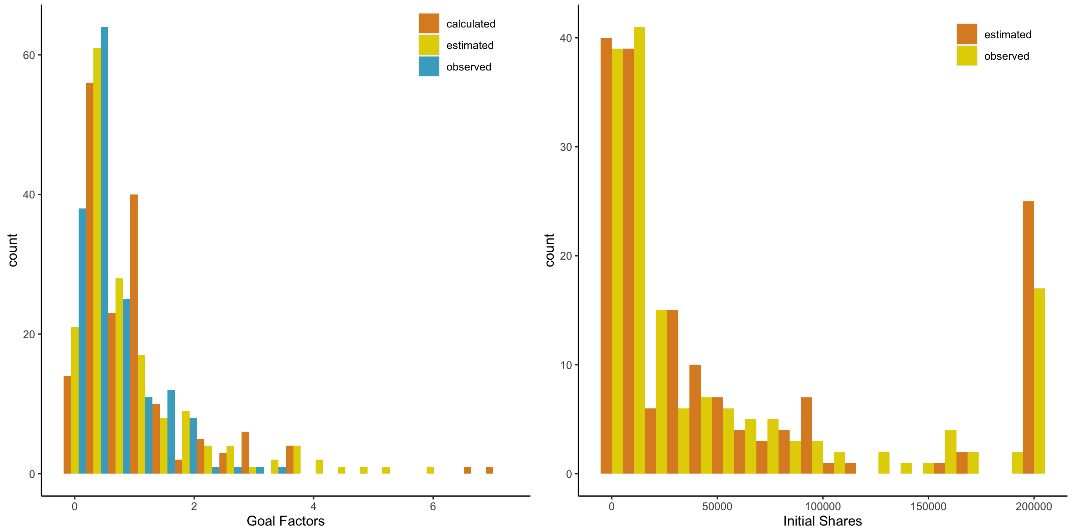

Note: These plots show histograms of Goal Factors and initial budgets across food banks. 'Observed' are the figures I received from Prendergast. 'Calculated' are the figures I calculated using poverty data and the formulae presented above. 'Estimated' are the figures I calibrated using the approach discussed above, that were then used in my counterfactual simulations.

200,000 share limit. 4. No food bank's budgets have a trend (positive or negative) of more than 100,000 shares over the period. 5. No food bank's budgets have a statistically significant (at the 5% level) trend. That is, we expect that their budgets should neither trend up nor downwards over time.

I perform this calibration as follows: Given proposed Goal Factors I find the initial share allocation that satisfies criteria 1-4. This is done by iteratively changing the initial allocation, simulating incomes (given observed expenditures) to find budgets, until the necessary initial allocation is converges. In an outer loop I find the Goal Factors that satisfy criteria 1,2, and 5. At each step I update Goal Factors by taking the average of the prior estimated Goal Factor and the implied new Goal Factor that satisfies the criteria. The process converges in around 100 iterations. For 95% of food banks Goal Factors change very little in relative terms. The largest change is seen by one food bank that consumes a large quantity of food on the Choice System, but has a low initial Goal Factor. Inspection reveals that this food bank exists in a so-called 'food desert', meaning they likely have very little access to local donors, so must rely on Feeding America for the majority of their food.

To validate this approach I compare the distributions of calibrated initial allocations and Goal Factors to those used in Prendergast, 2017, received from Prendergast, but that I was unable to link to my data. Importantly, these figures are around 5 years out of date relative to my data. The two sets of distributions are shown in A.1. The distribution of intial budgets are relatively similar, as are the distributions of Goal Factors, with the exception that my estimated Goal Factors have a larger right tail. Importantly, however, the distribution of my estimated Goal Factors fits the observed distribution better than my initially calculated Goal Factors.

The Goal Factor was designed to ensure that a food bank with a 1% higher goal factor received 1% more food. Prendergast, 2017 found that a 1% increase in Goal Factor was associated with a 0.45% increase in food won from the Choice System. I found that a 1% increase in estimated Goal Factor was associated with a 0.81% increase in consumption. Given that Prendergast's estimation was done on data with very different characteristics to mine, the inaccuracy of these estimated figures is unclear. The difference may be driven by the 15% of unknown Goal Factors in my data. If these food banks had relatively high goal factors this would drive the observed discrepancy, since we know these food banks choose not to consume much. Either way, the relationship between Goal Factor and consumption is not especially strong; the R^2 from a log-log regression is only 0.35. This weak correlation demonstrates the importance of food wealth in determining consumption behaviour. High Goal Factor food banks, who also happen to have many local donors, may not want to consume much food through the Choice System, weakening the correlation.

Appendix B

Proof of Proposition 1

In this Appendix I prove Proposition 1, which states that under the assumptions of the game, and under Conjecture 1, a Symmetric Markov Perfect Equilibrium exists.

The proof of Proposition 1 proceeds by first demonstrating that, conditional on Conjecture 1, a Pure Strategy Bayesian Nash Equilibrium exists in the stage game. I then show that the equilibrium pay-off in the stage game is consistent with the continuation value. I do so by employing Kakutani's fixed point theorem, which requires showing the existence, convex-valuedness, and upper hemicontinuity of the continuation value.

Proof: **Static Equilibrium of the entry game**: player i chooses their entry decision $\mathbf{d}$ to maximise their expected pay-off. If the player knew the entry decisions of other players in advance they would essentially choose an entry structure, then the expected pay-offs would just be the equilibrium expected payoffs associated with each entry structure. Instead, their expected pay-off is formed by taking expectation over the entry decisions of every other player, given the strategies of other players. The entry game is therefore a simple game of incomplete information.

A symmetric equilibrium in distributional strategies exists, thanks to Milgrom and Weber, 1985.[1] However this equilibrium may not be unique, a problem for continuity of the value function. Continuity can be restored by following Fudenberg and Maskin, 1991, entry strategies can be augmented to be a function of the realisation of a public random variable. The public random variable enables players to coordinate over which equilibrium will be played. Conditional on this public random variable the set of equilibrium pay-offs is convex (Aumann, 1974).

[1] Likewise, existence of a Pure Strategy equilibrium follows from their purification result, under the assumption that types are atomless. However, we will not require bidders to play pure strategies, and instead ensure our identification framework is consistent with distribution and/or pure entry strategies.

As in JP, equilibrium existence of the dynamic game then requires that the equilibrium pay-off in the stage game is consistent with the continuation value.[2] That is, can we write the ex-ante value function $\mathbf{V}_t^E$, stacked over states, as a function of $\mathbf{V}_{t+1}^E$, so that $\mathbf{V}_t^E = \Omega(\mathbf{V}_{t+1}^E)$ (existence). In addition, does the correspondence Ω have a fixed point such that $\mathbf{V}^E = \Omega(\mathbf{V}^E)$ (stationarity).

Existence of $\mathbf{V}_t^E = \Omega(\mathbf{V}_{t+1}^E)$: Writing the ex-ante value function in recursive form, substitute equation 3.2.3 for period $t+1$ into equation 3.2.3 into equation 3.2 back into equation 3.2.3 for period t. Existence then follows from the assumption that pay-offs are bounded. This ensures the set $\Omega(\mathbf{V}_{t+1}^E)$ is non-empty.

(non-)Uniqueness of $\Omega(\mathbf{V}_{t+1}^E)$: The possible existence of multiple equilibria in the entry game imply the value function is non-unique. therefore, the ex-ante value function is also non-unique. Fortunately Ω must be convex valued, as the set of equilibrium pay-offs, conditional on the public random variable, is convex.

Upper-hemi continuity of $\Omega(.)$: The continuation value must be continuous in $\mathbf{V}_{t+1}^E$, shown by equation 3.2.3. Next, consider the conditional value function, conditional on entry decision $\bar{\mathbf{d}}$:

$$\tilde{W}_i(\bar{\mathbf{d}}, \boldsymbol{v}_{it}, \mathbf{s}_t; \sigma_{-i}) = \max_{\mathbf{b}} \left\{ \Gamma_i(\mathbf{b}, \bar{\mathbf{d}}; \sigma_{-i})^T(\boldsymbol{v}_{it} - \mathbf{b}) + P_i(\mathbf{b}, \bar{\mathbf{d}}; \sigma_{-i})^T[J_i(\mathbf{s}_t) + \beta V_i(\mathbf{s}_t; \sigma_{-i})] \right\}$$

Continuity of $\tilde{\mathbf{W}}_t$ in $\mathbf{V}_{t+1}^E$ is guaranteed by conjecture 1, which requires equilibrium expected pay-offs are continuous in $J_i + \beta V_i$. The value function can then be written as $W_i(\boldsymbol{v}_{it}, \mathbf{s}_t; \sigma_{-i}) = \max_{\mathbf{d}} \{\tilde{W}_i(\mathbf{d}, \boldsymbol{v}_{it}, \mathbf{s}_t; \sigma_{-i})\}$. Upper-hemi continuity of $\mathbf{W}_t$ in $\tilde{\mathbf{W}}_t$, and hence in $\mathbf{V}_{t+1}^E$, arises from our public random variable (Fudenberg and Maskin, 2009).[3] Upper-hemi continuity of $\mathbf{V}_t^E$ arises from the ex-ante value function taking an expectation over states.

Existence of a stationary dynamic equilibrium: In order to show existence of a stationary equilibrium we must show that there exists a fixed point of the correspondence $\mathbf{V}^E = \Omega(\mathbf{V}^E)$. As $\Omega()$ is non-empty, convex valued, and upper-hemi continuous, we can apply Kakutani's fixed point theorem. Therefore, a Markov Perfect Equilibrium exists.

□

[2] This dynamic equilibrium will be symmetric since the equilibrium in the stage game is symmetric, so that strategies depend only on states, not player identities or time periods.

[3] Public randomisation ensures that the set of equilibrium pay-offs is convex. Public randomisation means $\mathbf{W}_t$ is the convex hull of possible equilibrium pay-offs from entry, $\tilde{\mathbf{W}}_t$. Therefore, so long as $\tilde{\mathbf{W}}_t$ is compact valued, $\mathbf{W}_t$ is upper hemicontinuous (Charalambos and Aliprantis, 2013). Compact valuedness comes from pay-offs being drawn from a compact set.

Appendix C

Extension of Proposition 1 from Jofre-Bonet and Pesendorfer (2003)

In this Appendix I essentially extend Proposition 1 from JP to the multi-object case. In Appendix C.1 I prove Proposition 3 from the main text. In Appendix C.2 I present (and prove) that the ex-ante value function also has an analytic expression. In Appendix C.3 I prove that even in the case of binding reservation prices the ex-ante value function still has an analytic expression. Finally, in Appendix C.4 I prove that the Inverse Bid System is strictly monotonic for bids strictly above the reservation price. This is an important proof as it allows me to employ the 'Law of the Unconscious Statistician'.

For the remainder of this section I regularly make use of the definition of $\mathbf{k} = C\mathbf{j} + \beta\mathbf{V}$, and equivalently $K(\mathbf{s}) = J(\mathbf{s}_i) + \beta V(\mathbf{s})$.

C.1 Proof of Proposition 3

Proposition 3 : Under assumptions 1 - 4, the expected stage pay-off is given by:

$$\begin{aligned}\tilde{\Pi}(\mathbf{b}^*|v;\mathbf{s}) = {} & \Gamma(\mathbf{b}^*|\mathbf{s})^T\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1}\Gamma(\mathbf{b}^*|\mathbf{s}) \\ & + [P(\mathbf{b}^*|\mathbf{s})^T - \Gamma(\mathbf{b}^*|\mathbf{s})^T\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1}\nabla_{\mathbf{b}}P(\mathbf{b}^*|\mathbf{s})]B_{\mathbf{s}}\mathbf{j} \\ & + [Q(\mathbf{b}^*|\mathbf{s})^T - \Gamma(\mathbf{b}^*|\mathbf{s})^T\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1}\nabla_{\mathbf{b}}Q(\mathbf{b}^*|\mathbf{s})]A_{\mathbf{s}}\beta\mathbf{V} \qquad \text{(C.1)}\end{aligned}$$

Proof: 1. Necessary First Order Conditions are given by:

$$\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})(v-\mathbf{b}^*) = \Gamma(\mathbf{b}^*|\mathbf{s}) - \nabla_{\mathbf{b}}P(\mathbf{b}^*|\mathbf{s})B_{\mathbf{s}}\mathbf{j} - \beta\nabla_{\mathbf{b}}Q(\mathbf{b}^*|\mathbf{s})A_{\mathbf{s}}\mathbf{V}$$

2. Left multiplying by $\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1}$

$$(v - \mathbf{b}^*) = \nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1}[\Gamma(\mathbf{b}^*|\mathbf{s}) - \nabla_{\mathbf{b}}P(\mathbf{b}^*|\mathbf{s})B_{\mathbf{s}}\mathbf{j} - \beta\nabla_{\mathbf{b}}Q(\mathbf{b}^*|\mathbf{s})A_{\mathbf{s}}\mathbf{V}]$$

3. Left multiplying by $\Gamma(\mathbf{b}^*|\mathbf{s})^T$

$$\begin{aligned}&\Gamma(\mathbf{b}^*|\mathbf{s})^T(v - \mathbf{b}^*)\\ &= \Gamma(\mathbf{b}^*|\mathbf{s})^T\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1}[\Gamma(\mathbf{b}^*|\mathbf{s}) - \nabla_{\mathbf{b}}P(\mathbf{b}^*|\mathbf{s})B_{\mathbf{s}}\mathbf{j} - \beta\nabla_{\mathbf{b}}Q(\mathbf{b}^*|\mathbf{s})A_{\mathbf{s}}\mathbf{V}]\end{aligned}$$

4. The Maximised Expected Pay-off is given by:

$$\tilde{\Pi}(\mathbf{b}^*|v;\mathbf{s}) = \Gamma(\mathbf{b}^*|\mathbf{s})^T(v - \mathbf{b}^*) + P(\mathbf{b}^*|\mathbf{s})B_{\mathbf{s}}\mathbf{j} + \beta Q(\mathbf{b}^*|\mathbf{s})A_{\mathbf{s}}\mathbf{V}$$

Substituting in $\Gamma(\mathbf{b}^*|\mathbf{s})^T(v - \mathbf{b}^*)$ gives the result.

□

C.2 Ex-ante Value Function

Building on Proposition 3 the ex-ante Value function can be written as:

$$\begin{aligned}V_i^E(\mathbf{s}_t) &= E_{\mathbf{b}}[\,\Gamma(\mathbf{b}|\mathbf{s}_t)^T\nabla_{\mathbf{b}}\Gamma(\mathbf{b}|\mathbf{s}_t)^{-1}\Gamma(\mathbf{b}|\mathbf{s}_t)]\\ &\quad + E_{\mathbf{b}}[Q(\mathbf{b}|\mathbf{s}_t)^T - \Gamma(\mathbf{b}|\mathbf{s}_t)^T\nabla_{\mathbf{b}}\Gamma(\mathbf{b}|\mathbf{s}_t)^{-1}\nabla_{\mathbf{b}}Q(\mathbf{b}|\mathbf{s}_t)\ |\mathbf{s}_t]K(\mathbf{s}_t) \qquad \text{(C.2)}\end{aligned}$$

Proof that this is the case should be trivial, however it requires that we apply a change of variable, changing from integrating over v to integrating over $\mathbf{b}$.

Proof:

1. To obtain the ex-ante value function from the equation presented in Proposition 3 we then take an expectation over both sides with respect to v for:

$$\begin{aligned}&E_v[\Gamma(\mathbf{b}^*|\mathbf{s}_t)^T(v - \mathbf{b}^*) + P(\mathbf{b}|\mathbf{s}_t)^T K(\mathbf{s}_t)]\\ &\qquad = E_v[\Gamma(\mathbf{b}^*|\mathbf{s}_t)^T\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s}_t)^{-1}\Gamma(\mathbf{b}^*|\mathbf{s}_t)]\\ &\qquad + E_v[Q(\mathbf{b}|\mathbf{s}_t)^T - \Gamma(\mathbf{b}^*|\mathbf{s}_t)^T\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s}_t)^{-1}\nabla_{\mathbf{b}}Q(\mathbf{b}^*|\mathbf{s}_t)]K(\mathbf{s}_t)\end{aligned}$$

2. By applying the Law of the Unconscious Statistician (change of variables for expectations) the right hand side of this equation is equal to

$$E_{\mathbf{b}}[\,\Gamma(\mathbf{b}|\mathbf{s}_t)^T\nabla_{\mathbf{b}}\Gamma(\mathbf{b}|\mathbf{s}_t)^{-1}\Gamma(\mathbf{b}|\mathbf{s}_t) + [Q(\mathbf{b}|\mathbf{s}_t)^T - \nabla_{\mathbf{b}}\Gamma(\mathbf{b}|\mathbf{s}_t)^{-1}\nabla_{\mathbf{b}}Q(\mathbf{b}|\mathbf{s}_t)]K(\mathbf{s}_t)\ |\mathbf{s}_t]$$

□

Importantly, in order to apply the Law of the Unconscious Statistician we require that the mapping $\xi(\mathbf{b}^*|K;\mathbf{s})$ is monotonic (the jacobian has non-zero determinant) in

b. I prove this in Appendix C.4. For those unfamiliar with the Law of the Unconscious Statistician, this intuitive law states that Given random variables U_l and $\mathbf{B}$ such that $U_l = h(\mathbf{B})$, where $||\nabla_{\mathbf{b}} h|| > 0$, the Law of the Unconscious Statistician states that $E_{U_l}[U_l] = \int v_l f_{U_l}(v_l) dv_l = \int_{b_1} ... \int_{b_L} h(\mathbf{b}) f_{\mathbf{B}}(\mathbf{b}) db_L ... db_1 = E_{\mathbf{B}}[h(\mathbf{b})]$

C.3 Ex-ante Value Function (Reservation Prices)

I now consider the case with binding reservation prices. In this case, I must define a partition of the bidding space according to which bids are at the reservation price, and which bids are strictly above the reservation price. This partition essentially partitions which first order conditions hold exactly (Lagrangian multiplier is zero), and which do not. Each component m, denoted $\mathbb{A}_m$, will consist of a set of bids above the reservation price $\mathbb{A}_m^+ = \{l : b_l > R\}$, a set at the reservation price $\mathbb{A}_m^- = \{l : b_l = R\}$, and a non-entered set $\mathbb{A}_m^c = \{l : b_l = \emptyset\}$. Importantly, for a given K, each component defines a corresponding set in lot-specific pay-off space, which I write as $v(\mathbb{A}_m)$. This set is known from the inequalities derived in section E.2.2. We can then write the ex-ante value function as follows:

$$\begin{aligned}
V_i^E(\mathbf{s}_t) = E_m[\sum_{l \in \mathbb{A}_m^-} \Gamma_l(R, d_l|\mathbf{s}_t)(E_{v_l}[v_l|\boldsymbol{v} \in v(\mathbb{A}_m)] - R) \\
+ \sum_{l \in \mathbb{A}_m^+} E_{\mathbf{b},\mathbf{d}}[\frac{\Gamma_l(b_l, d_l|\mathbf{s}_t)^2}{\nabla_{b_l}\Gamma_l(b_l, d_l|\mathbf{s}_t)} - \frac{\Gamma_l(b_l, d_l|\mathbf{s}_t)}{\nabla_{b_l}\Gamma_l(b_l, d_l|\mathbf{s}_t)} \nabla_{b_l} Q(\mathbf{b}, \mathbf{d}|\mathbf{s}_t) K(\mathbf{s}_t) | \mathbf{b} \in \mathbb{A}_m]] \\
+ E_{\mathbf{b},\mathbf{d}}[Q(\mathbf{b}, \mathbf{d}|\mathbf{s}_t)|\mathbf{b} \in \mathbb{A}_m]^T K(\mathbf{s}_t) \\
= \tilde{\Phi}(\mathbf{s}_t) + \tilde{\Omega}(\mathbf{s}_t) K(\mathbf{s}_t) \quad \text{(C.3)}
\end{aligned}$$

This expression is extremely convenient for the econometrician, who is then able to evaluate the ex-ante value function by taking integrals with respect to observed bids. The only difficulty is in calculating the truncated means $E_{v_l}[v_l|\boldsymbol{v} \in v(\mathbb{A}_m)]$.

I now derive equation C.3. For notational simplicity I drop the dependence on **s**.

Proof: 1. We can split an expectation into partitions of conditional expectations, which follows from the law of iterated expectations. This allows us to write:

$$\begin{aligned}&E_{\boldsymbol{v}}[\Gamma(\mathbf{b}^*,\mathbf{d}^*)^T(\boldsymbol{v}-\mathbf{b}^*)+P(\mathbf{b}^*,\mathbf{d}^*)^T K]\\ &=\sum_m P(\boldsymbol{v}\in v(\mathbb{A}_m))E_{\boldsymbol{v}}[\Gamma(\mathbf{b}^*,\mathbf{d}^*)^T(\boldsymbol{v}-\mathbf{b}^*)+P(\mathbf{b}^*,\mathbf{d}^*)^T K|\boldsymbol{v}\in v(\mathbb{A}_m)]\end{aligned}$$

2. As expectations are linear operators write the right hand side of this equation as:

$$\begin{aligned}=\sum_m P(\boldsymbol{v}\in v(\mathbb{A}_m))(\sum_l E_{\boldsymbol{v}}[\Gamma_l(b_l^*,d_l^*)(v_l-b_l^*)|\boldsymbol{v}\in v(\mathbb{A}_m)])\\ +E_{\boldsymbol{v}}[Q(\mathbf{b}^*,\mathbf{d}^*)^T K|\boldsymbol{v}\in v(\mathbb{A}_m)]\end{aligned}$$

3. Consider $E_{\boldsymbol{v}}[P(\mathbf{b}^*,\mathbf{d}^*)^T K|\boldsymbol{v}\in v(\mathbb{A}_m)]$. We will essentially apply the Law of the Unconscious Statistician to write this as an expectation over **b** and **d**. Apply the Law of Iterated Expectations, so that we take expectations first over v_l for $l\in\mathbb{A}_m^+$ first, and then over v_l for $l\notin\mathbb{A}_m^+$. We have:

$$\begin{aligned}&E_{\boldsymbol{v}}[P(\mathbf{b}^*,\mathbf{d}^*)^T K|\boldsymbol{v}\in v(\mathbb{A}_m)]\\ &=E_{v_l\, l\in\mathbb{A}_m^+}[E_{v_l\, l\notin\mathbb{A}_m^+}[Q(\mathbf{b}^*,\mathbf{d}^*)^T K|\boldsymbol{v}\in v(\mathbb{A}_m)\ \&\ v_l\ l\in\mathbb{A}_m^+]|\boldsymbol{v}\in v(\mathbb{A}_m)]\end{aligned}$$

4. By definition of $\mathbb{A}_m^+$ we have $E_{v_l\, l\notin\mathbb{A}_m^+}[Q(\mathbf{b}^*,\mathbf{d}^*)^T K|\boldsymbol{v}\in v(\mathbb{A}_m)\ \&\ v_l\ l\in\mathbb{A}_m^+]=Q(\mathbf{b}^*,\mathbf{d}^*)^T K$. This is because bids on lot $l\notin\mathbb{A}_m^+$ are constant for $\boldsymbol{v}\in v(\mathbb{A}_m)$.
5. Next, we are able to apply the law of the unconscious statistician, since the functions $b_l^*(\boldsymbol{v},K)$ for $l\in\mathbb{A}_m^+$ are monotonic in $\boldsymbol{v}$, as shown in appendix C.4.

$$E_{v_l\, l\in\mathbb{A}_m^+}[Q(\mathbf{b}^*,\mathbf{d}^*)^T K|\boldsymbol{v}\in v(\mathbb{A}_m)]=E_{b_l\, l\in\mathbb{A}_m^+}[Q(\mathbf{b}^*,\mathbf{d}^*)^T K|\boldsymbol{b}\in\mathbb{A}_m]$$

The idea here is that we couldn't apply this change of variables in general, because some of the random variables (bids) we tried to take expectations over are in regions of $\boldsymbol{v}$ space where the bids are non-monotonic in $\boldsymbol{v}$, and actually do not vary with $\boldsymbol{v}$ at all. But, precisely because they do not vary in these regions, we can just integrate out these dimensions, before applying the change of variables. Last of all, recognise that because b_l for $l\notin\mathbb{A}_m^+$ is constant, we have:

$$E_{b_l\, l\in\mathbb{A}_m^+}[Q(\mathbf{b}^*,\mathbf{d}^*)^T K|\boldsymbol{b}\in\mathbb{A}_m]=E_{\mathbf{b}}[Q(\mathbf{b}^*,\mathbf{d}^*)^T K|\boldsymbol{b}\in\mathbb{A}_m]$$

6. Next, focus on $E_{\boldsymbol{v}}[\Gamma_l(b_l^*,d_l^*)(v_l-b_l^*)|\boldsymbol{v}\in v(\mathbb{A}_m)]$. For $l\in\mathbb{A}_m^c$, so that $d_l^*=0$, we have $E_{\boldsymbol{v}}[\Gamma_l(b_l^*,d_l^*)(v_l-b_l^*)|\boldsymbol{v}\in v(\mathbb{A}_m)]=0$, since for this l we have $\Gamma_l(b_l^*,d_l^*)=0$ by definition.
7. Meanwhile, for $l\in\mathbb{A}_m^-$ we have

$$E_{\boldsymbol{v}}[\Gamma_l(b_l^*,d_l^*)(v_l-b_l^*)|\boldsymbol{v}\in v(\mathbb{A}_m)]=\Gamma_l(R,1)(E_{\boldsymbol{v}}[v_l|\boldsymbol{v}\in v(\mathbb{A}_m)]-R)$$

by definition of partition m.

8. Finally for $l \in \mathbb{A}_m^+$ we have

$$E_{\boldsymbol{v}}[\Gamma_l(b_l^*, d_l^*)(v_l - b_l^*)|\boldsymbol{v} \in v(\mathbb{A}_m)] = E_{\boldsymbol{v}}[\frac{\Gamma_l(b_l^*,1)^2}{\nabla_{b_l}\Gamma_l(b_l^*,1)} - \frac{\Gamma_l(b_l^*,1)}{\nabla_{b_l}\Gamma_l(b_l^*,1)}\nabla_{b_l}Q(\mathbf{b}^*,\mathbf{d}^*)|\boldsymbol{v} \in v(\mathbb{A}_m)]$$

This arises because the lth first order condition of the lagrangian given in equation E.4 is given by:

$$\nabla_{b_l}\Gamma_l(b_l^*, d_l^*)(v_l - b_l) - \Gamma_l(b_l^*, d_l^*) + \nabla_{b_l}Q(\mathbf{b}^*,\mathbf{d}^*)K + \lambda_l^* = 0$$

l such that $b_l^* > R$ we have $\lambda_l^* = 0$ (this solution is unconstrained). This equation then rearranges for

$$\Gamma_l(b_l^*, d_l^*)(v_l - b_l^*) = \frac{\Gamma_l(b_l^*,1)^2}{\nabla_{b_l}\Gamma_l(b_l^*,1)} - \frac{\Gamma_l(b_l^*,1)}{\nabla_{b_l}\Gamma_l(b_l^*,1)}\nabla_{b_l}Q(\mathbf{b}^*,\mathbf{d}^*)$$

9. Still considering $l \in \mathbb{A}_m^+$, the idea is that we now want to apply the Law of the Unconscious Statistician for

$$E_{\boldsymbol{v}}[\frac{\Gamma_l(b_l^*,1)^2}{\nabla_{b_l}\Gamma_l(b_l^*,1)} - \frac{\Gamma_l(b_l^*,1)}{\nabla_{b_l}\Gamma_l(b_l^*,1)}\nabla_{b_l}Q(\mathbf{b}^*,\mathbf{d}^*)|\boldsymbol{v} \in v(\mathbb{A}_m)] = E_{\mathbf{b}}[\frac{\Gamma_l(b_l^*,1)^2}{\nabla_{b_l}\Gamma_l(b_l^*,1)} - \frac{\Gamma_l(b_l^*,1)}{\nabla_{b_l}\Gamma_l(b_l^*,1)}\nabla_{b_l}Q(\mathbf{b}^*,\mathbf{d}^*)|\mathbf{b} \in \mathbb{A}_m]$$

The problem is that $\nabla_{b_l}Q(\mathbf{b}^*,\mathbf{d}^*)$ in general depends on $b_{l'}$ for $l' \notin \mathbb{A}_m^+$. However, just as we did in steps 3 to 5, we can first apply the law of iterated expectations and integrate out these other l' dimensions, before employing the change of variables.

10. Putting together steps 1 through 5, 6, 7, and 8 through 9, we can write:

$$\begin{aligned} &E_{\boldsymbol{v}}[\Gamma(\mathbf{b}^*,\mathbf{d}^*)^T(\boldsymbol{v} - \mathbf{b}^*) + Q(\mathbf{b}^*,\mathbf{d}^*)^T K] \\ &= \sum_m P(\boldsymbol{v} \in v(\mathbb{A}_m)) \sum_{l \in \mathbb{A}_m^-} \Gamma_l(R, d_l|\mathbf{s}_t)(E_{v_l}[v_l|\boldsymbol{v} \in v(\mathbb{A}_m)] - R) \\ &+ \sum_{l \in \mathbb{A}_m^+} E_{\mathbf{b},\mathbf{d}}[\frac{\Gamma_l(b_l, d_l|\mathbf{s}_t)^2}{\nabla_{b_l}\Gamma_l(b_l, d_l|\mathbf{s}_t)} - \frac{\Gamma_l(b_l, d_l|\mathbf{s}_t)}{\nabla_{b_l}\Gamma_l(b_l, d_l|\mathbf{s}_t)}\nabla_{b_l}Q(\mathbf{b},\mathbf{d}|\mathbf{s}_t)K(\mathbf{s}_t)|\mathbf{b} \in \mathbb{A}_m] \\ &+ E_{\mathbf{b},\mathbf{d}}[Q(\mathbf{b},\mathbf{d}|\mathbf{s}_t)|\mathbf{b} \in \mathbb{A}_m]^T K(\mathbf{s}) \end{aligned}$$

$\square$

C.4 Monotonicity of the Inverse Bid System

This argument proceeds in two parts, but is fundamentally just an application of the Envelope Theorem. First, I consider the second order conditions of the bidding problem, substituting in the first order conditions to find an particularly useful expression. I then consider the Jacobian of the inverse bidding system, substituting in the aforementioned expression to show that this jacobian matrix must have non-zero determinant.

Proof:

1. Differentiating equation E.4 we obtain a hessian matrix of second derivatives with entry ij given by:

$$H_{ln} = \begin{cases} \nabla^2_{b_l}\Gamma_l(b^*_l, d^*_l)(v_l - b^*_l) & \text{if} \quad l = n \\ \qquad -2\nabla_{b_l}\Gamma_l(b^*_l, d^*_l) + \nabla^2_{b_l}Q(\mathbf{b}^*, \mathbf{d}^*)K & \\ \nabla_{b_l}\nabla b_n Q(\mathbf{b}^*, \mathbf{d}^*)K & \text{if} \quad l \neq n \end{cases}$$

2. From the first order conditions we then substitute in $\nabla^2_{b_l}\Gamma_l(b^*_l, d^*_l)(v_l - b^*_l) = \frac{\nabla^2_{b_l}\Gamma_l(b^*_l,d^*_l)}{\nabla_{b_l}\Gamma_l(b^*_l,d^*_l)}(\Gamma_l(b^*_l, d^*_l) - \nabla_{b_l}Q(\mathbf{b}^*, \mathbf{d}^*)K)$:

$$H_{ln} = \begin{cases} \frac{\nabla^2_{b_l}\Gamma_l(b^*_l,d^*_l)}{\nabla_{b_l}\Gamma_l(b^*_l,d^*_l)}(\Gamma_l(b^*_l, d^*_l) - \nabla_{b_l}Q(\mathbf{b}^*, \mathbf{d}^*)K) & \text{if} \quad l = n \\ \qquad -2\nabla_{b_l}\Gamma_l(b^*_l, d^*_l) + \nabla^2_{b_l}Q(\mathbf{b}^*, \mathbf{d}^*)K & \\ \nabla_{b_l}\nabla b_n Q(\mathbf{b}^*, \mathbf{d}^*)K & \text{if} \quad l \neq n \end{cases}$$

3. Recognise that if reservation prices were not binding, this Hessian matrix would have to be negative definite as a condition for utility maximising. Even as it is, for the subset of bids that are strictly above the reservation price, and so whose Lagrangian multipliers are zero, the corresponding sub-Hessian must also be negative definite.
4. Next, consider the inverse bidding system for $b_l > R$:

$$\xi_l(\mathbf{b}, \mathbf{d}; K) = b_l + \frac{\Gamma_l(b_l, d_l)}{\nabla_{b_l}\Gamma_l(b_l, d_l)} - \frac{1}{\nabla_{b_l}\Gamma_l(b_l, d_l)}\nabla_{b_l}Q(\mathbf{b}^*, \mathbf{d}^*)K$$

Taking the Jacobian of this object yields:

$$J_{ln} = \begin{cases} 2 - \frac{\Gamma_l(b_l,d_l)\nabla^2_{b_l}\Gamma_l(b_l,d_l)}{\nabla_{b_l}\Gamma_l(b_l,d_l)^2} + \frac{\nabla^2_{b_l}\Gamma_l(b_l,d_l)}{\nabla_{b_l}\Gamma_l(b_l,d_l)^2}\nabla_{b_l}Q(\mathbf{b}^*, \mathbf{d}^*)K & \text{if} \quad l = n \\ \qquad -\frac{1}{\nabla_{b_l}\Gamma_l(b_l,d_l)}\nabla^2_{b_l}Q(\mathbf{b}^*, \mathbf{d}^*)K & \\ -\frac{1}{\nabla_{b_l}\Gamma_l(b_l,d_l)}\nabla_{b_l}\nabla_{b_n}Q(\mathbf{b}^*, \mathbf{d}^*)K & \text{if} \quad l \neq n \end{cases}$$

5. Substituting in our updated Hessian matrix considered previously, this yields the simple identity $J_{ln} = -\frac{1}{\nabla_{b_l}\Gamma_l(b_l,d_l)}H_{ln}$, or equivalently that $J = -\nabla_{\mathbf{b}}\Gamma_l(\mathbf{b}^*, \mathbf{d}^*)^{-1}H$. This is as expected given the envelope theorem, and definitely didn't come as an enormous surprise.
6. If every bid is strictly above the reservation price, or if reservation prices do not bind, then $det(J) = det(\nabla_{\mathbf{b}}\Gamma_l(\mathbf{b}^*, \mathbf{d}^*)^{-1})det(-H)$ which must be positive as H must be negative definite. Meanwhile, if only a subset of bids are strictly above the reservation price, we instead focus on just the sub-Hessian and Jacobians.

□

Appendix D

Proof of Proposition 4

In this Appendix I prove Proposition 4. The proposition is given by:

Proposition 4. *Under assumption 1 - 5* $\Psi(I_S - \beta T\Omega)^{-1}C$ has rank $S_i - 1$

The proof is split into three parts. First, I establish the rank of Ψ, and then find its null space. I then demonstrate that the intersection of this null space and the image of $(I_S - \beta T\Omega)^{-1}C$ only contains a single element.

D.1 Rank of Ψ

This proposition employs assumption 5, as well as two additional lemmas stated below. But first, we need some additional tools to help us establish this result.

D.1.1 Additional Definitions

Define the partial ordering $\succeq^*$ such that if $\mathbf{s}_i \succeq \mathbf{s}'_i$ then $\mathbf{s} \succeq^* \mathbf{s}'$. This simply extends the partial ordering of the individual state to the overall state.

Next, define a 'component', written $\mathbb{S}^c$, of the set $\mathbb{S}$ as follows:[1]

Definition D.1.1 (Component)**.** The component $\mathbb{S}^c \subset \mathbb{S}$ is such that two states $\mathbf{s}, \mathbf{s}' \in \mathbb{S}^c$ if and only if there exists a state $\bar{\mathbf{s}}$ such that one of the following holds:

$$\mathbf{s} \succeq^* \bar{\mathbf{s}} \ \& \ \mathbf{s}' \succeq^* \bar{\mathbf{s}} \quad \text{or} \quad \mathbf{s} \succeq^* \bar{\mathbf{s}} \succeq^* \mathbf{s}' \quad \text{or} \quad \mathbf{s}' \succeq^* \bar{\mathbf{s}} \succeq^* \mathbf{s} \quad \text{or} \quad \bar{\mathbf{s}} \succeq^* \mathbf{s} \ \& \ \bar{\mathbf{s}} \succeq^* \mathbf{s}' \tag{D.1}$$

[1] Recognise that this definition corresponds to the graph theoretic definition of a component if we instead state that there exists a directed path between 'nodes' $\mathbf{s}, \mathbf{s}'$ if and only if $\mathbf{s}' \succeq^* \mathbf{s}$

A component is essentially a subset of $\mathbb{S}$ that are 'connected' by this partial ordering. By definition $\mathbf{s}_0$ does not vary within a component, and in general there will be one component corresponding to each element of the set $\mathbf{s}_0 \in \mathbb{S}_0$. Components are mutually exclusive and exhaustive subsets of $\mathbb{S}$. Suppose there are S^c components.

Finally, denote $\tilde{\min}(\mathbb{S})$ as the subset of $\mathbb{S}$, such that $\forall \mathbf{s} \in \tilde{\min}(\mathbb{S})\ :\ \nexists\, \mathbf{s}' \in \mathbb{S}\ :\ \mathbf{s} \in \mathbb{S}^a(\mathbf{s}')$. This definition is primarily for notational convenience, and does not necessarily coincide with the set of minimal elements of $\mathbb{S}$. Instead, this is the (potentially empty) set of states that never occur as possible ex-post states. Intuitively, pay-offs from ending in these states will not be identified.

D.1.2 Useful Lemmas

Lemma D.1.1. *From any two distinct, non-maximal, states,* $\mathbf{s}$ *and* $\mathbf{s}'$*, if* $\mathbf{s}' \not\succeq^* \mathbf{s}$ *then there exists a state* $\mathbf{s}^a$ *such that* $\mathbf{s}^a \in \mathbb{S}^a(\mathbf{s})\ \&\ \mathbf{s}^a \notin \mathbb{S}^a(\mathbf{s}')$

This Lemma states that if one non-maximal state is not 'higher' in the partial ordering than another, then their set of ex-post states cannot perfectly overlap. Proof of the lemma is very simple, and focuses on whether the unique element of $\mathbb{S}^a(\mathbf{s})$ that consists of bidder i winning every lot, denoted $\mathbf{s}^{all_i}$, can be an element of $\mathbb{S}^a(\mathbf{s}')$. The lemma makes use of the following property of the partial ordering $\succeq$ from assumption 5: For any two incomparable states $\mathbf{s}_i, \mathbf{s}_i'$ and any $\mathbf{s}_0$ there must exist some $\mathbf{s}^a \in \mathbb{S}_i^a(\mathbf{s}_i, \mathbf{s}_0)$ such that $\mathbf{s}^a \notin \mathbb{S}_i^a(\mathbf{s}_i', \mathbf{s}_0)$. This result arises from the fact that the same set of lots are available in each state, ensuring that the sets $\mathbb{S}_i^a(\mathbf{s}_i, \mathbf{s}_0)$ and $\mathbb{S}_i^a(\mathbf{s}_i', \mathbf{s}_0)$ cannot totally overlap. This enables proof of lemma D.1.1:

Proof:

1. Suppose that $\mathbf{s}' \not\succeq^* \mathbf{s}$. This gives us 2 initial options: Either $\mathbf{s} \succeq \mathbf{s}'$, or the two states are incomparable.
2. If $\mathbf{s} \succeq \mathbf{s}'$ then they must lie in the same component, implying that $\mathbf{s}_0 = \mathbf{s}_0'$. In turn, this implies that exactly the same lots must be available in each state. The result follows trivially - it cannot be the case that $\mathbf{s}^{all_i} \in \mathbb{S}^a(\mathbf{s}')$ for non-maximal states.
3. If they are incomparable then we have another two options: Either $\mathbf{s}$ and $\mathbf{s}'$ belong to different components, or to the same component.
4. If they belong to different components then by definition $\mathbb{S}^a(\mathbf{s})$ and $\mathbb{S}^a(\mathbf{s}')$ must be mutually exclusive.
5. If they belong to the same component then by definition of the partial ordering $\succeq^*$ it must be that $\mathbf{s}_i$ and $\mathbf{s}_i'$ are incomparable under the ordering

$\succeq$. There must exist some $\mathbf{s}^a \in \mathbb{S}_i^a(\mathbf{s}_i, \mathbf{s}_0)$ such that $\mathbf{s}^a \notin \mathbb{S}_i^a(\mathbf{s}_i', \mathbf{s}_0)$. There must exist some $\mathbf{s}^a \in \mathbb{S}^a(\mathbf{s})$ such that $\mathbf{s}^a \notin \mathbb{S}^a(\mathbf{s}')$.

□

Lemma D.1.2. *$\Psi(\mathbf{s})A_\mathbf{s}$ has rank at least 2 if, for all non-maximal $\mathbf{s}, v$, $\Gamma_{il}(\mathbf{b}(v,\mathbf{s})|\mathbf{s}) \in (0,1)$ for each l*

The proof proceeds by first showing that $rank(\Psi(\mathbf{s}))$ is weakly greater than two, then using the full rank property of the transformation matrix $A_\mathbf{s}$.

Proof:
1. Consider the $L \times L(n-1)^{L-1}$ sub-matrix of $\Psi(\mathbf{s})$ that consists of only the columns of $\Psi(\mathbf{s})$ corresponding to outcomes in which player i wins exactly one lot. Call this matrix $\tilde{\Psi}$.
2. Row l, column a of $\tilde{\Psi}$ is strictly positive for columns corresponding to outcomes $\mathbf{w}^a$ in which bidder i wins lot l This arises because the probability that i wins lot l, and no other lot, is strictly increasing in b_l.
3. Row l, column a of $\tilde{\Psi}$ is strictly negative for columns corresponding to outcomes $\mathbf{w}^a$ in which bidder i does not win lot l. This arises because the probability that lot l is won, and no other lot is won, is strictly decreasing in b_m for $m \neq l$.
4. Any two rows of this matrix are linearly independent: Each row contains just one positive entry, each in a distinct column.[2] Therefore, $\tilde{\Psi}$, and hence $\Psi(\mathbf{s})$ has rank at least 2.
5. The matrix $A_\mathbf{s}$ is a rank n^L transformation matrix for non-maximal $\mathbf{s}$.
6. From steps 5 and 4, $\Psi(\mathbf{s})A_\mathbf{s}$ for non-maximal $\mathbf{s}$ has rank at least 2.

□

D.1.3 Rank(Ψ)

Proposition 7. *$Rank(\Psi) = S - S^c - |\tilde{\min}(\mathbb{S})|$*

Proving this proposition involves demonstrating that as we stack these $\Psi(\mathbf{s})A_\mathbf{s}$ matrices for non-maximal $\mathbf{s}$, the rank increases by *at least* two each time. By definition columns corresponding to elements in $\tilde{\min}(\mathbb{S})$ do not contain any non-zero entries, ensuring the rank must be deficient by at least $|\tilde{\min}(\mathbb{S})|$. Likewise, for each submatrix of Ψ made up of rows corresponding to states that are all within the same component (denoted by Ψ^c, a $|\mathbb{S}^c| \times S$ matrix), the rows all sum to zero. This ensures each Ψ^c is rank deficient by at least one, so Ψ is rank deficient by at least S^c.

[2]Note that this only holds for $L \geq 3$. For the case where $L = 2$ we must also assume $E[\Gamma_1 + \Gamma_2] \neq 1$, which holds generically.

Proof:

1. Arbitrarily order and label elements of $\mathbb{S}$ (hence the columns of Ψ) according to the partial ordering $\succeq^*$. Incomparable states can be ordered at random. This ensures that, for each $\mathbf{s}$, the furthest left non-zero column of $\Psi(\mathbf{s})A_{\mathbf{s}}$ is in the column corresponding to the ex-post state that corresponds to player i winning every lot having begun in state $\mathbf{s}$, $\mathbf{s}^{all_i}$.
2. Focus on one component, $\mathbb{S}^c$. Find the 'smallest' state within $\mathbb{S}^c$, smallest in terms of the ordering and labelling of step 1. Call this state $\mathbf{s}_1^c$. This element will be one of the minimal elements of $\mathbb{S}^c$, and, if the set is non-empty, will be an element of $\tilde{\min}(\mathbb{S}^c)$.
3. Next, find the second smallest state, which may also be a minimal element, and call this state $\mathbf{s}_2^c$. Vertically stack the two matrices $\Psi(\mathbf{s}_1^c)A_{\mathbf{s}_1^c}$ and $\Psi(\mathbf{s}_2^c)A_{\mathbf{s}_2^c}$, naming this matrix $\Psi^c_{\{1,2\}}$.
4. The matrix $\Psi^c_{\{1,2\}}$ has rank at least 4. Lemma D.1.2 ensures that both matrices have rank 2, while lemma D.1.1 ensures that each row of $\Psi(\mathbf{s}_1^c)A_{\mathbf{s}_1^c}$ is linearly independent of each row of $\Psi(\mathbf{s}_2^c)A_{\mathbf{s}_2^c}$. This last point arises because lemma D.1.1 ensures that since $\mathbf{s}_1^c \not\succeq^* \mathbf{s}_2^c$ there must be at least one column of non-zero entries in $\Psi(\mathbf{s}_2^c)A_{\mathbf{s}_2^c}$ that matches up to an all-zero column of $\Psi(\mathbf{s}_1^c)A_{\mathbf{s}_1^c}$.
5. Continue this process for each of the non-maximal states in component $\mathbb{S}^c$. At each stage, based on how we ordered the elements of $\mathbb{S}$ at step 1, and from lemmas D.1.2 and D.1.1, $\Psi(\mathbf{s}_n^c)A_{\mathbf{s}_n^c}$ must always contain at least one non-zero column that matches up to an all-zero column of $\Psi^c_{\{1,2\ldots n-1\}}$. In general, this will be the furthest left column, corresponding to the state $\mathbf{s}_n^{c\,all_i}$. Therefore, the rank will increase by at least 2 at each step.
6. The final matrix $\Psi^c_{\{1,2\ldots\}}$ has non-zero entries *somewhere* in each of the $|\mathbb{S}^c|$ columns corresponding to states in this set, except for columns correspond to elements of $\tilde{\min}(\mathbb{S}^c)$. These columns will contain all zeros, since there is always zero probability of ending a period in one of these states. As the rank of this matrix increased by at least two at each additional non-maximal state, and because we have at least as many non-maximal states as maximal states, this matrix must have rank at least $|\mathbb{S}^c| - |\tilde{min}(\mathbb{S}^c)| - 1$. The rank cannot be strictly greater than this. The rank must be strictly less than $|\mathbb{S}^c| - |\tilde{min}(\mathbb{S}^c)|$ because the row sum for each row of this final matrix must equal zero, a property inherited from the fact that $Q^T\boldsymbol{\iota} = 1$.
7. Any two components $\mathbb{S}^c$ and $\mathbb{S}^{c'}$ are mutually exclusive. Therefore, the two matrices for any two components $\Psi^c_{\{1,2\ldots\}}$ do not share any non-zero columns. Therefore, when we stack these matrices across different components, the ranks must sum together at each step.
8. Therefore $rank(\Psi) = \sum_{\mathbb{S}^c \subset \mathbb{S}} |\mathbb{S}^c| - |\tilde{min}(\mathbb{S}^c)| - 1 = S - |\tilde{\min}(\mathbb{S})| - S^c$

$\square$

One final thing of note: In general, the matrix $\Psi(\mathbf{s})A_{\mathbf{s}}$ has rank L. Essentially, that each state gives us L pieces of information, rather than just two pieces of information. Proof of this proposition has proven elusive, even though it is intuitively sensible and is observed empirically. This feature ensures that we only need $1/L$ of the states to be non-maximal states, rather than requiring at least half be non-maximal.

D.2 nullspace of Ψ

The aim is to find the $|\tilde{\min}(\mathbb{S})| + S^c$ elements of this null space. Intuitively, there are two types of elements in this null space: Those corresponding to the $|\tilde{\min}(\mathbb{S})|$ elements, and those corresponding to the remaining S^c elements.

D.2.1 The $|\tilde{\min}(\mathbb{S})|$ elements

Immediately obvious is that any vector $\mathbf{y}$ that only contains non-zero entries in rows corresponding to elements of the set $\tilde{\min}(\mathbb{S})$ is in this null space. This is because Ψ contains all zeros in columns corresponding to these states. Call this set of vectors $\mathbb{Y}^1$, which evidently contains $|\tilde{\min}(\mathbb{S})|$ distinct elements.

D.2.2 The S^c elements

Consider the vector $\mathbf{y}$ such that $y_{\mathbf{s}} = y_{\mathbf{s}'}$ if $\mathbf{s}$ and $\mathbf{s}'$ belong to the same component. Call this set of vectors $\mathbb{Y}^2$, which evidently contains S^c distinct elements.

As above, denote Ψ^c the $|\mathbb{S}^c| \times S$ submatrix of Ψ that takes rows corresponding to states within component c. As established previously, columns of this matrix that correspond to states in different components only contain zeros, which follows from the definition of a component.

Therefore, for any $\mathbf{y} \in \mathbb{Y}^2$ we have $\Psi^c\mathbf{y} = 0$. This is because entries of $\mathbf{y}$ are constant across rows that correspond to the non-zero entries of Ψ^c. Clearly this holds for any c. Therefore, as we stack the Ψ^cs into Ψ we will have $\Psi\mathbf{y} = 0$ for any $\mathbf{y} \in \mathbb{Y}^2$.

D.3 Image of $(I_S - \beta T\Omega)^{-1}C$

I previously established that the null space of Ψ is given by $\mathbb{Y}^1 \cup \mathbb{Y}^2$, containing $|\tilde{\min}(\mathbb{S})| + S^c$ elements. I now show that the intersection of this space and the image of $(I_S - \beta T\Omega)^{-1}C$ contains a single element - the constant vector, denoted ι_{S^i}. Establishing the necessary result requires three additional lemmas:

D.3.1 Three Additional Lemmas

Lemma D.3.1. *For any* $\mathbf{y} \in \mathbb{Y}^1$ *we have* $\Omega\mathbf{y} = 0$.

Proof:

1. Recall that $\Omega(\mathbf{s}) = E_{\mathbf{b}}[Q(\mathbf{b}^*|\mathbf{s})^T - \Gamma(\mathbf{b}^*|\mathbf{s})^T \nabla_{\mathbf{b}} \Gamma(\mathbf{b}^*|\mathbf{s})^{-1} \nabla_{\mathbf{b}} Q(\mathbf{b}^*|\mathbf{s})|\mathbf{s}] A_{\mathbf{s}}$
2. $A_{\mathbf{s}}\mathbf{y} = 0$ for $\mathbf{y} \in \mathbb{Y}^1$. This is because $A_{\mathbf{s}}$ selects the elements of $\mathbf{y}$ that correspond to possible ex-post states given beginning the period in state $\mathbf{s}$. But we know that $\mathbf{y}$ only contains non-zero entries in elements that correspond to states that are never observed as possible ex-post states.

□

Lemma D.3.2. *For any* $\mathbf{y} \in \mathbb{Y}^2$ *we have* $\mathbf{y} = \Omega\mathbf{y}$. *That is, every element of* $\mathbb{Y}^2$ *is an eignevector of the matrix* Ω *with eigenvalue* 1.

Proof:

1. Recall that $\Omega(\mathbf{s}) = E_{\mathbf{b}}[Q(\mathbf{b}^*|\mathbf{s})^T - \Gamma(\mathbf{b}^*|\mathbf{s})^T \nabla_{\mathbf{b}} \Gamma(\mathbf{b}^*|\mathbf{s})^{-1} \nabla_{\mathbf{b}} Q(\mathbf{b}^*|\mathbf{s})|\mathbf{s}] A_{\mathbf{s}}$
2. Note that for $\mathbf{y} \in \mathbb{Y}^2$ $A_{\mathbf{s}}\mathbf{y} = y_{\mathbf{s}}\iota_{2^L}$ Where ι_{2^L} is a 2^L dimensional vector of ones. This holds because $A_{\mathbf{s}}$ selects the elements of the vector $\mathbf{y}$ that correspond to states that are possible outcomes from an auction round beginning in state $\mathbf{s}$. By definition of a component every one of these elements of $\mathbf{y}$ lie in the same component.
3. As the rows of $Q(\mathbf{b}^*|\mathbf{s})^T$ sum to one, we have $E_{\mathbf{b}}[Q(\mathbf{b}^*|\mathbf{s})^T|\mathbf{s}]\iota_{2^L} = \iota_{2^L}$.
4. As the rows of $\nabla_{\mathbf{b}} Q(\mathbf{b}^*|\mathbf{s})$ sum to zero (derivative of a vector function with rows summing to one) we have:
$$E[\Gamma(\mathbf{b}^*|\mathbf{s})^T \nabla_{\mathbf{b}} \Gamma(\mathbf{b}^*|\mathbf{s})^{-1} \nabla_{\mathbf{b}} Q(\mathbf{b}^*|\mathbf{s})|\mathbf{s}]\iota_{2^L} = \mathbf{0}$$
5. Therefore $\Omega(\mathbf{s})\mathbf{y} = y_{\mathbf{s}}\iota_{2^L}$ for $\mathbf{y} \in \mathbb{Y}^2$
6. Therefore, as we stack our $\Omega(\mathbf{s})\mathbf{y}$s over $\mathbf{s}$ we are left with $\mathbf{y}$.

□

Finally, for $\mathbf{y} \in \mathbb{Y}^2$ we can write $\mathbf{y} = M\bar{\mathbf{y}}$ Where $\bar{\mathbf{y}}$ is an $S^c \times 1$ vector that contains each of the constant elements of $\mathbf{y}$, one for each component. Meanwhile M is an $S \times S^c$ dimensional matrix that contains a 1 in a row corresponding to state $\mathbf{s}$ and column corresponding to component c if $\mathbf{s} \in \mathbb{S}^c$, and zero otherwise. Therefore each row of M contains a single 1.

Lemma D.3.3. *Let the matrix* N *be any* $S^c \times S^c$ *submatrix of* $(I - \beta T)M$ *that is formed by selecting one row from each of the* S^c *components.* N *is non-singular.*

Proof:

1. Suppose we select S^c rows corresponding to states from different components, denoting this set of rows by $\mathbb{M}$. Therefore, the sub-matrix of interest is denoted $M_{\mathbb{M},.} - \beta T_{\mathbb{M},.}M$

2. We immediately have that $M_{\mathbb{M},.} = I$ from the definition of the matrix M. This is because we choose one row associated with each component. Each row of M contains a single 1, therefore so must $M_{\mathbb{M},.}$. Likewise, because every row is associated with a different component, each row contains a 1 in a different column.
3. Elements of the $S^c \times S^c$ sub matrix $T_{\mathbb{M},.}M$ are still just transition probabilities, so that $T_{\mathbb{M},.}M\iota_{S^c} = 1$. This is because right multiplying by M causes us to sum over states within a component. For row a particular row t we have element c of the row vector $T_{t,.}M$ is equal to $\sum_{\mathbf{s}\,:\,\mathbf{s}^c=\mathbf{s}^{\bar{c}}} P(\mathbf{s}|\mathbf{s}^t)$. That is, it is the probability, given ending a period in state $\mathbf{s}^t$, that they begin the next period in any state from component c.
4. Diagonal entries of the matrix $I - \beta T_{\mathbb{M},.}M$ must be strictly positive, as $\beta \times$ a probability must be strictly less than 1 (for $\beta < 1$). Likewise, off diagonal entries must be weakly negative, as we have $-\beta \times$ a probability. Last, rows must sum to $1 - \beta$. This is because rows of I evidently sum to 1, while rows of $T_{\mathbb{M},.}M$ also sum to 1.
5. This ensures this matrix is strictly diagonally dominant. Therefore, from the Levy–Desplanques theorem, the matrix must be non-singular.

□

D.3.2 Proof of Image Result

Proposition 8. *Image*$((I_S - \beta T\Omega)^{-1}C) \cap$ *null*$(\Psi) = \iota_{S^i}$

Proof of this proposition also makes use of the result that $T\iota_S = \iota_S$ — rows of a transition matrix sum to one. The proof proceeds by first demonstrating that the image of $(I_S - \beta T\Omega)^{-1}C$ doesn't intersect $\mathbb{Y}^1$ at all. I then prove that the intersection with $\mathbb{Y}^2$ only contains a single element - the constant vector.

Proof:

1. Suppose there exists an $\mathbf{x}$ such that for some $\mathbf{y} \in \mathbb{Y}^1$ we could write $\mathbf{y} = (I_S - \beta T\Omega)^{-1}C\mathbf{x}$.
2. This implies $(I_S - \beta T\Omega)\mathbf{y} = C\mathbf{x}$.
3. From Lemma D.3.1 this implies $\mathbf{y} = C\mathbf{x}$. In turn, from the definition of C this requires $\mathbf{x}$ contains zeros in every entry except the first.
4. However this cannot be the case, since we always normalise this first entry to zero. Therefore *image*$((I_S - \beta T\Omega)^{-1}C) \cap \mathbb{Y}^1 = \emptyset$
5. I will now show that this image does not intersect with $\mathbb{Y}^2$ apart from a single element. Suppose there exists an $\mathbf{x}$ such that for some $\mathbf{y} \in \mathbb{Y}^2$ we could write $\mathbf{y} = (I - \beta T\Omega)^{-1}C\mathbf{x}$. Or, equivalently, such that $(I - \beta T\Omega)\mathbf{y} = C\mathbf{x}$
6. From Lemma D.3.2 This requires $(I - \beta T)\mathbf{y} = C\mathbf{x}$, or $(I - \beta T)M\bar{\mathbf{y}} = C\mathbf{x}$. We can then write this in matrix form:

$$\begin{pmatrix} M - \beta\bar{T} & -C \end{pmatrix} \begin{pmatrix} \bar{\mathbf{y}} \\ \mathbf{x} \end{pmatrix} = 0$$

Where $\bar{T} = TM$, essentially summing over the probability of transitioning to any given component from any given state. Therefore, if $(M - \beta\bar{T}, -C)$, the $S \times (S_C + S_i)$ matrix has rank $S_C + S_i - 1$ then there must be a unique $\mathbf{y}$ and $\mathbf{x}$ such that this relationship holds.

7. Consider whether the first column of $-C$ is linearly independent of the columns of $(M - \beta\bar{T})$. $-C_{.,1}$ contains -1 in every element associated with states such that $\mathbf{s}_i = \mathbf{s}_i^1$ and zeros otherwise. I now show that no linear combination of the columns for the corresponding rows of $(M - \beta\bar{T})$ can match these zeros. Choose S_c rows of $(M - \beta\bar{T})$ such that each row is associated with a state from a different component. For example, we might choose rows such that in each component $\mathbf{s}_i = \mathbf{s}_i^{S_i}$ - the 'final' individual state. Call the $S_c \times S_c$ submatrix of $M - \beta\bar{T}$ made of these rows (and all columns) N.
8. From Lemma D.3.3 N must be non-singular. Therefore there does not exist an $S_c \times 1$ vector $\mathbf{z}$ such that $N\mathbf{z} = 0$. Therefore columns of $(M - \beta\bar{T})$ must be linearly independent of $-C_{.,1}$. So, by catenating on this new column, the rank must increase by one.
9. Repeat this process for columns $n = 1...S_i - 1$ of $-C$. That is, every column *except* the final column which is the only column to contain non-zeros in entries associated with $\mathbf{s}_i^{S_i}$.[3] Each of these columns must be linearly independent of $M - \beta\bar{T}$ - no linear combination of its columns can match the zero entries of $-C_{.,n}$, since any $S^c \times S^c$ submatrix that consists of one row from each component must be non-singular.
10. Likewise, columns of $-C$ are all linearly independent of each other.
11. Therefore, at each step $n = 1...S_i - 1$, the rank of our catenated matrix increases by 1. Therefore $rank(M - \beta\bar{T}, -C) \geq S_C + S_i - 1$.
12. The vector $(\bar{\mathbf{y}} = \iota_{S^c}, \mathbf{x} = (1-\beta)\iota_{S^i})$ lies in the null space of $(M - \beta\bar{T}, -C)$. This is evident since $(M - \beta\bar{T})\iota_{S^c} = (1-\beta)\iota_S$ while we also have $C(1 - \beta)\iota_{S_i} = (1-\beta)\iota_S$. Therefore, applying the rank-nullity theorem ensures that $Image((I_S - \beta T\Omega)^{-1}C) \cap null(\Psi) = \iota_{S^i}$

□

[3] This assumes that one individual state exists within each component (here, I used $\mathbf{s}_i^{S_i}$). This holds if for example, that $\mathbb{S} = \mathbb{S}^0 \times \prod_i \mathbb{S}_i$. However this is not strictly necessary. The only requirement is that at each step n I can select one row corresponding to a state from each component such that the corresponding rows of $-C_{.,n}$ are all zero.

Appendix E

Extensions

E.1 Second-Price Auctions

In this Appendix I show that the previous identification results extend, almost trivially, to the second price case. In Appendix E.1.1 I set up the bidder's optimisation problem in the second price framework. In Appendix E.1.2 I show how optimal bidding yields a set of First Order Conditions, and hence an Inverse Bid System, similar to the first-price case. As before, I then show F is identified, conditional on $\mathbf{j}$ and V, from the inverse bid system. In Appendix E.1.3 I extend Proposition 3 from the main text to the second-price case, showing that maximised expected pay-off can be written as a function of the observed distribution of bids. I then show that V is identified conditional on $\mathbf{j}$. In Appendix E.1.4 I prove that $\mathbf{j}$ is point identified from the same moment condition assumed in the main text.

I do not discuss estimation of the dynamic multi-object second price model. However the estimation procedure presented in Section 3.4 can be applied to the second price setting, making use of the inverse bid system presented below.

E.1.1 Setup

In the second price setting, player i wins lot l at time t if $b_{ilt} > \max_{i'} \{b_{i'lt}\}$. As in the text, let $\Gamma(\mathbf{b}|\mathbf{s})$ denote the $L \times 1$ equilibrium marginal probabilities of winning each lot. Define the vectors P and Q similarly. The Value Function is given by:

$$
\begin{aligned}
& W_i(v_{it}, \mathbf{s}_t; \sigma_{-i}) \\
&= \max_{\mathbf{b}} \left\{ \Gamma_i(\mathbf{b}; \sigma_{-i})^T (v_i - \tilde{\mathbf{b}}(\mathbf{b}; \mathbf{s}_t)) + P_i(\mathbf{b}; \sigma_{-i})^T J_i(\mathbf{s}_t) + \beta Q_i(\mathbf{b}; \sigma_{-i})^T V_i(\mathbf{s}_t; \sigma_{-i}) \right\}
\end{aligned} \tag{E.1}
$$

Where element a of the continuation value V_i is given by:

$$
V_{ia}(\mathbf{s}_t; \sigma_{-i}) = \int_{\mathbf{s}} \int_{v} W_i(v, \mathbf{s}; \sigma_{-i}) dF(v|\mathbf{s}) dT(\mathbf{s}|\mathbf{s}_t^a)
$$

$\tilde{\mathbf{b}}(\mathbf{b}; \mathbf{s}_t)$ gives the expected second highest bid, given that b_{ilt} is the highest. Since the cdf of the highest rival bids is $\Gamma_l(x|\mathbf{s})$, we can write $\Gamma_l(b_l|\mathbf{s})\tilde{b}_l(\mathbf{b}; \mathbf{s})) = \int_{\underline{b}_l}^{b_{ilt}} \bar{b}_l \nabla_{b_l} \Gamma_l(\bar{b}_l|\mathbf{s}) d\bar{b}_l$.

E.1.2 First Order Conditions and Inverse Bid System

Rearrange the maximand for:

$$
\Gamma(\mathbf{b}|\mathbf{s})^T v - \sum_l \int_{\underline{b}_l}^{b_l} \bar{b}_l \nabla_{b_l} \Gamma_l(\bar{b}_l|\mathbf{s}) d\bar{b}_l + P(\mathbf{b}|\mathbf{s}) J(\mathbf{s}) + \beta Q(\mathbf{b}|\mathbf{s}) V(\mathbf{s})
$$

differentiating for FOCs: $0 = \nabla_{\mathbf{b}} \Gamma(\mathbf{b}^*|\mathbf{s})(v - \mathbf{b}^*) + \nabla_{\mathbf{b}} P(\mathbf{b}^*|\mathbf{s}) J(\mathbf{s}) + \beta \nabla_{\mathbf{b}} Q(\mathbf{b}^*|\mathbf{s}) V(\mathbf{s})$.[1] We then invert the FOCs for the inverse bid system:

$$
\xi(\mathbf{b}_{it}|J, \beta V; \mathbf{s}) = \mathbf{b}_{it} - \nabla_{\mathbf{b}} \Gamma(\mathbf{b}^*|\mathbf{s})^{-1} [\nabla_{\mathbf{b}} P(\mathbf{b}^*|\mathbf{s}) B_{\mathbf{s}} \mathbf{j} + \nabla_{\mathbf{b}} Q(\mathbf{b}^*|\mathbf{s}) A_{\mathbf{s}} \beta \mathbf{V}]
$$

This is extremely similar to the inverse bid system presented in the main text, simply omitting the mark-up term $\nabla_{\mathbf{b}} \Gamma(\mathbf{b}^*|\mathbf{s})^{-1} \Gamma(\mathbf{b}^*|\mathbf{s})$. It is then clear that, conditional on $\mathbf{j}$ and $\beta \mathbf{V}$, the distribution of lot specific values F is point identified from the empirical quantiles of $\xi(\mathbf{b}_{it}|J, \beta V; \mathbf{s})$.

E.1.3 Extension of Proposition 3

I now extend Proposition 3 to the second price case. Note that there are many ways I could prove this general second price identification argument. I use this structure

[1] Interestingly, the condition for optimal bidding can equivalently be derived from the requirement that, at the optimum, b_{lt}^* must equal the marginal expected pay-off from winning lot l, conditional on bids for lots $m \neq l$.

for the purposes of outlining the similarity to the first price case.

Proposition 9. *Under assumptions 1 - 4, the expected stage pay-off is given by:*

$$\begin{aligned}\tilde{\Pi}(\mathbf{b}^*|v;\mathbf{s}) &= \Gamma(\mathbf{b}^*|\mathbf{s})^T(\mathbf{b}^* - \tilde{\mathbf{b}}(\mathbf{b}^*;\mathbf{s})) \\ &\quad + [P(\mathbf{b}^*|\mathbf{s})^T \ - \ \Gamma(\mathbf{b}^*|\mathbf{s})^T \nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1}\nabla_{\mathbf{b}}P(\mathbf{b}^*|\mathbf{s})]B_{\mathbf{s}}\mathbf{j} \\ &\quad + [Q(\mathbf{b}^*|\mathbf{s})^T \ - \ \Gamma(\mathbf{b}^*|\mathbf{s})^T \nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1}\nabla_{\mathbf{b}}Q(\mathbf{b}^*|\mathbf{s})]A_{\mathbf{s}}\beta\mathbf{V} \end{aligned} \tag{E.2}$$

This is similar to the expression given in Proposition 3, except that the optimal lot specific surplus term is given by $\mathbf{b}^* - \tilde{\mathbf{b}}(\mathbf{b}^*;\mathbf{s})$ instead of $\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1}\Gamma(\mathbf{b}^*|\mathbf{s})$. Proof is omitted due to its simplicity - simply substitute the inverse bid function $\xi(\mathbf{b}_{it}|J,\beta V;\mathbf{s})$ for $\boldsymbol{v}$ into the maximand of the value function in equation E.1.

From Proposition 9, employing the identity $P(\mathbf{b}|\mathbf{s})^T B_{\mathbf{s}} = Q(\mathbf{b}|\mathbf{s})^T A_{\mathbf{s}} C$, and taking an expectation over the observed bids, write the ex-ante value function as:

$$\begin{aligned} V^e(\mathbf{s}) &= \Phi(\mathbf{s}) + \Omega(\mathbf{s})[C\mathbf{j} + \beta\mathbf{V}] \\ \text{Where} \qquad \Phi(\mathbf{s}) &= E_{\mathbf{b}}[\Gamma(\mathbf{b}^*|\mathbf{s})^T(\mathbf{b}^* - \tilde{\mathbf{b}}(\mathbf{b}^*;\mathbf{s}))|\mathbf{s}] \\ \Omega(\mathbf{s}) &= E_{\mathbf{b}}[Q(\mathbf{b}^*|\mathbf{s})^T - \Gamma(\mathbf{b}^*|\mathbf{s})^T\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1}\nabla_{\mathbf{b}}Q(\mathbf{b}^*|\mathbf{s})|\mathbf{s}]A_{\mathbf{s}} \end{aligned}$$

Stacking this equation over $\mathbf{s}$ allows us to write the continuation value as: $\mathbf{V} = T\Phi + T\Omega[C\mathbf{j} + \beta\mathbf{V}]$ Which we can invert for: $\mathbf{V} = (I_S - \beta T\Omega)^{-1}[T\Phi + T\Omega C\mathbf{j}]$. This yields a stationary solution for the continuation value. This is precisely the equation derived in the text, except I have defined the matrix $\Phi(\mathbf{s})$ slightly differently. Therefore $\mathbf{V}$ is point identified condition on $\mathbf{j}$.

E.1.4 Identification

As in the main text I impose the mean zero property of $\boldsymbol{v}$ for:

$$\begin{aligned} 0 = E_{\mathbf{b}^*}[\xi(\mathbf{b}^*;\mathbf{s},(\mathbf{j},\mathbf{V}))|\mathbf{s}] &= E_{\mathbf{b}^*}[\mathbf{b}^*|\mathbf{s}] - E_{\mathbf{b}^*}[\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1}\nabla_{\mathbf{b}}Q(\mathbf{b}^*|\mathbf{s})|\mathbf{s}]A_{\mathbf{s}}[C\mathbf{j} + \beta\mathbf{V}] \\ &= \Upsilon(\mathbf{s}) - \Psi(\mathbf{s})[C\mathbf{j} + \beta\mathbf{V}] \end{aligned}$$

Again, using a slightly differently defined $\Upsilon(\mathbf{s})$ from the text. Stack over $\mathbf{s}$, and substituting in the expression for the $\mathbf{V}$ found in subsection E.1.3, we get:

$$0 = \Upsilon - \beta\Psi(I_S - \beta T\Omega)^{-1}T\Phi - \Psi(I_S - \beta T\Omega)^{-1}C\mathbf{j}$$

There exists a unique solution to this system of equations ($\mathbf{j}$ is point identified) if and only if the $LS \times S_i$ matrix $\Psi(I_S - \beta T\Omega)^{-1}C$ has rank $S_i - 1$. This matrix is exactly the same as in the main text. Proposition 4 holds in this case as well, ensuring the sufficient rank condition.

E.2 Binding Reservation Prices

In this Appendix I introduce binding reservation prices. A reservation price is binding if, in equilibrium, there is a non-zero probability of winning a lot at the reservation price. Binding reservation prices do not pose a substantive problem, though do introduce additional mathematical complexity. I also introduce endogenous entry with zero entry costs. When entry is costless reservation prices are necessary to prevent bidders submitting arbitrarily low bids. In Appendix E.3 I allow for costly entry where lot specific values are observed after entry.

In the presence of reservation prices a bidder with a low value may choose not to bid strictly above the reservation price. This results in corner solutions as bids clump at the reservation price. We lose point identification as the FOCs no longer point identify $\boldsymbol{v}_i$. This is a problem, even in a single object context.

The identification argument presented below diverges from the argument presented in the main text. Instead, it shares DNA with the estimation method presented in Section 3.4. Identification is demonstrated in an additional step. First I show that F is (partially) identified conditional on (J, V, β), but in particular it is partially identified conditional on $J + \beta V$. I then show that the object $j(\mathbf{s}_i) + \beta V(\mathbf{s})$ is partially identified, so that for some $\mathbf{s}_i$ it may only be bounded. This is shown using quantile moment conditions: Instead finding the $j + \beta V$ such that $E[\boldsymbol{\xi}(\mathbf{b}; \mathbf{s}, j + \beta V)|\mathbf{s}] = 0$ I find it such that $P(\xi_l(\mathbf{b}; \mathbf{s}, j + \beta V) \leq 0|\mathbf{s}) = 0.5$, imposing a zero conditional median assumption. Finally, I show that conditional on the identification of

F and $J + \beta V$, V is identified, and hence J can be backed out given an assumption about β.

E.2.1 Changes to the Model

Denote the reservation price as R. This could vary across lots, bidders, and time. Denote player i's entry decisions as the vector $\mathbf{d}_{it}$ with entry $d_{itl} = 1$ if they enter lot l, and zero otherwise. Adjust objects G, Γ, P and Q to be functions of bids and entry — if a player does not enter a lot, they trivially lose that lot with probability 1.[2]

Identification requires one additional assumption:

Assumption 10. $\frac{\partial \Gamma_{il}(\mathbf{b}_i|\mathbf{s})}{\partial b_{im}} = 0$ for $m \neq l$

I require that the probability an individual wins any given lot, conditional on the distribution of everyone else's bids, must depend only on their bid for that lot. This assumption implies that $\nabla \Gamma_i(\mathbf{b}_i|\mathbf{s})$ is a diagonal matrix. This assumption was not previously necessary for identification. If ties happen with zero probability or if tie breaking is exogenous, then this assumption will hold.

For mathematical convenience I assume ties occur in equilibrium with zero probability. This assumption was relaxed in a previous version of this thesis, when I allowed for ties at the reservation price. The only thing this changed was that it introduced a discontinuity in the inverse bidding system at the reservation price, so that as the bidder goes from bidding the reserve to just above it, their payoff changes discontinuously. This slightly changed how we identify F, as we must essentially introduce an additional discrete choice of whether the bidder bids the reservation compared to bidding just above it. This additional discrete choice then restores the (upper-hemi) continuity of equilibrium, payoffs. Finally, I assume the lot specific values have zero conditional median, replacing the previous zero conditional mean assumption. I am then able to prove the following proposition:

Proposition 10. *Given assumption 1, 9, 3, 4, and 10, both $F_i(.|\mathbf{s})$ and $K_i(\mathbf{s})$ are non-parametrically partially identified. $k(\mathbf{s}^a)$ is point identified if we observe the individual bidding $b > R$ on a lot that may yield pay-off $k(\mathbf{s}^a)$.*

[2] In principle this means that $\nabla_{\mathbf{b}} \Gamma$ will no longer be invertible, since it may contain zero rows/columns for non-entered auctions. However, this is not a problem as a simple generalised inverse will be used instead.

That is, we will point identify the truncated distribution $F_i(.|\boldsymbol{v} >= A_1(\mathbf{b}^*, \mathbf{s}); \mathbf{s})$, as well as the objects $F_i(A_1(\mathbf{b}^*, \mathbf{s}); \mathbf{s}) - F_i(A_2(\mathbf{b}^*, \mathbf{s}); \mathbf{s})$ and $F_i(A_2(\mathbf{b}^*, \mathbf{s}); \mathbf{s})$ for some (known) $A_1(\mathbf{b}^*, \mathbf{s}), A_2(\mathbf{b}^*, \mathbf{s})$.

I follow the same identification framework as in the text. While I assume players play pure strategies conditional on entry, I must allow for the possibility that players play mixed strategies in their entry decisions. However, I am able to use bidders' entry decisions to bound the pay-offs of un-entered auctions. I exploit the fact that, at the equilibrium mixing strategy, players can not *strictly* prefer to enter any other combination of auctions.

E.2.2 Identification of F, conditional on K.

Proposition 11. *Under assumptions 1, 9, 3, 4, and 10, and conditional on K being point identified, the cdf F is non-parametrically partially identified.*

I prove Proposition 11 by proving that, similar to case 6.3.1.2 described in Athey and Haile (2007), we can invert observed bids, point identifying v_l such that $b_l > R$. Meanwhile, for bids at the reservation price, so that $b_l = R$, we can use the first order conditions from the Lagrangian to find an upper bound on v_l. I then use the fact that they still prefer to enter this auction and bid low, than not enter, to find a lower bound on v_l. Lastly, for un-entered auctions we use the fact, at the margin, they prefer not to enter than to enter and bid low, to find an upper bound on these v_l. Therefore, I am able to partially identify F.

First, reformulate the problem to include the entry decisions. The player's problem is to decide which auctions to enter ($\mathbf{d}$), and then set their bids ($\mathbf{b}$) to maximise utility, subject to the constraint that bids are weakly above reservation prices. The Lagrangian for this problem, conditional on entry $\mathbf{d}^*$, is given as:

$$L(\mathbf{b}, \mathbf{d}^*, \boldsymbol{v}, \boldsymbol{\lambda}|\mathbf{s}) = \Gamma(\mathbf{b}, \mathbf{d}^*|\mathbf{s})^T(\boldsymbol{v} - \mathbf{b}) + P(\mathbf{b}, \mathbf{d}^*|\mathbf{s})^T K + \boldsymbol{\lambda}^T(\mathbf{b} - R) \tag{E.3}$$

At the optimum their first order necessary conditions (with respect to bids) yields:

$$\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*, \mathbf{d}^*|\mathbf{s})(\boldsymbol{v} - \mathbf{b}^*) - \Gamma(\mathbf{b}^*, \mathbf{d}^*|\mathbf{s}) + \nabla_{\mathbf{b}}P(\mathbf{b}^*, \mathbf{d}^*|\mathbf{s})^T K + \boldsymbol{\lambda}^* = 0 \tag{E.4}$$

Entry ll of $\nabla_{\mathbf{b}}\Gamma(\mathbf{b},\mathbf{d}|\mathbf{s})$ and entry la of $\nabla_{\mathbf{b}}P(\mathbf{b},\mathbf{d}|\mathbf{s})$ are as they were in section 3.3 if $d_l = 1$, and normalised to 0 otherwise. Rearrange this equation for:[3]

$$\boldsymbol{\xi}(\mathbf{b}^*,\mathbf{d}^*,\boldsymbol{\lambda}|K;\mathbf{s}) = \mathbf{b}^* + \nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*,\mathbf{d}^*|\mathbf{s})^{-1}[\Gamma(\mathbf{b}^*,\mathbf{d}^*|\mathbf{s}) - \nabla_{\mathbf{b}}P(\mathbf{b}^*,\mathbf{d}^*|\mathbf{s})K] - \nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*,\mathbf{d}^*|\mathbf{s})^{-1}[\boldsymbol{\lambda}^*] \quad \text{(E.5)}$$

At the true K we have $\xi_l(\mathbf{b}^*,\mathbf{d}^*,\boldsymbol{\lambda}^*|K;\mathbf{s}) = v_l$. But we do not observe $\boldsymbol{\lambda}^*$. Therefore, define $\boldsymbol{\xi}(\mathbf{b}^*,\mathbf{d}^*|K;\mathbf{s}) = \mathbf{b}^* + \nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*,\mathbf{d}^*|\mathbf{s})^{-1}[\Gamma(\mathbf{b}^*,\mathbf{d}^*|\mathbf{s}) - \nabla_{\mathbf{b}}P(\mathbf{b}^*,\mathbf{d}^*|\mathbf{s})K]$. Next, I consider what can be inferred for each of the four possible entry/bidding possibilities: *i*) $b_l > R$, *ii*) $b_l = R$, *iii*) $d_l = 0$, and the null case $l \notin \mathbb{L}$.

i) l **such that** $b_l^* > R$**:**

Any entry l such that $b_l^* > R_l$, $\lambda_l^* = 0$. By Assumption 10, entry l of $\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*,\mathbf{d}^*|\mathbf{s})^{-1}[\boldsymbol{\lambda}^*]$ equals zero, and $\xi_l(\mathbf{b}^*,\mathbf{d}^*|K;\mathbf{s}) = v_l$ is point identified.

ii) l **such that** $b_l^* = R$**:**

For entry l with $b_l^* = R_l$, $\lambda_l^* > 0$. From Assumption 10 entry l of $\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*,\mathbf{d}^*|\mathbf{s})^{-1}[\boldsymbol{\lambda}^*]$ is greater than zero, and we attain the following bound:

$$v_l \le \xi_l(\mathbf{b}^*,\mathbf{d}^*|K;\mathbf{s}) = R_l + \nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*,\mathbf{d}^*|\mathbf{s})^{-1}[\Gamma(\mathbf{b}^*,\mathbf{d}^*|\mathbf{s}) - \nabla_{\mathbf{b}}P(\mathbf{b}^*,\mathbf{d}^*|\mathbf{s})K]_l$$

Where, for vector M, $[M]_l$ denotes row l. Given that $(\mathbf{b}^*,\mathbf{d}^*)$ globally maximises expected payoffs, payoffs are (weakly) higher from playing $(\mathbf{b}^*,\mathbf{d}^*)$ than not entering auction l, playing $(\mathbf{b}^{l-},\mathbf{d}^{l-})$ (the only difference between these actions is that $d_l^{l-} = 0$). Therefore:

$$\Gamma(\mathbf{b}^*,\mathbf{d}^*|\mathbf{s})^T(\boldsymbol{v} - \mathbf{b}^*) + P(\mathbf{b}^*,\mathbf{d}^*|\mathbf{s})^T K \ge \Gamma(\mathbf{b}^{l-},\mathbf{d}^{l-}|\mathbf{s})^T(\boldsymbol{v} - \mathbf{b}^{l-}) + P(\mathbf{b}^{l-},\mathbf{d}^{l-}|\mathbf{s})^T K \quad \text{(E.6)}$$

Which rearranges for: $\Gamma_l(b_l^*,d_l^*|\mathbf{s})(v_l - R_l) + [P(\mathbf{b}^*,\mathbf{d}^*|\mathbf{s}) - P(\mathbf{b}^{l-},\mathbf{d}^{l-}|\mathbf{s})]^T K \ge 0$, and hence $v_l \ge R_l - \frac{1}{\Gamma_l(b_l^*,d_l^*|\mathbf{s})}[P(\mathbf{b}^*,\mathbf{d}^*|\mathbf{s}) - P(\mathbf{b}^{l-},\mathbf{d}^{l-}|\mathbf{s})]^T K$.

Therefore, we can bound v_l between these two cut-offs: $v_l \in [A_1(\mathbf{b}^{l-},\mathbf{s}), A_2(\mathbf{b}^*,\mathbf{b}^{l-},\mathbf{s})]$.

[3] Using the pseudo-inverse for $\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*,\mathbf{d}^*|\mathbb{L})^{-1}$, as non-entry causes rows to be all zeros.

iii) l **such that** $d_l^* = 0$**:**

Consider some l such that $d_l = 0$. They must attain a greater payoff from not bidding than from bidding the reservation price. Consider an alternate action $(\mathbf{b}^{l+}, \mathbf{d}^{l+})$ where the difference between this and their chosen action is that $b_l^{l+} = R_l$ and $d_l^{l+} = 1$. Therefore:

$$\Gamma(\mathbf{b}^*, \mathbf{d}^*|\mathbf{s})^T(v - \mathbf{b}^*) + P(\mathbf{b}^*, \mathbf{d}^*|\mathbf{s})^T K \geq \Gamma(\mathbf{b}^{l+}, \mathbf{d}^{l+}|\mathbf{s})^T(v - \mathbf{b}^{l+}) + P(\mathbf{b}^{l+}, \mathbf{d}^{l+}|\mathbf{s})^T K \tag{E.7}$$

rearranging for: $-\Gamma_l(b_l^{l+}, d_l^{l+}|\mathbf{s})(v_l - R_l) + [P(\mathbf{b}^*, \mathbf{d}^*|\mathbf{s}) - P(\mathbf{b}^{l+}, \mathbf{d}^{l+}|\mathbf{s})]^T K \geq 0$ Yielding the bound: $v_l < R_l - \frac{1}{\Gamma_l(b_l^{l+}, d_l^{l+}|\mathbf{s})}[P(\mathbf{b}^*, \mathbf{d}^*|\mathbf{s}) - P(\mathbf{b}^{l+}, \mathbf{d}^{l+}|\mathbf{s})]^T K$

Therefore, bids and entry decisions enable partial identification of F.

E.2.3 Identification of k under binding reservation prices

Proposition 12. *Under assumptions 1, 9, 3, 4, and 10, the function k is non-parametrically partially identified up to standard normalisations. $k(\bar{\mathbf{s}})$ is point identified at $\mathbf{s} = \bar{\mathbf{s}}$ if we observe bidding strictly above R on a combination of goods that would have the outcome $\mathbf{s}^a = \bar{\mathbf{s}}$.*

I prove this proposition by exploiting multiple observations for every state to establish a necessary rank condition, similar to the one presented in Section 3.3.6. This enables me to overcome the inherent order condition. Whereas the previous proof then employed a condition on the mean of $\xi(\mathbf{b}, \mathbf{d})$, this proof essentially employs a condition on the marginal quantiles of $\xi(\mathbf{b}, \mathbf{d})$. I set $k(\mathbf{s})$ such that the median (or some other quantile) is equal to zero.[4] As above, binding reservation prices cause our FOCs to break down, so that even at the true $\mathbf{k}$ $(= C\mathbf{j} + \beta\mathbf{V})$ we can only write:

$$v \leq \xi(\mathbf{b}, \mathbf{d}|k; \mathbf{s}) = \mathbf{b} + \nabla_{\mathbf{b}}\Gamma(\mathbf{b}, \mathbf{d}|\mathbf{s})^{-1}[\Gamma(\mathbf{b}, \mathbf{d}|\mathbf{s}) - \nabla_{\mathbf{b}}P(\mathbf{b}, \mathbf{d}|\mathbf{s})A_{\mathbf{s}}\mathbf{k}] \tag{E.8}$$

[4]Note that additional exclusions restrictions on F also aid identification, particular if we observe a large degree of constrained bidding. For example, we might impose that F is independent of the set of lots on offer (marginal distributions do not vary with $\mathbb{L}_t$). We gain power as we must also set k to ensure that the marginal distributions of inverse bids are invariant to $\mathbb{L}$. This argument appeared in a early draft of this thesis.

Which only holds with equality for rows l with $b_l > R$. Stack these over $\mathbf{s}$ for:

$$\underbrace{\underline{v}}_{LS\times 1} \leq \underline{\xi}(\underline{\mathbf{b}}, \underline{\mathbf{d}}|k) = \underbrace{\underline{\mathbf{b}}}_{LS\times 1} + \underbrace{\nabla_{\underline{\mathbf{b}}}\underline{\Gamma}(\underline{\mathbf{b}}, \underline{\mathbf{d}})^{-1}}_{LS\times LS}[\underbrace{\underline{\Gamma}(\underline{\mathbf{b}}, \underline{\mathbf{d}})}_{LS\times 1} - \underbrace{\nabla_{\underline{\mathbf{b}}}\underline{P}(\underline{\mathbf{b}}, \underline{\mathbf{d}})}_{LS\times S}\,\mathbf{k}] \tag{E.9}$$

Where:

$$\underline{\xi}(\underline{\mathbf{b}}, \underline{\mathbf{d}}|k) = \begin{pmatrix} \xi(\mathbf{b}_1, \mathbf{d}_1|k; \mathbf{s}_1) \\ \vdots \\ \xi(\mathbf{b}_S, \mathbf{d}_S|k; \mathbf{s}_S) \end{pmatrix} \qquad \underline{\mathbf{b}} = \begin{pmatrix} \mathbf{b}_1 \\ \vdots \\ \mathbf{b}_S \end{pmatrix}$$

$$\underline{\Gamma}(\underline{\mathbf{b}}, \underline{\mathbf{d}}) = \begin{pmatrix} \Gamma(\mathbf{b}_1, \mathbf{d}_1|\mathbf{s}_1) \\ \vdots \\ \Gamma(\mathbf{b}_S, \mathbf{d}_S|\mathbf{s}_S) \end{pmatrix} \qquad \nabla_{\underline{\mathbf{b}}}\underline{P}(\underline{\mathbf{b}}, \underline{\mathbf{d}}) = \begin{pmatrix} \nabla_{\mathbf{b}}P(\mathbf{b}_1, \mathbf{d}_1|\mathbf{s}_1)A_{\mathbf{s}_1} \\ \vdots \\ \nabla_{\mathbf{b}}P(\mathbf{b}_S, \mathbf{d}_S|\mathbf{s}_S)A_{\mathbf{s}_S} \end{pmatrix} \tag{E.10}$$

We will require a rank condition on $\nabla_{\underline{\mathbf{b}}}\underline{\Gamma}(\underline{\mathbf{b}}, \underline{\mathbf{d}})^{-1}\nabla_{\underline{\mathbf{b}}}\underline{P}(\underline{\mathbf{b}}, \underline{\mathbf{d}})$. If this has full rank then each possible $\underline{\xi}$ implies a unique $\mathbf{k}$, so that if we observed just one observation of $\underline{\xi}$ we could solve for $\mathbf{k}$. Note that $E[\nabla_{\underline{\mathbf{b}}}\underline{\Gamma}(\underline{\mathbf{b}}, \underline{\mathbf{d}})^{-1}\nabla_{\underline{\mathbf{b}}}\underline{P}(\underline{\mathbf{b}}, \underline{\mathbf{d}})] = \Psi$, the matrix presented in the main text. Importantly, the proof presented in D.1, that $Rank(\Psi) = S - S^c - |\tilde{\min}(\mathbb{S})|$ extends trivially to $\nabla_{\underline{\mathbf{b}}}\underline{\Gamma}(\underline{\mathbf{b}}, \underline{\mathbf{d}})^{-1}\nabla_{\underline{\mathbf{b}}}\underline{P}(\underline{\mathbf{b}}, \underline{\mathbf{d}})$. The proof never exploited the fact we had taken an expectation, and entirely used the partial ordering structure of the state space. However, we must normalise elements of $\mathbf{k}$ associated with the null space of Ψ, discussed shortly.

With binding reservation prices and entry, certain states may never be outcomes that *could have* occurred with positive probability, so the corresponding elements of $\mathbf{k}$ are not point identified. These entries of $\mathbf{k}$ do not appear in the above equation, having a coefficient of zero. These states will only be partially identified.

Next, fix an $LS \times 1$ vector of probabilities $\mathbf{p}$. Because of the median zero assumption each element will equal 0.5, but in principal we might use a higher quantile to identify as many elements of $\mathbf{k}$ as possible. By definition of the marginal CDF, the following relationship holds:

$$\begin{pmatrix} p_1 \\ \vdots \\ p_{LS} \end{pmatrix} = \begin{pmatrix} F_1(\tilde{v}_1|\mathbf{s}_1) \\ \vdots \\ F_L(\tilde{v}_{LS}|\mathbf{s}_S) \end{pmatrix} = \begin{pmatrix} E_{v_1}[\ \mathbb{I}[v_1 \leq \tilde{v}_1]\ |\mathbf{s}_1] \\ \vdots \\ E_{v_L}[\ \mathbb{I}[v_L \leq \tilde{v}_{LS}]\ |\mathbf{s}_S] \end{pmatrix} \tag{E.11}$$

Next, we employ a change of variables, taking expectations over the observed random variables $(\mathbf{B}, \mathbf{D})$ rather than v_l. This change is only valid for state-lot combinations such that when $v_l = \tilde{v}_l$, $b_l > R$, because only then $\xi_l(\mathbf{B}, \mathbf{D}; k) = v_l$ holds with equality, and so the mapping from $\mathbf{B}$ to v_l is continuous, smooth, and monotonic.[5] We must to drop rows where this condition fails, as we lose identifiability of corresponding elements of $\mathbf{k}$. The idea is that if, even when v_l is as large as $\tilde{v}_l$, the elements of $K(\mathbf{s})$ corresponding to winning lot l are so small that they never bid strictly above R on lot l, then these elements of $K(\mathbf{s})$ will not be identified.

This change of variables yields:

$$\mathbf{p} = \begin{pmatrix} E_{v_1}[\ \mathbb{I}[v_1 \leq \tilde{v}_1]\ |\mathbf{s}_1] \\ \vdots \\ E_{v_L}[\ \mathbb{I}[v_L \leq \tilde{v}_{LS}]\ |\mathbf{s}_S] \end{pmatrix} = \begin{pmatrix} E_{\mathbf{B},\mathbf{D}}[\ \mathbb{I}[\xi_1(\mathbf{B}_1, \mathbf{D}_1; k) \leq \tilde{v}_1]\ |\mathbf{s}_1] \\ \vdots \\ E_{\mathbf{B},\mathbf{D}}[\ \mathbb{I}[\xi_L(\mathbf{B}_S, \mathbf{D}_S; k) \leq \tilde{v}_{LS}]\ |\mathbf{s}_S] \end{pmatrix} \tag{E.12}$$

Proving point identification of $\mathbf{k}$ requires we show that this equation only holds at the true $\mathbf{k}$. We must show that the $\mathbf{p}$th quantiles of $\underline{\boldsymbol{\xi}}(\underline{\mathbf{B}}, \underline{\mathbf{D}}|k)$ equals $\tilde{\boldsymbol{v}}$ only at the true $\mathbf{k}$. But, from our rank condition, a unique $\underline{\boldsymbol{\xi}}(\underline{\mathbf{B}}, \underline{\mathbf{D}}|k)$ implies a unique $\mathbf{k}$. Therefore, only a unique $\mathbf{k}$ is such that the $\mathbf{p}$th quantiles of $\underline{\boldsymbol{\xi}}(\underline{\mathbf{B}}, \underline{\mathbf{D}}|k)$ equals $\tilde{\boldsymbol{v}}$. Therefore, there exists a unique $\mathbf{k}$ such that this equation holds.[6]

[5] This is essentially an application of the Law of the Unconscious Statistician. Monotonicity of the inverse bid function for bids strictly above the reservation price is proven in Appendix C.4.

[6] It should be noted that $\mathbf{k}$ is unique up to $|\tilde{min}(\mathbb{S})| + S^c$ elements of $\mathbf{k}$ that must be normalised to to the rank deficiency of the matrix Ψ. These elements are the entries associate with states $\mathbf{s} \in \tilde{min}(\mathbb{S})$ that are never observed as possible ex-post states, and one additional state from each component - associated with $\mathbf{s}_i = \mathbf{s}_i^1$. We will see in Appendix E.2.5 that these normalisations do not impact the identification of $\mathbf{j}$.

E.2.4 Non-Identified Elements

I now briefly discuss the elements of k that are not identified due to the binding reservation prices. There are three mutually exclusive sets of non-identified elements. First, those corresponding to outcome states $\mathbf{s}^a$ of various auctions that are never bid upon strictly above the reservation price, only at the reservation price. These elements are bounded above and below. If these elements were arbitrarily high then we could not rationalise the bidder only bidding the reservation price. If they were arbitrarily low then we could never rationalise them being bid on at all.

Second, there are elements corresponding to outcomes that are never bid on at all, but which there are occasions when one additional bid would make this outcome a possibility. These elements are bounded above. If they were arbitrarily high we could not rationalise the bidder never placing that one additional bid.

Third, there are the elements corresponding to $\mathbf{s}^a$ that are never bid on, and there are never occasions when, just by placing one additional bid, these outcomes become a possibility. These outcomes are neither bounded above, nor below. They are not bounded below because any arbitrarily low value would rationalise this outcome never being bid on. They are not bounded above because any arbitrarily high value can be rationalised by an arbitrarily low value corresponding to an 'adjacent' outcome from the second category, such that the risk of this arbitrarily low value is enough to put the bidder off from bidding on this arbitrarily high value.

Where available, these bounds can be found by rearranging the first order conditions for the set of bid and entry decisions such that $\boldsymbol{v} = \bar{\boldsymbol{v}}$. We do just as we did in Appendix E.2.2, albeit replacing $\boldsymbol{v} = \bar{\boldsymbol{v}}$, so that the only unknowns are the non-point identified elements of $\mathbf{k}$. Therefore, $\mathbf{k}$ is partially identified. Derivation of these bounds are omitted for brevity.

E.2.5 Identification of j and βV

Up to this point I have proven the non-parametric (partial) identification of both F_i and $K_i = J_i + \beta V_i$. I now demonstrate that the primitives from the dynamic game, namely J and V are non-parametrically point-identified up to β. That is, I am essentially extending the implications of Proposition 1 from JP to the multi-object, binding reservation price, case.

Proposition 13. *Under assumptions 1, 9, 3, 4, and 10, the objects $J_i(\mathbf{s})$, $V_i(\mathbf{s})$, and $F_i(|\mathbf{s})$ are non-parametrically (partially) identified up to $\beta = \bar{\beta}$.*

To prove Proposition 13 I only need to prove that, under these assumptions, J and V are separately identified conditional on both F and $J + \beta V$ being partially identified. I do so using a proof by construction. I show that V can be written as a function of only identified objects K and F. This is similar to JP's observation that in a single object context one can express the continuation value as a function of the distribution of bids alone.

First, consider maximised expected pay-off. Unlike the non-binding reservation prices case we cannot write down maximised expected utility as a function of bids and K only. In the presence of binding reservation prices the first order conditions break down when at least one bid is constrained. This means we then cannot substitute the inverse bid system into expected utility to find maximised expected utility. However, recognise that for any candidate $K_i = J_i + \beta V_i$ and any $\boldsymbol{v}_{it}$ we can still find maximised expected utility (Π^*) using numerical methods:

$$\Pi^*(\boldsymbol{v}_{it}, J_i + \beta V_i, \mathbf{s}_t) =$$
$$\max_{(\mathbf{b}_t, \mathbf{d}_t)} \left\{ \Gamma_i(\mathbf{b}_t, \mathbf{d}_t | \mathbf{s}_t)^T (\boldsymbol{v}_{it} - \mathbf{b}_t) + P_i(\mathbf{b}_t, \mathbf{d}_t | \mathbf{s}_t)^T [J_i + \beta V_i] \quad s.t. \;\; b_l \geq R_{lt} \;\; \forall l \right\}$$

Therefore, because the function $K_i(\mathbf{s}_t) = J_i(\mathbf{s}_t) + \beta V_i(\mathbf{s}_t)$ is partially identified then so is maximised expected utility, and so the empirical counterpart to the Value Function, given by $W_i(\boldsymbol{v}, \mathbf{s})$, is also at least partially identified for all $v_l \in [\underline{v}, \bar{v}]$. In fact, this function is point identified, rather than only partially identified, because elements of $J_i(\mathbf{s}_t) + \beta V_i(\mathbf{s}_t)$ that were never bid on are not identified, only bounded. We know that the bidder would never bid on these combinations, so they will never enter the maximised expected utility.[7]

[7] If there are values of $J_i(\mathbf{s}_t) + \beta V_i(\mathbf{s}_t)$ that were only ever bid on at the reservation price, then the value function actually will be only partially identified. However, this non-identified region will generally be very small. Likewise, elements of k corresponding to states which never appear as possible ex-post states will be trivially zeroed out in this equation, so it does not matter how they are normalised. Finally, the normalised elements corresponding to one (minimal, with $\mathbf{s}_i = \mathbf{s}_i^0$) element from each component $\mathbb{S}^c \subset \mathbb{S}$. These normalisations constitute location shifts of Π for all elements in that component, as we essentially made the normalisations because only marginal payoffs are identified. Finally, when we back out $\mathbf{j}$, we will normalise $j(\mathbf{s}_i^0) = 0$, in line with these location normalisations.

The ex-ante value function can be written as a function of F and K only:

$$V_i^E(\mathbf{s}_t) = \int_{\boldsymbol{v}_i} W_i(\boldsymbol{v}_{it}, \mathbf{s}_t) dF_i(\boldsymbol{v}_{it}|\mathbf{s}_t) \tag{E.13}$$

Importantly, as shown in Appendix C the ex-ante value function can be written analytically as a function of the distribution of bids and entry decisions $G_i(\mathbf{b}, \mathbf{d}|\mathbf{s}_t)$, and the identified objects K and F, extending proposition 3 to the binding reservation prices case. As above, even though F is only partially identified the ex-ante value function remains point identified. This is because the non-identified (truncated) region of F is the region in which the bidder is never observed bidding. Therefore, since they would never bid for v_l in this region their maximised expected utility will certainly just be the pay-off from not winning anything in this region.

We can then write the continuation value as follows:

$$V_{ia}(\mathbf{s}_t) = \int_{\mathbf{s}_{t+1}} V_i^E(\mathbf{s}) T_{\mathbf{s}}(\mathbf{s}|\mathbf{s}_t^a) \tag{E.14}$$

Therefore, given identification of $K_i(\mathbf{s})$ and $F(.|\mathbf{s})$ from the previous sections, we can evaluate this continuation value. Interestingly, however, V will be point identified even though F and K were not. The function J is then also non-parametrically identified as $J_i(\mathbf{s}) = K_i(\mathbf{s}) - \beta V_i(\mathbf{s})$. However, this will only be partially identified, given that K is only partially identified.

E.3 Endogenous Entry

In this Appendix I introduce endogenous entry in which entry is costly and v_{ilt} is not observed before entry, though I assume that the entry decisions of other players is observed before bidding.[8] I focus on the case with non-binding reservation prices, though it will be clear how the results from Appendix E.2 extend to this case.

The identification argument presents a minor generalisation on the one presented in the main text. The argument proceeds as follows: F is non-parametrically point

[8]Allowing the 'entry structure' to be unknown before bidding does not change anything substantive, since we simply have to alter the objects Γ_l P and Q to additionally take an expectation over the entry decisions of other players.

identified conditional on $\mathbf{k} = C\mathbf{j} + \beta\mathbf{V}$. As in the previous Appendix, $\mathbf{k}$ remains non-parametrically identified conditional on the identification of Γ and P using observed variation in $\mathbf{s}$, relying on our rank condition on the matrix Ψ. Given identification of $\mathbf{k}, \Gamma$, and P, Proposition 3 ensures that the expected payoff from each entry structure is also non-parametrically identified. Given these expected payoffs, the entry problem is then a multinomial discrete choice problem, so I rely on standard results for the identification of entry costs. Identification of expected entry payoffs and costs ensures the ex-ante value function, and hence the continuation value $\mathbf{V}$, is identified, thereby identifying $\mathbf{j} = C^{-1}(\mathbf{k} - \beta\mathbf{V})$.

I proceeds as follows: In Appendix E.3.1 I introduce changes to the main model, and demonstrate that the previous identification results for F and $\mathbf{k}$ also apply. In Appendix E.3.2 I show that the distribution of entry costs is non-parametrically identified, and finally that $\mathbf{V}$, and hence $\mathbf{j}$ are also identified.

E.3.1 Changes to the Model

All objects below should be treated as functions of the state $\mathbf{s}$. Conditional on an entry structure $\mathbb{D}$ and having observed the lot specific values $\boldsymbol{v}$ the agent places bids to maximise the following:

$$\Pi(\mathbf{b}|\boldsymbol{v};\mathbb{D}) = \Gamma(\mathbf{b}|\mathbb{D})^T(\boldsymbol{v} - \mathbf{b}) + P(\mathbf{b}|\mathbb{D})^T J + Q(\mathbf{b}|\mathbb{D})^T \beta V$$

Given the agent's behaviour conditional on entry, the agent's problem is to choose an entry structure $\mathbb{D}_i$ to maximise their expected pay-off. I assume that agent's entry costs, a $2^L \times 1$ vector $\mathbf{c}$, are drawn independently and privately from $C(.|\mathbf{s}_i)$ (independent of $\mathbf{s}_{-i}$). I assume that C is common knowledge.

The agent observes $\mathbf{s}$ and, given knowledge of F and $\mathbf{k}$ and their equilibrium beliefs about other players, forms and maximises an expected pay-off associated with any given entry structure:

$$W(\mathbb{D}_i|\mathbf{c}) = E_{\mathbb{D}_{-i}}[E_{\boldsymbol{v}}[\max_{\mathbf{b}} \{\Pi(\mathbf{b}|\boldsymbol{v};\mathbb{D})\}]|\mathbb{D}_i] - c_{\mathbb{D}_i}$$

The continuation value associated with ending the period in state $\mathbf{s}^a$ is then:

$$V(\mathbf{s}^a) = E_{\mathbf{s}}[E_{\mathbf{c}}[\max_{\mathbb{D}_i} \{W(\mathbb{D}_i|\mathbf{c})\} |\mathbf{s}]|\mathbf{s}^a]$$

Identification of F conditional on the identification of K

The Inverse Bid System, as given in equation 3.7, where the state variable has simply been augmented to include the entry structure. Hence F remains non-parametrically identified conditional on the identification of Γ, Q, and $\mathbf{k}$.

Identification of k

As in the main text, we can take a conditional expectation of the inverse bid system, setting this equal to zero: $E[\boldsymbol{\xi}|\mathbf{s}, \mathbb{D}] = 0$. We can then again stack this system of equations across states and entry structures for $0 = \Upsilon - \Psi\mathbf{k}$. Non-parametric point identification of $\mathbf{k}$ then requires the same rank condition on Ψ proven previously.[9]

Identification of $E_v[\tilde{\Pi}(\mathbf{b}^*|v; \mathbf{s}, \mathbb{D})]$

Recognise that Proposition 3 continues to hold, and so we can write the expected maximised payoff, conditional on $\mathbb{D}$, as

$$\bar{\Pi}(\mathbf{s}, \mathbb{D}) = E_v[\tilde{\Pi}(\mathbf{b}^*|v; \mathbb{D})] = \begin{array}{c} \Gamma(\mathbf{b}^*|\mathbb{D})^T \nabla_{\mathbf{b}} \Gamma(\mathbf{b}^*|\mathbb{D})^{-1} \Gamma(\mathbf{b}^*|\mathbb{D}) \\ +[Q(\mathbf{b}^*|\mathbb{D})^T \ - \ \Gamma(\mathbf{b}^*|\mathbb{D})^T \nabla_{\mathbf{b}} \Gamma(\mathbf{b}^*|\mathbb{D})^{-1} \nabla_{\mathbf{b}} Q(\mathbf{b}^*|\mathbb{D})] A_{\mathbf{s}} \beta \mathbf{k} \end{array}$$

E.3.2 Identification of C

At the entry stage, the agent sets their entry structure $\mathbb{D}_i$ such that:

$$E_{\mathbb{D}_{-i}}[E_v[\underset{\mathbf{b}}{max} \ \{\Pi(\mathbf{b}|v; \mathbb{D})\}]|\mathbb{D}_i] - c_{\mathbb{D}_i} \geq \underset{\mathbb{D}_i' \neq \mathbb{D}_i}{max} \left\{ E_{\mathbb{D}_{-i}}[E_v[\underset{\mathbf{b}}{max} \ \{\Pi(\mathbf{b}|v; \mathbb{D})\}]|\mathbb{D}_i'] - c_{\mathbb{D}_i'}) \right\}$$

Similar to how we identify G, because we observe entry decisions, we therefore observe the equilibrium distribution of $\mathbb{D}_i$ for all i. Therefore, following from the

[9] We must normalise elements of $\mathbf{k}$ that correspond to states which are either the minimal element of their component (connected subset of $\mathbb{S}$, connected by our partial ordering), or never appear as possible ex-post states. By definition, there will be $S^c + |min(\mathbb{S})|$ of these. In Appendix D.1 we found previously that Ψ has rank $S - S^c - |min(\mathbb{S})|$.

above, $E_{\mathbb{D}_{-i}}[E_{\boldsymbol{v}}[\max_{\mathbf{b}} \{\Pi(\mathbf{b}|\boldsymbol{v};\mathbb{D})\}]|\mathbb{D}_i]$ is non-parametrically point identified. Normalising that the entry cost of entering zero auctions is always zero, I now exploit the exclusion restriction that the distribution of $\mathbf{c}$ is independent of $\mathbf{s}_{-i}$. Therefore, variation in $\mathbf{s}_{-i}$ leads to known variation in $E_{\mathbb{D}_{-i}}[E_{\boldsymbol{v}}[\max_{\mathbf{b}} \{\Pi(\mathbf{b}|\boldsymbol{v};\mathbb{D})\}]|\mathbb{D}_i]$, thereby tracing out the distribution $C(.|\mathbf{s}_i)$, ensuring we have non-parametric identification.[10]

Following on from this, we then get that the ex-ante continuation value $V^e(\mathbf{s}) = E[\max_{\mathbb{D}_i} \left\{ E_{\mathbb{D}_{-i}}[E_{\boldsymbol{v}}[\max_{\mathbf{b}} \{\Pi(\mathbf{b}|\boldsymbol{v};\mathbb{D})\}]|\mathbb{D}_i] - c_{\mathbb{D}_i}) \right\}]$, and hence the continuation value $V(\mathbf{s})$ are non-parametrically identified, which in turn yields identification of the flow payoff function j.

E.4 Inter-temporal Budget Constraint

I now consider a setting where bidders face an inter-temporal budget constraint, rather than assuming utility is quasi-linear in wealth. I present two key results.

In subsection E.4.3 I show that in a stationary environment, the marginal utility of wealth is constant over time. Stationarity is needed to ensure that the distribution of bidding behaviour is independent of the time period. This is intuitive — in a stationary environment, the benefit of an extra unit of wealth is that it increases the lifetime budget. The extra wealth is smoothed over time. Conversely, the opportunity cost of spending wealth today is that it reduces the lifetime budget and less wealth is available for spending in future. Therefore, in the stationary environment, the opportunity cost of spending wealth should be the same in each period.

In subsection E.4.4 I show that when the marginal utility of wealth is constant over time, the model is observationally equivalent to a model with payoffs quasi-linear in wealth. This is a useful result as it ensures even when bidders are bidding subject to an intertemporal budget constraint, the econometrician need not worry themselves with this additional wrinkle, and can instead use a quasi-linear model.

I focus on the case without binding reservation prices, and zero probability ties, for simplicity. It is clear how the model carries over to these other cases. I also do not consider the problem of characterising nor proving existence of equilibrium in

[10]Technically, identification is only partial, since the set of states is finite, so we will only actually be point identifying $C(.|\mathbf{s}_i)$ at a finite set of points across its support. We can achieve full point identification either by assuming discrete support, or introducing one continuously varying element of $\mathbf{s}_{-i}$.

this adjusted game. I continue to rely on the assumption that equilibrium payoffs are continuous in values.

E.4.1 Budget Constraint

I will assume that bidders face a no-ponzi style budget constraint, such that their initial savings at $t+1$, A_{t+1} equal their savings at t, A_t plus their income y_t, less their realised expenditure in that period e_t. I assume bidders can save and borrow at interest rate r_t. This constraint is then given by:

$$\frac{A_{t+1}}{1+r_t} = A_t + y_t - e_t \tag{E.15}$$

For ease of notation I will treat r as constant and make the assumption that the equilibrium real interest rate satisfies $(1+r)\beta = 1$. This is a standard equilibrium result from undergraduate macro models. I also make the assumption that y is independent of $\mathbf{s}$, however this is not strictly necessary.

E.4.2 The Bellman Equation

Writing down the Bellman equation in standard terms prohibits the use of matrix notation. For the sake of notation, define the set of possible allocations as $\mathbb{W}(\mathbf{s})$. The Bellman equation for this problem is given as:

$$W(\boldsymbol{v}, \mathbf{s}, A) = \max_{\mathbf{b}} \{\Pi(\mathbf{b}; \boldsymbol{v}, \mathbf{s}, A)\}$$

Where

$$\Pi(\mathbf{b}; \boldsymbol{v}, \mathbf{s}, A) = \sum_{l \in \mathbb{L}(\mathbf{s})} \Gamma_l(\mathbf{b}|\mathbf{s}) v_l + \sum_{\mathbf{w}^a \in \mathbb{W}(\mathbf{s})} Q_a(\mathbf{b}|\mathbf{s}) j(\mathbf{s}^a)$$
$$+ \beta \sum_{\mathbf{w}^a \in \mathbb{W}(\mathbf{s})} Q_a(\mathbf{b}|\mathbf{s}) \int_{\tilde{\mathbf{s}}, \tilde{y}} \int_{\tilde{\boldsymbol{v}}} W(\tilde{\boldsymbol{v}}, \tilde{\mathbf{s}}, (1+r)[A - \sum_{l \in \mathbb{L}(\mathbf{s})} \mathbb{I}[w_l^a = i] b_l] + \tilde{y})) dF(\tilde{\boldsymbol{v}}|\tilde{\mathbf{s}}) dF(\tilde{\mathbf{s}}, \tilde{y}|\mathbf{s}^a)] \tag{E.16}$$

From now on I use the integrated value function $\bar{W}(\mathbf{s}, A) = \int_{\boldsymbol{v}} W(\boldsymbol{v}, \mathbf{s}, A) dF(\boldsymbol{v}|\mathbf{s})$. Existence of a unique solution to this equation holds from the standard results of

Bhattacharya and Majumdar (1989), under the assumptions that equilibrium expected flow payoffs are continuous in both actions and values, both of which come from compact bounded sets. Even though the state space $(\mathbf{s}, A) \in \mathbb{S} \times \mathbb{R}$ is no longer finite, this is not a problem so long as equilibrium expected payoffs remain continuous in types (which are a function of the states).

E.4.3 Stationarity and Constant Marginal Utility of Wealth

Proposition 14. *The unique solution to the recursive equation for the integrated value function can be written as $\bar{W}(\mathbf{s}, A) = \tilde{W}(\mathbf{s}) + \lambda A$ for some constant $\lambda > 0$, and for some recursive function $\tilde{W}$.*

Proving this proposition is a case of guess and verify that $\bar{W}(\mathbf{s}, A) = \tilde{W}(\mathbf{s}) + \lambda A$ is the solution, before showing that there exists a solution for $\tilde{W}$ and λ. Importantly, from Theorem 3.2 of Bhattacharya and Majumdar (1989), we know that the solution to the infinite horizon dynamic programming problem is unique.

The recursive integrated value function can be written as:

$$\bar{W}(\mathbf{s}, A) =$$
$$\int_{v} \max_{\mathbf{b}} \left\{ \begin{array}{c} \sum_{l \in \mathbb{L}(\mathbf{s})} \Gamma_l(\mathbf{b}|\mathbf{s}) v_l \\ + \sum_{\mathbf{w}^a \in \mathbb{W}(\mathbf{s})} Q_a(\mathbf{b}|\mathbf{s}) j(\mathbf{s}^a) + \beta \int_{\tilde{\mathbf{s}}} \tilde{W}(\tilde{\mathbf{s}}) dF(\tilde{\mathbf{s}}|\mathbf{s}^a) \\ + \sum_{\mathbf{w}^a \in \mathbb{W}(\mathbf{s})} Q_a(\mathbf{b}|\mathbf{s}) \beta \lambda ((1+r)[A - \sum_{l \in \mathbb{L}(\mathbf{s})} \mathbb{I}[w_l^a = i] b_l] + E_{\tilde{y}}[\tilde{y}]))] \end{array} \right\} dF(v|\mathbf{s})$$

Recognising that, for a constant M $\sum_{\mathbf{w}^a \in \mathbb{W}(\mathbf{s})} Q_a(\mathbf{b}|\mathbf{s}) M = M$, that

$\sum_{\mathbf{w}^a \in \mathbb{W}(\mathbf{s})} Q_a(\mathbf{b}|\mathbf{s}) \sum_{l \in \mathbb{L}(\mathbf{s})} \mathbb{I}[w_l^a = i] b_l = \sum_{l \in \mathbb{L}(\mathbf{s})} \Gamma_l(\mathbf{b}|\mathbf{s}) b_l$, and imposing $(1 + r)\beta = 1$, we can rewrite this equation as:

$$\bar{W}(\mathbf{s}, A) = \int_{v} \max_{\mathbf{b}} \left\{ \begin{array}{c} \sum_{l \in \mathbb{L}(\mathbf{s})} \Gamma_l(\mathbf{b}|\mathbf{s})(v_l - \lambda b_l) + \lambda E_{\tilde{y}}[\tilde{y}] \\ + \sum_{\mathbf{w}^a \in \mathbb{W}(\mathbf{s})} Q_a(\mathbf{b}|\mathbf{s}) j(\mathbf{s}^a) + \beta \int_{\tilde{\mathbf{s}}} \tilde{W}(\tilde{\mathbf{s}}) dF(\tilde{\mathbf{s}}|\mathbf{s}^a) \end{array} \right\} dF(v|\mathbf{s}) + \lambda A$$

And so, we have shown that $\bar{W}(\mathbf{s}, A) = \tilde{W}(\mathbf{s}) + \lambda A$ for $\tilde{W}(\mathbf{s})$ such that:

$$\tilde{W}(\mathbf{s}) = \int_{v} \max_{\mathbf{b}} \left\{ \begin{array}{c} \sum_{l \in \mathbb{L}(\mathbf{s})} \Gamma_l(\mathbf{b}|\mathbf{s})(v_l - \lambda b_l) + \lambda E_{\tilde{y}}[\tilde{y}] \\ + \sum_{\mathbf{w}^a \in \mathbb{W}(\mathbf{s})} Q_a(\mathbf{b}|\mathbf{s}) j(\mathbf{s}^a) + \beta \int_{\tilde{\mathbf{s}}} \tilde{W}(\tilde{\mathbf{s}}) dF(\tilde{\mathbf{s}}|\mathbf{s}^a) \end{array} \right\} dF(v|\mathbf{s})$$

Existence and uniqueness of this function is the same as existence of the integrated value function needed in the main text, and under the same requirements as above follows from Bhattacharya and Majumdar (1989).

E.4.4 Observational Equivalence with the Quasi-Linear Model

I now demonstrate that the intertemporally budget constrained model with constant marginal utility of wealth is observationally equivalent to the quasi-linear model. Given the similarity between the above value function, and the one presented in the main text for the quasi-linear model, this result should be unsurprising.

Consider the first order conditions for optimising behaviour associated with the above value function:

$$0 = \nabla_{b_l}\Gamma_l(\mathbf{b}^*|\mathbf{s})(v_l - \lambda b_l^*) - \lambda\Gamma_l(\mathbf{b}^*|\mathbf{s}) + \sum_{\mathbf{w}^a \in \mathbb{W}(\mathbf{s})} \nabla_{b_l} Q_a(\mathbf{b}|\mathbf{s}) j(\mathbf{s}^a) + \beta \int_{\tilde{\mathbf{s}}} \tilde{W}(\tilde{\mathbf{s}}) dF(\tilde{\mathbf{s}}|\mathbf{s}^a)$$

Which are identical to the first order conditions presented in equation 3.6, except for the λ term, which cannot be separately identified from the scale of v nor j, and so must be normalised to a constant scalar.

E.5 Stochastic Combination Value

In this Appendix I present two identification results for the case when the combination value is stochastic. I focus on the case when the object $j(\mathbf{s})$ is not a function but a probability distribution. I focus on the static setting. This is for two reasons. First, these cases with stochastic combination values are novel even in the static case. Second, as we have seen throughout this Appendix, identification of the primitives of a generalised static model (where primitives are allowed to depend on $\mathbf{s}_0$ etc), is sufficient for identification of the primitives of a dynamic model. This is because identification of the Pseudo-Static payoff function k implies identification of j.

I focus on the case with non-binding reservation prices and exogenous entry, since I established previously that these factors are only mathematically, not substantively, problematic. I focus on two cases: First, when objects are homogenous, so that J only depends on the number of objects won. Second, when J is a function of low-dimensional un-observables, such as stocks.[11] In both cases I consider several simple extensions that may be of interest to applied researchers in future.

These extensions both centre on the theme of finding some way to reduce the dimensionality of the unknowns. The key idea is this: Each observation of bidding on an auction yields L pieces of information. Therefore, in order to have any hope at point identifying unobservables, there cannot be more than L unobservables.

E.5.1 Case 1. Homogenous lots

First, I consider the case when J only depends on the number of lots consumed. Lots are allowed to differ (observably), but these differences cannot interact with the combinatorial pay-off. In this way, the pay-off associated with winning a combination of goods has two components: A component that depends on which lots are consumed, in an additively separable manner, and a component that depends just on how many lots are consumed. Expected utility is then given by:

$$EU(\mathbf{b}) = \Gamma(\mathbf{b})^T(\underbrace{\boldsymbol{\nu}}_{L\times 1} - \mathbf{b}) + \underbrace{P(\mathbf{b})^T}_{L+1\times 1}\mathbf{J}_t \tag{E.17}$$

Where ν give the deterministic lot-specific pay-off, and J_t gives the stochastic number of lots pay-off. Entry n gives the pay-off from winning $n-1$ lots.

We can rewrite expected utility as: $\Gamma(\mathbf{b})^T(\boldsymbol{\nu} - \mathbf{b}) + P(\mathbf{b})^T(\mathbf{J}_t - J_{1t}) + J_{1t}$ Which follows from $\sum_a P_a(\mathbf{b}) = 1$. Recognise that row 1 of $(\mathbf{J}_t - J_{1t})$ is equal to zero, and so we can focus on just rows 2 through to $L+1$. For ease of notation, define $\tilde{P} = P_{2:L+1}$ and $\tilde{\mathbf{K}} = \mathbf{K}_{2:L+1}$. Necessary first order conditions are given by:

$$0 = \nabla_{\mathbf{b}}\Gamma(\mathbf{b})(\boldsymbol{\nu} - \mathbf{b}) - \Gamma(\mathbf{b}) + \nabla_{\mathbf{b}}\tilde{P}(\mathbf{b})(\tilde{\mathbf{J}}_t - J_{1t})$$

[11] Importantly, in this case, for the dynamic setting, we will require that the continuation value, and hence k, does not depend on the unobservables of the other players. So, if we are thinking about an unobserved stock model, we will require that the continuation value for player i does not depend on the unobserved stocks of player j. This will be the case if, for example, the distribution of equilibrium winning bids is independent of player's stocks.

Importantly, $\nabla_{\mathbf{b}}\tilde{P}(\mathbf{b})$ is an $L \times L$ matrix with rank L.[12] Therefore, just as in the additively separable case we can invert these conditions for:

$$\tilde{\mathbf{J}}_t = J_{1t} + \nabla_{\mathbf{b}}\tilde{P}(\mathbf{b})^{-1}[\nabla_{\mathbf{b}}\Gamma(\mathbf{b})(\mathbf{b} - \boldsymbol{\nu}) + \Gamma(\mathbf{b})]$$

Therefore, conditional on J_{1t} and $\boldsymbol{\kappa}$, the joint distribution of $\tilde{\mathbf{J}}_t$ is non-parametrically point identified. Two important points are worth highlighting here: the fact that J_{1t} is not identified is not a problem. This is the same way that the pay-off from losing in the single-unit case is not identified. All we can identify is marginal pay-offs. Therefore, we will generally normalise the pay-off from losing to zero. Finally, $\boldsymbol{\nu}$ will be identified from variation in lot specific observables.

E.5.2 Case 2: Known function of low dimensional un-observables

Suppose the combinatorial value can be written as $\mathbf{J}(\mathbf{m}_t)$ where $\mathbf{m}_t \in \mathbb{M}$ is an unobserved (potentially) stochastic random variable of dimension $M \leq L$. I require that $\mathbf{J} : \mathbb{M} \to \mathbb{J}$ is a known function (with range $\mathbb{J} \subset \mathbb{R}^{2^L}$). Importantly, some elements of $\mathbf{m}$ may represent fixed parameters associated with the functional form J, in this respect, this is essentially a parametric identification argument.

Normalise the first element of this vector valued function (corresponding to player i losing every lot) to zero, so that I focus on the marginal combinatorial pay-off $\mathbf{J}(\mathbf{m})_{2:2^L} - \mathbf{J}(\mathbf{m})_1$. Expected utility is then $\Pi(\mathbf{b}) = P(\mathbf{b})^T\mathbf{J}(\mathbf{m}_t) - \Gamma(\mathbf{b})^T\mathbf{b}$. Necessary first order conditions are then given by:

$$0 = \nabla_{\mathbf{b}}P(\mathbf{b})\mathbf{J}(\mathbf{m}_t) - \nabla_{\mathbf{b}}\Gamma(\mathbf{b})\mathbf{b} - \Gamma(\mathbf{b})$$

Given that $\mathbf{m}$ is the only non-identified object in this equation, the question is then whether this is point identified. I make two assumptions about this function that are sufficient for $\mathbf{m}_t$ to be point identified:

Assumption 11.

1. $\mathbf{J}(\mathbf{m})$ is continuous and continuously differentiable for all $\mathbf{m}_t$.

[12]Rank of this matrix follows from the fact that the matrix $\nabla_{\mathbf{b}}P(\mathbf{b})$ presented in section 3.3 has rank L, proven in Appendix D.1. The matrix $\tilde{P}$ is just a collapsed version of this, summed over the different lots.

2. For any $\mathbf{m}$ and $\mathbf{m}'$ there exists a set $\mathbb{U} \subset \{1, 2, \ldots, 2^L\}$ with $|\mathbb{U}| = m$ that defines the vector value function $\mathbf{F}^{\mathbb{U}}$ where $F_n^{\mathbb{U}}(\mathbf{m}) = J_{U_n}(\mathbf{m})$ such that

$$(\mathbf{m} - \mathbf{m}')^T(\mathbf{F}^{\mathbb{U}}(\mathbf{m}) - \mathbf{F}^{\mathbb{U}}(\mathbf{m}')) > 0$$

The second part of this assumption is essentially an extension of strict monotonicity to the case of 2^L dimensional functions in M dimensional variables. The assumption states that for any two distinct $\mathbf{m}$ and $\mathbf{m}'$ we can find a set of rows of $\mathbf{J}(.)$ such that this inner product is strictly positive.[13] A key result of this property is that the function $\mathbf{J}(.)$ is a bijection: Each $\mathbf{m}$ maps onto a unique $\mathbf{J}$, and the condition ensures that for any two distinct $\mathbf{m}$ and $\mathbf{m}'$ it must be the case that $\mathbf{J}(\mathbf{m}) \neq \mathbf{J}(\mathbf{m}')$ (since otherwise we could not find a $\mathbb{U}$ such that $(\mathbf{m} - \mathbf{m}')^T(\mathbf{F}^{\mathbb{U}}(\mathbf{m}) - \mathbf{F}^{\mathbb{U}}(\mathbf{m}')) > 0$). This ensures that the inverse $\mathbf{J}^{-1}(.)$ exists, such that for all $\mathbf{m} \in \mathbb{M}$ $\mathbf{m} = \mathbf{J}^{-1}(\mathbf{J}(\mathbf{m}))$. Furthermore, because $\mathbf{J}(.)$ is continuous and continuously differentiable everywhere, so that $\mathbf{J}^{-1}(.)$ must be differentiable everywhere, $\mathbf{J}^{-1}(.)$ must also be continuous.

Proposition 15. *Under assumptions 1, 9, and 11, $\mathbf{m}_t$ is point identified up to normalisation.*

We may only identify $\mathbf{m}_t$ up to location and scale if, for example, the second and third elements of $\mathbf{m}_t$ are constant parameters describing the mean and standard deviation of m_{1t}.

The proof of Proposition 15 consists of arguing that we have L equations in only M unknowns, and that there exists a unique solution to the system of equations. The proof proceeds by recognising that the set of vectors $\mathbf{J}$ which satisfy the first order conditions must be convex. Which, from the continuity of the inverse function $\mathbf{J}^{-1}(.)$ and the (generalised) intermediate value theorem, implies that the set of $\mathbf{m}$ for which the first order conditions hold must be path connected. This implies there must be a point arbitrarily close to $\mathbf{m}_t$ for which the first order conditions hold. However, the fact that $\nabla_{\mathbf{b}} P(\mathbf{b})$ has rank L and the function $\mathbf{J}(.)$ is invertible implies that the system of equations must be locally unique.

Proof: 1. Consider the set of $2^L \times 1$ dimensional vectors which satisfy the system of equations $\nabla_{\mathbf{b}} P(\mathbf{b})\mathbf{K} - \nabla_{\mathbf{b}} \Gamma(\mathbf{b})\mathbf{b} - \Gamma(\mathbf{b}) = 0$. This set, which I will

[13] This property is satisfied when, for example, each element of J is weakly monotone in elements of $\mathbf{m}$, and strictly monotonic in at least one element.

refer to as $\tilde{\mathbb{J}}$, must be convex, and hence path-connected, as for any two vectors in this set $\mathbf{J}, \mathbf{J}'$ we have:

$$\begin{aligned}
\lambda \nabla_{\mathbf{b}} P(\mathbf{b})\mathbf{J} = \lambda(\nabla_{\mathbf{b}}\Gamma(\mathbf{b})\mathbf{b} + \Gamma(\mathbf{b})) \\
\& \qquad (1-\lambda)\nabla_{\mathbf{b}} P(\mathbf{b})\mathbf{J}' &= (1-\lambda)(\nabla_{\mathbf{b}}\Gamma(\mathbf{b})\mathbf{b} + \Gamma(\mathbf{b})) \\
\therefore\ \lambda \nabla_{\mathbf{b}} P(\mathbf{b})\mathbf{J} + (1-\lambda)\nabla_{\mathbf{b}} P(\mathbf{b})\mathbf{J}' &= \nabla_{\mathbf{b}}\Gamma(\mathbf{b})\mathbf{b} + \Gamma(\mathbf{b}) \\
\nabla_{\mathbf{b}} P(\mathbf{b})(\lambda\mathbf{J} + (1-\lambda)\mathbf{J}') &= \nabla_{\mathbf{b}}\Gamma(\mathbf{b})\mathbf{b} + \Gamma(\mathbf{b}) \qquad \text{(E.18)}
\end{aligned}$$

2. This implies that the image of the intersection of $\tilde{\mathbb{J}}$ and $\mathbb{J}$ defined by the continuous function $\mathbf{J}^{-1}(.)$, that is the set of $\mathbf{m}$ for which the first order conditions hold, must also be path connected. This result follows from the generalised intermediate value theorem, which states that for a continuous function $f : \mathbb{X} \to \mathbb{Y}$, if the set $\mathbb{X}$ is path-connected, then so is the image $f(\mathbb{X})$.
3. Therefore, if the intersection of $\tilde{\mathbb{J}}$ and $\mathbb{J}$ contains more than a single element, then for any $\mathbf{m}$ which satisfies the first order conditions, there must exist an arbitrarily nearby $\mathbf{m}'$ which also satisfies the first order conditions.
4. However, from the inverse function theorem, the first order conditions are locally unique. The Jacobian of the first order conditions, with respect to $\mathbf{m}$ are given by:
$$\nabla_{\mathbf{b}} P(\mathbf{b}) \nabla_{\mathbf{m}} \mathbf{J}(\mathbf{m})$$
which has rank M. This can be seen because $\nabla_{\mathbf{b}} P(\mathbf{b})$ has rank L, consisting of L linearly independent rows. Meanwhile, $\nabla_{\mathbf{m}} \mathbf{J}(\mathbf{m})$ has rank M, which arises because $\mathbf{J}(\mathbf{m})$ is invertible.
5. Therefore the set of $\mathbf{m}$ which satisfy the first order conditions must contain only a single element. Likewise, the intersection of $\tilde{\mathbb{J}}$ and $\mathbb{J}$ must also contain only a single element

□

Importantly, once $\mathbf{m}_t$ is point identified, we can identify its joint distribution over time. Information about previous winnings then allow us to pin down the location and scale of $\mathbf{m}_t$. For example, if $\mathbf{m}_t$ represents the bidder's stock of objects, which is assumed to follow an VAR(1) process along the lines of $\mathbf{m}_t = A\mathbf{m}_{t-1} + \mathbf{z}_{t-1} + \varepsilon_{t-1}$ then observation of winnings $\mathbf{z}_t$ immediately pin down the location and scale, while the distribution of ε_t and the auto-regressive matrix A can then be identified by running a simple auxiliary regression (assuming, for example, that ε_t is weakly exogenous with respect to $\mathbf{z}_t$).

Case 2: When $\mathbf{M} > \mathbf{L}$

We can also easily extend this case to allow $M > L$, essentially combining information across observations, rather than trying to identify everything from a single observation, so long as *enough* elements of M are constant across observations. This case is important when $\mathbf{m}_t$ can be decomposed into $(\mathbf{m}_t^1, \mathbf{m}^0)$, where $\mathbf{m}^0$ are fixed parameters that do not vary across observations. Suppose $M \leq 2L$, and in particular, $|\mathbf{m}_t^1| < L$. In this case, rather than considering a single set of first order conditions, we can look at a pair of first order conditions from two separate periods t_1 and t_2. Importantly, we still impose assumption 11. Combine the two sets of first order conditions as follows:

$$\begin{pmatrix} \nabla_{\mathbf{b}} P(\mathbf{b}_{t_1}) & 0 \\ 0 & \nabla_{\mathbf{b}} P(\mathbf{b}_{t_2}) \end{pmatrix} \begin{pmatrix} \mathbf{J}(\mathbf{m}_{t_1}) \\ \mathbf{J}(\mathbf{m}_{t_2}) \end{pmatrix} = \begin{pmatrix} \nabla_{\mathbf{b}} \Gamma(\mathbf{b}_{t_1})\mathbf{b}_{t_1} + \Gamma(\mathbf{b}_{t_1}) \\ \nabla_{\mathbf{b}} \Gamma(\mathbf{b}_{t_2})\mathbf{b}_{t_2} + \Gamma(\mathbf{b}_{t_2}) \end{pmatrix}$$

The fact there exists a unique pair $(\mathbf{m}_{t_1}, \mathbf{m}_{t_2})$ that solves this system of equations follows the same logic as the previous proof with the added note that $\nabla_{(\mathbf{m}_{t_1}, \mathbf{m}_{t_2})} \begin{pmatrix} \mathbf{J}(\mathbf{m}_{t_1}) \\ \mathbf{J}(\mathbf{m}_{t_2}) \end{pmatrix}$ has rank $2|\mathbf{m}_t^1| + |\mathbf{m}^0|$, ensuring that we are again able to appeal to the inverse function theorem for local uniqueness.

This result is important as it allows us to add a large number of additional parameters to the function $\mathbf{J}(.)$ which are identified by essentially using variation across observations. This makes use of a similar philosophy used to prove the identification results in the main sections of this chapter.

Appendix F

Stationarity

In this appendix I present evidence that the equilibrium stock process is stationary - that the distribution of stocks remains constant over time. I focus on two types of stationarity: First, whether stocks trend over time. Second whether stocks follow a random walk.

In Appendix F.1 I present suggestive evidence that stocks neither trend upwards nor downwards over time, by testing for structural breaks in bidding behaviour. In Appendix F.2 I discuss how an additional assumption about how observed winnings reacts to changes in the stock allows me to test whether the equilibrium stock process follows a random walk. In Appendix F.4 I discuss how the results of this analysis gives us information about food banks' unobserved stock process, giving us natural priors for $\boldsymbol{\mu}_i$ and Σ_i.

F.1 Trend Stationarity

If stocks trend over time we expect that bidding behaviour should follow a similar pattern. Therefore we can investigate the existence of trends by looking for evidence of trends in bidding behaviour.

If stocks have a linear trend it is ex-ante unclear whether average bids will also have a linear trend. To allow for the possibility of non-linear trends in bidding behaviour I focus on testing for the existence of more general structural breaks in behaviour. I focus on average monthly bids by food bank and food type, using the estimated α_{igt} parameters estimated from the tobit specification in Section 2.5.3. If these parameters do exhibit a linear trend a generalised test for structural breaks should pick this up. I omit the Fresh storage type due to the structural break caused on day 553 of my sample when fresh food ceased to be allocated on the Choice System. Even if the Fresh stock process does not exhibit any structural break (as remains my hypothesis) I cannot estimate average bids after this break, given that no fresh food was allocated. In Appendix K.2 I give additional details of how I model this structural break.

I test for a structural break in the series $\{\alpha_{igt}\}_{t\in\{1...T\}}$. For each t I split the sample into a before and after group, then run a t-test on the equality of means. This is performed separately for each food bank $\times$ storage type combination, and I allow the variances to differ in the before and after periods. I then plot the distribution of estimated t-statistics. Under the null hypothesis of no structural breaks, these statistics are t-distributed with 40 degrees of freedom.[1] Therefore I can compare the resulting distribution of t-statistics from their distribution under the null hypothesis.[2] In figure F.1 I plot the distribution of estimated t-statistics along side the distribution of these statistics under the null-hypothesis. Under the null we expect that 5% of test statistics will be above the critical values for a two-tailed test at 5% significance level ($= 2.02$). I find that 5.4% of test-statistics exceed the critical values. This gives some evidence that bidding behaviour does not exhibit trends or structural breaks, and so neither do stocks.

F.2 Cointegration

I now discuss how I can test whether the equilibrium stock process follows a random walk. I must maintain the assumption that stocks do not contain any sort of time trend, such as a linear trend. Fortunately the results presented in Appendix F.1 gives us evidence that equilibrium stocks are unlikely to exhibit time trends.[3]

[1] Although I have 44 months of data, I do not have all the data on the first and last month. So I only estimate 42 monthly means. I also adjust the t-statistics to account for sampling variation in the estimated α_{igt} using the law of total variance. Because I only have a finite number of months I assume normality of average monthly bids, resulting in this t-distribution.

[2] Using a Kolmogorv-Smirnov test I can reject the null hypothesis, at a 5% significance level, that the t-statistics come from this distribution. However, this is due to a lack of fit around the mode of the distribution, whereas the tails of the distribution fit much better.

[3] A linear trend is not identified from winnings data. Winnings are a measure of changes in the stock, so any linear trend is captured in the constant term. But, the constant term of average winnings also captures average net local donations μ_i. I am unable to disentangle these two objects using winnings data alone.

FIGURE F.1: Distribution of t-Test statistics

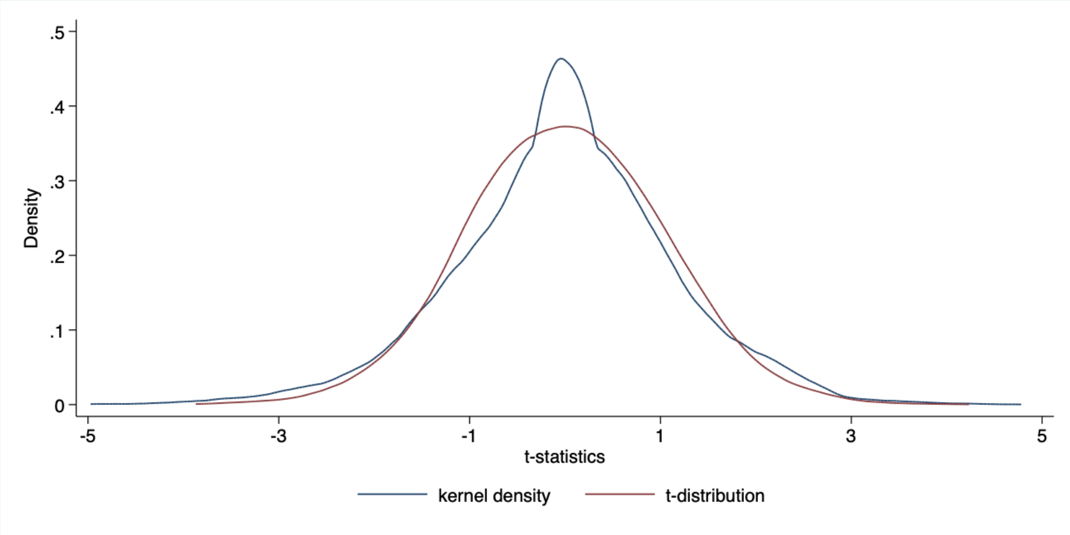

Note: This plot shows the distribution of t-test statistics, looking for structural breaks in average monthly bids. I focus on monthly bids between the 2nd and 43rd month. I also plot the standardised t-distribution with 40 degrees of freedom. I test for the presence breaks at each month between the 7th and 37th month. Tests are performed at the food bank × storage type level. Test-statistics are adjusted to account for sampling uncertainty in the estimates of average monthly bids.

In the main text I focus on the morning state, before any auctions take place. However for the purposes of this appendix it is most convenient to focus on the evening state, $\mathbf{s}^e$, after the final auction has taken place. The evening stock transition process is therefore given by:

$$\mathbf{s}^e_{it} = \mathbf{s}^e_{it-1} + \mathbf{w}^T_{it}\mathbf{z}^g_t + \mathbf{x}_{it}$$

The only difference is that the superscript on winnings is not lagged, as it is for the morning process. Evidently $\mathbf{w}^T_{it}\mathbf{z}^g_t$ depends on net daily donations $\mathbf{x}_{it}$, which can be considered short term changes in the stock. It will also depend on the previous stock $\mathbf{s}^e_{it-1}$. This is where the similarity to a cointegration framework arises. Food banks likely have an ideal level of stock they would maintain - they do not want the warehouse to be too empty, nor too full. Therefore, as well as reacting to short term changes in net daily donations, winnings should also respond to how far off optimal the previous stock is.

If I observed stocks as well as winnings I could easily test this relationship following the procedure of Engle and Granger, 1987. Instead, I make the additional assumption that the equilibrium stock process is given by:

$$\mathbf{s}^e_{it} = \delta \mathbf{s}^e_{it-1} + \alpha \mathbf{x}_{it} + \boldsymbol{\varepsilon}_{it}$$

This assumption states, on average, the evening stock ends up as some fraction of the previous evening stock, plus some fraction of net local donations. These are the fractions that equilibrium winnings could not offset. I do not impose that this process is stationary. For example, it is possible that $\delta = I$, so that the process follows a random walk with drift. I discuss $\boldsymbol{\varepsilon}_{it}$ in detail shortly. I must assume there is not a linear trend in this process. For simplicity I focus on the case in which δ and α are diagonal matrices, essentially focusing on one component of stocks at a time.

This assumption on the equilibrium stock process is really an assumption on the equilibrium winnings process. Equating the two previous equations and rearranging yields something similar to a standard error correction process:

$$\mathbf{w}^T_{it} \mathbf{z}^g_t = (\alpha - I)\mathbf{x}_{it} + (\delta - I)\mathbf{s}^e_{it-1} + \boldsymbol{\varepsilon}_{it} \tag{F.1}$$

This states that, on average, winnings offset some fraction of that day's net donations, as well as some fraction of the previous stocks. The residual random variable $\boldsymbol{\varepsilon}$ captures idiosyncrasies that affect the food bank's winnings. For example, how much is actually actioned that period, how many rival active bidders there are, and other attributes of the lots. It likely exhibits correlation over time, is non-normally distributed, lumpy, and may have non-zero mean. Importantly, this variable is assumed independent of $\mathbf{x}_{it}$.

If $\delta = 0$ and $\alpha = 0$ then winnings perfectly offsets changes in the stock and daily donations. This means that equilibrium stocks only vary with $\boldsymbol{\varepsilon}_{it}$. Instead, if $\delta = I$, so that the stock process followed a random walk, this equation would state that winnings do not depend on the previous stock. This asserts that winnings only react to $\mathbf{x}_{it}$ and $\boldsymbol{\varepsilon}_{it}$. In this way, a random walk stock process suggests that the food bank only reacts to daily stock changes, treating past stock changes as a lost cause, not trying to offset past losses. This means allowing previous losses to propagate completely over time, creating the random walk. This gives us an intuitive way to test for stationarity - test whether winnings depend on previous stocks.

Substituting the equilibrium stock process into Equation F.1 yields:

$$\mathbf{w}_{it}^T \mathbf{z}_t^g = (\alpha - I)\mathbf{x}_{it} + (\delta - I)\sum_{s=1}^{\infty} \delta^{s-1}[\alpha \mathbf{x}_{it-s} + \boldsymbol{\varepsilon}_{it-s}] + \boldsymbol{\varepsilon}_{it}$$

Now, suppose $\boldsymbol{\varepsilon}_{it}$ can be decomposed according to: $\boldsymbol{\varepsilon}_{it} = \gamma \mathbf{r}_{it} + \boldsymbol{\nu}_{it}$, where $\mathbf{r}_{it}$ is a vector of observables that impact the food bank i's winnings (all of which must be independent of $\mathbf{x}_i$). These can be considered non-stock factors that impact winnings. For example, I focus on the total supply in period t, by storage type, the number of other food banks who placed a bid on a lot of each storage type, as well as the minimum distance between food bank i and a lot of each storage type in period t. Importantly I require that $E[\mathbf{r}_{is} \otimes \boldsymbol{\nu}_{it}] = \mathbf{0}$. As above, I expect $\boldsymbol{\nu}_{it}$ to be lumpy, non-normal, and with possibly non-zero mean. We can then re-write the equilibrium winnings process as:

$$\mathbf{w}_{it}^T \mathbf{z}_t^g = \gamma \mathbf{r}_{it} + (\delta - I)\sum_{s=1}^{\infty} \delta^{s-1}\gamma \mathbf{r}_{it-s} + (\delta - I)\sum_{s=1}^{\infty} \delta^{s-1}[\alpha \mathbf{x}_{it-1} + \boldsymbol{\nu}_{it-s}] + (\alpha - I)\mathbf{x}_{it} + \boldsymbol{\nu}_{it}$$

Importantly, this is a regression equation that could, hypothetically, be consistently estimated. However, there are easier ways to consider a test of stationarity. Consider the simple null hypothesis that $\delta = I$, so that the equilibrium stock process is a random walk. Under this null hypothesis, the equilibrium winnings process does not respond to the previous period's evening state. Therefore, consider the following regression specification:

$$\mathbf{w}_{it}^T \mathbf{z}_t^g = \beta^0 + \beta^1 \mathbf{r}_{it} + \beta^2 \mathbf{r}_{it-1} + \vartheta_{it} \tag{F.2}$$

Under this null hypothesis $\beta^2 = \mathbf{0}$. This is essentially an Anderson-Rubin test. The intuition is that we consider whether winnings respond to lagged non-stock factors. These lagged factors likely impact lagged winnings. If they are found to impact present winnings this suggests current winnings responds to the lagged stock. If current winnings do not depend on the lagged stock the food bank ignores their previous stocks, so the equilibrium stock process follows a random walk.

Results for this test are presented in Appendix F.3. As well as presenting aggregated results, assuming every food bank has the same δ, I consider disaggregated results. However in this case I have significantly less power. Broadly, however, I find strong evidence of stationarity.

F.3 Results

Figure F.2 presents results from this regression. I run each regression separately for each storage method, focusing on food banks that won at least 100 lots. I also include a second lag of $\mathbf{r}_t$, as this allows us to interpret the coefficients on $\mathbf{r}_{t-1}$.[4] I include factors for all storage methods in every regression, but only present results for the matching storage types, since these are expected to be the most useful. I also include a dummy variable (and its lags) for whether no lots of a particular type were auctioned on a particular day.

There is evidence of stationarity, given that we can reject the null hypothesis $\delta = I$ at the 1% significance level ($p < 0.001$). Most of the coefficients have the expected signs. Non-lagged factors are almost all statistically significant, with winnings increasing in the amount of food allocated, decreasing in the minimum distance between the bidder and lots, and decreasing in the number of rival bidders bidding on lots of that type. Coefficients on the lagged factors generally have the opposite signs, as we expect. Lagged distance is never statistically significant, even though contemporaneous distance is, while lagged pounds and rival bidders are generally significant. Comparing coefficients across lagged and unlagged variables we can extract δ. Assume zero-off diagonal elements, then diagonal elements are estimated to be in the region of 0.8. This uses the formula $\hat{\delta} = 1 - \frac{\hat{\beta}^1}{\hat{\beta}^2}$.

I also consider results for individual food banks. I run a specification as above allowing for individual food bank coefficients. When considering winnings for storage type l I only include factors and lagged factors for storage type l. I can reject the null-hypotheses that $\delta = I$ for 25% of analysed food banks. These are predominantly the large food banks observed regularly bidding and winning on the Choice System. For many food banks the tests are under powered, with even 58% of the unlagged

[4] Conditional on $\mathbf{r}_{t-2}$ it is reasonable that $\mathbf{r}_{t-1}$ and $\mathbf{r}_t$ are uncorrelated with $\mathbf{r}_{t-s}$ for $s \geq 3$, the omitted variables.

factors being statistically insignificant at 5% level of significance. This is most likely due to having only a small amount of variation in winnings, given that many food banks do not win very frequently. Interestingly Distance and Lagged Distance are much more likely to be statistically significant when I run the individual analysis.

FIGURE F.2: Results: Stationarity

	Non-Food	Dried	Tinned/Bottled	Refrigerated	Fresh
Total Pounds	0.018	0.023	0.017	0.020	0.018
	(0.004)	(0.005)	(0.004)	(0.006)	(0.005)
Lagged Pounds	-0.001	-0.003	-0.001	-0.004	-0.001
	(0.001)	(0.001)	(0.001)	(0.003)	(0.001)
Minimum Distance	-0.856	-1.528	-0.884	-0.480	-0.288
	(0.285)	(0.450)	(0.264)	(0.751)	(0.107)
Lagged Distance	0.171	-0.024	0.187	0.116	0.030
	(0.118)	(0.160)	(0.107)	(0.320)	(0.074)
Active Bidders	-53.638	-71.039	-28.601	-593.755	-43.187
	(19.999)	(29.534)	(22.572)	(189.673)	(18.443)
Lagged Bidders	14.254	34.077	15.683	84.825	7.250
	(6.940)	(15.380)	(7.975)	(61.444)	(4.552)
FB Fixed Effects	✓	✓	✓	✓	✓
2nd Lags	✓	✓	✓	✓	✓

Note: Standard Errors clustered within food bank and period.
Regression for storage method l includes regressors for all other storage methods also.

F.4 Covariance Stationarity

I now demonstrate that this assumption about the equilibrium stock process yields two sets of intuitive information about parameters $\boldsymbol{\mu}_i$ and Σ_i.

F.4.1 μ_i

Take an expectation of the equilibrium winnings equation for:

$$E[\mathbf{w}_{it}^T \mathbf{z}_t^g] = (\alpha - I)\boldsymbol{\mu}_i + (\delta - I)\sum_{s=1}^{\infty} \delta^{s-1}[\alpha\boldsymbol{\mu}_i + E[\boldsymbol{\varepsilon}]] + E[\boldsymbol{\varepsilon}]$$

Recognise that $\sum_{s=1}^{\infty} \delta^{s-1} = (I - \delta)^{-1}$, so that we are left with:$E[\mathbf{w}_{it}^T \mathbf{z}_t^g] = -\boldsymbol{\mu}_i$ Therefore, on average winnings offset net local donation. I use this constraint to build prior means for $\boldsymbol{\mu}$. I do not impose this relationship on account of the difficulty of efficiently estimating $E[\mathbf{w}_{it}^T \mathbf{z}_t^g]$ in the presence of auto-correlation, meaning my estimates are likely to be imprecise.

F.4.2 Σ_i

Take the variance of the equilibrium winnings equation, recognising that we assumed $\mathbf{x}$ is uncorrelated over time, and uncorrelated with $\boldsymbol{\varepsilon}$:

$$Var(\mathbf{w}_{it}^T\mathbf{z}_t^g) = (\alpha - I)\Sigma_i(\alpha - I) + (\delta - I)[\sum_{s=1}^{\infty}\delta^{s-1}\alpha\Sigma_i\alpha\delta^{s-1}](\delta - I)$$
$$+ Var(\boldsymbol{\varepsilon}_t + (\delta - I)[\sum_{s=1}^{\infty}\delta^{s-1}\boldsymbol{\varepsilon}_{t-1})$$

$Var(\boldsymbol{\varepsilon}_t + (\delta - I)[\sum_{s=1}^{\infty}\delta^{s-1}\boldsymbol{\varepsilon}_{t-1})$ is evidently positive definite. Meanwhile, the infinite geometric series does not have a simple form. To simplify matters, focus on the case where δ, α, and Σ_i are diagonal matrices. This is relevant since I impose that Σ_i is diagonal in the empirical model. In this case, applying the rule for infinite geometric series, with $\delta_l < 1$ we have

$$Var(\mathbf{w}_{it}^T\mathbf{z}_{lt}^g) = (\alpha_l - 1)^2\Sigma_i^{ll} + \frac{(\delta_l - 1)^2}{1 - \delta_l^2}\alpha^2\Sigma_i^{ll} + Var(\varepsilon_{lt} + (\delta_l - 1)[\sum_{s=1}^{\infty}\delta_l^{s-1}\varepsilon_{lt-1})$$

Since $\alpha \in [0, 1]$, the first part of this expression varies between Σ_i^{ll} and $[(1 - \delta_l)/2]\Sigma_i^{ll}$, when $\alpha = (1 + \delta_l)/2$. Therefore, since $Var(\varepsilon_{lt} + (\delta_l - 1)[\sum_{s=1}^{\infty}\delta_l^{s-1}\varepsilon_{lt-1}) > 0$, I get: $Var(\mathbf{w}_{it}^T\mathbf{z}_{lt}^g) > \frac{1-\delta_l}{2}\Sigma_i^{ll}$

Therefore, I can bound Σ_i^{ll}, conditional on δ. I consider two benchmark cases, $\delta_l = 0$, so that, on average each period, winnings totally adapt to changes in previous stocks. This implies $\Sigma_i^{ll} < 2Var(\mathbf{w}_{it}^T\mathbf{z}_{lt}^g)$. I also consider $\delta = 49/50$, so that winnings do not strongly react to previous stocks. This implies $\Sigma_i^{ll} < 100Var(\mathbf{w}_{it}^T\mathbf{z}_{lt}^g)$.

I use the $\delta = 49/50$ case for a hard upper bound on Σ_i^{ll}, under the prior that $\delta < 49/50$. I use the $\delta = 0$ case for what is essentially the prior mean of Σ_i^{ll}, albeit with a very low prior weight, given by the degrees of freedom in the Normal-inverse-Gamma distribution. Full details of how I build these weakly informative priors is given in Appendix K.

Appendix G

Inverse Bid System

In this Appendix I demonstrate that, in addition to the transition equation given in Assumption 7, a food bank's optimisation problem yields the Observation and Censoring equations given in Equation 4.4. I focus on the case for quadratic parametrisation of k. The general case is presented in Chapter 3.

G.1 Set-up

Imposing the parametrisation given in section 4.3.4, and conditional on $\mathbf{d}_i^*$, the bidder's maximisation problem is given by:

$$\max_{\mathbf{b}} \left\{ \sum_l \Gamma_l(b_l, d_l^*; \mathbf{s})(v_l - b_l) + \sum_a P_a(\mathbf{b}, \mathbf{d}^*; \mathbf{s})[\Phi \mathbf{s}_i^{ah} + \mathbf{s}_i^{ag^T} \Psi \mathbf{s}_i^{ag}] \qquad s.t.\ b_l \geq R_l \right\}$$

G.1.1 Simplification

I now show how the maximand can be simplified to:

$$\sum_l \Gamma_l(b_l, d_l^*; \mathbf{s})(v_l - b_l + \Phi \mathbf{z}_l^h + 2\mathbf{z}_l^{g^T} \Psi \mathbf{s}_i^g + \mathbf{z}_l^{g^T} \Psi \mathbf{z}_l^g + \sum_{m \neq l} \Gamma_m(b_m, d_m^*; \mathbf{s}) \mathbf{z}_l^{g^T} \Psi \mathbf{z}_m^g) + \Phi \mathbf{s}_i^h + \mathbf{s}_i^{g^T} \Psi \mathbf{s}_i^{g^T}$$

First, recognise that $\sum_a P_a(\mathbf{b}, \mathbf{d}; \mathbf{s}) = 1$, as we sum over mutually exclusive and exhaustive events. We can then write, for example, $\sum_a P_a(\mathbf{b}, \mathbf{d}; \mathbf{s}) \mathbf{s}_i^{g^T} \Psi \mathbf{s}_i^g = \mathbf{s}_i^{g^T} \Psi \mathbf{s}_i^g$.

Second, $\mathbf{s}_i^a = \mathbf{s}_i + \mathbf{z}\mathbf{w}_i^a$, where $\mathbf{w}_i^a$ is the $L \times 1$ vector with entry l equal to 1 if i wins lot l in combinatorial outcome a and zero otherwise. The matrix z gives the size and composition of lots. Exploiting $\sum_a P_a = 1$, re-write the maximand as:

$$
\begin{aligned}
&\sum_l \Gamma_l(b_l, d_l^*; \mathbf{s})(v_l - b_l) + \sum_a P_a(\mathbf{b}, \mathbf{d}^*; \mathbf{s})[\Phi \mathbf{s}_i^{ah} + \mathbf{s}_i^{ag^T} \Psi \mathbf{s}_i^{ag}] \\
&= \sum_l \Gamma_l(b_l, d_l^*; \mathbf{s})(v_l - b_l) + \sum_a P_a(\mathbf{b}, \mathbf{d}^*; \mathbf{s})[\Phi \mathbf{z}^h \mathbf{w}_i^a + \mathbf{s}_i^{ag^T} \Psi \mathbf{s}_i^{ag} - \mathbf{s}_i^{g^T} \Psi \mathbf{s}_i^g] \\
&\qquad + \Phi \mathbf{s}_i^h + \mathbf{s}_i^{g^T} \Psi \mathbf{s}_i^g \\
&= \sum_l \Gamma_l(b_l, d_l^*; \mathbf{s})(v_l - b_l) + \sum_a P_a(\mathbf{b}, \mathbf{d}^*; \mathbf{s})[\Phi \mathbf{z}^h \mathbf{w}_i^a + \mathbf{w}_i^{a^T} \mathbf{z}^{g^T} \Psi (\mathbf{z}^g \mathbf{w}_i^a + 2\mathbf{s}_i^g)] \\
&\qquad + \Phi \mathbf{s}_i^h + \mathbf{s}_i^{g^T} \Psi \mathbf{s}_i^g
\end{aligned}
$$

Where the final line follows from quadraticness: $\mathbf{s}_i^{ag^T} \Psi \mathbf{s}_i^{ag} = (\mathbf{s}_i + \mathbf{z}\mathbf{w}_i^a)^T \Psi (\mathbf{s}_i + \mathbf{z}\mathbf{w}_i^a)$ and so $\mathbf{s}_i^{ag^T} \Psi \mathbf{s}_i^{ag} - \mathbf{s}_i^{g^T} \Psi \mathbf{s}_i^g = \mathbf{w}_i^{a^T} \mathbf{z}^{g^T} \Psi (\mathbf{z}^g \mathbf{w}_i^a + 2\mathbf{s}_i^g)$.

Finally, recognise that $\sum_a P_a(\mathbf{b}, \mathbf{d}; \mathbf{s}) \mathbf{s}_i^a = \mathbf{s}_i + \sum_l \Gamma_l(b_l, d_l; \mathbf{s}) \mathbf{z}_l$. This arises because stocks are additive in winnings.[1] This ensures that:

$$
\begin{aligned}
&= \sum_l \Gamma_l(b_l, d_l^*; \mathbf{s})(v_l - b_l + \Phi \mathbf{z}_l^h + 2\mathbf{z}_l^{g^T} \Psi \mathbf{s}_i^g) + \sum_a P_a(\mathbf{b}, \mathbf{d}^*; \mathbf{s}) \mathbf{w}_i^{a^T} \mathbf{z}^{g^T} \Psi \mathbf{z}^g \mathbf{w}_i^a \\
&\qquad + \Phi \mathbf{s}_i^h + \mathbf{s}_i^{g^T} \Psi \mathbf{s}_i^{g^T} \\
&= \sum_l \Gamma_l(b_l, d_l^*; \mathbf{s})(v_l - b_l + \Phi \mathbf{z}_l^h + 2\mathbf{z}_l^{g^T} \Psi \mathbf{s}_i^g + \mathbf{z}_l^{g^T} \Psi \mathbf{z}_l^g + \sum_{m \neq l} \Gamma_m(b_m, d_m^*; \mathbf{s}) \mathbf{z}_l^{g^T} \Psi \mathbf{z}_m^g) \\
&\qquad + \Phi \mathbf{s}_i^h + \mathbf{s}_i^{g^T} \Psi \mathbf{s}_i^{g^T}
\end{aligned}
$$

The final line follows from: (1) that $(\mathbf{z}^g \mathbf{w}_i^a)^T \Psi (\mathbf{z}^g \mathbf{w}_i^a) = \sum_l \sum_m (w_{il}^a \mathbf{z}_l^g)^T \Psi (w_{im}^a \mathbf{z}_m^g)$, which arises from quadraticness. Because this object only depends on pairs of winnings, we can marginalise out the probability of receiving a particular pair, so that: $\sum_a P_a(\mathbf{b}, \mathbf{d}; \mathbf{s})(\mathbf{z}^g \mathbf{w}_i^a)^T \Psi \mathbf{z}^g \mathbf{w}_i^a = \sum_l \sum_m Prob(\text{win } l \text{ and } m | \mathbf{b}, \mathbf{d}; \mathbf{s})(\mathbf{z}_l^g)^T \Psi (\mathbf{z}_m^g)$. (2), imposing part (iv) of Assumption 9: $Prob(\text{win } l \text{ and } m | \mathbf{b}, \mathbf{d}; \mathbf{s}) = \Gamma_l(b_l, d_l; \mathbf{s}) \Gamma_m(b_m, d_m; \mathbf{s})$ for $m \neq l$, and $\Gamma_l(b_l, d_l; \mathbf{s})$ otherwise.

[1] This result can be seen by focusing on the expectation of $\mathbf{w}_i^{a^T} \mathbf{z}_l$ for one particular l. As we sum across all the combinations in which they win lot l, the sum of these probabilities is just Γ_l, the marginal probability they win lot l.

G.2 First Order Conditions, conditional on entry

Written out in it's full simplified form, the lagrangian for this problem is given by:

$$L(\mathbf{b}|\mathbf{d}^*, \mathbf{s}, \boldsymbol{v}) = \sum_l \Gamma_l(b_l, d_l^*; \mathbf{s})(v_l - b_l + \Phi \mathbf{z}_l^h + 2\mathbf{z}_l^{gT} \Psi \mathbf{s}_i^g + \sum_m \Gamma_m(b_m, d_m^*; \mathbf{s}) \mathbf{z}_l^{gT} \Psi \mathbf{z}_m^g)$$
$$+ \Phi \mathbf{s}_i^h + \mathbf{s}_i^{gT} \Psi \mathbf{s}_i^{gT} - \sum_l \Lambda_l (R_l - b_l)$$

Where Λ_l are lagrangian multipliers. Necessary first order conditions are then:

$$0 = \nabla_l \Gamma_l(b_l^*, d_l^*; \mathbf{s})(v_l - b_l^* + \Phi \mathbf{z}_l^h + \mathbf{z}_l^{gT} \Psi \mathbf{z}_l^g + 2\mathbf{z}_l^{gT} \Psi \mathbf{s}_i^g + 2 \sum_{m \neq l} \Gamma_m(b_m^*, d_m^*; \mathbf{s}) \mathbf{z}_l^{gT} \Psi \mathbf{z}_m^g)$$
$$- \Gamma_l(b_l^*, d_l^*; \mathbf{s}) + \Lambda_l^*$$

Which rearranges for:

$$b_l^* + \frac{\Gamma_l(b_l^*, d_l^*; \mathbf{s})}{\nabla_b \Gamma_l(b_l^*, d_l^*; \mathbf{s})} - \Lambda_l^* = \Phi \mathbf{z}_l^h + \mathbf{z}_l^{gT} \Psi (\mathbf{z}_l^g + 2\mathbf{s}_i^g + 2 \sum_{m \neq l} \Gamma_m(b_m^*, d_m^*; \mathbf{s}) \mathbf{z}_l^g) + v_l = y_l$$

Let $y_l^* = b_l^* + \frac{\Gamma_l(b_l^*, d_l^*; \mathbf{s})}{\nabla_b \Gamma_l(b_l^*, d_l^*; \mathbf{s})} - \Lambda_l^*$. When we observe $b_l^* > R_l$, we can infer $\Lambda_l^* = 0$, so that $y_l^* = b_l^* + \frac{\Gamma_l(b_l^*, d_l^*; \mathbf{s})}{\nabla_b \Gamma_l(b_l^*, d_l^*; \mathbf{s})} = y_l$. In this case, the bidder is not constrained.

G.3 Reservation Price Bidding

When we observe $b_l^* = R_l$, the First Order Conditions break down, since as made clear in Section 4.3.3, beliefs are non-differentiable at the reservation price due to the non-negligible probability of ties. Therefore, consider the bidder's decision to bid at the reservation price, bidding vector b^*, compared to just above the reservation price at $R_l + 1$ playing vector b^+. Elements $m \neq l$ of these vectors will otherwise be equal. This implies that:

$$\sum_l \Gamma_l(b_l^*, d_l^*; \mathbf{s})(v_l - b_l^* + \Phi \mathbf{z}_l^h + 2\mathbf{z}_l^{g^T} \Psi \mathbf{s}_i^g + \sum_m \Gamma_m(b_m^*, d_m^*; \mathbf{s}) \mathbf{z}_l^{g^T} \Psi \mathbf{z}_m^g)$$
$$\geq \sum_l \Gamma_l(b_l^+, d_l^*; \mathbf{s})(v_l - b_l^+ + \Phi \mathbf{z}_l^h + 2\mathbf{z}_l^{g^T} \Psi \mathbf{s}_i^g + \sum_m \Gamma_m(b_m^+, d_m^*; \mathbf{s}) \mathbf{z}_l^{g^T} \Psi \mathbf{z}_m^g)$$

Therefore

$$\Gamma_l(R_l, d_l^*; \mathbf{s})(v_l - R_l + \Phi \mathbf{z}_l^h + \mathbf{z}_l^{g^T} \Psi(\mathbf{z}_l^g + 2\mathbf{s}_i^g + 2\sum_m \Gamma_m(b_m^*, d_m^*; \mathbf{s}) \mathbf{z}_m^g))$$
$$\geq \Gamma_l(R_l + 1, d_l^*; \mathbf{s})(v_l - R_l - 1 + \Phi \mathbf{z}_l^h + \mathbf{z}_l^{g^T} \Psi(\mathbf{z}_l^g + 2\mathbf{s}_i^g + 2\sum_m \Gamma_m(b_m^*, d_m^*; \mathbf{s}) \mathbf{z}_m^g))$$

Therefore

$$[\Gamma_l(R_l+1, d_l^*; \mathbf{s}) - \Gamma_l(R_l, d_l^*; \mathbf{s})](v_l - R_l + \Phi \mathbf{z}_l^h + \mathbf{z}_l^{g^T} \Psi(\mathbf{z}_l^g + 2\mathbf{s}_i^g + 2\sum_m \Gamma_m(b_m^*, d_m^*; \mathbf{s}) \mathbf{z}_m^g)) \leq \Gamma_l(R_l+1, d_l^*; \mathbf{s})$$
$$y_l = \Phi \mathbf{z}_l^h + 2\mathbf{z}_l^{g^T} \Psi \mathbf{s}_i^g + \mathbf{z}_l^{g^T} \Psi(\mathbf{z}_l^g + 2\mathbf{s}_i^g + 2\sum_m \Gamma_m(b_m^*, d_m^*; \mathbf{s}) \mathbf{z}_m^g) + v_l \leq R_l + \frac{\Gamma_l(R_l+1, d_l^*; \mathbf{s})}{\Gamma_l(R_l+1, d_l^*; \mathbf{s}) - \Gamma_l(R_l, d_l^*; \mathbf{s})} = y_l^*$$

Therefore, $y_l^* \geq y_l$

G.3.1 Bidding R_l vs Not Bidding

At the margin, the bidder must weakly prefer to enter and bid the reservation price, playing bidding/entry vector $\mathbf{b}^*, \mathbf{d}^*$, than to not enter at all, playing bidding/entry vector $\mathbf{b}^-, \mathbf{d}^-$. These vectors are identical apart from for lot l. This implies:

$$\sum_l \Gamma_l(b_l^*, d_l^*; \mathbf{s})(v_l - b_l^* + \Phi \mathbf{z}_l^h + 2\mathbf{z}_l^{g^T} \Psi \mathbf{s}_i^g + \sum_m \Gamma_m(b_m^*, d_m^*; \mathbf{s}) \mathbf{z}_l^{g^T} \Psi \mathbf{z}_m^g)$$
$$\geq \sum_l \Gamma_l(b_l^-, d_l^-; \mathbf{s})(v_l - b_l^- + \Phi \mathbf{z}_l^h + 2\mathbf{z}_l^{g^T} \Psi \mathbf{s}_i^g + \sum_m \Gamma_m(b_m^-, d_m^-; \mathbf{s}) \mathbf{z}_l^{g^T} \Psi \mathbf{z}_m^g)$$

Therefore $$\Gamma_l(R_l, d_l^*; \mathbf{s})(v_l - R_l + \Phi \mathbf{z}_l^h + \mathbf{z}_l^{g^T} \Psi(\mathbf{z}_l^g + 2\mathbf{s}_i^g + 2\sum_m \Gamma_m(b_m^*, d_m^*; \mathbf{s}) \mathbf{z}_m^g)) \geq 0$$

Therefore $$y_l = \Phi \mathbf{z}_l^h + 2\mathbf{z}_l^{g^T} \Psi \mathbf{s}_i^g + \mathbf{z}_l^{g^T} \Psi(\mathbf{z}_l^g + 2\mathbf{s}_i^g + 2\sum_m \Gamma_m(b_m^*, d_m^*; \mathbf{s}) \mathbf{z}_m^g) + v_l \geq R_l$$

Therefore $R_l \leq y_l \leq y_l^*$

G.3.2 Entry Decisions

If a food bank chooses not to enter the auction for lot l, then at the margin they must weakly prefer to not enter the auction, than to enter and bid the reservation price. This is just the complement of the previous inequality, allowing us to infer that $d_l^* = 0$ implies $y_l \leq R_l$.[2].

[2]Furthermore, negative definiteness of Ψ implies that if they prefer to not bid than enter and bid the reservation price, they also cannot prefer to bid strictly above the reservation price

Appendix H

Discriminatory Auctions

In this Appendix I discuss how the discriminatory auctions of homogenous lots are take into account. In Appendix H.1 I explain the rules of the discriminatory auctions. In Appendix H.2 I derive the Inverse Bid System in the presence of discriminatory auctions, presenting a generalisation of Appendix G. In Appendix H.3 I discuss how I use the inverse bid system in estimation.

H.1 Framework

In a discriminatory auction of R homogenous lots food banks place as many bids as they like. Lots are allocated to the R highest bidders, if at least R bids were placed, and bidders pay their bids.

7% of unique auctions that occur contain more than one homogenous good, and are auctioned in discriminatory fashion. Of all the lots allocated, 21% of lots are auctioned with at least one additional identical load. The vast majority of lots sold at the reservation price are lots from discriminatory auctions with a large number of homogenous loads being auctioned in this manner (the remainder are fresh loads). Food banks recognise that the lowest winning bid is likely to be very low. Loads allocated in this manner are always homogenous and come from the same source.[1]

Figure H.1 panel (A) shows a histogram of the number of loads included in each unique auction, conditional on at least two loads. As not every load may be sold on a particular date, Panel (B) shows the number of homogenous loads for each auction $\times$ date, conditional on at least two loads.

[1] However some homogenous lots, that are all donated by the same donors, are auctioned using simultaneous auctions. I do not analyse the determinants of this decision.

FIGURE H.1: Distribution of Homogenous Loads

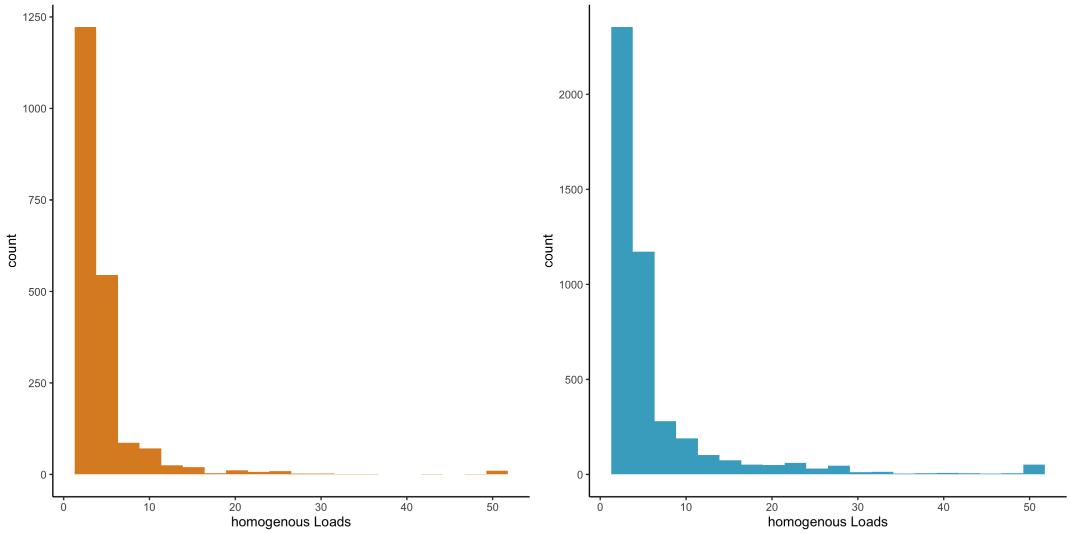

Note: These plots show histograms of the number of homogenous loads included in each auction, conditional on at least two homogenous loads. Panel (A) shows the number of loads for each unique auction. Panel (B) shows the number of loads for each date × auction, recognising that not every load is sold right away. More than 50 loads are grouped into the 50 category. This includes one auction with 181 loads, and 9 other auctions with between 60 and 80 Loads.

H.2 Adjusted Inverse Bid System

The only difference between simultaneous and discriminatory first-price auctions is that when the bidder wins on their rth lowest bid on lot l, they must also win on all higher bids. Write the rth bid on auction l as b_{lr}. Bids are ascending, such that $b_{lr} \leq b_{lr+1}$, up to $r = R_l$ the total number of loads contained in the lot. Food bank i's belief about the probability they win on their rth bid on lot l is given by $\Gamma_l(b_{lr})$, suppressing the entry decision and state variables for ease of notation. Their belief they win on bids r through to R, but no lower, is given by $\Gamma_l(b_{lr}) - \Gamma_l(b_{lr-1})$. This is the probability the lowest rival winning bid is between b_{lr} and b_{lr-1}.

Because lots are homogenous, I treat the lot specific value v_{il} as constant across the loads in auction l. Making use of the assumed parametrisation, the expected payoff is given by:

$$\pi(\mathbf{b}, \mathbf{d}) = \sum_{l}^{L} \sum_{r}^{R_l} \Gamma_l(b_{lr}, d_{lr}) v_l - b_{lr} + \Phi \mathbf{z}_l^h + \mathbf{z}_l^{g\,T} \Psi [\mathbf{z}_l^g + 2\mathbf{s}_t^g])$$

$$+ \sum_{l}^{L} \sum_{r}^{R_l} \Gamma_l(b_{lr}, d_{lr}) \mathbf{z}_l^{g\,T} \Psi [(R_l - r)\mathbf{z}_l^g + \sum_{m \neq l} \sum_{n}^{R_m} \Gamma_m(b_{mn}, d_{mn}) \mathbf{z}_m^g] + \sum_{n=1}^{r-1} \Gamma_l(b_{ln}, d_{ln}) \mathbf{z}_l^{g\,T} \Psi \mathbf{z}_l^g$$

The first line is the lot-specific component of the pay-off. The second line is the combinatorial component; how the pay-off varies with winnings from other auctions. In the simultaneous only case this is given by $\sum_l^L \Gamma_l(b_l, d_l)\mathbf{z}_l^{gT}\Psi[\sum_{m\neq l}\Gamma_m(b_m, d_m)\mathbf{z}_m^g]$ only. If they win with bid b_{lr}, then they win the $(R_l - r)$ higher lots with certainty. This is why the $(R_l - r)\mathbf{z}_l^g$ term does not have a probability multiplier. Likewise, with probability $\Gamma_l(b_{ln}, d_{ln})$ they also win on bid b_{lr} (for $n < r$). So, in expectation, they also gain this $\Gamma_l(b_{ln}, d_{ln})\mathbf{z}_l^{gT}\Psi\mathbf{z}_l^g$ combinatorial term.

Conditional on $d_{lr} = 1$, first order conditions with respect to bid b_{lr} are given by:

$$\begin{aligned}\Gamma_l(b_{lr}, d_{lr}) &= \nabla_{b_{lr}}\Gamma_l(b_{lr}, d_{lr})(v_l - b_{lr} + \Phi\mathbf{z}_l^h + \mathbf{z}_l^{gT}\Psi[\mathbf{z}_l^g + 2\mathbf{s}_t^g]) \\ &+ \nabla_{b_{lr}}\Gamma_l(b_{lr}, d_{lr})\mathbf{z}_l^{gT}\Psi[(R_l - r)\mathbf{z}_l^g + \sum_{m\neq l}\sum_n^{R_m}\Gamma_m(b_{mn}, d_{mn})\mathbf{z}_m^g + \sum_n^{r-1}\Gamma_l(b_{ln}, d_{ln})\mathbf{z}_l^g] \\ &+ \mathbf{z}^{gT}\Psi\nabla_{b_{lr}}\Gamma_l(b_{lr}, d_{lr})[\sum_{m\neq l}\sum_n^{R_m}\Gamma_m(b_{mn}, d_{mn})\mathbf{z}_m^g + \sum_{r+1}^{R_l}\mathbf{z}_l^g]\end{aligned}$$

The inverse bid system is given by:

$$\begin{aligned}\xi_{lr}(\mathbf{b}; k) = b_{lr} &+ \frac{\Gamma_l(b_{lr}, d_{lr})}{\nabla_{b_{lr}}\Gamma_l(b_{lr}, d_{lr})} - \Phi\mathbf{z}_l^h - \mathbf{z}_l^{gT}\Psi[\mathbf{z}_l^g + 2\mathbf{s}_t^g] \\ &- \mathbf{z}_l^{gT}\Psi[2(R_l - r)\mathbf{z}_l^g + 2\sum_{m\neq l}\sum_n^{R_m}\Gamma_m(b_{mn}, d_{mn})\mathbf{z}_m^g + \sum_n^{r-1}\Gamma_l(b_{ln}, d_{ln})\mathbf{z}_l^g]\end{aligned}$$

Next, we consider the decision to enter and bid the reservation price, rather than bid just above it. Importantly, if $b_{lr} = R + 1$ then all lower bids must be ether the same, at the reservation price, or not entered. Setting the utility at the reservation price greater than or equal to utility just above the reservation price we obtain the following upper bound on ξ_{lr}:

$$\begin{aligned}\xi_{lr}(\mathbf{b}; k) \leq b_{lr} &+ \frac{\Gamma_l(R+1, 1)}{\Gamma_l(R+1, 1) - \Gamma_l(R, 1)} - \Phi\mathbf{z}_l^h - \mathbf{z}_l^{gT}\Psi[\mathbf{z}_l^g + 2\mathbf{s}_t^g] \\ &- \mathbf{z}_l^{gT}\Psi[2(R_l - r)\mathbf{z}_l^g + 2\sum_{m\neq l}\sum_n^{R_m}\Gamma_m(b_{mn}, d_{mn})\mathbf{z}_m^g + \sum_n^{r-1}\Gamma_l(b_{ln}, d_{ln})\mathbf{z}_l^g]\end{aligned}$$

Next, consider the decision to enter and bid the reservation price, versus not bidding at all. We obtain the following lower bound for ξ_{lr}:

$$\xi_{lr}(\mathbf{b};k) \geq b_{lr} - \Phi\mathbf{z}_l^h$$
$$- \mathbf{z}_l^{gT}\Psi[\mathbf{z}_l^g + 2\mathbf{s}_t^g + 2(R_l - r)\mathbf{z}_l^g + 2\sum_{m\neq l}\sum_{n}^{R_m}\Gamma_m(b_{mn}, d_{mn})\mathbf{z}_m^g + \sum_{n}^{r-1}\Gamma_l(b_{ln}, d_{ln})\mathbf{z}_l^g]$$

As in Appendix G, the decision not to enter then yields the same upper bound as this lower bound. There is a clear similarity between this system of equations and those presented previously.

H.3 Computation

First Stage

Food bank i wins lot l, load r, with the probability that their bid exceeds the lowest winning bid on that lot. Therefore it is their belief about the distribution of the lowest winning bid that matters. When fewer food banks place bids than loads are being auctioned, this is just the reservation price. Therefore, when estimating beliefs I ignore all higher winning bids.

H.3.1 Second Stage

The difficulty in the second stage is that the lot specific idiosyncratic pay-off is assumed to be perfectly correlated within a discriminatory auction. Allowing these objects to vary, even if highly correlated, does not make sense given loads are perfectly homogenous. By definition of the homogenous lots, there can be no unobserved variation in lot characteristics across loads.

When performing the data augmentation step in the second stage of the procedure I treat discriminatory auctions correctly. That is, I correctly sample censored bids from their posterior distribution and sample states from their conditional posterior. However, in the Gibbs Sampling step I only use information from the highest bid (which may also be censored). If the idiosyncratic terms are perfectly correlated then I do not gain any additional information by using lower bids.[2] Finally, when considering model fit by simulating the Choice System, I do simulate these auctions (albeit limiting food banks to only be able to place up to 5 bids on a single auction). Therefore if my simplification does lead to inaccuracy or bias, this should become evident.

H.3.2 Third Stage

Finally, in the third step I treat these auctions properly in how I evaluate the ex-ante continuation value. The derivation presented in Appendix J extends easily to allow for the discriminatory case, and only requires summing over the probability of these combinatorial wins.

[2]In practice specification error will means I cannot rationalise the model with identical lot specific pay-offs. I could treat them as correlated, and estimate a correlation parameter, but in practice I will gain very little additional information from doing so. First, if they do not place bids at all in a particular auction I gain no additional information from considering additional bids. If they place one bid, but not a second, some information can be gained by considering why they didn't place this second bid. However, even given specification error we expect a very high degree of correlation, meaning I do not gain a large amount of information beyond that contained in the first bid.

Appendix I

Non-parametric Identification

I now prove Proposition 5, that the model is non-parametrically identified. In Chapter 3 I proved that conditional on identifying i) the equilibrium distribution of bids (conditional on the state) and ii) the state transition process, that observed variation in the state is sufficient for non-parametric identification of both the deterministic flow payoff function j, and the lot specific idiosyncratic payoff distribution F^v.[1] Therefore, it suffices to prove that the equilibrium conditional bid distribution and state transition process are both non-parametrically point identified.[2]

The argument builds on the key result in Hu and Shum (2012), closely following their set up and proof using spectral decomposition techniques for linear operators. The key distinction between our arguments consists of the introduction of an additional intermediate variable (winnings) that acts as an observed shifter, or instrument, of the unobserved process. This requires that conditional on $\mathbf{s}_t$, bids are independent of previous winnings, and that conditional on $\mathbf{b}_t$, the winnings from that period are independent of $\mathbf{s}_t$.[3] This significantly reduces the assumptions required for identification. It also yields clear intuition behind identification of the

[1] While I proved this for the case where states are discrete, it is clear how this intuition extends to the setting with continuous states. Likewise, the identification argument presented here also holds when the state space is finite. Furthermore, this argument and the proof contained also applies to games with latent states. If we think of the vector $\mathbf{s}_t$ as being the stocks of every player, $\mathbf{b}_t$ the bids of every player, and $\mathbf{w}_t$ the winnings of every player, this argument still ensures the full conditional distribution of bids and transition processes are non-parametrically identified. This is not an approach I take in the main text.

[2] I do not explicitly consider identification of the λ_i parameters, which give the marginal value of wealth for each bidder. These parameters are identified by comparing bidding behaviour across bidders, under the required normalising assumption (to allow inter food bank comparisons) that the variance of lot-specific values is constant across bidders. See Section 4.2.5 for detailed discussion.

[3] Identification of my model is also driven by variation in choice sets over time, and the argument presented in this Appendix can easily be extended to include this type of variation, even allowing it to not be exogenous. However, introducing these additional observed states significantly increases the notation required. Meanwhile, as we will see, this type of variation is not necessary for identification, as variation in winnings alone is shown to be sufficient.

latent state: Variation in this observed shifter of the unobserved state pins down the relationship between bids and the unobserved state. Then variation in bids over time, holding constant the observed shifter, enables identification of the state transition process. Furthermore, depending on how the model is defined, this argument only requires 2 periods of data, not 4 periods as in Hu and Shum (2012). The final distinction is that this framework easily extends to the case when the latent state is multivariate, as the signals of the latent state, bids and winnings, are both also multivariate. I focus on the stationary setting, though it is clear how the arguments can easily be extended to the non-stationary setting considered in Hu and Shum (2012).

I make two additional simplifying assumptions: I assume that reservation prices do not bind, and that $\mathbf{z}_t^g$ has full rank. Binding reservation prices ensure the first order conditions do not hold with equality. However, as discussed in Appendix E.2, reservation prices are not a first-order issue, not substantially alter the identification problem. In the way that a censored regression model, which requires a Tobit or MAD specification, does not substantially alter the regression identification problem. The key intuition garnered from this simplified approach extends to the case with reservation prices. Meanwhile the rank condition on $\mathbf{z}_t^g$, the size and composition of lots auctioned each day, simply ensures that at least one of each food type is auctioned each period. This ensures that bidding in every period is informative of every dimension of the bidders' stocks. As it turns out, identification actually only requires this rank condition holds for 3 consecutive periods, which holds.

The identification argument proceeds as follows: In Appendix I.1 I restate the model in more general terms, before detailing sufficient assumptions on the model primitives for non-parametric point identification. In Appendix I.2 I prove non-parametric point identification, following the structure of Hu and Shum, 2012.

I.1 Assumptions

The econometrician observes bids $\mathbf{b}_t$ and winnings $\mathbf{w}_t$, but not the latent stock $\mathbf{s}_t$. The aim is to non-parametrically point identify the equilibrium conditional distribution of bids $f_{\mathbf{b}_t|\mathbf{s}_t}(\mathbf{b}_t|\mathbf{s}_t)$, and the stock transition process $f_{\mathbf{s}|\mathbf{w}_{t-1},\mathbf{s}_{t-1}}(\mathbf{s}_t|\mathbf{w}_{t-1},\mathbf{s}_{t-1})$. I abstract away from additional observed state variables that change over time, which in my application may includes the set of lots being auctioned each day. These variables can be included without changing the central arguments. For notational convenience, except in cases when it is necessary to avoid confusion, I denote density functions using just the subscript term, writing the conditional bid distribution as $f_{\mathbf{b}_t|\mathbf{s}_t}$. The full transition process is given by $f_{\mathbf{w}_t,\mathbf{b}_t,\mathbf{s}_t|\text{history}_{t-1}}$.

Assumption 12. *Limited Feedback:*

$i)$ $f_{\mathbf{w}_t,\mathbf{b}_t,\mathbf{s}_t|\text{history}_{t-1}} = f_{\mathbf{w}_t,\mathbf{b}_t,\mathbf{s}_t|\mathbf{w}_{t-1},\mathbf{b}_{t-1},\mathbf{s}_{t-1}}$

$ii)$ $f_{\mathbf{w}_t,\mathbf{b}_t,\mathbf{s}_t|\mathbf{w}_{t-1},\mathbf{b}_{t-1},\mathbf{s}_{t-1}} = f_{\mathbf{w}_t|\mathbf{b}_t} f_{\mathbf{b}_t|\mathbf{s}_t} f_{\mathbf{s}_t|\mathbf{w}_{t-1},\mathbf{s}_{t-1}}$

Part $ii)$ of this assumption is a restricted version of the equivalent assumption in Hu and Shum (2012). It imposes a specific causal structure with three parts: $a)$ Conditional on $\mathbf{b}_t$ the winnings from that period are independent of $\mathbf{s}_t$, $b)$ Conditional on $\mathbf{s}_t$ bids are independent of previous winnings, and $c)$ conditional on previous winnings $\mathbf{w}_{t-1}$ (and stocks $\mathbf{s}_{t-1}$) stocks are independent of previous bids. Part $b)$ is the key assumption here, ensuring that $\mathbf{w}_{t-1}$ acts like an excluded instrumental variable.

I.1.1 Linear Operators

Denote the supports of $\mathbf{s}_t, \mathbf{w}_t, \mathbf{b}_t$ as $\mathbb{S}, \mathbb{W}$, and $\mathbb{B}$ respectively.[4] Next, define a 'linear operator' $L_{x,y}$ as the map from the $L^{|y|}$ space of functions of y to the $L^{|x|}$ space of functions of x, such that for a function $g : \mathbb{R}^{|x|} \to \mathbb{R}^n$ (or equivalently $g \in L^{|x|}$):

$$(L_{x,y}g)y = \int f_{x,y}(x,y)g(x)dx$$

Likewise, define the diagonal operator $D_{x,y}$ as follows:

$$(D_{x,y}g)y = f_{x,y}(x,y)g(x)$$

Linear operators are close to infinite dimensional counterparts to matrices, with similar properties. They are regularly used in statistical completeness arguments, as has been used to argue non-parametric identification of instrumental variables

[4] Both $\mathbb{W}$, and $\mathbb{B}$ could be allowed to vary observably over time, but I drop this dependence on t for notational convenience, as it requires introducing additional observed state variables.

models (Newey and Powell, 2003). In particular, for an Instrumental Variable model of the form $y = g(x) + \varepsilon$, the function g is non-parametrically identified if there exists an instrumental variable z, mean independent of ε, that is complete for x, such that for all possible functions $\tilde{g}(x)$:

$$0 = \int (g(x) - \tilde{g}(x))f(x|z)dx \qquad \text{implies} \qquad g(x) = \tilde{g}(x) \forall x \in \mathbb{X}$$

Therefore, the only function $\tilde{g}(x)$ which satisfies $E[\varepsilon|z] = 0$ is the true g. In operator notation this can be written as $(L_{x|z}h)z = 0$ implies $h(x) = 0$ for all x, with $h(x) = g(x) - \tilde{g}(x)$. This condition holds if $L_{x|z}$ is injective, so that the inverse operator $L_{x|z}^{-1}$ exists. The injectivity property for linear operators is similar to the full rank property of a matrix, ensuring that inverses exist for both. This exemplifies the relationship between statistical completeness and instrument relevance in instrumental variables, in that z being complete for x is similar to z being a relevant instrument for x. In other words, variation in z creates 'enough' variation in x to enable us to fully pin down the relationship between x and y.

Assumption 13. *Invertibility:*

$i)$ For any $\mathbf{w}_t \in \mathbb{W}$ there exist a $\mathbf{b}_t \in \mathbb{B}$ and a neighbourhood $\bar{\mathbb{B}}$ around $\mathbf{b}_t$ such that for any $\bar{\mathbf{b}}_t \in \bar{\mathbb{B}}$, $L_{\mathbf{b}_{t+1},\mathbf{w}_t,\bar{\mathbf{b}}_t,\mathbf{w}_{t-1}}$ is injective.

$ii)$ For any $\mathbf{w}_t \in \mathbb{W}$ $L_{\mathbf{b}_{t+1}|\mathbf{w}_t,\mathbf{s}_t}$ is injective.

This assumption is essentially a pair of instrument relevance conditions. Part $i)$ requires that variation in past winnings $\mathbf{w}_{t-1}$ creates 'enough' variation in future bids $\mathbf{b}_{t+1}$ for some values of $\mathbf{w}_t$ and $\mathbf{b}_t$. Part $ii)$ then similarly requires that current stocks create 'enough' variation in future bids. The idea behind these conditions is that variation in $\mathbf{w}_{t-1}$ and $\mathbf{s}_t$ are, to an extent, substitutable. We know how $\mathbf{b}_{t+1}$ varies with $\mathbf{w}_{t-1}$, so we can then learn something about $\mathbf{s}_t$ from changes in $\mathbf{b}_{t+1}$.

This assumption is where the simplification about reservation prices and the rank of $\mathbf{z}^g$ has bite. When reservation prices bind, for some values of $\mathbf{w}_t$, they will place no bids irrespective of $\mathbf{b}_t$ and $\mathbf{w}_{t-1}$, and so variation in $\mathbf{w}_{t-1}$ creates no variation in $\mathbf{b}_{t+1}$. More generally, if a food bank never bids on a particular type of food (due to the reservation prices), we never learn about the food bank's preferences for that type of food. Meanwhile, when $\mathbf{z}_{t-1}^g$ is rank deficient, then so must be winnings

$\mathbf{w}_{t-1}$, so we do not gain information about every element of $\mathbf{s}_t$. For example, if a particular type of food is never auctioned, it is intuitive that we cannot learn about food banks' preferences for that type of food.

I.1.2 Assumptions on the Conditional Bid Distribution

Assumption 13 will be used to rely on a spectral decomposition type argument. Therefore, just like Hu and Shum (2012) I make an assumption to ensure the uniqueness of the spectral decomposition of my linear operators:

Assumption 14. *Uniqueness of Spectral Decomposition:*

For any $\mathbf{w}_t \in \mathbb{W}$ and any $\mathbf{s}_t \neq \bar{\mathbf{s}}_t \in \mathbb{S}$, there exists a $\mathbf{b}_t$ and corresponding neighbourhood $\bar{\mathbb{B}}$ satisfying Assumption 13 $i)$, such that for some $\bar{\mathbf{b}}_t \in \bar{\mathbb{B}}$:

$i)$ $0 < \frac{f_{\mathbf{w}_t|\mathbf{b}_t}(\mathbf{w}_t|\mathbf{b}_t)}{f_{\mathbf{w}_t|\mathbf{b}_t}(\mathbf{w}_t|\bar{\mathbf{b}}_t)} \frac{f_{\mathbf{b}_t|\mathbf{s}_t}(\mathbf{b}_t|\mathbf{s}_t)}{f_{\mathbf{b}_t|\mathbf{s}_t}(\bar{\mathbf{b}}_t|\mathbf{s}_t)} < C < \infty$, *for some finite constant C*

$ii)$ $\frac{f_{\mathbf{w}_t|\mathbf{b}_t}(\mathbf{w}_t|\mathbf{b}_t)}{f_{\mathbf{w}_t|\mathbf{b}_t}(\mathbf{w}_t|\bar{\mathbf{b}}_t)} \frac{f_{\mathbf{b}_t|\mathbf{s}_t}(\mathbf{b}_t|\mathbf{s}_t)}{f_{\mathbf{b}_t|\mathbf{s}_t}(\bar{\mathbf{b}}_t|\mathbf{s}_t)} \neq \frac{f_{\mathbf{w}_t|\mathbf{b}_t}(\mathbf{w}_t|\mathbf{b}_t)}{f_{\mathbf{w}_t|\mathbf{b}_t}(\mathbf{w}_t|\bar{\mathbf{b}}_t)} \frac{f_{\mathbf{b}_t|\mathbf{s}_t}(\mathbf{b}_t|\bar{\mathbf{s}}_t)}{f_{\mathbf{b}_t|\mathbf{s}_t}(\bar{\mathbf{b}}_t|\bar{\mathbf{s}}_t)}$

These two assumptions will ensure uniqueness of the eigenvalues for a spectral decomposition of some function of the linear operator $L_{\mathbf{b}_{t+1},\mathbf{w}_t,\bar{\mathbf{b}}_t,\mathbf{w}_{t-1}}$. Part $i)$ of the assumption is relatively innocuous, ensuring that for any winnings $\mathbf{w}$ and stocks $\mathbf{s}$ we can find two distinct (but probably nearby) bids $\mathbf{b}$ that jointly occur with non-zero and finite probability. The second part then requires that variation in stocks causes variation in bids for each pair of possible stocks. This assumption will hold if, for example, the derivative of the pseudo-static payoff function k is globally invertible, as assumed in the text.[5]

Assumption 15. *Perfect Substitutes and Normalisation:*

$i)$ $f_{\mathbf{s}_t|\mathbf{w}_{t-1},\mathbf{s}_{t-1}} = f_{\mathbf{s}_t|\mathbf{w}_{t-1}+\mathbf{s}_{t-1}}$

$ii)$ $E[\mathbf{s}_t] = 0$

[5]This is because invertibility of $\nabla_{\mathbf{s}} k(\mathbf{s})$, which (combined with the assumption that $\mathbf{z}$ has full rank) ensures the 2^L vector of potential payoffs $K(\mathbf{s})$ is globally invertible. Optimising behaviour (assuming no reservation prices) ensures that bids must satisfy $\mathbf{b} + \nabla\Gamma(\mathbf{b})^{-1}\Gamma(\mathbf{b}) = \nabla\Gamma(\mathbf{b})^{-1}\nabla P(\mathbf{b})K(\mathbf{s}) + \boldsymbol{v}$. Because $\boldsymbol{v}$ is independent of $\mathbf{s}$, the object $\mathbf{b} + \nabla\Gamma(\mathbf{b})^{-1}\Gamma(\mathbf{b})$ is invertible if and only if $\nabla\Gamma(\mathbf{b})^{-1}\nabla P(\mathbf{b})$ has rank $|\mathbf{s}|$, from the inverse function theorem. Finally, invertibility of $\mathbf{b}$ follows from monotonicity of $\nabla\Gamma(\mathbf{b})^{-1}\Gamma(\mathbf{b})$ in each dimension l.

The first part of this assumption states that past winnings and past stocks are perfect substitutes.[6] Normalisations are necessary in non-parametric identification arguments for latent states, since the model will only ever be identified up to one to one transformations of the unobserved states (as the eigenfunctions recovered from a spectral decomposition are only ever identified up to one to one transformation). Therefore, we need additional assumptions to pin down the values of the unobserved state. We need part ii), the location normalisation, because of the additivity assumption — we can add a constant to both $\mathbf{s}_t$ and $\mathbf{s}_{t-1}$ without changing the observable implications.[7]

One final remark: This perfect substitutes assumption also allows us to fully test Assumptions 13 and 14. This is because invertibility and monotonicity in $\mathbf{s}_t$ can be tested by testing for invertibility and monotonicity in $\mathbf{w}_t$. That is, if variation in $\mathbf{w}_t$ creates enough variation in $\mathbf{b}_{t+1}$ for identification, then because Assumption 15 imposes that $\mathbf{w}_t$ and $\mathbf{s}_t$ have the same effect on $\mathbf{b}_{t+1}$, this must also be invertible in $\mathbf{s}_t$. This is essentially the story shown in Figure 2.11. Both $\mathbf{w}_t$ and $\mathbf{w}_{t-1}$ (conditional on $\mathbf{w}_t$ and $\mathbf{b}_t$) cause statistically significant variation in $\mathbf{b}_{t+1}$.

I.2 Proof of Proposition 5

I now prove the following proposition:

Proposition 16. *Non-parametric Identification Under Assumptions 12 - 15, $f_{\mathbf{b}_t|\mathbf{s}_t}$ and $f_{\mathbf{s}_t|\mathbf{w}_{t-1},\mathbf{s}_{t-1}}$ are non-parametrically identified.*

The proof proceeds along the same lines as Hu and Shum, 2012. First, I prove that if both $f_{\mathbf{b}_{t+1}|\mathbf{w}_t,\mathbf{s}_t}$ and $f_{\mathbf{b}_t|\mathbf{s}_t}$ are identified, then so is $f_{\mathbf{s}_t|\mathbf{w}_{t-1},\mathbf{s}_{t-1}}$. Next, I prove that $f_{\mathbf{b}_{t+1}|\mathbf{w}_t,\mathbf{s}_t}$ and $f_{\mathbf{b}_t|\mathbf{s}_t}$ are identified using a spectral decomposition type argument, in which $f_{\mathbf{b}_t|\mathbf{s}_t}$ and $f_{\mathbf{b}_{t+1}|\mathbf{w}_t,\mathbf{s}_t}$ make up the eigenvalues and eigenfunctions respectively

[6] If the perfect substitutes assumption is unpalatable, as it should be in some applications, we can use a less restrictive separability assumption that $f_{\mathbf{s}_t|\mathbf{w}_{t-1},\mathbf{s}_{t-1}} = f_{\mathbf{s}_t|h(\mathbf{w}_{t-1},\mathbf{s}_{t-1})}$ where h is a known function. Otherwise, we can instead impose Assumption 4 as given in Hu and Shum (2012). For this reason, in the proof presented in Section I.2 I continue to work with $f_{\mathbf{s}_t|\mathbf{w}_{t-1},\mathbf{s}_{t-1}}$

[7] We do not need scale, or other monotonic normalisations because of how the value of stocks is pegged to $\mathbf{w}$. If we were to rescale stocks, multiplying $\mathbf{s}_t$ and $\mathbf{s}_{t-1}$ by the same large constant, then $\mathbf{w}_{t-1}$ will have a (relatively) much smaller effect on $\mathbf{b}_t$, and so have observable implications. The idea then, is that the scale (and higher order moments) of the transition process are identified by the relative importance of $\mathbf{w}_{t-1}$ in determining future bidding behaviour. In the text, I impose this location normalisation through my informative priors.

of some known linear operation on $f_{\mathbf{b}_{t+1},\mathbf{w}_t,\mathbf{b}_t,\mathbf{w}_{t-1}}$. That is, the density $f_{\mathbf{b}_{t+1},\mathbf{w}_t,\mathbf{b}_t,\mathbf{w}_{t-1}}$ completely determines the densities $f_{\mathbf{b}_t|\mathbf{s}_t}$ and $f_{\mathbf{b}_{t+1}|\mathbf{w}_t,\mathbf{s}_t}$.

I.2.1 Identification of $f_{\mathbf{s}_t|\mathbf{w}_{t-1},\mathbf{s}_{t-1}}$ given $f_{\mathbf{b}_{t+1}|\mathbf{w}_t,\mathbf{s}_t}$ and $f_{\mathbf{b}_t|\mathbf{s}_t}$

Lemma I.2.1. *If both $f_{\mathbf{b}_{t+1}|\mathbf{w}_t,\mathbf{s}_t}$ and $f_{\mathbf{b}_t|\mathbf{s}_t}$ are identified, then so is $f_{\mathbf{s}_t|\mathbf{w}_{t-1},\mathbf{s}_{t-1}}$.*

Proof:

1. By Assumption 12, $f_{\mathbf{b}_{t+1}|\mathbf{w}_t,\mathbf{s}_t} = \int f_{\mathbf{b}_{t+1}|\mathbf{s}_{t+1}} f_{\mathbf{s}_{t+1}|\mathbf{w}_t,\mathbf{s}_t} d\mathbf{s}_{t+1}$
2. In operator notation, for a fixed value of $\mathbf{w}_t \in \mathbb{W}$, this can be written as:
$$L_{\mathbf{b}_{t+1}|\mathbf{w}_t,\mathbf{s}_t} = L_{\mathbf{b}_{t+1}|\mathbf{s}_{t+1}} L_{\mathbf{s}_{t+1}|\mathbf{w}_t,\mathbf{s}_t}$$
3. By Assumption 13 *ii*), the inverse operator $L^{-1}_{\mathbf{b}_{t+1}|\mathbf{w}_t,\mathbf{s}_t}$ exists. Implicitly, this also ensures that $|\mathbf{b}| \geq |\mathbf{s}|$, which ensures that $L_{\mathbf{b}_{t+1}|\mathbf{s}_{t+1}}$ must also be invertible.
4. Therefore we can write $L_{\mathbf{s}_{t+1}|\mathbf{w}_t,\mathbf{s}_t} = L^{-1}_{\mathbf{b}_{t+1}|\mathbf{s}_{t+1}} L_{\mathbf{b}_{t+1}|\mathbf{w}_t,\mathbf{s}_t}$, and so identification of the transition process $f_{\mathbf{s}_t|\mathbf{w}_{t-1},\mathbf{s}_{t-1}} (= f_{\mathbf{s}_{t+1}|\mathbf{w}_t,\mathbf{s}_t})$ follows from the identification of $f_{\mathbf{b}_{t+1}|\mathbf{w}_t,\mathbf{s}_t}$ and $f_{\mathbf{b}_t|\mathbf{s}_t}$.

□

I.2.2 Identification of $f_{\mathbf{b}_{t+1}|\mathbf{w}_t,\mathbf{s}_t}$ and $f_{\mathbf{b}_t|\mathbf{s}_t}$ by spectral decomposition

Lemma I.2.2. *The density $f_{\mathbf{b}_{t+1},\mathbf{w}_t,\mathbf{b}_t,\mathbf{w}_{t-1}}$ completely determines the densities $f_{\mathbf{b}_t|\mathbf{s}_t}$ and $f_{\mathbf{b}_{t+1}|\mathbf{w}_t,\mathbf{s}_t}$.*

Proof:

1. By Assumption 12 $f_{\mathbf{b}_{t+1},\mathbf{w}_t,\mathbf{b}_t,\mathbf{w}_{t-1}} = \int f_{\mathbf{b}_{t+1}|\mathbf{w}_t,\mathbf{s}_t} f_{\mathbf{w}_t,\mathbf{b}_t|\mathbf{s}_t} f_{\mathbf{s}_t,\mathbf{w}_{t-1}} d\mathbf{s}_t$
2. In Operator notation, for fixed $(\mathbf{w}_t, \mathbf{b}_t) \in \mathbb{W} \times \mathbb{B}$, this can be written as:
$$L_{\mathbf{b}_{t+1},\mathbf{w}_t,\mathbf{b}_t,\mathbf{w}_{t-1}} = L_{\mathbf{b}_{t+1}|\mathbf{w}_t,\mathbf{s}_t} D_{\mathbf{w}_t,\mathbf{b}_t|\mathbf{s}_t} L_{\mathbf{s}_t,\mathbf{w}_{t-1}} \tag{I.1}$$
3. By Assumption 13, there exists a neighbourhood near this fixed $\mathbf{b}_t$ (and $\mathbf{w}_t$), $\bar{\mathbb{B}}$, such that for all $\bar{\mathbf{b}}_t \in \bar{\mathbb{B}}$, $L_{\mathbf{b}_{t+1},\mathbf{w}_t,\bar{\mathbf{b}}_t,\mathbf{w}_{t-1}}$ is invertible, where $L_{\mathbf{b}_{t+1},\mathbf{w}_t,\bar{\mathbf{b}}_t,\mathbf{w}_{t-1}} = L_{\mathbf{b}_{t+1}|\mathbf{w}_t,\mathbf{s}_t} D_{\mathbf{w}_t,\bar{\mathbf{b}}_t|\mathbf{s}_t} L_{\mathbf{s}_t,\mathbf{w}_{t-1}}$.

4. Therefore, employing Assumption 13 *ii*) and 14 (which also ensures $D_{\mathbf{w}_t,\bar{\mathbf{b}}_t|\mathbf{s}_t}$ is invertible) we get:

$$L_{\mathbf{s}_t,\mathbf{w}_{t-1}} = D^{-1}_{\mathbf{w}_t,\bar{\mathbf{b}}_t|\mathbf{s}_t} L^{-1}_{\mathbf{b}_{t+1}|\mathbf{w}_t,\mathbf{s}_t} L_{\mathbf{b}_{t+1},\mathbf{w}_t,\bar{\mathbf{b}}_t,\mathbf{w}_{t-1}} \tag{I.2}$$

5. Substituting Equation I.2 into I.1 yields:

$$L_{\mathbf{b}_{t+1},\mathbf{w}_t,\mathbf{b}_t,\mathbf{w}_{t-1}} = L_{\mathbf{b}_{t+1}|\mathbf{w}_t,\mathbf{s}_t} D_{\mathbf{w}_t,\mathbf{b}_t|\mathbf{s}_t} D^{-1}_{\mathbf{w}_t,\bar{\mathbf{b}}_t|\mathbf{s}_t} L^{-1}_{\mathbf{b}_{t+1}|\mathbf{w}_t,\mathbf{s}_t} L_{\mathbf{b}_{t+1},\mathbf{w}_t,\bar{\mathbf{b}}_t,\mathbf{w}_{t-1}}$$

6. And so, by Assumption 13 *i*), we can write:

$$L_{\mathbf{b}_{t+1},\mathbf{w}_t,\mathbf{b}_t,\mathbf{w}_{t-1}} L^{-1}_{\mathbf{b}_{t+1},\mathbf{w}_t,\bar{\mathbf{b}}_t,\mathbf{w}_{t-1}} = L_{\mathbf{b}_{t+1}|\mathbf{w}_t,\mathbf{s}_t} D_{\mathbf{w}_t,\mathbf{b}_t,\bar{\mathbf{b}}_t,\mathbf{s}_t} L^{-1}_{\mathbf{b}_{t+1}|\mathbf{w}_t,\mathbf{s}_t} \tag{I.3}$$

   Where the diagonal operator

$$(D_{\mathbf{w}_t,\mathbf{b}_t,\bar{\mathbf{b}}_t,\mathbf{s}_t} h)(\mathbf{s}_t) = (D_{\mathbf{w}_t,\mathbf{b}_t|\mathbf{s}_t} D^{-1}_{\mathbf{w}_t,\bar{\mathbf{b}}_t|\mathbf{s}_t} h)(\mathbf{s}_t) = \frac{f_{\mathbf{w}_t|\mathbf{b}_t}(\mathbf{w}_t|\mathbf{b}_t) f_{\mathbf{b}_t|\mathbf{s}_t}(\mathbf{b}_t|\mathbf{s}_t)}{f_{\mathbf{w}_t|\mathbf{b}_t}(\mathbf{w}_t|\bar{\mathbf{b}}_t) f_{\mathbf{b}_t|\mathbf{s}_t}(\bar{\mathbf{b}}_t|\mathbf{s}_t)} h(\mathbf{s}_t)$$

7. Equation I.3 states that the left hand side has an eigenvalue-eigenfunction decomposition given by the right hand side equation, which can be found using a spectral decomposition. I now argue that this decomposition is unique. Assumption 14 *i*) ensures the eigenvalues are bounded, and so the operator on the left hand side is similarly bounded (linear operators are bounded by their largest eignevalue). This ensures we can apply Theorem XV.4.3.5 from Dunford and Schwartz (1971).[8] Likewise 14 *ii*) ensures that eigenvalues for different values of $\mathbf{s}_t$ are all distinct (for some $\mathbf{w}_t \in \mathbb{W}$ and $\mathbf{b}_t \neq \bar{\mathbf{b}}_t \in \mathbb{B}$).

8. Finally, both eigenvalues and eigenfunctions are unique up to invertible transformations of $\mathbf{s}_t$, so we need 15 to pin values of $\mathbf{s}_t$ to observables $\mathbf{w}_t$. I prove that this normalisation is sufficient by contradiction: Suppose

[8] This theorem only ensures uniqueness of the eigenfunctions up to a scalar multiple. The requirement that the eigenfunctions are also proper densities that integrate to one pin down the precise value of the scalars.

there exists an invertible function $g : \mathbb{R}^{|s|} \to \mathbb{R}^{|s|}$ such that

$$f(\mathbf{b}_{t+1}|\mathbf{w}_t + \mathbf{s}_t) = f(\mathbf{b}_{t+1}|\mathbf{w}_t + g(\mathbf{s}_t))$$

For any $\mathbf{b}_{t+1}$, $\mathbf{s}_t$, $\mathbf{w}_t$ it must be that for any $\mathbf{x}$ we also have:

$$\begin{aligned} f(\mathbf{b}_{t+1}|\mathbf{w}_t + \mathbf{s}_t) =& f(\mathbf{b}_{t+1}|(\mathbf{w}_t - \mathbf{x}) + (\mathbf{s}_t + \mathbf{x})) \\ =& f(\mathbf{b}_{t+1}|(\mathbf{w}_t - \mathbf{x}) + g(\mathbf{s}_t + \mathbf{x})) \end{aligned}$$

This is the crux of the perfect substitutes assumption. Because g is invertible, the only function that satisfies perfect substitutes is $g(\mathbf{s}) = \mathbf{s} + \boldsymbol{\mu}$, so that stocks are identified up to location. Then, imposing that stocks have long run mean of zero ensures we can write $0 = E[\mathbf{s}] = E[g(\mathbf{s})] = E[\mathbf{s} + \boldsymbol{\mu}] = E[\mathbf{s}] + \boldsymbol{\mu}$ which holds only for $\boldsymbol{\mu} = 0$.

9. Therefore, the density $f_{\mathbf{b}_{t+1}|\mathbf{w}_t,\mathbf{s}_t}$ and the relative density $\frac{f_{\mathbf{w}_t|\mathbf{b}_t}(\mathbf{w}_t|\mathbf{b}_t) f_{\mathbf{b}_t|\mathbf{s}_t}(\mathbf{b}_t|\mathbf{s}_t)}{f_{\mathbf{w}_t|\mathbf{b}_t}(\mathbf{w}_t|\bar{\mathbf{b}}_t) f_{\mathbf{b}_t|\mathbf{s}_t}(\bar{\mathbf{b}}_t|\mathbf{s}_t)}$ are identified. Because $\mathbf{w}$ and $\mathbf{b}$ are observed, $f_{\mathbf{w}_t|\mathbf{b}_t}$ and $f_{\mathbf{w}_t|\bar{\mathbf{b}}_t}$ are already identified, ensuring that the ratio $\frac{f_{\mathbf{b}_t|\mathbf{s}_t}(\mathbf{b}_t|\mathbf{s}_t)}{f_{\mathbf{b}_t|\mathbf{s}_t}(\bar{\mathbf{b}}_t|\mathbf{s}_t)}$ is identified. We then separately identify $f_{b|s}$ by requiring that densities integrate to 1, integrating over $\mathbf{b}_t$.

$\square$

Appendix J

Proof of Proposition 6.

In this Appendix I prove Proposition 6. This essentially just extends Arcidiacono and Miller, 2011 to the continuous choice case. Meanwhile, the form for the expected payoff follows directly from Propositon 3 in Chapter 3.

The proposition to be proven, as stated in the main text, is as follows:

Proposition 6. *The ex-ante Value Function can be expressed as:*

$$E[W(\boldsymbol{v}_{it}, \mathbf{s}_i, \mathbf{s}_0)|\mathbf{s}_i, \mathbf{s}_0] = \frac{E[q_t(\mathbf{s}_i^g)\pi(\mathbf{b}_{it}, \mathbf{d}_{it}|\mathbf{s}_i^g, \mathbf{s}_0)|\mathbf{s}_0]}{E[q_t(\mathbf{s}_i^g)|\mathbf{s}_0]}$$

Where $q_t(\mathbf{s}_i^g)$ gives the posterior probability that $\mathbf{s}_{it}^g = \mathbf{s}_i^g$ and

$$\pi(\mathbf{b}, \mathbf{d}|\mathbf{s}_i^g, \mathbf{s}_0) = \sum_l \lambda_i \frac{\Gamma_l(b_{il}, d_{il}; \mathbf{s}_0)^2}{\nabla_b \Gamma_l(b_{il}, d_{il}; \mathbf{s}_0)} - \sum_{m \neq l} \Gamma_l(b_{il}, d_{il}; \mathbf{s}_0) \mathbf{z}_l^{gT} \Psi_i \mathbf{z}_m^g \Gamma_m(b_{im}, d_{im}; \mathbf{s}_0) + \mathbf{s}_i^{gT} \Psi_i \mathbf{s}_i^g$$

In the presence of reservation prices and endogenous entry, however, this final line should be written as:

$$\pi(\mathbf{b}, \mathbf{d}|\mathbf{s}_i^g, \mathbf{s}_0) = \mathbf{s}_i^{gT} \Psi_i \mathbf{s}_i^g + \sum_l \mathbb{I}[b_l > R_l] \left(\lambda \frac{\Gamma_l(b_l, d_l)^2}{\nabla_b \Gamma_l(b_l, d_l)} - \sum_{m \neq l} \Gamma_l(b_l, d_l) \mathbf{z}_l^{gT} \Psi_i \mathbf{z}_m^g \Gamma_m(b_m, d_m) \right)$$
$$+ \mathbb{I}[b_l = R_l] \left(\Gamma_l(R_l, 1)(E[v_l|b_l = R_l, \mathbf{b}_{-l}] - \lambda R_l + \Phi \mathbf{z}_l^h + \mathbf{z}_l^{gT} \Psi[\mathbf{z}_l^g + 2\mathbf{s}^g + \sum_{m \neq l} \Gamma_m(b_m, d_m) \mathbf{z}_m^g)] \right)$$

Which can be derived by applying the result presented in the extension of Proposition 3 to allow for reservation prices in Appendix E.2, employing the identity derived in Appendix C.3.

To simplify notation, I drop the i subscripts and dependence on the observe state $\mathbf{s}^0$. This can be trivially introduced by multiplying objects by $\mathbb{I}[\mathbf{s}_t^0 = \mathbf{s}^0]$. I also drop dependence on the observed discrete action $\mathbf{d}$, also trivially introduced by multiplying objects by $\mathbb{I}[\mathbf{d}_t = \mathbf{d}]$ and summing over possible actions, just as in the discrete choice case. I now prove the following:

$$E_{v_t}[W(v_t, \mathbf{s})|\mathbf{s}] = \frac{E_{\mathbf{b}_t}[q_t(\mathbf{s})\pi(\mathbf{b}_t|\mathbf{s})]}{E[q_t(\mathbf{s})]} \tag{J.1}$$

Where $q_t(\mathbf{s}) = f_{\mathbf{s}_t}(\mathbf{s}|\mathbb{O}_T)$ gives the posterior density, that the unobserved state is $\mathbf{s}$ at time t.

The proof makes use of the Dirac Delta function, defined for continuous random variable $\mathbf{B}$ with frequency density $f_{\mathbf{B}}$ such that $E_{\mathbf{B}}[\delta(\mathbf{B}-\mathbf{b})] = f_{\mathbf{B}}(\mathbf{b})$ and with the property that $\int_{\mathbf{B}} \delta(\mathbf{B}-\mathbf{b})d\mathbf{B} = 1$. I also make use of the fact that $\delta((\mathbf{B},\mathbf{S}) - (\mathbf{b},\mathbf{s})) = \delta(\mathbf{B}-\mathbf{b})\delta(\mathbf{S}-\mathbf{s})$.

Proof: 1. First, I prove that $f_{\mathbf{b}_t}(\mathbf{b}|\mathbf{s}) = \frac{E_{\mathbb{O}_T}[\delta(\mathbf{b}_t-\mathbf{b})|q_t(\mathbf{s})]}{E_{\mathbb{O}_T}[q_t(\mathbf{s})]}$:

$$\begin{aligned}
f_{\mathbf{b}_t}(\mathbf{b}|\mathbf{s}) &= \frac{f_{\mathbf{b}_t,\mathbf{s}_t}(\mathbf{b},\mathbf{s})}{f_{\mathbf{s}_t}(\mathbf{s})} && \text{Bayes' rule}\\
&= \frac{E_{\mathbf{b}_t,\mathbf{s}_t}[\delta((\mathbf{b}_t,\mathbf{s}_t)-(\mathbf{b},\mathbf{s}))]}{E_{\mathbf{s}_t}[\delta(\mathbf{s}_t-\mathbf{s})]} = \frac{E_{\mathbf{b}_t,\mathbf{s}_t}[\delta(\mathbf{b}_t-\mathbf{b})\delta(\mathbf{s}_t-\mathbf{s})]}{E_{\mathbf{s}_t}[\delta(\mathbf{s}_t-\mathbf{s})]} && \text{Definition of } \delta\\
&= \frac{E_{\mathbb{O}_T}[E_{\mathbf{b}_t,\mathbf{s}_t}[\delta(\mathbf{b}_t-\mathbf{b})\delta(\mathbf{s}_t-\mathbf{s})|\mathbb{O}_T]]}{E_{\mathbb{O}_T}[E_{\mathbf{s}_t}[\delta(\mathbf{s}_t-\mathbf{s})|\mathbb{O}_T]]} && \text{Law of Iterated Expectations}\\
&= \frac{E_{\mathbb{O}_T}[\delta(\mathbf{b}_t-\mathbf{b})E_{\mathbf{s}_t}[\delta(\mathbf{s}_t-\mathbf{s})|\mathbb{O}_T]]}{E_{\mathbb{O}_T}[E_{\mathbf{s}_t}[\delta(\mathbf{s}_t-\mathbf{s})|\mathbb{O}_T]]} && \text{as } \mathbf{b}_t \text{ is part of } \mathbb{O}_T\\
&= \frac{E_{\mathbb{O}_T}[\delta(\mathbf{b}_t-\mathbf{b})q_t(\mathbf{s})]}{E_{\mathbb{O}_T}[q_t(\mathbf{s})]} && \text{Definition of } q
\end{aligned} \tag{J.2}$$

2. Next recognise that we can write $E_{v_t}[W(v_t,\mathbf{s})|\mathbf{s}] = E_{v_t}[\pi(\mathbf{b}(v_t;\mathbf{s}),\mathbf{s})|\mathbf{s}]$ as $\mathbf{b}$ is set to maximise the period payoff, given v_t and $\mathbf{s}$. Here, π is just some known function.

3. Applying a change of variables (the law of the unconscious statistician) ensures this equals $E_{\mathbf{b}_t}[\pi(\mathbf{b}_t,\mathbf{s})|\mathbf{s}]$. This requires that the mapping $\mathbf{b}_t = \mathbf{b}(v_t;\mathbf{s})$ is monotonic (has a positive definite jacobian). This result is proven in Appendix C.4.

4. Applying the result from step 1.:

$$E_{\mathbf{b}_t}[\pi(\mathbf{b}_t,\mathbf{s})|\mathbf{s}] = \int_{\mathbf{b}} \pi(\mathbf{b},\mathbf{s})f_{\mathbf{b}_t}(\mathbf{b}|\mathbf{s})d\mathbf{b} = \int_{\mathbf{b}} \pi(\mathbf{b},\mathbf{s})\frac{E_{\mathbb{O}_T}[\delta(\mathbf{b}_t-\mathbf{b})q_t(\mathbf{s})]}{E_{\mathbb{O}_T}[q_t(\mathbf{s})]}d\mathbf{b}$$

5. Recognise that the denominator is not a function of the random variable $\mathbf{b}$, so we can pull it out of the integral. Then, move $\pi(\mathbf{b}, \mathbf{s})$ into the expectation for:

$$= \frac{\int_{\mathbf{b}} \pi(\mathbf{b}, \mathbf{s}) E_{\mathbb{O}_T}[\delta(\mathbf{b}_t - \mathbf{b}) q_t(\mathbf{s})] d\mathbf{b}}{E_{\mathbb{O}_T}[q_t(\mathbf{s})]} = \frac{\int_{\mathbf{b}} E_{\mathbb{O}_T}[\pi(\mathbf{b}, \mathbf{s}) \delta(\mathbf{b}_t - \mathbf{b}) q_t(\mathbf{s})] d\mathbf{b}}{E_{\mathbb{O}_T}[q_t(\mathbf{s})]}$$

6. From the definition of the delta function we recognise that the expectation equals zero for $\mathbf{b} \neq \mathbf{b}_t$, so that I can replace $\pi(\mathbf{b}, \mathbf{s})$ with $\pi(\mathbf{b}_t, \mathbf{s})$. Then, swap the order of integration, moving the integral into the expectation for:

$$= \frac{\int_{\mathbf{b}} E_{\mathbb{O}_T}[\pi(\mathbf{b}_t, \mathbf{s}) \delta(\mathbf{b}_t - \mathbf{b}) q_t(\mathbf{s})] d\mathbf{b}}{E_{\mathbb{O}_T}[q_t(\mathbf{s})]} = \frac{E_{\mathbb{O}_T}[\int_{\mathbf{b}} \pi(\mathbf{b}_t, \mathbf{s}) \delta(\mathbf{b}_t - \mathbf{b}) q_t(\mathbf{s}) d\mathbf{b}]}{E_{\mathbb{O}_T}[q_t(\mathbf{s})]}$$

7. Within the expectation, $\mathbf{b}_t$ and $\mathbf{s}$ are constant, so pull $\pi(\mathbf{b}_t, \mathbf{s}) q_t(\mathbf{s})$ out of the integral, before applying the definition of the delta function:

$$= \frac{E_{\mathbb{O}_T}[\int_{\mathbf{b}} \delta(\mathbf{b}_t - \mathbf{b}) d\mathbf{b}\ \pi(\mathbf{b}_t, \mathbf{s}) q_t(\mathbf{s})]}{E_{\mathbb{O}_T}[q_t(\mathbf{s})]} = \frac{E_{\mathbb{O}_T}[\pi(\mathbf{b}_t, \mathbf{s}) q_t(\mathbf{s})]}{E_{\mathbb{O}_T}[q_t(\mathbf{s})]}$$

$\square$

Appendix K

Estimation Details

In this Appendix I give additional details of the estimation procedure outlined in Section 4.3. I outline my specification of priors, as well as computational details of how each step of the estimation procedure is performed. Appendix K.1 outlines details of the first estimation step, Appendix K.2 the second step, Appendix K.3 the third step, and finally Appendix K.4 details the model specification used for the Type 2 food banks.

K.1 Step 1.

In the first estimation step I estimate food banks beliefs about the probability they win a given lot given their bid. While I assume there is zero probability of ties above the reservation price, I allow for the possibility of ties at the reservation price. I begin by discussing how I conceptualise food banks' beliefs in Appendix K.1.1. I discuss ties in Appendix K.1.2. I then detail my parametrisation in Appendix K.1.3, before discussing how estimation is performed K.1.4.

K.1.1 Maximum Rival Bid

I do not explicitly parameterise bid distributions and use this to form food banks' beliefs about equilibrium win probabilities, as in Jofre-Bonet and Pesendorfer, 2003 or Gentry, Komarova, and Schiraldi, 2023. Instead I take an approach closer to that in Backus and Lewis, 2016 and estimate the distribution of equilibrium winning bids. If the winning bid on auction l at time t was $\bar{b}_{lt}$, then food bank i knows they would have won the lot had they bid $b_{ilt} > \bar{b}_{lt}$. If no food bank placed a bid on lot l then food bank i knows they would have won if they had bid the reservation price. If a food bank won lot l at the reservation price, i knows they would have drawn had they bid the reservation price.

A food bank's ex-ante belief about the probability of winning given a bid is given by $P(\bar{b}_{lt} < b_{ilt}|\mathbf{s}_t)$, which requires the conditional cdf of the random variable $\bar{b}_{lt}$. However, this object is subject to censoring at the reservation price.

K.1.2 Ties

Ties are observed very rarely in the data, in around 0.02% of auctions, and all at the reservation price. Due to the continuity of bids, ties happen above the reservation price with probability zero.[1] However, because winning bids are observed more frequently at the reservation price, in around 20% of auctions, food banks must consider the much larger chance of a tie if they bid the reservation price and no higher. Furthermore, food banks appear to recognise this, and often bid just above the reservation price. This means we get high density of winning bids just above the reservation price which food banks presumably also recognise, and so must be taken into account in the model.

The bidder wins lot l given bid b_{ilt} if $b_{ilt} > \bar{b}_{lt}$. If $b_{ilt} = \bar{b}_{lt}$ they win with probability 0.5.[2] Like i's bids, $\bar{b}_{lt}$ is censored both at R_l (when the maximum rival bid equals the reservation price) and below it (when no rivals place bids). Therefore I introduce the latent random variable $\bar{b}^*_{lt}$, with cdf $G_l(b^*|\mathbf{s}_{0t})$, such that:

$$\bar{b}_{lt} = \begin{cases} \bar{b}^*_{lt} & \text{if } \bar{b}^*_{lt} > R_{lt} \\ R_l + \epsilon_{lt} & \text{if } \bar{b}^*_{lt} \in [R_l, \bar{R}_l) \\ R_l & \text{if } \bar{b}^*_{lt} \in [\bar{R}_l, \underline{R}_l) \\ \emptyset & \text{if } \bar{b}^*_{lt} \leq \underline{R}_l \end{cases}$$

Where $(\bar{R}_l, \underline{R}_l)$ are a category specific cutoff to be estimated. This is not dissimilar to cutoffs estimated in an ordered logit model, enabling me to capture the varying

[1] In practice bids must be integer amounts, but because bids tend to range between −2000 and 4000 I treat this as continuous.

[2] In other words I assume they tie with at most one other bidder, and the tie is broken with the flip of a coin. In practice, bidders believe that ties only occur with more than one bidder at the rate that ties are observed in the data, in 0.02% of auctions, which I deem as negligible.

likelihood of winning bids at the reservation price across categories. This latent variable structure states that if the 'true', latent, winning bid $\bar{b}^*_{lt}$ is extremely low ($\leq \underline{R}_l$), then a competing food bank would win if it bid the reservation price. If it is somewhat higher $\bar{b}^*_{lt} \in [\bar{R}_l, \underline{R}_l)$ then the observed winning bid is just the reservation price - a competing food bank would draw if it bid the reservation price. The competing food bank may not value the lot enough to bid much above the reservation price, but may be willing to bid just one or two additional shares to ensure it doesn't risk a tie. Finally, if $\bar{b}^*_{lt}$ is just below the reservation price $\bar{b}^*_{lt} \in [R_l, \bar{R}_l)$ then the observed winning bid is actually just above the reservation price, where $\epsilon_{lt} \sim exponential(\alpha)$ and α is a parameter to be estimated. This means that a competing food bank must take into account the excess mass just above the reservation price - if it bids just one share above the reservation price it may lose out to equally strategic food banks.

This modelling approach is unusual, but enables the model to rationalise both the excess mass of winning bids at the reservation price, and also just above it. I assume that food banks do not internalise the probability of tieing at just one share above the reservation price (and likewise two, three, etc shares). Importantly, $\underline{R}_l$ is identified by the excess mass of winning bids at the reservation price (and how this varies across categories). $\bar{R}_l$ is identified from the excess mass just above the reservation price, and α is identified from how the excess mass diminishes as we move further from the reservation price.

Given the distribution of $\bar{b}^*_{lt}$, and implied distribution of $\bar{b}_{lt}$, Food bank i's beliefs are given by:

$$P(i \text{ wins } l | b_{ilt}; \mathbf{s}_{0t}) = \Gamma_l(b_{ilt}|\mathbf{s}_{0t}) = \begin{cases} G_l(b_{ilt}|\mathbf{s}_{0t}) - f(b_{ilt}) & \text{if } \ b_{ilt} > R_{lt} \\ \frac{1}{2}G_l(\underline{R}^c|\mathbf{s}_{0t}) + \frac{1}{2}G_l(\bar{R}^c|\mathbf{s}_{0t}) & \text{if } \ b_{ilt} = R_{lt} \\ 0 & \text{otherwise} \end{cases} \tag{K.1}$$

Where $f(b_{ilt}) = [G_l(R_l|\mathbf{s}_{0t}) - G_l(\bar{R}^c|\mathbf{s}_{0t})]e^{-\alpha b_{ilt}}$ capture the probability that i loses out to a food bank bidding just above the reservation price.

K.1.3 Parameterisation

I normalise all bids by the reservation price, so that the estimated distributions can be considered the distribution of the difference between the winning bid and the reservation price. Therefore from here on, we can replace $\bar{b}^*_{lt}$ with $\bar{b}^*_{lt} - R_l$. Reservation prices are known to be -2000 for all lots except for fresh produce and Maroon lots which have $R_l = 0$.

As in Assumption 9 I assume that the distribution of $\bar{b}^*_{lt}$ is a function of $\boldsymbol{\vartheta}(\{\mathbf{s}_i\}_N)$, aggregate statistics of states only, so that it does not depend on the states of each individual food bank. In particular, I assume that it only depends on the previous 30 day supply of food from each storage type, as well as the supply of food allocated at time t from each storage type. This is intended to capture how prices vary with supply. This also ensures that food banks do not need to take into account exactly which food bank wins which combination of lots each period. Instead, i only needs to consider which lots they themselves win.

I assume that the latent random variable $\bar{b}^*_{lt}$, on lot l given common state $\mathbf{s}_{0t}$, follows a generalised extreme value distribution, with:

$$P(\bar{b}^*_{lt} \leq b|\mathbf{s}_{0t}) = exp(-t(\frac{b - \nu(\mathbf{s}_{0t})}{\zeta(\mathbf{s}_{0t})})) \quad \text{Where}: \ t(x) = \begin{cases} (1 + \xi(\mathbf{s}_{0t})x)^{-\frac{1}{\xi(\mathbf{s}_{0t})}} & \text{if} \ \ \xi(\mathbf{s}_{0t}) \neq 0 \\ exp(-x) & \text{if} \ \ \xi(\mathbf{s}_{0t}) = 0 \end{cases}$$

The shape parameters ξ are category specific for categories with at least 500 loads, and the remainder are constrained to be equal to one another. Shape parameters are constrained to be > -1 to ensure bids are monotonic in values. This constraint does not bind.

The scale parameters ζ are also category specific. In addition, within a category if the subcategory is listed as “unspecified", “mixed" or “miscellaneous" these receive an additional fixed effect on their scale parameter. This is to allow me to capture additional variation due to uncertainty over the goods included in the lot. I constrain scale parameters to be strictly positive. Finally, I also allow the scale parameter for lot l to vary depending on whether the lot has also been auctioned in a previous period. If this is the case bidders gain information about rival bidders values for this lot, making it intuitive that the variance of rival bids is expected to decrease.

Each lot can contain up to four distinct categories, subcategories and storage types. Therefore, for both the shape and scale parameters, if the lot contains a mixture of categories, I use an average over the different categories / subcategories.

The location shifter ν varies with both lot specific covariates and the common state variable. I include subcategory fixed effects, as well as dummies for whether the lot includes free delivery, geographic restrictions, any unobserved notes about the lot contents, whether the lot is a "Maroon load" (category specific for categories with at least 50 maroon loads), which US region the lot originates in,[3] is shelf-stable, the number of distinct categories included in the lot, whether the lot has been auctioned previously, and the number of homogenous loads being auctioned simultaneously. It also varies with the log of the sum of the previous 30 day's supply of that type of food, up to $t-1$, by storage type, and the log of the sum of food of that storage type being auctioned simultaneously that day. As with the shape and scale parameters, if the lot includes multiple categories, subcategories, or storage types, I use the average location shifter.

The threshold cutoffs $\bar{R}_l$ and $\underline{R}_l$ are allowed to vary across categories, but only for categories in which there are at least 100 lots won at the reservation price. This includes Beverages, Cereal, Condiments, Fresh Produce, Meals, and Snacks. The remaining categories are grouped together. The exponential parameter α is constrained to be positive.

K.1.4 Computation

In the first stage I estimate 268 parameters, from 26,000 auctions, dropping the first 60 days to enable construction of the sum of the previous 30 days' supplies.

I initially estimate parameters using maximum likelihood, using the implied distribution of $\bar{b}_{lt}$. I set initial shape parameters to 0, and initial scale parameters to the observed standard deviation of winning bids. I set all the location shifters to zero.

Having maximised the likelihood function, I then draw samples from the posterior distribution using the Metropolis Hastings algorithm. I use the inverse hessian from the maximum likelihood procedure for my proposal variance, and also

[3] Using the 8 US economic regions + Canada

adaptively tune this variance using the procedure of Atchadé and Rosenthal, 2005 to ensure that on average 23.4% of proposed draws are accepted.

K.2 Step 2.

In the second estimation step I estimate the distribution of lot specific values, the distribution of net local donations, and the pseudo-static pay-off function. In Appendix K.2.1 I restate my parametrisation as in the main text. In Appendix K.2.2 I set out my assumptions on prior distributions. In Appendix K.2.3 I discuss the data augmentation algorithm, and in Appendix K.2.4 I discuss the gibbs sampling algorithm for drawing parameters from their conditional posteriors. In Appendix K.2.5 I discuss additional computational details. Note that this appendix makes heavy use of the results on the inverse bid system given in Appendix G.

K.2.1 Parametrisation

As stated in Section 4.3.4, the main parametric assumptions are as follows:

$$k(\mathbf{s}_i) = \Phi \mathbf{s}_i^h + \mathbf{s}_i^{g^T} \Psi_i \mathbf{s}_i^g \qquad v_{ilt} \sim N(\alpha_i \text{distance}_{ilt}, \sigma_l^2) \qquad \mathbf{x}_{it} \sim N(\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i)$$

The 164 subcategory parameters of Φ are constant across food banks, and constrained to be positive. This is to allow us to interpret these subcategory weights as the benefit food banks receive from holding the various subcategories to give out to the clients. This constraint doe not bind. Ψ_i is a symmetric matrix that is allowed to vary across food banks, with 15 unique elements for each food bank (5 storage types and 10 interaction terms). I do not impose other constraints on this matrix, given that with probability 1 each draw of the matrix will have full rank.

α_i gives food bank i's cost, in shares, of transporting a lot an additional kilometre. These parameters are constrained to be negative (costs), and this does not bind. I do not allow this cost to vary with either the size of the lot, or the type of food. I allow the lot specific variance to vary depending on the combination of goods auctioned together in the lot. In particular, I find the 60 most common category combinations (e.g. $\frac{2}{3}$ dairy $\frac{1}{3}$ cereal), and associate each combination with a unique variance parameter. I also include an 'other' variance parameter, which covers the remaining 5.5% of combinations. I parametrise the lot specific variances in this manner as it makes the problem of sampling from their posterior significantly easier. These parameters are constrained to be positive, and are not allowed to vary across food banks. The λ_i parameters, which capture the opportunity cost to food bank *i* of spending a share essentially capture variation in the variance of bids across food banks. These parameters are also constrained to be positive, and the parameter for the median consuming Type 1 food bank is constrained to 1.

Finally, for the distribution of net local donations I impose that Σ_i is a diagonal matrix. Informed by the analysis from Appendix F I impose that diagonal entries are strictly above $0.01 Var[\mathbf{w}_{it}^T \mathbf{z}_t]$ and below $100 Var[\mathbf{w}_{it}^T \mathbf{z}_t]$. For all parameters, if not otherwise constrained I impose an upper limit of e^{50} and a lower limit of $-e^{50}$.

K.2.2 Priors and Hierarchical Distributions

Write $\boldsymbol{\psi}_i$ as the vector of unique elements of the matrix Ψ_i. I assume these come from the hierarchical distribution, such that

$$\psi_i \sim N(\boldsymbol{\psi}, \Sigma^{\psi})$$

The hierarchical framework reduces the posterior variance of estimated parameters at a cost of bias, as estimated parameters are drawn together. Observations with a lot of identifying variation place little weight on the hierarchical parameters, whereas observations with little identifying variation place more weight on hierarchical parameters. Any bias caused by this framework causes parameters to be drawn together, so that my estimates will be biased in favour of the Old System rather than the Choice System. I assume weak inverse-Wishart priors for the hierarchical parameters $(\boldsymbol{\psi}, \Sigma^{\psi})$. The prior mean for $\boldsymbol{\psi}$ is -1 for diagonals and zero for off-diagonals of Ψ. The prior mean for Σ^{ψ} is set to be arbitrarily small. The two shape parameters are each set to 2.

I assume independent normal priors for Φ and α_i with means $\frac{1}{40}$ and 1 respectively, and prior variance 10000 to reflect my prior ignorance over these parameters. I assume weak inverse-gamma priors for the lot specific variance, with prior-mean set to the observed variance of bids, and shape parameter set to 2. For the lambda parameters I assume that λ_i^2 takes a prior gamma distribution with shape and rate parameters $T_\lambda^0 = 100$. This ensures λ_i^2 has prior mean of 1, and confidence about this prior mean equal to approximately 1/100th the weight placed on the data.

I assume normal-inverse-gamma priors for $\boldsymbol{\mu}_i$ and Σ_i, with prior means $\boldsymbol{\mu}_i^0 = \frac{1}{T}\sum_{t=1}^T \mathbf{w}_{it}^T \mathbf{z}_{it}$ and $\Sigma_i^0 = \frac{2}{T-1}\sum_{t=1}^T (\mathbf{w}_{it}^T \mathbf{z}_{it} - \frac{1}{T}\sum_{t=1}^T \mathbf{w}_{it}^T \mathbf{z}_{it})^2$ respectively. Given that in my estimation sample I have $T = 1075$, I set prior 'shape' parameters for these distributions to 107, essentially meaning that I place ten times as much weight on the data as I do on my priors.

K.2.3 Data Augmentation step

Given parameters and unobserved states I form the inverse bid system as in Appendix G. For observation ilt such that $b_{ilt} > R_{lt}$ the inverse bid system gives me a conditional observation of v_{ilt}. For observations of bids at or below the reservation price I am only able to bound v_{ilt}. I augment my data by drawing these conditional observations from their conditional posterior, the truncated normal distribution, using the sampling procedure of Botev, 2017. I then revert these inverse bids into 'observations' y_{ilt}, essentially observations of bids from below the reservation price.

To sample the unobserved states from their posteriors I run a standard Kalman filter using the current draw of parameters and the current draw of censored observation. I begin the filter on day 61, as I do not have estimates of beliefs from before this point. I set the initial state to 0, essentially normalising the first set of stocks, with initial variance of zero.

I then run the Carter-Kohn algorithm (Carter and Kohn, 1994) to backwards sample the unobserved states from their conditional posterior. I only run it backwards until day 101, essentially discarding an extra 40 days of the filter. I do this to reduce the reliance on the initial state assumption. This is because even though the initial state is not identified, if there is significant bayesian shrinkage due to the hierarchical model, the initial state actually may be identified.

Finally, I also take into account the observed change in the supply of fresh produce that occurs on day 553 in my sample. Thereafter fresh produce stops being allocated through the Choice System and is instead allocated to food banks outside the system. Each food bank has two separate mean local donations for fresh food, one for before this period and one after this period. I fix the parameter for after period 553 to 0, so that on average food banks give out as much fresh food as they receive in net local donations. Anything else would lead stocks either to trend upwards or downwards indefinitely. Therefore I only estimate the expected net donation for fresh food using sampled states from before the break.

K.2.4 Gibbs Sample step

Given a sample of states $\{\mathbf{s}_{it}\}_{t\in\{1\ldots T\}}$ I back out a sample of net donations by writing the transition equation as a function of $\mathbf{x}_{it}$. Write x_{itm} for the mth element of $\mathbf{x}_{it}$. The conditional posterior distribution is normal-inverse-gamma, and given as:

$$
\begin{aligned}
(\mu_{im}, \Sigma_{imm})|\, \{x_{itm}\}_{t\in\{1\ldots T\}} &\sim N - IG(A, B, C, D) \\
A = \frac{T^0\mu^0_{im} + T\bar{x}_{im}}{T^0 + T} \qquad & B = T^0 + T \qquad C = \frac{T^0 + T}{2} \\
D = \frac{T^0\Sigma^0_{imm}}{2} + \frac{1}{2}\sum_{t=1}^{T}(x_{itm} - \bar{x}_{im})^2 &+ \frac{TT^0}{T + T^0}\frac{(\bar{x}_{im} - \mu^0_{im})^2}{2}) \\
\bar{x}_{im} &= \frac{1}{T}\sum_{t=1}^{T} x_{itm}
\end{aligned}
$$

I then move on to the parameters of k and F^v. I focus on Ψ_i first, rewriting the observation equation (using both sampled states and censored observations) as:

$$
\begin{aligned}
Y_{ilt} = \lambda_i y_{ilt} - \Phi\mathbf{z}^h_{tl} - \alpha_i dist_{ilt} &= \mathbf{z}^{gT}_{tl}\Psi_i(\mathbf{z}^g_{tl} + 2\mathbf{s}^g_{it} + 2\sum_{m\neq l}\Gamma_{im}(b_{itm}\mathbf{z}^g_{tm})) + v_{itl} - \alpha_i dist_{ilt} \\
&= \mathbf{X}_{itl}\boldsymbol{\psi}_i + \varepsilon_{itl}
\end{aligned}
$$

Stacking Y_{itl} and $\mathbf{X}_{itl}$ over tl, the conditional posterior distribution of ψ_i is then multivariate normal, and given as:

$$\boldsymbol{\psi}_i|(\mathbf{Y}_i,\mathbf{X}_i),(\psi,\Sigma^{\psi}) \sim N(M,V)$$
$$M = V^{-1}(\Sigma^{-1\psi}\psi + \mathbf{X}_i^T\mathbf{Y}_i) \qquad \& \qquad V = (\Sigma^{-1\psi} + \mathbf{X}_i^T\mathbf{X}_i)^{-1}$$

The hierarchical parameters (ψ,Σ^{ψ}), which I sample after sampling Ψ_i, have normal-inverse Wishart distribution, with conditional posterior:

$$(\boldsymbol{\psi},\Sigma^{\psi})|\{\boldsymbol{\psi}_i\}_{i\in\{1...N\}} \sim N - IW(A,B,C,D)$$
$$A = \frac{N\bar{\boldsymbol{\psi}} + \boldsymbol{\psi}^0}{N+1} \qquad B = N+1 \qquad C = N+1$$
$$D = \Sigma^{0\psi,-1} + \sum_i(\boldsymbol{\psi}_i - \bar{\boldsymbol{\psi}})(\boldsymbol{\psi}_i - \bar{\boldsymbol{\psi}})^T + \frac{N}{N+1}(\bar{\boldsymbol{\psi}} - \boldsymbol{\psi}^0)(\bar{\boldsymbol{\psi}} - \boldsymbol{\psi}^0)^T)$$
$$\bar{\boldsymbol{\psi}} = \frac{1}{N}\sum_i \boldsymbol{\psi}_i$$

I jointly sample the distance parameters α_i and subcategory weights Φ using standard bayesian regression, given normal priors, Ψ_i, sampled states and censored observations. I sample lot-specific variances just as in bayesian regression, given regression coefficients and $\{\lambda_i\}_{i\in\{1...N\}}$.

Finally, rewriting the observation equation as: $\lambda_i y_{ilt} = \mathbf{Z}_{itl}\delta + \varepsilon_{ilt}$, the conditional posterior pdf of λ_i is proportional to:

$$f(\lambda_i|\mathbf{Z}_i,\delta,\sigma) \quad \propto \quad (\frac{\lambda_i^{LT+2(T_\lambda^0-1)}}{\prod_{lt}\sigma_l\sqrt{2\pi}})exp(-\frac{1}{2}(\lambda_i^2\sum_{tl}\frac{y_{ilt}}{\sigma_{lt}}^2 - \lambda_i\sum_{tl}2\frac{y_{ilt}\mathbf{Z}_{itl}\delta}{\sigma_{lt}^2} + 2T_\lambda^0))$$

I draw samples from this posterior distribution using metropolis hastings. I divide the λ_is by that of the median food bank, ensuring the relevant normalisation.

K.2.5 Computation

I focus on data from only the highest 25 bids placed each day by each food bank. Even type 1 food banks rarely place more than between 5 and 10 bids each day - the 90th percentile food bank only bids on 4 lots each day. However on 50% of days with at least one auction there are more than 25 unique lots being auctioned simultaneously. In principle by ignoring that food banks also choose not to bid on any more than the first 25 lots I may bias my results towards food banks being willing to bid on a higher proportion of auctioned lots than in fact. However, the degree of this bias is unlikely to be large, since I am already taking into account that food banks only bid on maybe the first 10 lots, then choose not to bid on the next 15 lots. Furthermore, to show robustness to this assumption in Appendix M.2.4 I present results from considering 50 unique auctions each day. This assumption is useful in ensuring results converge relatively more quickly, since the higher the degree of censoring, the slower results are expected to converge.

The order of my data augmentation and Gibbs Sampling procedure is as it was presented in the main text. Every tenth iteration I draw a new sample of beliefs using five repetitions of Metropolis Hastings. At the very beginning of the procedure I run the data augmentation step 30 times without running the gibbs sampling step. This is to reduce the sensitivity to the initial draw of augmented data, in which it is assumed that states do not vary at all.

I run the full procedure for 300,000 iterations, and burnout the first 200,000 draws. I run 4 independent chains. For parameters with informative priors, initial points are drawn from the prior distribution. For parameters with diffuse priors I sample uniformly between 0 and 2 $\times$ the prior mean. I uniformly sample 250 points from each of the chains, so that I keep 1,000 parameter draws in total. I then use these parameters in evaluating both the third stage, and the Choice System simulations. I estimate around 1780 parameters across 1.1 million observations observations, 0.95 of which are censored (i.e. a bid is not placed).

K.3 Step 3.

In the third estimation step I evaluate the continuation value as a function of observed bids and the pseudo-static pay-off, before backing out the combination flow pay-off.[4]

[4]One thing to note: The marginal welfare from consuming a lot with subcategory composition $\mathbf{z}^h_{tl}$ is just $\Phi\mathbf{z}^h_{tl}$, and does not depend on $\mathbf{s}^h_{it}$. I do not need to worry about the fact that Φ is a 'pseudo-static' object, not a present discounted sum of expected future flow payoffs from winning $\mathbf{z}^h_{tl}$. The reasons

In Appendix K.3.1 I describe how I form the posterior probabilities that $\mathbf{s}_t = \mathbf{s}$, given by $q_t(\mathbf{s})$. In Appendix K.3.2 I outline how I evaluate the expression for the maximised expected pay-off. In Appendix K.3.3 I discuss how I pool information across food banks in this estimation stage. In Appendix K.3.4 I justify and detail the polynomial approximation used for the ex-ante value function. Finally, in Appendix K.3.5 I describe how I take the expectation of the ex-ante value function over states, yielding the continuation value, before backing out the combination flow payoff.

K.3.1 Posterior Probabilities

Because states are continuous I must evaluate the continuation value over a finite set of states. For each food bank I form a 20^5 dimensional grid of states, so that each dimension of stocks is split into 20 evenly spaced points. For the minimum and maximum points I take the 2.5 and 97.5 percentiles of all their sampled states.

For each of these states I form the posterior probability density that at any given time this was the true state of their stocks, using my $1,000$ draws of states for each time period. I use an independent normal kernel, with Silverman's rule of thumb to calculate bandwidth h:

$$\hat{q}_{it}(\mathbf{s}) = \hat{p}(\mathbf{s}_{it} = \mathbf{s}_i|\text{data}) = \prod_{m=1}^{5} \frac{1}{1000} \sum_{r=1}^{1000} \frac{1}{h_m} \phi\left(\frac{s^r_{itm} - s_{im}}{h_m}\right)$$

K.3.2 Maximised Payoff

I evaluate the maximised pay-off at each time period $\pi(\mathbf{b}_t, \mathbf{d}_t|\mathbf{s})$ using the reservation price adjusted formula discussed in Appendix J. This expression is evaluated once for each parameter draw, using the sample counterparts to the expectation operators given.

In principle I ought to take into account sampling variation in these finite sample expectations. However, given the large number of time periods we expect fairly little variation. One possibility is to use a bootstrap procedure when evaluating

---

for this are simple - when flow payoffs are affine in $\mathbf{s}^h_{it}$, so is the pseudo-static payoff. Furthermore, bidding behaviour does not depend on $\mathbf{s}^h_{it}$, and so future bidding behaviour does not depend on $\mathbf{z}^h_{tl}$. Suppose the flow payoff is given by $\tilde{\Phi}\mathbf{s}^h_{it} + j(\mathbf{s}^g_{it})$, and allow that $\mathbf{s}^h_{it} = \delta\mathbf{s}^h_{it-1} + \mathbf{w}^T_{it-1}\mathbf{z}^h_{t-1} + \mathbf{x}^h_{it}$. In this case, the marginal welfare from winning a lot with subcategory composition $\mathbf{z}^h_{tl}$ (focusing only on the subcategory component) is just $\tilde{\Phi}\sum_{s=0}^{\infty}\beta^s\delta^s\mathbf{z}^h_{tl} = \tilde{\Phi}(I-\beta\delta)^{-1}\mathbf{z}^h_{tl}$. Therefore $\Phi = \tilde{\Phi}(I-\beta\delta)^{-1}$.

these averages to ensure that we introduce sampling variation alongside the variation in parameters from our draws. The difficulty is that this does not account for the correlations between sampled parameters and the sample expectations, so will overestimate posterior variances. This procedure is performed in Appendix M.3.3.

K.3.3 Information Pooling

I also pool information across food banks, giving me additional observations when evaluating the expectation. Similar to the bayesian hierarchical model, this is expected to bias my estimates in favour of the Old System, pushing food banks' flow pay-offs closer together. I use two adjustments to minimise this bias.

When constructing the ex-ante value function for food bank i given parameter draw θ_i, I find the probability density (using the same independent normal kernel as used above) that food bank j draws these parameters given their posterior distribution. The parameters I compare are the estimates for Ψ_i, α_i, and λ_i. I do not need to compare parameters for the net donation process, since the ex-ante value function is evaluated conditional on the state. This density yields a weight for food bank j. I normalise the weights so they sum to 1. When summing posterior probabilities across t I also multiply the probabilities by the associated food bank weight.

Finally, I also use a first-order adjustment to account for different food banks' bids, and hence maximised payoffs, are determined by different parameters. Write:

$$\pi_i(\mathbf{b}_t, \mathbf{d}_t|\mathbf{s}_{it};\theta_i) \approx \pi_j(\mathbf{b}_t, \mathbf{d}_t|\mathbf{s}_{it};\theta_i) + \nabla_{\Psi,\alpha,\lambda}\pi_j(\mathbf{b}_t, \mathbf{d}_t|\mathbf{s}_{it};\theta_i)\begin{pmatrix}\Psi_i - \Psi_j\\ \alpha_i - \alpha_j\\ \lambda_i - \alpha_j\end{pmatrix}$$

Where the derivative $\nabla_{\Psi,\alpha,\lambda}\pi_j(\mathbf{b}_t, \mathbf{d}_t|\mathbf{s}_i;\theta_i)$ comes from differentiating maximised expected pay-off with respect to the parameters, employing the envelope theorem. This should reduce the bias caused by different food banks bidding subject to different parameters.

In Appendix M.3.2 I consider my results differ when I do not pool information across food banks.

K.3.4 Approximation

Having evaluated the ex-ante value function across the grid of states I fit a polynomial function of the states to the ex-ante value function. I include all interaction terms. The fit is performed using a standard weighted least squares procedure, weighting by the sum of posterior probabilities. This is to ensure that state observations that are more likely receive greater weight.

The main version uses a simple quadratic function. This is done primarily because my counterfactuals occasionally require extrapolation (given many of my counterfactual mechanisms do not allow food banks to maintain their balanced level of stocks). A quadratic polynomial has the appealing property that changes in the extrapolated values are constrained to be linear. The difficulty with higher-ordered polynomial (e.g. cubics, quintics etc) is that extrapolated values can be much further from interpolated values.

To validate the quadratic approximation I consider several measures of fit: The R^2 from the regression, as well as the results from considering higher order polynomials (up to order 6). These results are presented in Appendix M.3.1.

K.3.5 Continuation Value

Given the approximated ex-ante value function I evaluate the continuation value by taking an expectation of the polynomial function, given the distribution of $\mathbf{s}_{it+1}$ given $\mathbf{s}_{it}^{a}$. This is done using standard recursive formulae for the higher order moments of normally distributed random variables. With the continuation value in hand I evaluate the combinatorial flow-payoff using the definition of the estimated pseudo-static payoff function. All the above analysis is performed separately for time periods from both before and after period 553, when there is a structural break in the supply of fresh produce. I then average (weighting appropriately) the estimated flow payoffs from either side of the break. In future I will test whether continuation values and hence estimated combinatorial flow-payoffs are constant over the break.

K.4 Type 2 Food Banks

I now discuss the model of Type 2 food banks, the food banks who do not bid, nor win, regularly. This means I do not have significant identifying variation to allow estimation of their model parameters without a large degree of noise. Furthermore, many of these food banks never win certain types of food at all, meaning their parameters are not separately identified.

The Type 2 food banks consist of those food banks who win fewer than 200 lots over the sample period, and excludes the food banks who's locations are unknown or consume fewer than 30 lots over the period (which make up 2.5% of total consumption).

K.4.1 Differences to Type 1s

The key difference is that Type 2 food banks are assumed to be myopic bidders. That is, they are not forward looking. Aside from this assumption, I estimate the model using the same specification and estimation procedure as I used to estimate the pseudo-static payoff function for Type 1 food banks. I continue to recognise that I not observe food banks' stocks, and that variation in stocks is likely to be a key source of variation in bidding behaviour.

Due to the lack of variation in winnings and bids (an even more extreme degree of censoring), I assume that the combinatorial pay-off function for Type 2 food banks comes from the same hierarchical distribution as the pseudo-static payoff from Type 1 food banks. This means Type 2 food banks have the same Φ parameters, and their Ψ_i parameters are drawn from the same hierarchical distribution.[5] I also assume their lot specific values v_{itl} have the same variance as Type 1 food banks. I also assume they have the same beliefs as Type 1 food banks.

Therefore at each iteration of the estimation procedure I do the following: First, given the previous draw of parameters and unobserved states, draw censored observations from their conditional posterior. Second, given the previous draw of parameters and censored observations, use the Carter-Kohn algorithm to draw unobserved

[5] I previously considered imposing that they all have the same Ψ_i, however due to the large amount of data, and the large degree of censoring convergence is impractically slow.

states from their conditional posterior distribution. Third, draw Φ, σ_l, and hierarchical parameters $(\boldsymbol{\psi}, \Sigma^{\psi})$ from the unconditional posterior distribution of Type 1 food banks. Fourth, use the Gibbs Sampling algorithm described above to draw Ψ_i, then α_i and λ_i parameters from their conditional posterior distribution. Finally, every 10th iteration, draw beliefs from their posterior distribution using Metropolis Hastings.

I use the same specification of priors as type 1 food banks and the same distributional/functional form assumptions. Again, I focus on just the first 25 unique auctions each period. I run 200,000 iterations, burning out the first 100,000 draws, and perform 4 independent chains. I keep 250 parameter draws from each chain, sampled uniformly, maintaining the correlations with the sampled parameters from Type 1 food banks. I estimate around 2400 parameters across 1.6 million observations, 1.45 million of which are censored.

K.4.2 Discussion

In reality even the Type 2 food banks are likely forward looking. And in principle, rather than interpret what I estimate as a static payoff, I could interpret it as another pseudo-static payoff function. Therefore I could apply the third stage estimation procedure. While this procedure is likely to produce fairly imprecise results, due to the lack of variation in bidding behaviour and imprecise estimates, I will consider this approach in a future robustness exercise.

However, the cost of misinterpreting their pseudo-static payoff function as a payoff function is potentially large. This is despite that Type 2 food banks only consume a relatively small amount of food under the Choice System, and even less under the Old System. This is because pseudo-static payoffs take into account expected future flow payoffs. Therefore when summing over time periods we essentially double count flow payoffs, skewing results towards these (typically lower priority) food banks. To alleviate this issue (and only in my final welfare calculations) I make the simplification that the exante value function for Type 2 food banks can be written as $K + \mathbf{s}_i^{g^T} \Psi_i \mathbf{s}_i^g$. Comparing this to the ex-ante value function for Type 1 food banks, the quadratic term is just the pseudo-static payoff from winning no lots each period. This simplification therefore asserts that the ex-ante *marginal* value function, given here by K, is independent of $\mathbf{s}_i^g$. Given that type 2 food banks bid so infrequently this simplification is plausible. We can then write the flow payoff function as:

$$
\begin{aligned}
j(\mathbf{s}_{it}^{g}) = \mathbf{s}_{it}^{gT}\Psi_i\mathbf{s}_{it}^{g} - \beta(K + E[\mathbf{s}_{it+1}^{gT}\Psi_i\mathbf{s}_{it+1}^{g}|\mathbf{s}_{it}^{g}])& \\
= (1-\beta)\mathbf{s}_{it}^{gT}\Psi_i\mathbf{s}_{it}^{g} - 2\beta\boldsymbol{\mu}_i^T\Psi_i\mathbf{s}_{it}^{g} - \beta[\boldsymbol{\mu}_i^T\Psi_i\boldsymbol{\mu}_i + Trace(\Psi_i\Sigma_j)]&
\end{aligned}
$$

The constant term drops out because we normalise $j(0) = 0$. This process ensures that we do not double count flow payoffs. It is true that this simplification will impact my welfare calculations, but the effects are expected to be minor given that Type 2 food banks only consume a relatively small amount of food under the Choice System.

Even though I use this approach for welfare calculations, I continue to assume to assume myopia for counterfactual simulations. That is, their accept/reject decisions are based the estimated pseudo-static payoff under the Choice System, rather than the estimated flow payoffs and a counterfactual equilibrium continuation value. Therefore, my simulated counterfactual equilibrium will be invalid. Using the pseudo-static payoff essentially assume that the mechanism reverts to the Choice System in the following period. These food banks rarely bid, and rarely win, and are more likely than others to be picky. Therefore they are likely to reject more often than they should, as their value function incorrectly assumes they can be picky next period. In practice, for Type 1 food banks, I find that the discounted continuation value is not a large component of accept/reject decisions under the Old System. This is because the continuation value is extremely flat - much more so than the continuation value under the Choice System. Therefore this inaccuracy is likely to be minor.

Appendix L

Additional Estimation Results

In this Appendix I report additional estimation results, adding to those in section 4.4. This includes tables and plots of parameter estimates, Gelman-Rubin Convergence tests, and model fit.

L.1 First Stage

Shape Parameters

Differ if more than 500 loads. Category specific shape parameters are given in column 1 of Figure L.1. Only the 8 most common categories (with more than 500 loads auctioned over the period, excluding fresh produce) have category specific parameters, the remainder are constrained to be equal. The estimated shape parameters for mixed loads is 0.0538 (0.0383,0.0685). These parameters all lie within the interval (-0.05,0.1) with the exception of Condiments, Cereal, and Meals which have values exceeding 0.1, suggesting that winning bids on these types of food have larger right tails, likely due to subcategories, such as peanut butter, that attract extremely high bids. Only Dairy has a shape parameter that is estimated to be significantly below zero, meaning that winning bids on Dairy are bounded above. This is perhaps due to the storage requirements for Dairy products.

L.1.1 Scale Parameters

The standard deviation of the winning bid on lot l is given by $\frac{\sigma_l}{\xi_l}$. Column 2 of Figure L.1 give the estimated scale parameters. These typically lie between 2000 and 5000, with the exception of fresh Produce, for which winning bids are typically clustered around the reservation price and have small standard deviation, and Pasta, which receives a small number of very high bids.

Column 2 of Figure L.1 gives the additional scale fixed effect from a load being from the "other", "mixed" or "assorted" subcategories. If this value is negative it suggests these subcategories have a smaller standard deviation than other subcategories within the category. Most often, these parameters are not estimated to be significantly different from zero. I also estimate an additional parameters for when the load has already been auction previously, since food banks have additional information about previous bids on the load. This parameter is estimated at 3370 (3030,3710).

FIGURE L.1: Category Specific First Stage Parameters

Category	Shape	Scale	Scale (other)	Maroon	Loads	Threshold 1	Threshold 2
Baby	0.0886	3100	-1,200	-846	-10,500	-3,020	-175
	(0.0659,0.113)	(2790,3410)	(-1,970,-173)	(-1,520,-239)	(-17,700,-5,890)	(-3,230,-2,800)	(-216,-136)
Bev	0.0174	2170	-370	-756	-5,200	-3,470	-320
	(-0.00205,0.0388)	(2080,2280)	(-812,130)	(-2,470,835)	(-5,880,-4,590)	(-4,090,-2,850)	(-439,-203)
Baked	0.0886	2160	240	-846	-10,400	-3,020	-175
	(0.0659,0.113)	(1850,2460)	(-728,1370)	(-1,520,-239)	(-16,100,-6,040)	(-3,230,-2,800)	(-216,-136)
Cereal	0.133	4560	-520	-2,600	-2,750	-3,020	-175
	(0.0935,0.175)	(4320,4810)	(-886,-179)	(-4,710,-486)	(-4,700,-1,290)	(-3,230,-2,800)	(-216,-136)
Condiment	0.342	3660	-226	623	-9,140	-3,020	-175
	(0.277,0.417)	(3380,3910)	(-916,569)	(-1,070,2370)	(-11,800,-6,750)	(-3,230,-2,800)	(-216,-136)
Dairy	-0.0421	2340	-224	-846	-5,600	-2,680	-186
	(-0.0726,-0.00883)	(2220,2480)	(-945,817)	(-1,520,-239)	(-6,470,-4,660)	(-3,410,-1,940)	(-317,-64.5)
Fresh	0.0886	576	-3.88	-124	-2,940	-564	-474
	(0.0659,0.113)	(516,635)	(-73.8,63.4)	(-1,440,1230)	(-3,290,-2,600)	(-1,020,-122)	(-600,-356)
Frozen	0.0204	2560	1110	-1,390	-5,930	-3,020	-175
	(-0.0268,0.0762)	(2390,2780)	(-791,3930)	(-3,030,225)	(-9,350,-3,500)	(-3,230,-2,800)	(-216,-136)
H/B	0.0886	3570	-1,460	-846	-5,720	-3,020	-175
	(0.0659,0.113)	(3290,3900)	(-1,890,-977)	(-1,520,-239)	(-9,310,-2,590)	(-3,230,-2,800)	(-216,-136)
Meals	0.149	4020	-587	-846	-4,670	-3,020	-175
	(0.113,0.194)	(3830,4220)	(-1,090,-49.4)	(-1,520,-239)	(-5,970,-3,480)	(-3,230,-2,800)	(-216,-136)
Meat	0.0886	4360	1550	1240	-10,400	-3,020	-175
	(0.0659,0.113)	(3930,4840)	(577,2660)	(-1,260,3660)	(-15,300,-6,530)	(-3,230,-2,800)	(-216,-136)
Cleaning	0.173	2730	137	2110	-4,730	-3,020	-175
	(0.116,0.235)	(2520,2970)	(-393,692)	(426,3690)	(-6,190,-3,500)	(-3,230,-2,800)	(-216,-136)
Pasta	0.0886	5730	-612	-846	-6,790	-3,020	-175
	(0.0659,0.113)	(4360,7390)	(-2,600,1090)	(-1,520,-239)	(-11,900,-2,490)	(-3,230,-2,800)	(-216,-136)
Snack	0.0374	2220	-353	-767	-6,680	-3,830	-239
	(0.0211,0.057)	(2160,2310)	(-572,-119)	(-2,230,653)	(-7,980,-5,590)	(-4,490,-3,210)	(-331,-147)
Vegetables	0.0886	3540	-374	1230	-2,870	-3,020	-175
	(0.0659,0.113)	(3260,3850)	(-1,070,409)	(-761,3090)	(-5,150,-1,310)	(-3,230,-2,800)	(-216,-136)

Note: 95% Credible Intervals are given in parentheses.

L.1.2 Location Parameters

Figure L.2 plots the subcategory specific fixed effects. Results are strongly correlated with the results presented in figure 2.7 ($R^2 = 0.74$). Figure L.3 plots the slope coefficients on the log of aggregate supply across the five "Use" types of food. Aggregate supply is given by both the previous 30 day aggregate supply and supply being auctioned that period. Coefficients are standardised, so that a one standard deviation increase in log 30 day supply of "Ingredients" decreases winning bids by 0.007

standard deviations. Estimated parameters are typically very small, and only significantly negative for the 30 day supply of Meals, Ingredients, and Snacks. I estimate a significantly positive and small (0.006 standard deviations) effect for contemporaneous supply of condiments, however this is difficult to interpret. It is possible that the small estimated coefficients are due to the extremely coarse food groupings I have used.[1] In practice it is unlikely that food banks keep track of the food supply from particularly detailed food groups. However, this remains a weakness of this analysis.

Figure L.1 column 4 displays the coefficients on maroon dummy variables. Parameters are constrained to be equal for categories with fewer than 50 maroon loads. Importantly, Maroon loads have a reservation price of zero, rather than -2000. The model I estimate focuses on difference from the reservation price, rather than raw winning bids. Therefore, with the exception of Fresh produce (which also has a reservation price of 0), I estimate that winning bids on maroon loads are systematically higher than winning bids on non-maroon loads when estimated parameters are significantly greater than -2000. On average Maroon loads attract winning bids around 1000 shares higher than non-maroon loads. Figure L.1 column 5 shows the linear slope coefficient on the number of homogenous loads auctioned simultaneously. Estimates are all significantly below zero, so that multiple loads attract lower bids than single loads. This is sensible — there is less competition for each load.

For the 8 different economic regions (and Canada), only loads from the "MidEast" region attracts significantly different winning bids, but the effect is small at around 0.001 standard deviations. Previously auctioned loads attract significantly lower winning bids, with a point estimate of -11,700 (-12,400,-11,000), around a third of a standard deviation. Loads that contain different categories of food attract lower bids, but the magnitude is small (less than 0.001 standard deviations), and significantly different from zero only when the lot contains four distinct types of food. Loads with free delivery, additional notes, or additionally shelf stable products do not attract significantly different winning bids. Finally, loads with restrictions on where they can be sent, or how the food must be picked up attract significantly lower winning

[1] however in a previous version of the model I estimated category specific slopes, and found similarly small effects.

bids, with a point estimate of -3,490 (-4,360,-2,700).

FIGURE L.2: Estimated Subcategory Fixed Effects

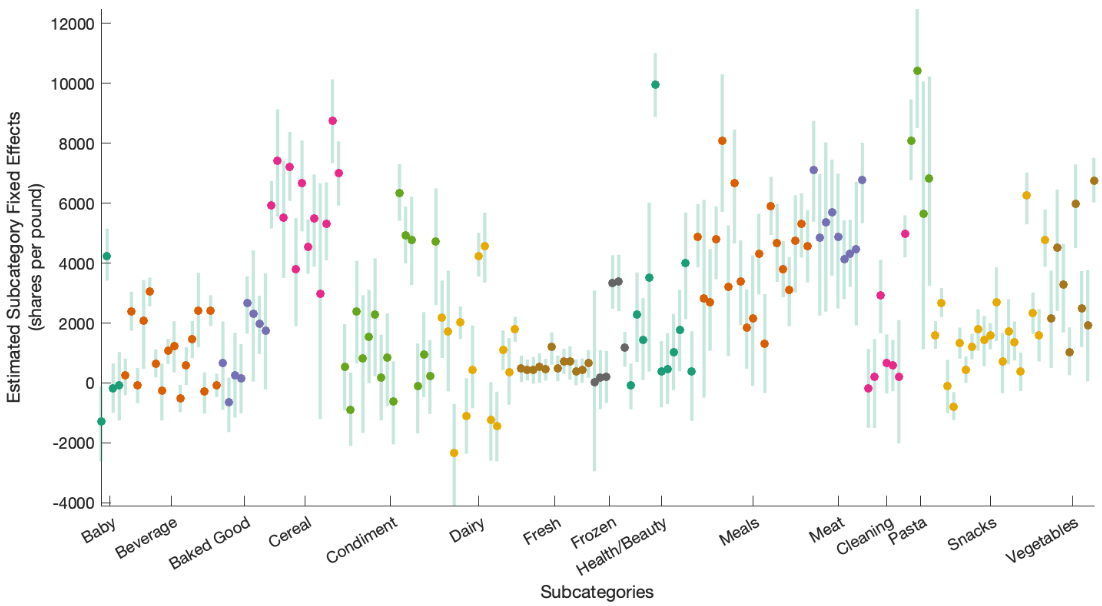

Note: Plot shows coefficients on location shifting subcategory specific dummy variables. Points give posterior means, and 95% Credible Intervals are given by the shaded lines. To interpret magnitudes in terms of standard deviations of winning bids, one must multiply by the associated shape parameter, and divide by the associated scale.

L.1.3 Threshold Parameters

Figure L.1 column 6 gives the estimated threshold parameters $\underline{R}^c$. These are all estimated to be far from zero, indicative of the high likelihood of observing winning bids at the reservation price. Column 7 gives the other threshold parameters $\bar{R}^c$. To interpret the parameters in terms of the excess mass just above the reservation price, they must be divided by the standard deviations (around 30,000), so that on average there is around 0.5% more mass just above the reservation price than expected. The estimated exponential parameter is 0.395 (0.343,0.462), so that the model predominantly rationalises bids within 5 shares of the reservation price in this way.

FIGURE L.3: Estimated Effect of Aggregate Supply on prices

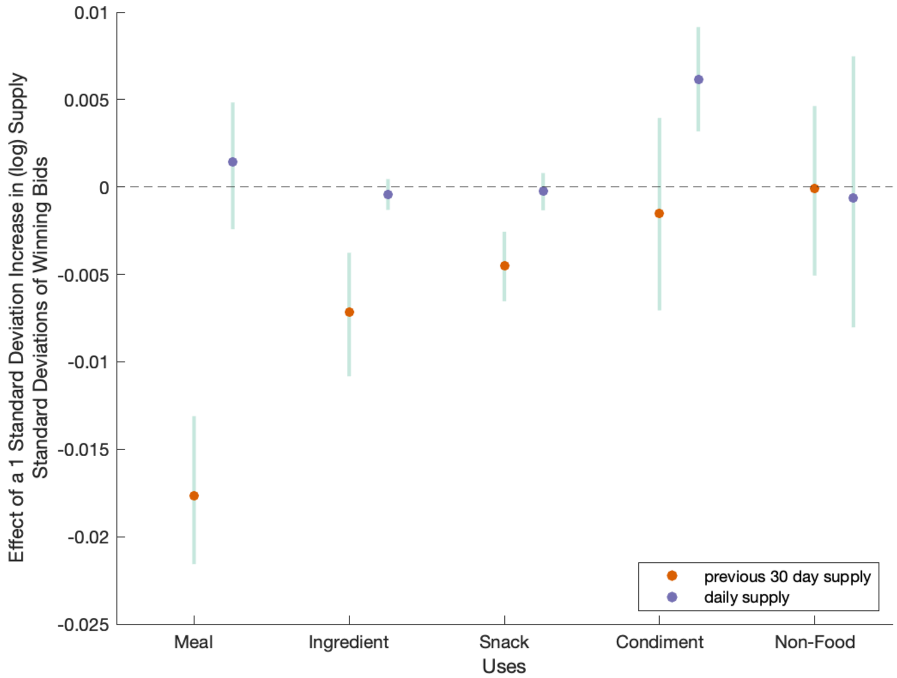

Note: Plot shows coefficients on aggregate supply, by food use for both daily supply and the previous month's supply. Points give posterior means, and shaded lines show 95% Credible Intervals. In non-standardised terms at the mean of 1600 tons of meals (food that can be consumed as a meal in itself) per month, an increase in the previous 30 day supply of meals by 1000 tons ($\approx$ 50 loads decreases the expected winning bid by around 500 shares.

L.2 Second Stage

L.2.1 Lot Specific Pay-off Parameters

Figure L.4 panel (A) plots food banks' estimated transportation costs, measured in consumer surplus. Coefficients cannot be interpreted as willingness to pays as they are not divided by marginal value of wealth (λ_i). Coefficients are positive, suggesting transportation is costly, and we see significant differences across food banks. Figure L.4 panel (B) plots the estimated log marginal value of wealth across food banks, $\hat{\lambda}_i$. Estimates above zero are relatively more budget constrained than the median food bank, and I estimate significant variation in these parameters. This suggests that shares are not allocated correctly, since to achieve efficiency the social planner would equate marginal values of wealth.[2]

[2]Note that identification of these parameters, which comes from variation in the variance of bids across food banks, strongly rests on the assumption that the lot specific payoff has the same variance across food banks. Therefore it is not possible to determine whether variation in these parameters is actually due to unmodelled variation in lot specific variances.

Figure L.5 panel (A) plots the estimated standard deviations of the lot specific idiosyncratic payoffs across category combinations. These tend to be between 10,000 and 30,000, which is significantly higher than the observed standard deviation of bids (around 3,000). They are large due in part due to the much larger variance of bids plus the markup term $b_l + \frac{\Gamma_l(b_l)}{\nabla_b \Gamma_l(b_l)}$, and also to rationalise the large degree of censoring. If only 2% of bids are observed, then the observed variation is only the variation in the far right hand tail.

L.2.2 Combinatorial Pay-off Parameters

Figure L.5 panel (B) plots the estimated Φ parameters, essentially how $k(\mathbf{s}_i, \mathbf{s}_0)$ vary with stocks of subcategories. Estimated parameters are strongly correlated with with the estimated first stage subcategory parameters plotted in figure L.2 ($R^2 = 0.82$)

FIGURE L.4: Estimated distance and marginal value of wealth parameters

(A) (B)

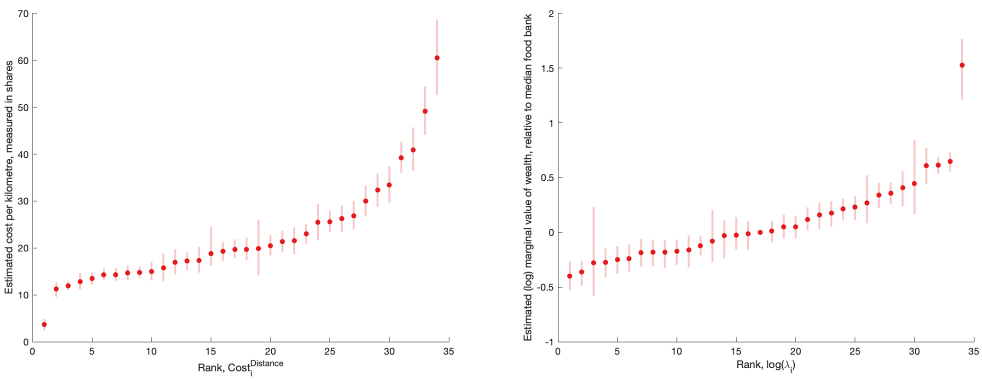

Note: Plot shows posterior means and 95% credible intervals of estimated distance coefficients (in km, panel (A)) and log marginal value of wealth (panel B). The Marginal value of wealth is normalised to 1 for the Type 1 food banks with median consumption.

FIGURE L.5: Estimated lot specific standard deviations and subcategory parameters (Φ)

(A) (B)

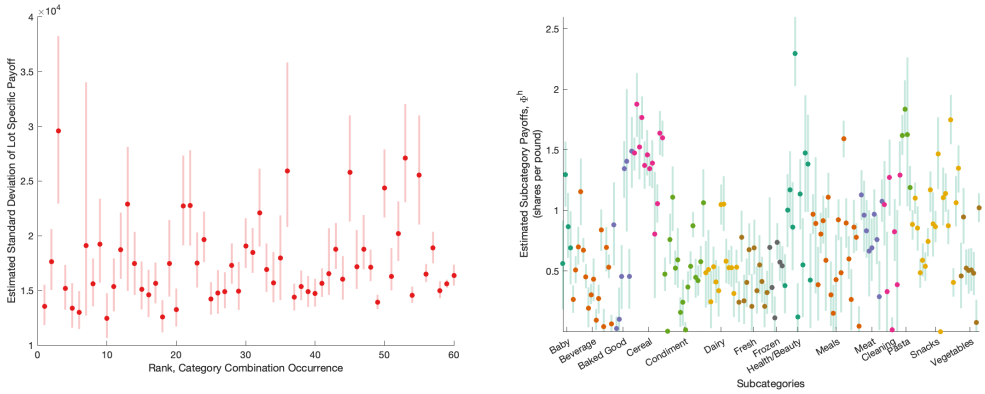

Note: Plot shows posterior means and 95% credible intervals of estimated standard deviations of v_{ilt} (panel A), and of estimated Φ subcategory pay-off parameters (panel B). For panel (A) The x-axis shows different combinations of categories included in the same lot, as each of the 59 most common unique combinations receives their own parameter. More common combinations appear further to the right. The furthest right parameter corresponds to the remaining 423 observed unique combinations, which make up 7% of the data. For panel (B) The scale can be interpreted as consumer surplus, measured in shares. A coefficient of 1 can be interpreted as every addition pound of food increasing consumer surplus by one share.

L.3 Third Stage

L.4 Type 2s

L.4.1 Unobserved State

Figure L.6 panel (A) presents estimated mean net donations for Type 2 food banks. Estimates are generally greater (less negative) than for Type 1 food banks. This is to be expected given these food banks typically win less food, suggesting they need less food in the first place. Likewise Figure L.6 panel (B) presents estimated standard deviations of net donations for Type 2 food banks. Estimates are generally smaller than for Type 1 food banks. This is partly surprising, since one might expect that each period larger food banks give out a more predictable amount to their clients, and receive a more predictable amount from local donors. This is essentially a law of large numbers given Type 1 food banks are expected to give out and receive more food than Type 2 food banks. However, Type 2 food banks are not necessarily smaller than Type 1s, and many have larger Goal Factors. The fact they rely on the Choice System less is likely due to having many local donors, which may lead to them receiving a more constant supply of local donations over time. It is also possible that these results are driven mechanically by my priors (as food banks who win more food may mechanically have a larger variance of their winnings), making it important that my priors about the variance are informative.

FIGURE L.6: Estimated unobserved state parameters (Type 2)

(A) (B)

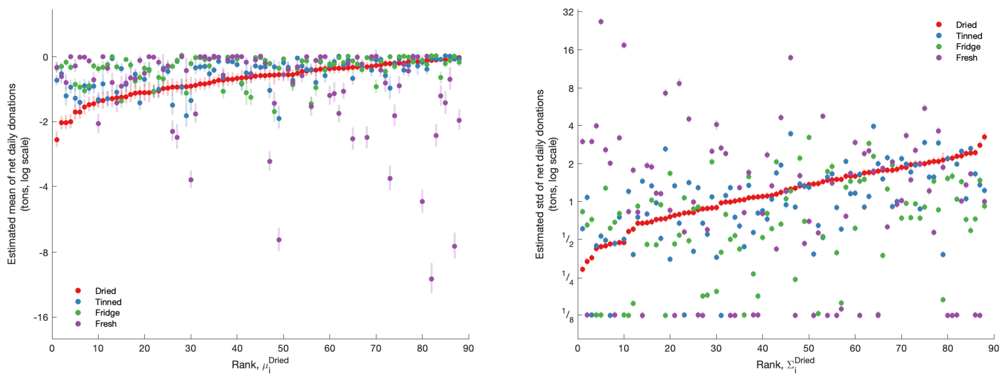

Note: The figure plots posterior means for the mean and standard deviations of net local donations, as well as 95% credible intervals. Results are sorted according to the estimates for the Dried storage type. The plot excludes the 'non-food' type, to improve graphability.

L.4.2 Lot Specific Payoff

Figure L.7 panel (A) plots the distance costs for Type 2 food banks. Some of the estimated coefficients are in the ranges of the type 1 food banks, these are likely the large Type 2 food banks with access to many local donors. Mostly, Type 2 food banks have much higher distance costs. Figure L.7 panel (B) plots the log of the estimated marginal value of wealth, which is still measured relative to the median Type 1 food bank. Type 2 food banks are estimated to generally have much lower estimates, suggesting that shares are not as valuable to Type 2 food banks as they are to Type 1 food banks. This is unsurprising given that Type 2 food banks are known to already need less food, either due to access to local donors, or as they have a smaller amount of poverty in their local area.

FIGURE L.7: Estimated distance and marginal value of wealth parameters (Type 2)

(A) (B)

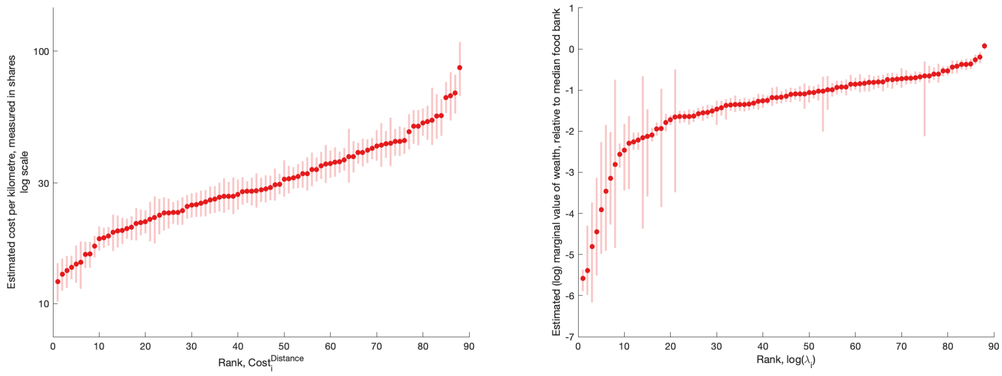

Note: Plot shows posterior means and 95% credible intervals of estimated distance coefficients (in km, panel (A)) and log marginal value of wealth (panel B). The Marginal value of wealth is normalised to 1 for the Type 1 food banks with median consumption.

L.4.3 Combinatorial Payoff

Figure L.8 plots the distribution of Type 2 food banks' willingness to pays for a single lot from each storage type, given stocks of zero. Estimates are most often negative, suggesting that storing food is costly. Estimates are generally higher (lower storage costs) for Type 2 food banks than Type 1 food banks. This is most likely because the state is normalised to zero for the first date in my sample period. The zero, normalised, state is likely higher for Type 1 food banks as they typically win more food. Therefore it is sensible that these food banks are more capacity constrained to begin with.

L.5 Diagnostics

This sub-Appendix reports Gelman-Rubin statistics, allowing us to assess model convergence. I follow the approach to constructing the test statistics laid out in Gelman et al., 1995. Figure L.9 reports the results of this analysis. For each set of parameters I report the proportion of statistics below the recommended cutoffs of 1.2 and 1.1. I report results for both types of food banks, except in cases when the relevant parameters are the same for both types.

FIGURE L.8: Estimates of Ψ_i (Type 2)

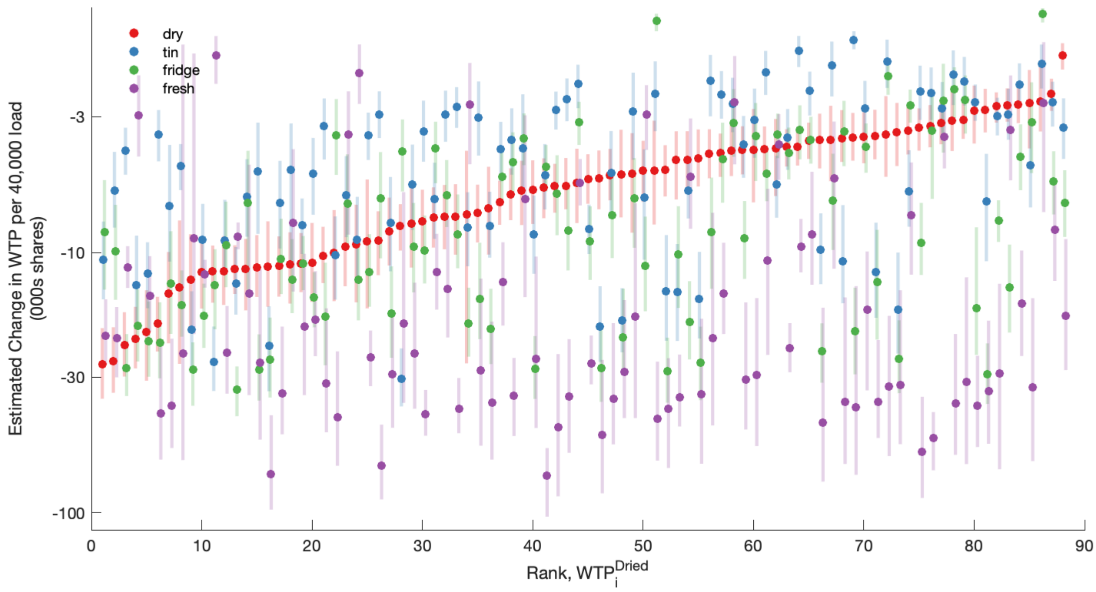

Note: Figure plots posterior mean equilibrium willingness to pay for a 40,000 load for each storage type. Bars give the 95% credible intervals. Estimates are ordered according to the estimates for Dried loads. The plot excludes estimates for non-food storage type. WTPs are evaluated when stocks are zero.

Broadly I have evidence of convergence, though not as strong as one might hope for the second stage parameters.

For the first stage parameters every parameter is found to converge except one of the 'other subcategory' scale parameters (likely due to a lack of observations). For the parameters of the second stage evidence of convergence is less strong, though only when we focus on the more stringent cutoffs. The lack of convergence is clustered within the four least regularly bidding Type 1 food banks, and the ten least regularly bidding Type 2 food banks. Each individual food bank has a relatively small effect on my counterfactuals (with the exception of the 5 largest food banks, all of whom converged). Therefore the effect on my main results of this non-convergence is likely to be very minor. For the sake of posterity, I now discuss this lack of convergence in more detail.

My sampler likely did not fully converge to the target distribution for certain parameters. Convergence was already expected to be slow due to the large degree of censoring, and the 300,000 iterations were likely insufficient to achieve full convergence for every parameter.[3]

FIGURE L.9: Gelman-Rubin Convergence Statistics

	Type 1 food banks		Type 2 food banks	
Parameters	Prop < 1.1	Prop < 1.2	Prop < 1.1	Prop < 1.2
Γ	0.996	1	-	-
μ_i	0.947	0.976	1	1
$\boldsymbol{\Sigma}_i$	0.9	0.935	1	1
$\boldsymbol{\Phi}$	0.976	1	-	-
λ_i	0.909	0.97	0.784	0.943
σ_l	0.933	1	-	-
Distance	0.971	1	0.966	1
ψ_i	0.902	0.975	0.92	0.986
ψ	1	1	-	-
$\boldsymbol{\Sigma}^{\psi}$	1	1	-	-

L.6 Fit

L.6.1 First Stage

Figure L.10 plots the estimated and empirical probability a food bank wins a lot given their bid, where bids are measured in distance from the reservation price. The estimation drops the first 60 days, and the final 150 days, and then randomly samples 95% of the remaining data. The other 5% is used as a validation dataset. Probabilities are plotted by taking an expectation over covariates. The discontinuity in probability occurring at zero is due to the non-negligible probability of ties at the reservation price.

The model does a good job of matching the probability of winning at the reservation price, and just above it. The fit worsens around zero (typically 2000 shares above the reservation price) due to excess mass at zero, and hence excess cumulative

[3] In future more iterations should be used. In particular, just as I presently only sample beliefs every 10th iteration due to computational cost, in future I will only run the Carter-Kohn algorithm (drawing stocks from their posterior) every 10th iteration. These Carter-Kohn steps were by far the most computationally intensive, taking around 90% of all the computation time. This will allow me to run the sampler for around five times as many iterations in the same space of time. This should improve convergence by allowing me to overcome the large degree of censoring.

probability above zero. The model is unable to rationalise food banks' bids being anchored around zero. However this inaccuracy is not large, even if it is statistically significant - the vertical distance between the two lines never exceeds 0.05.

FIGURE L.10: First Stage Fit, actual vs simulated

(A) (B)

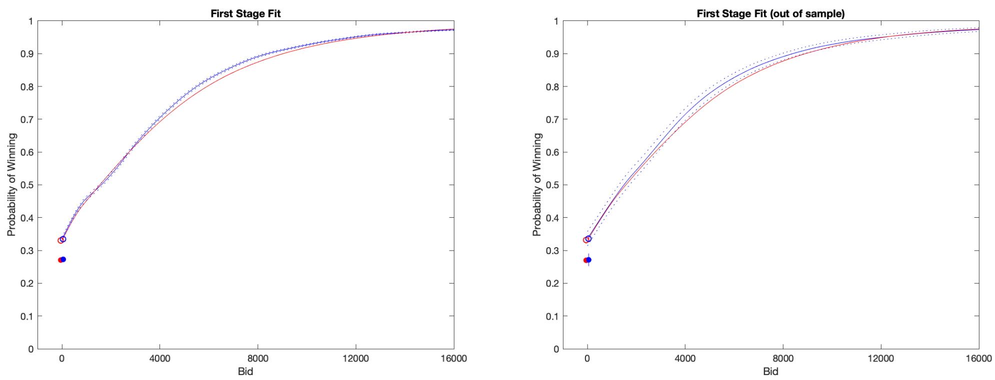

Note: probability of winning given bid, i.e. cdf of winning bids. discont due to ties. simulated values from estimated distribution vs empirical distribution. averaged over covariates.

L.6.2 Second Stage

Figure L.11 presents several observed moments, comparing them to simulated moments. For the simulated moments I present the posterior mean moment, as well as the 2.5th and 97.5th percentiles. In sample moments are calculated on the training sample (in-sample), which cuts off the first 60 and final 150 days. The validation sample (out-of-sample) uses only the final 150 days.

Most of the moments I consider are self-explanatory, except for the number of lots won, relative to the observed number. For each simulation this takes the number of lots a food bank wins, dividing this by the number of lots they were observed winning. I then consider the average of this proportion across food banks. It mechanically equals one in the observed data. This moment is typically fit quite well by the model, with both Type 1 and Type 2 food banks winning similar amounts of food in y simulations compared to the true data. Type 2 food banks perhaps win too many lots and Type 1 food banks perhaps slightly too few. This is likely due to the simulated over bidding of Type 2 food banks. This moment is important because it shows that the model correctly predicts food banks equilibrium allocations, even if it does not correctly predict bidding behaviour. A similar pattern is seen when we consider equilibrium expenditure (essentially weighting winnings by value) as well as food won by type of food.

Model fit, in terms of bidding behaviour, is poor for Type 2 food banks. They are predicted bidding more often than they actually do, and bidding too aggressively conditional on bidding. This is typically due to parameters failing to converge properly. In particular, the estimated marginal value of wealth for these food banks are far too low. Fortunately given that these food banks consume a relatively small proportion of total food, and this is predicted well by the model, these food banks have a small total impact on welfare, and so these inaccuracies will not majorly impact my counterfactuals.

The model fits much better for Type 1 food banks, though still estimates them bidding too aggressively, with average bids around 70% larger than observed. However, relative to the already large standard deviation of bids, this is not major. These inaccuracies seem to be caused by simulation error, as for each simulation I must numerically find optimum entry and bidding decisions, for which I use a simple hill-climbing heuristic that need not necessarily find a global optimum.

FIGURE L.11: Estimation Moments: Observed vs Simulated

		In-Sample				Out-of-Sample			
Moment	FB Type	Observed	Mean	p2.5	p97.5	Observed	Mean	p2.5	p97.5
Average No. bids	1	1.33	1.32	1.25	1.43	0.899	1.08	0.938	1.22
(per period)	2	0.252	0.596	0.554	0.633	0.146	0.525	0.476	0.57
P $\geq$ 1 bids	1	0.457	0.488	0.472	0.504	0.344	0.437	0.412	0.455
(per period)	2	0.157	0.363	0.345	0.38	0.0997	0.33	0.307	0.346
Mean bid	1	1640	2560	2260	2770	1010	2710	2350	3050
(given enter)	2	2590	6120	5600	6640	1590	6480	5790	7190
Std of bids	1	3890	4600	4420	4790	3020	4540	4290	4720
(given enter)	2	5290	8040	7380	8780	3260	7880	7080	8810
No. of lots won	1	1	0.829	0.562	1.16	1	0.887	0.429	1.53
relative to observed	2	1	1.54	0.979	2.32	1	2.62	1.04	6.02

Note: Moments are calculated for Type 1 and Type 2 food banks separately. 'Out of sample' refers to the final 150 days which are dropped from estimation. The final set of moments considers the number of lots won by each food bank in each simulation, and how this compares to the number of lots actually won. For each food bank and each simulation I consider the ratio of lots actually won and to the lots won in the simulation. Values closer to 1 are closer to the observed data. Values above 1 have food banks winning too many lots on average.

Appendix M

Robustness

This Appendix investigates how robust my results are to certain key assumptions and simplifications. Robustness exercises are split across the three stages of my estimation procedure in Appendices M.1, M.2, and M.3 respectively.

M.1 First Stage

In this appendix I consider how robust is estimation to the specific assumptions I make about equilibrium beliefs. At present, I focus on the assumptions made on beliefs in Assumption 9. Specifically, that every food bank faces the same distribution of maximum rival bids, and that beliefs are not a function of individual food banks' states, but rather aggregate states. In Appendix M.1.3 I consider the assumption that winning bids are conditionally independent across auctions.

M.1.1 Food bank Specific Beliefs

Assumption 9 part (v) imposed that $\Gamma_i = \Gamma$, so that every food bank is assumed to face the same distribution of rival bids. This allowed me to estimate Γ on the distribution of winning bids only. I can test this assumption, testing whether the distribution of food bank i's rival's highest bids is significantly different from the distribution of winning bids. In practice, this involves replacing food bank i's winning bids with the second highest bid in each of these auctions. This permits a simple hypothesis test for food bank i: We construct this alternative dataset and consider whether the estimated Γ_i from this dataset is significantly different from the estimated Γ constructed from winning bids only. This can be done using a simple Score Test.

In figure M.1 panel (A) I present the distribution of test statistics across food banks. Under the null hypothesis these statistics take a χ^2 distribution with 268 degrees of freedom (the number of first stage parameters). None of these tests can reject the null hypothesis at the 10% significance level.

M.1.2 Dependence on Aggregate Supply

Assumption 9 part (v) also requires that beliefs do not depend on any individual food banks' state, and instead only depend on aggregate statistics, such as the aggregate supply of various types of food. The argument is that equilibrium is sufficiently competitive that no individual food bank's behaviour is able to significantly shift the distribution of equilibrium winning bids. If this is the case, then no individual food bank's state will be able to significantly shift this distribution either.

To test this assumption I consider whether any food bank has an individually significant effect on the distribution of equilibrium winning bids. The results presented in Appendix M.1.1 act as evidence in favour of this hypothesis, essentially presenting the distribution of equilibrium winning bids when each food bank is 'removed' from the system in turn. To go further, I can also consider whether the distribution of equilibrium winning bids changes when data from food bank i *and* the auctions they won are removed from the system. If food bank i has a significant effect on the distribution of winning bids, we would expect that the distribution of winning bids is different when we drop all the data from food bank i. In figure M.1 panel (B) I present the distribution of score test statistics across food banks. Under the null hypothesis these statistics take a χ^2 distribution with 268 degrees of freedom (the number of first stage parameters). None of these hypothesis tests can reject the null hypothesis at the 10% significance level.

FIGURE M.1: Robustness: Stage 1

(A) (B)

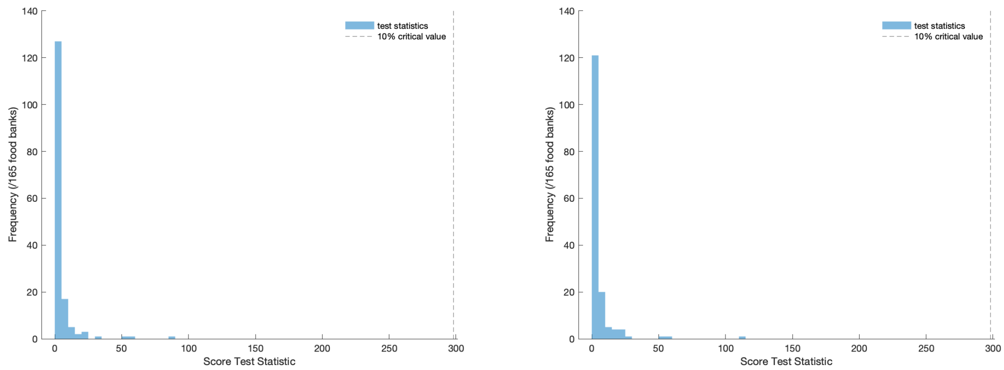

Note: The figure plots Score Test statistics from two robustness checks. Panel (A) relaxes the restriction that every food bank has the same equilibrium beliefs, while panel (B) tests whether any individual food bank's bidding behaviour has a significant effect on the distribution of winning bids.

M.1.3 Independence of Winning Bids

In this Appendix I investigate the assumption that winning bids within a period are conditionally independent across auctions. This assumption is necessary to ensure the joint probabilities of combinatorial outcomes $P(\mathbf{b}_t, \mathbf{d}_t|\mathbf{s}_t)$ can be written as products of the marginal distributions. Given the linear demand, or quadratic pseudo-static payoff parametrisation I make, I only require that winning bids are pairwise independent.

I investigate the validity of this assumption by investigating the degree of pairwise correlation in winning bids within a period. While a lack of correlation is not sufficient to infer independence, it at least suggests that one winning bid is not informative of another winning bid. This ensures that food banks' beliefs about joint probabilities of pairwise outcomes should be close to the product of the marginal win probabilities, meaning that errors from this misspecification are expected to be minor. Importantly, I only need to test for the presence of conditional correlation. Two winning bids are allowed to be conditionally correlated (conditional on covariates), as it is assume that food banks form beliefs conditional on the state. Correlation in winning bids is most likely to arise from the complementarity terms Ψ that also creates correlation in food banks' bids across different lots (given I assumed v_{ilt} are uncorrelated across l). For example, since we expect lots to be substitutes a food bank's bids are likely to exhibit negative correlation (conditional on the state). Therefore we might also expect winning bids to exhibit negative correlation.

Writing $\bar{b}_{lt}$ for the winning bid on lot l at time t I investigate correlation between winning bids using the following regression specification:

$$\bar{b}_{lt} = \beta^1 \bar{b}_{l't} + \beta^2 \mathbf{x}_{lt} + \beta^3 \mathbf{x}_{l't} + \beta^4 \mathbf{s}_t^0 + \varepsilon_{lt}$$

$\mathbf{x}_{lt}$ give lot specific covariates, using the covariates included in the first stage of the estimation procedure such as subcategory fixed effects. $\mathbf{s}_t^0$ give time specific common state variables that do not vary across lots, just as in the first estimation stage. I include every pair of auctions (l, l') that occur simultaneously, giving me around $800,000$ observations (essentially including each pair twice, once on each side of the regression). Under the null hypothesis of independence β^1 should equal zero. However, this specification imposes that the relationship between every pair of winning bids is the same. One might expect negative correlation between substitutes and positive correlation between complements. To allow for differential correlations I also consider a specification that interacts $\bar{b}_{l't}$ with all three sets of covariates, allowing the correlation to depend on observable characteristics of the lots. I also consider a specification that includes triple interactions between $\bar{b}_{l't}$, $\mathbf{x}_{lt}$, and $\mathbf{x}_{mt}$. This predominantly involves including dummy variables for whether both lots come from the same subcategory, the same region, etc.

I consider statistical significance of the $\bar{b}_{l't}$ coefficients using asymptotic F-tests. However, it is also worthwhile to consider how much variation in $\bar{b}_{lt}$ $\bar{b}_{l't}$ is able to explain. If $\bar{b}_{l't}$ has very little explanatory power, then the extent of the dependence is expected to be minor. This means any departure from independence is unlikely to cause much inaccuracy in my results, since the true joint probabilities are expected to be very close to the product of the marginal probabilities.

Results are presented in Figure M.2. I can reject the null hypothesis of independence at the 1% significance level in all of my specifications. Therefore I have evidence that the independence assumption is invalid. However, as is evident from examining the R^2 values, the degree of dependence is extremely small. The covariates alone account for 41.2% of the variation in winning bids. Including the $\bar{b}_{l't}$

interactions is then only able to explain an additional 0.3% of the variation in winning bids. This suggests that winning bids are very close to being independent, even though we can reject independence. Therefore while I have found this independence assumption to be invalid, I have also found that it is likely to be a very good approximation to food banks' beliefs. This analysis does not presently account for censoring of winning bids at the reservation price, but will do in future.

One final point worth considering is how this simplification might impact my results. One difficulty with simultaneous auctions is that it makes it difficult for food banks to win precise numbers of loads. They might bid on two loads of cereal wanting precisely one, but there is a non-negligible risk they win both or neither. The more dependence there is between winning bids, the more information the food bank has about how they should bid, giving them even more control over which lots they win. In the cereal example, if winning bids are negatively correlated, the food bank knows they can place two middling bids and they will likely win precisely one lot. If they are positively correlated they know they should place one high and one low (if at all) bid. Therefore, this assumption is expected to bias my results against the Choice System.

FIGURE M.2: Robustness: Independence of Winning Bids

Specification	Covariates	F test df	p-value	R^2
$\bar{b}_{l't}$		1	0	0.08734
Covariates only	✓			0.4121
$\bar{b}_{l't}$	✓	1	0	0.4124
$\bar{b}_{l't} \times (\mathbf{x}_{l't}, \mathbf{x}_{lt}, \mathbf{s}_{0t})$	✓	457	0	0.414
$\bar{b}_{l't} \times (\mathbf{x}_{l't}, \mathbf{x}_{lt}, \mathbf{s}_{0t}, [\mathbf{x}_{l't} \times \mathbf{x}_{lt}])$	✓	681	0	0.4153

Note: The F test degrees of freedom and p-value refer to the hypothesis tests that all coefficients on $\bar{b}_{l't}$ are equal to zero, where the degrees of freedom gives the number of coefficients being considered.

M.2 Second Stage

I consider four alternate model specifications for the second stage. These are designed to test the model's robustness to relaxing key simplifications made in the main model. In Appendix M.2.1 I allow the value function to depend on common state variables. In Appendix M.2.2 I account for endogeneity in the observation equation using an control function procedure. In Appendix M.2.3 I consider robustness to the assumption of normally distributed lot specific payoffs by allowing the lot specific idiosyncratic payoff to follow a normal-inverse-gamma distribution. In Appendix M.2.4 I consider how the simplification to only consider data from 25 auctions each period impacts my results, by estimating the model using 50 auctions from each period.

M.2.1 Incorporating the Common State

In general food banks' continuation values depend on the common state variables, which contain information about future prices. Common state variables are captured by the demand index estimated in the first stage, mapping common states onto parameters of the distribution of winning bids.

To allow continuation values to depend on common states, when estimating the pseudo-static payoff function k in the second estimation step, we must allow k to vary with the demand index. Importantly the index must be interacted with food bank specific state variables, so that the marginal payoff also depends on the index. Otherwise dependence on the index will not be identified from bidding behaviour alone. I introduce the demand index by specifying k as follows:

$$k(\mathbf{s}_i, \mathbf{s}^0) = \Phi(I + D^0)\mathbf{s}_i^h + \mathbf{s}_i^{g^T}\Psi_i\mathbf{s}_i^g$$

Where D^0 is a diagonal matrix with entry $D^0_{hh} = \sum_u \delta_u d^{0u}\mathbb{I}[h \in u]$, where $\mathbb{I}[h \in u]$ is a dummy variable for whether subcategory h has usage type u. d^{0u} is the demand index for food from usage type u, and δ_u are parameters to be estimated. These parameters describe how strongly bidding behaviour changes given changes in aggregate supply. For example, when supply is high, and so d_t^{0u} is low, winning bids are expected to be low. If supply is also positively correlated over time, the opportunity cost from losing a lot today is low, as winning bids are also likely to be low in future. Therefore bidding will be *even* less aggressive today, and so we expect $\delta_u > 0$. This specification is natural - if there is dependence on common states, we would expect to see evidence of it to show up in a linear term. I interact the index with the subcategory stock term, rather than the storage type term, because the subcategory term reflects a food banks' 'wants', while the storage term is intended to reflect the costs that the food bank must put up with. This is relevant because the index affects how easily the food bank can win the types of food it wants, on behalf of their clients.

Different sized food banks, with different budgets and storage capacities, are expected to respond differently to variation in common states. For example, a food bank that is not heavily reliant on the Choice System for their staples is unlikely to be responsive to common states. Therefore I allow the δ parameters to vary across food banks, but again employ a bayesian hierarchical model to ensure a degree of shrinkage for food banks for whom identifying variation is scarce. I assume that $\delta_i \sim N(\delta, \Sigma^{\delta})$, where priors for $N(\delta, \Sigma^{\delta})$ are weak normal-inverse-wishart. The parameters are identified using variation in the demand indices, which arise from variation in the common states, and seeing how this translates into variation in bidding behaviour.

In figure M.3 panel (A) I plot estimates of δ_u across food banks. None are significant at the 5% significance level. This is predominantly caused by the lack of variation in the demand indices - as we saw in Figure L.3 winning bids do not vary much with variation in the common state variables. This explains the extremely large credible intervals relative to the scale: if $\delta^g = 1$ this means that a one unit increase in d_t^g (associated with a one share expected increase in the winning bids) is associated with a $\Phi \mathbf{z}_{tl}^h$ unit increase in bids.

Therefore we have evidence that food banks' continuation values also do not vary with common state variables. Consequently, when evaluating maximised expected payoffs and the ex-ante value function in the third estimation step, we do not need to explicitly consider dependence on the common states, as this will not impact estimates of the flow payoffs j backed out in the final step.

M.2.2 Endogeneity of the Inverse Bid System

I now consider the endogeneity of the observation equation, caused by non-additivities across lots. In essence, I re-estimate the second stage of my estimation procedure using a control function approach. The observation equation is given by:

$$\lambda_i y_{ilt} = \Phi \mathbf{z}_{tl}^h + \mathbf{z}_{tl}^{gT} \Psi_i (\mathbf{z}_{tl}^g + 2\mathbf{s}_{it}^g + 2 \sum_{m \neq l} \Gamma_m(b_{itm}) \mathbf{z}_{tm}^g) + v_{ilt}$$

Given that this step is essentially estimated using a bayesian regression, estimation requires that the error term v_{ilt} is independent of the regressors. In general there exists a dependency between v_{ilt} and b_{itm} (for $m \neq l$), because optimum bids (and entry decisions) are a function of every lot specific payoff. That said, we have reason to think this dependency might be small, since typically $\Gamma_m(b_{itm})$ will depend much more strongly on things other than v_{ilt}. However, allowing for this endogeneity is relatively easy. The endogenous regressor is $\mathbf{z}_{tl}^g(\mathbf{z}_{tl}^g + 2\mathbf{s}_{it}^g + 2\sum_{m \neq l} \Gamma_m(b_{itm})\mathbf{z}_{tm}^g)^T$, and there exists an obvious instrument for this regressor: $\mathbf{z}_{tl}^g(\mathbf{z}_{tl}^g + 2\mathbf{s}_{it}^g)^T$. This is the same type of instrument proposed in Chapter 3. What makes this instrumental variable procedure even easier is that our first stage is actually known, and given by the structure of the model.

Estimation is done using a control function approach. The basic idea is that we have a regression model along the following lines:

$$\begin{array}{ccc} y_t = \mathbf{x}_t^T \beta + u_t & & u_t | \mathbf{v}_t \sim N(\mathbf{v}_t \rho, \sigma^2) \\ & \& & \\ \mathbf{x}_t = \mathbf{z}_t + \mathbf{v}_t & & u_t | \mathbf{z}_t \sim N(0, \sigma^2) \end{array} \tag{M.1}$$

This is a standard case of endogeneity with an available instrument, except with a known first stage. In this setting β can be estimated consistently using the regression equation:

$$y_t = \mathbf{x}_t^T \beta + \mathbf{v}_t^T \rho + e_t$$

Because $y_t | \mathbf{x}_t, \mathbf{v}_t \sim N(\mathbf{x}_t^T \beta + \mathbf{v}_t^T \rho, \sigma^2)$. In my setting we specify the observation equation as we did previously, but include $\mathbf{z}_{tl}^g(2\sum_{m \neq l} \Gamma_m(b_{itm})\mathbf{z}_{tm}^g)^T$ as an additional regressor. The coefficient of this regressor is essentially an estimate of the endogeneity.

I specify weak normal priors for $\boldsymbol{\rho}_i$.

In practice, the endogeneity is unlikely to be linear. However this remains a useful starting point for two reasons. First, even when the endogeneity is non-linear, so that $E[u_t|\mathbf{v}_t] = f(\mathbf{v}_t)$, the posterior distribution for β, marginalised over ρ, remains correct. This is because, conditional on $\mathbf{v}_t$, $\mathbf{x}_t$ is independent of any non-linear functions of $\mathbf{v}_t$ that remain in the error term. Second, if there are non-linearities and the effect of endogeneity on $\mathbf{x}_t$ is large, the linear term should still pick up evidence of endogeneity.

Estimation remains as it was in the main text, except that before sampling ψ_i parameters, I form the conditional joint posterior distribution for (ψ_i, ρ_i), before marginalising over ρ_i, and sampling ψ_i from the marginal posterior for ψ_i.

In Figure M.3 panel (B) I present posterior means of ρ_i across food banks. Estimates are presented on the same scale as the results from Figure 4.2. Only 16% of estimates are individually significant at the 5% significance level. Furthermore, the estimated magnitudes of the bias are small - between 2% and 8% of the relevant entries of Ψ_i.

FIGURE M.3: Robustness: Stage 2 (1)

(A) (B)

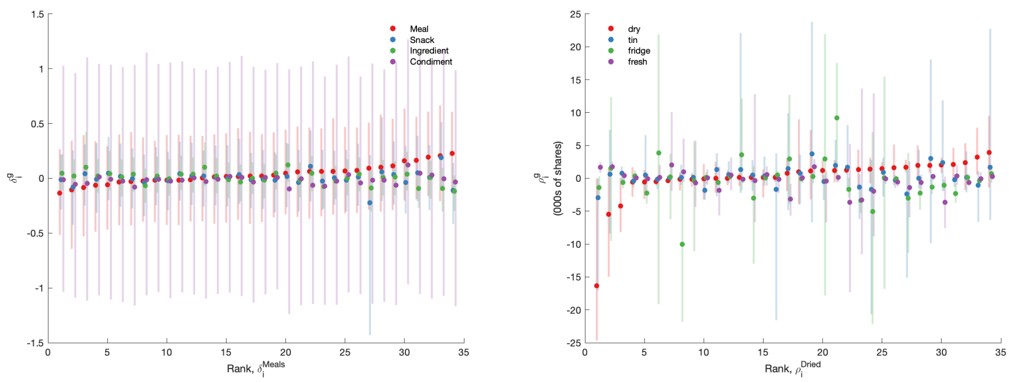

Note: The figure plots deviation statistics from two robustness checks. Panel (A) plots posterior means and 95% credible intervals of the coefficients on the demand indices across food usage types. To interpret the scale of the coefficients, if $\delta^g = 1$ this means that a one unit increase in d_t^g (associated with a one share expected increase in the winning bids) is associated with a $\Phi\mathbf{z}_{tl}^h$ unit increase in bids. We expect δs to be small and positive. Panel (B) plots posterior means and 95% credible intervals of the ρ parameters, which can be interpreted as estimates of the bias in the Ψ_i parameters caused by endogeneity of the observation equation. The estimated ρs are presented on the same scale as estimates of Ψ presented in Figure 4.2.

M.2.3 Normal-Inverse-Gamma Idiosyncratic Payoff

In this appendix I relax the assumption that the lot specific idiosyncratic terms v_{ilt} are normally distributed. Instead, I allow for the possibility that they take a normal inverse-gamma distribution:

$$v_{ilt} \sim N(0, \sigma_l^2 U_{ilt}) \qquad \text{where} \qquad U_{ilt} \sim IG(\alpha, \alpha)$$

This distribution has heavier tails than the normal distribution. The distribution can be interpreted as taking into account unobserved variation in lot specific attributes that affect the variance of the payoff. For example, I do not take into account different varieties of apples. The quality of some types of apples may be significantly more variable than other.

The procedure described shortly can also be extended to net donations $\mathbf{x}$. However the assumption that net donations are normally distributed is significantly more reasonable, given that these net donations are the sum of many local donations and many loads sent out to food pantries.

Sampling v_{ilt}s from their censored distributions and sampling $\mathbf{s}_{it}^g$ using the Carter-Kohn algorithm both rest strongly on the normal distribution assumption. The posterior distributions of $\mathbf{s}_{it}^g$ is intractable when v_{ilt} is non-normal. However, conditional on $\{U_{ilt}\}_{ILT}$, we have normality again. Therefore I use an additional data augmentation step in which I sample $\{U_{ilt}\}_{ITL}$ conditional on $\{v_{ilt}\}_{ITL}$, α, and σ_l^2. Given known difficulties associated with estimating the shape parameters of these types of distributions I fix $\alpha = 5$, ensuring the first four moments of the distribution exist. The σ_l parameters are just a rescaling of those presented in the main text.

This data augmentation step uses the following conditional posterior:[1]

$$U_{ilt} | v_{ilt}, \sigma_l^2, \alpha \sim \textit{scaled-inv-}\chi^2(2\alpha + 1, \frac{2\alpha + \frac{v_{ilt}^2}{\sigma_l^2}}{2\alpha + 1})$$

In Figure M.4 I present the results of Wald tests from different groups of parameters, considering how this alteration to the model changes the estimated model

[1] This result comes from the fact that the marginal distribution of the normal inverse-gamma distribution for v_{ilt} only is a t-distribution with 2α degrees of freedom. I then use the standard result that t-distributions can be written as a scale mixture of normals, for which the conditional posterior is readily available.

parameters. I can reject that the parameters have the same posterior means for an overall Wald test. However it is useful to see where the main differences are coming from. I find that we can only reject the null hypothesis that posterior means are equal for the lot specific variance, and marginal value of wealth parameters.[2] Estimated variance parameters are on average lower than those form the baseline specification, and λ_i parameters higher. This is because this specification does not need an excessively large variance in order to rationalise the heavy right tail of bids. In future I will investigate how these differences lead to different welfare effects from my simulations.

FIGURE M.4: Robustness tests, differences in posterior means

Alternate Model	Statistic	Ψ_i	Φ	λ_i	σ_l	Distance	μ_i	Σ_i
NIG v_{ilt}	χ^2	570	10.2	85.2	97.1	4.24	1.91	37.7
	p-val	0.0344	1	1.7e-06	0.00172	1	1	1
50 auctions	χ^2	2490	31.8	51.9	314	8.52	0.761	19
	p-val	0	1	0.0192	0	1	1	1

Note: This table presents test statistics and p-values from Wald tests for differences in posterior means across several alternate model specifications. Tests are performed separately across groups of parameters. The test statistic has an asymptotic chi^2 distribution with degrees of freedom given by the number of parameters of that type.

M.2.4 Accounting for Additional Non-Entered Lots

In the main text I estimated the model using only 25 unique auctions held each day. Although no food bank was every observed placing more than 25 bids, on around 50% of days there were more than 25 unique auctions. As many as 87 auctions unique auctions were observed being held simultaneously in my data. This simplification risks introducing bias as it does not recognise food banks' decisions not to bid on these additional lots. This bias is similar to the possible bias in a standard tobit model from simply dropping half the censored observations. However, given that food banks rarely place more than 10 bids each period, so that I am already taking into account their decision not to bid on 15 lots, these additional observations

[2] I can also reject the hypothesis for the Ψ_i parameters at the 5% significance level, however this appears to come from a small number of food banks, and preliminary investigation suggests it may relate to non-convergence of their parameters.

are unlikely to yield much additional information. The simplification was made to speed up the convergence of my Gibbs Sampler, as the large degree of censoring will typically harm this.

To investigate robustness to this decision I estimate the model using data on 50 auctions held each day. To get around convergence problems I begin estimation from the final iteration(s) of the main model. I only observe days with more than 50 auctions on 5% of days. If I find that my results are generally robust to doubling the number of auctions considered each day, it is unlikely that including the remaining auctions will change the results either.

In Figure M.4 I present the results of Wald tests from different groups of parameters, considering how this alteration to the model changes the estimated model parameters. I can reject that the parameters have the same posterior means for an overall Wald test. I can also reject that posterior means are equal for the pseudo-static payoff function parameters Ψ_i, marginal values of wealth λ_i, and the lot specific variances. I need a larger variance to rationalise the lower probability of bidding on any given auction. For the Ψ_i parameters, as expected, a number of them failed to converge properly. This warrants additional investigation in future.

M.3 Third Stage

M.3.1 Quadratic Approximation

After evaluating the ex-ante value function across a 3.2×10^6 grid of states, I take a quadratic approximation across this grid, weighted by each states' estimated relevance. This is necessary because I must evaluate the ex-ante value function for each food bank and each draw, and I am unable to save all 34×1000 grids. A legitimate concern is whether this approximation is accurate. Figure M.5 panel (A) presents a histogram of the R^2s from forming this approximation. 100% of these value lie between 0.99 and 1. The fit is strong because of the quadratic term which appears in the equation presented in Proposition 6.

As an additional robustness test I consider whether fitting higher order polynomials improves the fit significantly. I consider up to a 4th order polynomial. Figure

M.5 panel (B) plots the results from this analysis. None of these test statistics exceeds 2, far below the critical values. This is evidence that including higher order polynomials does not yield better fit than using a quadratic approximation.

FIGURE M.5: Robustness: Stage 1

(A)

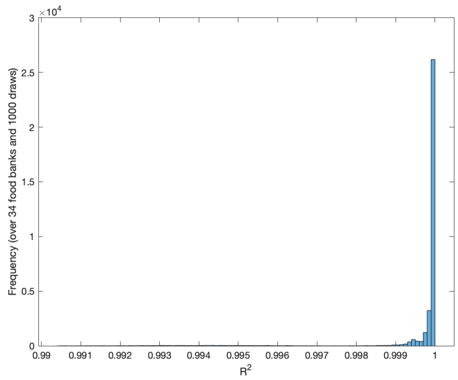

(B)

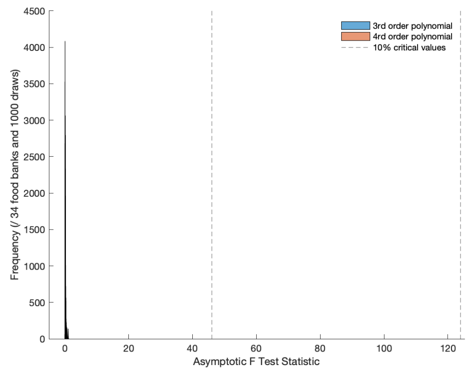

Note: The figure considers the accuracy of my quadratic approximation. Panel (A) presents R^2 statistics from the least squares quadratic fit, while panel (B) considers third and fourth order polynomials, presenting F statistics for whether the additional parameters are significant at the 10% level.

M.3.2 No Information Pooling

As discussed in Appendix K.3.3 I pool information across food banks when evaluating the ex-ante value functions. This may introduce bias, by drawing food banks' estimated flow payoffs together. I already use a first-order correction for this bias, however as an additional robust test I consider how estimated parameters vary when I do not use this information pooling. Figure M.6 panel (A) displays the results from this analysis, presenting estimates of the marginal flow payoff across food banks just as in Figure 4.2. My estimates remain in line with the previous results.

M.3.3 Sampling Variation in Means

As discussed in Appendix K.3.2 I do not take into account sampling variation in my finite sample evaluations of the expectation terms in Proposition 6. Therefore I likely underestimate the posterior variance of the flow payoffs. In this appendix I consider how results change when I employ a bootstrap resampling procedure on the estimated expectations. When estimating these means, for each draw from my posterior distribution of second stage parameters I randomly draw (with replacement) the time periods used to evaluate the means. This procedure should overestimate the posterior variance, as it does not account for covariance between the sampled parameters and the sampled time periods. Figure M.6 panel (B) displays the results from this analysis, presenting estimates of the marginal flow payoff across food banks just as in Figure 4.2. Credible intervals become somewhat larger, but not by much, and the plot remains similar. This is unsurprising given the long panel, and that I am pooling information across food banks.

FIGURE M.6: Robustness: Stage 2

(A) (B)

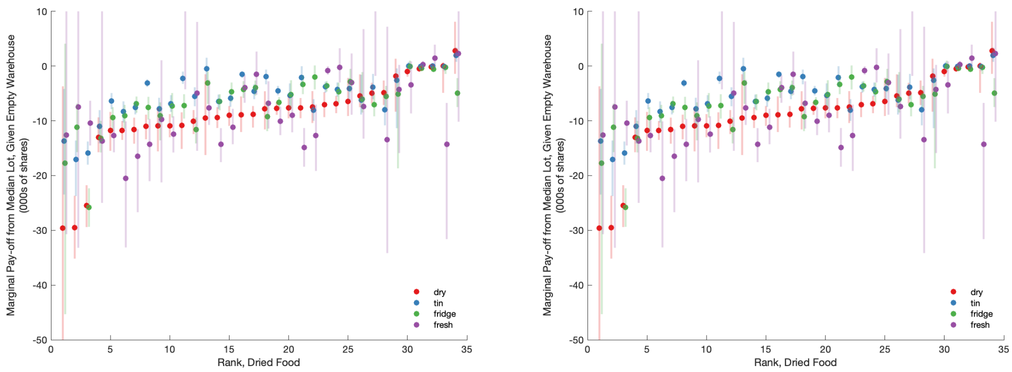

Note: The figure plots two sets of estimates for the marginal flow payoff, just as in Figure 4.2 panel (B). Here, panel (A) estimates the marginal flow payoff without pooling information on bidding behaviour across food banks, while panel (B) estimates this object taking into account sampling variation in the finite sample evaluations of the expectation terms in Proposition 6.

Appendix N

Simulation Details

In this Appendix I describe how the counterfactual simulations are performed. Appendix N.1 focuses on the Old System, detailing how I numerically solve for the equilibrium value function. Appendix N.2 outlines how I simulate the Choice System. Appendices N.3 - N.6 detail how I simulate equilibrium allocations under the remaining counterfactuals. The additional counterfactuals all use the same basic continuous time set up as the Old System. They only differ in the offer and acceptance processes. Due to computational constraints equilibria are evaluated only at the posterior means of my parameter draws. This simplification is unlikely to have a major impact on my results as equilibrium accept/reject decisions are much more strongly determined by the flow payoffs than the continuation values.

For all the counterfactuals there is a risk that stocks trend downwards for some food banks, and may trend upwards for others (particularly in the random allocation). If so, I would have to extrapolate estimated flow-payoffs into regions of the state space that were never visited under the Choice System. This is mostly a problem under the random allocation, the limited offer Old System, and the limited offer Closest mechanism. To alleviate these concerns I make two simplifications. First, if stocks exceed the highest sampled stock for a given food bank for a particular type of food, the food bank turns down all subsequent offers for that type of food until stocks return to levels within the sampled space. Likewise if stocks stray below the minimum sampled stock, all additional loads are unconditionally accepted. Second, whenever stocks exceed the maximum sampled stock (or stray below the minimum) by more than one standard deviation of sampled stocks, I do not extrapolate the flow payoff to that state. Instead, I fix the flow payoff to the minimum flow payoff from across the sampled states. Without these simplifications, payoffs are significantly lower under these three counterfactual mechanisms.

N.1 Old System

In this Appendix I detail how the simulations of the Old System are performed. I use the same procedure for both the Old System with only 10 offers, and the Old System with food offered to every food bank. I treat time as continuous, and each day is of length 1. This means I assume local donations and offers of food from Feeding America are received continuously during the day. However, to ensure that results are easily comparable across the Choice System and Old System simulations, when evaluating welfare I treat local donations as only altering stocks at the end of the day. Likewise, that flow payoffs only accrue at the end of the period. In evaluating the equilibrium value function, however, I treat both these objects as continuous.

N.1.1 Set Up

Arrivals

Food is donated to Feeding America at some exogenous rate. Conditional on arriving, the load has various characteristics. The rate and probability of these characteristics are taken from the empirical distribution. What matters in the agent's problem is their belief about the rate at which they are offered food, and the probabilities of characteristics they are offered.

Priorities

Food is offered to whichever food bank is at the head of a queue. A food bank's position in the queue is given by their rank in a priority ordering. The priority ordering at time t is given by the difference between the total amount of the food bank has been offered up to t, and that food bank's target amount at t. Food bank i's target amount at time t is given by $\frac{GF_i}{\sum_j GF_j} \times$ The total amount of food allocated up to time t.

Because the ordering is a function of the amount of food offered to food bank i, not the amount of food actually allocated to them, once a food bank is offered a load, their new priority is independent of whether they accept or reject the load.

I set initial priorities equal to long-run average priorities, with minor perturbation to ensures initial priorities differ across simulation draws. I also drop results from the first 100 days in my sample period, reducing the dependence of my results on the initial priorities.

Net Local Donations

Estimated local donations arrive at discrete intervals, however I must translate this into continuous time. I assume that local donations arrive at exogenous Poisson rate $\mathbf{q}_i$, with one element for each of the storage methods. Denote Net Local Donations at time t from storage method l by $\tilde{x}_{il}(t)$. This is non-zero with rate q_{il}. Conditional on being non-zero, this donation X_{ilt} is drawn from distribution F^X. The assumption I make on F is discussed shortly.

Daily local donations of type l are given by $\int_0^1 \tilde{x}_{il}(t)dt$. Given the assumptions of our model, this is expected to be normally distributed with mean μ_{il} and variance Σ_{il}. This is a sum of i.i.d. random variables, ensuring that the choice of F^X should not matter if q_{il} is sufficiently high so that I can apply the central limit theorem. In general, both F^X and q are not jointly identified from discrete data. However, the requirement that total net daily local donations has mean $\boldsymbol{\mu}_i$ and variance Σ_i, plus a functional form assumption on F^X, can be used to pin down F^X and $\mathbf{q}_i$.

I assume that $X_{ilt} \in \{\underline{X}_{il}, \bar{X}_{il}\}$ with probabilities $1 - r_{il}$ and r_{il} respectively. Conditional on $\underline{X}_{il}, \bar{X}_{il}$ (which I discuss shortly), I set q_{il} and r_{il} to ensure that the mean and variance of daily donations equal $\boldsymbol{\mu}_i$ and Σ_i.[1]

Payoffs and States

If food bank i accepts lot l they receive lot specific flow-payoff v_{ilt}. If they are in state $\mathbf{s}_i$, then they also receive combination specific flow-payoff $\Phi\mathbf{z}_{lt}^h + j(\mathbf{s}_i + \mathbf{z}_{lt}^g)$. If they reject the lot, they only receive flow payoff $j(\mathbf{s}_i)$. Importantly, this means that at every continuous moment in time (with density zero), the food bank receives flow payoff $j(\mathbf{s}_i)$.

[1] I also impose that $r_{il} \in [0,1]$. Average net daily donations are then given by $q_{il}(\underline{X}_{il} + r_{il}(\bar{X}_{il} - \underline{X}_{il}))$. The variance is given by $q_{il}^2 r_{il}(1 - r_{il})(\underline{X}_{il} - \bar{X}_{il})^2$. Under the additional restriction that $\bar{X}_{il} = -\underline{X}_{il}$, it can be shown that $q_{il} = \sqrt{\frac{\Sigma_{il}+\mu_{il}^2}{-\bar{X}_{il}\underline{X}_{il}}}$ and $r_{il} = \frac{\mu_{il}/q_{il}-\underline{X}_{il}}{\bar{X}_{il}-\underline{X}_{il}}$

I discretise the individual state space in the same way as done in Section 4.3.5, using a grid formed of 20 evenly spaced points from each dimension of the state. This means that accepting a lot, or receiving a local donation, can only move the state in a finite number of ways, which I now discuss.

From state $\mathbf{s}_i$ if they accept lot l their stocks would increase by $\mathbf{z}_l^g$. Therefore, for each lot $\times$ state combination I find the nearest discretised state that minimises the euclidean distance to $\mathbf{s}_i + \mathbf{z}_l^g$. This allows me to define the $20^5 \times 20^5$ matrix Z_l containing a single 1 in each row (corresponding to a particular state) in the column that corresponds to this closest state from accepting lot l.

I do a similar thing for the local donations. I set $\bar{X}_{il} = -\underline{X}_{il}$ equal to the distance between my grid points, so that with rate q_{il} the food bank moves up or down a grid point, with probabilities r_{il} and $1 - r_{il}$ respectively. This allows me to define the transition matrix Q_l containing two non-zero values (r_{il} and $1 - r_{il}$) in each row, in the columns corresponding to the states one discrete notch above and below (along dimension l) of each state.

N.1.2 Equilibrium

The Agent's Problem

Write the agent's value function as $V(t, \mathbf{s}_i, \mathbf{s}_0)$. This gives their presented discounted value from state $(\mathbf{s}_i, \mathbf{s}_0)$ at time t. I augment the common state to include the newly defined priorities and Goal Factors. If the food bank is offered a load at t they must be at the head of the queue, and so have the highest priority. If they are offered load l characterised by $(v_{ilt}, \mathbf{z}_{lt}^g, \mathbf{z}_{lt}^h)$, they will accept if $v_{ilt} + \Phi \mathbf{z}_{lt}^h + V(t, \mathbf{s}_i + \mathbf{z}_{ilt}^g, \mathbf{s}_0) > V(t, \mathbf{s}_i, \mathbf{s}_0)$.

Beliefs

The agent believes that Feeding America will offer them a load at Poisson rate $p_i(t, \mathbf{s}_0)$. In principle this should depend on the state of every food bank, including i, however I will assume that food banks do not observe each others' states. The agent then believes that, conditional on receiving an offer, the load will have characteristics $(v_i, \mathbf{z}^g, \mathbf{z}^h)$ with probability density $f_i^c(v_i, \mathbf{z}^g, \mathbf{z}^h; t, \mathbf{s}_0)$.

Equilibrium

I assume a Markov Perfect Equilibrium in symmetric strategies, as defined in section 4.2.4. This requires that food banks make optimal accept/reject decisions given their beliefs about p and f^c, and that their beliefs about p and f^c are consistent with the observed realisation of the rates at which Feeding America offers them loads.[2] As I have assumed a stationary equilibrium, I require that p and f are conditionally independent of t.

Because equilibrium value functions must be calculated over a large state space I make a number of simplifying assumptions about equilibrium beliefs. I assume food banks do not observe other food banks' stocks, nor when loads are offered to other food banks (hence aggregate supply is also unobserved). They only observe when Feeding America offers them a load. I therefore assume the only objects used to form their beliefs are $\mathbf{s}_i$, their own (relative) Goal Factor, and the time since they were last offered a load τ. I assume that f^c, the distribution of lot characteristics, depends only on GF_i. I assume that the offer rate p also depends only on GF_i. In principle I could allow p to depend on τ, however for simplicity I assume it does not.[3] I will consider this dependence as a robustness check in future. I also assume food banks beliefs do not change conditional on the previous history of offers. That is, food banks do not infer from frequent offers that offers will be more frequent in future. I will allow for this dependence in a future robustness check.

I assume parametric forms for both these objects. Broadly, for f_i^c, I split lots into the same 60 discrete category combinations used for the lot specific variances σ_l, detailed in Appendix K.2. Therefore f_i^c can be interpreted as conditional probabilities. Then, conditional on the category combination, I assume food banks believe that,

[2]One problem with assuming a stationary equilibrium is that food banks' stocks may not be stationary under a counterfactual. If $\mu_i < 0$ then all else equal their stocks trend down over time. I previously made the assumption that food banks are able to consume enough food from the Choice System to ensure their unobserved stock process remains stationary. Priors suggest the Old System is worse at allowing food banks to access the food they need. Therefore there is no guarantee their unobserved states will be stationary. However, this does not rule out the possibility of a stationary equilibrium. If stocks quickly trend down to the point where food banks accept any load they are offered, this equilibrium is still stationary. It is therefore important that I drop the first 100 days in my analysis (and equilibrium belief update), to ensure food banks are in this equilibrium.

[3]This is unlikely to be a problem, since, for most food banks, loads are offered to them so frequently it is unlikely they will ever have to wait particularly long before receiving another offer. I cannot allow f^c to depend on τ for computational reasons.

in equilibrium, the distance between the lot and a given food bank is normally distributed with some mean and variance. I also assume that, conditional on category combination, food banks believe $\Phi\mathbf{z}^h$ is also normally distributed. I then assume a multinomial logit form for the probabilities of each category combination being offered, with probabilities allowed to vary with GF_i. Finally, I assume that p_i is (logit) linear in GF_i.

N.1.3 The Optimal Control Problem

Under the assumptions outlined above, we can write the value function as a function of $\mathbf{s}_i$ and τ, and GF_i. It does not depend on t due to the stationarity assumption. For numerical convenience I absorb GF_i into the individual specific value function. Therefore write the value function as $V_i(\tau, \mathbf{s}_i)$.

Accept/Reject decision

Food bank i, that is offered load l, accepts the load if $v_{il} + \Phi\mathbf{z}_l^h + V_i(0, \mathbf{s}_i + \mathbf{z}_l^g) \geq V_i(0, \mathbf{s}_i)$. Regardless of whether they accept or reject the load, τ resets to 0.

HJB Equation

The Hamilton-Jacobi-Bellman differential equation is given by:

$$(\rho + p_i + \sum_{l=1}^{5} q_{il})V_i(\tau, \mathbf{s}_i) = p_i \int_{\mathbf{c}_l} \max\left\{v_l + \Phi\mathbf{z}_l^h + V_i(0, \mathbf{s}_i + \mathbf{z}_l^g), V_i(0, \mathbf{s}_i)\right\} dF_i^c(v_l, \mathbf{z}_l^h, \mathbf{z}_l^g)$$
$$+ \sum_{l=1}^{5} q_{il} \int V_i(\tau, \mathbf{s}_i + X)dF^X(X_l) + j(\mathbf{s}_i) + \frac{\partial V_i(\tau, \mathbf{s}_i)}{\partial \tau} \quad \text{(N.1)}$$

Where ρ gives the discount rate $(= (1-\beta)/\beta)$. To solve this differential equation, write V_i in vector form, stacking over all the possible individual states $\mathbf{s}_i$ (i.e. our discretised states). Also discretise the category combinations across c. The equation can then be written as:

$$(\rho + p_i + \sum_l q_{il})\mathbf{V}_i(\tau) = p_i\mathbf{H}_i(\tau) + \sum_l q_{il}Q_l\mathbf{V}_i(\tau) + \mathbf{j} + \nabla_\tau\mathbf{V}_i(\tau)$$

$$\text{Where} \quad H_i(\tau, \mathbf{s}_i) = \sum_c f_c^i E[\max\left\{v_l + \Phi\mathbf{z}_l^h + Z_c^g\mathbf{V}_i(0), V_i(0, \mathbf{s}_i)\right\}|c, \mathbf{s}_i] \quad \text{(N.2)}$$

Where Z_c^g gives the transition matrix defined by the pounds from a load of category combination c, and Q_l gives the transition matrix formed from the net local donations. The expectation is taken over $v_l + \Phi\mathbf{z}_l^h$. This vector differential equation does not have an analytic solution. However, recognising that $\mathbf{H}_i(\tau) = \mathbf{H}_i(0)$ it is clear that there exists a solution for $\mathbf{V}$ which is independent of τ, for which we can solve using numerical methods.

Numerical Solution

For a given $\mathbf{V}_i^k$ and beliefs (p^k, f^{ck}) I compute $\mathbf{H}_i^k$, then evaluate $\mathbf{V}_i^{k+1} = ([\rho + p_i]I + \sum_l q_{il}[I - Q_l])^{-1}(\mathbf{j} + p_i^k H_i^k)$, repeating until the magnitude of the normal vector $|\mathbf{V}_i^{k+1} - \mathbf{V}_i^k|$ is less than 1. I use these successive approximations, and switch to Newton-Kantorovich algorithm as in Rust, 1987 when progress slows. Inverting the matrix Q_l is not feasible due to its size. However multiplying by Q_l is trivial given its sparsity. Therefore in evaluating this matrix inverse, and the inverse used in Newton-Kantorovich, I use the Neumann formula for matrix inversion. This procedure generally converges in around 100 iterations.

I then simulate the Old System using these value functions, before updating beliefs. I update p_i^{k+1} by running a Poisson regression on the number of offers each food bank receives each day, conditional on goal factor, and dropping the first 100 days. I then update f^{ck+1} by estimating a multinomial logit model on the category combination that composes each offer. I repeat this process until the rates and estimated probabilities change by a total less than 10^{-4}. Beliefs converge extremely quickly, generally around 4 iterations, as value functions are relatively insensitive to beliefs.

N.2 Choice System

I now detail how I simulate the Choice System. I simulate the mechanism as described in section 2.1.2. I use these simulations both as a comparison for my counterfactuals and to assess model fit.

N.2.1 Basics

I simulate the system once for each of the 1000 posterior parameter draws. For each of these draws I use the associated draw of net donations (given by the unobserved stocks less their observed winnings). The set of objects being allocated each period is taken as given. As I observe and estimate my model on equilibrium bidding data under the Choice System, I do not need to solve for equilibrium beliefs or continuation values. Instead, estimated beliefs Γ can be used as equilibrium beliefs in my simulations, and the estimated pseudo-static payoff function $k(\mathbf{s}_i, \mathbf{s}_0)$ can be used in place of flow payoffs plus the equilibrium discounted continuation value. This approach would not be valid if I wanted to consider changes to the Choice System, such as changes in food banks' budgets.

I treat maroon pounds as exogenously determined, so I continue to not model food banks decisions to sell their local donations. I also treat joint bidding as exogenous - if a bid is placed jointly I have each food bank optimally set their bid *taking as given* the other player's bid. Food banks continue to split winnings evenly. This is a major simplification, but one that I would not expect to significantly impact the results, particularly given that joint bidding makes up a small fraction of bids. I treat discriminatory auctions of multiple loads correctly, though only allow food banks to place up to 5 bids on each set of these auctions. This is done using the adjustment to payoffs and first order conditions discussed in Appendix H.

The central problem in estimating the Choice System concerns the bidding function, as this involves a complex combinatorial problem of deciding which combination of lots to bid on. It is made more complicated by the possibility of multiple local optima.

N.2.2 The Bidding Function

I describe how I find optimal bids by first discussing how I optimise bids conditional on an entry decision $\mathbf{d}_{it}$, before discussing how I find the optimal entry decisions. I then discuss how I validate any optimised bids. The key simplification I make is assuming that Ψ_i is negative definite. This would imply that the payoff function is concave in bids and entry decisions, allowing me to exploit standard results from convex optimisation. If Ψ_i is indefinite, the problem of finding an optimum is NP-hard. In practice my sampled Ψ_is for Type 1 food banks are not always negative definite, but they are in the vast majority of cases. Even when they are not negative definite, most often they are fairly 'close' to negative definite, in that the largest (positive) eigenvalue is orders of magnitude smaller than the smallest (most negative) eigenvalue.

Conditional on an entry decision $\mathbf{d}$ I use standard interior point methods to numerically maximise payoffs subject to reservation prices. I begin the maximisation process at $b^0_{ilt} = R_l + 1$. In principal there may exist multiple local optima when Ψ_i is non-negative definite. However even for simulated Ψ_i matrices that exhibited large complementarities I was unable to find evidence of multiple optima. This is likely on account of my quadratic assumption, ensuring that for any $\mathbf{b}_{-l}$ (any vector of bids excluding b_l) payoffs are quasi-concave in b_l.

Considering every permutation of entry decisions is not feasible. Instead I find initial optimum entry decisions $\mathbf{d}^*$ using a hill climbing procedure. Under this procedure there are three options for each l, either $d_l = 0$, $d_l = 1$ and $b_l = R_l$, or $d_l = 1$ and $b_l = R_l + 1$. I begin with every $d_l = 0$ and run a hill climb until reaching a local optimum. If Ψ_i is negative definite this is guaranteed to be the global optimum, and any optimal bids found after this procedure are also guaranteed to be optimal. In simulations I found numerous occasions in which there were multiple local optima, but only in cases when the simulated Ψ_i exhibited sufficiently strong complementarities. Sampled (non-negative definite) Ψ_i do occasionally admit multiple optima, particularly when the number of desirable lots is large.

Finally, I use the First Order Conditions to 'check' my optimum. For $b_{ilt} > R_l$ I check that the partial derivatives are less than 10^{-5}. For $d_{ilt} = 0$ I check that the food bank does not strictly prefer $d_{ilt} = 1$, and for $b_{ilt} = R_l$ I check the food bank does not strictly prefer either $b_{ilt} > R_l$ nor $d_{ilt} = 0$. If Ψ_i is negative definite then solutions found by hill climb followed by numerical optimisation are guaranteed to satisfy these conditions. If either of these conditions fails (and if so, it is always one of the latter two conditions) I repeat the hill climb from this point, and repeat the process until I find a solution that does satisfy these conditions. By construction, each time I repeat this process the expected payoff increases, ensuring that this process terminates in a finite number of iterations. This occurs in a limited number of cases for Type 1 food banks only.

Two final things are worth mentioning. First, the problem of multiple optima is unlikely to be significantly impacting my results as it is only a problem for certain draws of Ψ_i. Even then the model fit is typically fairly good (at least for Type 1 food banks, for whom multiple optima is a problem). Second, food bank managers likely do not solve the full combinatorial problem, and instead likely use heuristics. It is also not impossible that they also get stuck at local optima. A hill climbing heuristic is possibly even more sophisticated than they might use in practice (as it can require many iterations to find a solution). Therefore this algorithm could be considered a reasonable approximation to their behaviour.

N.3 Random Allocation

I now detail how I perform the random allocation. This mechanism is fundamentally the same as the Old System in which food is offered to every food bank, using the same queueing system. The only difference is that food banks are not given a choice to reject the lot. The only time a food bank is not offered a lot is if their stocks are above the maximum of the stocks sampled under the Choice System. This is in order to prevent having to make large extrapolations, and keep the resulting welfare comparable to that under the Choice System and other counterfactuals.

N.4 Closest Mechanism

The closest mechanism offers food to the nearest food bank first and, in this case of the 'all offers' version, then works down food banks in order of distance. Strategically it is very similar to the Old System, except that offers (and characteristics conditional on an offer) will be much more food bank specific, rather than determined by Goal Factor. It is also much more likely that these objects do not depend on the time since the previous offer.

I follow the continuous time modelling approach used for the Old System, so that the Hamilton-Bellman-Jacobi equation remains fundamentally the same. Food banks form beliefs about the rate p_i at which they receive offers of food. Conditional on an offer, the load has characteristics $\mathbf{c}$ with probability f_i^c. As for the Old System I group food into the 60 category combinations used for the lot specific variances. Conditional on a category $\times$ food bank combination, I assume the distance between the food bank and the lot, as well as $\Phi \mathbf{z}_l^h$, is normally distributed.

As with the Old System I numerically solve the Hamilton-Bellman-Jacobi equation. For the 'single offer' Closest mechanism I can directly estimate p_i, f_i^c, and the means and variances of the normally distributed lot characteristics by considering the set of lots for which they are the closest food bank. For the 'all offer' version, given initial beliefs, I must repeatedly evaluate the value function and simulate the system until beliefs about these objects converge.

N.5 Like Mechanism

Details of the Like mechanism come from Walsh, 2015. Under the Like mechanism each load is offered to every food bank simultaneously. The load is then randomly assigned, with some probability, among the food banks that 'Liked' it. This assignment probability is given by a food banks' Goal Factor, divided by the sum of Goal Factors of the food banks which 'Liked' the load.

Once more, I model this allocation problem in continuous time. I assume food banks form beliefs about the probability that food is offered p, and the characteristics of food being offered f_i^c.[4] Neither of these objects depend on the actions of other

[4]The i subscript here is just to recognise that food banks face different distributions of distances.

food banks. I also assume they form beliefs about the probability π_i^c of winning any given lot conditional on 'Liking' it and characteristics c. I assume these equilibrium probabilities does not depend on time, nor on other aspects of the state. While food banks have more information than under the Old System, I assume they still do not observe which food bank wins the food.[5]

Under this set up, the Hamilton-Bellman-Jacobi differential equation is given by:

$$
\begin{aligned}
&(\rho + p + \sum_{l=1}^{5} q_{il}) V_i(\tau, \mathbf{s}_i) = \\
&\quad p \sum_c f^c \pi_i^c(\tau) E[\max \left\{ v_l + \Phi \mathbf{z}_l^h + V_i(\tau, \mathbf{s}_i + \mathbf{z}_c^g) - V_i(\tau, \mathbf{s}_i), 0 \right\} | c, \mathbf{s}_i] + V_i(\tau, \mathbf{s}_i) \\
&\qquad + \sum_{l=1}^{5} q_{il} \int V_i(\tau, \mathbf{s}_i + X) dF^X(X_l) + j(\mathbf{s}_i) + \frac{\partial V_i(\tau, \mathbf{s}_i)}{\partial \tau}
\end{aligned} \tag{N.3}
$$

As in the main text I will assume a symmetric Markov Perfect Equilibrium, so that V and π_i^c are independent of τ. I solve for equilibrium just as I did for the Old System.

N.6 Efficient Sequential Mechanism

Under the Efficient Sequential mechanism each load is allocated to the food bank with the highest value, or discarded if no food bank has a weakly positive marginal value. Value includes both flow payoffs and the continuation value. Strategically this mechanism is very similar to the Old System, except by construction food banks will always accept any load they are offered.

To evaluate the equilibrium value function I again assume food banks form beliefs about the rate p of food being donated to Feeding America, and the probabilities of loads coming from each category combination f^c. They then believe they have the highest marginal value for that lot with probability Γ. I assume this probability function takes the same generalised extreme value form as in the specification of beliefs for the Choice System. This object will be a function of their marginal value from winning the lot. As above, I again assume that food banks also form beliefs about

[5] It is plausible that these beliefs should depend on aggregate states, just as win probabilities did in the Choice System model. However given the small effects I found in the first stage estimation it is unlikely there would be any economically meaningful changes with the state.

the distribution of distances and $\Phi\mathbf{z}_l^h$ conditional on food from category combination c, which again I treat as normally distributed. The Hamilton-Bellman-Jacobi equation is given by:

$$(\rho+p+\sum_{l=1}^{5}q_{il})V_i(\tau,\mathbf{s}_i) = pH(\tau,\mathbf{s}_i)+\sum_{l=1}^{5}q_{il}\int V_i(\tau,\mathbf{s}_i+X)dF^X(X_l)+j(\mathbf{s}_i)+\frac{\partial V_i(\tau,\mathbf{s}_i)}{\partial\tau}$$

$$\text{Where}\quad H(\tau,\mathbf{s}_i) = V_i(\tau,\mathbf{s}_i)+\sum_c f^c E[\Gamma(B(v_l+\Phi\mathbf{z}_l^h,\tau,\mathbf{s}_i))B(v_l+\Phi\mathbf{z}_l^h,\tau,\mathbf{s}_i)|c,\mathbf{s}_i]$$

$$B(v_l+\Phi\mathbf{z}_l^h,\tau,\mathbf{s}_i) = v_l+\Phi\mathbf{z}_l^h+V_i(\tau,\mathbf{s}_i+\mathbf{z}_c^g)-V_i(\tau,\mathbf{s}_i) \quad \text{(N.4)}$$

Once more I assume a symmetric Markov Perfect Equilibrium, and fit the same empirical specification to Γ estimating food bank specific separate shape, scale, and location parameters for each category combination. Unlike in the previous mechanisms this conditional expectation does not have a closed form solution, so I find it through simulation. I then numerically solve this equation as I did for the mechanisms as described previously. I simulate the system and estimate the function Γ just as in the main text, repeating this procedure until beliefs converge.