

Ergodic Equilibria in Stochastic Sequential Games

Jeremy H. Large*

JEREMY.LARGE@OXFORD-MAN.OX.AC.UK

*Oxford-Man Institute of Quantitative Finance, University of Oxford, and
AHL, Man Investments,
Blue Boar Court, Alfred Street, Oxford, OX1 4EH, U.K.*

Thomas W. L. Norman

THOMAS.NORMAN@MAGD.OX.AC.UK

*Magdalen College, University of Oxford,
Oxford, OX1 4AU, U.K.*

April 23, 2009

Abstract

This paper characterizes, and provides existence conditions for, geometrically ergodic and stationary Markov equilibria in a type of stochastic sequential game where players have two available actions, increasing or (weakly) decreasing the integer state. For example, offering goods for sale at a market increases market-wide stock; buying decreases that stock. Similarly, joining a queue lengthens it, unlike balking. Stationarity allows us to derive the invariant strategy of players ignorant of the state, without the need for dynamic programming. Applications of this result include a microfoundation for market-clearing price adjustment.

JEL classification: C73, D41.

Key Words: Stochastic games, sequential games, ergodicity, market games.

*We thank Margaret Meyer, Nicolas Vieille and three anonymous referees for important contributions. We are also grateful for helpful comments from Joe Perkins and participants at Games 2008, and for the support of All Souls College, Oxford.

1 Introduction

Consider a market where agents arrive in random sequence and decide whether to buy or sell a unit of a single good. The agents differ in their private valuation of the good, but a buyer must be matched with a seller before trade can occur, which may require an agent to wait for a number of arrivals subsequent to his own choice.¹ This would seem an appealing description of trade, and indeed a fairly minimal departure from a traditional Marshallian market. And yet, solving for equilibrium in such a market is a considerable challenge: the optimal choice of each agent depends, not just on his valuation, but on the number of buyers and sellers present when he arrives, and on his forecast of subsequent arrivals and choices.

The market is an example of a stochastic game (1953), where rules can evolve through time and depend on past actions. The richness of this powerful and general framework naturally raises issues of tractability. Even time-homogeneous equilibria in Markovian strategies—whose dynamic simplicity has made them the focus of existence theorems in the area (e.g. Shapley 1953, Parthasarathy 1982, Nowak 1985, Rieder 1979)—can be difficult to find. Duffie, Geanakoplos, Mas-Colell, and McLennan (1994) point out that it is natural to seek equilibria where the state variable is furthermore ergodic, and prove their existence under various conditions.

In this paper we show, in a particular class of stochastic *sequential* games (Fudenberg and Tirole 1991), that while seeking this ergodicity makes proving existence more involved, it subsequently enhances tractability. More specifically, we provide conditions such that there exist ergodic equilibria whose stationarity allows us to derive the invariant strategy of certain players without the need for dynamic programming. Our class of stochastic sequential games has the following properties: heterogeneity in the players' types; an ordered state space that a binary action set, $\{up, down\}$ say, can increase or decrease; and imperfect observability of the current state. These are applicable in the above market game, where we provide a microfoundation for market-clearing prices; and in the context of a queue, where we find invariances in balking choices when service opportunities are persistent, which provide a rationale for not excessively censoring queue length information.²

The tractability of this framework is derived from the presence of 'uninformed' players, by which we mean players who do not observe the current state (or at least cannot act upon

¹We remain quite agnostic as to the exact mechanism used in the market to match counterparties.

²We would anticipate that a number of further applications may be available, such as (after transformation) the dynamic limit order book game in Large (2009).

it). These players have a complicated forecasting problem, since they must consider the continuation games beginning in every state in the support of the ergodic measure. However, we show that this problem is an average of the problems faced by all players, so that they select an average of all players' strategies. Under the conditions of our model, this average strategy is known *a priori*, as the only one consistent with a zero-drift condition that obtains once we know that the state is ergodic.³

To prove that equilibrium exists we first use a familiar single crossing condition from the literature on monotone comparative statics, obtaining an interim equilibrium existence result (among players who observe the state) which exploits duality with a result of Rieder (1979). We then introduce two conditions—called *Self-Regulation* and *Separable Preferences*—that allow us to prove that a stationary Markov equilibrium exists where the state is geometrically ergodic and stationary. Self-Regulation implies that as the state becomes arbitrarily high (low), on average it becomes attractive for players to decrease (increase) it. Intuitively, in the context of the market game, when the number of buyers far exceeds the number of sellers, choosing to sell becomes more attractive. Using results in Tweedie (1983), this gives the existence of an ergodic measure with a first moment. Separable Preferences, meanwhile, is a strengthening of the single crossing condition; it requires that the effect of a player's type on her payoff can be separated from that of the state in an additive or a multiplicative way (or a hybrid of the two).

In equilibrium, agents foresee the consequences of their actions on future dynamics. Solving for equilibrium typically replicates this, potentially complicated, inference. The key simplification exploited in the paper is to deduce best responses in an easier indirect way that bypasses entirely this dynamic forecasting problem: Under the single crossing condition players use monotone cutoff strategies; then, under Separable Preferences and an assumption about the distribution of types (which implies that it is uniform in our queuing and market applications), uninformed players play *up* and *down* in the *same* ratio as the overall population. Using the zero-drift condition, this finally identifies the cutoff chosen by uninformed players.

This result has interesting invariance and welfare consequences. First, the uninformed best response is invariant to alterations in the economic environment, including changes in the

³Since our focus is on uninformed players, it may appear *prima facie* that informed players appear in the model only incidentally. However, their presence is in fact a necessary ingredient for the state to be ergodic or stationary.

observability of the state, in impatience levels, and in many aspects of equilibrium dynamics. In essence, equilibrium effects precisely compensate for and defeat various appealing intuitions about how to change uninformed players' incentives and behavior. Second, in some examples (including our queuing application), because we know which type of agent is indifferent between *up* and *down*, this leads to an expression for uninformed expected surplus, which has the same invariances as those just mentioned. It also places a lower bound on welfare, since uninformed players do weakly worse than average.

In Markov equilibrium of our market game then, the state—which includes the last traded price—adopts a stationary distribution. Therefore it induces buying (via the action *up*) and selling (via *down*) in the ratio required for the stock of the good to be stationary, thereby clearing the market and equating expected supply to demand. Thus, price does not settle on any single market-clearing level, but rather occupies a market-clearing *distribution* of levels. Nevertheless, we show that our model gives a microfoundation for an intuitive ideal market-clearing price level (see Gale 1985, 2000, and Rubinstein and Wolinsky, 1985). Specifically we show, under an *even-handedness* condition, that the ideal market-clearing price is at the mean of our game's equilibrium distribution of prices; and we argue that an efficient market has a low variance of prices around that mean.

Our second application is to the classic queueing model of Naor (1969). We are able to characterize the case where service opportunities resemble durable goods and can persist without decay until consumed. Specifically, our modelling technique allows us to solve for the best response and surplus of players who, like in Edelson and Hilderbrand (1975) and Hassin (1986), must decide irrevocably to join a queue before observing its length. These turn out to be invariant to a wide range of toll levels and queuing disciplines.

The paper is organized as follows: Section 2 lays out the model; Section 3 then deals with equilibrium existence and characterization; Section 4 outlines and solves the market-clearing example; Section 5 develops the queueing application; and Section 6 concludes. Detailed proofs are held in an Appendix.

2 The Model

Let there be an infinite sequence of players, indexed by $i \in \mathbb{N}$, that arrive in order $\{1, 2, 3, \dots\}$ (for example, consider the arrival of agents at a queue). Player i independently draws and

observes her type, θ_i , from a distribution with continuous, increasing, and differentiable CDF F , of support $[0,1]$ (for example, her valuation of service at the queue). Let ω_i be a state variable at the time of the i th player, that inhabits a state space representable with integer value states, $\Omega \subseteq \mathbb{Z}$ (e.g. queue length). Let $\omega_1 \in \Omega$ be drawn from an initial distribution μ_1 . For times, i , in $\mathbb{N}$, let $\{d_i\} \subset \mathbb{Z} \times \mathbb{N}$ represent a bounded i.i.d. stationary time series, which will describe a possible drift in ω_i ; and let $\bar{d} = E[d_i]$ (d can represent the arrival of service opportunities at a queue). Write $\mathcal{F}_i$ for the entire history of all random variables up to and including time i .

The following sequence of events occurs for each player in turn:

1. Player i receives the signal s_i . With probability $\pi \in (0, 1]$ this informs her of the history of the game prior to her arrival, so that $s_i = (\mathcal{F}_{(i-1)}, \omega_i)$. Otherwise $s_i = \emptyset$, where $\emptyset$ is the empty signal, and player i is *uninformed*. Hence, the state is imperfectly observable.
2. Conditioning on i , θ_i and s_i , she then makes an action, $a_i \in \{\underline{a}, \bar{a}\} \subset \mathbb{Z}$.

We assume that $\underline{a} < \bar{d} < \bar{a}$; and that $\bar{a}$ is in the support of d_i .⁴

3. For all $h = 1, 2, \dots, i - 1$, player $(i - h)$ receives a payoff as a (measurable) function,

$$f_h(\theta_{i-h}; \{(\omega_j, a_j) : (i - h) \leq j \leq i\}), \quad (1)$$

of her type and the play since time $(i - h)$.

4. Player i does not act again.
5. The next state, ω_{i+1} , evolves according to the rule

$$\omega_{i+1} = \omega_i + a_i - d_i. \quad (2)$$

In addition to the structure of stochastic games in general then (see, e.g., Mertens 2002), our setup has incomplete information, a countable, ordered and imperfectly monitored state, and sequential binary actions that either increase or decrease the state. Because $|d_i|$ is bounded, we may define Y to be a uniform upper bound on $|\omega_{i+1} - \omega_i|$.

All players have the same time preferences, represented by the same discount factor per period, $\delta \leq 1$. Player j therefore has payoff

$$f_0(\theta_j; \omega_j, a_j) + \sum_{k=1}^{\infty} \delta^k f_k(\theta_j; \cdot),$$

⁴So $\bar{a}$ corresponds to *up*, as mentioned in the Introduction, and $\underline{a}$ is *down*.

where ‘ $\cdot$ ’ contains the history of actions and states from time j to time $(j + k)$, inclusive. Suppose that there exists an upper bound, $R > 0$, such that for all h , R bounds the range of $|f_h|$ above. Assume that if $\delta = 1$ then the number of h such that f_h is non-zero is uniformly bounded above.

Under this condition, for any Lebesgue measurable σ , the expectation

$$V_\sigma(\theta_j; s_j, a_j) = \mathbb{E} \left[f_0(\theta_j; \omega_j, a_j) + \sum_{k=1}^{\infty} \delta^k f_k(\theta_j; \cdot) \mid s_j, \theta_j \right] \quad (3)$$

exists.⁵ So player j may choose a_j so as to maximize it.

2.1 Equilibrium

We will be interested in symmetric equilibria where all players employ the same strategy, σ , so that for all i , $a_i = \sigma(s_i, \theta_i, i)$.⁶ In equilibrium, it is the case that for all θ_i and s_i , $V_\sigma(\theta_i; s_i, \dagger)$ is maximized by setting $\dagger = \sigma(s_i, \theta_i, i)$. We shall further confine our attention to equilibria in Markov and stationary strategies, where $\sigma(s_i, \theta_i, i)$ turns out to be invariant to i , and to $\mathcal{F}_{i-1}$, so that it only depends on the current state, ω_i . To shorten notation, in the stationary Markov equilibria whose existence and properties we consider in detail, we therefore write σ as a function of only s_i and θ_i , i.e. $\sigma(s_i, \theta_i)$; and we drop the information $\mathcal{F}_{i-1}$ from s_i : so that $s_i = \omega_i$ or $\emptyset$. Thus

$$a_i = \begin{cases} \sigma(\omega_i, \theta_i) & : \text{player } i \text{ is informed} \\ \sigma(\emptyset, \theta_i) & : \text{player } i \text{ is uninformed.} \end{cases}$$

2.2 Ergodicity

A stationary Markov equilibrium induces Markov transition probabilities in the process $\omega := \{\omega_i\}_{i \in \mathbb{N}}$, as given by:

$$P^n(x, y) := \Pr(\omega_{1+n} = y \mid \omega_1 = x).$$

If $\pi < 1$, then we must solve for the uninformed agents’ best response $\sigma(\emptyset, \theta_i)$. This task is greatly simplified if we can establish that ω is stationary. To show this we consider whether ω is ergodic, specifically whether

$$\lim_{n \rightarrow \infty} |P^n(x, y) - \mu(y)| = 0,$$

⁵Because of uninformed players, the conditional expectation must encompass both first and second parts of this expression.

⁶This concept of symmetry is rather weak since σ depends on the rank, i , of the player.

so that the measure μ is then the *ergodic distribution* of ω . If we then draw the initial state from μ (i.e. set $\mu_1 = \mu$), ergodicity of ω implies that μ is the stationary distribution of ω , so that for any state $\omega' \in \Omega$, uninformed players believe they are in state ω' with probability $\mu(\omega')$.

2.3 Monotonicity

We will be interested throughout in conditions under which each player uses a *monotone pure strategy*. This is to say that he choose his action according to a *cutoff* $c(s_i)$ in his type:

$$\sigma(s_i, \theta_i) = \underline{a} \Leftrightarrow \theta_i \leq c(s_i).$$

Such cutoffs are implied by there being an ordering on types, where lower types prefer $\underline{a}$, all other things being equal; and higher types prefer $\bar{a}$. In the case of a queue where players value service independently, for example, players with higher reservation value for service join a queue of given length; while players with lower value for service do not. The next section will propose conditions under which such monotonicity obtains.

3 Existence and Properties of Equilibrium

In this section, we first provide conditions for the existence of equilibrium, and then characterize the equilibrium strategies of uninformed players, for whom $s_i = \emptyset$. These players need particularly careful attention; hence, we begin by fixing their behavior so that we may defer treatment of their decision problem.

Given this restriction, together with a single crossing property in payoffs, we consider a dual game and use a result in Rieder (1979) to prove the existence of a stationary Markov equilibrium in cutoff strategies among the informed players but not the uninformed, where the state ω_i is Markov.

Under a condition of Self-Regulation, we are then able to demonstrate that ω_i is ergodic and potentially stationary, whereupon uninformed players' decision problems are constant. We close the argument for equilibrium existence by introducing Separable Preferences. If this condition holds, and a further assumption about the distribution of types is satisfied, then uninformed players' best response is to act precisely as we had restricted them to act from the start. The section closes with some remarks about the welfare and invariance properties of equilibrium.

3.1 Single Crossing Property and Cutoffs

Definition 1 *Let the zero-drift quantile of F , denoted τ , be defined by:*

$$\tau = F^{-1} \left(\frac{\bar{a} - \bar{d}}{\bar{a} - \underline{a}} \right).$$

Consider a player, i , who uses cutoff equal to the zero-drift quantile τ : that is, $c(s_i) = \tau$. It is quite easily shown that

$$E[a_i | c(s_i) = \tau] = \bar{d},$$

so that a_i is on average $\bar{d}$.⁷ From (2), this has the consequence that the conditional expectation of $(\omega_{i+1} - \omega_i)$ is zero, and ω has zero conditional drift. So we might say that in ignorance of this player's type, we would expect her to have zero effect on the level of the state. We will show that, in stationary Markov equilibrium, uninformed players use the zero-drift quantile τ as their cutoff, and therefore do not in isolation cause the state to drift.

The first step in showing this is to establish that the use of cutoffs (i.e. monotone pure strategies) is optimal, for which purpose we can exploit the lattice-theoretic results of Milgrom and Shannon (1994). Note that, for informed players and a given state, ω' , the objective function V maps from the product space $[0, 1] \times \{\underline{a}, \bar{a}\}$ to the real line. Note also that $\{\underline{a}, \bar{a}\}$ is a lattice, while $[0, 1]$ is a partially (indeed totally) ordered set with the natural order on $[0, 1]$.

Condition 1A (Strict Single Crossing Property, Milgrom and Shannon 1994) *For every pair of types $\theta' > \theta''$, and given any strategy profile, σ , we have that: $V_\sigma(\theta''; \omega, \bar{a}) \geq V_\sigma(\theta''; \omega, \underline{a})$ implies that $V_\sigma(\theta'; \omega, \bar{a}) > V_\sigma(\theta'; \omega, \underline{a})$.*

From the definition of V_σ , this implies that if playing the high action $\bar{a}$ is optimal for informed type θ'' , then in the same state it is strictly optimal for the higher type θ' .

Remark 1 *In applications we show that Condition 1A holds by appeal to the structure of f_h , $h \in \mathbb{N}$ (and without reference to σ , which may vary widely).*

For example, f_h might separate into functions $f_h^\dagger$ and $f_h^\sharp$, so that

$$f_h(\theta_{i-h}; \{(\omega_j, a_j) : (i-h) \leq j \leq i\}) = f_h^\dagger(\theta_{i-h}) \times f_h^\sharp(\{(\omega_j, a_j) : (i-h) \leq j \leq i\}), \quad (4)$$

⁷This can be seen from the following equalities:

$$\Pr(a_i = \underline{a} | c(s_i) = \tau) = \Pr(\theta_i < \tau) = F(\tau) = \frac{\bar{a} - \bar{d}}{\bar{a} - \underline{a}}.$$

where $f_h^\dagger$ is increasing; and $f_h^\sharp$ is non-negative; and f_h is not always trivially zero. Condition 1A then follows.

Lemma 1 *Under Condition 1A, all best responses imply monotone pure strategies.*

Proof. Since $\{\underline{a}, \bar{a}\}$ is totally ordered, V_σ is (trivially) quasisupermodular in a_j . Hence, under strict single crossing, $\arg \max_{a \in \{\underline{a}, \bar{a}\}} V_\sigma(\theta; \omega, a)$ is monotone nondecreasing in θ by the Monotone Selection Theorem 4' of Milgrom and Shannon (1994). Moreover, indifference has measure zero under the Strict Single Crossing Property. ■

As we will later see, Condition 1A can be reasonable in applications. However, it rules out non-monotone behavior, which can also be of great interest but is not studied here.

For the rest of the paper, we will take Condition 1A as holding. The next lemma takes an important step towards proving equilibrium existence for this game.

Lemma 2 *Restrict players to act such that $c(\emptyset) = \tau$ (so uninformed players act non-strategically). Then (among the informed players) there exists a stationary Markov equilibrium. The induced stochastic process, ω , is Markov (but may not be stationary, or ergodic).*

Proof. The proof uses a result in Rieder (1979), and considers a dual game where players select cutoffs rather than actions. See Appendix A.1. ■

3.2 Self-Regulation and the existence of ergodic equilibrium

As uninformed players' actions have been fixed by the cutoff τ , the question of whether their decision problem is well-defined has not yet arisen. In fact, in the equilibrium of Lemma 2, they differ from informed players because in general they have a measure over the model's random variables only if they know their arrival time, i . This rules out the existence of a stationary Markov equilibrium which includes strategic uninformed players.

This section first provides a weak condition under which the process ω has ergodic distribution μ , then shows that setting the distribution of ω_1 to be μ makes ω stationary. Like informed players, uninformed players then face a stationary decision problem which does not depend on i .

Ergodicity is therefore a desirable condition, but establishing it requires some work. We have an *irreducible* process—for, any given pair of states has positive probability of eventually interchanging. We also have a (*strongly*) *aperiodic* process, since $P^1(x, x) > 0$ for some $x \in \mathbb{Z}$.

This is the case because $\bar{a}$ is in the support of d_i . But given the infinite state space these properties are not sufficient to establish ergodicity.

However, there is a simple “drift condition” (Tweedie 1976) (stemming from Foster 1953), for countably infinite-space Markov chains to be ergodic. This can be shown to hold in our model under Condition 2 below, but in fact we go further, proving *geometric ergodicity*: namely, that there are $\rho < 1$ and $k_{xy} < \infty$ such that, for all x, y, n ,

$$|P^n(x, y) - \mu(y)| \leq k_{xy}\rho^n.$$

A geometrically ergodic process converges to its ergodic distribution at a uniform geometric rate. This is reassuring, in that ergodicity does not require an unreasonably long time to become relevant. It also provides a rationale for the assumption we will make that the distribution of the initial condition is equal to μ , i.e. that the process is stationary: for, regardless of the initial condition, this stationary process is approached in distribution by the ergodic process at a geometric rate.

A necessary and sufficient drift condition for geometric ergodicity was provided by Popov (1977): Suppose there is a function $g \geq 1$ over states, a number $q < 1$ and a finite set A of states such that

$$\mathbb{E}(g(\omega_2) \mid \omega_1 = x) \leq qg(x), \quad x \in A^c,$$

and

$$\max_{x \in A} \mathbb{E}(g(\omega_2) \mid \omega_1 = x) < \infty.$$

We can show that this requirement is satisfied in equilibrium under the next condition.

Condition 2 (Self-Regulation) *We will say that a game (and, by extension, its equilibria) are Self-Regulating if there exists a (small bound) $\lambda > 0$ and a (possibly wide) interval $A = (\underline{\omega}, \bar{\omega})$ such that if player i is informed, so she observes ω_i , then:*

- *when $\omega_i > \bar{\omega}$, it is dominant for any type $\theta_i < (\tau + \lambda)$ to use action $\underline{a}$. Likewise,*
- *when $\omega_i < \underline{\omega}$, it is dominant for any type $\theta_i > (\tau - \lambda)$ to use action $\bar{a}$.*

Recalling the definition of τ as the zero-drift quantile, a Self-Regulating game is one where, whenever the state is sufficiently high (resp., low), the next agent’s best response, together with d_i , will on average cause it to fall (resp., rise). This is the case regardless of the strategy profile, σ , and hence it is a property of the game as a whole. More precisely, it is a limiting property of the game, because Self-Regulating behavior only manifests itself above and below

an arbitrarily wide interval $(\underline{\omega}, \bar{\omega})$ of the state space. Because it is a limiting property, it is reasonably easy to demonstrate in our applications, as we later do.

At this stage, we concern ourselves with showing that, in equilibrium of a Self-Regulating game, the state has sufficient centric drift for Popov's condition to hold. For the results of the next subsection to hold, however, we will require slightly more than simple ergodicity: namely, that the ergodic distribution has a first moment. Fortunately, having established geometric ergodicity, we can exploit another drift condition of Tweedie (1983) to give us the existence of a first moment.

Lemma 3 *Restrict players to choose a strategy σ such that $c(\emptyset) = \tau$ (so uninformed players act non-strategically), and fix ω_1 . Under Self-Regulation, ω is geometrically ergodic, and the ergodic first moment, $\int x d\mu(x)$, exists.*

Proof. See Appendix A.2. ■

Proposition 1 *Restrict players to choose σ such that $c(\emptyset) = \tau$; then under Self-Regulation a monotone pure-strategy stationary Markov equilibrium exists, where ω has stationary distribution μ with a first moment.*

Proof. Consider the equilibrium in Lemma 2. Lemma 3 shows that under Self-Regulation ω has ergodic measure μ , whose first moment exists. Now suppose that the initial condition ω_1 is chosen randomly from the measure μ : as ω_1 could have been fixed at any value, this does not disrupt the optimality of players' strategies. Then, ω is stationary, with stationary distribution μ . ■

3.3 Separable Preferences and characterizing the uninformed best response

This section contains the central result of the paper, Theorem 1. The theorem identifies conditions under which the restriction in Proposition 1, that $c(\emptyset) = \tau$, may be lifted while retaining the existence of equilibrium. In fact, it goes further, showing that that doing so leaves τ as the optimal cutoff for uninformed players.

The parts of this section proceed as follows: first we derive some properties of the zero-drift quantile; we then provide a preliminary result to motivate the main theorem; next we provide a new condition, Separable Preferences; before using it in the next part to prove the main theorem. We conclude with a discussion of the theorem's implications; and with a discussion of its robustness.

3.3.1 Properties of the zero-drift quantile, τ

Consider an equilibrium where ω has stationary distribution μ with a first moment. Then the law of ω_i is constant as i varies; and we have the following important property: $E[\omega_i] = E[\omega_{i+1}]$ for all i . Put otherwise, over time ω_i must increase just as much as it decreases. So from (2), for all i ,

$$E[a_i] = \bar{d}.$$

Using this fact, we have the following lemma:

Lemma 4 *In a monotone pure-strategy stationary Markov equilibrium where ω has stationary distribution μ with a first moment,*

$$\tau = F^{-1} E[F(c(s_i))]. \quad (5)$$

This is the case whether or not uninformed players' actions are restricted so that $c(\emptyset) = \tau$.

Proof. $E[a_i] = E[E[a_i|s_i]] = \underline{a} E[F(c(s_i))] + \bar{a}(1 - E[F(c(s_i))])$; equate to $\bar{d}$. (Note that s_i equals $\emptyset$ if player i is uninformed, and equals ω_i if player i is informed). ■

This lemma shows that for equilibrium to be ergodic with first moment, an average of players' cutoffs must be the zero-drift quantile, τ , of F . Unless F is linear, a deformation in the average is induced by the cumulative distribution function F . However, in the case of a uniform distribution, where F is an increasing linear function, no such deformation arises, and (5) reduces to the identity $\tau = E[c(s_i)]$. This will be the case in the dynamic market example of Section 4 and the queuing game in Section 5.

3.3.2 A preliminary limiting result

By way of an introduction to the main idea, we now provide a corollary which draws a consequence from Lemma 4 in the limit where the proportion of uninformed players converges to 1. First note that (by an accounting identity)

$$E[F(c(s_i))] = \pi E[F(c(\omega_i))] + (1 - \pi)F(c(\emptyset)). \quad (6)$$

As $\pi \downarrow 0$, the expression $(1 - \pi)F(c(\emptyset))$ dominates the right-hand side of (6), so that given Lemma 4, in the limit $c(\emptyset) \rightarrow \tau$. Consequently:

Corollary 1 *Consider an infinite sequence of stationary sequential equilibria under the conditions of Proposition 1, where ω has stationary distribution μ with a first moment. Suppose that in the limit of this sequence $\pi \downarrow 0$; then in the limit*

$$c(\emptyset) \rightarrow \tau.$$

So we have shown that in equilibria where a large majority of players are uninformed, they choose a cutoff near to the zero-drift quantile, τ . This is intuitive: if they did otherwise they would induce drift in the state, ruling out the ergodicity that is implied by Self-Regulation.

3.3.3 Separable Preferences

While the limiting result of Corollary 1 is useful, two features are unsatisfactory: first, the uninformed best response is only identified in the limiting case where $\pi \downarrow 0$; second, the existence of the sequence of unrestricted equilibria is not proven. To resolve both of these difficulties, we strengthen Condition 1A so that payoffs can be given more structure, and we relate this structure to the shape of the players' type distribution, as given by F .

On the one hand, strengthening Condition 1A allows us to cash-out formally the intuition that uninformed players, because their decision problem is the average of the problems faced by all players, select an average of other players' cutoffs. On the other hand, relating this to the shape of F allows us to control the nature of this average (whether it is algebraic, geometric, etc.) in such a way that it coincides with the type of average in Lemma 4 discussed above, thereby giving us the identification that we seek.

Condition 1B (Separable Preferences) *Assume that there exists a function J such that, for any strategy profile σ , there exists a representation of preferences in terms of further functions H_σ and G_σ such that*

$$V_\sigma(\theta_i; \omega_i, \bar{a}) - V_\sigma(\theta_i; \omega_i, \underline{a}) = H_\sigma(\theta_i)[J(\theta_i) - G_\sigma(\omega_i)],$$

for all i , where $H_\sigma : [0, 1] \mapsto \mathbb{R}^{++}$; and $J : [0, 1] \mapsto \mathbb{R}$ is increasing and continuously differentiable. Further, $\text{Range}(G_\sigma) \subset \text{Range}(J)$. Normalize J , G_σ and H_σ so that $J(0) = 0$ and $J(1) = 1$.

Comments about the Separable Preferences condition This condition implies Condition 1A, and hence implies by Lemma 1 that optimal strategies are monotone pure strategies.

The market game example of Section 4 provides an example satisfying Condition 1B. If H_σ is constant in θ_i (as it will be in the market game) then payoffs separate additively into a contribution due to the type, $J(\theta_i)$, and a contribution due to the state, $G_\sigma(\omega_i)$. On the other hand, if J is constant, then payoffs separate multiplicatively between these two effects. If both J and H_σ vary, then a richer hybrid relationship is possible. In our experience, it has often not been possible to find G_σ at equilibrium, but this is not a hindrance to our solution strategy.

It is important to stress that the functions G_σ and H_σ are dependent on the strategy profile σ . While noting this, it will help in the upcoming exposition to suppress their subscripts.

Under Separable Preferences, in response to any σ , an expected utility maximizer prefers $\underline{a}$ iff $J(\theta_i) < G(\omega_i)$. This is so iff $\theta_i < J^{-1}(G(\omega_i))$ (inverting J). Therefore

$$c(\omega_i) = J^{-1}(G(\omega_i)). \quad (7)$$

Meanwhile, the expected benefit to an uninformed player i of $\bar{a}$ over $\underline{a}$, conditional on θ_i , is

$$H(\theta_i)(J(\theta_i) - \mathbb{E}[G(\omega_i)]). \quad (8)$$

So uninformed players would use a cutoff to determine action if their behavior were unrestricted.

3.3.4 The main theorem: lifting the restriction that $c(\emptyset) = \tau$

To consider what uninformed players would do, if free, we introduce a piece of notation: let $c^*(\emptyset)$ be the unrestricted optimal choice of cutoff for an uninformed player, i . In principle, $c^*(\emptyset)$ may differ from the restricted cutoff $c(\emptyset) = \tau$. From (8) we have that

$$c^*(\emptyset) = J^{-1} \mathbb{E}[G(\omega_i)]. \quad (9)$$

So, combining (7) and (9),

$$c^*(\emptyset) = J^{-1} \mathbb{E}[J[c(\omega_i)]]. \quad (10)$$

This evidently bears a close similarity to the unrelated result of Lemma 4 (which stemmed from the stationary state), that

$$\tau = F^{-1} \mathbb{E}[F[c(s_i)]]. \quad (11)$$

The equality in (11) describes (without pinning them down) informed players' best responses when uninformed players select the cutoff $c(\emptyset) = \tau$; whereas (10) describes the best responses of uninformed players when informed players use cutoffs $\{c(\omega_i) : \omega_i \in \Omega\}$.

Under Separable Preferences then, (10) formalizes the intuition that the uninformed player prefers an average cutoff because she faces the average problem. Like the average in Lemma 4, this average is deformed, but by the function J rather than the CDF F .

The next theorem will show that setting the background distribution F equal to J equates the nature of these ‘deformed’ averages and causes them to coincide: so uninformed players use as a cutoff the quantile τ , and their best response is identified—without solving their forecasting problem.⁸ This would be immediate if both (10) and (11) referred to s_i . However, as (10) refers to ω_i in place of s_i , there remains work for the theorem to do (see the proof).

Theorem 1 *Suppose that Conditions 1B (Separable Preferences) and 2 (Self-Regulation) apply, and that $F \equiv J$. Setting μ_1 to the ergodic distribution μ , a monotone pure-strategy stationary Markov equilibrium exists.*

While equilibrium may not be unique, it is the case that in all equilibria

$$c(\emptyset) = \tau.$$

Proof. Restrict players to choose a strategy σ such that $c(\emptyset) = \tau$. By Proposition 1, a monotone pure-strategy stationary Markov equilibrium exists, where ω has stationary distribution μ with a first moment. To complete the proof, we show that in this equilibrium the restriction $c(\emptyset) = \tau$ is not binding because it is the optimal choice for uninformed players. Formally, we show that $c^*(\emptyset) = c(\emptyset)$.

From Lemma 4, $E[F[c(s_i)]] = F(\tau)$. Replacing J with F in (10) gives $F(c^*(\emptyset)) = E[F[c(\omega_i)]]$. Finally we have imposed the restriction that $F(c(\emptyset)) = F(\tau)$. Substituting these three equalities into (6) gives

$$F(\tau) = \pi F(c^*(\emptyset)) + (1 - \pi)F(\tau). \quad (12)$$

As $\pi > 0$ it then follows that $c^*(\emptyset) = \tau$. ■

Remark 2 *Note that μ_1 is set equal to the ergodic distribution μ for convenience, to give immediate stationarity. Were the initial condition set another way, say at a fixed value, then by Lemma 3 this ergodicity would be approached at a geometric rate. Indeed, Theorem 1 could*

⁸There may at first sight appear to be a circularity in this: uninformed players’ behavior can be characterized because it is the average of all players’ particular behavior; however, to know all players’ particular behavior, one must already know uninformed players’ particular behavior. Equation (12) illustrates how the apparent dilemma is bypassed: it encapsulates the putative circularity, in that $F(\tau)$ is present on both sides of the equation. Nonetheless, it can be solved for $F(c^*(\emptyset))$ provided $\pi > 0$.

also be stated as holding in the time-limit as a not-necessarily stationary state approached its ergodic distribution. Some other robustness properties of Theorem 1 are discussed in Subsection 3.3.6.

3.3.5 Implications of Theorem 1

Invariance to the game's dynamics Under Conditions 1B and 2, it has been possible to identify the cutoff employed by uninformed players ($c(\emptyset) = \tau$), while bypassing any study of the dynamic properties of the model. Therefore, the best response of uninformed players is invariant to a range of features of the dynamic environment: in particular it is invariant to the primitives δ and π , as well as to the endogenous dynamics of equilibrium as given by the function G and the ergodic measure μ .

Uninformed welfare In some settings, one can go further: inferring from the nature of the uninformed cutoff the expected *ex ante* surplus of uninformed players, again without resolving equilibrium dynamics. This can be informative about the welfare properties of the game. The following lemma shows that informed players do better than uninformed players, so that the level of surplus of uninformed players provides a lower bound on welfare per player:

Lemma 5 *Suppose that Conditions 1B and 2 apply and that $F \equiv J$. In a stationary Markov equilibrium, the expected surplus of uninformed players does not exceed that of informed players.*

Proof. An informed player can choose to ignore the signal, s_i , in which case she can obtain expected surplus equal to that of uninformed players. She can also choose not to ignore the signal, and may thereby do better. ■

The key to characterizing uninformed players' *ex ante* surplus is to exploit the fact that the type τ is indifferent between $\underline{a}$ and $\bar{a}$, when uninformed. The following proposition shows that in order to derive the uninformed surplus, it is enough that the payoff from only one of the two actions be known or easy to derive. This might be the case, for example, if one action were the act of not participating in the game, carrying a payoff of zero.

Proposition 2 *Suppose that Conditions 1B (Separable Preferences) and 2 (Self-Regulation) apply, and that $F \equiv J$. If type θ_i 's expected payoff from action $\underline{a}$ is a constant, $\bar{U}$, then*

$$\bar{U} \leq \mathbb{E}[H(\theta_i)\{J(\theta_i) - J(\tau)\}^+]$$

is the equilibrium *ex ante* expected surplus of an uninformed player. This is invariant to the primitives δ and π , as well as to the dynamics of the game as given by the function G and the ergodic measure μ . It is exceeded by welfare per player.

In this proposition, $\{X\}^+$ means $\max\{X, 0\}$.

Proof. See Appendix A.3. ■

3.3.6 Robustness of Theorem 1

These results draw on a distributional assumption linking dynamic preferences (via J) to the distribution of types (via F). However, because J and F are continuously differentiable, we can prove the following corollary, giving some reassurance that Theorem 1 is robust to perturbations in J or F :

Corollary 2 *Consider a sequence of equilibria such that players are restricted to choose σ such that $c(\emptyset) = \tau$ and where in each case $\mu_1 = \mu$. Suppose that Conditions 1B (Separable Preferences) and 2 (Self-Regulation) apply throughout.*

If in this sequence of equilibria $(J - F)$ converges uniformly to zero on $[0, 1]$, then in the limit $c^(\emptyset)$ converges to τ .*

Proof. In each of the sequence of equilibria, by (10) and Lemma 4,

$$J(c^*(\emptyset)) - F(\tau) = E[(J - F)[c(\omega_i)]]. \quad (13)$$

Hence $(J(c^*(\emptyset)) - F(\tau))$ converges to zero as J converges uniformly to F . Because J and F are continuous and monotonic, then also $(c^*(\emptyset) - \tau) \rightarrow 0$ as $(J - F) \rightarrow 0$. ■

4 A Market with Sequential Trade

We turn now to applications of the general results. The first of these studies the dynamic centralized market with which we introduced the paper, deriving a microstructure foundation for the classical prediction that in such a competitive market, the price settles near to its market-clearing level. We begin from a simple definition of this level which we call $\hat{p}$. After using Theorem 1 to solve for the uninformed best response, welfare is analyzed, and it is shown that efficiently-designed markets have transaction prices (net of frictional asymmetries) that are close to $\hat{p}$.

This section adds to a number of game-theoretic foundations for competitive equilibrium that have been explored since Cournot’s (1838) famous limiting result. Notably, Shapley–Shubik games (Shubik 1973, Shapley and Shubik 1977) show equilibrium convergence to Walrasian allocations as the number of agents becomes arbitrarily large in a general-equilibrium setting. This involves a centralized, simultaneous allocation of goods via a bidding mechanism, and implicitly a single price for any given good. In *dynamic matching and bargaining games* (Rubinstein and Wolinsky 1985, Gale 1995, 2000), by contrast, agents engage in random search for trading partners, with whom they then bargain. Such models can give delicate convergence to competitive Marshallian allocations in the limit as the players become arbitrarily patient and/or frequent. Our market game occupies a similar Marshallian setting, but has sequential trade in a centralized market institution as opposed to pairwise bargaining. Moreover, we find competitive market-clearing, not in any limit, but rather as a property of the stationary distribution of the market state and prevailing price.

Section 5 provides a second application, to a queue for a durable good or service. Here we can go further than in the market game in one direction: by providing a simple closed-form expression for uninformed players’ *ex ante* surplus.

Consider again then our dynamic Marshallian market which trades only one good, and where informed players decide whether to buy or sell a unit quantity of the good on observing the state of the market, while uninformed players must make this decision in ignorance of the market state.

Players arrive at the market sequentially at times $1, 2, 3, \dots$. Each player, i , is informed of her reservation value for the good, θ_i , which she draws independently from an identical uniform distribution of CDF, F . This distribution’s support is normalized to $[0, 1]$. She may either buy or sell one unit of the asset. In addition, at each time i a natural number d_i of further units of the good arrive exogenously at the market—from a factor of production for the good, say—which for simplicity we suppose to be of zero cost. The sequence $\{d_i\}$ is a bounded non-negative i.i.d. random process such that $\Pr(d_i = 1) > 0$ (ensuring we have an aperiodic process). Writing $\bar{d} = E[d_i]$, we assume that $\bar{d} < 1$, which implies that goods do not arrive at the market weakly faster than they can be consumed.

At price $p \in (0, 1)$ the expected demand, D , for the good per unit time is

$$D(p) = \Pr[\theta_i \geq p],$$




Figure 1: Marshallian benchmark

so that $D(p) = 1 - p$. Adding together the exogenous supply of the good with the endogenous selling decisions of players, expected supply per unit time, S , is given by

$$S(p) = \bar{d} + p.$$

Note that all players belong to a single population, but may contribute either to the demand or to the supply for the good; hence $D(p) + S(p)$ is invariant to p .

4.1 Marshallian benchmark

Before specifying further the model dynamics, we consider a ‘flow equilibrium’ benchmark (Gale 1985), where the market-clearing price is the one that equates expected demand per unit time with supply per unit time. This permits an analysis of the extent to which equilibrium dynamics depart, if at all, from the case without any market frictions.

Figure 1 illustrates how supply and demand schedules intersect where $D(p) = S(p)$, and gives the Marshallian market-clearing price, denoted $\hat{p}$. Laying out a few properties of this equilibrium: $\hat{p} = \frac{1}{2}(1 - \bar{d})$; a proportion $\hat{p}$ of players sell and a proportion $(1 - \hat{p})$ buy; and

the average surplus per trade is the shaded area in Figure 1 to the left of both demand and supply schedules, of size $(\frac{1}{2} - \hat{p}^2)$.

4.2 Payoffs and frictions in the dynamic market

Since players do not arrive simultaneously at the market, sometimes players must wait to trade. Because of opportunity cost, waiting effort, or impatience to trade, this wait is costly for market participants. To model this we introduce a friction cost, $f(t)$, which enters additively into players' payoffs if they wait a time t to trade. Assume that $f \geq 0$ is increasing, convex and bounded, and that $f(0) = 0$.

The payoffs are determined as follows: if player j buys at price p at time i she gets a payoff of

$$\theta_j - p - f(i - j); \tag{14}$$

while if she sells she gets a payoff of

$$p - \theta_j - f(i - j).$$

4.3 Sequence of moves

Let B_i denote the number of buyers waiting to trade at time i , and let S_i indicate the stock of the good awaiting sale at time i : so S_i consists of those goods provided by sellers (one unit each), together with those goods provided by the exogenous factor. The state of the market at that time, ω_i , contains the number of buyers and goods still waiting for trade:

$$\omega_i = (B_i, S_i),$$

and ω_1 is a random variable on $(\mathbb{N} \times \mathbb{N})$.

The following events take place at time i :

1. Player i receives the signal s_i . With probability $\pi \in (0, 1]$ this informs her of the entire history of the game. Otherwise, $s_i = \emptyset$, where $\emptyset$ is the empty signal, and player i is *uninformed*.
2. She makes an action a_i : she either becomes a buyer ($a_i = +1$) or a seller ($a_i = -1$).
3. A quantity d_i of the good arrives at the market from the exogenous factor.

A *market mechanism* matches some or all available buyers to goods, and determines corresponding purchase prices which may differ from one another. Admissible traded prices lie within $(0, 1)$, on an interval $[p^-, p^+] \subseteq (0, 1)$

At each time i , whether and at what price each player trades are measurable with respect to events subsequent to that player arriving at the market.

4. Player i does not act again.
5. The next state, $\omega_{i+1} = (B_{i+1}, S_{i+1})$, gives the number of remaining buyers and goods at the market.

Expected utilities from the two actions (+1, buying; and -1, selling) are given by

$$V_\sigma(\theta_i; s_i, +1) = E[\theta_i - p_i^b - f(T_i^b)|s_i], \quad (15)$$

$$V_\sigma(\theta_i; s_i, -1) = E[p_i^s - \theta_i - f(T_i^s)|s_i], \quad (16)$$

where p_i^b is the random variable giving the price at which player i would buy if she became a buyer, and T_i^b gives the time from i until her trade (p_i^s and T_i^s are the equivalent quantities for selling).⁹

4.4 Parameter restrictions

For the conditions of our general model to hold in this setting, frictional costs must be of comparable magnitude to realized prices. This is so that they can become large enough to help regulate the market, but not so large that they dominate the decision-making of players, so destabilizing the market. Thus, f must not be too great, so that

$$f(\infty) < \min(p^-, 1 - p^+). \quad (17)$$

However, f must also not be too small, so that $f(\infty) > \frac{1}{2}(p^+ - p^-)$.

Finally, the range of admissible prices, $[p^-, p^+]$ must surround the market-clearing price $\hat{p}$ quite closely, so that $\hat{p}$ is nearer than a distance $f(\infty)$ from that interval's boundaries:

$$p^- + f(\infty) > \hat{p} > p^+ - f(\infty). \quad (18)$$

⁹This market game can be viewed as a model of portfolio managers trading equities for idiosyncratic purposes, where for simplicity traded quantities are fixed. If tactical order submission decisions are made not by the managers but rather by traders who act for them as agents, then the managers themselves are left with a binary decision in any given period whether to buy or sell, as in the current case.

4.5 Solution: uninformed best response

Given the last subsection, it can be checked that this dynamic market game has the properties of the more general stochastic sequential game presented in Section 2, but with a richer, bivariate state variable ω_i . The zero-drift quantile, τ , of Definition 1 has the property that if $c(s_i) = \tau$ (so that player i uses the cutoff τ) then $E[a_i] = \bar{d}$. In this case, it is easy to show the following.

Lemma 6 *The zero-drift quantile of the dynamic market game, τ , is given by $\tau = \hat{p}$.*

Define the notation $\Delta f_i = f(T_i^s) - f(T_i^b)$, which gives the relative waiting frictions of selling versus buying, and let

$$m_i = \frac{p_i^b + p_i^s}{2},$$

be the ‘midquote’ faced by player i at the time of trade.

The next proposition provides conditions under which the equilibrium existence and characterization results of Section 3 can be directly applied. In particular, it specializes to market mechanisms using a first-in-first-out (FIFO) discipline, where buyers (sellers) who arrive earlier at the market buy (sell) earlier.

Proposition 3 *Suppose that the market mechanism matches counterparties as quickly as possible so that for all i*

$$\min\{B_i, S_i\} = 0,$$

and suppose that it uses a FIFO discipline. Then there exists a distribution for ω_1 such that a monotone pure-strategy stationary Markov equilibrium exists, and $c(\emptyset) = \hat{p}$.

Proof. If $\forall i \min\{B_i, S_i\} = 0$, then ω_i can also be represented without loss of information as $(B_i - S_i)$. With this representation it obeys the law of motion that

$$\omega_{i+1} = \omega_i + a_i - d_i.$$

We may write

$$V_\sigma(\theta_i; s_i, +1) - V_\sigma(\theta_i; s_i, -1) = 2 \left\{ \theta_i - E \left[m_i - \frac{1}{2} \Delta f_i \mid \omega_i \right] \right\},$$

so that the game conforms to Separable Preferences, with $H(\theta_i) \equiv 2$; and $J(\theta_i) = \theta_i$. Therefore $J \equiv F$. Also,

$$G(\omega_i) = E \left[m_i - \frac{1}{2} \Delta f_i \mid \omega_i \right]. \tag{19}$$

From (17), the image of G is inside that of J . Furthermore, when $\omega_i > 0$ then $T_i^s = 0$. Since the market is FIFO, $T_i^b \geq \omega_i$. Therefore $\lim_{\omega_i \rightarrow \infty} \Delta f_i = f(\infty)$, and

$$\lim_{\omega_i \rightarrow \infty} G(\omega_i) > [p^- + f(\infty)] > \hat{p}.$$

By Lemma 6, $\hat{p} = \tau$. Therefore, after rearranging, there exists a $\lambda > 0$ such that $c(\omega_i) > (\tau + \lambda)$. As a similar condition holds for large negative ω_i , ω satisfies Self-Regulation.

So the conditions of Theorem 1 are satisfied and the result follows. ■

As well as providing an existence result, this result identifies the best response of uninformed players to the dynamic market: it shows that if their reservation value for the asset, θ_i , is greater than the *theoretical* market-clearing price, $\hat{p}$, as derived in the flow equilibrium benchmark, then uninformed traders buy, but if it is lower than $\hat{p}$, they sell. This cutoff, which determines whether they buy or sell, does not depend on realized prices at all.

4.6 A direct argument to our main result

The proof of Proposition 3 draws heavily on the notation and results of Section 3; however, a more direct argument for it in this context can be organized as follows: Call the market *self-regulating* if the market mechanism causes the market state process ω not to explode but to have an ergodic distribution. Under the conditions of Proposition 3 the market is self-regulating, but there are other self-regulating markets which for example do not use a FIFO discipline. In any such case, the process $(B_i - S_i)$ has an ergodic distribution. Note that, for any i ,

$$\mathbb{E}[(B_{i+1} - S_{i+1}) - (B_i - S_i)] = \mathbb{E}[a_i] - \bar{d}.$$

For ω to be ergodic, $(B_i - S_i)$ should have zero drift, meaning that $\mathbb{E}[a_i] = \bar{d}$. But then,

$$\Pr[\underline{a}] = \frac{1}{2}(1 - \mathbb{E}[a_i]) = \frac{1}{2}(1 - \bar{d}) = \hat{p}.$$

So, players are induced to sell a proportion $\hat{p}$ of the time; and buy on a proportion $(1 - \hat{p})$ of occasions.

But now consider uninformed players only. We can show that they *also* sell a proportion $\hat{p}$ of the time. This can be the case only if $c(\emptyset) = \hat{p}$ —which identifies the best response of uninformed players to this complex dynamic game, regardless of the precise self-regulating mechanism implemented by the market, and regardless of the form of the friction function, f .

To see this, consider the following: Each possible signal, s_i , is associated with a cutoff, $c(s_i)$, such that player i buys iff $\theta_i > c(s_i)$ and sells otherwise. From the linearity of (15) and (16), it follows that

$$c(\emptyset) = E[c(\omega_i)].$$

So, uninformed players play the *average* of the informed best responses, because they face the average of their problems. It then follows that

$$c(\emptyset) = E[c(s_i)],$$

where $s_i \in \{\emptyset, \omega_i\}$ is the signal given to player i . So uninformed players play the *overall* average best response. But, averaging across cutoffs, $\Pr(\underline{a}) = E[c(s_i)]$, which (as observed above) is $\hat{p}$, because $(B_i - S_i)$ has zero drift. Consequently,

$$c(\emptyset) = \hat{p}.$$

4.7 Realized prices in a self-regulating market

The last two subsections showed that the market-clearing price in flow equilibrium, $\hat{p}$, is central to the decision-making of uninformed players. This subsection builds on that result to establish that $\hat{p}$ also has an intimate relationship to *realized* prices in the self-regulating dynamic market game.

Let us denote by $\bar{p}$ the quantity $E[m_i]$. As $\bar{p}$ is the average midquote faced by uninformed players, we may think of it as the ‘uninformed midquote’. Additionally, let us call a market *even-handed* if in equilibrium there is no frictional advantage to uninformed selling over buying, i.e. if $E(\Delta f_i)$ is zero. When a self-regulating market is even-handed, it follows from (19) that $c(\emptyset) = E(G(\omega_i)) = \bar{p}$. But from Proposition 3, $c(\emptyset) = \hat{p}$. Therefore, the average midquote faced by players, $\bar{p}$, is the market-clearing price in flow equilibrium, $\hat{p}$:

$$\bar{p} = \hat{p}.$$

So, in an even-handed self-regulating market, if it were the case that the uninformed midquote, $\bar{p}$, exceeded (was exceeded by) the market-clearing price $\hat{p}$, then $(B_i - S_i)$ would have downwards (upwards) drift, contradicting the stationarity of ω .

It is fairly immediate to show that even if the self-regulating market is not even-handed,

$$\bar{p} = \hat{p} + \frac{1}{2}E[\Delta f_i].$$

So for the market price—defined as the uninformed midquote—to deviate systematically from the frictionless market-clearing price of the flow equilibrium, $\hat{p}$, it must be the case that the market distributes frictional costs to buyers and sellers in uneven amounts.

4.8 Welfare and surplus per trade in a self-regulating market

Recall that the surplus per trade in the flow equilibrium benchmark is given by

$$\frac{1}{2} - \hat{p}^2.$$

This includes the surplus obtained from the exogenous factor, which produces $\bar{d}$ goods at zero cost, but sells them for $\hat{p}$. Because welfare is invariant to cash transfers, another way of describing welfare per trade, is by

$$E[\theta_i|\bar{a}] \Pr[\bar{a}] - E[\theta_i|\underline{a}] \Pr[\underline{a}],$$

which gives the utility gained by buyers who acquire the asset, balanced against the utility lost by sellers who relinquish it.

Moving from this to the dynamic market game, we add the expected frictional cost per player, say $\bar{f}$, to obtain a similar expression for surplus per trade:

$$E[\theta_i|\bar{a}] \Pr[\bar{a}] - E[\theta_i|\underline{a}] \Pr[\underline{a}] - \bar{f}.$$

Appendix A.4 shows that this may be written as

$$\left(\frac{1}{2} - \hat{p}^2\right) - (\bar{f} + \text{Var}[c(s_i)]). \quad (20)$$

Expression (20) shows that welfare per trade is below that of the flow-equilibrium benchmark, $(\frac{1}{2} - \hat{p}^2)$. Welfare losses relative to this benchmark arise from two components: from waiting frictions, $\bar{f}$; and from ‘erroneous’ transactions where a player buys when it would be efficient for her to sell, or vice versa. This latter component is represented by $\text{Var}[c(s_i)]$. In an ideal case, player i would sell iff $\theta_i < \hat{p}$. Otherwise, players i and j exist where $\theta_i < \theta_j$ but (because they face differing market conditions), i chooses to buy while j chooses to sell. This eventuality does arise, because cutoffs deviate from $E[c(s_i)]$, i.e. from $\hat{p}$. The extent of the deviation can be described by $\text{Var}[c(s_i)]$.

Therefore, noting that here $G(s_i) = c(s_i)$, an efficiently-designed market is one with low

$$\bar{f} + \text{Var}[G(s_i)].$$

From (19), in an even-handed market $G(s_i)$ is interpretable as the expected mid-quote of player i . So in an efficiently-designed market, players' expectations about prices (specifically about mid-quotes) fluctuate in the vicinity of $\hat{p}$.

5 A Queue with Durable Service Opportunities

This section applies our general analysis to the case of a queue. Once queue-length has settled down into a stationary distribution we are in a position to characterize welfare invariances. This allows us: 1) to develop results regarding Hassin (1986)'s 'second best problem', which compares a system whose players can observe queue length before joining the queue, to a system where they can't; and 2) to analyze the level of an efficient toll for the queue.

We follow Edelson and Hilderbrand (1975) and Hassin (1986) in studying queues where some players must decide irrevocably to join before observing the state of the system. We focus on these uninformed queuers. Uninformed queuers have a surplus which is invariant to the queue's observability, as well as to the level of toll. Thus, attempts to enhance uninformed surplus through improved observability, or by changing the queue's toll level or indeed discipline (provided this continues to lead to a self-regulating queue) are defeated by equilibrium effects which emerge from the endogenous dynamics of the queue as it returns to stationary behavior.

In fact, under our assumptions, uninformed surplus depends on only one primitive parameter: namely, the system's utilization factor, which is the ratio of the arrival rate of service opportunities to the arrival rate of players. This provides a lower bound on welfare, while an ideal 'frictionless' benchmark gives an upper bound. By restricting welfare to this range we identify welfare-improving measures that can be taken by a toll-collecting agency and by the administrators of the queue.

5.1 Modelling assumptions; durable service opportunities

In our model, players either queue for service, or balk from the queue, departing with zero payoff; and there is no reneging from the queue. In line with Naor (1969), Edelson and Hilderbrand (1975), and Hassin (1986) there are additive waiting costs and a toll may be charged to players who queue. Like Chen and Frank (2001) our model allows for players who are heterogenous in their valuation of service.

However, our setting differs from these papers because service opportunities are not

wasted if they arise when no buyer is present: rather they resemble durable goods, so that they can persist without decay at the head of the queue until matched with a queuer.

Since otherwise they would accumulate unsustainably ultimately rendering the queue irrelevant, we assume that service opportunities arise less often than players arrive, and are therefore rationed by the queue.

In our setting, partial observability ensures that the queue length is stable (‘Self-Regulating’). On the other hand partial unobservability gives us tractability, because it introduces uninformed players like those of Theorem 1. So rather than, as Hassin (1986), comparing full observability to a full absence of observability, we investigate the continuum of intermediate cases where the proportion $(1 - \pi) \in (0, 1)$ of uninformed players can be varied.¹⁰

5.2 Model of a queue

Adapt/reinterpret the dynamic market game of Section 4 in the following five ways:

1. Service opportunities arise when a good arrives. They are always seized if a queuer is present. So, at all times,

$$\min\{B_i, S_i\} = 0. \quad (21)$$

2. The ‘*down*’ action corresponds to balking, so it leaves the state of the game unchanged:

$$\underline{a} = 0.$$

Players have no goods to sell and must simply leave if they do not wish to become a buyer.

3. If player i balks by selecting $a_i = 0$, then she gets a payoff of zero; $V_\sigma(\theta_i; s_i, 0) = 0$.
4. All purchases of the good happen at the same price, say $p^* \in (0, 1)$. This price has the interpretation of being a *toll*, which is charged to each player as they acquire a unit of the good (see Naor 1969).

¹⁰In Hassin (1986), even if information on the queue’s state is always suppressed, there exists a toll level for which equilibrium exists. Because this equilibrium has the property that players join the queue less often than service opportunities arise, it would not be in steady state here: service opportunities, which we assume persist at the queue until used, would accrue unsustainably. By contrast, in our setting complete suppression of queue length data leads to instability. Nevertheless, we show that ergodic equilibrium does exist if information about the state is *partially* suppressed, so that some but not all players do observe queue length before deciding to join.

5. Frictional waiting costs can rise sufficiently high for the queue's length to be self-regulating; but when there is no friction cost because the queue is empty, buyers arrive often enough for the stock of service opportunities not to accumulate arbitrarily. We provide a formalization of this condition below in (23), which replaces (18).

We can relate these adaptations to Figure 1: they shift the supply curve, $S(p)$, so that it is vertical and inelastic: $S(p) = \bar{d}$. It follows that, returning momentarily to the notation of the previous section, $\hat{p} = (1 - \bar{d})$.

Since time is implicitly normalized so that players arrive once per unit of time, the quantity $\bar{d}$ has the interpretation of being the *utilization factor*: it describes the expected proportion of queuers who get service.¹¹

It is straightforward that, as before, the zero-drift quantile of the queuing game, τ , is equal to the same quantity as $\hat{p}$, namely here

$$\tau = (1 - \bar{d}). \quad (22)$$

With this notation in place, conditions for Adaptation 5 can be given by:

$$p^* < (1 - \bar{d}) < p^* + f(\infty). \quad (23)$$

Interpretation of (23) These inequalities place constraints on the level of the toll, p^* so as to ensure stability in the queue. Were p^* greater than $(1 - \bar{d})$, then no player with $\theta_i \leq (1 - \bar{d})$ would choose to queue for the good. Queuers per unit time would therefore arrive at a rate slower than the goods, which arrive at rate $\bar{d}$, meaning that the stock of the good could accumulate arbitrarily—a contradiction of the premise that the queue is self-regulating. The other inequality in (23) indicates that for individual players facing a long queue, waiting frictions can exceed $(1 - \bar{d} - p^*)$, which, as will be shown in the proof of the next proposition, is a great enough threshold to imply Self-Regulation so that the queue cannot grow in length unsustainably.

5.3 Equilibrium existence and characterization

Proposition 4 *Suppose that the queue uses a FIFO discipline. Then there exists a distribution for ω_1 such that a monotone pure-strategy stationary Markov equilibrium exists, and*

¹¹An alternative semantics defines the utilization factor to be $\frac{1}{\bar{d}}$, the number of potential queuers per utilized service opportunity, so that a queue which rations service is over-utilized with utilization factor greater than 1. Our results are amenable to either semantics.

$$c(\emptyset) = \tau.$$

Proof. We may write

$$V_\sigma(\theta_i; s_i, +1) - V_\sigma(\theta_i; s_i, 0) = \theta_i - (p^* + E[f(T_i^b)|\omega_i]).$$

So we have Separable Preferences with $H(\theta_i) \equiv 1$ and $J(\theta_i) = \theta_i$. The proof then follows that of Proposition 3 closely. $\blacksquare$

As in the dynamic market game, so in this queuing model if ω is to be ergodic, then $(B_i - S_i)$ must have zero drift, meaning here that $E[a_i] = \bar{d}$. We can again imagine that queuing disciplines other than FIFO may generate ergodicity. But in any stationary and ergodic queue,

$$\Pr[\bar{a}] = E[a_i] = \bar{d}.$$

So, players are induced to buy on a proportion $\bar{d}$ of occasions; and balk a proportion $(1 - \bar{d})$ of the time. The proposition establishes that, as in the dynamic market game, uninformed players *also* balk a proportion $(1 - \bar{d})$ of the time. This then implies directly that

$$c(\emptyset) = (1 - \bar{d}).$$

Note that the uninformed player's decision to join the queue is *invariant* to the level of the toll, p^* , levied on players at service: if uninformed player i 's type, θ_i , is less than $(1 - \bar{d})$, then she balks, while if $\theta_i > (1 - \bar{d})$ then she joins the queue. Her action is therefore also invariant to the dynamic features of the queue, such as might result from variation in the queuing discipline or in the functional form of the waiting cost f .

5.4 Surplus, welfare and efficiency

Using Proposition 2 we can establish that uninformed expected surplus in this queuing game is $\frac{1}{2}\bar{d}^2$, being given by

$$\bar{U} + E[H(\theta_i)\{J(\theta_i) - J(\tau)\}^+] = E[(\theta_i - \tau)^+] = \frac{1}{2}\bar{d}^2,$$

since here the payoff from balking is zero: $\bar{U} = 0$. So, uninformed welfare is also invariant to the choice of toll, p^* , as well as to the degree to which queue length information is suppressed, π . As advertised in the introduction to this section, it depends only on the utilization factor, $\bar{d}$.

Waiting costs undergone by uninformed queuers The proof of Proposition 2 exploits the indifference of the marginal type—who draws θ_i equal to the zero-drift quantile, τ —to *up* versus *down*. Here, the indifference is between joining the queue and balking for zero payoff. So this implies that a queuer with $\theta_i = \tau = (1 - \bar{d})$ expects zero surplus from queuing. Consequently, writing $\bar{f}_u$ for the mean frictional cost undergone by an uninformed queuer,

$$(1 - \bar{d}) - p^* - \bar{f}_u = 0.$$

This identifies $\bar{f}_u$, while bypassing any characterization of the queuing dynamics that entail it:

$$\bar{f}_u = (1 - \bar{d} - p^*).$$

Choosing a more efficient toll Unlike in the dynamic market game, in a queuing setting players can formulate their decision problem as being to minimize expected waiting costs (since their valuation of the good, net of the toll, is known to them). By Lemma 5 average waiting costs, as given by $\bar{f}$, are exceeded by $\bar{f}_u$, and therefore are exceeded by $(1 - \bar{d} - p^*)$.

Because waiting costs are the only source of inefficiency in this queuing game, we arrive at an interesting consequence: Consider any choice of toll, p_0^* , and consequent mean equilibrium frictions per queuer $\bar{f}_0$. Then a more efficient equilibrium may be found by selecting a higher toll level, p^* , so that $(1 - \bar{d} - p^*) < \bar{f}_0$. Raising the toll turns out to be efficient in this setting. However, as an empirically-grounded argument this reasoning has a limit: for, as previously noted, the stock of the good approaches explosive, non-stationary behavior, as p^* approaches close to $(1 - \bar{d})$ from below.

6 Conclusion

In sum then, we have provided conditions under which a stochastic sequential game has an ergodic equilibrium whose stationarity can be exploited to simplify analysis of equilibrium behavior and welfare.

The usefulness of this simplification has been illustrated in the context of a market game and a queuing model. The former gave a model of sequential trade where price adjustments sort players between buying and selling in the precise proportions required for market-clearing. This is a model, not of an explicit market-clearing equilibrium *price*, but of an implicit market-clearing equilibrium price *process*. The system is thus free to generate rich short-run price dynamics in the process of generating long-run market-clearing.

The second application allowed us to understand queues for durable goods or services. In this model players act so as to induce zero drift in the queue length. We showed that uninformed players, who cannot observe the queue length, also induce zero drift in deciding whether to queue. This identifies the uninformed best response, and indirectly uninformed surplus—both of which are invariant to a wide range of levels of the queue’s toll and observability, as well as to queuing disciplines.

We anticipate that our general results can be applied to a variety of other examples. Indeed, Large (2009) exploits the results to model the market microstructure of trade on a limit order book, a popular mode of exchange for financial assets. There are also a number of possible avenues for further research concerning our general model. For example, one might further specify the model to the point where the ergodic distribution of the state could be explicitly obtained, and the implications for short-run price dynamics investigated.

A Appendix

A.1 Proof of Lemma 2

Consider the dual game where informed players' actions are to select cutoffs, $c \in [0, 1]$. After acting, player i 's type θ_i is revealed, and she automatically plays $\underline{a}$ if $\theta_i \leq c$ and $\bar{a}$ if $\theta_i > c$. As in the original game, actions are made conditional on $s_i = (\mathcal{F}_{(i-1)}, \omega_i)$. The dual game evolves otherwise just as the original game. Since player j does not observe her type θ_j , she does not observe the payoff in (3). Rather, let her maximize

$$\mathbb{E}[V_\sigma(\theta_j; \omega_j, a_j) \mid s_j],$$

the expected value of her payoff across possible types. This 'expects out' player types, so that θ_j is not a state variable. Finally, take ω_1 as a given.

Rieder (1979) proves the existence of a stationary equilibrium with Markov strategies when there are a countable number of players, a countable state space, and a compact action space. Interpreting the state variable to be $s_i = (\mathcal{F}_{(i-1)}, \omega_i)$, this result can be applied to this stochastic game, giving equilibrium existence.

Since player i is not indifferent to her choice of cutoff by Strict Single Crossing, and since previous values of a_i and ω_i do not enter her stream of future payoffs (as determined by $\{f_j : j \in \mathbb{N}\}$), such values do not affect her action. So she conditions action only on ω_i . Therefore the process $\{\omega_i\}$ is Markov.

Moreover, the optimal cutoff for player j ,

$$\begin{aligned} \arg \max_{c \in [0,1]} \mathbb{E}[V_\sigma(\theta_j; s_j, a_j) \mid s_j = s] &= \arg \max_{c \in [0,1]} \int_0^1 V_\sigma(\theta; s, a_j) dF(\theta) \\ &= \arg \max_{x \in [0,1]} \left\{ \int_0^x V_\sigma(\theta; s, \underline{a}) dF(\theta) + \int_x^1 V_\sigma(\theta; s, \bar{a}) dF(\theta) \right\}, \end{aligned}$$

is clearly the critical type where $\arg \max_{a \in \{\underline{a}, \bar{a}\}} V_\sigma(\theta; s, a)$ changes from $\underline{a}$ to $\bar{a}$ under Strict Single Crossing. But this is just the optimal cutoff in the original game where agents select actions after their type θ_i is revealed. There is thus a 1–1 mapping between equilibria of the original and the dual games under Strict Single Crossing, so that Rieder's result gives existence in our case also. $\blacksquare$

A.2 Proof of Lemma 3

Since the game is Self-Regulating, in the equilibrium of Lemma 2 there exists a (small bound) $\lambda > 0$ and a (possibly wide) interval $A = (\underline{\omega}, \bar{\omega})$ such that whenever $\omega_i > \bar{\omega}$, we have that

$c(\omega_i) > (\tau + \lambda)$. Likewise, $\omega_i < \underline{\omega} \Rightarrow c(\omega_i) < (\tau - \lambda)$.

Tweedie (1983) (specialized to the countable case) provides the following sufficient condition for the existence of a first ergodic moment of the absolute value of the state (which implies that the first moment itself exists): There is a set A with $\mu(A) > 0$ and $\sum_{x \in A} x\mu(x) < \infty$, and some function g with $g(x) \geq |x|$, $x \in A^c$,

$$\sum_{y \in A^c} P(x, y)g(y) \leq g(x) - |x|, \quad x \in A^c,$$

and

$$\max_{x \in A} \sum_{y \in A^c} P(x, y)g(y) < \infty.$$

To exploit this condition, we first establish geometric ergodicity (and thus the existence of the relevant positive-measure set A) using Popov's condition.

Suppose that $x > Y$ and $g(x) = e^{\alpha|x|}$ for some $\alpha > 0$ which we will later choose conveniently. (A similar argument applies to the case where $x < -Y$.) For any such x ,

$$\mathbb{E}(e^{\alpha|\omega_2|} \mid \omega_1 = x) = e^{\alpha\omega_1} \mathbb{E}(e^{\alpha(\omega_2 - \omega_1)} \mid \omega_1 = x) \leq q e^{\alpha\omega_1} \quad (24)$$

$$\Leftrightarrow \mathbb{E}(e^{\alpha(\omega_2 - \omega_1)} \mid \omega_1 = x) \leq q. \quad (25)$$

Taking a Taylor expansion of the LHS, the condition becomes

$$\mathbb{E}(1 + \alpha(\omega_2 - \omega_1) + (\alpha^2/2)(\omega_2 - \omega_1)^2 + \dots \mid \omega_1 = x) \leq q. \quad (26)$$

Since F is continuously differentiable and increasing, Self Regulation implies that there exists a $\gamma > 0$ and a (possibly wide) interval $A = (\underline{\omega}, \bar{\omega})$ such that whenever $\omega_i > \bar{\omega}$, we have that $F(c(\omega_i)) > F(\tau) + \gamma/(\bar{a} - \underline{a})$. Likewise, $\omega_i < \underline{\omega} \Rightarrow F(c(\omega_i)) < F(\tau) - \gamma/(\bar{a} - \underline{a})$.

Using this γ , for x sufficiently high,

$$\mathbb{E}(\omega_2 - \omega_1 \mid \text{informed}, \omega_1 = x) = \mathbb{E}(a_0 - d_0 \mid \text{informed}, \omega_1 = x) \quad (27)$$

$$\leq (F(\tau) + \gamma/(\bar{a} - \underline{a}))\underline{a} + (1 - F(\tau) - \gamma/(\bar{a} - \underline{a}))\bar{a} - \bar{d} \quad (28)$$

$$\leq ((\bar{a} - \bar{d} + \gamma)\underline{a} + (\bar{d} - \underline{a} - \gamma)\bar{a})/(\bar{a} - \underline{a}) - \bar{d} \quad (29)$$

$$= -\gamma, \quad (30)$$

by (2) and as ω satisfies Self-Regulation. Furthermore, $\mathbb{E}(\omega_2 - \omega_1 \mid \text{uninformed}, \omega_1 = x) = 0$. Therefore, by the Law of Iterated Expectations,

$$\mathbb{E}(\omega_2 - \omega_1 \mid \omega_1 = x) \leq -\pi\gamma. \quad (31)$$

Also note that from the definition of Y ,

$$\mathbb{E}((\omega_2 - \omega_1)^2 | \omega_1 = x) \leq Y^2. \quad (32)$$

In this case, a sufficient condition for (26) to hold is

$$1 - \alpha\pi\gamma + (\alpha^2/2)Y^2 + \cdots \leq q, \text{ or} \quad (33)$$

$$-\pi\gamma + \alpha Y^2/2 + \cdots \leq (q - 1)/\alpha. \quad (34)$$

For any $\alpha > 0$ we can choose $q < 1$ such that

$$-\pi\gamma/4 \leq (q - 1)/\alpha.$$

Fixing $\alpha < \pi\gamma/Y^2$, we then have

$$-\pi\gamma + \alpha Y^2/2 + \cdots < -\pi\gamma + \pi\gamma/2 + \cdots < -\pi\gamma/4 \leq (q - 1)/\alpha.$$

Therefore, we can choose a set of states A such that the required inequality holds for some $q < 1$ and $x \in A^c$. Since $\max_{x \in A} \sum_{y \in A^c} P(x, y)g(y) < \max_{x \in A} e^{\alpha(x+Y)} < \infty$, geometric ergodicity then follows by Popov's drift inequality.

With geometric ergodicity thus established, we may now use Tweedie's (1983) drift condition to guarantee the existence of the first moment.

Suppose now that $x > Y$ and $g(x) = e^{\alpha|x|}/\alpha$ for some $\alpha > 0$ which we will later choose conveniently. (A similar argument applies to the case where $x < -Y$.) For any such x ,

$$\sum_{y \in A^c} P(x, y) e^{\alpha|y|} / \alpha = \mathbb{E} (e^{\alpha|\omega_{i+1}|} / \alpha \mid \omega_i = x) \quad (35)$$

$$= e^{\alpha\omega_i} \mathbb{E} (e^{\alpha(\omega_{i+1} - \omega_i)} / \alpha \mid \omega_i = x). \quad (36)$$

So, if we define $A = [-Y, Y]$, then Tweedie's (1983) drift condition requires that for any $x > Y$:

$$e^{\alpha\omega_i} \mathbb{E} (e^{\alpha(\omega_{i+1} - \omega_i)} / \alpha \mid \omega_i = x) \leq e^{\alpha\omega_i} / \alpha - \omega_i \quad (37)$$

$$\Leftrightarrow \mathbb{E} (e^{\alpha(\omega_{i+1} - \omega_i)} \mid \omega_i = x) \leq 1 - \alpha\omega_i e^{-\alpha\omega_i}. \quad (38)$$

Taking a Taylor expansion of the LHS of the last line above, the condition becomes

$$\mathbb{E} (1 + \alpha(\omega_{i+1} - \omega_i) + (\alpha^2/2)(\omega_{i+1} - \omega_i)^2 + \cdots \mid \omega_i = x) \leq 1 - \alpha\omega_i e^{-\alpha\omega_i}. \quad (39)$$

But, for x sufficiently large, $E(\omega_{i+1} - \omega_i | \omega_1 = x) \leq -\pi\gamma$, by (2) and the fact that ω satisfies Self-Regulation. In this case, a sufficient condition for (39) to hold is

$$1 - \alpha\pi\gamma + (\alpha^2/2)Y^2 + \dots \leq 1 - \alpha\omega_i e^{-\alpha\omega_i}, \text{ or} \quad (40)$$

$$-\pi\gamma + \alpha Y^2/2 + \dots \leq -\omega_i e^{-\alpha\omega_i} / \alpha. \quad (41)$$

Fixing $\alpha < \pi\gamma/Y^2$, if ω_i is sufficiently large then $\omega_i e^{-\alpha\omega_i} / \alpha \leq \pi\gamma/4$, and hence

$$-\pi\gamma + \alpha Y^2/2 + \dots < -\pi\gamma + \pi\gamma/2 + \dots < -\pi\gamma/4 \leq -\omega_i e^{-\alpha\omega_i} / \alpha.$$

Therefore, we can choose a set of states A such that $\mu(A) > 0$, $\sum_{x \in A} x\mu(x) < \infty$ and the inequality above holds for $x \in A^c$. Since $\max_{x \in A} \sum_{y \in A^c} P(x, y)g(y) < \max_{x \in A} e^{\alpha(x+Y)} / \alpha < \infty$, the result then follows by Tweedie's (1983) Theorem 1. $\blacksquare$

A.3 Proof of Proposition 2

From Theorem 1 there exists an initial condition, μ , such that a monotone pure-strategy stationary Markov equilibrium exists, and $c(\emptyset) = \tau$.

Before observing her type in this equilibrium, an uninformed player i foresees that with probability τ she will play $\underline{a}$ and with probability $(1 - \tau)$ she will play $\bar{a}$. So her *ex ante* welfare is

$$\tau E[V_\sigma(\theta_i; \emptyset, \underline{a}) | \theta_i < \tau] + (1 - \tau) E[V_\sigma(\theta_i; \emptyset, \bar{a}) | \theta_i > \tau].$$

This is

$$E[V_\sigma(\theta_i; \emptyset, \underline{a})] + (1 - \tau) E[V_\sigma(\theta_i; \emptyset, \bar{a}) - V_\sigma(\theta_i; \emptyset, \underline{a}) | \theta_i > \tau].$$

By Condition 1B, this may be written as

$$E[V_\sigma(\theta_i; \emptyset, \underline{a})] + (1 - \tau) E[H(\theta_i)\{J(\theta_i) - G(\omega_i)\} | \theta_i > \tau].$$

Substituting for $\bar{U}$ and using the Law of Iterated Expectations, this expression becomes

$$\bar{U} + (1 - \tau) E[E[H(\theta_i)\{J(\theta_i) - G(\omega_i)\} | \theta_i] | \theta_i > \tau] \quad (42)$$

$$= \bar{U} + (1 - \tau) E[H(\theta_i)\{J(\theta_i) - E[G(\omega_i) | \theta_i]\} | \theta_i > \tau] \quad (43)$$

$$= \bar{U} + (1 - \tau) E[H(\theta_i)\{J(\theta_i) - J(\tau)\} | \theta_i > \tau] \quad (44)$$

$$= \bar{U} + E[H(\theta_i)\{J(\theta_i) - J(\tau)\}^+], \quad (45)$$

where we combine the properties that $J(\tau) = E(G(\omega_i))$ and that because player i is uninformed $E(G(\omega_i)) = E[G(\omega_i) | \theta_i]$. $\blacksquare$

A.4 Derivations of Surplus in the Self-Regulating Market Game

The expected surplus per player is

$$E[\theta_i|\bar{a}] \Pr[\bar{a}] - E[\theta_i|\underline{a}] \Pr[\underline{a}] - \bar{f}.$$

Recall that $\Pr[\bar{a}] = (1 - \hat{p})$ and $\Pr[\underline{a}] = \hat{p}$. So the expected surplus per player is

$$(1 - \hat{p}) - E[1 - \theta_i|\bar{a}] \Pr[\bar{a}] - E[\theta_i|\underline{a}] \Pr[\underline{a}] - \bar{f}.$$

This is

$$(1 - \hat{p}) - \frac{1}{2} E[1 - c(s_i)|\bar{a}] \Pr[\bar{a}] - \frac{1}{2} E[c(s_i)|\underline{a}] \Pr[\underline{a}] - \bar{f}.$$

Applying Bayes Law to the middle two terms, this is

$$(1 - \hat{p}) - \frac{1}{2} E[(1 - c(s_i))^2] - \frac{1}{2} E[c(s_i)^2] - \bar{f}.$$

And this is

$$(1 - \hat{p}) - \frac{1}{2} \text{Var}[c(s_i)] - \frac{1}{2} E[1 - c(s_i)]^2 - \frac{1}{2} \text{Var}[c(s_i)] - \frac{1}{2} E[c(s_i)]^2 - \bar{f}.$$

Or,

$$(1 - \hat{p}) - \text{Var}[c(s_i)] - \frac{1}{2}(1 - \hat{p})^2 - \frac{1}{2}\hat{p}^2 - \bar{f}.$$

References

- CHEN, H., AND M. FRANK (2001): “State dependent pricing with a queue,” *IIE Transactions*, 33, 847–860.
- COURNOT, A. (1838): *Recherches sur les Principes Mathématiques de la Théorie des Richesses*. English edition: *Researches into the Mathematical Principles of the Theory of Wealth*, ed. by N. Bacon. Macmillan, New York, 1897.
- DUFFIE, D., G. GEANAKOPOLOS, A. MAS-COLELL, AND A. MCLENNAN (1994): “Stationary Markov Equilibria,” *Econometrica*, 62, 745–781.
- EDELSON, N., AND D. HILDERBRAND (1975): “Congestion Tolls for Poisson Queuing Processes,” *Econometrica*, 43, 81–92.

- FOSTER, F. G. (1953): “On the Stochastic Matrices Associated with Certain Queuing Processes,” *Annals of Mathematical Statistics*, 24, 355–360.
- FUDENBERG, D., AND J. TIROLE (1991): *Game Theory*. MIT Press, Cambridge, Massachusetts.
- GALE, D. (1987): “Limit Theorems for Markets with Sequential Bargaining,” *Journal of Economic Theory*, 43, 20–54.
- (2000): *Strategic Foundations of General Equilibrium*. Cambridge University Press, Cambridge.
- HASSIN, R. (1986): “Consumer information in markets with random product quality: the case of queues and balking,” *Econometrica*, 54, 1185–1196.
- LARGE, J. (2009): “A Market-Clearing Role for Inefficiency on a Limit Order Book,” *Journal of Financial Economics*, 91, 102–117.
- MERTENS, J.-F. (2002): “Stochastic Games,” in *Handbook of Game Theory*, ed. by R. J. Aumann, and S. Hart, vol. 3, pp. 1809–1832. Elsevier, Amsterdam.
- MILGROM, P., AND C. SHANNON (1994): “Monotone Comparative Statics,” *Econometrica*, 62, 157–180.
- NAOR, P. (1969): “The regulation of queue size by levying tolls,” *Econometrica*, 37, 15–24.
- NOWAK, A. (1985): “Existence of Equilibrium Stationary Strategies in Discounted Non-cooperative Stochastic Games with Uncountable State Space,” *Journal of Optimization Theory and Applications*, 45, 591–603.
- PARTHASARATHY, T. (1982): “Existence of Equilibrium Stationary Strategies in Discounted Stochastic Games,” *Sankhya*, 44, 114–127.
- POPOV, N. N. (1977): “Conditions For Geometric Ergodicity of Countable Markov Chains,” *Soviet Math. Dokl.*, 18, 676–679.
- RIEDER, U. (1979): “Equilibrium Plans for Non-Zero-Sum Markov Games,” in *Game Theory and Related Topics*, ed. by O. Moeschlin, and D. Pallaschke, pp. 91–101. North-Holland, Amsterdam.

- RUBINSTEIN, A., AND A. WOLINSKY (1985): “Equilibrium in a Market with Sequential Bargaining,” *Econometrica*, 53, 1133–1150.
- SHAPLEY, L. S. (1953): “Stochastic Games,” *Proceedings of the National Academy of Sciences of the USA*, 39, 1095–1100.
- SHAPLEY, L. S., AND M. SHUBIK (1977): “Trade Using One Commodity as a Means of Payment,” *Journal of Political Economy*, 85, 937–968.
- SHUBIK, M. (1973): “Commodity Money, Oligopoly, Credit and Bankruptcy in a General Equilibrium Model,” *Western Economic Journal*, 11, 24–38.
- TWEEDIE, R. L. (1976): “Criteria for Classifying General Markov Chains,” *Advances in Applied Probability*, 8, 737–771.
- (1983): “The Existence of Moments for Stationary Markov Chains,” *Journal of Applied Probability*, 20, 191–196.