Incentives, Reputation and Learning

Elwyn Davies

Balliol College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy in Economics

December 2016

Incentives, Reputation and Learning

Elwyn Davies

Balliol College
University of Oxford

Abstract Many economic interactions are characterized by imperfect information and imperfect enforcement. This thesis analyzes the use and the effectiveness of incentives, reputation and learning and how economic agents use them to mitigate the problems caused by both.

Chapter 1 presents experimental evidence on the use of price incentives and reputational incentives from gift-exchange game lab experiments in Ghana and the United Kingdom. While in the United Kingdom, like in other experiments, participants assigned the role of employers reward high effort by increasing wages and punish low effort by decreasing wages, it appears that Ghanaian participants refrain from doing so. This results in lower average effort in Ghana compared to the United Kingdom. Introducing competition for workers or a reputation mechanism does not improve effort.

Chapter 2 further examines low worker productivity in Ghana by empirically testing two other incentive mechanisms: the employer giving feedback and the worker making a pledge in advance. While these incentives on average do not affect effort, they do increase conditional reciprocity of the worker's behalf. Workers receiving a high wage are more likely to exert high effort when feedback and promises are possible.

Chapter 3 presents a partial-equilibrium model showing how business networks can lead to higher quality firms overall. The presence of such networks leads to a *filtering effect*, which means that fewer low-quality firms produce, but also to a *lock-out effect*, implying that networks redistribute demand to their own members and make production outside the network less viable.

Chapter 4 presents a model and experimental evidence on learning and strategic experimentation in the context of a repeated ultimatum game. The model predicts that Proposers make low offers in early periods, hoping that learning will benefit them in later periods. However, when Responders also have strategic considerations, the Responders will strategically reject low offers, which will optimally lead to higher offers. A series of online experiments finds evidence that Responders reject low offers, but cannot disentangle this from distributional preferences and reciprocity concerns.

Supervisors This thesis was supervised by Alan Beggs (University of Oxford) and Marcel Fafchamps (Stanford University).

Total words This thesis contains approximately 69141 words.

Statement of authorship Chapters 1 and 2 are joint work with Marcel Fafchamps. Marcel and I designed and planned the experiments together. Marcel acquired funding for the UK experiments conducted as part of Chapter 1. I programmed the experiments and conducted them in Ghana and the United Kingdom. I have implemented the analysis and have written up the results, as included in this thesis, with Marcel providing comments and guidance. Chapters 3 and 4 are fully my own work.

Acknowledgement of funding I acknowledge funding from the Economic and Social Research Council (ESRC), grant number ES/J500112/1, and the University of Oxford Social Sciences Doctoral Training Centre (DTC). The experiments in Ghana, part of Chapter 1 and 2, have been funded by the UK Department for International Development under the Improving Institutions for Pro-Poor Growth (iiG) Phase 2 programme. Furthermore, I acknowledge financial support for conference funding and research expenses from the George Webb Medley Fund, Balliol College, Wadham College and the Royal Economic Society. I am very grateful for the support I have received.

Ethical approval For Chapters 1, 2 and 4 ethical approval has been sought and obtained from the University of Oxford CUREC committee.

Corrections This version incorporates the minor corrections required after the viva. The examiners of this thesis are Johannes Abeler (University of Oxford) and Rocco Macchiavello (LSE). These corrections were completed on 24 March 2017.

To Peter and Leonie

Acknowledgements

There are many people I would like to thank for supporting me and encouraging me during the writing of this thesis. I am hugely indebted to my supervisors, Alan Beggs and Marcel Fafchamps. Both have been extremely helpful in the preparation of this thesis, by providing advice and guidance. I would like to thank Alan for our many meetings and discussions how the experimental results fit in the theory. I would like to thank Marcel for introducing me to development economics and behavioural economics as well as for being able to always present a new perspective on my findings. It has been a pleasure to work with them.

There are also many others in the Economics department who I would like to thank. One person who deserves to be mentioned separately is Simon Quinn. I would like to thank him for his advice and comments and especially for creating a research group on firms and development, which has become very much my academic "home" here in Oxford. This research group has been invaluable for me as it provided a forum to present and discuss new ideas. Second, I would like to thank the researchers and administrative staff at the Centre for the Study of African Economies for supporting me throughout my research. Without the administrative help from Julia Coffey, Suzanne George, Richard Payne, Rose Page and Gail Wilkins my fieldwork would have been so much harder. I would like to thank Stefano Caria, Kate Orkin and Karlijn Morsink, with whom I attended so many conferences and who have provided me with valuable insights and feedback.

I would also very much like to thank my fellow DPhil students here in Oxford. Not only the members of the research group in firms and development (Jono, Muhammad, Rachel, Daniela, Anisha, Lukas, Emma, Michael, Marc, Pavel, Chris), but also students working in other areas of development economics and economics in general (Viviana, Dan, Thomas, Andreas, Matej, Damian, Irem, Sébastien, Germán, Igor). I have made many new friends here and they have made my days in the open space area so much nicer and brighter.

This work would also not have been possible without the fieldwork team in Ghana. I would like to thank Andrew Kerr, Denise Gray, Martin Černý, Václav Těhle and Moses-Awoonor Williams for helping out with the organization of the fieldwork. Furthermore, I would like to thank my team of enumerators for their assistance with running the experiments: Marzu Abdulai, Ibrahim Lawal Musah, Anna Consomer, Dan Kumi Ntow, Eric Agyekum and Bismark Owusu Tetteh Nortey. I would like to thank CESS in Oxford for their assistance with running the Oxford experiments.

Outside the department there are many others I would like to acknowledge and thank. First of all, my parents, Peter Davies and Leonie van der Valk (who beat me in finishing her PhD thesis by a couple of months). They have always been supporting me in whatever I have been doing and wherever I have been going. Their love and dedication have brought me here. Being home is always a pleasure. I dedicate this thesis to them. And of course, I would like to thank my brother, Dylan, and his girlfriend Laura, for their many spontaneous visits to Oxford. No idea is too crazy for them. Dylan's always positive mood should be an example for many. Furthermore I would like to thank the rest of the family for their support and visiting me in Oxford (Liesbeth, Karel, David, Aad, Irene & Jasper). This thesis is also in memory of my grandmother, Corrie Davies-Scheij, and my uncle, Dik van der Valk, who passed away during the writing of this thesis.

I made many new friends, both within and outside Oxford. Balliol College has provided me with a warm and welcoming academic environment. I would like to thank the members of the Balliol "German (and Dutch) breakfast club", Carla, Emily, Anatolij, Antonia, Nina, David, for sharing so many breakfasts (and dinners) together as well as many other amazing people I met there (Arjun, Ibrahim, Cramer, Vera and many others). Also, I met so many interesting people during conferences and summer schools, with whom I became friends. The development "selfie group" (Lisette, Zlata, Łukasz, Heidi, Martin and Sanne) deserves a special mention. Another person deserving a special mention is Christopher Jones, with whom I made so many spontaneous last minute trips to destinations far away. I very much value our friendship and hope that there are still many other spontaneous trips to come!

I also would like to thank my friends from back home. Michiel, who I have known since I was three days old and with whom I remained close over all these years (and of course, his fiancée Tessa and his parents, Hans and Mieke, who have been like second parents to me). And of course, the friends I made during my undergraduate degree and with whom I have remained in touch during all these years here in Oxford (among others, Lucas, Jacintha, Adrienne and Niels).

Last, *but not least*, I would like to thank Hannah, to whom I owe so much for her love and understanding, and not to forget, her support throughout the writing of this thesis. I hope that one day I will be able to reciprocate this! I am very glad to have met her and am looking forward to being able to spend more time together!

To all of you, I would like to say:

Bedankt! *Thanks!* *Danke!*

Elwyn Davies
Oxford, December 2016

Contents

Introduction

This thesis investigates how firms cope with imperfect information and imperfect enforcement and in particularly how they use incentives and reputation, how they learn and how this behaviour is related to firm productivity and development. The main motivation for this thesis comes from two big economic questions. The first question is why differences in labour productivity across countries are so large, even after correcting for differences in human and physical capital (Hall & Jones, 1999; Caselli & Coleman, 2006). One often mentioned potential driver are differences in how workers are managed and motivated (Bloom & Reenen, 2010). Incentives play an important role, but not all incentives work equally well in every environment (Prendergast, 1999). The second question is how firms and consumers deal with what Greif (2000) has called the "fundamental problem of exchange": how do trading partners recognize mutually beneficial trading relationships and how do they make sure that the other party keeps to their promises?

In this thesis I present three experimental studies on the use of incentives and learning in repeated interactions and one theoretical perspective on reputation and business networks.

Contributions My thesis makes several contributions. The first contribution relates to the use and effectiveness of incentives across different contexts. In this thesis I compare the use of incentives in both a developed country setting, the United Kingdom, and a developing country setting, Ghana. I show that the use of incentives, and in particular wage incentives, in experiments is not universal. Even though the use and effectiveness of incentives have been studied widely, both in the lab (e.g., Fehr et al., 2007; Brown et al., 2004; Fehr et al., 1997), the field (e.g., Falk, 2007; Gneezy & List, 2006; Gneezy & Rustichini, 2000; Kube et al., 2006; DellaVigna et al., 2016; Bandiera et al., 2005) and the lab-in-the-field (e.g., Fehr & List, 2004), many of these studies have been done in developed countries, often with student populations. It can be questioned whether these results translate to other contexts (Henrich et al., 2010; Cardenas & Carpenter, 2008; Falk et al., 2015). While results from experiments in high income countries (e.g., Brown et al., 2004) and from my own experiments in the United Kingdom show that monetary incentives, such as wage incentives, encourage high effort by workers, I find in Ghana that such incentives are not as commonly adopted. Low effort is punished little and high effort is not much rewarded. This lack of adoption of wage incentives in Ghana is robust across different experimental specifications as well as when using a second sample of actual entrepreneurs.

The second contribution of my thesis relates to the use of monetary and nonmonetary incentives in the workplace. I show that in the Ghanaian context, where my earlier experiments showed a lack of adoption of monetary incentives, nonmonetary incentives can increase effort, but *only* when wages are high. Many earlier studies have shown that promises (Charness & Dufwenberg, 2006; Vanberg, 2008), revealing identity (Gächter & Fehr, 1999) or feedback (Masclet et al., 2003; Cooper & Lightle, 2013) can increase cooperative behaviour, but that these incentives are less effective than monetary incentives or repeated interactions (Falk et al., 1999). The experiments presented in the thesis confirm that nonmonetary incentive can increase effort in Ghana, but only when wages are high. Positive feedback is related to higher improvements in effort than negative feedback. This suggests that the channel through which nonmonetary incentives like communication work is that they increase reciprocity. However, it also suggests that monetary and nonmonetary incentives are the most effective when used in tandem.

The third contribution of my thesis relates to reputation. I show that in the Ghanaian context, multilateral reputational incentives are less strong than bilteral reputational incentives. Furthermore, I show theoretically that mechanisms to establish reputation, such as business networks, can lead to an overall increase in quality of products and services, but also lead to an inefficient redistribution of demand. When enforcement is imperfect and relying on bilateral exchange cannot get you all the products or services you need as a firm or consumer, reputations can be crucial to help in choosing which seller to buy from (Greif et al., 1994; MacLeod, 2007). Experiments have shown that sharing information on past performance of employees with other employers (multilateral reputation) can increase effort (Falk et al., 2005). My thesis contributes to this literature by studying the relative strength of multilateral and bilateral reputation in a series of experiments. I show that employers are twice as responsive to low effort in their wage offers when this low effort was exerted when the worker was working for them instead of another employer. Besides this, my thesis studies theoretically how firms can use business networks as ways to establish reputations. I show that firms do not only join networks for informational advantages, but also for protectionist reasons: networks allow firms to redistribute demand to them. In some case the cost that this imposes outweighs the informational benefit that the network provides.

The fourth contribution relates to strategic learning behaviour. I show that, compared to a rational utility maximizing model, people underestimate the opportunity costs of experimentation in an ultimatum game. Furthermore, I show that the evidence on strategic rejection is mixed: there is an identifiable group of participants who rejects strategically, but this does not affect the average offers or reservation prices in an ultimatum game. This contributes to a larger literature on experimentation. When experimenting, economic agents often face a trade-off between short-run losses as a result from making a suboptimal choice and the long-term benefits of learning what the optimal offer is (see e.g. Bolton & Harris, 1999; Keller & Rady, 1999). This means that agents sometimes stop experimenting, even though there is still scope for learning (Aghion et al., 1991). In this thesis, I drop the assumption that one party is myopic and show how this diminishes opportunities for learning. Besides this, I conduct a series of online experiments to test the model of learning.

My final contribution is methodological. As part of the work for the thesis I developed a platform for running realtime interactive experiments on tablets or online, *LabBox*. This platform allows users to set up mobile labs in places where there were none. This platform has the potential to allow us to conduct interactive experiments with subject pools who are otherwise difficult to bring into a lab. The lightweight infrastructure, which can operate independently from mains electricity or an internet connection, means that these experiments can also be conducted in areas where infrastructure constraints make computer use more difficult. Compared to similar platforms like Z-Tree (Fischbacher, 2007) and O-Tree (Chen et al., 2016), LabBox is less demanding in its requirements of specific software on both the client-side and the server-side.[1]

Overview This thesis consists of four chapters, each presented as a separate research paper. Chapter 1 studies the adoption of wage incentives and the use of relational contracting in principal-agent experiments in Ghana and the United Kingdom. The experiment is an adapted version of the gift-exchange game (Fehr et al., 1993), repeated over five periods. In this game, each period, employers make offers to workers, who can choose to accept or reject this offer. Next, the workers who have accepted an offer can choose a level of effort. While in the UK high effort is rewarded and low effort punished, like in earlier experiments in high-income countries, I find limited evidence for this in Ghana. Overall effort provision is low, but this seems to be driven by employers not setting the workers incentives rather than the workers being more opportunistic or less trustworthy in Ghana.

[1]More details on the platform, including demonstrations, can be found online, at `http://www.elwyndavies.com/tools/labbox/`.

Additional experiments, including with actual entrepreneurs instead of students, show similar patterns of behaviour. Introducing competition among workers and employers does not change the result. Employers in the United Kingdom seem to be responsive to reputation, but overall this does not increase effort.

Chapter 2 builds further on this result. The lack of usage of wage incentives by our Ghanaian subject pool motivated us to study alternative incentives. In informal relationships nonmonetary incentives like praising and shaming can be important (see e.g. Fafchamps, 2004). In this chapter we study the role of feedback and promises in an additional series of principal-agent experiments in Ghana, to see if these incentives are more effective. In some treatments of the experiments employers have the opportunity to send a feedback message after the worker has chosen effort. In other treatments, workers can make a promise about what level of effort they will exert in the next period. The results show that feedback increases effort, but only when wages are high. I do not find a strong effect for promises.

Chapter 3 moves away from the topic of incentives and instead focuses on reputation and networks. The chapter presents a model of how firms can use business networks, such as business or trading associations, to gain information on the quality of suppliers within the network. The chapter also studies the impact of the network's existence on the quality that firms provide in the entire economy, including those outside the network. I show that the presence of such a network leads to two effects: a *filtering effect*, which means that fewer low quality firms produce, but also to a *lock-out effect*, where networks redistribute demand from firms outside the network towards their own members and make production outside the network less viable. The lock-out effect occurs even when the informational benefits from joining the network are low. In this case the cost of networking outweighs the informational benefits.

Chapter 4 presents a study on experimentation, learning and strategic considerations in the context of an ultimatum game. In this chapter I describe how a proposer in the ultimatum game will experiment strategically with their offers and how a responder can use this experimentation of the proposer to their advantage, by feigning higher than true costs and rejecting low but positive offers. I test the predictions of the model experimentally and show that there is more downward experimentation than the model predicts, suggesting that people underestimate the opportunity costs of experimenting. The results on strategic rejection are more mixed: overall, I do not find significant evidence that strategic rejection consideration affect offers or reservation prices, but a heterogeneity analysis suggests that there is a subset of responders who reject offers due to strategic reasons.

The motivation for this chapter followed from a conjecture on the results from Chapter 1, namely that the limited use of incentives was the result of learning. If an employer is unsure about what wage the worker needs to exert high effort, the employer could take low effort as a sign that the wage offer was to low and therefore decide to increase the wage instead of decreasing it. The observed differences between the UK and Ghana might therefore be the result from differences in bargaining strategies.

In the conclusion I discuss the implications for policy of this thesis as well as directions for future research.

Chapter 1

When No Bad Deed Goes Punished: Relational Contracting in Ghana versus the United Kingdom

with **Marcel Fafchamps**, Stanford University

Abstract In relational contracting the threat of punishment in future periods provides an incentive to not to cheat. However, to what extent do people actually carry out this punishment? We compare relational contracting patterns in Ghana and the United Kingdom by conducting a repeated principal agent lab experiment, framed in a labour market setting. Each period, employers make offers to workers, who can choose to accept or reject this offer and, after accepting and being paid, what effort to exert. The employers and workers interact repeatedly over several periods. While in the UK, in line with theoretical predictions and previous experiments (e.g. Brown et al., 2004, 2012), high effort is rewarded and low effort punished, we do not find evidence for the use of such incentives in Ghana. As a result, employers fail to discipline a subgroup of "selfish" workers, resulting in a low average effort and low and often negative employers' earnings. Structural estimates of Fehr-Schmidt preferences find that in both countries the choices by workers are not explained by a model of purely self-interested behaviour. The estimated parameter for advantageous inequality aversion ("fairness") is slightly higher in the UK than in Ghana. However, set identification shows that the share of "selfish" workers in Ghana is not substantially different from the UK or from earlier experiments. Introducing competition for workers or a reputation mechanism does not significantly improve workers' effort.

1.1 Introduction

Worker performance is a key determinant of firm productivity. In many sub-Saharan African countries worker productivity is low, even after accounting for differences in physical and human capital (Hall & Jones, 1999; Caselli & Coleman, 2006). Firm surveys and other studies have shown that employers in developing countries have substantial difficulties in managing their workers and making sure that their employees perform (e.g. Bloom & Reenen, 2010; Fafchamps & Söderbom, 2006). At the same time, competition for jobs is high: employers often face many applicants per opening (e.g. Falco & Teal, 2012). This seems at odds with predictions from labour market models (Shapiro & Stiglitz, 1984; Bowles, 1985), as well as empirical studies in the United States and the United Kingdom (Green & Weisskopf, 1990; Wadhwani & Wall, 1991), in which competition for jobs and threat of dismissal drive higher worker performance.

Relationships and reputation are important when the enforcement of formal contracts is problematic: legal institutions in developing countries are often weak and wage contracts are often incomplete, due to the complexity of many labour tasks. This means that employers have to rely on other incentives. Interacting repeatedly with one party can provide a viable alternative: if both parties gain enough value out of a relationship, the threat of losing the relationship or having its value reduced can provide an incentive to perform well. This is the essence of relational contracting. Many studies have been conducted studying its importance, especially in developing country context (e.g, McMillan & Woodruff, 1999; Macchiavello & Morjaria, 2015) Similarly, the ability to develop a public reputation can work as an enforcement mechanism: if individuals gain enough value out of having a good reputation, the threat of losing this reputation can deter cheating (Milgrom et al., 1990; Greif, 2000).

This chapter looks at relational contracting in Ghana and compares this with the United Kingdom. We conduct a labour market lab experiment to assess the use of relational contracting and reputation in labour relations. In this experiment, based on the gift-exchange games of Fehr et al. (1993) and Brown et al. (2004), participants interact as workers and employers in a principal-agent setting: employers start out by making an offer to a worker, specifying a wage in return for effort. The worker then chooses to accept or reject this offer and what effort to exert. Crucially, the employer cannot condition the wage on effort. However, the workers and employers interact repeatedly, and the employers can adjust their offers in the next periods. In five different treatments, we vary the size of the market, the contractual completeness and the information available to employers.

We find a substantial group of workers exerting low effort, despite high average wages, in both Ghana and the United Kingdom. This does not differ from previous experiments. What differs in Ghana from earlier experiments and our experiments in the UK is a *lack of punishment* by employers: employers keep offering high wages to workers choosing low effort, resulting in low and often negative payoffs for employers. This is at odds with relational contracting models where low effort is punished by termination or by a lowering of the wage (see e.g. Shapiro & Stiglitz, 1984; Kranton, 1996; Ghosh & Ray, 1996). This is in stark contract with the experiments of Brown et al. (2004, 2012), conducted in Switzerland, and experiments in other countries, where low effort was strongly punished, even when there was a shortage of labour. Furthermore, we find that introducing reputation or increasing the number of market participants does not change this.

Our paper contributes to an emerging literature on how norms, preferences and strategic behaviour differ across cultures (Henrich et al., 2006; Cardenas & Carpenter, 2008; Henrich et al., 2010). The most extensive work in this area, the Global Preference Survey conducted by Falk et al. (2015), showed large differences in preference measures across representative samples in 76 countries. Not only did they find large differences between countries, but especially also within countries. While they elicited preferences through simple survey questions, we compare behaviour in a more extensive lab game, where players interact repeatedly and can respond to each other's behaviour.

Our experiment intends to capture labour market heuristics influenced by prevailing norms and expectations. Compared to many high income countries, the labour market in Ghana is characterized by a large degree of informality. Informal labour, such as contributing to a family firm, as well as self-employment are more common than wage employment: in 2012-13 only 20.2 percent of the country's working population and 32.5 percent of the urban working population was wage employed (Ghana Statistical Service, 2014). This lower importance of wage employment is likely to influence the expectations of both workers and employers in wage relationships. Understanding these labour market norms and expectations in a developing country is crucial, especially now given that the share of wage employment in Ghana has been rising over the last decade. Our results show that certain actions predicted by theory, such as lowering wages of low effort workers, are not commonly used by our Ghanaian subjects. This is in line with earlier empirical findings that have shown that labour management is inadequate and often problematic in developing countries (Bloom et al., 2014; Fafchamps & Söderbom, 2006).

The paper is structured as follows: Section 2 discusses related experiments. Section 3 presents the experiment and predictions based on the theory. Section 4 presents the results and tests the predictions. Section 5 concludes and discusses possible explanations for the behaviour found in this experiment.

1.2 Related literature

The gift-exchange game has been used widely to document the role of labour market institution and contract relations (Charness & Kuhn, 2011; Fehr et al., 2009). The first gift-exchange game lab experiment was conducted by Fehr et al. (1993). Their experiment consisted of one-shot interactions of workers and employers who were randomly rematched after each period. Despite the lack of explicit incentives or future interactions, they showed that employers offered wages above the market clearing level and workers exerted a higher level of effort in return to this higher wage. This provides evidence for the fair wage-effort hypothesis for involuntary employment as formulated by Akerlof (1982): employers and workers engage in a "gift exchange". This pattern can sustain even when there are more workers than employers and workers bid for wage contracts: Fehr & Falk (1999) find that some employers refuse to hire employees who undercut wages, fearing that they are more likely to shirk.

In a repeated interaction, these concerns can be reinforced. Gächter & Falk (2002) introduce a treatment of the gift-exchange game experiment in which an employer repeatedly interacts with the same worker for ten periods. They find that a repeated interaction makes the positive wage-effort relation stronger, compared to a treatment in which workers and employers are randomly rematched every period. Furthermore, they find that the repeated interaction can function as a disciplining device: "selfish" types imitate "reciprocal" types in the first periods of the game. This is in line with the predictions of Kreps et al. (1982) and Fehr & Schmidt (1999) on cooperation in finite games.

In Brown et al. (2004) this is taken further, by introducing multiple employers and workers who can contract with each other in the same marketplace. They allow the employers to both make public offers, made to all workers and visible to all participants, and private offers, only made to one specific worker. They show that relational contracting can emerge naturally in this environment: employers keep offering high wage contracts to workers that exerted high effort in the past. This leads to a higher effort than in a treatment in which the identifiers of workers and employers are scrambled every period so that employers cannot recognize their past workers. Follow-up experiments have replicated these results (see e.g. Brown et al., 2012; Altmann et al., 2014; Wu & Roe, 2007a). In these experiments, employers, when faced with a low level of effort, choose to either terminate the relationship (contingent contract renewal) or lower the wage in the next period.[1] The threat of economic loss if the relationship is ended or its value reduced makes these bilateral relationships feasible.

While these studies only revealed past performance to the past employers, and therefore only allowed for bilateral reputations to emerge, Falk et al. (2005) introduce a treatment in which past performance of workers is made publicly available to all employers. They find that effort in the public past history treatment is higher than in the treatment in which employers cannot keep track of past performance of workers. However, the effect of reputation is limited: bilateral relationships still play an important role.

In these experiments there is only contractual incompleteness on the side of the worker: the worker is always ensured of receiving the promised wage, while the employer is uncertain of what effort to receive. Several other experiments introduced the possibility of ex post wage adjustments, such as giving a bonus to a worker or withholding some of the promised wage (a malus) (see e.g. Fehr et al., 1997, 2007; Wu & Roe, 2007b,a; Falk et al., 2008). This is generally effective: Falk et al. (2008) find that bonus systems can work of a substitute for long-term contracts and Wu & Roe (2007b) note that the trading pattern is very close to complete contracts.

The gift-exchange game has been conducted in many OECD countries: Austria (Fehr et al., 1993, 1997), Germany (Abeler et al., 2010; Altmann et al., 2014; Falk et al., 2008; Fehr et al., 2007), Switzerland (Fehr & Falk, 1999; Brown et al., 2004; Falk et al., 2005), the Netherlands (Van Der Heijden et al., 2001), Portugal (Pereira et al., 2006), Spain (Brandts & Charness, 2004), the United Kingdom (Gächter et al., 2016) and California, Ohio and Florida in the United States (Charness, 2004; Cooper & Lightle, 2013; Wu & Roe, 2007a,b). All these experiments have found evidence for a general pattern of high effort in return for high wages. Similar results have been found in experiments in former communist countries, such as Hungary (Falk et al., 1999) and Russia (Fehr et al., 2014). However, to the best of our knowledge very few gift-exchange experiments have been conducted in a developing country and none in a sub-Saharan African country.[2]

Another group of experiments to which our experiment is related are trust games (Glaeser et al., 2000; Karlan, 2005; Bohnet et al., 2010). The gift-exchange game is very

[1]Brown et al. (2012) show that on their 10-point scale of effort, a one point increase in effort leads to an increase of the wage of 5.527 (Table 4, column 6). This is significant at a 1 percent level. Their wage is bounded between 1 and 100.

[2]Siang et al. (2011) conducted a bilateral gift exchange experiment in Malaysia.

closely related to the trust game (see Berg et al., 1994; Camerer & Weigelt, 1988), with the crucial difference is that in the gift-exchange game the size of the surplus is determined by the choice of effort by the worker (the receiver), while in the trust game the size of the surplus is determined by the initial offer by the employer (the proposer). Just like in the trust game, trust plays an important role in the gift-exchange game: a high wage encourages high effort, but an employer offering a high wage is making himself or herself vulnerable to a worker choosing low effort, resulting in a potentially high loss for the employer.

1.3 The experiment

1.3.1 Experimental setup

The experiment is a multi-period gift exchange game, based on Brown et al. (2004) and the original gift-exchange game of Fehr et al. (1993). This experiment mimics labour relations in a principal agent setting. At the beginning of the experiment, participants are randomly assigned the role of a worker or an employer. Markets consists of either one worker and one employer or three workers and three employers, depending on the treatment. This game is played for five periods, after which the employers and workers are rematched. Each participant plays four games of five periods.[3] The timing of the game in each period is as follows:

- **Contracting.** In this stage employers and workers contract with each other, in a virtual marketplace. Each of the workers is listed with their identification number clearly visible. This stage consists of three steps:
  - First, the employer can make offers to every worker individually. An offer specifies the payment that the employer will make to the worker w and the desired level of effort e from the worker.
  - Second, the employer observes all the offers that are made by other employers and has then a chance to revise the offers to the workers. The initial offers are not shown to the workers.
  - Third, in a randomly determined sequence, each of the workers gets to see their offers and then can choose to accept one of the offers. In case the worker rejects all offers or no offers are available, the worker does not receive any payment. If an offer of an employer is accepted by a worker, then no other worker can

[3] The worker-employer pairs or groups stay the same during the five periods. Participants can recognize workers and employers using a random identification letter, which is randomized at the beginning of each game.

Table 1.1: The payoff structure in the experiment.

		Effort level (e)		
		Low	Medium	High
$B(e)$	Employer's benefit	5	20	40
$c(e)$	Worker's cost	0	2	6
$S(e)$	Surplus ($B(e) - c(e)$)	5	18	34

accept an offer from this employer. This ensures that each worker has at most one employer and that each employer has at most one worker.

- **Supply.** In this stage, the worker first receives the payment w from the employer and then chooses the amount of effort e to exert. This effort can either be low, medium or high ($e \in \{L, M, H\}$) and costs the worker $c(e)$. This level of effort gives the employer a benefit of $B(e)$. The payoffs to the employer π_E and to the worker π_W are:

$$\pi_E = B(e) - w \tag{1.1}$$

$$\pi_W = w - c(e) \tag{1.2}$$

  The choice of effort is not revealed to the employer until the end of the next stage.

- **Rehiring.** Before moving to the next period, we ask the employer to make a contingent choice for contract renewal in the next period. At this stage the employer does not know the chosen level of effort yet. For each possible effort choice we ask the employer whether he or she wants to rehire the worker in the next period and for what kind of contract (specifying a wage and effort level, as before). Similarly, we ask the worker what the *minimum* payment needs to be in order to accept the offer. If the offer of the employer is equal or higher to this minimum payment, the employer and worker will contract with each other in the next period, and proceed directly to the supply stage of the next period.

The values of $c(e)$ and $\Pi(e)$ for the three different levels of effort are described in Table 1.1.[4] High effort costs the worker 6 points and gives the employer a benefit of 40 points. This is a desirable outcome for the employer, but it requires trust: the wage needs to be high enough for the worker to cover the 6 point costs, but a high wage makes the employer vulnerable to the risk that the worker chooses low effort. Medium effort costs the worker

[4]The payoff structure is designed in such a way that the marginal cost of effort is strictly increasing. The effort levels are equal to three out of the ten effort levels in Brown et al. (2004). The number has been limited to three to allow the strategy elicitation in the rehiring stage and to simplify the game.

2 points and give the employer a benefit of 20 points. Finally, low effort is costless to the worker, but gives the employer a benefit of only 5 points.

The rehiring stage is novel to this experiment. The strategy method applied allows us to analyze the conditional responses to effort on behalf of the employer and the expectations on behalf of the worker. When no agreement is reached, both parties are given a new chance to contract with each other. The rehiring stage also allows for the development of "long-term" relationships in the treatments with three workers and three employers: if an agreement is reached both parties will continue to trade with each other. Due to the random order in which workers choose offers, there is no such guarantee in the normal contracting sequence.

Our experiment differs from Brown et al. (2004) in several ways, besides introducing a rehiring stage. First, contracting takes place in stages and not continuously. We were worried that differences in the participants' technical ability would impact the result. For example, high wage offers might be taken by those who can respond the most quickly, which might be correlated with relevant personal characteristics (e.g. age, competitiveness and so on). Second, we reduced the number of effort levels to three, which simplifies the game and allows for a simpler application of the strategy method. Third, the number of workers and employers in our experimental markets is smaller and there is no market imbalance: there is no excess labour demand or supply. In Brown et al. (2004) there were 7 employers and 10 workers, which makes unemployment more costly for workers. However, a follow-up study showed that introducing an excess of labour supply did not matter for the prevalence and pattern of relational contracting, but only for the division of surplus (Brown et al., 2012). Fourth, we reduced the number of periods of each interaction, and introduced a sequential treatment design, where each participant plays four games of five periods in a row. In a series of follow-up experiments we increased the number of periods to ten and did not find significant differences. Finally, our experiment is explicitly phrased in labour market terms, instead of describing work as a "good", an employer as a "buyer" and a worker as a "seller". Our pilot tests showed that using more neutral terms decreased the understanding of the game.[5] Fehr et al. (2007) found in an earlier gift-exchange game experiment with students in Munich that labour market framing did not produce different results from a neutral framing.

[5]Especially the role of the buyer and seller confused participants: as it is the buyer (i.e., the employer) who specifies the price of the product, while normally in transactions the price is often determined by the seller (i.e., the worker).

Table 1.2: Overview of the five treatments.

Treatment	Number of workers and employers	Effort choice	Information on past worker's actions
(1C)	1 / 1	no	only own information
(1E)	1 / 1	yes	only own information
(3C)	3 / 3	no	only own information
(3E)	3 / 3	yes	only own information
(3ES)	3 / 3	yes	shared between employers

Table 1.3: The number of participants in each treatment in each part of the sequence in Ghana.

	Ghana				**United Kingdom**			
Treatment	Game 1	Game 2	Game 3	Game 4	Game 1	Game 2	Game 3	Game 4
(1C)	304	28	28	28	192	–	–	–
(1E)	–	138	48	48	–	192	60	60
(3C)	–	138	30	30	–	–	–	–
(3E)	–	–	198	90	–	–	132	54
(3ES)	–	–	–	108	–	–	–	78
Total	304	304	304	304	192	192	192	192

As we will discuss later in the paper, despite our changes to the design, our results for the United Kingdom are very similar to earlier experiments conducted in OECD countries. Arguably, some of the changes in our paper make conditional reciprocity more likely: the introduction of the rehiring stage requires participants to explicitly think about their conditional response.

1.3.2 Treatments

We introduce five treatments that vary in the number of market participants, whether compliance is enforced and the available information on past actions. Table 1.2 presents the five treatments. In the control treatments (1C) and (3C) workers can only choose the level of effort that is demanded by the employer. The difference between these two treatments is that in (1C) markets consist of one worker and one employer, while markets in (3C) consist of three workers and three employers. These treatments are essentially an adapted version of the ultimatum game, because the worker can only choose to accept or reject the proposed split.

In treatment (1E) markets consist of one worker and one employer and the worker has full choice of effort. This treatment is similar to bilateral gift-exchange games with fixed partners (see e.g., Kirchler et al., 1996; Gächter & Falk, 2002). In treatments (3E) and (3ES) markets consist of three workers and three employers and workers have full choice of

Table 1.4: The seven treatment sequences of the experiment.

	Game 1 (5 periods)	Game 2 (5 periods)	Game 3 (5 periods)	Game 4 (5 periods)	**Ghana** No. of subjects	**UK** No. of subjects
I	(1C)	(1C)	(1C)	(1C)	28	–
II	(1C)	(1E)	(1E)	(1E)	48	60
III	(1C)	(1E)	(3E)	(3E)	36	54
IV	(1C)	(1E)	(3E)	(3ES)	54	78
V	(1C)	(3C)	(3C)	(3C)	30	–
VI	(1C)	(3C)	(3E)	(3E)	54	–
VII	(1C)	(3C)	(3E)	(3ES)	54	–
Total					304	192

Note: After each game of five periods employers and workers are rematched randomly.

effort. These treatments are the closest to the multilateral gift-exchange games conducted by Brown et al. (2004). The difference between (3E) and (3ES) is that in treatment (3E) employers only have information on whether a worker complied with their effort demand while working for them, while in treatment (3ES) this information is shared among all three employers. Treatment (3ES) allows for a multilateral reputation mechanism, while the other effort choice treatments only allow for bilateral reputations.

Each participant plays four games of five periods. Table 1.4 shows the seven treatments sequences and the number of participants. This setup is designed to allow for comparisons within and between subjects as well as a gradual introduction of the more complicated treatments. In the United Kingdom, we only conducted three out of the seven treatment sequences.[6]

These treatments allow us to compare the impact of imperfect enforcement, the role of competition (increasing the number of employers and workers) and the role of sharing information between employers. In treatment (1C) and (3C) the worker has to comply with the demanded effort and by comparing these treatments with treatments (1E) and (3E), where a choice of effort is possible, we can analyze the impact of imperfect enforcement, following Brown et al. (2004). In treatments (3C), (3E) and (3ES), the employment markets consist of three workers and three employers, creating a situation of (balanced) competition. By comparing these treatments with treatment (1C) and (1E), we can analyze the impact of having a larger market on wage offers and effort. Finally, comparing treatments (3E) and (3ES) allows us to see whether having more information on past enforcement

[6]Our aim of the United Kingdom experiments was to compare behaviour in the effort choice treatments. We therefore only replicated the three treatments which featured predominantly effort choice treatments. See Table 1.4 for an overview which treatments were selected.

leads to employers offering higher wages to high effort workers that they did not employ themselves.

1.3.3 Implementation

The participants of this study were students from colleges and universities in Accra, Ghana, and Oxford, United Kingdom. In Ghana a total 16 sessions were held, with 18 to 20 participants each and a total of 304 participants. In the UK we held 13 sessions, with 192 participants in total. The points that participants earned during the session were converted to Ghana cedis or British pounds at the end of the session, with an exchange rate of 0.05 Ghana cedi and 0.03 British pound for every point. Sessions lasted between 1.5 and 2 hours. Including show-up fees, the average earnings were 32.2 Cedis (about 10 British pounds) in Ghana and about 18 pounds in the UK.

For the experiments we developed our own tablet-based mobile lab, LabBox. This platform can operate completely independently from mains electricity or existing network structures. The experiments run on 7-inch Android tablets with a custom-built app, which collects the user input and communicates with the LabBox central server using a wireless connection. Each session started with a 15 minute instruction on how to use the touch screen of the tablet, followed by an extensive demonstration of how the game is played. The experiment was completely conducted in English, the common language of instruction higher education in both countries. Visual on-screen aids were used to tell participants of the prospective earnings every time they were about to make a specific choice (such as making an offer), to make sure that participants were aware of the payoff structure of the game.

1.3.4 Theory

The gift-exchange game is a *sequential* prisoner's dilemma game. The game is sequential in the sense that the worker makes a choice only after observing the employer's choice. In infinite games, the Folk Theorem tells us that grim trigger strategies can sustain cooperation. Also other strategies, such as tit-for-tat strategies can sustain cooperation in games like these (see e.g. Axelrod, 1984). The main difference between the grim trigger and the tit-for-tat strategy is that tit-for-tat allows for redemption, while for the grim trigger the punishment lasts forever. However, *conditional reciprocity* is an element of both: a player will only cooperate if the other player cooperated in the past as well.

Our experiment is a finite period game. Backwards induction, assuming that players only maximize their own earnings, predicts non-cooperation as the subgame perfect equilibrium: in the last period workers will choose low effort, regardless of the wage offered by the employer. Anticipating this, employers will choose a minimum wage of zero, such that the worker is indifferent between accepting or rejecting the offer. By backwards induction, the subgame perfect equilibrium is that the employer offers a wage of zero points in every period, regardless of past behaviour of the worker, and that the worker chooses low effort in every round, regardless of the amount offered by the employer: cooperation unravels.

However, many experimental studies have shown that cooperation is possible in a finite game. Selten & Stoecker (1986) and Andreoni & Miller (1993) show that in a finite period prisoner's dilemma game subjects tend to cooperate for some period of time, until the last couple of periods, when people defect. Theoretical models have tried to explain this, for example by assuming that players have incomplete information on the other player's options or motivations. Kreps et al. (1982) show that if a player has a belief that there is a small probability that another player is acting "irrationally", i.e., not playing according to the subgame perfect equilbrium, but for example to a tit-for-tat strategy instead, cooperation can be feasible in a finite game. Fudenberg & Maskin (1986) use a similar argument to show that in a finite repeated game an equilibrium is possible where in the first periods of the game the game is played like an infinitely repeated game, and where players play a cooperation with punishment strategy, while switching to behaviour consistent with backward induction in the last periods of the game.

A similar result as Kreps et al. (1982) can also be achieved by incorporating other-regarding preferences, such as inequity aversion and fairness concerns into the utility function of some of the players. In the Appendix we present a model in which there are two types, a social type (S-type) and a rational self-interested type (R-type). The S-type's effort choice is solely dependent on the wage offered: higher wages attract higher levels of effort. The R-type is purely self-interested and aims to achieve the highest payoff. In the Appendix we show that when the share of S-types in the population is high enough, a Perfect Bayesian equilibrium exists in which the R-type's optimal response is to mimic the S-type until the second to last period and for the employer to keep offering a high wage as long as the worker chooses high effort. In the last period the R-type will choose low effort. If the worker chooses low effort, the employer infers that the worker must be the R-type and will subsequently offer a zero wage. This punishment by the employer following low effort provides an incentive for the R-type to choose high effort, as long as the high wage outweighs the cost of high effort. In this equilibrium, the R-type tries to build a reputation for being a S-type, allowing him or her to capture parts of the surplus of exerting high effort.

Depending on the wage parameters and the distribution of types in the population, other types of Perfect Bayesian equilbria are possible as well, for example in which the employer offers high wages in the first periods of the game and then a medium-level wage in the last period or in which the employer offers medium level wages throughout.

The behaviour of the social type can be rationalised by other-regarding preferences, such as models of altruism or inequality aversion. For example, in the Fehr-Schmidt model of inequality aversion (Fehr & Schmidt, 2002) players receive a disutility from both earning more than the other player (advantageous inequality) or from earning less than the other player (disadvantageous inequality). The utility function in the two-player case can be written as

$$U_i(x) = x_i - \alpha_i \max\{x_j - x_i, 0\} - \beta_i \max\{x_i - x_j, 0\}. \tag{1.3}$$

In this equation α_i represents the disutility from disadvantageous inequality and β_j represents the disutility from advantageous inequality. Generally, it is assumed that $\alpha_i \geq \beta_i$: a player has a higher disutility from disadvantageous inequality than from advantageous inequality. We will use the Fehr-Schmidt model of inequality aversion later in this paper to structurally estimate the reciprocity parameters of our participants in Ghana and the UK.

Competition and reputation

In treatments (3C), (3E) and (3ES) there are three workers and three employers. In our setup there is no market imbalance, as each worker can in principle find an employer and vice versa. A purely self-interested profit-maximizing principal will prefer to hire a worker that is willing to exert high effort for a low wage. If employers had full information, they would compete with each other to hire this worker, which can result in increasing the wages to outbid the other employers. In Davies & Fafchamps (2016) we show that if there is heterogeneity in the minimum wage that a worker needs to accept an offer, competition between employers to hire the worker with the lowest minimum wage will lead to employers increasing the wages offered to these workers to the level of the worker with the highest minimum wage.

Introducing reputation in treatment (3ES) can have multiple effects. First, by revealing past compliance, an employer can identify more reliable workers, which can lead to increased competition to hire this worker and therefore to wage increases. Second, the sharing of information on compliance can also function as an incentive, if employers use this information in their offers. Non-compliance for a worker can now become more costly: not only will the employer who he or she was working for lower their offer, also other employers will lower their offers. Greif et al. (1994) argue that such a multilateral reputation mechanism can function as an enforcement device and deter cheating.

1.3.5 Predictions

On the basis of the theoretical model as well as the results from earlier experiments, we predict three main patterns of contracting to happen.

First, we expect employers to offer high wages in the initial periods, rather than the zero wage predicted by the subgame perfect equilibrium of purely self-interested rational agents. The theoretical model shows that if there is a sizeable group of reciprocating workers, offering a high wage is rationally the optimal choice for the employer.

Second, we expect conditional reciprocity on behalf of the workers. Workers will reciprocate high wages with high effort. This behaviour is for example implied by other-regarding models of altruism or inequality aversion. As we show in the Appendix, according to the Fehr-Schmidt model of inequality aversion, a worker with inequality aversion coefficients of $\alpha_i = 0.5$ and $\beta_i = 0.5$ will choose high effort when the wage is 19 points or higher, choose medium effort for wage offers between 8 and 19 points and choose low effort for lower offers. This positive reciprocity by the workers has been shown consistently across many gift-exchange game experiments and happens even when there are no future interactions present Fehr et al. (1997, 1993, 1998); Fehr & Falk (1999).

Third, we expect conditional reciprocity on behalf of the employers. This means that we expect employers to reduce wages of workers choosing low effort or to cut the relationship completely. In the finite period theoretical model with types, this is the outcome of screening, as low effort reveals the type of the worker. The threat of lowering the wage provides an incentive even for rationally self-interested workers to exert high effort (provided their incentive compatibility constraint is satisfied).

Earlier gift-exchange game experiments have confirmed this presence of conditional reciprocity on behalf of employers. In the experiment of Brown et al. (2004, 2012) employers implemented a strategy of contingent contract renewal, in which they were more likely to renew contracts of workers choosing high effort. Workers choosing low effort were either fired or saw their wage reduced. These three predictions can be tested by analyzing the contracting behaviour in treatment (1E), where one employer can make an offer to one worker only, and where the worker can choose effort.

Furthermore, we anticipate that introducing competition will increase wages, due to the wage pushing up effect of bidding. This can be tested by comparing wage offers in treatment (3E) with wage offers in treatment (1E). Moreover, we anticipate that sharing information on the past workers' actions increases compliance as the availability of a multilateral reputation mechanism increases the scope for collective punishment. This can be tested by comparing wage offers following low and high effort in treatment (3ES) with wage offers in treatment (3E).

1.4 Results

In our discussion of the results we start by discussing the main result of the paper, by discussing the three predictions for relational contracting: first, employers offer higher wages than what the the classical subgame perfect equilibrium with purely self-interested agents predicts; second, workers reciprocate conditionally by exerting high effort after receiving a high wage; and third, employers reciprocate conditionally by offering a high wage following high effort. This conditional reciprocity on both sides makes deviations from cooperation costly and sustains the pattern of cooperation. While we find evidence for the first two predictions in both Ghana and the UK, we do not find evidence for the third prediction in Ghana: employers are not rewarding high effort or punishing low effort.

Next, we discuss the impact of competition, by comparing treatment (3E) with treatment (1E), and reputation, by comparing treatment (3ES) with treatment (3E). We do not find any difference when it comes to effort exerted. Employers in the UK are slightly more likely to increase wages to workers with whom they interacted in the past or to workers with a good reputation.

1.4.1 Patterns of relational contracting in the bilateral treatment

In this section we describe the extent we find these patterns in the bilateral treatment (1E) in our Ghanaian and UK experiments and the differences between the two countries. Furthermore, we discuss how these difference relate to effort over time and earnings.

Offers

We start by presenting evidence for the first prediction, that employers make wage offers that are above the minimum wage predicted by the subgame perfect equilibrium.

Wages We confirm the first finding from previous experiments for the bilateral treatment (1E) in both our UK and Ghanaian sessions. Figure 1.1 shows the average wage offers

Table 1.5: The average wage, the share accepted, compliance and the average earnings in the various treatments.

Treatment		Ghana					United Kingdom				
		Average wage offer	Share accepted	Average compliance	Average employer's payoff	Average worker's payoff	Average wage offer	Share accepted	Average compliance	Average employer's payoff	Average worker's payoff
(1C)	Game 1	20.0	76.7%	100.0%	11.0	16.8	18.6	77.9%	100.0%	17.0	14.3
(1C)	Game 2	20.0	75.7%	100.0%	18.3	14.2	-	-	-	-	-
(1C)	Game 3	20.2	87.1%	100.0%	16.1	15.3	-	-	-	-	-
(1C)	Game 4	21.4	84.3%	100.0%	16.1	17.1	-	-	-	-	-
(1E)	Game 2	15.3	86.1%	42.1%	0.2	14.5	12.9	78.8%	56.1%	5.9	12.4
(1E)	Game 3	13.3	92.0%	40.2%	−1.1	12.4	13.6	80.7%	60.3%	6.3	12.7
(1E)	Game 4	13.1	88.0%	38.6%	−1.6	12.9	13.9	86.0%	72.1%	9.7	11.7
(3C)	Game 2	20.4	30.1%	100.0%	11.7	17.2	-	-	-	-	-
(3C)	Game 3	20.6	29.3%	100.0%	12.0	17.5	-	-	-	-	-
(3C)	Game 4	19.8	28.9%	100.0%	11.4	17.3	-	-	-	-	-
(3E)	Game 3	14.7	30.7%	45.3%	−0.0	14.3	12.9	30.4%	62.5%	6.5	12.6
(3E)	Game 4	11.7	30.4%	40.2%	1.5	12.0	13.5	29.9%	76.0%	10.3	13.0
(3ES)	Game 4	13.2	30.4%	53.9%	1.7	13.1	13.2	29.4%	75.6%	9.4	12.7

Note: The above figures are averaged over all five periods in each game. *Compliance* is defined as whether a worker chose the level of effort that was demanded, or a higher level of effort. For the multilateral treatments, *Share accepted* indicates the share of offers that were taken up, regardless of whether that offer was presented to any of the workers. As employers can make offers to three workers and only one can be taken up, this figure is likely to be smaller in the multilateral treatments.

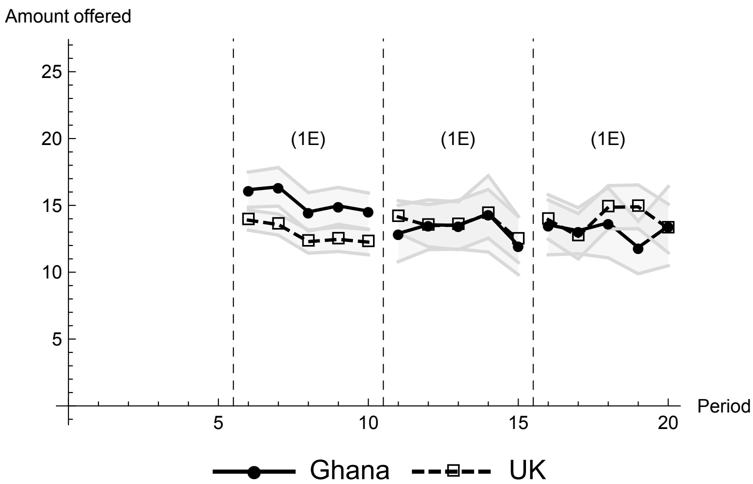

Figure 1.1: Average wage offers in treatment (1E) in Ghana and the UK.

across all periods in treatment (1E) for both Ghana and the UK. Table 1.5 shows the average offered wages, the share accepted offers, compliance with demanded effort and the average earnings for all treatments in the United Kingdom and Ghana. The average offered wage in game 2 of treatment (1E) is 15.3 in Ghana and 12.9 in the United Kingdom. In game 3 and 4 the average offered wages are respectively 13.3 and 13.3 in Ghana and respectively 13.6 and 13.9 in the United Kingdom (see also Table 1.5). These average wage offers are significantly higher than the wage predicted by the subgame perfect equilibrium of the model of rational and purely self-interested agents (0 or 1 points). This is in line with findings from earlier bilateral gift-exchange experiments (e.g., Kirchler et al., 1996; Fehr et al., 1998; Gächter & Fehr, 2001).

Result 1.1. *In both countries, the average wage offers in treatment (1E) are higher than the wage predicted by the subgame perfect equilibrium of the model of rational and purely self-interested agents.*

This confirms our first prediction.

Wage offers like these expose employers to losses if a worker chooses low effort. The average wage offer is significantly higher than the employer's benefit of low effort (5 points), which means that unless the worker chooses medium or high effort an employer that offers this average wage will face a loss. Again, this is in line with the results from earlier gift-exchange game experiments as well as trust game experiments, which showed that participants are willing to make themselves vulnerable to choices made by others in order to encourage reciprocity.

We see few differences between our Ghanaian and UK sessions when comparing average wage offers in treatment (1E). The differences are small, especially in the later games, and mostly insignificant. In Ghana, compared to the United Kingdom, average wage offers are 2.4 points higher in game 2, 0.4 points higher in game 3 and 0.8 points higher in game 4. However, only the difference in game 2 is significant and only at the 10% level (the p value is 0.066).[7]

Result 1.2. *The differences between the UK and Ghana in average wage offers in treatment (1E) are small. Average wage offers are slightly higher in Ghana, but this is only statistically significant in game 2, at the 10% level.*

[7] The t test of the difference in means in game 2 yields a p value of 0.066 when clustering at the individual level. For games 3 and 4 the corresponding p values are 0.878 and 0.738.

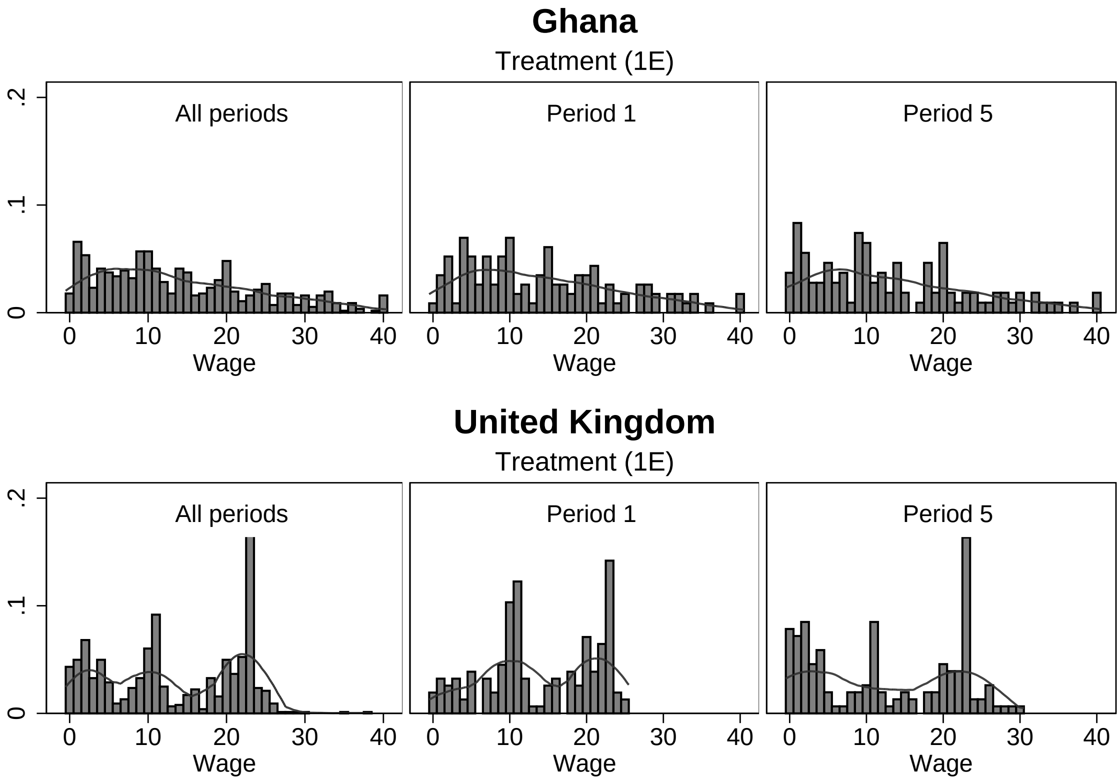

Figure 1.2: Histograms of the wages offered in treatment (1E), games 2-4, in Ghana and the United Kingdom. The line represents a kernel density estimate, calculated using a Epanechnikov kernel.

This means that we do not find much difference between the two countries to the extent that the first prediction holds.

The averages mentioned so far pool data from all periods and therefore partially incorporate the employer's response to the worker's choices. In contrast, the wage offered in period 1 cannot, by construction, depend on worker effort and is therefore more informative of the employer's initial expectation. However, the results are very similar and we do not find evidence for a significant difference between the UK and Ghana. The initial average wage offers in games 2, 3 and 4 of treatment (1E) are respectively 16.2, 12.9 and 13.6 in Ghana and respectively 13.9, 14.1 and 13.9 in the United Kingdom. None of the differences between the UK and Ghana are significant.[8]

Figure 1.2 shows the distribution of wage offers in treatment (1E). In Ghana the distribution of wage offers is more spread out than in the United Kingdom, with only small peaks. The distribution in the United Kingdom is multimodal, concentrated around 3, 11 and 23 points. These levels correspond to more or less equal earnings when the worker

[8]The offer in period 1 of game 2 is 2.3 points higher in Ghana than in the United Kingdom, but this difference is marginally insignificant (the p value of the t test is 0.107). In later games the differences are smaller.

chooses respectively low, medium or high effort.

However, non-parametric equality of distribution tests show that the differences in distributions are mostly insignificant. Both the Kolmogorov-Smirnov and the Mann-Whitney U rank-sum tests fail to reject the null hypothesis of equal distributions (e.g., in game 2, the Kolmogorov-Smirnov test gives $p = 0.393$ and the U-test gives $p = 0.234$, with $Z = -1.189$). The Epps-Singleton test rejects the null hypothesis of equal distributions in game 2, but only at the 10% level ($p = 0.076$).[9]

We also confirm the finding of Brown et al. (2004) that contractual incompleteness leads to lower offers. Compared to treatment (1C), where the workers cannot choose effort, average wage offers in treatment (1E) are lower. The average wage offer in treatment (1C) is 20.2 in Ghana and 18.6 in the UK, while the average wage offer in treatment (1E) is 14.5 in Ghana and 13.2 in the United Kingdom.[10] The difference between treatment (1C) and (1E) in average wages are significant for both countries (for both countries, a *t*-test of the difference in means yield *p*-values smaller than 0.001). In Appendix Section 1.B.1 we calculate the effect of imperfect enforcement using a within-subject, between-subject and difference-in-difference approach. Most tests indicate that the difference in wage offers between treatment (1E) and (1C) is significant.

Effort demands Figure 1.3a shows the effort demanded by the employers in treatment (1E). For a majority of offers in both Ghana and the United Kingdom (respectively 51.4% and 51.3% of the offers), employers demand high effort. Remarkable is the substantial share of employers in both countries that demand low effort in their offers. 12.3% of the offers in Ghana and 14.2% of the offers in the United Kingdom demand low effort. It might seem surprising that these employers do not ask for a higher level of effort, given that effort demand is a form of "cheap talk" and that the worker cannot choose a lower level of effort than low effort. However, in most cases, low effort is only demanded when a low wage wage is offered. One explanation for this behaviour is that employers anticipate that a

[9] All tests were conducted with unmatched data pairs, with the offer averaged across the five periods for each employer, such that each employer counts as one observation in the test. Wherever appropriate, the tests were two-sided and exact *t* statistics were used. Wage offers are discrete and therefore we expect the Kolmogorov-Smirnov test to be more conservative, as ties in the distribution are more likely to exist. As we averaged the offers across five periods, this becomes less of a concern (out of 155 combined observations, there are 96 unique observations). The discreteness of the data is not a problem for the U-test or the Epps-Singleton test. The Kolmogorov-Smirnov test detects differences in the empirical distribution functions, while the U-test detects locational shifts. The Epps-Singleton test detects differences in the distribution characteristic function, based on the Fourier transformation of the distribution (see e.g. Georg & Kaiser, 2009).

[10] These figures are averaged over all games (see Table 1.5 for the breakdown by game). Note that treatment (1C) only occurred in game 1 of the experiment in the UK. The difference between the UK and Ghana in average wage offers in treatment (1C) is significant at the 10% level (a *t*-test yields $p = 0.080$, clustered at the employer level).

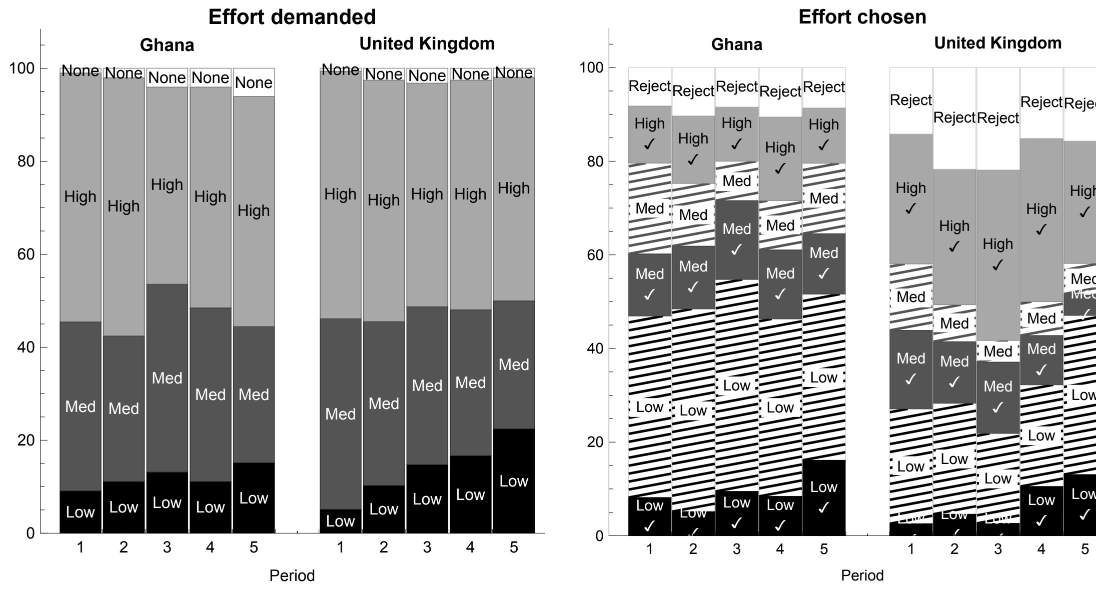

(a) Effort demanded **(b)** Effort chosen

Figure 1.3: The distribution of demanded effort and chosen effort in the (1E) treatment. The dashed areas indicate that the effort chosen was lower than the effort that was demanded. *Note:* In the left graph, *None* means that no offer was made.

worker is more likely to reject a low wage offer asking for high effort than a similar offer asking for low effort.[11]

Response of the workers

We now present evidence for our second prediction of relational contracting, *conditional reciprocity on behalf the workers*: workers reciprocate high wages with high effort and low wages with low effort or a rejection of the offer. We first focus on the acceptance of offers and then on effort choice.

Acceptance Figure 1.3b shows a breakdown of rejection and effort choice in treatment (1E) for both countries. Across all periods and games, 12.3% of the offers in Ghana and 19.5% of the offers in the United Kingdom in treatment (1E) are rejected. Figure 1.4 shows non-parametric local polynomial regressions of the rate of acceptance and compliance on the wage offered. This Figure shows that workers are more likely to reject low wage offers than high wage offers. The rejection rate of offers with a wage of five or less points is 23.6% in Ghana and 45.7% in the United Kingdom. Table 1.6 presents a linear probability model of acceptance and compliance on wage, for both Ghana and the UK as well as for

[11] There is some evidence for this behaviour on behalf of workers in the UK. Of the wage offers of five points or less, workers reject 64.5% of the offers asking for high effort, but 51.4% and 28.6% of the offers asking for respectively low or medium effort. We do not find large differences in Ghana: the rejection rates for low wage offers asking for low, medium and high effort are respectively 24.0%, 21.2% and 25.0%.

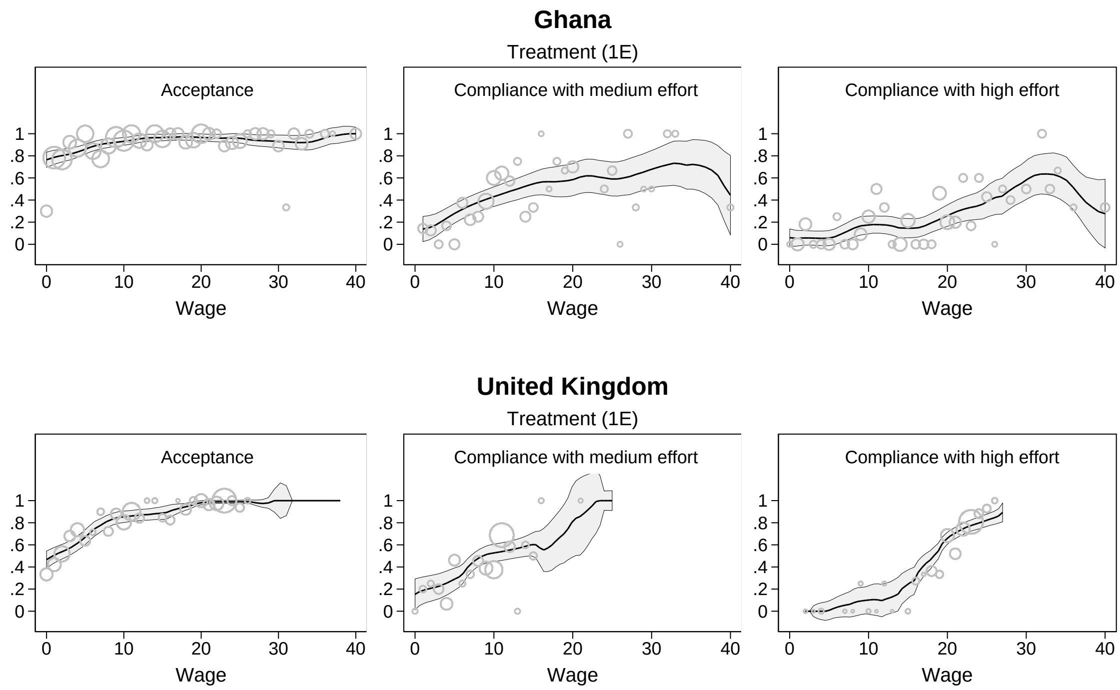

Figure 1.4: The relationship between the wage offered and acceptance, compliance with offers demanding medium effort and compliance with offers demanding high effort, in games 2-4 of treatment (1E). Compliance is defined as having chosen the demanded level of effort or a higher level of effort. These figures include data from all periods in each game. The size of the bubbles are scaled with the amount of offers made at this wage level. The solid line is a non-parametric local polynomial regression of the outcome variable on the wage. The confidence levels are at the 95% level.

both countries pooled. The linear probability model in Table 1.6 confirms this positive relation between wage offers and acceptance. This relation is stronger in the UK than in Ghana: according to the results of the linear model, a one point wage increase is related to an increase in the probability of acceptance by 0.96 percentage points in Ghana and 2.20 percentage points in the UK. The difference between Ghana and the UK in the size of this coefficient is significant at the 1% level, as the interaction term in Column (7) in Table 1.6 shows.

Result 1.3. *High wage offers are more likely to be accepted. This relationship between wages and acceptance is stronger in the United Kingdom than in Ghana.*

This result is in in line with the second prediction of conditional reciprocity on behalf of the workers.

Compliance with effort demands The dashed areas in Figure 1.3b indicate non-compliance with the employer's effort demand. In these cases, the worker chose a lower effort than what the employer demanded. As these areas show, there is considerable amount of non-compliance in both Ghana and the UK. We define compliance as a worker choosing the demanded effort level or a higher level of effort. Compliance is higher in the UK than in Ghana: in Ghana the level of compliance is respectively 42.1%, 40.2% and 38.6% in games 2, 3 and 4 of treatment (1E), while the corresponding figures are 56.1%, 60.3% and 72.1% in the United Kingdom (see Table 1.5).

For both countries, we confirm the second prediction of conditional reciprocity on behalf of the worker. We find a positive wage-effort relation found in earlier studies. Both the non-parametric regressions in Figure 1.4 and the linear probability model in Table 1.6 show a positive relationship between wage and compliance.[12] The linear probability model shows that a one point wage increase is related to an increase in the probability of complying with a high effort demand by 1.56 percentage points in Ghana and 2.60 percentage points in the United Kingdom.[13] Both coefficients are significantly different from zero, at the 5% level. However, despite the coefficient for the UK being higher, the difference between Ghana and the UK in the size of the coefficient is not significant, as the interaction

[12] As can be seen in Figure 1.4, the relationship between wage and compliance resembles an inverted U-shape. However, the decline in compliance following higher wage offers is only driven by a few observations (there are few offers above 35 points) and is mainly driven by low effort choice following offers of 40 points. The offers of 40 points are difficult to rationalize: even in the best scenario for the employer, when the worker chooses high effort, the employer's payoff is zero.

[13] The figure and the tables pool data from all periods within each game and might therefore mask certain dynamic effects. However, inclusion of an interaction term between the wage and the period dummy in the regressions in Table 1.6 does not yield significant coefficients at the 5% level.

Table 1.6: Compliance and acceptance in relation to the wage in treatment (1E).

	Ghana			United Kingdom			Both countries (pooled)		
	(1) Acceptance	(2) Compliance *(medium demand)*	(3) Compliance *(high demand)*	(4) Acceptance	(5) Compliance *(medium demand)*	(6) Compliance *(high demand)*	(7) Acceptance	(8) Compliance *(medium demand)*	(9) Compliance *(high demand)*
Wage	0.00959^{**}	0.0100	0.0156^{**}	0.0220^{***}	0.0322^{**}	0.0260^{***}	0.00959^{***}	0.0100	0.0156^{***}
	(0.00239)	(0.00717)	(0.00495)	(0.00235)	(0.0136)	(0.00628)	(0.00224)	(0.00678)	(0.00460)
Wage × UK							0.0124^{***}	0.0221	0.0104
							(0.00322)	(0.0150)	(0.00775)
Constant	0.784^{***}	0.398	-0.0330	0.527^{***}	0.178	0.0519	0.778^{***}	0.356^{**}	-0.0944
	(0.0395)	(0.244)	(0.154)	(0.0389)	(0.201)	(0.161)	(0.0512)	(0.145)	(0.105)
N	478	160	224	763	217	358	1241	377	582
Adjusted R^2	0.225	0.513	0.297	0.316	0.398	0.450	0.238	0.235	0.433
Fixed effects	Worker	Worker	Worker	Worker	Worker	Worker	Worker	Worker	Worker
Period dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes†	Yes†	Yes†

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
† These regressions also include the interaction of the period dummies with the country dummy.

Note: This is a fixed effects linear probability model regression of acceptance and compliance on the wage offered. *Acceptance* is an indicator variable and is equal to one if the worker accepted the offer. *Compliance* is an indicator variable and is equal to one if the worker chooses the effort level specified by the employer, or a higher effort level. For the *Compliance (medium demand)* and *Compliance (high demand)* columns the regression only includes observations where medium or high effort is demanded by the employer. The regressions include worker fixed effects. Columns (7), (8) and (9) also include the interaction of the period dummies with the country dummy. Standard errors are clustered by session.

terms in Column (9) shows ($p = 0.196$).[14] The coefficients in the regression of compliance with medium effort demands show a similar pattern, even though not all coefficients are significant.

Result 1.4. *In both countries, compliance with effort demands is higher for higher wage offers.*

We therefore find in both countries strong evidence for the second prediction of conditional reciprocity on behalf of the worker.

Wage revision by employers

We now discuss the evidence for the third prediction, *conditional reciprocity on behalf the employer*. Models of relational contracting predict that employers continue to offer a high wage as long as the worker chooses high effort (reward behaviour) and low effort leads to the employer ceasing to offer high wages (punishment behaviour). This behaviour forms an incentive for workers to exert high effort, since the punishment following low effort becomes too costly.

[14]The regression specifications in Table 1.6 include worker fixed effects and therefore control for individual differences in the likelihood of accepting or complying. The coefficients should therefore be interpreted as the worker's response to wage variations that a worker faces throughout the experiment. Excluding fixed effects in the regression does lead to a significant difference in the size of the coefficient between Ghana and the UK, as can be seen in Appendix Table 1.19.

Table 1.7: Response to the previous period's effort for above median wages in treatment (1E)

	Ghana				United Kingdom			
	Worker's effort in period $t-1$ for above median wages ($w_{t-1} \geq 15$)				Worker's effort in period $t-1$ for above median wages ($w_{t-1} \geq 15$)			
	Low	Medium	High	All	Low	Medium	High	All
Panel A. Response in period 2								
Share of effort in period 1	37.0%	39.1%	23.9%	100.0%	10.1%	33.3%	56.5%	100.0%
Employer's response in period $t = 2$								
Decrease wage	41.2%	38.9%	54.6%	43.5%	71.4%	60.9%	10.3%	33.3%
Same wage (± 2 points)	29.4%	38.9%	36.4%	34.8%	0.0%	26.1%	84.6%	56.5%
Increase wage	29.4%	16.7%	9.1%	19.6%	14.3%	8.7%	5.1%	7.3
No offer	0.0%	5.6%	0.0%	2.2%	14.3%	4.4%	0.0%	2.9%
Panel B. Response in all periods								
Share of effort in period $t-1$	41.3%	32.3%	26.4%	100.0%	11.0%	20.9%	68.1%	100.0%
Employer's response in period t								
Decrease wage	50.7%	29.6%	45.5%	42.5%	70.0%	70.2%	6.5%	26.7%
Same wage (± 2 points)	24.6%	46.3%	47.7%	37.7%	10.0%	19.3%	88.2%	65.2%
Increase wage	21.7%	20.4%	4.6%	16.8%	13.3%	8.8%	4.3%	6.2%
No offer	2.9%	3.7%	2.3%	3.0%	6.7%	1.8%	1.1%	1.8%

Note: Panel B includes the responses from periods 2-5. Only responses to effort following above median wages are included in this table. Rejected offers are not included in this table. We use a 2 point margin in the wage classification. Wages that were more than 2 points lower are classified as a wage decrease, and wage increases of more than 2 points are classified as a wage increase. Wages within these 2 points margin as classified as being the same wage.

We first focus on wage revision following a high wage, with the median wage of 15 points (inclusive) as cutoff.[15] Table 1.7 shows the employer's response in treatment (1E) to low, medium and high effort, provided that the employer offered an above median wage in the previous period. The top rows (Panel A) relate to the response in the second period, while the bottom rows (Panel B) pool data from periods 2 to 5. The Table highlights two main differences between our British and Ghanaian sessions in how employers revise wages. First, following low effort, Ghanaian employers are less likely to *decrease* their wages than the British employers.[16] In Ghana, 41.2% of the wages are lowered in period 2 following low effort in period 1, while in the UK the corresponding figure is 71.4%. Across all periods, these figures are respectively 50.7% and 70.0%.

Second, following high effort, British employers are more likely to *keep offering the same wage* than the Ghanaian employers. In the UK, 84.6% of the wages are kept the same in the second period after high effort in period 1, while the corresponding figure in Ghana is 36.4%. Ghanaian employers are more likely to *lower* wages following high effort: in Ghana 54.6% of the wages were lowered in period 2 following high effort in period 1, while in the

[15]The median wage is 15 when pooling the data from the UK and Ghana together. The median wage in Ghana is 14 and the median wage in the United Kingdom is 16.5.

[16]Just like in Table 1.7 we adopt a 2 point margin in the figures presented in the text (i.e., a wage of 24 following a wage of 22 is considered as the "same" wage). This is to account for potential imprecise input from the tablet's touch screen. None of the results change qualitatively when no margin or a margin of 1 point is used.

Table 1.8: Linear regression of wage offers on rejection and previous compliance in treatment (1E).

	Ghana		**United Kingdom**		**Both countries (pooled)**	
Dependent variable: Wage offer in period t	(1)	(2)	(3)	(4)	(5)	(6)
Rejection in period $t-1$?	-1.192	-1.087	2.487^{***}	2.645^{***}	-1.172	-1.043
	(1.376)	(1.339)	(0.773)	(0.811)	(1.309)	(1.292)
Compliance in period $t-1$?	0.283		6.260^{***}		0.221	
	(0.902)		(1.228)		(0.923)	
Compliance in period $t-1$? (high effort demanded)		0.744		10.21^{***}		0.678
		(2.356)		(1.460)		(2.285)
Compliance in period $t-1$? (medium effort demanded)		0.875		4.906^{***}		0.822
		(0.945)		(1.384)		(0.973)
UK $\times$ Rejection in period $t-1$?					3.684^{**}	3.676^{**}
					(1.521)	(1.528)
$\times$ Compliance in period $t-1$?					6.092^{***}	
					(1.502)	
$\times$ Compliance in period $t-1$? (high effort demanded)						9.544^{***}
						(2.750)
$\times$ Compliance in period $t-1$? (medium effort demanded)						4.047^{**}
						(1.689)
Constant	16.79^{***}	16.66^{***}	11.05^{***}	10.45^{***}	13.25^{***}	12.85^{***}
	(1.239)	(1.414)	(0.828)	(0.631)	(0.662)	(0.629)
Observations	380	380	608	608	988	988
R^2	0.0475	0.0496	0.158	0.266	0.0669	0.0882
Adjusted R^2	0.0136	0.0132	0.139	0.249	0.0545	0.0742
Fixed effects	Employer	Employer	Employer	Employer	Employer	Employer
Period dummies	Yes	Yes	Yes	Yes	Yes	Yes

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: This is a fixed effects regression of wage offered on previous compliance and rejection. *Rejection* is an indicator variable and equal to one if the worker rejected the offer. *Compliance* is an indicator variable and equal to one if the worker chose the effort level demanded by the employer, or a higher level. Only wage offers in period 2-5 are included. Standard errors are clustered on the session level.

United Kingdom only 10.3% of the wages were lowered.

The lower likelihood of wage cuts following low effort and the higher likelihood of wage cuts following high effort in Ghana is also reflected in the coefficients in the linear regression of wage offers on past compliance and past rejection. Table 1.8 shows this regression. We see that in the United Kingdom, compliance is related to a wage increase of 6.260 points in the next period. This figure is significantly different from zero at the 1% level. In Ghana the corresponding coefficient is 0.283, meaning that, on aggregate, compliance is associated with a wage increase of only 0.283 in the next period. However, this coefficient is not significantly different from zero and small in magnitude, given that wages can range between 0 and 40 points. The difference between the UK and Ghana in the employer's response to compliance is significant: in the last two columns of this Table we pool the observations from Ghana and the UK together and add an interaction term between compliance and the country indicator, which is significant at the 1% level.

Result 1.5. *There is a strong and positive relation between the offered wage and compliance in the previous period in the United Kingdom. There is no such relation in Ghana. Employers in the United Kingdom are more likely to decrease wages following low effort than employers in Ghana.*

The behaviour in the United Kingdom is in line with the third prediction of conditional reciprocity on behalf of the employer. However, the lack of a relation between compliance and wage revision in Ghana stands in a stark contrast with this as well as earlier results from experiments in high income countries.[17]

Consequences

We continue by discussing how, following this lack of punishment and reward in Ghana, effort evolves over time in Ghana and the United Kingdom. We also discuss the employers' and workers' earnings and compare these with the earnings in treatment (1C).

Evolution of effort over time Table 1.9 shows the transition matrix of effort for treatment (1E) for both Ghana and the United Kingdom. This table omits observations from period 5 to exclude final period effects. In the United Kingdom, there are few transitions out of high effort. Of the workers that chose high effort in one period, 74.0% will choose high effort again in the next period. The corresponding figure in Ghana is 48.5%. This shows that

[17] For example, in Brown et al. (2012) the coefficient on the relation between effort chosen and the subsequent rent that is offered is positive and significant (Table 4, Column (6)). The coefficient has a magnitude of 5.527 and is significant at the 1% level. In their experiment, effort can take 10 values and offered wages can vary between 0 and 100. This coefficient is based on data from both the ICF-S and ICF-D treatments. The data for the ICF-S treatment, with an excess supply of labour, is the same data as used in Brown et al. (2004).

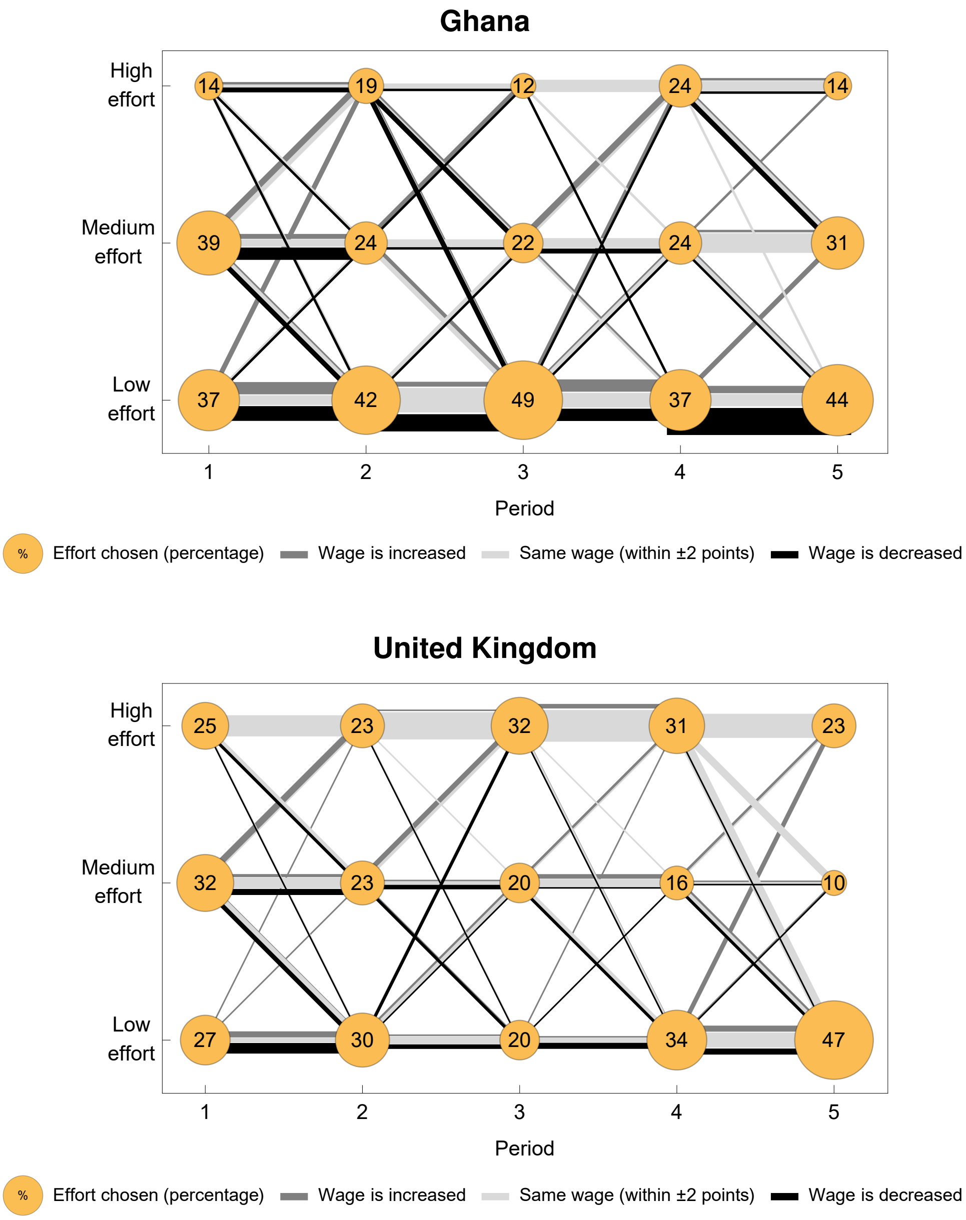

Figure 1.5: Transitions between effort levels of workers in game 2 of the (1E) treatment. The numbers and the size of the circles indicate the percentage of workers choosing this effort level in that particular period. The thickness of the lines indicates the share of transitions between effort levels. The shading of the lines indicate whether the employer increased the wage (with more than 2 points), kept it the same (within a ±2 point range) or decreased the wage (with more than 2 points). Rejections are not included in this graph.

Table 1.9: Evolution of effort in treatment (1E)

Ghana

$\downarrow t-1$ \ $t \rightarrow$	Low	Medium	High	No offer	Rejected
Low	69.6	10.9	5.4	4.3	9.8
Medium	25.9	37.0	22.2	5.6	9.3
High	24.2	24.2	48.5	0.0	3.0
No offer	37.5	25.0	25.0	12.5	0.0
Rejected	30.0	20.0	5.0	10.0	35.0

United Kingdom

$\downarrow t-1$ \ $t \rightarrow$	Low	Medium	High	No offer	Rejected
Low	47.3	12.2	6.8	4.1	29.7
Medium	23.6	38.9	20.8	2.8	13.9
High	11.7	9.1	74.0	1.3	3.9
No offer	40.0	40.0	20.0	0.0	0.0
Rejected	30.0	16.7	8.3	1.7	43.3

Note: These are right stochastic (transition) matrices (each row adds up to 100%). The numbers represent percentages. Data from period 5 has been omitted. The rows represent the choices in period $t-1$ and the columns represent the choice in the next period (period t).

the Ghanaian participants acting as employers are failing to maintain high effort provision by their workers, while the British employers are better at this. As discussed earlier, the higher likelihood of employers in our Ghanaian sessions to lower wages after high effort could play a role.

Figure 1.5 shows the transitions between efforts graphically. The size of the bubbles and the number inside the bubbles indicate the percentage of workers choosing a particular level of effort. The thickness of the lines indicate the likelihood of a transition. This graph confirms the observation from the Table that in the UK, apart from period 5, there are few transitions out of high effort. Furthermore, the graphs show some evidence for final period effects, which are the strongest in the United Kingdom. In period 5, the share of workers choosing low effort is similar in both countries (44% against 47%).

The shading of the lines indicate whether wages increased, decreased or remained the same. For our British sessions we see that offering the same wage (the light gray line) after high effort is almost always followed by high effort again in the next period, apart from period 5. Increasing the wage (the dark gray line) is likely to increase effort, while decreasing the wage (the black line) is likely to decrease effort. In Ghana we see a similar pattern, even though employers in our Ghanaian sessions are less likely to keep offering the same wage.

Earnings

These differences in contracting and effort levels between Ghana and the UK also have consequences for the earnings of the employers and the workers in our experiment. Under the (1C) perfect enforcement condition, the earnings of the worker and employer are close to equal, especially in the later games (see Table 1.5). However, introducing imperfect enforcement, as is done in treatment (1E), reduces the employers' average earnings substantially. Compared to game 1 of treatment (1C), the average employer's earnings in game 2 of treatment (1E) drop from 11.0 to 0.2 points in Ghana and from 17.0 to 5.9 points in the United Kingdom. These differences are highly significant (see Appendix Table 1.17). The average employer's earnings in treatment (1E) are higher in the UK and this difference is significant. In later games the difference in average earnings between the UK and Ghana increases even further: while in Ghana, the average employer's earnings are negative in games 3 and 4 (at respectively -1.1 and -1.6 points), the average employer's earnings in the UK increase (to respectively 6.3 and 9.7 points). Just like in game 2, the difference between the UK and Ghana is significant.

Result 1.6. *Imperfect enforcement reduces average employer's earnings. The average employer's earnings in treatment (1E) are lower in our Ghanaian sessions than in our British sessions.*

In contrast, we do not see any significant differences in average worker's earnings, both when we compare treatment (1E) with (1C) and when comparing the United Kingdom with Ghana. Compared to treatment (1C) the average worker's earnings are slightly lower in treatment (1E) in both Ghana and the United Kingdom, but this difference is not significant (see also Table 1.5 and Appendix Table 1.17). This suggests that in Ghana it's especially the employers who are "paying the price" of the low exerted effort.

1.4.2 The role of competition

We now discuss the impact of competition, by analyzing the treatment effect of being assigned a treatment with multiple workers and employers. In treatments (3C), (3E) and (3ES) there is competition between employers to find the best worker and between workers to receive the best offers. In our design there is no market imbalance: the number of workers is equal to the number of employers.

We find no evidence for a difference in wage offers when comparing treatment (3E) with treatment (1E). For example, in game 3 of treatment (3E) the average wage is 14.7 in Ghana and 12.9 in the United Kingdom, while in game 3 of treatment (1E) the average was is 13.3 in Ghana and 13.6 in the United Kingdom (see Table 1.5).[18] The difference between treatment (3E) and treatment (1E) is not significant. Also the treatment effect calculated using a within-subject, between-subject and a difference-in-difference estimation are not significant (see Appendix Tables 1.16, 1.17 and 1.18).

[18] We focus on game 3, since treatment (3E) is introduced from that game onwards.

Figure 1.6 shows the distribution of wage offers made in the (1E), (3E) and (3ES) treatments in Ghana and the United Kingdom. For both countries, non-parametric tests do not find significant differences in the distribution of offers between these treatments.[19] This leads to the following result:

Result 1.7. *There are no significant difference in wage offers between treatment (3E) and (1E).*

This suggests that there is no strong treatment effect of introducing competition.

However, we do find a difference when comparing the offers in the competition treatments between the United Kingdom and Ghana. The average wage offer is significantly higher in the United Kingdom and non-parametric tests reject the null hypothesis of having the same distribution of offers in Ghana and the United Kingdom.[20]

Wage setting One of the key results in models of relational contracting and from earlier experiments is that employers offer higher wages to employees with whom they have successfully interacted in the past. In Brown et al. (2004), where employers could make both private and public offers, private offers were higher than public offers and employers based their wages on how successful the past interactions were.

We see a similar pattern in the United Kingdom, but not in Ghana. Table 1.10 shows how the employer's wage offer relate to the number of past interactions with a worker, regardless of whether the worker complied with the demanded effort. In the United Kingdom a higher number of past interactions is associated with a significantly higher wage offer. For example, in period 5, the wage offer is 4.85 points higher when an employer has interacted with a worker for three periods. In Ghana we do not find a significant relation between the number of past interactions and the wage offered.

We see a similar difference between Ghana and the United Kingdom when we also consider whether the worker in the past interaction complied with the demanded effort. Columns (1) and (3) of Table 1.11 show a fixed effects regression of how the employer's wage offer depends on whether the employer contracted with this worker in the past period and whether the worker complied with the demanded effort. We find significant coefficients for our British sessions, but not for our Ghanaian sessions. In the United Kingdom,

[19] A Wilcoxon (Mann-Whitney) rank-sum test and a Kolmogorov-Smirnov test do not reject the null hypothesis that there is a significant difference in the distributions between treatment (1E) and (3E) (e.g. in game 3, the Wilcoxon rank-sum test yields for Ghana $Z = -0.966$ and $p = 0.334$, and for the United Kingdom $Z = -0.022$ and $p = 0.9822$).

[20] A Wilcoxon rank-sum test rejects the null hypothesis of the same distribution in Ghana and the United Kingdom for game 4 ($Z = 2.90$, $p = 0.0037$), but not for game 3 ($Z = -0.986$, $p = 0.324$). A Kolmogorov-Smirnov test rejects the null hypothesis of the same distribution for game 3 at the 10% level ($p = 0.093$) and for game 4 at the 1% level ($p = 0.000$).

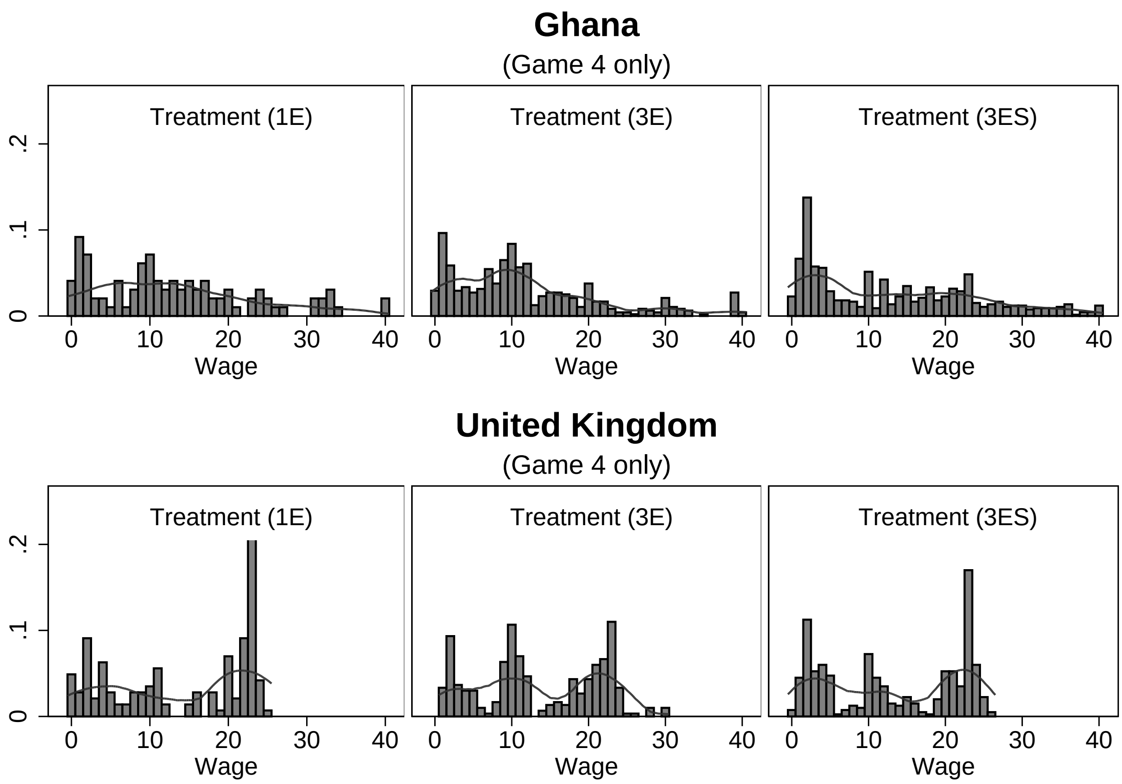

Figure 1.6: Histograms of the wages offered in treatment (1E), (3E) and (3ES), in Ghana and the United Kingdom. To allow for comparisons, only game 4 data has been included, because treatment (3ES) was only conducted as game 4. The line represents a kernel density estimate, calculated using a Epanechnikov kernel.

Table 1.10: Wages and past interactions in treatment (3E)

	Ghana					United Kingdom				
Period	1	2	3	4	5	1	2	3	4	5
Wage without past interactions	15.18	14.38	13.92	12.47	13.40	14.23	13.69	13.43	11.02	11.11
Wage increase/decrease after										
. . . 1 past interaction		−0.60	0.18	1.21	−1.56		−1.18	−1.96	−0.83	−1.72
. . . 2 past interactions			−0.92	1.98	−0.94			2.73*	−0.03	−0.91
. . . 3 past interactions				−1.11	−1.72				8.94**	4.85***
. . . 4 past interactions					−1.41					6.45**
Observations	341	329	316	310	313	273	197	200	173	180

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: The reported numbers are from a fixed effects linear regression of the wage on the number of interactions. Fixed effects are taken on the employer level. For the tests, standard errors (not reported) are clustered on the session level.

compliance is associated with a 7.512 higher wage in the next period. The corresponding coefficient for Ghana is small, at at 0.965, and not significant at a 10% level.

Result 1.8. *In the United Kingdom, employers offer higher wages to workers with whom they have successfully interacted in the past. We do not find evidence for this in Ghana.*

The result from the United Kingdom is very much in line with the findings by Brown et al. (2004), who find that employers offer higher wages to workers they successfully interacted with in previous periods, and theories of relational contracting. The results from Ghana stand in a stark contrast to this.

1.4.3 The role of reputation

Next, we discuss the treatment effect of being assigned the reputation treatment, treatment (3ES). In this condition employers can see past compliance not only for their own workers, but also for the workers that contracted with a different employer. This provides the employers with an additional mechanism for screening, which in itself can increase the incentive for workers to comply with the effort demanded.

Wages We find little evidence that reputation increases the average wage offer in Ghana, but find some evidence for this for the United Kingdom. Table 1.12 compares treatment (3ES) with treatment (3E) for a number of outcome variables, using a within subject, between groups, a (fixed effects) difference-in-difference analysis (see also Appendix Section 1.B.1). For Ghana, the within subject estimator shows a significant drop in the wage as a result of information sharing, but when we control for a time trend, as the difference-in-difference and the fixed effects estimators do, the effect disappears. For the United Kingdom, the within subject, difference-in-difference and fixed effects estimators show a

Table 1.11: Wage offers in treatment (3E) and (3ES)

	Ghana		United Kingdom	
Dependent variable: Wage offer by employer i in period t	(1) Treatment (3E) only	(2) Both treatments	(3) Treatment (3E) only	(4) Both treatments
Contracted with employer $i?_{t-1}$	-0.872	-0.682	-3.070^{***}	-3.292^{***}
	(0.837)	(0.892)	(0.885)	(0.951)
Contracted with *other* employer$?_{t-1}$	-0.614	-0.442	0.525	0.518
	(0.830)	(0.877)	(0.913)	(0.950)
Compliance with employer i's offer$?_{t-1}$	0.965	0.813	7.512^{***}	7.851^{***}
	(0.630)	(0.743)	(1.039)	(0.981)
Compliance with *other* employer's offer$?_{t-1}$	0.543	0.401	-0.0199	0.103
	(0.467)	(0.484)	(0.765)	(0.748)
Treatment (3ES)?		0.133		2.085
		(1.658)		(1.655)
× Contracted with employer $i?_{t-1}$		-0.307		1.618
		(0.955)		(0.973)
× Contracted with *other* employer$?_{t-1}$		-0.662		-0.925
		(0.965)		(1.127)
× Compliance with employer i's offer$?_{t-1}$		0.102		0.0279
		(1.245)		(0.995)
× Compliance with *other* employer's offer$?_{t-1}$		0.414		4.060^{***}
		(1.020)		(0.981)
Constant	14.92^{***}	15.14^{***}	12.43^{***}	11.67^{***}
	(0.719)	(0.799)	(0.499)	(0.576)
Observations	1268	1795	750	1038
R-squared	0.710	0.699	0.616	0.624
Adjusted R-sq	0.687	0.681	0.572	0.592
Fixed effects	Employer	Employer	Employer	Employer
Period dummies	Yes	Yes	Yes	Yes

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: This is a linear regression with the wage offered by employer i in period t as dependent variable. All independent variables are indicator variables and relate to period $t-1$, i.e. the period before the offer. Standard errors are clustered on the session level.

Table 1.12: Treatment effect of treatment (3ES) compared with treatment (3E)

	Ghana				United Kingdom			
	(1) Within	(2) Between	(3) DID	(4) FE	(5) Within	(6) Between	(7) DID	(8) FE
Offered wage								
(3ES) vs. (3E)	-2.145^{***}	1.506	-0.306	-0.262	1.681^{**}	-0.305	3.040^{***}	3.264^{***}
	(0.653)	(2.309)	(0.902)	(0.984)	(0.573)	(1.837)	(0.926)	(0.962)
Acceptance								
(3ES) vs. (3E)	-0.0112	-0.0125	-0.0143	-0.0144	0.0178	0.0267	0.0443	0.0409
	(0.0149)	(0.0164)	(0.0145)	(0.0160)	(0.0235)	(0.0433)	(0.0295)	(0.0313)
Compliance								
(3ES) vs. (3E)	0.0529	0.136	0.0643	0.0517	0.113^{**}	-0.00452	-0.0310	-0.0406
	(0.0304)	(0.0813)	(0.0526)	(0.0598)	(0.0379)	(0.0406)	(0.0643)	(0.0686)
Surplus								
(3ES) vs. (3E)	0.283	1.214	1.022	1.022	3.718^{***}	-1.377	2.785^{*}	2.785
	(1.710)	(2.171)	(1.971)	(2.076)	(1.018)	(2.909)	(1.452)	(1.528)
Employer's earnings								
(3ES) vs. (3E)	1.996	0.156	1.179	1.179	2.723^{***}	-0.943	0.0712	0.0712
	(1.625)	(1.200)	(1.697)	(1.787)	(0.538)	(1.365)	(0.836)	(0.880)
Worker's earnings								
(3ES) vs. (3E)	-0.278	1.072	1.651	2.543	-1.916	-2.644	2.139	0.623
	(1.349)	(5.109)	(2.208)	(2.271)	(1.346)	(1.838)	(2.681)	(2.889)

Note: Each cell corresponds to a separate regression. Standard errors are given in parentheses.

positive and significant effect of reputation on wage offers, while the between group estimator shows a small and insignificant negative effect.

We also find that in our British sessions employers are basing their wage offers on the information provided by the reputation mechanism in treatment (3ES), while we do not find evidence for this in our Ghanaian sessions. Columns (2) and (4) of Table 1.11 show that compliance with another's employer's offer is associated with a wage increase of 4.060 in the UK, while in Ghana the corresponding figure is only 0.414. The Ghanaian figure is not significant, while the British figure is.

Result 1.9. *In the United Kingdom employers in treatment (3ES) base their wages on the multilateral reputation of the worker. We do not find evidence for this in Ghana.*

We see that the employers in the British sessions put more importance on the worker's compliance with their own offers than on compliance with other employer's offers. The coefficient for compliance with own offers is almost double in size, at 7.851. One explanation for this could be that employers value their own experiences more and therefore "punish" their own non-complying workers more than the non-complying workers of other employers. Another explanation is that employers, despite the reputation mechanism, have better information on their own workers. The only information that is shared is compliance and not the actual wage offered or the effort that was demanded. This information is important for evaluating whether the worker was truly at fault. For example, even an honest worker might choose to not comply with a high effort demand if the wage is unreasonably low.

Effort choice and compliance We have seen that in the UK, employers base their wages on the worker's reputation. However, does this incentivize workers to comply with high effort? We do not find strong evidence that this is the case. As can be seen in Table 1.5 the rate of compliance in Ghana is 53.9% in treatment (3ES) and 40.2% in treatment (3E). In the United Kingdom rate of compliance is 75.6% in treatment (3ES) and 76.0% in treatment (3E). However, as the treatment effects estimations in Table 1.12 show, the difference in compliance between treatment (3E) and treatment (3ES) is not significant.

Also when we correct for the wage offered, we find limited evidence that having a reputation mechanism significantly increases compliance. Table 1.13 presents a linear probability model of acceptance and compliance as a function of the wage and whether the participant is in treatment (3ES). While the coefficient on the wage is positive and significant in both Ghana and the United Kingdom, the coefficient on treatment (3ES) is not significant in most columns. Only the coefficient in the regression of compliance with medium effort demands in Ghana is significant at the 10% level.[21]

Result 1.10. *We do not find significant evidence that reputation increases compliance in both Ghana and the United Kingdom.*

This means that we do not find significant evidence that multilateral reputation works as a stronger "disciplining mechanism" compared to bilateral reputation, even though multilateral reputation leads to higher wage offers.

1.5 Discussion and conclusion

We find striking differences in relational contracting patterns between the United Kingdom and Ghana. In this section we will discuss a number of alternative specifications as well as results from a series of additional experiments we ran. In the additional experiments we find a stronger response of the wage to past effort, but the relationship is weaker than in the United Kingdom. Next, we dive deeper into the question whether the differences in effort are due to different distributional preferences on behalf of the workers. To do this, we estimate a structural model of worker's preferences and categorize workers according to

[21]Even though mostly not significant, some of the effect sizes themselves are large: for example, Column (5) in Table 1.13 implies that multilateral reputation increases compliance with medium by 14.5 percentage points, when controlling for the wage.

Table 1.13: Compliance in treatments (3E) and (3ES)

	Ghana			United Kingdom		
	(1) Acceptance	(2) Compliance *(medium demand)*	(3) Compliance *(high demand)*	(4) Acceptance	(5) Compliance *(medium demand)*	(6) Compliance *(high demand)*
Wage	0.00579^{***}	0.0137^{*}	0.0101^{**}	0.0187^{***}	0.0548^{***}	0.0229^{**}
	(0.000924)	(0.00703)	(0.00424)	(0.00203)	(0.0157)	(0.00887)
Treatment (3ES) ?	−0.00566	0.217^{*}	0.127	−0.00719	0.145	−0.0264
	(0.0194)	(0.104)	(0.100)	(0.0288)	(0.0936)	(0.111)
Constant	0.226^{***}	0.437^{**}	0.188	0.0329	0.190	0.235
	(0.0256)	(0.151)	(0.120)	(0.0411)	(0.174)	(0.158)
Observations	2270	288	396	1423	186	352
R-squared	0.0283	0.617	0.664	0.153	0.779	0.643
Adjusted R-sq	−0.0137	0.445	0.559	0.106	0.645	0.549
Fixed effects	Worker	Worker	Worker	Worker	Worker	Worker
Period dummies	Yes	Yes	Yes	Yes	Yes	Yes

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: This is a fixed effects linear probability model regression of acceptance and compliance on the wage offered, with data from games 3 and 4 from treatments (3E) and (3ES). *Acceptance* is an indicator variable and is equal to one if the worker accepted the offer. *Compliance* is an indicator variable and is equal to one if the worker chooses the effort level specified by the employer, or a higher effort level. *Treatment (3ES)* is an indicator variable and equal to one if the treatment is (3ES). The *medium demand* and *high demand* columns only include observations where medium or high effort was demanded by the employer. Fixed effects are taken on the worker level. Standard errors are clustered by session.

this model. Finally, we discuss the implications of our results as well as several hypotheses of what could be driving the observed differences.

1.5.1 Robustness

Alternative specifications

Our result of the lack of wage response to past non-compliance in Ghana is robust to a number of alternative specifications. In Appendix Table 1.21 we present two additional specifications of the regression of the wage on past compliance. Including a lagged value of the wage or taking the change in the wage as the dependent variable does not change the results in either magnitude or significance.

In Appendix Table 1.22 we regress the wage offered on past effort and whether the chosen effort was a positive or negative "surprise" (i.e., higher or lower than demanded), in a similar way as Brown et al. (2012) present their results from their German experiments.[22] We find that in both Ghana and the United Kingdom higher effort is rewarded with a higher wage in the next period, but the coefficient is (significantly) higher in the UK than in Ghana. Also the response to negative surprises is significantly different: in the UK a

[22] See columns (4) and (5) of Table 4 of Brown et al. (2012).

negative surprise is related to a lowering of the wage of -2.1 points, while in Ghana this is related to a wage increase of 1.7 points.

Out-of-sample validation

In Chapter 2 we run a second series of experiments, where we focus on the role of feedback. However, these experiments also allows us to validate the experiments conducted for this paper. For these experiments, we simplified the game and also *invited real entrepreneurs* from small and medium-sized entreprises. We reduced the effort levels from three to two (only low and high effort) and made high effort the default demanded effort level. The payoff function was kept equal, only in some sessions we increased the employer's benefit from low effort from 5 to 10 or 15 points. In total we held 31 sessions with a total of 61 entrepreneurs and 559 students.

In this series of experiments we actually find a significant relationship between wage and past effort. Table 1.14 replicates the regression of wage on past compliance with high effort and past rejection. Unlike in the first series of the experiment, the coefficient on compliance in this experiment is significantly different from zero, at 1.976 points. However, this coefficient is small, given that wages can range between 0 and 40, and that the employer's loss of low effort ranges between 25 and 30 points.

We find evidence for a stronger degree of conditional wage setting by the entrepreneur participants. The coefficient for entrepreneurs is higher, at 3.523, and this figure is significantly higher than the corresponding figure for students, as the regression with the interaction term in Column (4) of Table 1.14 shows. However, compared to the British subjects from the first series of experiments, this figure is still low (the coefficient on compliance with high effort was 10.21 as Table 1.8 shows).

1.5.2 Heterogeneity of workers: a structural model

The finite period theoretical model relies on the presence of a group of "social" type workers, who act reciprocally. The "opportunistic" workers will mimic the behaviour of the social types and act reciprocally until the last period, in which they make a purely selfish choice. According to this model, the higher the share of "social" type workers is, the more likely a pattern of relational contracting can emerge.

In this section we use two different estimation strategies of a structural model of inequality aversion to show that there is little difference between the UK and Ghana when it comes to the distribution of "social" types in the population. We do not find evidence for Ghanaian workers acting more opportunistically than their British counterparts, in fact

Table 1.14: Linear regression of wage offes on rejection and previous compliance in treatment (1E) of the additional experiments in Ghana.

	Students	Entrepreneurs	Both groups (pooled)	
Dependent variable: Wage offer in period t	(1)	(2)	(3)	(4)
High effort in period $t-1$?	1.762^{***} (0.381)	3.523^{**} (0.855)	1.976^{***} (0.377)	1.762^{***} (0.380)
Rejection in period $t-1$?	1.788^{***} (0.407)	0.805 (0.825)	1.602^{***} (0.370)	1.788^{***} (0.406)
Entrepreneur?				
× High effort in period $t-1$?				1.762^{**} (0.858)
× Rejection in period $t-1$?				-0.983 (0.846)
Constant	15.30^{***} (0.278)	12.39^{***} (0.382)	14.68^{***} (0.273)	14.69^{***} (0.270)
Observations	2496	458	2954	2954
R^2	0.0269	0.0612	0.0250	0.0318
Adjusted R^2	0.0207	0.0272	0.0197	0.0212
Fixed effects	Employer	Employer	Employer	Employer
Period dummies	Yes	Yes	Yes	Yes

$^{*}\ p < 0.1$, $^{**}\ p < 0.05$, $^{***}\ p < 0.01$

Note: This is a fixed effects regression of wage offered on previous high effort choice and rejection, for the additional experiments conducted in Ghana. *Rejection* is an indicator variable and equal to one if the worker rejected the offer. *High effort* is an indicator variable and equal to one if the worker chose high effort. Only wage offers in period 2-5 are included. Standard errors are clustered on the session level.

some of our results point in the opposite direction. The structural model is derived from the model of inequality aversion by Fehr & Schmidt (1999), which is more extensively discussed in the Appendix. The utility function is given by Equation 1.3. The key parameter here is β_i, which represents the disutility from advantageous inequality, meaning that the agent received more than the agent's partner. A higher value of β_i corresponds to a higher degree of advantageous inequality aversion. We also estimate α_i, which represents disadvantageous inequality, but our experiment is underpowered to provide precise estimates.[23] A fully rational self-interested agent would have parameters $\alpha_i = \beta_i = 0$.

We use two methods to estimate these parameters: a set identification by considering which values of α_i and β_i are consistent with the choices made and a maximum likelihood estimation of a multinomial logit model. Set identification gives an idea of the potential bounds on the parameters, while maximum likelihood gives a point estimate of the parameters that predicts the observed data the best.

Set identification For the set identification, we construct bounds on the parameters α_i and β_i by calculating for each realized choice in each period the set of values of α_i and β_i for which this choice maximizes utility. Next, for each worker in each five period game, we calculate the overlap in the sets from each period. For 33.9% of the workers in Ghana and 23.9% of the workers in the UK we find that the sets overlap in all periods. For 71.2% of the workers in Ghana and 61.4% of the workers in the UK we find that the sets overlap in at least four out of the five periods. For the rest of the analysis, we focus on the group of workers with an overlap in at least four out of five periods.

This allows us to construct the potential range of cumulative distribution functions (CDFs) for α_i and β_i. Figure 1.7 shows the potential cumulative distribution function for β_i. Figures 1.7a and 1.7b show the CDFs for the values that overlap in four out of the five periods, while Figures 1.7c and 1.7d show the CDFs for the sets in period 5 only. In the Appendix 1.B.7 we show the corresponding CDFs of α_i, but as noted before, our experiment is not particularly informative on the values of this parameters.

These CDFs shows that there is substantial heterogeneity in the β_i values of the workers in both countries. From Figures 1.7a and 1.7b we can see that in Ghana for 45% of the workers the value of β_i is 0.14 or higher. The corresponding share in the United Kingdom is 67%. These workers could be typed as reciprocal: workers with a β_i of 0.14 will choose

[23] As noted by Charness & Haruvy (2002), in the gift-exchange game most of the worker's choices involve payoff distributions where the worker earns more than the employer, i.e. where the worker faces advantageous inequality. These choices will therefore be informative about β_i, but not very much about α_i.

medium or high effort if the wage is sufficiently high. For 33% of the workers in Ghana and 12% of the workers in the United Kingdom we find a β_i that is lower than 0.13. This value of β_i corresponds to a choice of low effort, even for high wages, and these workers could be typed as being more "selfish". From these figures follow that the share of workers acting reciprocally is larger in the United Kingdom than in Ghana.

However, this analysis masks strategic changes in behaviour. As discussed earlier, in the Kreps model purely self-interested agents mimic the behaviour of "social" types up until the last period and in that period the agent will lower effort. In our analysis, workers acting this way will be typed as "fair", because they are behaving reciprocally in four out of five periods, while in reality they might not have reciprocal preferences, but are just acting out of their pure self-interest. Only in period 5 their "true" social preferences are revealed. Figures 1.7c and 1.7d are based on period 5 behaviour, and we see that the differences between the UK and Ghana disappear. In fact, the pattern is reversed and the share of workers acting reciprocally is higher in Ghana than in the United Kingdom: 44% of the workers in Ghana can be classified as having a β_i parameter of 0.14 or higher and 17% as having a β_i of 0.13 or lower. For the UK, these figures are respectively 36% and 23%.[24]

Maximum likelihood As an alternative approach to estimating α_i and β_i we can use a random utility model with multinomial logit choice probabilities and estimate α_i and β_i using maximum likelihood. The multinomial logit choice probabilities are based on the four options that a worker has: rejecting an offer or choosing low, medium or high effort. Table 1.15 shows the maximum likelihood population estimates of α and β for both countries. While the value of α is not significantly different from zero (at the 5% level), we find that the estimate of β is significantly larger than zero, at 0.147 in Ghana and at 0.213 in the UK. The estimate for the UK is significantly larger than in Ghana, even though the difference is smaller in period 5.

We repeat this estimation, but now for each worker individually, using the observations from all five periods of the game.[25] Figure 1.8 shows the distribution of β_i estimated using this procedure, for both Ghana and the UK. Just like in the CDF estimated by our

[24]However, in both countries there is a substantial group of workers for whom we cannot bound their β_i parameter. For this group, either their behaviour is not captured by the Fehr-Schmidt model or the employer's wage offer does not allow for the identification of β_i. For example, for a low wage, choosing low effort is consistent with low values of β_i as well as very high values of β_i.

[25]Unfortunately, our data is not rich enough to estimate β_i for a particular worker in a particular period. We can therefore not detect changes in behaviour, as is possible using the set identification method.

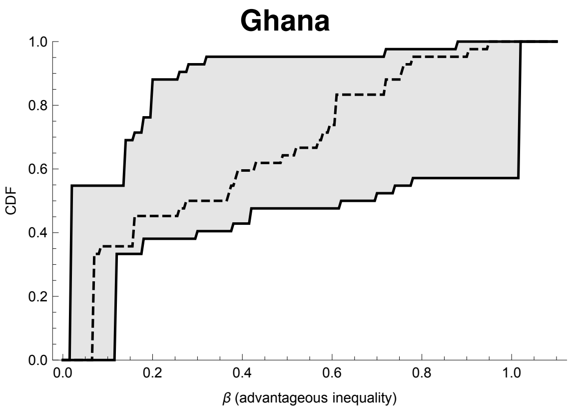

(a) All periods

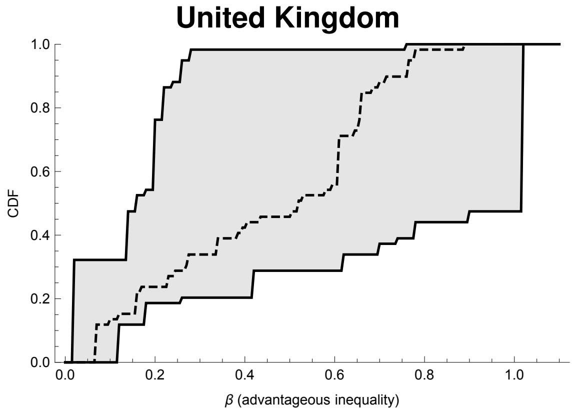

(b) All periods

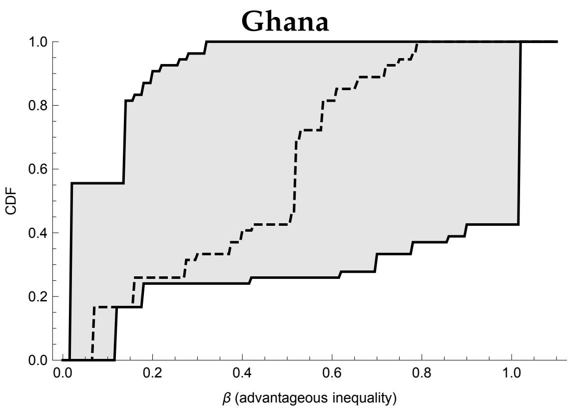

(c) Period 5 only

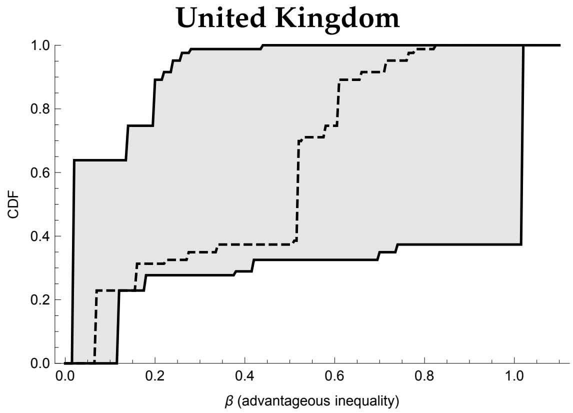

(d) Period 5

Figure 1.7: Cumulative distribution of β_i in treatment (1E). The dashed line represents the cumulative distribution of the β_i value of the geometric centre (i.e., the two-dimensional average) of each choice set.

Table 1.15: Maximum likelihood population estimates of α and β of the Fehr-Schmidt model in Ghana and the United Kingdom

	Ghana			**United Kingdom**			**Both countries (pooled)**		
	(1) All periods	(2) Period 1	(3) Period 5	(4) All	(5) Period 1	(6) Period 5	(7) All periods	(8) Period 1	(9) Period 5
α (disadvantageous inequality aversion)	-0.0291^{*} (0.0171)	-0.0195 (0.0285)	-0.0393 (0.0296)	0.0238 (0.0165)	-0.00935 (0.0303)	0.0263 (0.0194)	-0.0291^{*} (0.0163)	-0.0195 (0.0271)	-0.0393 (0.0281)
UK?							0.0529^{**} (0.0230)	0.0101 (0.0403)	0.0656^{*} (0.0340)
β (advantageous inequality aversion)	0.147^{***} (0.00777)	0.158^{***} (0.0106)	0.141^{***} (0.0117)	0.213^{***} (0.00886)	0.208^{***} (0.0125)	0.181^{***} (0.0125)	0.147^{***} (0.00738)	0.158^{***} (0.0101)	0.141^{***} (0.0111)
UK?							0.0663^{***} (0.0114)	0.0497^{***} (0.0159)	0.0399^{**} (0.0166)
Observations	562	115	108	763	155	153	1325	270	261

$^{*}\ p < 0.1$, $^{**}\ p < 0.05$, $^{***}\ p < 0.01$

Note: The parameters have been estimated using a random utility multinomial logit model, using all observations of the entire subject pool. Positive numbers indicate inequality aversion.

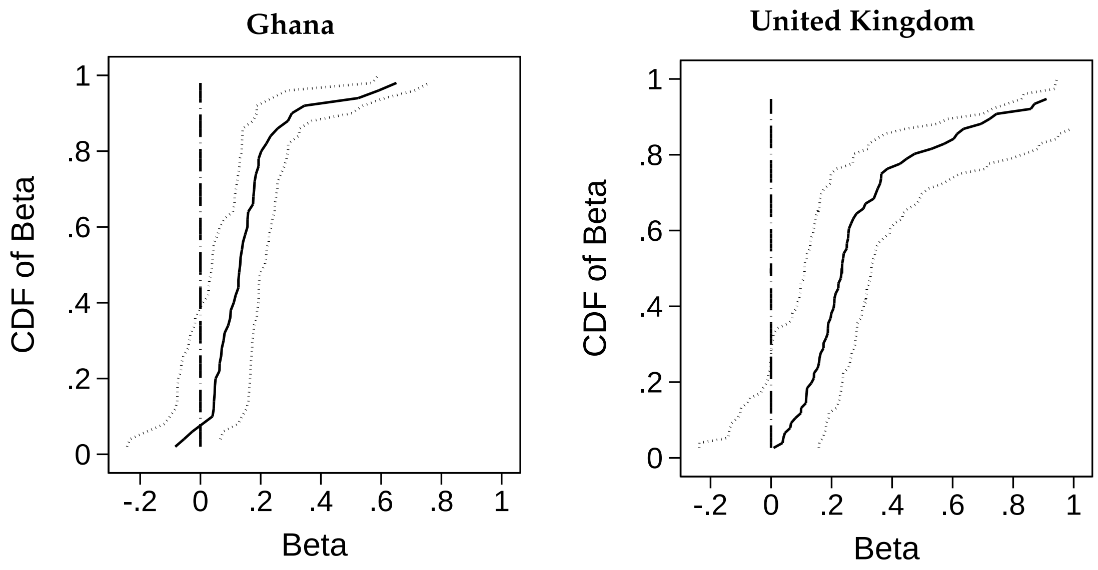

Figure 1.8: Cumulative distribution function of the estimated β_i using maximum likelihood estimation, based on treatment (1E).

set identification, the estimated β_i values tend to be higher in the UK than in Ghana. Nevertheless, in both countries we find a group of workers with a positive estimate of β_i that is significantly different from zero. Therefore, a conclusion that the workers are purely self-interested across the board is unwarranted.

Response of employers We can use the classification of workers with a low value of β_i and workers with a high value of β_i to look at the contracting behaviour of employers to these types across time. Figure 1.9 shows the wage offered and the earnings of the employers, grouped by whether we can classify them as having a β_i below 0.13 ("selfish") or a β_i above 0.13 ("fair"). We see that in the United Kingdom employers facing with a high β_i type lower their wage offers from period 3 onwards, while in Ghana the wage offers to these types remain high in all periods. Another difference can be seen in the bottom two graphs, which show the employer's payoff. In the United Kingdom employers facing low β_i types increase their payoffs from period 3 onwards, while Ghanaian employers keep making losses.[26]

[26] In these graphs, reverse causality is a potential issue: low wage offers by employers might lead to more selfish behaviour by the worker. We can find some evidence for this in the United Kingdom, where the low β_i workers faced lower wages in the first period. Nevertheless, if we redo the analysis where we base the classification of β_i on the behaviour in game 2 and look at the response of the employer in game 3 or 4, we see very similar patterns.

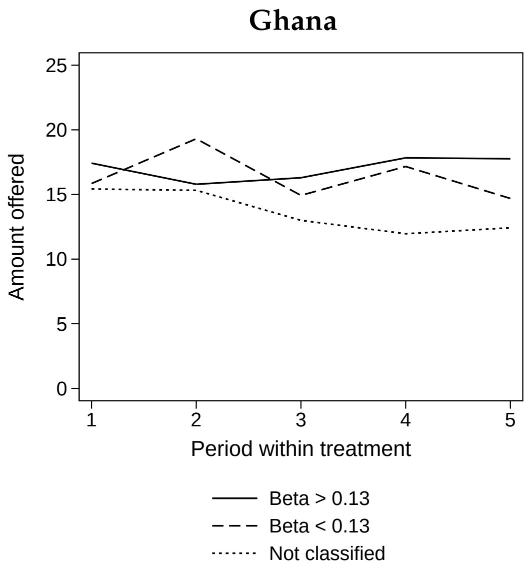

(a) Amount offered

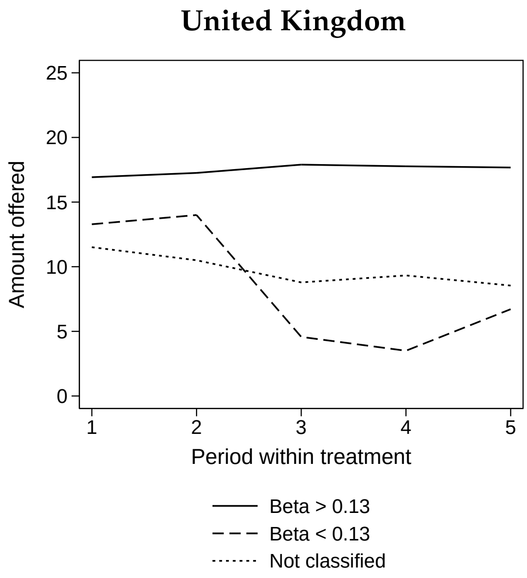

(b) Amount offered

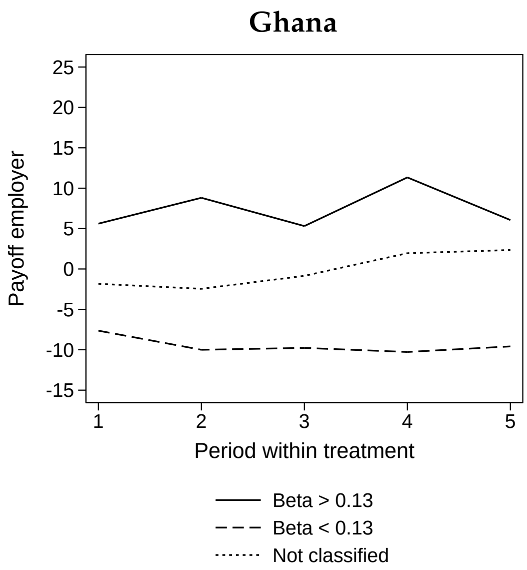

(c) Per period employer's earnings

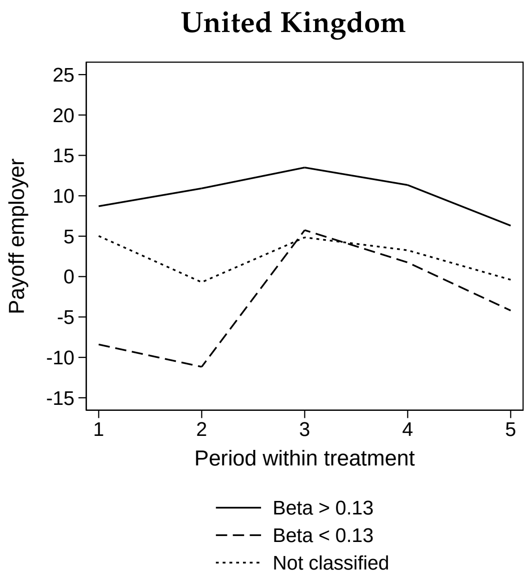

(d) Per period employer's earnings

Figure 1.9: The amount offered by the employer and the per period employer's earnings when facing the "fair" type ($\beta_i \geq 0.13$) and the "selfish" type ($\beta_i < 0.13$). *Not classified* are workers for whom the potential range of β_i covers both values below 0.13 and above 0.13 or worker for whom we could not find an overlapping set for α_i and β_i in four out of five periods.

1.5.3 Concluding remarks

Earlier experiments with gift-exchange games in developed countries have shown strong support for cooperating behaviour and relational contracting, based on reciprocity. Our results in the United Kingdom support this. However, in our Ghanaian sessions we find low levels of effort and low levels of cooperation, especially on the side of the worker. A substantial group of workers chooses low effort, even after receiving "fair" offers, where a large share of surplus is given to the worker. We find a substantial degree of worker heterogeneity. We find little evidence for punishment behaviour on behalf of the employers: while their economic loss of a worker choosing low effort instead of high effort is 35 points, there is no significant decrease in the wage they offer to a worker choosing low effort in the next period. On average, despite low compliance, employers keep offering high wages, proposing an equal or better split of surplus.

Even though the low worker productivity found in our experiment in Ghana is in line with firm surveys in developing countries, the lack of punishment on behalf of the employers in our Ghanaian experiments is surprising. It differs from the behaviour found in the UK and is at odds with both predictions of prominent theoretical models of cooperation and other earlier experiments done in developed countries: just like in our UK treatments, in Brown et al. (2004) and Brown et al. (2012) employers punish poorly performing workers, even in a situation of a shortage of labour.

We believe that our results are not the result of a lack of understanding by the participants. In the game, we tried to make the payoffs under the different effort levels as salient as possible, by explicitly showing the possible payoffs under different levels of effort while making the offer. Also, in the rehiring stage we explicitly ask the employers to make conditional offers based on the effort chosen of the worker, encouraging the employers to think about what they would offer in the next period. Despite this encouragement, we do not find strong evidence for conditional contracting in the offer stage of the game.[27]

The next question is what our experiments capture from developing country labour markets. As noted in the introduction, the Ghanaian economy is characterized by a large informal employment sector. Recruitment for workers often happens informally, through social networks. This social dimension might mean that employers could be less likely to punish workers, since employment also serves a social function. Even though theory and some studies suggests that recruitment through social networks can increase effort

[27] Note that in the rehiring stage, we do find some evidence for conditional reciprocity in some treatments. However, the difference between offers after low and high effort provision is small compared to the losses faced by the employer. See also appendix section 1.B.6.

(e.g., Montgomery, 1991; Bandiera et al., 2009), other studies have shown that incentives which result in workers receiving different wages can lead to opposite effects. For example, Breza et al. (2015) find that the effort of workers in India decreased after they received wages based on their past effort. This effect seems to be stronger when pay is relative and the effort of one worker imposes a negative externality on another worker (Bandiera et al., 2005). One interpretation of our result is that employers internalize this norm while playing the game.

Another interpretation has to do with status and the expected levels of transfers that people with a certain status are expected to make to others in their surroundings. Our games were explicitly framed in labour market terms. However, introducing such a frame comes with strong cultural connotations of employers providing a living to employees. Even though roles are randomly assigned, by giving somebody the identity of being an employer and somebody else the identity of a worker, these identities could affect their fairness concerns. An employer could be seen as a person who is well off, which then could reduce the reciprocity motivations of the worker.[28]

Another question is to what extent wage cuts as opposed to termination of the relationship are effective as a punishment. In our experiments, firing is not a salient option, but it can be achieved by not making an offer to a worker. In the one-to-one treatments, there is no outside option for both the employer and the worker. Not making an offer essentially means that the employer will not have a worker for that period and players might refrain from that for that reason. Making firing a more attractive or salient option could increase the incentive to exert high effort. Nevertheless, the gift-exchange game experiment in Brown et al. (2012) showed that even when there is excess demand for labour (and firing is therefore relatively unattractive), there was little difference in contracting patterns compared to treatments with an excess supply of labour.

A further question to study is that employers rely on nonmonetary incentives instead of monetary incentives. They could vent their frustration with the worker's effort in different ways, for example by cursing at the person or spreading a negative reputation in their social networks. Chapter 2 further explores the role of feedback.

Acknowledgements This chapter is an output from research funded by the UK Department for International Development (DFID) as part of the iiG, a research programme to study how to improve institutions for pro-poor growth in Africa and South-Asia. The views expressed are not necessarily those of DFID. Elwyn Davies gratefully acknowledges funding from the Economic and Social Research Council (ESRC). Many thanks to Václav Těhle for excellent research assistance,

[28] In a pilot experiment, we switched roles halfway the experiment. The group that were employers after the switch did not behave differently from the group that were employers before the switch.

and also thanks to Andrew Kerr and Denise Gray for their assistance during trial sessions of this experiment. Furthermore, we would like to thank Moses Awoonor-Williams for his role in helping out with the logistical matters, as well as our local team in the field, in particular Eric Agyekum and Bismark Owusu Tetteh Nortey. We would like to thank Johannes Abeler, Alan Beggs, Douglas Bernheim, Stefano Caria, Johannes Haushofer, Karine Marazyan, Fola Malomo, Muriel Niederle, Simon Quinn, Yogita Shamdasani, Hannah Uckat as well as participants of the Oxford Firms & Development Research Group, the Stanford Behavioral and Experimental Lunch, the CSAE-iiG Conference, the Stanford Development Workshop, the CSAE Research Workshop, the Paris DIAL conference, the World Bank Poverty and Applied Micro Seminar, NEUDC, PacDev and SEEDEC for helpful comments.

Appendices of Chapter 1

1.A Model of cooperation

1.A.1 Basic model of cooperation

For our model, we follow Kreps et al. (1982) as well as the Appendix of Brown et al. (2004). Let w be the wage offer. We assume there are two types in the population, a social type and a purely self-interested type:

Assumption 1.1. *There exist two types, a "social" type $S(\bar{w}_M, \bar{w}_H)$ and a rationally self-interested type R. The share of social types in the population is σ. The social type S-type will choose*

- *high effort if $w \geq \bar{w}_H$,*
- *medium effort if $\bar{w}_M \leq w < \bar{w}_H$,*
- *and low effort if $w < \bar{w}_M$.*

The threshold wages $\bar{w}_H$ and $\bar{w}_M$ are fixed and $\bar{w}_H > \bar{w}_M > 0$.

The rationally self-interested R-type will maximize monetary earnings across the game. The utility function of earnings in a particular period is

$$\pi_W(w, e) = w - c(e), \tag{1.4}$$

with $c(e)$ the cost of effort. We assume no discounting.

The employers are risk-neutral:

Assumption 1.2. *Employers are risk-neutral and maximize their monetary earnings, i.e. choose the wage w such that their payoff $\pi_E(w, e)$ is maximized, with*

$$\pi_E(w, e) = \Pi(e) - w. \tag{1.5}$$

Furthermore, we assume full knowledge of the relevant parameters:

Assumption 1.3. *The share social types σ and the threshold wages $\bar{w}_H$ and $\bar{w}_M$ are public knowledge.*

High effort equilibrium

Proposition 1.1. *Under certain conditions, a perfect Bayesian equilibrium exists in which the employer offers $\bar{w}_H$ in all periods, unless the worker deviates from the S-type behaviour, after which the employer will offer $w_L = 0$. The S-type worker chooses effort according to the offered wage throughout. The R-type mimics the actions of an S-type in all periods except for the final period, in which the R-type will choose low effort.*

We can use backward induction to prove the existence of such an equilibrium under certain conditions:

- In the last period, T, the R-type will choose low effort and the S-type will choose effort dependent on the offered wage, as described above. When offered a wage of zero, both types of workers accept the offer, as they are indifferent between rejection and acceptance, and both types will choose low effort.

- In the last period, the employer will choose the wage that gives the highest expected payoff. Let X_T be the history of play up until period T. The expected employer's payoff when offering w is

$$\begin{aligned} P(S|X_T)\Pi(H) + (1 - P(S|X_T))\Pi(L) - w & \quad \text{if } w \geq \bar{w}_H, \\ P(S|X_T)\Pi(M) + (1 - P(S|X_T))\Pi(L) - w & \quad \text{if } \bar{w}_M \leq w < \bar{w}_H, \\ \Pi(L) - w & \quad \text{if } w < \bar{w}_M, \\ 0 & \quad \text{if no offer is made.} \end{aligned}$$

  In the above $P_T(S|X_T)$ is the probability that the worker is a social type given history X_T.

  If the R-type mimics the social type until this period, the belief of the employer that the worker is of the S-type equals $P_T(S|X_T) = \sigma$. (For any other behaviour $P_T(S|X_T) = 0$ and the employer will offer $w_L = 0$.) The lowest wage possible to encourage high effort from the S-type is $\bar{w}_H$ and the lowest wage to encourage medium effort from the S-type is $\bar{w}_M$. The employer will offer $\bar{w}_H$ over $\bar{w}_M$ if

$$\Pi(H) - \Pi(M) \geq \frac{\bar{w}_H - \bar{w}_M}{\sigma}, \tag{1.6}$$

  and over $w_L = 0$ if

$$\Pi(H) - \Pi(L) \geq \frac{\bar{w}_H - w_L}{\sigma} = \frac{\bar{w}_H}{\sigma}. \tag{1.7}$$

  If the employer is presented with a worker that does not choose according to how an S-type chooses, $P_T(S|X_T) = 0$, and the employer will offer a wage of $w_L = 0$.

- In period $T-1$ the R-type worker will mimic the actions of the S-type if the sum of payoffs in period $T-1$ and period T of doing so outweighs the payoffs of choosing low effort. The incentive compatibility constraint for the R-type worker in period $T-1$ becomes, when $\bar{w}_H$ is offered,

$$\bar{w}_H - c(H) + \bar{w}_H - c(L) \geq \bar{w}_H - c(L) + w_L - c(L).$$

so

$$\bar{w}_H \geq c(H).^{29} \tag{1.8}$$

- Given that the R-type will mimic the S-type in period $T-1$, the earnings in this period are respectively $\Pi(H) - \bar{w}_H$, $\Pi(M) - \bar{w}_M$ and $\Pi(L)$ when offering $\bar{w}_H$, $\bar{w}_M$ and $w_L = 0$. The employer will offer $\bar{w}_H$ if

$$\Pi(H) - \Pi(M) \geq \bar{w}_H - \bar{w}_M \qquad \text{and} \qquad \Pi(H) - \Pi(L) \geq \bar{w}_H \tag{1.9}$$

Note that these conditions are satisfied when conditions (1.6) and (1.7) are satisfied.

- We can repeat this exercise for the earlier periods and show that conditions (1.8) and (1.9) need to hold in these periods as well.[30]

Provided conditions (1.6), (1.7), (1.8) and (1.9) hold, the Perfect Bayesian equilibrium as described in Proposition 1.1 holds.

Alternative Perfect Bayesian equilibria

For other conditions, other Perfect Bayesian equilibria exist, for example in which the employer offers $\bar{w}_H$ in the first periods and then offers $\bar{w}_M$ in the last period, or in which the employer offers $\bar{w}_M$ throughout.

The equilibrium in which $\bar{w}_M$ is offered in the last period, and $\bar{w}_H$ in the periods before that, can exist under the following conditions:

[29]When this condition holds, the worker will also mimic the S-type when offered $\bar{w}_M$ in this period, because if $c(H) > c(M)$,

$$\bar{w}_M - c(M) + \bar{w}_H - c(L) \geq \bar{w}_M - c(L) + w_L - c(L).$$

[30]E.g., the R-type worker in period t will choose high effort if

$$\bar{w}_H - c(H) + (T-t)\bar{w}_H - (T-t-1)c(H) - c(L) \geq \bar{w}_H - c(L) + (T-t)w_L - (T-t-1)c(L) - c(L),$$

which is equivalent to

$$\bar{w}_H \geq c(H).$$

- In the final period, the employer will offer $\bar{w}_M$ provided this yields a higher payoff than offering $\bar{w}_H$ or zero, so if

$$\Pi(H) - \Pi(M) \leq \frac{\bar{w}_H - \bar{w}_M}{\sigma} \qquad \text{and} \qquad \Pi(M) - \Pi(L) \geq \frac{\bar{w}_M}{\sigma}, \tag{1.10}$$

- In period $T-1$, the R-type worker will mimic the S-type worker and choose high effort following $\bar{w}_H$ if

$$\bar{w}_H - c(H) + \bar{w}_M - c(L) \geq \bar{w}_H - c(L) + w_L - c(L),$$

$$\bar{w}_M > c(H). \tag{1.11}$$

and choose medium effort following $\bar{w}_M$ if

$$\bar{w}_M - c(M) + \bar{w}_M - c(L) \geq \bar{w}_M - c(L) + w_L - c(L),$$

$$\bar{w}_M > c(M). \tag{1.12}$$

- In period $T-1$ $\bar{w}_H$ is the optimal wage if

$$\Pi(H) - \Pi(M) \geq \bar{w}_H - \bar{w}_M \qquad \text{and} \qquad \Pi(H) - \Pi(L) \geq \bar{w}_H \tag{1.13}$$

- For periods $t < T-1$ the R-type will choose high effort following $\bar{w}_H$ if

$$\bar{w}_H - c(H) + (T-t-1)(\bar{w}_H - c(H)) + \bar{w}_M - c(L) \geq \bar{w}_H - c(L) + (T-t)(w_L - c(L)),$$

$$\Rightarrow \qquad \bar{w}_H \geq \frac{(T-t)c(H) - \bar{w}_M}{T-t-1}, \tag{1.14}$$

which is implied by condition (1.11).[31]

Proposition 1.2. *Under conditions* (1.10), (1.11) *and* (1.13) *a Perfect Bayesian equilibrium is possible in which the employer offers $\bar{w}_H$ in all but the final period and $\bar{w}_M$ in the final period to workers that choose according to the wage-effort schedule of the S-type and zero to other workers. R-type workers will choose according to the S-type schedule up until the second to last period and choose low effort in the last period.*

[31] We can rewrite the left had side of (1.14) using condition (1.11):

$$\frac{(T-t-1)w_H + w_M}{T-t-1} > \frac{(T-t-1)w_M + w_M}{T-t-1} = \frac{(T-t)w_M}{T-t-1} \geq \frac{(T-t)\,c(H)}{T-t-1}$$

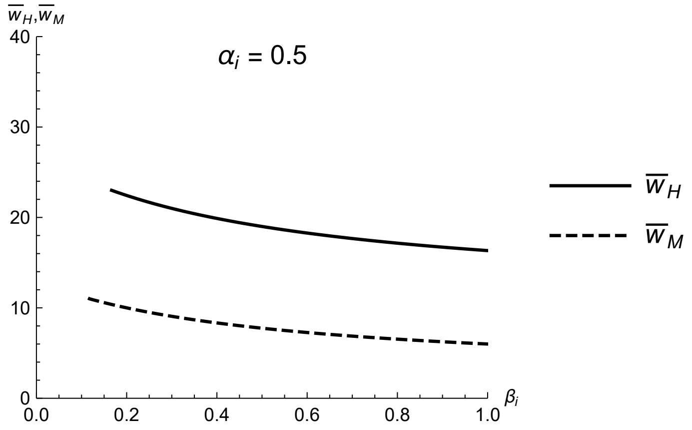

Figure 1.10: The levels of $\bar{w}_H$ and $\bar{w}_M$ as a function of β_i in the Fehr-Schmidt framework (using $\alpha_i = 0.5$ and the parameters from the game).

The main difference between the all high wage equilibrium and this equilibrium is that the expected payoff in the last period from offering $\bar{w}_H$ is not high enough, for example if there are too few S-types in the population. Offering $\bar{w}_M$ in the last period will yield a higher payoff in expectation. In the periods before that offering $\bar{w}_H$ is - in expectation - the most profitable choice, as both the R- and the S-type will reciprocate this with high effort (see condition (1.13)).

There are also equilibria possible where the employer offers $\bar{w}_M$ throughout. This will be the case when condition (1.13) is violated, and the wage differential $\bar{w}_H - \bar{w}_M$ is higher than the profit differential $\Pi(H) - \Pi(M)$. For these wage levels, even if the entire population were S-types, offering $\bar{w}_M$ would be more profitable than offering $\bar{w}_H$.

1.A.2 Microfoundations of other-regarding behaviour

The model in the previous section did not provide a microfoundation for the behaviour of the social S-type, but rather assumed that this type would choose according to a schedule. A common model to provide such a microfoundation is the Fehr-Schmidt model of inequality aversion, where the individual utility function is defined by

$$U_i(x_i, x_j) = x_i - \alpha_i \max(x_j - x_i, 0) - \beta_i \max(x_i - x_j, 0).$$

In this equation α_i represents the disutility from having a lower payoff than the other player and β_j represents the disutility of having a higher payoff than the other player. Generally, it is assumed that $\alpha_i \geq \beta_i$: a player has a higher disutility from disadvantageous inequality (represented by α_i) than from advantageous inequality (represented by β_i).

Figure 1.10 shows the relation between β_i and the threshold wages $\bar{w}_M$ and $\bar{w}_H$, taking α_i fixed at 0.5 and using the parameters from our experiment. When both α_i and β_i are fixed at 0.5, the corresponding threshold wages are $\bar{w}_M = 8$ and $\bar{w}_H = 19$.

Other classes of utility functions exist that result in similar threshold wages. For example, a model of altruism with a linear weighting of individual utility functions, e.g.,

$$U_i(x_i, x_j) = \gamma v(x_i) + (1-\gamma)v(x_j)$$

will lead to the worker choosing high effort for high wages, medium effort for medium-level wages and low effort for low wages, provided the individual value functions $v(x_i)$ are sufficiently concave.[32]

1.A.3 Screening and incentives

In the Perfect Bayesian equilibrium an employer observing low effort following a wage offer of $\bar{w}_M$ or $\bar{w}_H$ will infer from this that the worker is a R-type worker and subsequently reduce the wage to zero. This reduction of the wage following low effort provides the basis of the incentive compatibility constraint (1.8). The screening mechanism used by the employer provides the incentive for the R-type worker to exert high effort.

In the model, screening is perfect and instant: if the worker chooses low effort following a high wage, the worker is singled out as an R-type and will be offered zero wages for the remaining periods. There is no rationale for still offering a high wage following low effort. Such a rationale could be provided by models where the effort choice is not perfect. Suppose we introduce a small margin of error in the decision of the S-type.

Assumption 1.4. *With probability $\rho < 1$ each period, a S-type worker will choose low effort.*

For simplicity, assume we play the game for only two periods. Suppose that in period 1 the worker received a high wage, $\bar{w}_H$, but chose low effort. The probability that the worker is a S-type is

$$P(S|\bar{w}_H, L) = \frac{P(L|\bar{w}_H, S)P(S|\bar{w}_H)}{P(L|\bar{w}_H, S)P(S|\bar{w}_H) + P(L|\bar{w}_H, R)P(R|\bar{w}_H)} = \frac{\rho\sigma}{\rho\sigma + P(L|R, \bar{w}_H)(1-\sigma)}.$$

Now suppose the R-type always chooses low effort, so $P(L|R, \bar{w}_H) = 1$, does it make sense to offer $\bar{w}_H$ again? The employer will offer a high wage if

$$\begin{aligned} \bar{w}_H - \bar{w}_M &\leq \frac{\rho\sigma}{\rho\sigma + (1-\sigma)}(1-\rho)\left(\Pi(H) - \Pi(M)\right), \\ \text{and} \qquad \bar{w}_H &\leq \frac{\rho\sigma}{\rho\sigma + (1-\sigma)}(1-\rho)\left(\Pi(H) - \Pi(L)\right). \end{aligned}$$

[32]Given that the payoff functions for both the worker and employer are linear in w, using a linear value function instead of a convex value function will result in either always the same effort choice, regardless of wage, or in a corner solution, or in indifference between two effort levels for all wages.

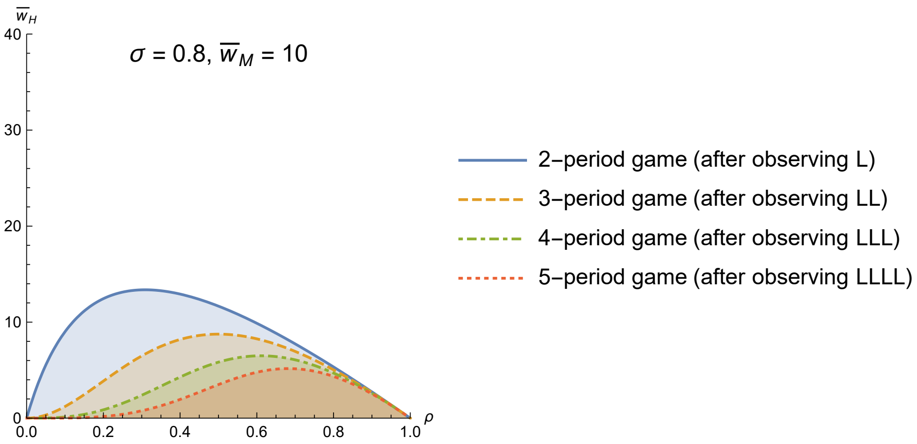

Figure 1.11: The maximum level of $\bar{w}_H$ for which it is rational for an employer to offer this wage, after observing low effort for one period (solid line), two periods (dashed line), three periods (dot-dashed line) and four periods (dotted line), as a function of ρ.

If ρ and σ are sufficiently high and $\bar{w}_H$ sufficiently low, these conditions indeed hold, and it is optimal for the employer to offer high effort following low effort. Given this, it is optimal for the R-type to choose low effort in the first period.

If we add another period to the model, and look at whether it is still rational to offer high wage following two choices of low effort, we see that the conditions become even stricter: we require a higher level for ρ and σ and a lower level for $\bar{w}_H$:

$$\begin{aligned} \bar{w}_H - \bar{w}_M &\leq \frac{\rho^2\sigma}{\rho^2\sigma + (1-\sigma)}(1-\rho)\left(\Pi(H) - \Pi(M)\right), \\ \text{and} \qquad \bar{w}_H &\leq \frac{\rho^2\sigma}{\rho^2\sigma + (1-\sigma)}(1-\rho)\left(\Pi(H) - \Pi(L)\right). \end{aligned}$$

Figure 1.11 shows these conditions graphically, with $\sigma = 0.8$ and $\bar{w}_M = 10$, for the two-period game (solid line), the three-period game (dashed line), the four-period game (dot-dashed line) and the five-period game (dotted line), using the payoff parameters of our game. From the graph it can be seen that the maximum level of $\bar{w}_H$ that supports such an equilibrium is low for low values of ρ (if an employer sees low effort, the probability is high that the worker is an R-type), then increases for medium levels of ρ and finally decreases again (for high values of ρ offering $\bar{w}_H$ becomes to risky, because the probability that S-types choose low effort is high).

The Figure shows that when $\sigma = 0.8$, $\bar{w}_M = 10$, the highest level of $\bar{w}_H$ that can sustain offering $\bar{w}_H$ following low effort in a 2-period game is 13, if $\rho = 0.3$. As discussed earlier, most other-regarding models point to values of $\bar{w}_H$ that are higher than this and therefore cannot sustain an equilibrium where high wages are offered following low effort. The equilibria where employers offer high effort following two instances of low effort require high values of ρ (i.e., the probability that an S-type makes a mistake must be high) and also $\bar{w}_H$ must be low (S-types are willing to exert high effort following a low wage).

Other equilibria exist, for example in which the employer lowers the wage from $\bar{w}_H$ to $\bar{w}_M$ following low effort. These equilibria can exist for higher levels of $\bar{w}_M$ and $\bar{w}_H$ and lower levels of ρ and σ.

1.B Further regressions and alternative specifications

1.B.1 Treatment effects

This section presents the treatment effects of the various treatments. The staggered sequence of treatments allows us to use both a within-subject as well as between-groups analysis. Furthermore, we can also apply difference-in-difference methods to measure the treatment effects.

We compare two treatments across two games at a time. One of the treatments functions as "control" treatment. Our notation is as follows: T_{ig} equals one if participant i receives the treatment in game g. D_i is time-invariant and equals one if participant i is in the treatment group and will receive the treatment at some point. P_g indicates the game and is equal to one for observations from the second game.

For example, if we compare treatments (1C) and (1E) across games 1 and 2, we treat (1C) as the control treatment and (1E) as the treatment of interest. The control group ($D_i = 0$) in this case consists of participants who were given treatment (1C) in both games, while the treatment group ($D_i = 1$) consists of the participants that received treatment (1C) in game 1 and treatment (1E) in game 2. T_{ig} equals one for observations from treatment (1E). P_g equals one for observations from game 2. Due to the incremental setup of our sequences (see Table 1.4), T_ig is equal to the interaction of D_i and P_g.

1. **Within-subject test.** For this test, we compare the outcome variable y of the same participant in the control in the first game and the treatment in the second game. Only participants from the treatment group are included, as they received the control in the first game and the treatment in the second game. This corresponds to running the following fixed-effect regression on the treatment group:

$$y_{igt} = \alpha_i + \beta_1 T_{ig} + \epsilon_{igt} \tag{1.15}$$

The coefficient β_1 corresponds to the treatment effect.

2. **Between groups test.** For this test, we compare the outcome of the control group and the treatment group in the same game ($g = \bar{G}$), while receiving different treatments. We ignore the observations from the first game in which both groups received the control treatment. This corresponds to the following regression:

$$y_{i\bar{G}t} = \beta_0 + \beta_1 \, T_i\bar{G} + \epsilon_{i\bar{G}t} \tag{1.16}$$

3. **Difference-in-difference.** In the difference-in-difference test we correct for the common time trend of both groups across the two games. This corresponds to the following regression:

$$y_{igt} = \beta_0 + \beta_1 \, T_{ig} + \beta_2 \, P_g + \beta_3 \, D_i + \epsilon_{igt} \tag{1.17}$$

   Here D_i is a time-invariant indicator on whether an individual is in the treatment group, P_g is an indicator for the second game. Note that in our setup, $T_{ig} = D_i \times P_g$.

4. **Fixed effects.** The fixed-effects regression is in its essence similar to the difference-in-difference regression, but exploits the full panel dimension of the data by controlling for individual characteristics. This approach has potentially more power. Just like in the difference-in-difference regression we rely on a common trend assumption for our estimation of the treatment effect. This corresponds to the following regression:

$$y_{igt} = \alpha_i + \beta_1 \, T_{ig} + \beta_2 \, P_g + \epsilon_{igt} \tag{1.18}$$

   The main difference between this equation and Equation (1.17) is that in Equation (1.17) β_0 and β_3 are estimated on the basis of the entire sample, while in this equation α_i is estimated for each individual separately.

Tables 1.16, 1.17 and 1.18 present the effects of the treatments on the offered wages, surplus, the earnings of the worker and the employer, the rate of acceptance and the rate of compliance.

Table 1.16: Treatment effects.

	Ghana				United Kingdom			
	(1) Within	(2) Between	(3) DID	(4) FE	(5) Within	(6) Between	(7) DID	(8) FE
Panel A. Offered wage								
(1E) vs. (1C)								
Game 2 / Game 1	-4.025^{**}	-4.624^{*}	-2.207	-1.891	-5.698^{***}			
	(1.526)	(2.309)	(2.027)	(2.110)	(0.862)			
Game 3 / Game 1	-7.033^{***}	-6.896^{**}	-5.383^{***}	-5.140^{***}	-4.406^{**}			
	(1.184)	(2.354)	(1.282)	(1.448)	(1.473)			
Game 4 / Game 1	-7.128^{***}	-8.235^{**}	-6.722^{*}	-6.402^{*}	-4.406^{**}			
	(1.711)	(2.873)	(3.059)	(3.274)	(1.473)			
(3C) vs. (1C)								
Game 2 / Game 1	0.452	0.444	2.030	2.585				
	(0.795)	(2.112)	(1.580)	(1.646)				
Game 3 / Game 1	-2.513^{***}	0.439	-1.309	-0.621				
	(0.570)	(4.702)	(0.632)	(0.738)				
Game 4 / Game 1	-3.097	-1.587	-3.335	-2.372				
	(1.952)	(3.429)	(2.081)	(2.163)				
(3E) vs. (1E)								
Game 3 / Game 2	-0.617	1.428	1.903	1.765	0.911	-0.712	0.585	1.190
	(0.457)	(1.836)	(1.815)	(1.350)	(0.958)	(1.460)	(1.572)	(1.643)
Game 4 / Game 2	-2.079	-1.380	0.892	0.368	1.547	-0.386	0.862	1.653
	(1.705)	(3.057)	(2.835)	(1.785)	(1.245)	(1.180)	(1.606)	(1.753)
(3E) vs. (3C)								
Game 3 / Game 2	-5.146^{***}	-5.908	-5.374^{***}	-4.752^{***}				
	(0.951)	(4.503)	(1.196)	(1.045)				
Game 4 / Game 2	-6.998^{***}	-8.028^{**}	-7.187^{**}	-6.392^{**}				
	(1.448)	(2.838)	(2.445)	(2.604)				
(3ES) vs. (3E)								
Game 4 / Game 3	-2.145^{***}	1.506	-0.306	-0.262	1.681^{**}	-0.305	3.040^{***}	3.264^{***}
	(0.653)	(2.309)	(0.902)	(0.984)	(0.573)	(1.837)	(0.926)	(0.962)
Panel B. Surplus								
(1E) vs. (1C)								
Game 2 / Game 1	-7.675^{***}	-11.96^{**}	-9.017	-9.017	-10.06^{***}			
	(1.803)	(4.397)	(6.493)	(6.835)	(0.845)			
Game 3 / Game 1	-10.57^{***}	-16.97^{***}	-14.74^{***}	-14.74^{***}	-8.087^{***}			
	(1.285)	(1.528)	(1.577)	(1.658)	(1.297)			
Game 4 / Game 1	-11.06^{***}	-18.00^{***}	-15.77^{***}	-15.77^{***}	-8.087^{***}			
	(1.098)	(1.462)	(1.639)	(1.723)	(1.297)			
(3C) vs. (1C)								
Game 2 / Game 1	4.313^{***}	1.577	2.970	2.970				
	(0.954)	(4.105)	(6.469)	(6.811)				
Game 3 / Game 1	4.387	-1.522	0.215	0.215				
	(3.874)	(1.889)	(4.810)	(5.055)				
Game 4 / Game 1	3.293^{*}	-3.158	-1.421	-1.421				
	(1.746)	(2.503)	(2.665)	(2.800)				
(3E) vs. (1E)								
Game 3 / Game 2	-1.949^{**}	2.683^{*}	1.144	-0.789	3.258^{**}	2.127	2.604	2.604
	(0.826)	(1.295)	(2.639)	(2.625)	(1.395)	(2.403)	(2.305)	(2.427)
Game 4 / Game 2	-2.544	2.324	0.189	-0.894	6.267^{***}	2.433	2.527	2.527
	(1.910)	(2.091)	(3.860)	(3.071)	(2.010)	(2.136)	(2.585)	(2.721)
(3E) vs. (3C)								
Game 3 / Game 2	-12.47^{***}	-12.76^{***}	-17.21^{***}	-15.54^{***}				
	(1.576)	(2.553)	(3.456)	(3.831)				
Game 4 / Game 2	-13.23^{***}	-12.52^{***}	-16.29^{***}	-15.21^{***}				
	(2.855)	(2.509)	(2.281)	(3.580)				
(3ES) vs. (3E)								
Game 4 / Game 3	0.283	1.214	1.022	1.022	3.718^{***}	-1.377	2.785^{*}	2.785
	(1.710)	(2.171)	(1.971)	(2.076)	(1.018)	(2.909)	(1.452)	(1.528)

$^{*}\ p < 0.1$, $^{**}\ p < 0.05$, $^{***}\ p < 0.01$

Note: Each cell corresponds to a separate regression. Standard errors are given in parentheses.

Table 1.17: Treatment effects (continued).

	Ghana				United Kingdom			
	(1) Within	(2) Between	(3) DID	(4) FE	(5) Within	(6) Between	(7) DID	(8) FE
Panel C. Employer's earnings								
(1E) vs. (1C)								
Game 2 / Game 1	−7.939***	−13.70***	−12.30***	−12.30***	−8.652***			
	(1.326)	(2.617)	(3.034)	(3.194)	(0.590)			
Game 3 / Game 1	−8.350***	−15.05***	−12.88***	−12.88***	−7.480***			
	(0.522)	(2.195)	(1.271)	(1.337)	(0.677)			
Game 4 / Game 1	−8.730***	−14.96***	−12.79***	−12.79***	−7.480***			
	(1.229)	(2.556)	(1.741)	(1.830)	(0.677)			
(3C) vs. (1C)								
Game 2 / Game 1	2.103***	−3.279	−2.254	−2.254				
	(0.652)	(2.594)	(3.162)	(3.329)				
Game 3 / Game 1	5.627	−3.537	1.098	1.098				
	(3.982)	(3.630)	(4.953)	(5.205)				
Game 4 / Game 1	4.960*	−3.732	0.903	0.903				
	(2.621)	(2.379)	(1.404)	(1.475)				
(3E) vs. (1E)								
Game 3 / Game 2	−2.703***	0.983	−2.522	−3.843	1.218	0.897	0.498	0.498
	(0.750)	(2.100)	(3.606)	(3.571)	(0.877)	(1.461)	(1.432)	(1.508)
Game 4 / Game 2	−2.256**	2.764	−2.310	−3.016	4.059***	0.855	0.0459	0.0459
	(0.765)	(2.618)	(2.819)	(2.729)	(0.890)	(1.194)	(1.712)	(1.801)
(3E) vs. (3C)								
Game 3 / Game 2	−10.76***	−10.53**	−15.71***	−14.56***				
	(1.569)	(4.028)	(2.769)	(3.234)				
Game 4 / Game 2	−10.06***	−8.459***	−13.89***	−13.19***				
	(2.387)	(1.897)	(1.552)	(2.472)				
(3ES) vs. (3E)								
Game 4 / Game 3	1.996	0.156	1.179	1.179	2.723***	−0.943	0.0712	0.0712
	(1.625)	(1.200)	(1.697)	(1.787)	(0.538)	(1.365)	(0.836)	(0.880)
Panel D. Worker's earnings								
(1E) vs. (1C)								
Game 2 / Game 1	−0.520	4.455	1.164	1.205	−2.410			
	(0.630)	(3.657)	(3.368)	(3.957)	(1.882)			
Game 3 / Game 1	−3.421	7.279	−3.051	−4.265	1.354			
	(3.698)	(5.142)	(3.342)	(3.787)	(4.015)			
Game 4 / Game 1	−2.483	4.035	−3.216	−4.779	1.354			
	(2.451)	(4.227)	(4.096)	(3.419)	(4.015)			
(3C) vs. (1C)								
Game 2 / Game 1	2.641*	5.771	4.969	4.365				
	(1.294)	(3.327)	(3.287)	(3.283)				
Game 3 / Game 1	1.321	9.840	0.428	0.478				
	(2.463)	(3.670)	(4.914)	(3.970)				
Game 4 / Game 1	0.818	7.111**	1.023	−1.478				
	(6.853)	(1.534)	(4.804)	(7.535)				
(3E) vs. (1E)								
Game 3 / Game 2	0.840	−2.620	1.602	2.036	3.015	−0.290	0.139	0.261
	(2.461)	(3.164)	(4.295)	(5.360)	(2.111)	(3.411)	(3.680)	(3.724)
Game 4 / Game 2	0.996	−4.343	−3.295	−0.329	5.821**	2.760	8.113	7.706
	(1.786)	(5.453)	(4.775)	(2.544)	(2.014)	(3.428)	(4.599)	(5.101)
(3E) vs. (3C)								
Game 3 / Game 2	−1.250	−5.181	1.963	1.042				
	(1.960)	(6.178)	(1.552)	(2.074)				
Game 4 / Game 2	−4.221***	−7.419	−2.008	−3.683				
	(0.616)	(5.339)	(5.379)	(4.536)				
(3ES) vs. (3E)								
Game 4 / Game 3	−0.278	1.072	1.651	2.543	−1.916	−2.644	2.139	0.623
	(1.349)	(5.109)	(2.208)	(2.271)	(1.346)	(1.838)	(2.681)	(2.889)

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: Each cell corresponds to a separate regression. Standard errors are given in parentheses.

Table 1.18: Treatment effects (continued).

	Ghana				United Kingdom			
	(1) Within	(2) Between	(3) DID	(4) FE	(5) Within	(6) Between	(7) DID	(8) FE
Panel E. Acceptance								
(1E) vs. (1C)								
Game 2 / Game 1	0.132***	0.144	0.185	0.176	0.0153			
	(0.0343)	(0.140)	(0.198)	(0.207)	(0.0189)			
Game 3 / Game 1	0.137*	0.0417	0.0497	0.0476	0.0376			
	(0.0686)	(0.0522)	(0.0674)	(0.0687)	(0.0358)			
Game 4 / Game 1	0.101*	0.0551	0.0631	0.0664	0.0376			
	(0.0501)	(0.0388)	(0.0735)	(0.0776)	(0.0358)			
(3C) vs. (1C)								
Game 2 / Game 1	−0.436***	−0.399**	−0.408*	−0.393				
	(0.0185)	(0.141)	(0.205)	(0.210)				
Game 3 / Game 1	−0.460***	−0.564***	−0.569**	−0.549**				
	(0.0797)	(0.0257)	(0.0823)	(0.0951)				
Game 4 / Game 1	−0.455***	−0.511**	−0.516**	−0.490**				
	(0.0453)	(0.0585)	(0.0740)	(0.0826)				
(3E) vs. (1E)								
Game 3 / Game 2	−0.533***	−0.597***	−0.539***	−0.534***	−0.361***	−0.401***	−0.379***	−0.361***
	(0.0232)	(0.0202)	(0.0354)	(0.0373)	(0.0428)	(0.0388)	(0.0529)	(0.0558)
Game 4 / Game 2	−0.517***	−0.554***	−0.492***	−0.478***	−0.406***	−0.499***	−0.510***	−0.484***
	(0.0465)	(0.0284)	(0.0546)	(0.0554)	(0.0730)	(0.0379)	(0.0580)	(0.0611)
(3E) vs. (3C)								
Game 3 / Game 2	−0.0235	0.00854	−0.0201	−0.0187				
	(0.0170)	(0.0312)	(0.0377)	(0.0399)				
Game 4 / Game 2	−0.0158	0.0122	−0.00618	−0.0171				
	(0.0207)	(0.0314)	(0.0344)	(0.0270)				
(3ES) vs. (3E)								
Game 4 / Game 3	−0.0112	−0.0125	−0.0143	−0.0144	0.0178	0.0267	0.0443	0.0409
	(0.0149)	(0.0164)	(0.0145)	(0.0160)	(0.0235)	(0.0433)	(0.0295)	(0.0313)
Panel F. Compliance								
(1E) vs. (1C)								
Game 2 / Game 1	−0.585***	−0.579***	−0.579***	−0.585***	−0.443***			
	(0.0401)	(0.0325)	(0.0325)	(0.0399)	(0.0334)			
Game 3 / Game 1	−0.590***	−0.598***	−0.598***	−0.590***	−0.413***			
	(0.0971)	(0.0884)	(0.0886)	(0.0935)	(0.0629)			
Game 4 / Game 1	−0.617***	−0.614***	−0.614***	−0.617***	−0.413***			
	(0.0919)	(0.0829)	(0.0831)	(0.0885)	(0.0629)			
(3E) vs. (1E)								
Game 3 / Game 2	−0.0353	0.0512	−0.0980	−0.113	0.0431	0.0213	−0.0473	−0.0524
	(0.0643)	(0.0958)	(0.157)	(0.168)	(0.0410)	(0.0672)	(0.0811)	(0.0830)
Game 4 / Game 2	−0.178	0.0161	−0.209	−0.232	0.186***	0.0394	−0.0167	−0.0168
	(0.145)	(0.0964)	(0.159)	(0.175)	(0.0262)	(0.0538)	(0.0682)	(0.0736)
(3E) vs. (3C)								
Game 3 / Game 2	−0.521***	−0.547***	−0.547***	−0.521***				
	(0.0632)	(0.0421)	(0.0421)	(0.0627)				
Game 4 / Game 2	−0.584***	−0.598***	−0.598***	−0.584***				
	(0.157)	(0.0729)	(0.0731)	(0.156)				
(3ES) vs. (3E)								
Game 4 / Game 3	0.0529	0.136	0.0643	0.0517	0.113**	−0.00452	−0.0310	−0.0406
	(0.0304)	(0.0813)	(0.0526)	(0.0598)	(0.0379)	(0.0406)	(0.0643)	(0.0686)

Note: Each cell corresponds to a separate regression. Standard errors are given in parentheses.

1.B.2 Excluding fixed effects

Table 1.19: Acceptance and compliance in treatment (1E), no fixed effects

	Ghana			United Kingdom			Both countries (pooled)		
	(1) Acceptance	(2) Compliance *(medium demand)*	(3) Compliance *(high demand)*	(4) Acceptance	(5) Compliance *(medium demand)*	(6) Compliance *(high demand)*	(7) Acceptance	(8) Compliance *(medium demand)*	(9) Compliance *(high demand)*
Wage	0.00688^{**}	0.0145^{**}	0.0139^{**}	0.0205^{***}	0.0371^{***}	0.0453^{***}	0.00688^{***}	0.0145^{***}	0.0139^{***}
	(0.00183)	(0.00482)	(0.00424)	(0.00215)	(0.00983)	(0.00361)	(0.00171)	(0.00449)	(0.00395)
UK							-0.244^{***}	-0.130	-0.374^{***}
							(0.0589)	(0.211)	(0.127)
Wage $\times$ UK							0.0136^{***}	0.0226^{*}	0.0314^{***}
							(0.00273)	(0.0108)	(0.00534)
Constant	0.803^{***}	0.317	-0.0512	0.559^{***}	0.187	-0.425^{***}	0.809^{***}	0.314^{***}	-0.125^{*}
	(0.0490)	(0.168)	(0.0830)	(0.0373)	(0.142)	(0.101)	(0.0384)	(0.0842)	(0.0700)
N	478	160	224	763	217	358	1241	377	582
Adjusted R^2	0.225	0.513	0.297	0.316	0.398	0.450	0.238	0.235	0.433
Fixed effects	None	None	None	None	None	None	None	None	None
Period dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes†	Yes†	Yes†

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

† Columns (7), (8) and (9) also include the interaction of the period dummies with the country dummy.

Note: This is a linear probability model regression of acceptance and compliance on the wage offered, without fixed effects (see Table 1.6 for the table with fixed effects). *Acceptance* is an indicator variable and is equal to one if the worker accepted the offer. *Compliance* is an indicator variable and is equal to one if the worker chooses the effort level specified by the employer, or a higher effort level. For the *Compliance (medium demand)* and *Compliance (high demand)* columns the regression only includes observations where medium or high effort was demanded by the employer. Standard errors are clustered by session.

Table 1.20: Linear regression of wage offers on rejection and previous compliance in treatment (1E), without fixed effects.

	Ghana		United Kingdom	
Dependent variable: Wage offer in period t	(1)	(2)	(3)	(4)
Rejection in period $t-1$?	-1.195	-1.093	1.677^{**}	1.690^{*}
	(1.448)	(1.397)	(0.712)	(0.944)
Compliance in period $t-1$?	0.789		7.558^{***}	
	(0.807)		(1.034)	
Compliance in period $t-1$? (high effort demanded)		1.799		13.01^{***}
		(2.357)		(1.000)
Compliance in period $t-1$? (medium effort demanded)		1.252		5.308^{***}
		(0.969)		(1.314)
Constant	13.82^{***}	13.12^{***}	8.602^{***}	7.572^{***}
	(2.640)	(3.521)	(1.319)	(1.021)
Observations	380	380	608	608
R^2	0.0475	0.0496	0.158	0.266
Adjusted R^2	0.0136	0.0132	0.139	0.249
Fixed effects	None	None	None	None
Period dummies	Yes	Yes	Yes	Yes

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: This is a linear regression of wage offered on previous compliance and rejection, without fixed effects. *Rejection* is an indicator variable and equal to one if the worker rejected the offer. *Compliance* is an indicator variable and equal to one if the worker chose the effort level demanded by the employer, or a higher level. Only wage offers in period 2-5 are included. Standard errors are clustered on the session level.

1.B.3 Alternative specifications

Table 1.21: Alternative specifications of Table 1.8

	Ghana		United Kingdom	
	(1) Wage$_t$	(2) Δ Wage$_t$	(3) Wage$_t$	(4) Δ Wage$_t$
Wage in period $t-1$	0.0850		0.326^{***}	
	(0.0995)		(0.0740)	
Compliance in period $t-1$?	0.254	-0.653	6.065^{***}	5.582^{***}
	(0.896)	(1.599)	(1.073)	(1.188)
Rejection in period $t-1$?	-0.767	5.198^{**}	4.751^{***}	9.855^{***}
	(1.772)	(1.612)	(1.116)	(1.522)
Constant	15.46^{***}	0.0808	6.145^{***}	-3.998^{***}
	(2.878)	(1.227)	(1.126)	(0.685)
Observations	371	371	595	595
R-squared	0.0594	0.0565	0.239	0.205
Adjusted R-sq	0.0224	0.0221	0.221	0.187
Employer fixed effects	Yes	Yes	Yes	Yes
Period dummies	Yes	Yes	Yes	Yes

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 1.22: Alternative specifications of Table 1.8

	Ghana		United Kingdom		Both countries (pooled)	
Dependent variable: Wage offer in period t	(1)	(2)	(3)	(4)	(5)	(6)
Medium effort$_{t-1}$	2.502*	3.167**	3.989***	3.072**	2.502**	3.167***
	(1.065)	(0.916)	(1.195)	(1.110)	(0.995)	(0.856)
High effort$_{t-1}$	2.951	4.231*	9.583***	7.823***	2.951	4.231**
	(1.831)	(1.791)	(1.275)	(1.508)	(1.711)	(1.673)
Positive surprise$_{t-1}$		0.887		0.270		0.887
		(1.043)		(1.985)		(0.974)
Negative surprise$_{t-1}$		1.725***		-2.132^{**}		1.725***
		(0.339)		(0.976)		(0.317)
UK $\times$ Medium effort$_{t-1}$					1.487	-0.0950
					(1.545)	(1.392)
$\times$ High effort$_{t-1}$					6.632***	3.592
					(2.126)	(2.242)
$\times$ Positive surprise$_{t-1}$						-0.616
						(2.193)
$\times$ Negative surprise$_{t-1}$						-3.857^{***}
						(1.017)
Constant	14.49***	13.34***	5.677***	7.416***	9.221***	9.796***
	(2.197)	(2.106)	(1.447)	(1.499)	(1.189)	(1.188)
Observations	416	416	619	619	1035	1035
R-squared	0.0741	0.0818	0.263	0.273	0.182	0.191
Adjusted R-sq	0.0369	0.0402	0.243	0.251	0.156	0.162
Employer fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Period dummies	Yes	Yes	Yes	Yes	Yes	Yes

Standard errors in parentheses

1.B.4 Extensive-form representation

Figure 1.12 shows an extensive form representation of the game.

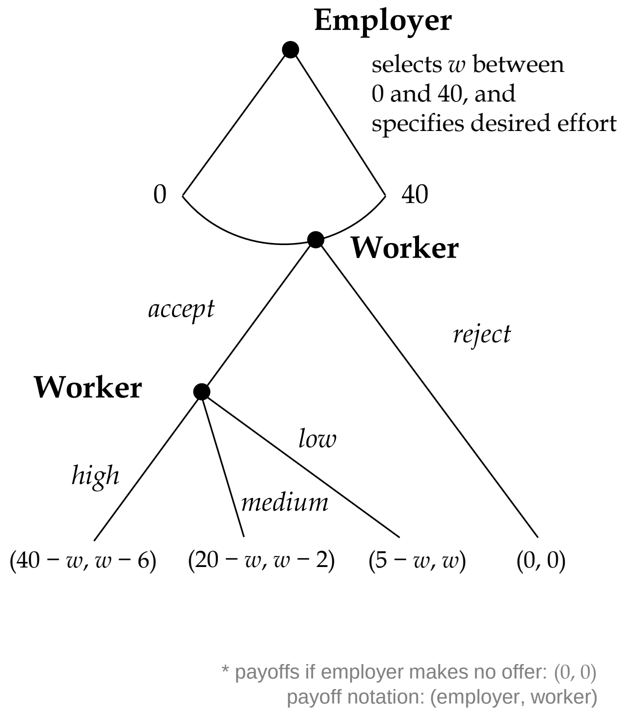

Figure 1.12: Extensive-form representation of the game

1.B.5 Fairness and effort choice

Figure 1.14 shows the effort choice corresponding to the α_i and β_i parameters of the Fehr-Schmidt model in a one-shot interaction.

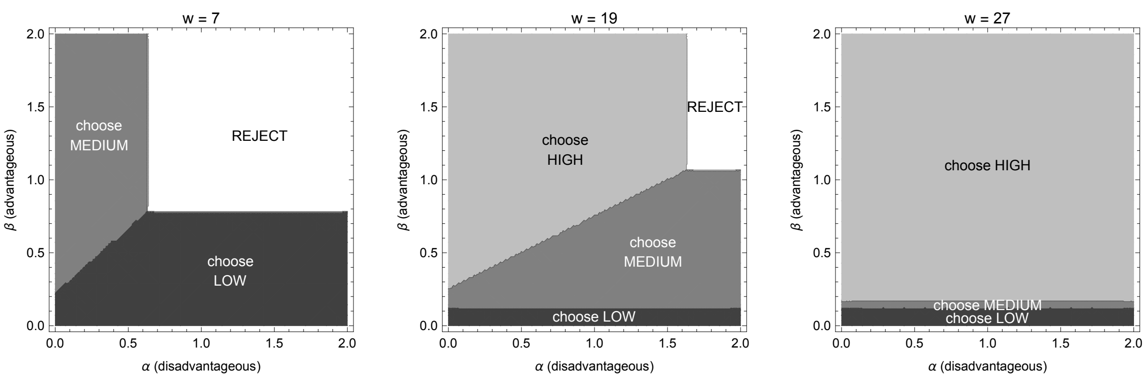

Figure 1.13: Expected behaviour by worker in the fairness model for different parameters of α_i and β_i, as response to the wage w offered by the employer. Figure 1.14 shows expected behaviour for all possible values of the wage w.

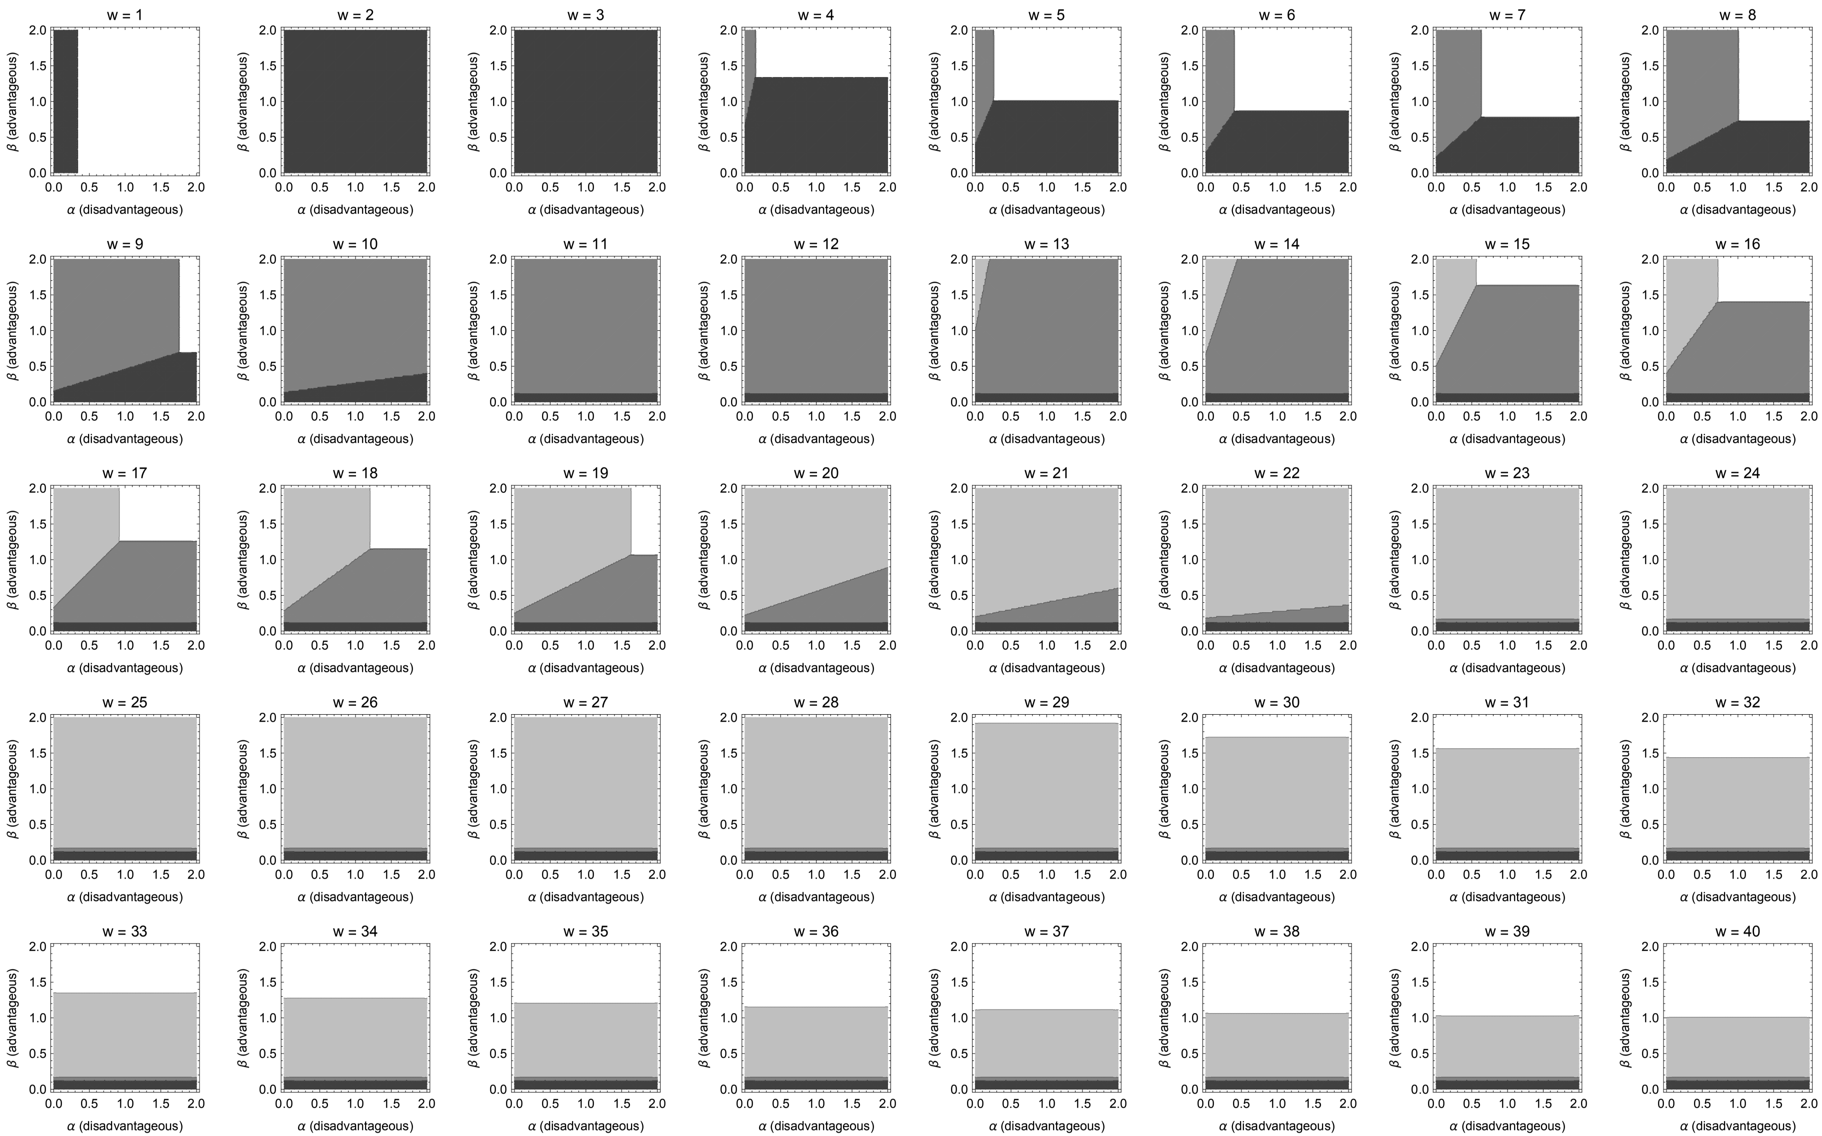

Figure 1.14: Expected behaviour by worker in the fairness model for different parameters of α_i and β_i. The shading indicate different effort levels (see also Figure 1.13.

1.B.6 Rehiring behaviour of employers

Dependent variable: Wage offered	(1) (1E)	(2) (3E)	(3) (3ES)
Medium effort exerted	4.139**	0.0258	0.650
	(1.230)	(0.343)	(1.521)
High effort exerted	6.749*	0.173	-0.640
	(2.855)	(0.320)	(1.305)
Constant	13.41***	13.38***	12.95***
	(1.750)	(0.169)	(0.626)
Observations	133	579	309

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 1.23: The amount offered by the employer to the worker in the rehiring stage in treatments (1E), (3E) and (3ES). Standard errors are clustered on the session level.

In the rehiring stage, we ask the employers for each of the possible effort levels whether they would like to make an offer to the worker again, and what the offer would be. Similarly, the workers are asked for their minimum wage in order to accept an offer. On average, 27.4% of the employers specify offers in the rehiring stage. Only a small share of trades is initiated through this rehiring: around 2-3 percent.

The Table above shows a regression of the amount offered by the employer to the worker, with the level of effort chosen by the worker as one of the regressors. In the (1E) treatment, we find that high effort is rewarded with a wage that is 6.7 points higher than a low effort offer (significant at the 10% level). A medium effort offer is rewarded with a wage that is 4.2 points higher and is significant at the 5% level. We do not find significant differences in the (3E) and (3ES) treatments.

These results show that there is some presence of conditional contacting by employers. However, the differences in wages are small and not sufficient to cover the losses incurred by low effort provision (the difference in the payoff to the employer between high and low effort provision is 35 points).

1.B.7 Classifying worker types

This appendix section describes the classification of types for the estimation of the structural model in a more detailed way. As a starting point, we assume that in every period the workers choose according to Fehr-Schmidt preferences:

$$U_i(x) = x_i - \alpha_i \max(x_j - x_i, 0) - \beta_i \max(x_i - x_j, 0)$$

Figure 1.14 shows the areas of α_i and β_i that are consistent with making a particular choice. For each worker, using these sets of values, we can take his or her choices, and plot the value sets of α_i and β_i that are consistent with this choice for the five periods in each game:

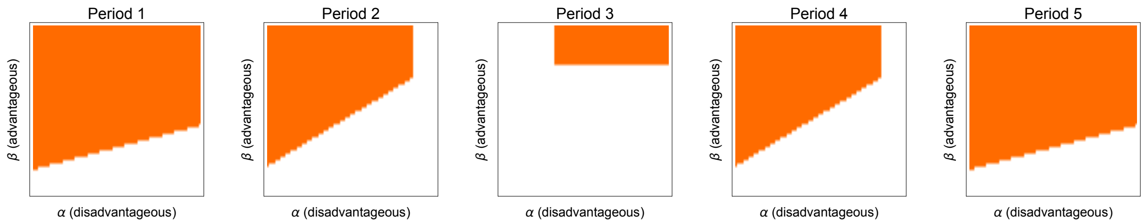

Figure 1.15: The value sets of α_i and β_i consistent with the choice made by one worker in the five periods of a game.

Next, we can overlap the areas with each other and then consider the set of values which are consistent with at least four out of the five choices made by a worker:

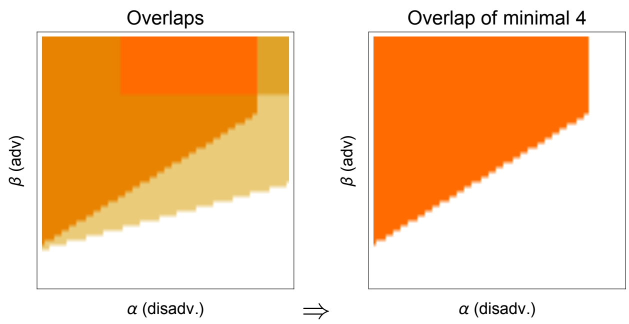

Figure 1.16: The overlapping area of α_i and β_i for one worker.

We repeat this process for each worker, and using this way, derive bounds for α_i and β_i.

1.C Instructions

Script

Below is an excerpt of the script used in the experiment for our British sessions. For our Ghanaian sessions the same script was used, with the exception of the references to the currency used for payment.

"Good morning. / Good afternoon. My name is Elwyn Davies. I am a researcher at the University of Oxford. This is [...], who will be helping out today.

Welcome to this session of the experiment. This experiment is part of a wider study done by the University of Oxford on entrepreneurship and firms. The goal is to see how people behave in a virtual marketplace. Please remain silent during the entire duration of the experiment. *Please turn off any mobile phones.*

During this session you will earn points. These points will be converted to pounds at the end of the session. *100 points is equal to 3 pounds.* You will receive this amount in private at the end of the session. You will also receive a show-up payment of 4 pounds. All earnings will be rounded up to the next 10p.

I will explain more about this experimental session later on. First, I will talk about how to use the tablet. You will use the tablet to make choices in this experiment. Please have a look at your tablet and read the message. If you are done reading, press OK.

Make sure to touch the screen gently. Do not press it too hard.

If, at any point, you have a question about the experiment, please raise your hand and we will come to your desk to answer your question in private.

(Wait.)

We are now going to practice how to make offers. In this game you will be either an employer or a worker. Employers make offers to the workers. I will talk about that more later on. First we will practice how to make an offer.

Please press the green Make Offer button. You will see a new screen in which you can make an offer. You always have to indicate two things. First you always need to select the level of effort you demand from the worker, by pressing Low, Medium or High.

Please press Low, Medium or High.

Your selected choice becomes yellow.

Now try to change your choice. For example: from High to Low, and from Low to Medium.

Second, you need to select an amount of compensation, on the right part of the screen. Please touch here, and select a number. The number has to be between 0 and 40. After selecting a number, please change the number, for example from 0 to 5, or from 32 to 37. And then change the number again.

On the bottom left of the screen you can see graphs. These graphs show you how much you can earn if the worker accepts the offer. Please change the number on the right hand side, and see how the graphs change as well. Yellow means that you are earning money, gray means that points will be deducted.

Click on OK if you want to make the offer. You can also click on Cancel, or on Delete offer, if you don't want to make an offer to the worker. After clicking on OK, you will see the offer you have made. If you want to change the offer, click on Change offer.

If you have any questions, please ask them. We will come to your desk to answer them.

(Wait until everybody is done.)

I will now talk about the experiment itself and what we are going to do.

Some of you will be employers, some of you will be workers. We will determine by chance which role you get.

Each period the employers start by making offers to the workers. Each offer must specify a level of effort and an amount of compensation. The workers then choose to accept or reject an offer. If they reject an offer, both get zero points.

If the worker accepts the offer, the worker will work for the employer. This work gives a profit to the employer, but working hard is tiring for the worker, so the worker will get less.

There are three levels of effort: high, medium and low.

- High effort means that the employer gets 40 points. The worker will lose 6 points.
- Medium effort means that the employer gets 20 points. The worker will lose 2 points.
- Low effort means that the employer gets only 5 points, but the worker will not lose anything.

We will hand out a paper to remind you about this. You can also use the graphs to help you with the calculation.

The effort level chosen is the same as what the employer demanded: the worker can only choose the effort level demanded by the employer.

So the points you get each round are as follows:

- As a worker, you get the amount offered *minus* the cost of effort, which is 6, 2 or 0 points.
- As an employer, you get the profit of 40, 20 or 5 points *minus* the payment to the worker.

Remember, the graphs will remind you of what you can earn when making an offer.

Please let us know if you have any questions about this. We will have two rounds of practice. **(Walk around and check for questions.)**

(Initialize the main game by pressing the Start button on the admin screen.)

The screen will now tell you whether you are a worker or an employer. Please press OK to continue. We will first play two rounds of practice. No points can be lost or earned.

Please press OK to continue. If you see a waiting screen, please wait. You will see this screen a couple of times during the experiment. You will have to wait until everybody is finished making their choices.

(Wait until everybody has clicked OK twice)

We are now in the first practice period. If you are an employer, please make an offer by selecting a number on the gray bar. And then press Submit. Make sure to do this before the time runs out.

(Wait until the workers can choose.)

Now the workers can choose to accept or reject the offer. If you want to accept the offer, press Select. Then click OK. For this practice round, make sure to accept the offer, so that you know how this works.

Make sure to do this before the time runs out. At the top of the screen you can see how many minutes are still remaining.

If you have accept an offer, we will ask you how many points you need in the next round to accept the offer. For example, if you select 30, you will accept all offers of 30 and higher, and reject the offers of 29 and lower.

We will also ask the employers what they would like to offer to you. If both of you agree, you will automatically accept the offer in the next period. If you don't accept, the employer will make you an offer again in the next period, just as before.

(At the end of the two practice periods.)

We will now play for real points, that will be converted to pounds at the end of the session. If you have any questions, please raise your hand and we will come to your desk to answer them.

You will have the same worker or employer for the next five periods. The letter on your screen *does not* correspond to the letter on the desk: they are different. You cannot tell who in the room your worker or employer is.

You can check your number of points by pressing the **Show history** button during the experiment. You can then also see what has happened in the previous rounds.

Press *I am ready* to continue.

(Part 2: Treatment (1E))

We have now finished the first set of five periods.

For the next part of the experiment, we are going to do the same thing. However, now the workers can choose their level of effort. They do not have to do what you demanded as an employer. So for example, if you ask for high effort, the worker can also choose medium or low effort. If you ask for low effort, the worker can also choose medium or high effort.

As an employer you will now see graphs for all the three options.

During the making of an offer, you will see bullets indicating what the worker has done in previous rounds. Gray means that the worker rejected the offer. Green means that the worker chose the level of effort that was demanded, or higher. Red means that the worker choose a lower level of effort than you demanded.

You will have a different worker than before. You will have the same worker or employer for the next five periods.

(Part 3: Treatment (3E))

For the next five periods, workers can still choose their level of effort. However, some of the employer can now make offers to three workers at the same time. Similarly, workers can accept offers from different employers. Every worker can only choose one offer. An employer can only have one worker.

Workers will take turns in choosing their offer. Sometimes you might be the first one, and have all the offers available. Sometimes you might be the last one, and there might be fewer offers available. We will determine the order in which you choose by chance.

As an employer, when you make your first offer, this will be shared with the other employers that can offer to this worker.

Again, as an employer, you can see what the workers have done in the previous periods when they were working for you. You *cannot* see what workers did when they were working for someone else.

(Part 4: Treatment (3E))

For the next five periods, you will have new workers or employers. There is no practice period.

(Part 4: Treatment (3ES))

For the next five periods, you will have new workers or employers. There is no practice period.

As an employer, you can now see what the workers have done in the previous periods, not only when they were working for you, but also what they did when they were working for someone else.

(Questionnaire)

Ok, we are now almost done with the experiment. Please fill in the questions in the questionnaire. In the meantime we will prepare your payment. We will call you when we are done with making the payments."

Chapter 2

Pledging, Praising, Shaming: Experimental Labour Markets in Ghana

with **Marcel Fafchamps**, Stanford University

Abstract Firm surveys have shown that labour management in developing countries is often problematic. Earlier experimental research (Davies & Fafchamps, 2017) has shown that managers in Ghana are reluctant to use monetary incentives to motivate workers. This chapter presents the results from a gift-exchange game experiment in Ghana in which the worker can make a promise to the employer before a contract is offered (ex ante communication) and in which the employer can send negative or positive feedback to the worker after the worker has chosen effort (ex post communication). The results indicate that feedback can help sustain cooperate behaviour (high effort provision), but only if the wage offered is high enough. Feedback reinforces reciprocity concerns on the behalf of the worker. In particular positive messages (praising) leads to higher effort provision, no significant relation between negative feedback and effort can be found. Promises are related to higher effort, but do not necessarily lead to higher wages.

2.1 Introduction

Monetary incentives, either implicit or explicit, are often seen as effective in encouraging employees to exert high effort (Prendergast, 1999; DellaVigna et al., 2016). However, it can be questioned to what extent such incentives are effective in countries where informal employment is the norm. In Chapter 1, we showed that participants in a gift-exchange game were less likely to adopt wage incentives in Ghana than in the United Kingdom. Other studies have shown that monetary incentives can lead to a decrease in morale and effort, because workers dislike the implied income inequality (e.g. Breza et al., 2015) or the negative externality they impose on others in case of relative pay incentives (Bandiera et al., 2005). In such environments, it could be more effective for employers to resort to other means of motivating their workers. A natural alternative is to motivate the worker in a different way, by for example providing feedback on their performance or to verbally agree on what level effort is put in, by for example letting the worker make a promise. Lab experiments have shown that promises ex ante (Charness & Dufwenberg, 2006; Vanberg, 2008) or providing a feedback message ex post (Masclet et al., 2003) can increase trustworthiness in a trust game.

This chapter studies the use and effectiveness of feedback and promises as nonmonetary incentives in labour relationships in Ghana. Employment in Ghana is characterized by a by a huge informal sector: only 20.2 percent of the country's working population is in wage employment (Ghana Statistical Service, 2014). In informal employment, contracts are rare and many of the employment conditions are established verbally. Even though labour surveys have shown that many formal and informal firms in Ghana offer bonuses (Barr & Serneels, 2009), differentiating bonuses on the basis of performance is seen as controversial.[1] The lack of usage of monetary incentives by employers in the Ghanaian experiments of Chapter 1 corresponds with such a story.

The raises the question if nonmonetary social incentives could function as an alternative to monetary incentives. To test for this, we implement an amended version of the gift-exchange game (Fehr et al., 1993). In one treatment, we allow the employers to send either positive or negative feedback to the worker after the worker has exerted effort. The employer can choose from three prespecified positive feedback messages and three prespecified negative feedback messages. The employer can also choose not to send feedback. In another treatment we also allow the workers to make a promise about future effort to the employer.

Our results show that allowing for promises and feedback does not increase average effort significantly. However, allowing for feedback does increase reciprocity on behalf of the workers. In the feedback treatments workers who received a high wage are more likely to exert high effort compared to the non-feedback treatments. Furthermore, we see that in particular positive messages are related to an increase in effort. Workers who have made a promise are also significantly more likely to exert high effort, provided the wage is

[1] Many firms choose to give the bonus to all workers. Debrah & Mmieh (2009) describe the problems a large formal insurance company in introducing a bonus system. Bonuses were initially given to everyone and workers felt entitled to receiving them. Making these bonuses dependent on performance led to protests from long-serving staff who were underperforming. The CEO eventually succumbed to the pressure after two years and cancelled the scheme.

high enough. This suggest that the success of communication as an incentive is very much dependent on the wage that is offered.

This paper is structured as follows. Section 2 of this chapter discusses related literature. Section 3 presents the experiment and the treatments. Section 4 discusses the implementation of the experiment in Ghana. Section 5 presents the results. Section 6 concludes.

2.2 Related literature

Our study focuses on whether positive and negative feedback as well as promises can encourage cooperation. Earlier experimental studies suggest that communication can indeed be effective in increasing trust and cooperation, but the extent of the effect is often dependent on the type of communication that the experiment allows for as well as the social distance between the players. The behavioural reasons underlying this behaviour are often explained as a preference for people to avoid guilt, a preference for consistency or concerns about social approval. However, generally, communication incentives are found to be less effective than monetary incentives. We will now discuss this in more detail.

Promises and cheap talk Even though promises are "cheap talk" and do not directly affect earnings, earlier studies found that they can increase cooperation. In trust game experiments that allowed for hidden actions, promises generally lead to increased levels of trusting and trustworthiness (Charness & Dufwenberg, 2006, 2010; Ismayilov & Potters, 2015). In these games a sender can choose whether she puts the responder in a position in which the responder can make a choice that is either very beneficial or costly for the sender. With some probability the reverse choice of the responder is implemented, giving the responder some level of plausible deniability when choosing the harmful option. Promises generally improve both the level of trust and the level of trustworthiness. Charness & Dufwenberg (2006) gave responders the possibility to write a non-binding message on a piece of paper before the game. This made senders more likely to trust the responder and also responders more likely to reciprocate.

We can think of several reasons why responders might prefer to keep to their promises. The first one, as argued by Charness & Dufwenberg, is people keep promises due to an aversion to the experience of guilt. They show that making a promise shifts beliefs on the other side and therefore raises expectations of what choice the responder makes. Breaking these expectations leads to a feeling of guilt that people want to avoid by not breaking their promises. Another view for the effectiveness of promises is a desire of consistency and that making a promise commits you to keeping it. Vanberg (2008) randomly swapped some of the participants in a two-sided dictator game, putting them in the shoes of someone else who had exchanged messages with another participants. Even though the other party knew that a swap was possible, they were kept in the dark whether a swap was actually made and therefore beliefs and expectations were not changed. Nevertheless, participants predominantly only kept to promises they made themselves and not to promises made by the person in whose shoes they were put. They were mainly reluctant to break expectations that they set themselves, but not the ones set by others. Vanberg argues that this provides evidence that people are keeping to their promises out of concern for consistency.

The effectiveness of promises are found to be dependent on whether communication is constrained. In our experiment we use pre-specified "canned" messages. Earlier experiments have shown that such "canned" messages (e.g. "I promise to choose Roll", where "Roll" corresponds to the responder's reciprocating choice) are generally less effective than free-form messages (Charness & Dufwenberg, 2010; Brandts et al., 2015). Furthermore, other experiments have shown that trustworthy behaviour just does not follow from sending the message, but that it is important that the promise is received at the other side. Ismayilov & Potters (2015) show that receivers behave less trustworthy if they made a promise that was undelivered than if that promise was delivered. This could be interpreted that people do not necessarily have a preference to for consistency with respect to themselves, but they might have a preference to be perceived as consistent by the other party.

Feedback and social approval Positive and negative feedback gives an idea of whether the principal approves of the action of the agent after the choice has been made. Earlier studies of social approval find that this can increase contributions in a public goods game or effort in a gift-exchange game, as discussed below. Many of these studies test social approval by *revealing the identity of the participants* at the end of the experiment. Gächter & Fehr (1999) reveal the participants' identities and contributions after playing ten rounds of a public goods game and find that this increases contributions in the treatment where participants engaged in a group activity before the start of the experiment. Having a basic level of social familiarity with each other is key: the contributions in the treatment where identities are revealed but no group activity took place were not higher than in the treatment with full anonymity. Falk et al. (1999) reveal the identity of the participants in a gift-exchange game, by letting them play the game face-to-face and letting them openly discuss the outcomes of the game in some of the treatments. They find only a weak effect of revealing identity on effort. Letting participants play each other repeatedly but anonymously is almost as effective. This suggests that reciprocity is more important than social approval.

When feedback is public, reputational concerns start playing a role. Rege & Telle (2004) conduct a similar experiment, but make the revelation of the contribution and identity more salient, by letting the participants open the envelop with their contribution in front of the entire group immediately after making their choice. They find that even if the participants do not know each other before the start of the experiment, such a social approval incentive can increase contributions. Similar results have been found in studies in developing countries, where participants who had to reveal their contributions in front of a group significantly increased their contributions (e.g. Barr, 2001).

Another approach for testing for social approval is *allowing for feedback* after a choice is made, as we do in our experiment. Masclet et al. (2003) test the effectiveness of such a mechanism in a public goods game. In their experiment, participants can give each other "punishment points", which do not carry any monetary value. Even though these points do not have monetary consequences, they do lead to higher contributions. However, the contributions are lower than in the treatment where these punishment points have monetary consequences.

However, there could also be internal motivations to give feedback. It can function as a way of venting frustration and therefore mitigate the need to use costly punishment. Xiao & Houser (2005) showed that responders in an ultimatum game were less likely to costly reject low offers if they had the opportunity to convey their feelings through a message.

The paper closest to our setup is Cooper & Lightle (2013), who conducted a gift-exchange game with communication in the United States. In this experiment the workers could send a free-form message to the employers *after* the interaction. They found that communication increased both wages and effort. The increase in effort was mainly due to employers increasing their wages. Positive and negative responses were less effective than advice given by the worker to increase the wage. Our experiments differ from their design in three ways: first, feedback is sent by the employer and not by the worker; second, we allow for promises while their experiment only allowed for feedback (players were randomly rematched after each period, so no reputational effects could emerge); and third, our messages are pre-specified while their experiment allowed for free-form communication.

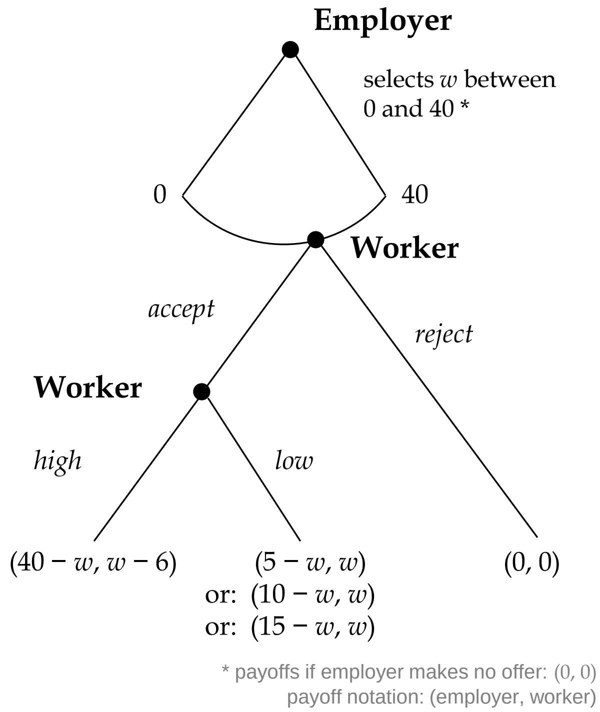

Figure 2.1: Extensive form representation of a single period in the game.

2.3 Design of the experiment

2.3.1 The gift-exchange game

The experiment is an adaptation of the gift exchange game and mimics labour relations in a principal agent setting. Participants are divided into workers and employers. Employers and workers are randomly matched to one other, in such a way that each employer is matched to one worker and each worker is matched to one employer. Each treatment consists of five or ten trading periods, in which the matched employer and worker can enter a labour contract with each other. The worker can choose two levels of effort: high effort or low effort. A higher level of effort is more beneficial to the employer, but comes at a higher cost to the workers. For the five or ten trading periods, the same employer and worker are contracting with each other. Each trading period consists of three stages: the contracting stage, the supply stage and the rehiring stage:

The contracting stage. In this stage the employer can choose to make an offer to the worker. The worker can then choose to accept or reject the offer. If the employer does not make an offer or the worker chooses to reject the offer, both the worker and the employer receive a payoff of zero points.

Effort level	low	high
Benefit to employer	5, 10 or 15	40
Cost to worker	0	6
Surplus	5, 10 or 15	34

Table 2.1: Payoff structure. The benefit to the employer is varied between sessions.

The supply stage. In this stage, the worker first receives the payment w from the employer and then chooses either a low or a high level of effort. A high level of effort is costly to the worker, but gives the employer a higher benefit. The choice of effort is not revealed to the employer until the end of the next stage.

The rehiring stage. If the worker accepts the offer, the employer is asked what wage she would offer if the worker chose low effort and what she would do if the worker chose high effort. The worker is asked what the minimum payment needs to be to accept the offer. If the offer of the employer matches the minimum desired payment, the worker and the employer will contract again with each other in the next period, skipping the contracting stage. If there is no match, contracting in the next period will happen as before. At the end of this stage the effort choice of the worker and the final payoffs are revealed.

A choice of high effort gives the employer a payoff of 40 points, from which the wage is deducted. The cost of providing high effort for the worker is 6 points, which is deducted from the wage. In case the worker chooses low effort, the worker does not lose any points, but the employer receives a lower payoff. This payoff is varied across the sessions, and is either 5, 10 or 15 points. Figure 2.1 shows an extensive-form representation of a single period in the game, and the payoffs are summarized in Table 2.1.

2.3.2 Treatments

The starting point of our experiment is treatment (1E), in which workers have a choice between providing high and low effort. To assess the effect of limited enforceability of contracts, we introduce control treatment (1C), in which workers can only choose high effort after accepting: in other words, a worker has to comply by choosing high effort.[2] This version of the game is equivalent to the ultimatum game. In both treatment (1E) and

[2]In Brown et al. (2004) employers can ask for other levels of effort than just high effort. This possibility is removed here, just like in perfect contracting treatments of earlier gift-exchange experiments, such as Fehr & Falk (1999)

Table 2.2: The four different treatment sequences.

Sequence no.	Game 1 (5 periods)	Game 2 (5 periods)	Game 3 (5 periods)	Game 4 (10 periods)	No. of participants
1	(1C)	(1C)	(1C)	(1C)	40
2	(1C)	(1E)	(1E)	(1E)	100
3	(1C)	(1E)	(1EM)	(1EM)	100
4	(1C)	(1E)	(1EMP)	(1EMP)	80

Note: Treatment (1C) is the control treatment, in which the worker can only choose high effort. In treatments (1E), (1EM) and (1EMP) the worker has the choice of both low and high effort. In treatments (1EM) and (1EMP) the employer can send a positive or negative feedback message to the worker. In treatment (1EMP) the worker can make a conditional promise.

(1C) no communication between workers and employers can take place, apart from the offers posted.

In the third treatment, treatment (1EM), the employer can send a feedback message to the worker, after the worker's effort choice has been revealed. Three positive and three negative feedback messages are available, in varying degrees of strength. Some of the messages have a direct link with the previous action of the worker (e.g. "Thank you" or "Why did you do this to me?"), while the stronger messages allow the employer fully praise or shame the character of the worker in more general terms (e.g. "You are a very reliable person" or "I pity your children").[3] The messages are sent anonymously. The workers see the messages directly on their screens after they have been sent.

In the fourth treatment, treatment (1EMP), not only the employer, but also the worker can communicate by making a promise. This promise is made at the end of each period for the next period. The promise is predefined: "If you offer me a high payment, I will choose high effort." If a promise has been made, this is revealed to the employer. If no promise has been made, nothing is shown to the employer.

Besides this we vary the benefit of low effort provision to the employer, as an exogenous way of stimulating employers to make more varied wage offers. This benefit is either 5, 10 or 15 points (see also Table 2.1).

Each session consists of four subsequent treatments. The first three treatments consist of five trading periods, while the last treatment consists of ten trading periods. The exact treatment sequences can be found in table 2.2. After each treatment, the workers and employers are rematched randomly for the next five or ten periods.[4]

[3] The messages available were selected from a wider range of messages to suit the Ghanaian context. They were set up in consultation with our local team, and were piloted in early sessions of the experiment.

[4] In the five or ten periods of the treatment, no rematching takes place: the worker keeps facing the same employer and vice versa. Random rematching also takes place for the practice rounds.

2.3.3 Predictions

Both feedback messages and promises are "cheap talk" in the sense that they do not directly influence the participant's monetary earnings. However, if we assume that participants have *preferences over processes* and not only care about the monetary payoffs but also about the way these payoffs are achieved, feedback and promises can influence behaviour. Guilt aversion or a preference for consistency deters workers who have made a promise from breaking that promise. We therefore anticipate that the possibility for promises increases conditional cooperation and reciprocity: workers who receive a high wage will be more likely to exert high effort when they have made a promise. Feedback can increase effort levels if workers care about the approval or disapproval by the employer.

Other potential channels could be that allowing some degree for communication makes the transaction less "impersonal". The feedback messages can make it more salient that the participants are playing with another human being. This realization can in itself increase reciprocity.[5] Furthermore, the text of the promise ("If you offer me a high payment, I will choose high effort.") strongly suggests conditional play by the workers and might prime some participants into playing conditionally.

2.4 Implementation

For this experiment, we invited 320 students and entrepreneurs from Accra, Ghana, to participate. Students were recruited from the main universities and polytechnic schools in Accra, and the entrepreneurs were recruited from participants of firm surveys conducted earlier. We used *LabBox*, a custom-made platform developed by one of the co-authors. Participants were given Android tablets running a custom-made HTML/JavaScript platform as participant terminals. A PHP web server was used to facilitate the communication between the participant terminals. All participants were able to read and write and extensive instructions were given to make sure that participants were comfortable with using the technology. The experiment was conducted in English, which is the main language of instruction at tertiary institutions in Ghana and often used in business interactions, especially within larger firms. Each session lasted for 1.5 to 2 hours and the average payoff was 30-35 cedis (about 10 dollars).

While making offers, employers are presented graphs indicating the potential consequences of their offers, showing the amount of points they and their worker would earn

[5] It should be noted that a large share of the subjects were unfamiliar with lab experiments. A way of testing this channel would be to allow for communication unrelated to the transaction in the game.

Table 2.3: The average wage, the share accepted, compliance and the average earnings in the various treatments.

Treatment		Mean wage	Share accepted	Mean compliance	Mean employer's payoff	Mean worker's payoff
(1C)	Game 1	21.1	83.4%	100.0%	18.0	16.0
(1C)	Game 2	23.7	87.0%	100.0%	15.3	18.7
(1C)	Game 3	24.7	94.0%	100.0%	15.1	18.9
(1C)	Game 4	24.0	97.0%	100.0%	15.8	18.2
(1E)	Game 2	15.9	88.5%	36.1%	4.1	14.5
(1E)	Game 3	15.3	85.6%	32.7%	3.4	14.4
(1E)	Game 4	15.9	85.6%	34.1%	3.2	15.0
(1EM)	Game 3	15.5	87.6%	42.9%	6.6	13.7
(1EM)	Game 4	13.3	88.0%	38.0%	7.2	11.9
(1EMP)	Game 3	17.3	95.0%	33.2%	2.2	15.7
(1EMP)	Game 4	15.2	90.8%	35.5%	4.7	13.8

Note: The above figures are averaged over all five or ten periods in each game. *Compliance* is defined as whether a worker chose high effort. For the multilateral treatments, *Share accepted* indicates the share of offers that were taken up.

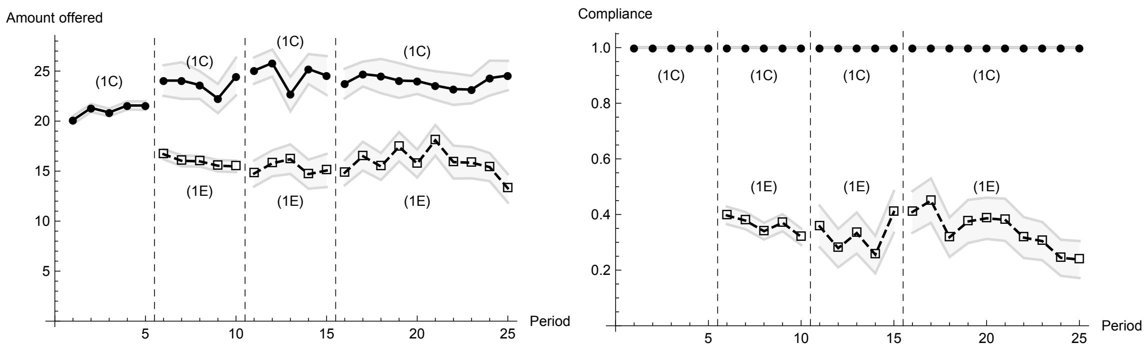

(a) Wage offers

(b) Compliance (high effort choice)

Figure 2.2: Average wage offers and compliance (i.e., high effort choice) for treatment (1C) and treatment (1E)

in case of low and high effort choice by the worker. Similarly, workers are shown graphs with the consequences for the payoffs when making the choice between high or low effort. The screens and exact instructions can be found in the Appendix.

2.5 Results

2.5.1 Wages and effort when communication is not allowed

Just like earlier gift-exchange game experiments, we find that employers offer higher than minimum wages and a positive relationship between wages and effort. Table 2.3 shows the average wage, the share of accepted offers, the mean compliance with wages and the average employer's and worker's payoff. Compliance is defined whether a worker chooses high effort. We can see that both wages and the share of offers that are accepted drop

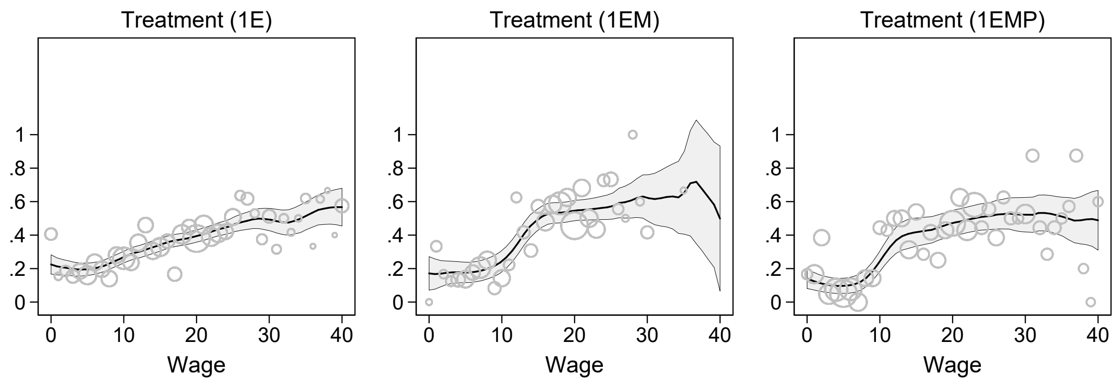

Figure 2.3: The relationship between wage and effort in treatment (1E), (1EM) and (1EMP). The dependent variable is share of compliance (i.e., high effort choice).

when enforcement is no longer perfect. In games 2, 3 and 4 of treatment (1E) the workers choose high effort in 36.1%, 32.7% and 34.1% of the cases, which is significantly lower than the perfect compliance we see in treatment (1C). However, as in previous gift-exchange games, there is a pattern of positive reciprocity: effort is higher when wages are higher. Figure 2.2 shows the average amount offered and the average share of workers choosing high effort in treatment (1C) and (1E), breaking these figures down by period. The lower level of compliance at the end of treatment (1E) suggests that there are some end game effects.

Figure 2.3 shows the relationship between wage and high effort. Nevertheless, in treatment (1E), even for high wages, effort provision is low. For example, in treatment (1E), for offers higher than 23 points the average compliance rate is 40.4%. These offers can be seen as generous, because even when choosing high effort, the worker receives at least half of the surplus. This results in a unequal distribution of surplus: the workers capture most of the surplus. Making fair offers close to 23 points (which corresponds to an equal sharing of surplus if the worker chooses high effort) is costly for employers: offers between 20 and 25 points lead to an average loss of 4.7 points, while in the full compliance case this led to a profit of 18.8 points. We see that the introduction of effort choice is costly for both the employer, whose average income drops from 15.7 points to only 3.0, and for the worker, whose average income drops from 18.2 points to 15.0 points (see Table 2.3 for a breakdown by game). A t-test shows that the differences are highly significant ($p = 0.000$ in both cases). In this case, limited enforcement hurts both parties.

Table 2.4: Feedback messages sent in treatments (1EM) and (1EMP)

Feedback	**Low effort**	**High effort**
No feedback	20.7%	9.1%
Negative feedback		
"Why did you do this to me?"	30.2%	3.0%
"I pity your children"	6.5%	0.0%
"You are a disgrace to your family"	10.7%	0.7%
Positive feedback		
"Thank you"	22.9%	49.0%
"You are hardworking"	5.7%	17.6%
"You are a reliable person"	3.4%	20.6%
Total	100.0%	100.0%

Note: This table pools the observations from games 3 and 4 of treatment (1EM) and treatment (1EMP).

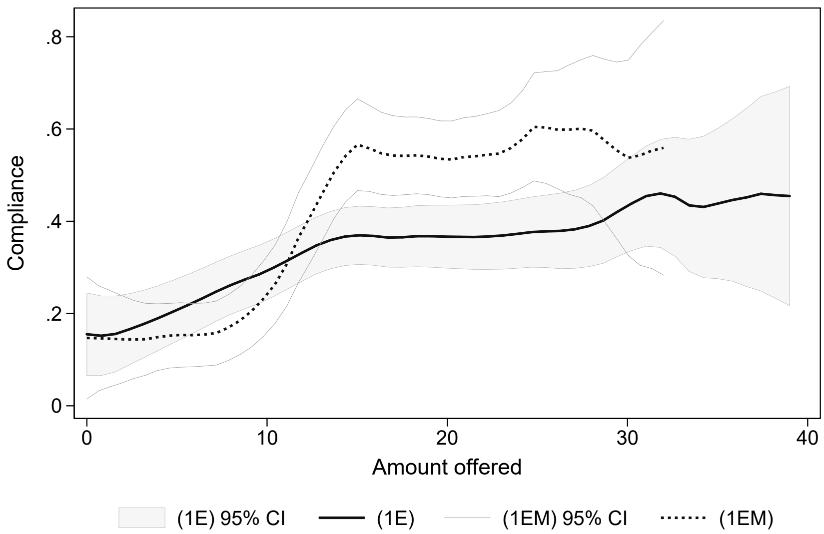

Figure 2.4: A local polynomial regression of compliance as a function of the wage offered. Compliance is defined as choosing high effort after accepting an offer. The figure compares the treatment with messages (1EM) with the treatment without messages (1E).

Table 2.5: Share of high effort choice in treatments (1C), (1E), (1EM) and (1EMP)

	Game 3						Game 4										
Period	1	2	3	4	5	All	1	2	3	4	5	6	7	8	9	10	All
Panel A. All offers																	
High effort choice in																	
... treatment (1C)	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
... treatment (1E)	0.36	0.28	0.33	0.26	0.41	0.33	0.41	0.45	0.32	0.38	0.39	0.38	0.32	0.30	0.24	0.24	0.34
... treatment (1EM)	0.47	0.42	0.47	0.47	0.49	0.43	0.44	0.44	0.26	0.40	0.42	0.43	0.39	0.44	0.43	0.37	0.38
... treatment (1EMP)	0.34	0.36	0.31	0.34	0.31	0.33	0.43	0.45	0.29	0.47	0.36	0.43	0.26	0.34	0.30	0.21	0.35
Panel B. Above-median wages ($w_t \geq 16$)																	
High effort choice in																	
... treatment (1C)	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
... treatment (1E)	0.33	0.27	0.33	0.26	0.53	0.34	0.39	0.50	0.45	0.37	0.40	0.39	0.50	0.44	0.26	0.31	0.38
... treatment (1EM)	0.58	0.46	0.63	0.58	0.65	0.58	0.48	0.70	0.33	0.68	0.52	0.65	0.50	0.56	0.60	0.50	0.56
... treatment (1EMP)	0.39	0.42	0.38	0.50	0.45	0.43	0.56	0.67	0.42	0.63	0.59	0.64	0.33	0.60	0.50	0.21	0.46
Panel C. Treatment (1EM)																	
% Sent positive feedback	0.55	0.53	0.55	0.45	0.50	0.52	0.43	0.55	0.33	0.50	0.43	0.50	0.45	0.50	0.53	0.53	0.47
High effort after positive msg	-	0.52	0.71	0.57	0.67	0.62	-	0.65	0.42	0.75	0.63	0.56	0.58	0.71	0.60	0.58	0.61
% Sent negative feedback	0.18	0.23	0.30	0.33	0.25	0.26	0.28	0.23	0.40	0.23	0.38	0.33	0.30	0.28	0.28	0.28	0.30
High effort after negative msg	-	0.33	0.11	0.17	0.46	0.27	-	0.20	0.00	0.29	0.22	0.36	0.18	0.10	0.20	0.18	0.19
Panel D. Treatment (1EMP)																	
% Made promises	0.50	0.45	0.43	0.73	0.68	0.56	0.68	0.68	0.63	0.73	0.60	0.73	0.65	0.70	0.63	0.63	0.66
High effort after promise	-	0.37	0.29	0.44	0.32	0.36		0.54	0.32	0.57	0.39	0.46	0.32	0.36	0.30	0.22	0.39
High effort after no promise	-	0.35	0.32	0.27	0.25	0.30	-	0.25	0.23	0.31	0.30	0.40	0.10	0.31	0.30	0.20	0.27

Note: The *All* column shows the average choice of high effort or the average share of workers or employers sending feedback or making a promise.

2.5.2 Positive and negative feedback

In treatments (1EM) and (1EMP) employers are given the possibility to send out feedback. They have a choice out of six pre-specified messages. In a majority of the cases employers make use of this opportunity. Table 2.4 shows the feedback messages sent by the employer following low and high effort. High effort is followed by positive feedback in 87.2% of the cases, negative feedback in 3.7% of the cases and no feedback in 9.1% of the cases. This suggests that high effort generally is rewarded with a positive message. After low effort employers send negative feedback in 47.4% of the cases, positive feedback in 32.0% of the cases and no feedback in 20.7% of the cases. The high share of positive feedback following low effort is remarkable, but this could indicate a sarcastic or ironic response.

The average wages in the feedback treatments are similar to the other treatments. In treatment (1EM), the average wage offer is 15.5 and 13.3 points in games 3 and 4, while in treatment (1E), the average wage offers in games 2, 3 and 4 are 15.9, 15.3 and 15.9 points. These differences are not significant. The acceptance rates are very similar in the two treatments and not different significantly from each other.[6]

[6] When averaging across all games, the average in treatment (1E) is 85.6% and in treatment (1EM) is 88.0%. A *t*-test of the differences yields a p-value of 0.263.

The rate of compliance is higher in treatment (1EM) than in treatment (1E). The share of high effort choices is 36.1%, 32.7% and 34.1% in games 2, 3 and 4 of treatment (1E) and 42.9% and 38.0% in games 3 and 4 of treatment (1EM). However, these differences are not significant.[7]

Figure 2.4 compares the relationship between compliance and wage for both treatments (1E) and (1EM). As can be seen in the figure, the relationship between wages and effort is *stronger* in treatment (1EM) than in treatment (1E): for low wage offers, high effort is less likely, but for higher wage offers, high effort has become more likely. This differential response could explain the lack of a significant overall effect of feedback. The decrease in high effort provision at low wages partially cancels the increase in high effort provision at high wages.

If we consider wages below and above the median (16 points) separately, we see that, compared to treatment (1E), high effort is about 4-7% less likely for below-median wages in treatment (1EM) than in treatment (1E), even though the differences are not significant.[8] However, for wages at or above the median we find a significant difference. High effort is 23.8% more likely in game 3 of treatment (1EM) compared to game 3 of treatment (1E). This difference is statistically significant at the 5% level ($p = 0.024$). The corresponding difference for game 4 is 15.0%, but this figure is statistically insignificant ($p = 0.151$).

Result 2.1. *Feedback leads to a stronger relationship between wages and effort. In game 3, for above-median wages, high effort is 23.8% more likely.*

These results suggest that reciprocity on the worker's side becomes stronger when feedback is possible. However, the lower difference for game 4 suggests that this effect wears off and decreases over time.

Result 2.2. *The possibility of feedback does not lead to significantly lower or higher offers. Positive messages are significantly related to higher wage offers in the next period.*

We now ask the question whether this is mainly driven by positive or by negative feedback. Table 2.5 shows that compliance is generally higher following positive feedback. However, this is not surprising, because positive feedback generally follows after a worker chose high effort in the previous round and these workers might be more likely to choose

[7] A t-test yields $p = 0.112$ when comparing the game 3 values, $p = 0.451$ when comparing the game 4 values.

[8] For game 3 high effort is 7.2% less likely for below-median wages in treatment (1EM), while for game 4 high effort is 4.2% less likely. However, these differences are not significant.

Table 2.6: Linear probability model of effort

Dependent variable: Choice of high effort in period t	(1)	(2)	(3)	(4)
Feedback possible?	0.0726	-0.0407	0.0194	-0.0194
	(0.0716)	(0.0630)	(0.0787)	(0.0792)
Negative feedback sent?$_{t-1}$			0.0242	-0.00892
			(0.0375)	(0.0635)
Positive feedback sent?$_{t-1}$			0.0802^{**}	-0.0661
			(0.0363)	(0.0394)
Promise possible?	-0.0459	-0.0525	-0.0856	-0.0239
	(0.0757)	(0.0560)	(0.0841)	(0.0772)
Promise made?$_{t-1}$			0.102^{***}	-0.000969
			(0.0158)	(0.0211)
Wage$_t$	0.00964^{***}	0.00679^{***}	0.00943^{***}	0.00653^{***}
	(0.00115)	(0.00111)	(0.00120)	(0.00115)
$\times$ Feedback possible?		0.00684^{***}		0.00189
		(0.00220)		(0.00327)
$\times$ Negative feedback sent?$_{t-1}$				0.00251
				(0.00341)
$\times$ Positive feedback sent?$_{t-1}$				0.00918^{***}
				(0.00308)
$\times$ Promise possible?		0.000310		-0.00366
		(0.00244)		(0.00283)
$\times$ Promise made?$_{t-1}$				0.00560^{***}
				(0.00124)
Constant	0.223^{***}	0.272^{***}	0.197^{***}	0.246^{***}
	(0.0354)	(0.0333)	(0.0307)	(0.0310)
Observations	3,136	3,136	2,640	2,640
R-squared	0.058	0.063	0.064	0.075
Period dummies	Yes	Yes	Yes	Yes
Worker fixed effects	Yes	Yes	Yes	Yes
Session dummies	Yes	Yes	Yes	Yes

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: This is a linear probability model, with high effort as the dependent variable and worker fixed effects. Observations from treatments (1E), (1EM) and (1EMP) are included. Standard errors (in brackets) are clustered by session.

high effort again, for example because the employer keeps offering a high wage. We therefore not claim a causal relationship between positive messages and effort. To derive some meaningful conditional correlations, we control for other factors besides feedback influencing effort choice, such as the wage offered, and include worker fixed effects, which will absorb time-invariant characteristics of the worker.

Table 2.6 shows a linear probability model of high effort choice. Column (1) and (2) show the intention to treat regressions, including coefficients for whether the employer had the possibility of sending feedback. Column (3) and (4) include variables for whether feedback was actually sent and break this down into positive and negative feedback. In columns (2) and (4) we interact these terms with the wage.

In all regressions there is a positive and significant coefficient for the wage. From Column (1) we can see that an increase of a wage offer by 1 point is related to an increase the probability of choosing high effort by 0.964 percentage points. The coefficient on the possibility of feedback by the employer is not significant. The coefficient of the interaction between the wage and whether feedback is possible is positive and significant. This confirms the earlier finding that the possibility of feedback is related to a *stronger* relationship between wage and effort.

In Columns (3) and (4) feedback is broken down in positive and negative feedback. We can see a differential pattern. Column (3) shows that receiving negative feedback is associated with a 2.4 percentage point higher probability of choosing high effort and positive feedback is related to a 8.0 percentage point increase in the probability of choosing high effort. Only the coefficient for positive messages is significant. When we include the interaction with wage in Column (4), we see that the interaction coefficient of wage and these terms is positive, but again only significant for positive feedback.

These coefficients suggest that especially positive messages play a role in strengthening reciprocity. However, as can be seen in Table 2.4, positive messages are mainly sent out when the worker already chose high effort. The results therefore suggest that for these relationships were the worker already exerted high effort, positive feedback makes the workers more likely to choose high effort *again*. Feedback therefore mainly seems to work as a positive reinforcement mechanism.

Negative messages are predominantly sent out following low effort. The lack of a significant coefficient for negative feedback therefore suggests that workers who already exerted low effort are not much more likely to choose high effort following negative feedback. However, workers are also not more likely to choose low effort following negative

Table 2.7: The relation between wages, feedback and promises.

Dependent variable: Wage offer in period t	(1)	(2)	(3)	(4)
Wage offer$_{t-1}$		0.198^{***}		0.198^{***}
		(0.0413)		(0.0413)
Rejection?$_{t-1}$	1.169^{**}	1.966^{***}	1.445^{***}	2.190^{***}
	(0.442)	(0.388)	(0.497)	(0.428)
High effort?$_{t-1}$	3.000^{***}	2.575^{***}	2.995^{***}	2.464^{***}
	(0.561)	(0.467)	(0.613)	(0.504)
Feedback possible?	-0.680	-0.474	-1.324	-1.101
	(1.182)	(1.006)	(1.321)	(1.146)
Negative feedback sent?$_{t-1}$			0.912	0.569
			(0.732)	(0.660)
Positive feedback sent?$_{t-1}$			0.825^{**}	0.981^{**}
			(0.341)	(0.363)
Promises are possible?	1.353	0.841	1.302	0.842
	(1.190)	(0.922)	(1.254)	(0.974)
Promise made?$_{t-1}$			0.00606	-0.0817
			(0.338)	(0.290)
Observations	2,954	2,932	2,954	2,932
R-squared	0.582	0.600	0.582	0.601
Period dummies	Yes	Yes	Yes	Yes
Employer fixed effects	Yes	Yes	Yes	Yes
Session dummies	Yes	Yes	Yes	Yes

$^{*}\ p < 0.1$, $^{**}\ p < 0.05$, $^{***}\ p < 0.01$

Note: This is a linear regression, with the wage offer as the dependent variable and employer fixed effects. Observations from treatments (1E), (1EM) and (1EMP) are included. Standard errors (in brackets) are clustered on the session level.

feedback. Negative feedback therefore seems to have little impact on the worker's behaviour.

Result 2.3. *Of the feedback sent, only positive feedback is significantly related to higher effort levels and a higher degree of conditional reciprocity on behalf of the worker.*

We now discuss how feedback influence wages offers. Table 2.7 shows a series of fixed effect regressions of wage offers. Columns (1) and (2) provide an intention-to-treat analysis. Column (3) and (4) include indicator variables whether positive or negative feedback was sent. From all regressions it follows that a rejection in the previous period is related to a higher wage in the next period. The next period wage is 1.1 to 2.2 points higher. High effort is also related to a higher wage in the next period: the coefficients suggest that the wage is 2.4 to 3.0 points higher. Including a lagged dependent variable does not change the variables significantly.

As discussed in Chapter 1, the coefficient for high effort, despite being significant, is small. This means that there is relatively little change in wages offered to workers that chose high effort and workers that chose low effort. This is surprising, especially if we take the difference in the employer's payoff into account. High effort yields the employer a benefit of 40 points, while low effort only yields the employer 5, 10 or 15 points, depending on the session.

The possibility of sending messages as feedback does in itself not significantly influence wage offers. The coefficients on whether feedback is possible and whether promises are possible are both statistically insignificant in Column (1). However, when we include variables to indicate a negative or positive message was sent, we see that negative feedback is related to wages of 0.6 or 0.9 points higher and positive feedback is related to wages of 0.8 to 1.0 points higher. Despite that the coefficient for negative and positive feedback have similar magnitudes, only the figure for positive feedback is significant.[9]

2.5.3 Promises

We next ask how promises affect these results. In treatment (1EMP) the workers can send out a pre-specified promise ("If you offer me a high payment, I will choose high effort."), in addition to the possibility for feedback by the employer. In 73.0% of the cases workers make this promise.

The possibility of making promises does not increase average effort. As we can see in Table 2.3, in game 3 of treatment (1EMP) the workers choose high effort 33.2% of the cases. This is lower than in the corresponding treatment (1EM) figure, 42.7%, and only barely higher than the share of high effort in game 3 of treatment (1E), which was 32.7%. None of these differences are significant.[10] Also the rate of acceptance does not differ significantly between treatment (1EMP) and treatment (1EM).

If we consider wages above and below the median separately, we do not find a significant difference in effort level either. For below-median wages ($w < 16$), the average share of high effort provision is 19.5% in treatment (1EMP) and 23.8% in treatment (1EM). For wages at or above the median ($w \geq 16$), the shares are 48.7% for treatment (1EMP) and 56.1% for treatment (1EM). These differences are not significant $p = 0.624$ and $p = 0.419$). This suggests that promises, at least in addition to feedback, do not lead to a stronger or weaker relationship between wages and effort. The linear probability model of high effort

[9]These figures should not be given a causal interpretation, because the employer is both sending the message and subsequently setting the wage offer. There is scope for omitted variable biases: for example, a positive message could indicate that the employer deems the worker to be trustworthy, which in itself could lead to higher offers. In the regression we control for past effort choice, but this variable might not capture all relevant information to establish whether the employer deems the worker trustworthy.

[10]A t test comparing the shares in treatments (1EMP) and (1EM) yields $p = 0.146$.

in Table 2.6 confirms this. Column (2) shows that the the coefficient of interaction term between the wage and whether promises are possible is not significant. The estimate is also small compared to the coefficient of interaction between wages and the possibility for feedback.

Result 2.4. *The possibility of sending out promises, in addition to feedback, does not significantly impact average wages and average effort or the relationship between wages and effort.*

However, even though the *possibility* for making promises, on aggregate, does not affect effort, we might still see a difference in behaviour between workers who made a promise and workers who did not make a promise. In Column (4) of the linear probability regression of effort in Table 2.6, the coefficient of the interaction term of wage and having made a promise is positive and significant. This suggests that there is a positive correlation between making a promise and exerting high effort, at least when the wage is high enough. This might seem at odds with the lack of effect of promises we found earlier, but could suggest a sorting effect: for example, if promises are only made by the workers who, without the possibility for promises, were already responding conditionally to wages and the other workers do not make promises, then there is no overall change in behaviour, but the coefficient for the interaction term between wages and having made a promise will still be positive and significant.

Promises do not seem to affect wage offers significantly. In Table 2.7 the coefficients for whether promises are possible and whether a promise was actually made are not significant.

2.6 Discussion

2.6.1 Comparison between students and entrepreneurs

In our experiments both students and entrepreneurs participated. Most of these entrepreneurs were managers from small and medium sized enterprises. The vast majority of entrepreneurs in the game were given the role of employer, to reflect their actual role of employer within the business. Overall, we find little evidence for strong differences between our student subjects and our employer subjects. Appendix Table 2.8 breaks down the feedback messages by whether the employer in the game was an entrepreneur and shows that behaviour is similar. Appendix Table 2.9 compares the offer behaviour, broken down by player type, and suggest that there are some minor differences between the groups, but this behaviour is not significant.[11]

2.6.2 Concluding remarks

As we already have seen in the experiments of Chapter 1, and confirmed in the experiments in this paper, the average level of worker's effort in this experiment is low. In Chapter 1 we saw that participants acting as employers in our experiments Ghana were less likely to use wage incentives to motivate workers, compared to the United Kingdom, which contributed to the lower level of effort. In this chapter we introduced the possibility for feedback as well as for promises in additional experiments run in Ghana, arguing that because of the high prevalence of informal employment in Ghana, social incentives could be more effective than monetary incentives. Communication, and in particular feedback, indeed seems to matter. Feedback increases conditional reciprocity on behalf of the worker. For higher wages, feedback, and in particular positive feedback, increases effort. This suggests that communication is not an substitute for monetary incentives, but rather a complement.

The absence of a separate effect for promises, even at high wages, could suggest that a motivation to avoid guilt or a motivation to be consistent across the choices have little role in influencing effort. However, we did not study promises in isolation of feedback, and it might very well be that promises in isolation have an effect.

As discussed in Section 2.2, a preference for social approval is often suggested as an underlying motive for why feedback can function as an effective way to encourage higher effort. However, we find that feedback is mainly effective at high wage levels. At low wage levels, there is no significant difference in the choice of effort. This begs the question why workers would be less receptive to social approval incentives when facing an employer offering a low wage. The literature on social approval suggest that it is important to take into account who the social approval is coming from (see e.g., Gächter & Fehr, 1999). In our experiment the social distance between the worker and the employer does not play a role, since the worker and employer are interacting anonymously. However, the actions of employers could reveal something to the worker about their personality or identity, and this could make the worker care more about gaining their approval. An example would be that a worker observing a high wage sees the employer as being of a "noble" or a

[11] In the robustness checks in Chapter 1, part of the sample of this experiment is used to analyze the use of incentives by the two groups. As discussed in Chapter 1, here there is some difference between the student participants and the actual entrepreneurs: the entrepreneurs' wage offers are responding more to the worker's effort choice than the wage offers by the student participants.

"just" personality type, with whom the worker would like to identify (see e.g., Akerlof & Kranton, 2000).

Another interpretation of the results is that feedback reduces the uncertainty surrounding distributional concerns. A worker with distributional concerns does not only care about their own utility but also about the utility of others. However, if the worker faces imperfect information regarding the other employer's utility, feedback allows the worker to learn about the preferences of the employer. Receiving a positive feedback message indicates that the worker is behaving in line with the employer's interest.

Alternatively, the increase in effort after a positive message could be a result of positive reciprocity. If the worker gains utility from simply receiving a positive message, then reciprocity models suggest that the worker is more likely to reciprocate. However, the credibility of the message is important. If we analyze "sarcastic" messages (positive messages that follow after low effort) and include these in the linear probability regression of Table 2.6, we find that these messages are not effective in increasing effort.

Another question is how anonymity influences our results. Earlier experimental studies indicate that social approval incentives became more effective when there is some degree of social familiarity (Gächter & Fehr, 1999) or when there is a face-to-face interaction (Rege & Telle, 2004; Falk et al., 1999). Real life employment situations are far from being anonymous and particularly in informal unemployment social networks play an important role. Nevertheless, many firm surveys still suggest that despite the higher degree of social proximity between workers and employers in many developing countries, worker management is problematic (Bloom & Van Reenen, 2007).

A further unexplored issue is how feedback can function when it is more public. Praising and, in particular, shaming are often inherently public. Part of the deterrence effect of shaming is that other people know. A relevant question here is whether publicly shaming someone of an action, which does not only tell third parties about the action itself but also how the action was perceived by the other party, is more effective than just spreading the information on the action.

Acknowledgements This chapter is an output from research funding by the UK Department for International Development (DFID) as part of the iiG, a research programme to study how to improve institutions for pro-poor growth in Africa and South-Asia. The views expressed are not necessarily those of DFID. These experimental sessions are part of a wider study on Ghanaian entrepreneurship. Funding from the Economic and Social Research Council (ESRC) is gratefully acknowledged. Many thanks to Martin Černý, Ibrahim Lawal Musah, Anna Consomer, Dan Kumi Ntow, Eric Agyekum and Marzu Abdulai for excellent research assistance. Also thanks to Václav Těhle, Andrew Kerr and Denise Gray for helping out during earlier sessions of this experiment. Furthermore, we would like to thank Moses Awoonor-Williams for his role in helping out with the logistical matters. We would like to thank Johannes Abeler, Stefano Caria, Ritam Chaurey, Catharine Eckel, Simon Quinn, Hannah Uckat, Sébastien Walker and the participants of the Oxford Firms & Development seminar (2014), the Oxford Post-doc & DPhil Seminar (2014), the Cambridge Summer School on Social Economics (2014), the PODER Sumer School on Behaviour Economics in Development (Cape Town 2014), the Oxford Development Workshop (Oxford 2014), the NEUDC Conference (Boston 2014), the EUDN PhD Workshop (Hannover 2014) and the SEEDEC Conference (Bergen 2014) for helpful comments and suggestions.

Appendices of Chapter 2

2.A Additional graphs and figures

2.A.1 Payoffs of workers and employers

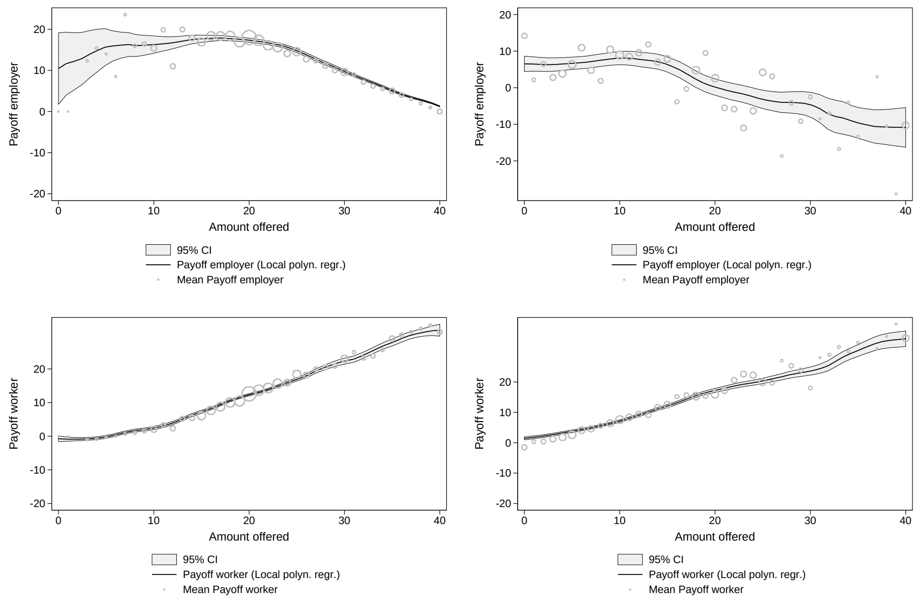

(a) Treatment (1C) **(b)** Treatment (1E)

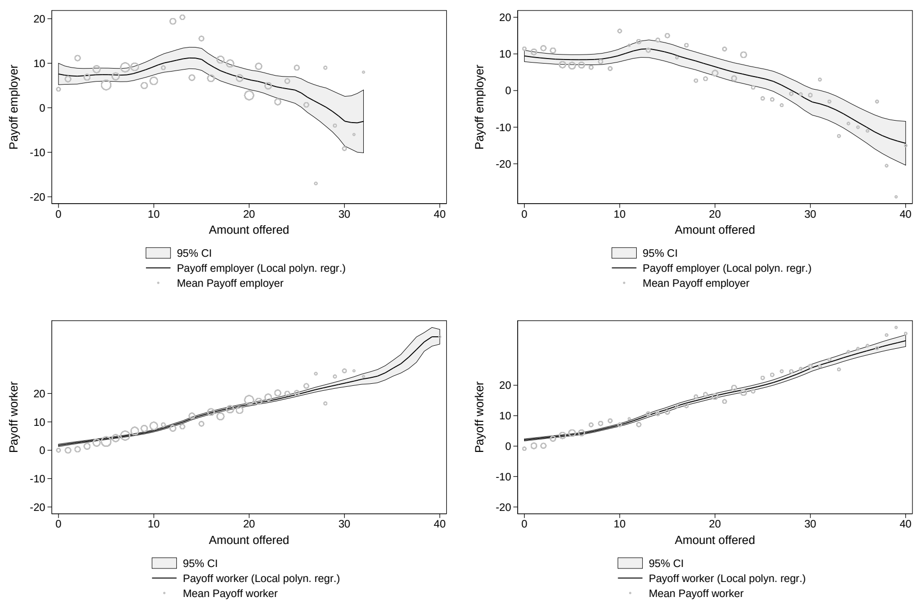

(c) Treatment (1ET) **(d)** Treatment (1ETP)

Figure 2.5: The payoffs of workers and employers, and the offers made by employers in the four treatments.

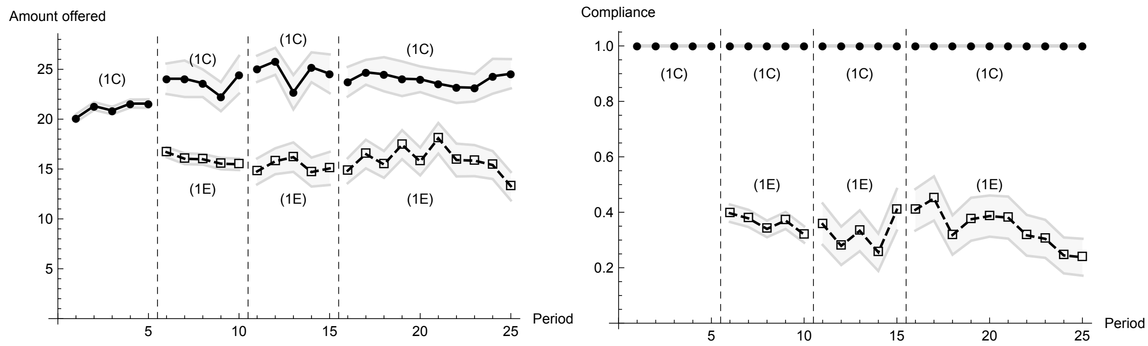

(a) Wage offers **(b)** Compliance

Figure 2.6: Average wage offers and compliance for treatment (1C) and treatment (1E)

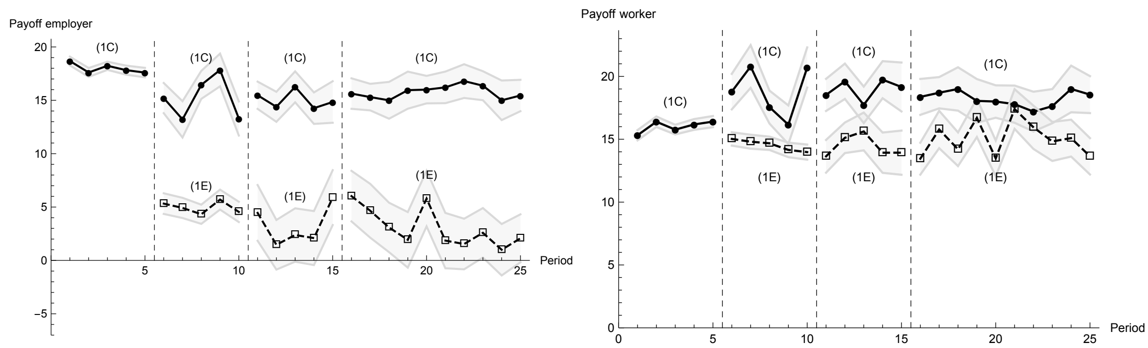

(a) Employer's earnings **(b)** Worker's earnings

Figure 2.7: Average earnings in treatment (1C) and (1E)

2.A.2 Further analysis of the impact of incomplete contracting

Figures 2.6 and 2.7 show the average wage, average compliance and the payoffs for each period of each game. We see that the offered wages are lower than in the perfect enforcement treatment. However, we do not find a significant difference in acceptance rates, apart from the last game. This could be related to the lower wages, which are linked to lower acceptance rates. The average compliance rate is low and appears to decrease over time. This low level of compliance leads to a highly unequal distribution of the surplus, where workers capture most of it.

Amounts offered The bottom graphs of Figure 2.8 shows the distribution of offers. When comparing the full compliance treatment (1C) with the effort choice treatment (1E) we see that in the case of the control treatment a substantial part of the offers lie around 20 points. A clear single peak pattern can be seen here, while in the effort choice treatments two peaks in the distribution of amounts offered appear: one around five or six points and one around 20 points. The peak around 20 point corresponds to a scenario where the surplus is almost equally shared between workers and employers (the amount corresponding to an exactly equal sharing of surplus if the worker chooses high effort is 23 points). In case of the full compliance (1C) model, these offers are in line with the predictions made by the fairness models and reject the predictions of the model assuming perfectly self-interested agents.

At an offer of five or six points and a choice of low effort most of the surplus is captured by the worker (depending on whether the cost of low effort is 5, 10 or 15 points). When accepting an offer of five or six points, it is not (strictly) profitable for a worker to choose high effort, because a high effort offer will cost the worker six points. Employers offering this amount are probably anticipating that the worker is choosing low effort in the next round. Still, the offer is higher than what a purely self-interested worker would need to accept the offer (zero or one), indicating that either the employers still care about fairness at these low wage levels or that they anticipate that the worker will still care about this (and reject an offer that is too low).

Acceptance of offers The top graphs in Figure 2.8 show the relation between wage and acceptance of the offer. When looking at acceptance rates, we see that for the same wage acceptance rates are higher in the effort choice treatments, especially for offers above 10 points: here almost all offers are accepted. The model assuming solely self-interested subjects predicts that any offer of six or higher in case of full compliance treatments and any offer of zero or higher in case of the effort choice treatments should be accepted by the worker. We find that participants still reject offers above these levels. Acceptance rates are higher for the effort choice treatments. This is in line with predictions from the fairness models, as workers now can punish low offers costlessly by accepting them and then choosing low effort. This is of course not possible in the full compliance treatment (1C).

Effort choice This is discussed in further detail in the main text.

2.A.3 Comparison between entrepreneurs and students

Table 2.8 shows the messages sent by the student and the entrepreneur samples (replicating Table 2.4). The pattern in both groups is very similar: low effort is followed by a negative message in the overall majority of the cases, while high effort is followed by a positive message. Entrepreneurs are less likely to send stronger messages to their workers.

Table 2.9 compares wage setting by the student and entrepreneur samples, replicating columns (2) and (4) of Table 2.7. The differences are mostly insignificant.

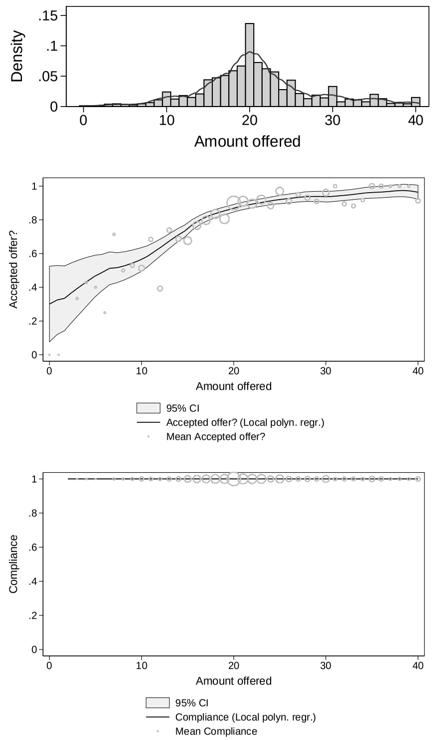

(a) Treatment (1C)

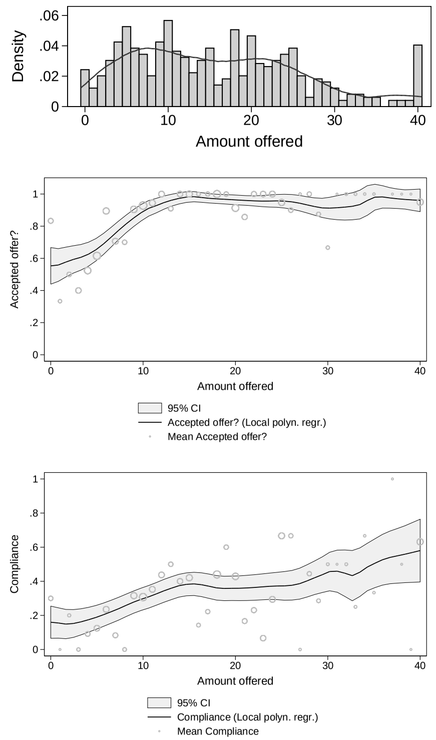

(b) Treatment (1E)

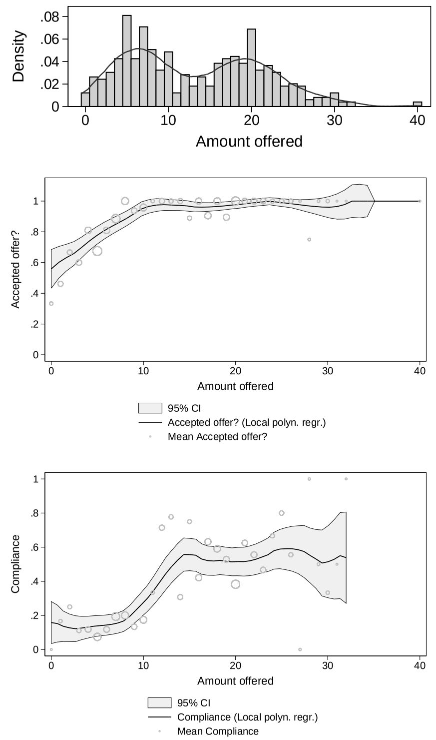

(c) Treatment (1EM)

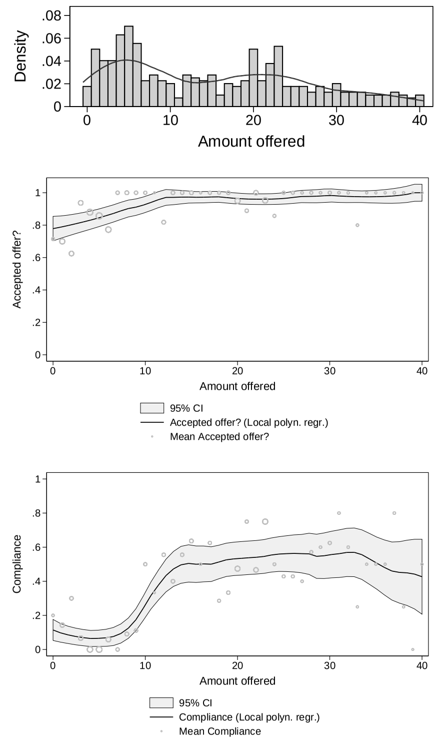

(d) Treatment (1EMP)

Figure 2.8: Acceptance rates, compliance (choosing high effort) and offers made by employers in the four treatments. Only results from Part IV (the last part in the sequence, see table 2.2) are shown. The horizontal axis shows the amount offered by the employer.

Table 2.8: Feedback messages sent in treatments (1EM) and (1EMP) (comparison between entrepreneurs and students)

	Students		**Entrepreneurs**	
Feedback	Low effort	High effort	Low effort	High effort
No feedback	19.6%	8.9%	19.8%	3.5%
Negative feedback				
"Why did you do this to me?"	28.0%	3.6%	42.6%	0.0%
"I pity your children"	6.4%	0.4%	7.9%	0.0%
"You are a disgrace to your family"	10.0%	0.5%	13.9%	0.0%
Positive feedback				
"Thank you"	26.1%	44.3%	10.9%	79.3%
"You are hardworking"	6.6%	18.4%	1.0%	6.9%
"You are a reliable person"	3.2%	23.9%	4.0%	10.3%
Total	100.0%	100.0%	100.0%	100.0%

Note: This table pools the observations from games 3 and 4 of treatment (1EM) and treatment (1EMP).

Table 2.9: The relation between wages, feedback and promises (comparison between entrepreneurs and students).

	Students		Entrepreneurs		Both groups (pooled)	
	(1)	(2)	(3)	(4)	(5)	(6)
Wage offer$_{t-1}$	0.197^{***}	0.197^{***}	0.202	0.203	0.197^{***}	0.196^{***}
	(0.0432)	(0.0431)	(0.154)	(0.154)	(0.0431)	(0.0430)
High effort?$_{t-1}$	2.438^{***}	2.358^{***}	3.671^{**}	3.265^{***}	2.448^{***}	2.371^{***}
	(0.498)	(0.548)	(0.809)	(0.677)	(0.502)	(0.552)
Rejection?$_{t-1}$	2.074^{***}	2.301^{***}	1.519	1.758	2.069^{***}	2.304^{***}
	(0.430)	(0.473)	(1.014)	(1.157)	(0.433)	(0.472)
Feedback possible?	-0.673	-1.272	0.716^{***}	-0.275	-0.350	-0.962
	(1.168)	(1.314)	(0.0601)	(0.227)	(1.053)	(1.181)
Promise possible?	0.713	0.723			0.715	0.725
	(0.982)	(1.032)			(0.980)	(1.037)
Negative feedback sent?$_{t-1}$		0.574		0.796		0.601
		(0.756)		(0.399)		(0.759)
Positive feedback sent?$_{t-1}$		0.867^{**}		2.241^{***}		0.880^{**}
		(0.346)		(0.433)		(0.349)
Promise made?		-0.0439				-0.0443
		(0.290)				(0.290)
Entrepreneur?						
× Amount offered$_{t-1}$					0.00660	0.00906
					(0.139)	(0.138)
× High effort?$_{t-1}$					1.174	0.811
					(1.029)	(0.926)
× Rejection?$_{t-1}$					-0.500	-0.537
					(0.950)	(1.108)
× Feedback possible?					-0.505	-0.753
					(0.625)	(0.821)
× Negative feedback sent?$_{t-1}$						-0.0311
						(0.947)
× Positive feedback sent?$_{t-1}$						1.258^{**}
						(0.553)
Constant	11.77^{***}	11.79^{***}	11.38^{**}	11.37^{**}	11.71^{***}	11.72^{***}
	(1.040)	(1.054)	(2.499)	(2.488)	(0.935)	(0.945)
Observations	2477	2477	455	455	2932	2932
R-squared	0.613	0.613	0.509	0.511	0.601	0.601
Adjusted R-sq	0.567	0.567	0.421	0.421	0.553	0.552
Period dummies	Yes	Yes	Yes	Yes	Yes	Yes
Employer fixed effects	Yes	Yes	Yes	Yes	Yes	Yes

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: This is a linear regression, with the wage offer as the dependent variable and employer fixed effects. Observations from treatments (1E), (1EM) and (1EMP) are included. Note that treatment (1EMP) was only given to student participants and not to entrepreneur participants. Standard errors (in brackets) are clustered on the session level.

2.A.4 Messaging and the wage-effort relationship

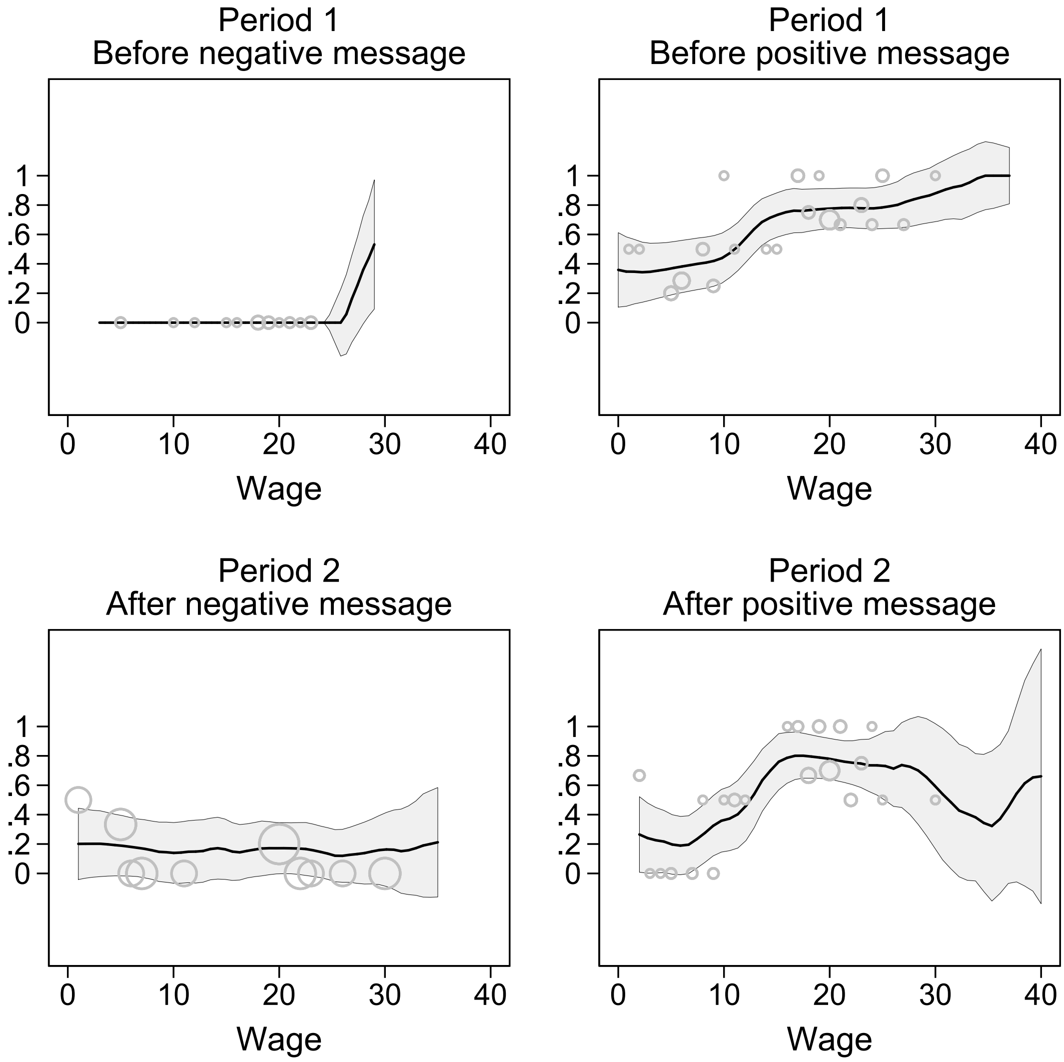

Figure 2.9: The wage-effort relationship before and after a positive and negative message, in treatment (1EM).

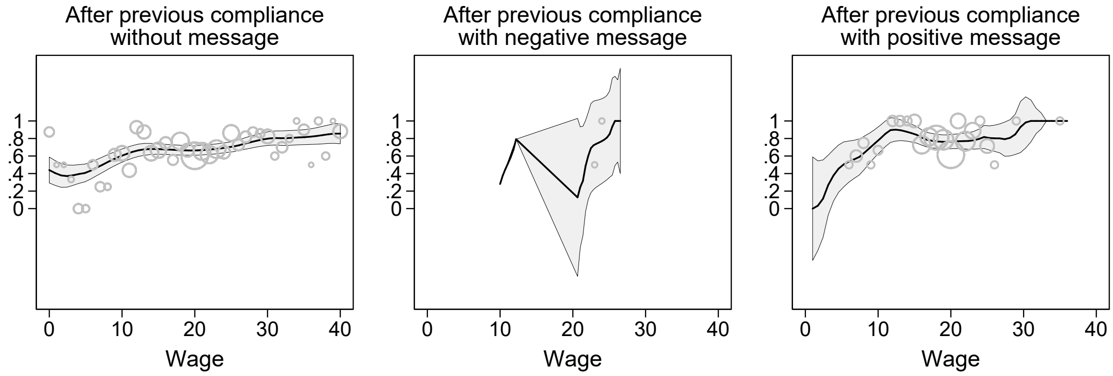

Figure 2.10: The wage-effort relationship change after no message (treatment (1E)), a negative or positive message (treatment (1EM)), given that the worker chose high effort in the previous period.

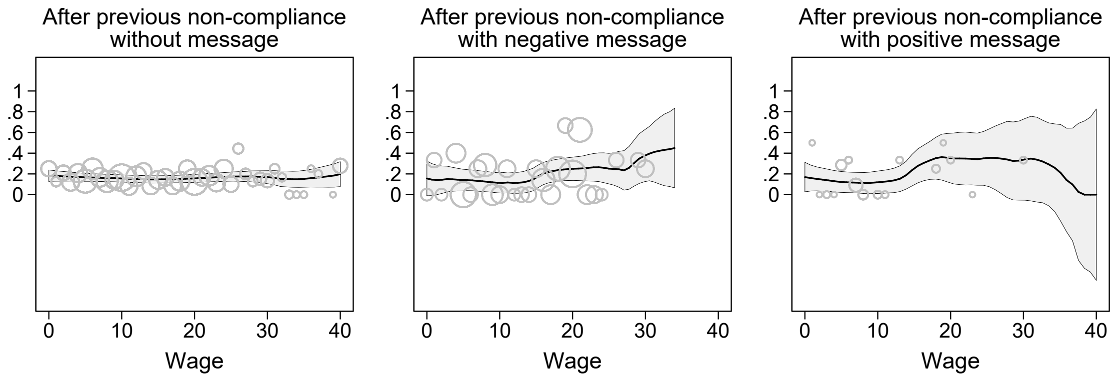

Figure 2.11: The wage-effort relationship change after no message (treatment (1E)), a negative or positive message (treatment (1EM)), given that the worker did not chose high effort in the previous period.

2.B Instructions

2.B.1 Screens

Participants used tablets to make their choices. In Figure 2.12 shows the offer screen for the employer. The user can both slide the slider or use the plus and minus buttons. A popup appears while adjusting the offer (see Figure 2.13), showing the payoffs for the worker and the employer.

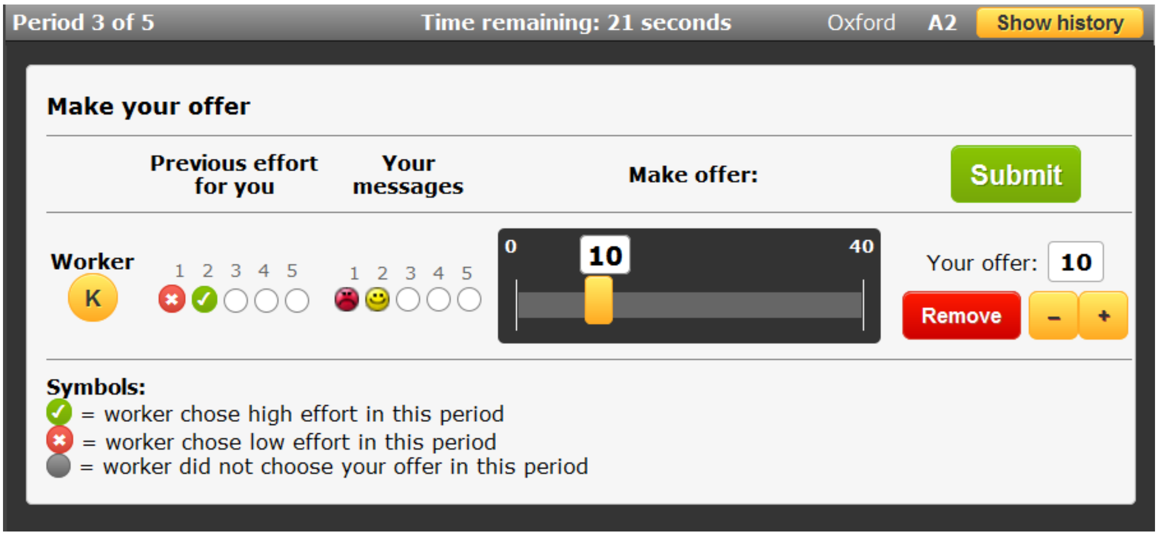

Figure 2.12: Employers make offers.

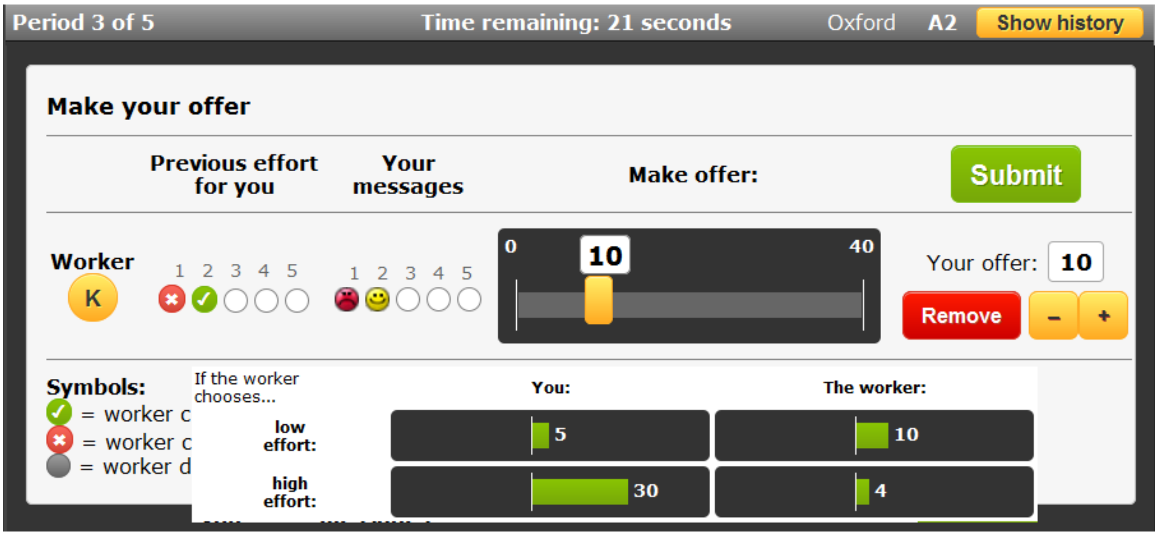

Figure 2.13: The pop-up indicating potential payoffs.

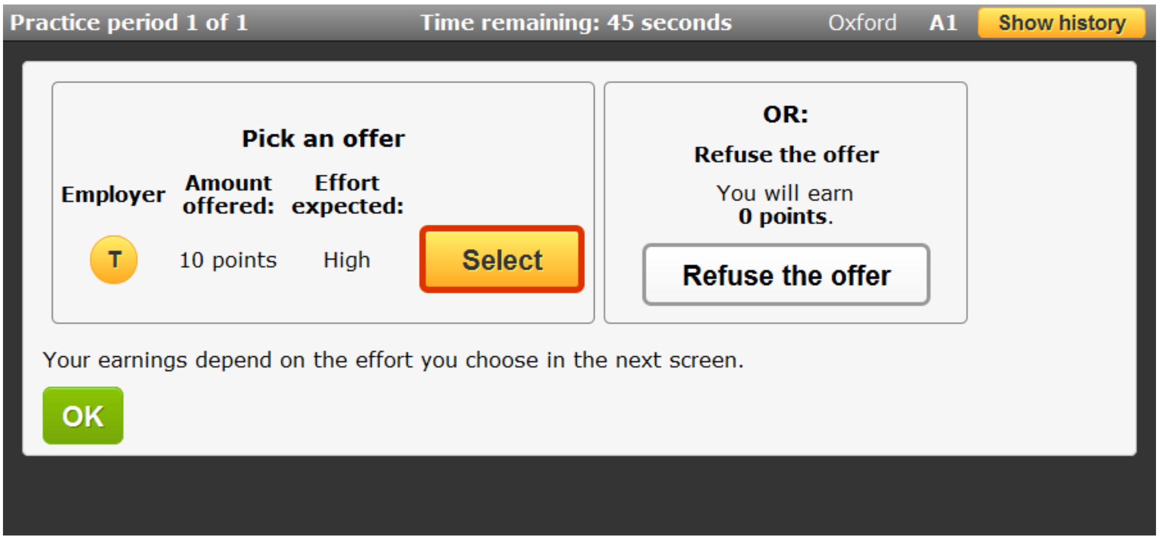

Figure 2.14: Workers choose to accept or reject the offer.

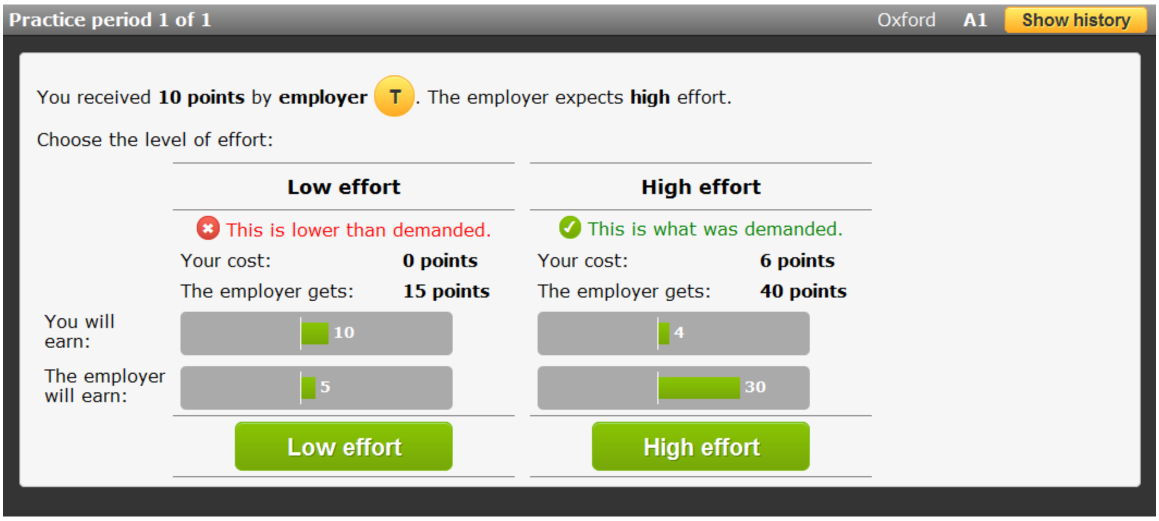

Figure 2.15: Worker chooses the level of effort.

Figure 2.16: Messages in the (1EM) and (1EMP) treatments.

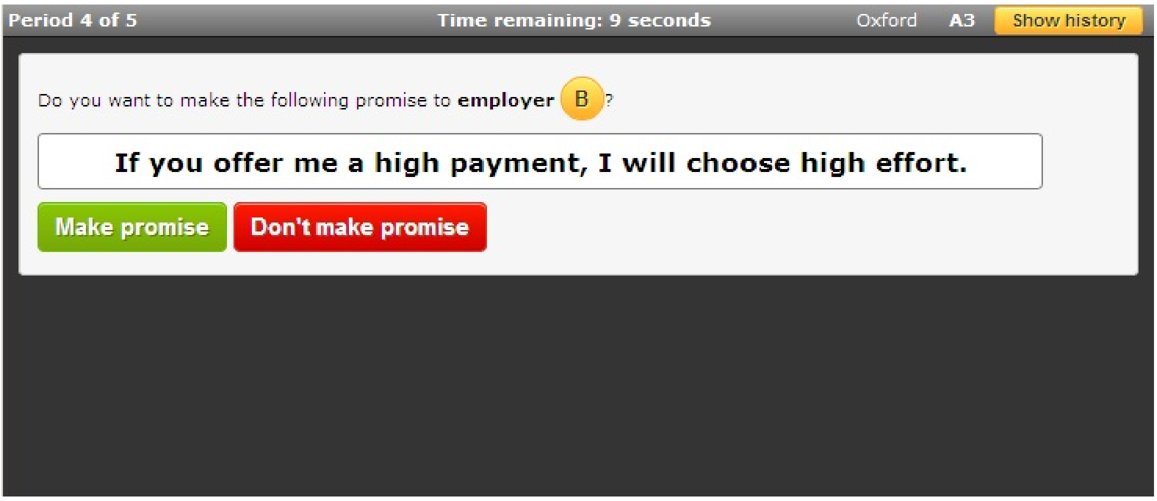

Figure 2.17: Promises in the (1EMP) treatment.

2.B.2 Script

Below is an excerpt of the script used in the experiments. First, we make sure that participants feel comfortable using the tablets, without going through the details of the game. After this, we explain the game in an interactive way, without using the tablets. Finally, we let the participants practice the game for two periods on the tablets, before starting the first treatment.

"Good morning. / Good afternoon. My name is Elwyn Davies. I am a researcher at the University of Oxford. This is Martin Černý, who is also working on this project. Besides this, we have *(insert names of local team)*.

Welcome to this session of the experiment. This experiment is part of a wider study done by the University of Oxford on Ghanaian entrepreneurship and firms. The goal is to see how people behave in a virtual marketplace. Please remain silent during the entire duration of the experiment and keep your mobile phones switched off. I don't want to hear any mobile phones during this session.

During this session you will earn points. These points will be converted to Cedis at the end of the session. 100 points is equal to five Cedis. If you arrived on time, you will earn 10 cedis in addition to this.

I will explain more about this experimental session later on. First, I will talk about how to use the tablet. You will use the tablet to make choices in this experiment. Please have a look at your tablet and read the message. If you are done reading, press OK.

Make sure to touch the screen gently. Do not press it too hard.

(Wait.)

We are now going to practice how to make offers. In this game you will be either an employer or a worker. Employers make offers to the workers. I will talk about that more later on. First we will practice how to make an offer.

Please press somewhere on the gray bar. You will see a number appear. You can select a number between 0 and 40. Please try to select different numbers: 0, 5, 10, 15 and so on. Try to change from 0 to 5, or from 32 to 37. And then change the number again.

You can use the plus and minus signs to change the number. Please try this as well.

If you don't want to make an offer, you can click on *Remove*. The number then disappears. You can get the number back by clicking on the bar. Please try this as well.

If you have any questions, please ask them. We will come to your desk to answer them.

(Wait until everybody is done.)

I will now talk about the experiment itself and what we are going to do. Some of you will be employers, some of you will be workers. We will determine by chance which role you get. Each period the employers start by selecting the payment they would like to offer to the workers. You have to select a number on the gray bar. The workers then choose to accept or reject the payment. If they reject this payment, both get zero points. If the worker accepts the offer, the worker will get paid the amount and will then work for the employer. The worker is working hard for the employer, and this gives a profit of 40 points to the employer. Working hard is tiring for the worker, so the worker will get less, and lose 6 points. But of course, the worker received a payment from the employer.

Let's give you an example. Suppose I am an employer, and Martin is a worker. Workers are indicated by a letter. You cannot tell who in the room this worker is, but it is somebody in this room. Let's assume that Martin is worker A. **(Show the sign of worker A)** I am an employer and can make an offer to Martin, by selecting a number between 0 and 40. **(Ask somebody in the public.)** Can I get a number between 0 and 40 from you? **(Write down the number)**

Ok, I am now going to make an offer of X points to Martin. Martin, do you accept or reject this offer?

(Martin rejects.)

Martin rejects, we now both get 0 points.

Now let's say he accepts the offer.

(Martin gets the offer and keeps it.)

I pay him the X points and he is going to work for me. That gives me a profit of 40 points. But I have paid him X points. So how much have I earned?

(Ask public and check understanding.)

Martin was working hard, and that is tiring for him. He lost 6 points, but got my payment of X points. So how much has he earned?

(Ask public and check understanding.)

So I have earned ... points and Martin has earned ... points.

Please let us know if you have any questions about this. **(Walk around and check understanding with participants.)**

As long as you touch the bar, you can see graphs on top of the screen. These graphs show you how much you can earn if the worker accepts the offer. Please change the number, and see how the graphs change as well. Green means that you are earning money, red means that points will be deducted. Try to offer X points, and check the graphs. These numbers are the same as we just calculated together.

(Initialize the main game by pressing the Start button on the admin screen.)

We will now tell you whether you are a worker or an employer. Please press OK to continue. We will first play two rounds of practice. No points can be lost or earned. Please press OK to continue. If you see a waiting screen, please wait. You will see this screen a couple of times during the experiment. You will have to wait until everybody is finished making their choices.

(Wait until everybody has clicked OK twice)

We are now in the first practice period. If you are an employer, please make an offer by selecting a number on the gray bar. And then press Submit. Make sure to do this before the time runs out.

(Wait until the workers can choose.)

Now the workers can choose to accept or reject the offer. If you want to accept the offer, press Select. Then click OK. For this practice round, make sure to accept the offer, so that you know how this works. Make sure to do this before the time runs out. If you have accept an offer, we will ask you how many points you need in the next round to accept the offer. For example, if you select 30, you will accept all offers of 30 and higher, and reject the offers of 29 and lower. We will also ask the employers what they would like to offer to you. If both of you agree, you will automatically accept the offer in the next period. If you don't accept, the employer will make you an offer again in the next period, just as before. **(At the end of the practice period.)** We will now play for real points, that will be converted to Cedis at the end of the session. Press *I am ready* to continue. If you have any questions, please raise your hand and we will come to your desk to answer them."

Chapter 3

The Use and Misuse of Business Networks: Overcoming Information Asymmetries and Locking Out

Abstract This chapter presents a model of how business networks, such as business associations, can lead to higher quality firms overall. There are two types of firms: competent firms, who provide high quality services, and inept firms, who provide low quality services. Consumers and firms can choose to join a network. Members of this network know exactly what the type of other members is and can therefore assure themselves of consuming high quality services. However, joining the network is costly and the cost of joining differs heterogeneously between agents. The presence of such network leads to a *filtering effect* and fewer inept firms will produce. At the same time, these networks redistribute demand to their own members and make production outside the network less viable. This occurs even when the informational benefits from joining the network are low. This is the *lock out effect*.

3.1 Introduction

One fundamental challenge for firms is to find a high quality inputs, whether it is capital, intermediate goods, services or labour. For many transactions it is hard to objectively verify the quality of what is provided, meaning that specifying quality in a contract is often impossible. Firms might therefore have to rely on other means to verify or ensure quality of their suppliers. Entering into a long-term relationships is one way, but this is problematic when the product or service sought is only a one-off transaction. This is where business networks can play a role. Through the contacts in business networks, reputations can be established and members can gain information that is normally hidden in the market.

This chapter explores the scope for business networks as a way to overcome informational asymmetries. Business networks can take many shapes, from informal networks such as trading or business associations to a formal hierarchy of firms belonging to the same conglomerate, such as the keiretsu in Japan (Rauch, 2001; Morck et al., 2005). Central to this chapter is the assumption that some institution is necessary to allow reputations to be established. The main aim of this chapter is to discuss the implications of the existence of such a network for the composition of firms in the economy.

In this chapter I present a stylized model of networking, that aims to mimic the interactions in existing business networks, such as business associations, or certain social clubs with a strong networking function (e.g., local charity organizations). There are two types of firms: *competent* firms who provide high quality services and *inept* firms who provide low quality services. Besides this there is a group of final consumers, who only consume. Consumers and firms can choose to join a business network.[1] Inside the business network, information about quality flows freely and members of this network know exactly what the quality of other members is and can therefore assure themselves of consuming high quality products. However, joining the network is costly and the cost of joining differs heterogeneously between agents.

This network structure leads to two effects. First, the network causes sorting in quality: if the costs of receiving bad quality products becomes too high, more firms will see a benefit for joining the network, which makes it harder for inept firms outside the network to produce. This improves the quality composition of the economy, filtering out the lower quality firms. This is the *filtering effect*. The higher the cost of receiving a bad quality product is, the stronger the need to join the network. At very high costs, practically all consumers and producing firms will be part of the network, so that for a firm outside the network it is no longer profitable (in expectation) to produce. Because the proportion of inept firms outside the network is relatively higher, fewer inept firms are producing.

However, the firms use networks not only as a way of improving the quality of their consumption, but also as a way of coordinating supply and demand: it provides an implicit mechanism for competent firms inside the network to agree to only consume each other's services, while still providing to firms outside the network. This means that firms outside the network receive fewer customers than they would have received when network opportunities were not present. However, this *lock-out effect* bears a cost, as firms are joining the network, even when the informational benefit from networking is small. Even when the cost of receiving a bad quality product gets close to zero, the presence of a

[1] Usually, "pure" consumers are not members of business associations; however, they might be members of networks that are formed around social clubs, such as charity organizations. Alternatively, the "pure" consumers can be thought of as firms in a different segment in the market (e.g. manufacturing goods instead of services) or by their nature having a customer base outside the local network (e.g. exporting firms).

network leads to fewer firms entering and producing compared to the case without such network.

The model is especially applicable to products where the quality is difficult to measure or to be included in a contract. One particular group of transactions for which this often holds true are services. Services do not have a physical appearance like goods have. They are intangible and non-storable, which means that they are consumed while they are produced. It is for this reason often not possible to try out the service, in the way that this is possible for goods, where it is often easier to try out a sample of the good. Moreover, because of their intangible nature, it is difficult to draw up complete contracts where every detail of the service is captured and, even if a contract is detailed, it is often difficult to objectively determine whether the services were provided according to the terms of the contract and what the quality of the services was. Even if indicators of quality are available, they are often affected by other factors. For instance, in order to measure the quality of an advice of a consultancy firm it could be possible to use all kinds of output and profitability indicators, but these indicators are also affected by the extent that the advice is actually followed and implemented or by other external factors. This could lead to a potential moral hazard problem, where service providers have an incentive to minimize effort and quality to the standard which can be enforced.

Second, many services are tailor-made and customized to the needs of the service consumer. This makes every provided service different from other services. The possibility of standardizing a service will differ from sector to sector. For software producers, the product provided to each consumer is more or less the same, while consultancy or lawyer firms will provide a different service to each of their clients. This great degree of differentiation means that a centralized or organized market is often absent. The matching of consumers and producers often more resembles the matching of employers and employees than that of standardized goods: consumers look around for a potential provider and bargain over the price and what service should be delivered.

Services are often locally produced and consumed. While this is necessary in case of "contact services" where physical proximity is required (e.g. plumbing or cleaning), is actually a wide range of services where this need for physical presence is not necessary (e.g. Jones & Ruane, 1990). Thanks to improvements in communication technologies, many services could actually be provided from a different location than the consumer's. To some extent this is already happening, think of the increase in the use of teleconferencing in businesses or the outsourcing of customer services call centres. However, many services are still traded locally to a large extent, which has been called the "proximity burden" of services (Christen & Francois, 2010). Several reasons can be given for this: legal and institutional barriers to trade in services, the need for face-to-face interactions (Storper & Venables, 2004) or the difficulty in finding customers or suppliers outside your own area.

The model focuses on the last reason. Reputations are often strongly local: a well-known firm in one region might be unknown in a different region. Building a reputation requires time and effort, where firms have to make sure that they get noticed in a market and are able to reach potential customers. Reputations can be spread through a variety of ways, in which recommendations and social networks have an important role. Personal and often informal interactions are important in these recommendations and social networks. Since these are often organized on a local level, it means that this kind of matching primarily takes place on local levels.

3.2 Related literature

The matching problem is essentially one of incomplete information. In this situation the provider has more information about the service than the consumer. Models often consider two types of information asymmetries: the difficulty in finding a high quality provider or the difficulty in finding the right complimentary provider. In the first case, services are differentiated *vertically*: some firms provide high quality services, while others provide lower quality services, but this ordering in high and low quality services is the same for all agents. The consumer does not know the quality of the service in advance. In the second case, products are heterogeneous and differentiated *horizontally*: specific customers prefer one specific type of service to another type, but this preference will differ between customers.

In many ways the matching problem of trade in services is similar to other matching problems in the literature. In this chapter I focus on the case of vertical differentiation, where customers try to find a high quality service provider. However, in the literature a large part of the models focus on either horizontal differentiation, where agents try to match their complementary type, or on improving matching from the producer's point of view. This section provides an overview of some of these models and discusses their applicability to services provision.

3.2.1 Models of reputation

Reputation is one of way of dealing with the information asymmetry of unknown quality. Based on the history of qualities provided in the past, firms can gain a good reputation and use this reputation as a signal for service consumers to convince them that their products are of a high quality. The existence of such a reputation mechanism also means that firms have an incentive to provide higher quality products: providing a low quality product would harm your reputation and therefore also your chances of attracting customers. Shapiro (1983) provides a model of perfect competition, but imperfect consumer information, where firms are homogenous and incur costs to be able to provide high quality products. Consumers are heterogeneous and differ in their taste regarding quality and their willingness to pay. A firm with an established good reputation can charge a price above the cost of production, but firms with no established reputation will initially charge a price below the cost of production, to convince consumers to buy their product (an introductory offer). Reputation can be seen as an asset, in which firms first invest by offering products at low prices, and in return it offers the possibility to charge higher prices as soon as the reputation has been established. In the case where the reputation is determined by the quality of the product in the last period, firms that have established a reputation can charge a price above cost:

$$p(q) = c(q) + r\left[c(q) - c(q_0)\right].$$

The price of the product $p(q)$ is equal to the cost of the product $c(q)$ plus a term representing the returns of having established a good reputation, which is equal to the interest rate r times the investment made for providing a high quality product, which equals the cost of providing a high quality product $c(q)$ minus the cost of a minimum standard product $c(q_0)$.

The prospects of being able to charge a price above cost, provides an incentive to provide high quality products. In this way reputations exist as both a way of overcoming informational asymmetries as well as an incentive mechanism.

However, reputations can form an informational barrier to trade: the reputation that well-established firms have built up make it difficult for entrants to compete. Grossman & Horn (1988) consider the situation where products in the domestic market are supplied by a well-established foreign sector and the lack of reputation puts domestic entrants at a comparative disadvantage. This informational barrier to entry is sometimes used as a justification to legitimate infant industry protection. In the first period, the quality of the foreign firms is known, while the quality of local firms is unknown. By introducing a tariff on foreign products only in the first period, people can be convinced to also consume local products and get to know their quality. In the case with rational expectations, they pay a price which makes them indifferent whether they consume a foreign product or an average quality domestic product. However, this leads to a *fly-by-nights problem*: the local firms who need to invest more to produce high quality products have an incentive to provide a minimum quality product and thereby gaining a profit in the first period. Grossman & Horn (1988) show that infant-industry protection reduces welfare in this case. Only when expectations of local firms are very pessimistic, which means that entry is difficult, a minor tariff can raise welfare.

They note that these informational issues especially matter when it comes to trade in services, because of the tailor-made nature of services: the producer and consumer must enter into a purchase contract before the performance of the service is actually delivered. However, even though the infant industry argument might be used to justify protection and this might be relevant for some countries, for most countries the situation is actually the opposite: local firms providing services in their home markets, while it is difficult for foreign firms to enter the market and compete with them.

These two models assume perfect competition, but imperfect information: customers have an overview of all the firms that are present and the price they offer for a specific good, but the quality is unknown. While this assumption might be realistic for trade in goods, where often either reference prices or organized markets exists, this assumption might be less realistic for trade in services: since services are often tailor-made to the customer, the customer must specify its needs to the producer, which is costly and this means that customers will limit the number of potential providers they start negotiations with. As is discussed later on, the search process in that sense is more similar to matching between employers and potential employees than to these perfect competition models, where prices play an important role.

3.2.2 Matching and geography

The occurrence of matching is often used in economic geography as an explanation of why certain spatial patterns develop. Agglomeration means that population is dense and close to each other and this improves matching. Duranton & Puga (2004) give an overview of models explaining the interplay between matching and agglomeration and discuss the two main effects of agglomeration on the matching process: it improves the *chance of being matched* and it improves the *quality of matches*. They apply these ideas to matching on the labour market, but in principle these ideas can also be applied to matching between consumers and producers.

For agents engaging in matching agglomeration has the benefit that there are plenty of opportunities nearby. This can either be a disadvantage, because having more opportunities and applicants leads to more frictions (resulting in decreasing to returns to scale in the matching process), but can also be an advantage: when more opportunities are present, chances are lower that none of the opportunities work out, leading to a higher probability of being matched (Coles & Smith, 1998).

Another advantage of agglomeration is that it can not only improve the matching process itself, but also the quality of the matches themselves. In a situation of horizontal differentiation, where workers have heterogeneous skills and firms differ in their preferences for skills, agglomeration allows firms seeking for workers to choose from a more diversified range of skills. Increasing the number of workers means that firms can find a closer match between worker skills and the skills they need for a certain job, leading to less mismatch costs. This agglomeration effect is mitigated by an urban crowding effect, creating a trade-off between these two effects.

In the case of vertical differentiation, where agents care about quality, agglomeration can also improve the quality of the matching. Venables (2011) presents a model where the costs of urban crowding functions as a signalling mechanism: just like education in the signalling model of Spence (1973) could function as a signal of ability, location could function as a signal of ability and lead to self-selection: in an expensive city there will be more agents with higher abilities. In this model agents match with another agent to undertake a joint project, which gives them a joint benefit. It is assumed that good matches are of greater value for high ability type agents than for lower type agents. As a result of this assumption, a separating equilibrium can be sustained, where low type agents stay in a low-cost place, while high type agents are willing to incur the cost of living in an expensive city that makes it possible for them to match with another high type agent to undertake their joint project.

Using a similar model of two types of agents undertaking joint projects, Storper & Venables (2004) analyze a model where agents can form "in-groups" where past performance can be tracked. The probability of success of a project, which can either succeed or fail, is determined by the ability of each of the partners and the (costly) effort they make. If a project fails, the two matched agents inside the group face the probability of being ejected from the group. This leads to a situation where agents who are a member of the in-group have a higher than average ability. Members will therefore have higher earnings, because they match more often with a more able agent. This model does not only target the information asymmetry, but also the incentive structure: because these in-groups reveal more information about their members, there is an incentive for agents to make a greater effort to provide high quality output. Storper & Venables (2004) relate the formation of these in-groups to the importance of face-to-face interactions, because this allows agents to judge each other and to be judged.

To a large extent the ideas of these models can also be applied to service provision. However, this assumption of jointly undertaking a project ignores the aspect of the asymmetry of the relation between customers and producers in service provision: in service provision, it is often one party (the producer) that incurs the costs and the other party (the customer) enjoys the benefits. The customer compensates the producer for its costs by a transfer in the form of the price paid for the service. This leads to a different incentive structure: without any reputational effects or other implicit or explicit incentives to induce effort, the producer will have an incentive to minimize its effort to the minimal objectively verifiable level. Imposing an informational structure, such as reputations that can be spread or "in-groups" that can be formed, will have a different effect on a consumer than on a producer and this difference should be taken into account in a more complete specification.

3.2.3 Matching on the labour market

The matching problem of service firms is similar in many ways to the matching problem between employers and employees. Employers try to hire the employee with the right set of skills and abilities, while potential employees look around for employers that will give them the highest personal utility, both in monetary as in non-monetary terms. This search process has been studied quite extensively by labour economists (Rogerson et al., 2005).

Central to these models is the notion that *frictions* are important. Because there is no centralized market where workers and employers can meet and trade at a single price, agents have to search for a suitable match. When a worker meets a possible prospective employer who gives him an offer, the worker will compare this with the expected benefit of waiting some more time and receiving an offer from another employer. This could lead to situations where workers of similar abilities end up being paid different wages.

To a large extent these ideas could also be applied to trade in services: because services are often tailor-made to wishes of the consumer, there is no centralized market where service providers and consumers meet. The matching between service providers and consumers resembles the matching of workers and employers in the way that both parties have to engage in a purchase contract and negotiate over the terms. However, most labour matching models assume that workers can only match one employee at a time, while most service providers will be able to provide their services to a second customer.

Role of social networks

Sociological studies have shown that a substantial part of the workforce have found their job through personal contacts (Granovetter, 2005). Several models have been set up that emphasize the role of personal networks in finding a job.

Montgomery (1991) studies the role of social networks in labour-market outcomes by focusing on the role of referrals. In his model, production takes place in two periods, but workers live only for one period and after this period the firms can observe the productivity of the first period worker. At the end of the period, the firm can choose to engage in referral hiring or not. In case of referral hiring, the first period worker can relay the employer's offer to a second period worker in his own network, who can accept or reject his offer. Agents that do not engage in referral hiring or who refuse the referred offer, will be matched on the market.

Crucial for the working of the model is the assumed social structure, which depends on two parameters: *network density* and *inbreeding bias*. A high network density means that first period workers and second period workers are well connected and a high inbreeding bias means that high quality first period workers are more connected to high quality second period workers than low quality first period workers are. Because of this assumed inbreeding bias, employees will only choose to engage in referral hiring when the first period worker was highly productive. In a situation with a high networking density and a high inbreeding bias, high quality second period workers are mostly found by referral hiring, implying that mostly low quality workers are left over on the market, lowering the wage offered on the market, which is a "lemons effect".

Calvó-Armengol & Jackson (2004) focus less on the ability of the worker, but more on how a specific network pattern between agents affect unemployment. In his model, agents receive job offers from either an employer or from another agent, which they can choose to accept or pass on to other agents to which they are linked. This creates a situation where an agent who has been unemployed for a long time is more likely to remain unemployed than an agent who has just became unemployed, because being unemployed for a long time means that it is more likely that the agents to which you are connected are unemployed as well, meaning that it less likely that you receive a relayed offer from them.

3.2.4 Networks and firms

These ideas of social networks have also been applied to the relation between firms. Several empirical studies have shown the importance of social relationships and networks in markets. Uzzi (1996) finds that social relationships in apparel markets are important in cases where firms have repeated interactions. These social relationships increase trust between the parties to overcome hold-up problems, improve the transfer of information between parties and encourages joint problem solving to deal promptly with problems that arise (Jackson, 2008). In their study of the fish market of Marseille Weisbuch et al. (2000) show that a substantial share of fish buyers are loyal to a single seller, while other fish buyers purchase from many sellers. For these "tied" sellers, predictability of supply is important, while price is more important for the untied sellers. In this way networks offer a possibility of insurance: a high cuisine restaurant that is famous for a certain fish dish will consider it to be more important that this fish dish is served than that the price is not too high.

Kranton & Minehart (2001) examine exchange of goods through networks as an alternative to markets. Sellers produce one unit of indivisible output and buyers desire one unit, but have a private, uncertain valuation of the produced output. Goods can only be obtained if a link exists between the buyer and the seller. Within the established link pattern, the goods are auctioned by an ascending-bid auction (which can be regarded as a "fictional Walrasian auction"), so the price is determined by the valuation of the seller who drops out last in each auction. It is assumed that there are more buyers than sellers. Kranton & Minehart (2001) point out two reasons why networks will be formed: one economic reason for suppliers and a strategic reason for buyers. For suppliers the network offers a possibility to *pool uncertainty in demand*: when a seller is connected to multiple buyers, a drop in the valuation of a single buyer will have less effect, since the firm can now rely on more sellers. For the buyers, being well-connected means that you have access to more sources of supply and therefore can *assure better terms of trade* for sellers. However, the formation of networks is costly, so buyers and sellers have to balance their benefits with costs.

The role of social networks in international trade is assessed by Rauch (2001). A central question is whether networks will grow or shrink as a consequence of international trade. This is difficult to answer, because networks are often created because of trade. In an earlier paper, Rauch (1999) argues that a distinction should be made between goods with reference prices and those without reference prices. Some commodities are sufficiently differentiated so that prices cannot convey enough information for international trade. Networks are in particular important for differentiated products (and therefore also for most kinds of services).

In this setting networks exist as an *informal barrier* to restrict market access. An example of such a strong network would be the Japanese *keiretsu*: a group of companies with interconnected business relationships and shareholdings. Companies like Mitsubishi are often involved in a wide range of services and goods provision, varying from bank and shipping services to cars and petroleum.

These networks can be transnational. Much attention in the literature has been paid to the role of ethnic networks facilitating trade, for example the role of the overseas Chinese Diaspora in facilitating trade. Rauch & Casella (2003) set up a model where firms need to find a producer of the right complementary type to form a joint venture. The firms can choose to produce with a domestic or a foreign producer. Domestic firms knew exactly of other the domestic producers which firm is of the right complementary type, while most domestic firms do not know this of foreign firms. However, some domestic firms are exogenously "tied" and they know one foreign producer of the right complementary type. Labour, necessary for the production, is immobile, leading to wage differences between the countries. It is assumed that wages are lower in the foreign country. If domestic firms can make a good match, it is more profitable for them to produce with a foreign producer. Tied firms obviously have an advantage in this, highlighting the importance of these networks.

The idea of networks playing an important role in how firms match with each other has also been applied on issues of development and how markets develop. In these models agents can choose to match on the market or to match with an agent in an exclusive bilateral relationship. Kranton (1996) sets up a model where agents are consuming one good and can choose to find this good on the market (as a result of which search costs are incurred) or to form a reciprocal relationship with another agent and consume the goods that the other agent in the reciprocal relationship is producing. The downside of reciprocal exchange is that this sometimes means that the agent does not desire the good that the other agent is producing. As long as market exchange is not very attractive, agents will exert effort and incur costs in their production for the other agent, because the punishment of not doing so (which means breaking up the reciprocal exchange and matching on the market) is too severe. However, as market exchange becomes more attractive, it becomes more difficult to sustain a reciprocal exchange relationship, because the punishment of breaking up is less severe. There is a *thick market externality*: as more agents leave their reciprocal relationships and match on the market, market exchange becomes even more attractive, meaning that more agents will switch. Other models have built on these ideas and are used to explain switches from personal economic relationships, which were the main mode of trade in the past and are still prevalent in many developing countries, to a more impersonal market exchange (e.g. Fafchamps, 2002).

3.3 Model

This section presents the model. There are two categories of agents in the model: *final consumers*, who only consume services, and *service-producing firms*, who both consume and produce services. A part of the produced services is used as intermediate inputs in the production of services: a service producer needs to consume one service to be able to produce services itself. Service producers can produce multiple services.

The central matching problem in this model is one of *quality*: there are two categories of quality on offer (high and low quality) and this quality cannot be observed in advance. There are two types of service-producing firms: *competent* firms who produce high quality products, and *inept* firms who produce low quality services. Consuming a low quality service incurs a cost θ for the consumer. There is no cost when consuming a high quality service. It is assumed that it is not possible to write a contract offering a compensation in case that the firm produces a low quality product and that other forms of insurance are also unavailable. This assumption means that even though the service consuming firm can observe the quality of the service provided, this quality cannot be observed objectively and high quality service provision is therefore not enforceable. Because of this, price cannot be dependent on quality on offer.

In this model, to circumvent this problem of asymmetric information, firms and final consumers can choose to join a network. The cost of entering the network is heterogeneous and differs across firms. Inside the network information about the characteristics of firms inside the network is freely available: a member of the network knows exactly the quality on offer provided by another member of the network. Members of the network can choose to procure a service inside the network or outside the network. If they choose to find a provider within the network, they will match randomly with a high quality provider inside the network. This is called *network matching*. The firms and final consumers outside the network consume randomly from any producing service firm, *including* those inside the network. This is called *random matching*. Given that network matching always results in a match with a high quality firm, while random matching could lead to a match with a low

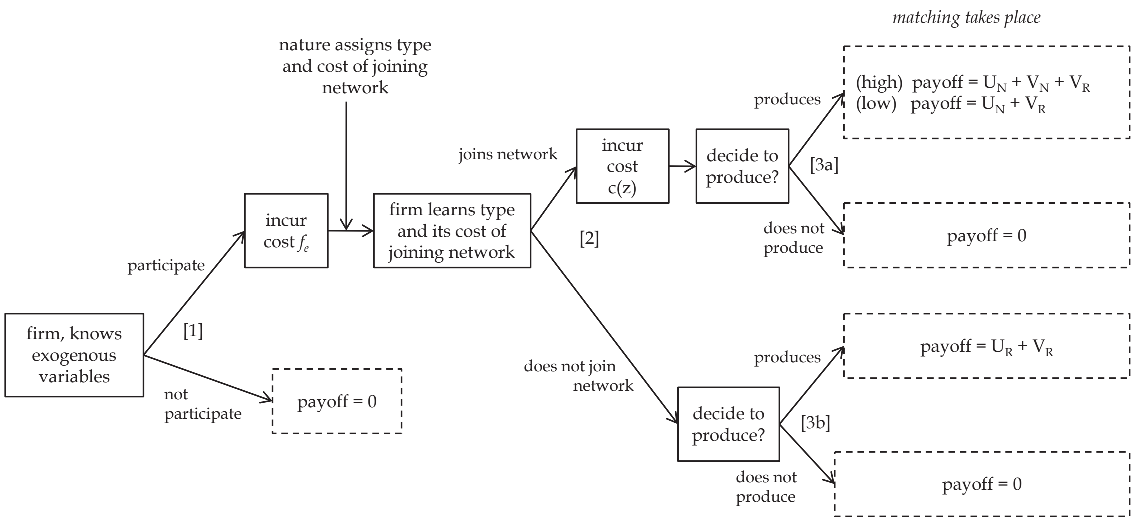

Figure 3.1: Timing and decisions of the competent and inept firms in the model. First, the firm has to choose whether to enter the game (decision [1]), then whether to join the network or not (decision [2]) and finally whether to produce (decision [3a] and [3b]).

quality firm, firms inside the network will only buy a service through network matching. However, firms inside the network will still provide services through random matching to firms outside the network.

For all members, both firms and final consumers, the network offers the benefit of being assured of high quality service consumption, which means that they do not have to incur the cost associated with consuming a low quality service. For competent firms the network provides an additional benefit, since they can benefit from a larger consumer base. The matching process between consumers and producers is described by a bargaining process similar to many labour market search models: a consumer and a potential producer match randomly (either inside or outside the network) and bargain over the price of their service. The Nash bargaining solution is used to model the outcome of this bargaining process.

3.3.1 Timing of the model

There is a fixed number of final consumers F, who will always consume one service to fulfil their needs.[2] The number of services firms is variable and will depend on several

[2]The model requires a number of outside consumers. When there are no final consumers and the cost of receiving a bad quality service is positive (i.e. $\theta > 0$), no firm outside the network will produce, because you sell fewer products than you buy when there are firms inside the network (who will only buy within the network). Incompetent firms will never join the network, because their is no market for their products (outside the network). This means that the network then only consist of competent firms inside the network. These firms will buy one product and sell one product. However, the revenue made by selling a service (p) will not outweigh the cost of production (k), the cost of procuring a service (cost p) and the cost of joining the network ($c(z)$). When facing the decision to enter at the beginning of the game, a firm will therefore choose not to enter.

conditions that will be discussed later on. Both types of firms and consumers exist in a continuum and are, for convenience, indexed separately: competent firms are indexed by z_H, inept firms by z_L and final consumers by z_F. The index z_j is used to refer to an agent in general, without specifying its type or role (so $j = H, L, F$).[3] Firms are heterogeneous in two ways: they differ in the quality of the services they produce and they differ in their cost of joining the network $c_j(z_j)$. The distribution of high and low quality firms is fixed: of the firms entering a share μ is competent, providing high quality services, and a share $1 - \mu$ is inept, providing low quality services. The timing of the model consists of the following stages:

1. **Entry stage.** There is a large supply of potential firms. These potential firms know all the exogenous variables (see Appendix Section 3.A), but do not know their own quality or their cost of joining the network. An entry cost f_e has to be incurred for a firm to be able to produce and to learn their own characteristics. The number of firms that choose to enter is S_e.

2. **Network joining stage.** All the firms that have chosen to enter incur the entry cost and learn the quality of their services (high or low) and their cost of entering the network $c_j(z_j)$ (with $j = H$ for competent firms and $j = L$ for inept firms). Also the final consumers learn their cost of entry $c_F(z_F)$. Now, the firms and final consumers have to make the choice whether they will join the network or not. The number of competent firms that choose to join is N_H, the number of inept firms that choose to join is N_L and the number of final consumers that choose to join is N_F.

3. **Decision to produce.** The firms and final consumers joining the network incur the joining cost $c_j(z_j)$ (with $j = H, L, F$, depending whether the agent is a competent firm, an inept firm or a final consumer). Firms can choose whether they will engage in the matching process and produce or to not produce at all. The number of firms producing is S_p.

4. **Matching stage.** Matching takes place. The services firms that have chosen to produce and the final consumers will need to consume one service. They look for producers, either through matching with a high quality provider inside the network (*net-*

[3]Technically, because firms and consumers exists in a continuum, it does not make sense to speak about the "number of firms" or the "number of consumers". However, the assumption of a continuum instead of discrete firms has been made to simplify the analysis; the intuition behind the model and the mechanisms are the same in the discrete version. Also, some of the graphs in this paper have been made using numerical methods, which rely on a discrete version of the model, because an analytical solution could not be found in these cases. A discussion of these numerical methods can be found in appendix section 3.D.2.

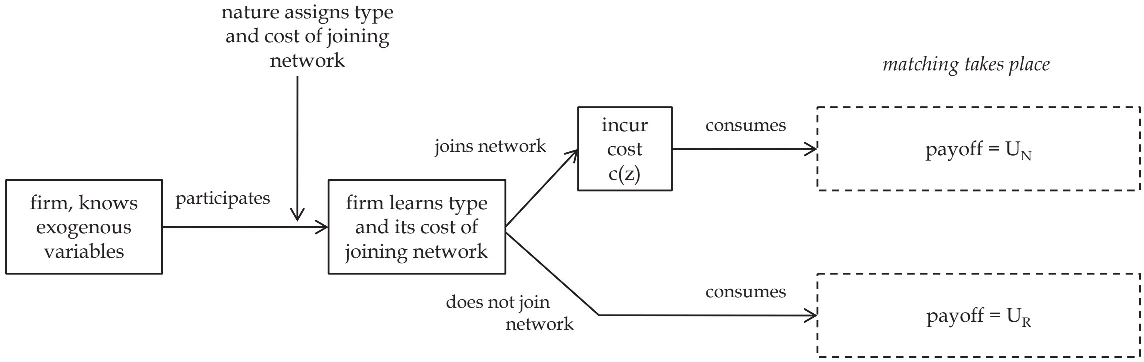

Figure 3.2: Timing and decisions of the final consumers in the model. Final consumers always consume one good and can choose whether to consume that good inside the network or outside.

Type		**Quantity**	**Provides through**	**Consumes through**
competent service firm	inside network	N_H	random matching **and** network matching	network matching
	outside network	R_H	random matching	random matching
inept service firm	inside network	N_L	random matching	network matching
	outside network	R_H	random matching	random matching
final service consumers	inside network	N_F		network matching
	outside network	R_F		random matching

Table 3.1: The matching patterns of the two types of services firms (given that they are producing) and the final consumers. Final consumers only consume services, while services firms both produce and consume services.

work matching) or matching entirely random (*random matching*). If a match is made, the consumer pays price p for the service and the producer incurs a fixed marginal cost k. (It is assumed that $p > k$).

5. **Consumption and production stage.** Firms produce their services and consumers consume their bought services simultaneously. Final payoffs are realized. If it turns out that the service consumed is of bad quality, a cost θ is incurred.

This means that there are two types of matching: *network matching* and *random matching*. Firms and final consumers inside the network will buy a service randomly from a competent firm inside the network (network matching), while firms and final consumers outside the network match randomly with *any* firm that is producing (random matching). This firm could either be inside or outside the network. As can be seen in Table 3.1, this means that competent firms inside the network will provide their services through both network matching *and* random matching.

Firms inside the network always produce, while under certain conditions some firms outside the network will choose not to produce. The body of producing firms consists therefore of the competent and inept firms inside the network (their number being respectively N_H and N_L) and the competent and inept firms outside the network that are producing (their number being respectively R_H and R_L). By definition, S_p equals the sum of producing firms:

$$S_p \equiv N_H + N_L + R_H + R_L. \tag{3.1}$$

The number of final consumers inside the network is N_F and the number of final consumers consuming outside the network is R_F. The total number of consumers F is equal to the sum of these:

$$F \equiv N_F + R_F. \tag{3.2}$$

The outcome of this matching process is determined by considering Nash equilibria of this game: given the strategy of other firms, firms cannot do better. This analysis is restricted to equilibria that are *subgame perfect*: the strategy of each firm must be a Nash equilibrium in every subgame of the game. The solution can then be found by backwards induction.

This means that we first need to analyze the payoff from the last two stages, the matching and the consumption stage. From this an expression for the expected payoffs of both types of matching is derived. Second, using these expected payoffs we can consider which firms inside and outside the network are actually going to produce. Third, we consider which firms will join the network by looking at the marginal firm that is indifferent between joining and not joining. Finally, we need to consider the number of firms entering.

3.3.2 Matching process and payoffs

As discussed earlier, firms engage in two kinds of matching: *random matching*, where the sellers are outside the network, and *network matching*, where the sellers and the buyers are inside the network. We analyze these two kinds of matching separately. To simplify the analysis, we consider the payoffs of both kinds of matching separately from a buyer's perspective and a seller's perspective.

Random matching

In random matching, the buyer can match with any firm that is producing, including the firms inside the network. The benefit for buyer depends on whether the buyer is consuming a high or low quality service. In case the service of the seller is of low quality the buyer incurs not only the price p, but also the disutility θ of receiving a low quality service (by assumption, $\theta > 0$). So the benefit to buyer z_j, $U_R(z_j)$, is given by

$$U_R(z_j) \equiv \begin{cases} -p & \text{if seller competent (high quality);} \\ -p-\theta & \text{if seller inept (low quality).} \end{cases} \tag{3.3}$$

The probability of matching an inept seller is equal to the share of inept sellers in the population of all sellers. The number of inept sellers equals $R_L + N_L$ (with N_L inept sellers inside the network and R_L inept sellers outside the network) and the number of all sellers is S_p, the number of producing firms. The expected benefit to the buyer is given by

$$\mathbb{E}\left[U_R\right] = -p - \frac{R_L + N_L}{S_p}\theta.^{4} \tag{3.4}$$

The benefit to seller z_j, $V_R(z_j)$, equals the number of buyers of its service times the benefit from each transaction $v(p)$:

$$V_R(z_j) \equiv \text{number of random matching buyers of } z_j\text{'s service} \times v(p).$$

The benefit of each transaction equals price p minus marginal cost k, so $v(p) = p - k$. Since matching happens randomly, the expected number of buyers depends on the total number of buyers through random matching divided by the number of sellers through random matching, so expected benefit equals

$$\mathbb{E}\left[V_R\right] = \frac{\text{number of random matching buyers}}{\text{number of random matching sellers}} \times v(p).$$

All firms producing will sell their services through random matching, but only firms and final consumers outside the network will buy through random matching. The number of buyers therefore equals the number of firms and final consumers outside the network, so:

$$\mathbb{E}\left[V_R\right] = \frac{R_H + R_L + R_F}{S_p}(p-k) = \frac{(S_p + F) - (N_L + N_H + N_F)}{S_p}(p-k). \tag{3.5}$$

For a final consumer outside the network, the total expected payoff is only given by the buyer's side utility $\mathbb{E}\left[U_R\right]$. For a competent or inept firm outside the network, the total expected payoff is the sum of the buyer's side utility and the seller's side utility, minus incurred cost. This total expected payoff is given by:

$$\mathbb{E}\left[\pi_R\right] = \mathbb{E}\left[U_R + V_R\right] - f_e.$$

[4] In the case that the proportion of quality of the producing firms equals the ex ante proportion of quality (this is the case when $S_e = S_p$), a share $1-\mu$ of the firms is inept and provides low quality services. In this case expected benefit to the buyer is given by

$$\mathbb{E}\left[U_R\right] = -p - (1-\mu)\theta.$$

Type	Inside network	Outside network
competent firm	$\mathbb{E}\left[\pi_N^H(z_H)\right] = \mathbb{E}\left[V_N + V_R + U_N\right] - c_H(z_H) - f_e$	$\mathbb{E}\left[\pi_R\right] = \mathbb{E}\left[V_R + U_R\right] - f_e$
inept firm	$\mathbb{E}\left[\pi_N^L(z_L)\right] = \mathbb{E}\left[V_R + U_N\right] - c_L(z_L) - f_e$	$\mathbb{E}\left[\pi_R\right] = \mathbb{E}\left[V_R + U_R\right] - f_e$
final consumer	$\mathbb{E}\left[\pi_N^F(z_F)\right] = \mathbb{E}\left[U_N\right] - c_F(z_F) - f_e$	$\mathbb{E}\left[\pi_R^F\right] = \mathbb{E}\left[U_R\right] - f_e$

Table 3.2: Expected total profits $\mathbb{E}\left[\pi(z_j)\right]$ for the two types of services firms (if they are producing) and the final consumers, including the cost of joining the network $c_j(z_j)$ and entry cost f_e.

Network matching

Buyers inside the network will only choose to consume high quality products. The only cost they are incurring is the price of the service p. This cost is given by:

$$U_N \equiv -p. \tag{3.6}$$

Only competent firms inside the network will provide services through network matching, and will do this in addition to provision through random matching. The benefit of sellers inside the network equals the number of buyers of their product times the benefit of each transaction $v(p) = p - k$:

$$V_N(z_j) \equiv \text{number of network matching buyers of } z_j\text{'s service} \times v(p).$$

Each firm inside the network has the same probability of being matched with a customer. The expected number of buyers is therefore equal to the total number of buyers inside the network divided by the number of other firms that are also selling through the network. Expected benefit from selling through the network is therefore equal to:

$$\begin{aligned}\mathbb{E}\left[V_N\right] &= \frac{\text{number of network matching buyers}}{\text{number of network matching sellers}} \times v(p) \\ &= \frac{N_H + N_L + N_F}{N_H}(p - k).\end{aligned} \tag{3.7}$$

The final payoff for a final consumer inside the network is just given by $\mathbb{E}\left[U_N\right] - c_F(z_F)$. A competent firm inside the network will provide services through both random matching and network matching, and its total expected payoff including incurred costs is therefore given by:

$$\mathbb{E}\left[\pi_N^H(z_H)\right] = \mathbb{E}\left[U_N + V_R + V_N\right] - c_H(z_H) - f_e.$$

An inept firm inside the network will consume its service through network matching, but provide its service through random matching. So its expected payoff including incurred costs is given by:

$$\mathbb{E}\left[\pi_N^L(z_L)\right] = \mathbb{E}\left[U_N + V_R\right] - c_L(z_L) - f_e.$$

Bargaining process

Prices and costs are constant in this model, resulting in a constant mark-up $p - k$ per service delivered. As discussed by Rauch (2001), differentiated products such as services often lack organized markets and reference prices. In the case of services, price setting should be seen more as the result of a search and bargain process between the buyer and seller, rather than full price competition on the market. In this section a bargaining model is described, which provides a justification for a constant price p for the services without any strategic role for prices as present in reputation models (e.g. Grossman & Horn, 1988; Shapiro, 1983).

This bargaining model is similar to what is often found in labour matching models (Rogerson et al., 2005): consumers and producers are matched randomly with each other each period and bargain over the price of the service with each other, just like employers and potential employees are matched with each other and negotiate over the wage in labour matching models. The consumer's value of the bargaining outcome is determined by their outside option, which in turn is determined by the value of the outcome of a similar bargain with another, but similar, producer, just like the negotiated wage of a potential employee is determined by his or her expectations of what a bargain with another employer would yield. For this bargaining model, the Nash bargaining axioms and the resulting Nash bargaining solution are used to describe the outcome of the bargaining process, as in Rogerson et al. (2005). It is assumed that bargaining is costly from the consumer's perspective, since the consumer needs to specify to the potential producer what kind of service it needs. This bargaining cost adds an additional trading friction and can result in prices above marginal costs.

Consumers and producers meet randomly, except for consumers inside the network: they match randomly with one of the competent firms inside the network. The outside option for the producer is to not provide the service to a client, so $u_0 = 0$.[5] The firms that have chosen to produce[6] and the final consumers each consume exactly one service. The outside option for these consumers is therefore equal to $u_0 = u(p^*) - \xi$, where ξ is a renegotiation cost and $u(p^*)$ is the outcome of a bargain with another provider (for a

[5]An important difference with many labour market matching models is that workers, because they can only produce one unit of labour, have a benefit to work for the employer that values their output the most. In this case the outside option of the worker depends on the expectation of the value of other potential employers. In this model, producers can supply to multiple consumers and for this reason this consideration is not important. The outside option now is not to produce for this particular customer, i.e. the outside option is a payoff of zero.

[6]It is assumed that if a firm has chosen to produce, it cannot withdraw from the bargaining process. Even if a firm that has chosen to produce cannot find a consumer, it must still consume one service.

consumer outside the network this other provider is randomly chosen among all firms, but for a consumer inside the network this is another competent firm inside the network).

The Nash bargaining solution maximizes

$$\max_{p} \left[v(p) - 0\right] \left[u(p) - u(p^{*}) + \xi\right],$$

where $v(p)$ is the benefit of a transaction for the producer (i.e. $v(p) = p - k$) and $u(p)$ represents the cost for the consumer, at price p (i.e. $u(p) = -p$).[7] The first order condition is

$$v'(p) \left[u(p) - u(p^{*}) + \xi\right] + v(p)u'(p) = 0.$$

The price p that results from this bargain, is the same what the consumer would get if it would bargain with another producer (since this would result in the same bargain), and should therefore equal p^{*}. This means that in the above first order condition $p = p^{*}$:

$$\begin{aligned} \xi v'(p) + v(p)u'(p) &= 0 \\ \xi - (p - k) &= 0 \\ p = p^{*} &= k + \xi \end{aligned}$$

This bargaining model provides a justification for, given constant marginal costs k, a constant price $p > k$, where the mark-up in this case is determined by the renegotiation cost ξ that the consumer has to incur every period when it starts bargaining with a new potential producer.

3.3.3 Decision to produce

Inside the network A firm inside the network will choose to engage in matching and produce if the expected profits of doing so are larger than its outside option not to produce. The cost of entry f_e and cost of joining the network $c_j(z_j)$ are sunk costs at this point. So a competent firm inside the network will produce if the expected benefits of doing so are larger than zero:

$$\mathbb{E}\left[U_N + V_R + V_N\right] \geq 0. \tag{3.8}$$

Similarly, an inept firm inside the network will produce if

$$\mathbb{E}\left[U_N + V_R\right] \geq 0. \tag{3.9}$$

[7] At this stage the bargaining cost for the consumer is a sunk cost and is therefore not included.

These expected benefits of production will depend on choices made by other firms. Firms inside the network will always produce, because they have incurred costs to join the network. A strategy where a firm is joining the network, but not producing cannot be a Nash equilibrium. So these firms always produce:

Proposition 3.1. *Firms inside the network always produce.*

See the appendix for a proof.

Outside the network The expected profits excluding incurred costs for a firm outside the network when producing are $\mathbb{E}[U_R + V_R]$. A firm outside the network will find it profitable to produce if

$$\mathbb{E}[U_R + V_R] \geq 0. \tag{3.10}$$

The more firms are selling through random matching, the lower its expected profits are. This is stated in the following proposition:

Proposition 3.2. *Expected profits for a firm outside the network excluding sunk costs* $\mathbb{E}[U_R + V_R]$ *are decreasing in the number of producing firms* S_p*, given that these expected profits are weakly positive (*$\mathbb{E}[U_R + V_R] \geq 0$*) and that* $\frac{N_H}{\mu} > \frac{N_L}{1-\mu}$.

A proof is given in the appendix. Later it will be shown that $\frac{N_H}{\mu} > \frac{N_L}{1-\mu}$, so that the second condition is satisfied. The equilibrium value of expected profits excluding sunk costs for a firm outside the network, $\mathbb{E}[U_R + V_R]$, is connected to the number of firms producing through random matching:

- If $\mathbb{E}[U_R + V_R]$ is *strictly positive* in the equilibrium, it must be that all firms that have entered are also producing ($S_p = S_e$). A situation where no firms that have entered are producing ($S_p < S_e$) and $\mathbb{E}[U_R + V_R] > 0$ cannot be a Nash equilibrium, because in this case a firm that is not producing can do better by producing.

- If $\mathbb{E}[U_R + V_R]$ is *zero* in the equilibrium, it is possible to have an equilibrium where some firms that have entered are not producing ($S_p < S_e$). Proposition 3.2 implies that if $\mathbb{E}[U_R + V_R] = 0$, a firm that is not producing cannot do better by choosing to produce, because then $\mathbb{E}[U_R + V_R]$ would decrease and be negative (and therefore lower than the outside option of not producing).

- If $\mathbb{E}[U_R + V_R]$ is *strictly* negative in the equilibrium ($\mathbb{E}[U_R + V_R] < 0$), then it must be the case that *no* firm is selling through random matching ($R \equiv R_H + R_L = 0$). If a firm would still produce through random matching in this case, then this firm would earn negative profits and could do better by not producing, which yields zero profits in the equilibrium (ignoring the sunk costs).

These conditions determine the total number of firms outside the network that are producing R, but not the *quality composition* of these firms (i.e. R_H and R_L, where $R \equiv R_H + R_L$). Suppose that $\mathbb{E}[U_R + V_R] = 0$ and not all firms that have entered are producing ($S_p < S_e$). All firms outside the network are then indifferent between producing and not producing. It is assumed that every firm, regardless of its type, has an equal probability of being one of the producing firms outside the network. Therefore, the (expected) ratio of firms that choose to produce is the same ratio as the quality composition of firms outside the network.[8] This leads to the following expected values of R_H and R_L:

Proposition 3.3. *Given that the probability of producing is the same for each firm outside the network, regardless of its quality type, the expected number of competent firms producing through random matching R_H and the expected number of inept firms producing through random matching R_L is equal to*

$$\mathbb{E}[R_H] = (\mu S_e - N_H)\frac{S_p - N_H - N_L}{S_e - N_H - N_L}, \tag{3.11a}$$

$$\mathbb{E}[R_L] = ((1-\mu) S_e - N_L)\frac{S_p - N_H - N_L}{S_e - N_H - N_L}. \tag{3.11b}$$

A derivation of these expressions is given in the appendix.

Decision to join the network

Competent firms will join the network if the expected profits from joining the network minus the cost of joining are higher than the outside option of selling through random matching or not selling at all:[9]

$$\mathbb{E}[U_N + V_N + V_R] - c_H(z_H) \geq \max\{0, \mathbb{E}[U_R + V_R]\}.$$

Similarly, inept firms will join the network if

$$\mathbb{E}[U_N + V_R] - c_L(z_L) \geq \max\{0, \mathbb{E}[U_R + V_R]\}.$$

Final consumers will always consume and therefore will join the network if

$$\mathbb{E}[U_N] - c_F(z_F) \geq \mathbb{E}[U_R].$$

[8] The ratio of the quality composition of firms outside the network is equal to $\frac{\mu S_e - N_H}{(1-\mu)S_e - N_L}$. The term in the numerator is the number of competent firms that have entered, but have not joined the network. The term in the denominator is the number of inept firms that have entered, but not have joined the network.

[9] Firms that are indifferent between joining and not joining, are assumed to join the network.

Firms and final consumers are heterogeneous in their cost of joining the network. Following Storper & Venables (2004), this cost is distributed in a linear way between $\underline{c}$ and $\overline{c}$, depending on the firm's index z_j. This leads to the following joining cost functions for respectively competent firms, indexed by z_H; inept firms, indexed by z_L; and final consumers, indexed by z_F:

$$c_H(z_H) = \underline{c} + \frac{\overline{c} - \underline{c}}{\mu S_e} \times z_H \tag{3.12a}$$

$$c_L(z_L) = \underline{c} + \frac{\overline{c} - \underline{c}}{(1-\mu) S_e} \times z_L \tag{3.12b}$$

$$c_F(z_F) = \underline{c} + \frac{\overline{c} - \underline{c}}{F} \times z_F \tag{3.12c}$$

These functions are set up in such a way that the lowest indexed firm has a joining cost of $\underline{c}$ and the highest indexed firm has a joining cost of $\overline{c}$.[10]

The number of firms and final consumers joining the network is determined by the marginal firm that is indifferent between joining and not joining the network. Firms or final consumers with a higher index than this marginal one, have a higher cost of joining the network and will not have an incentive to enter. This gives the following conditions for the number of firms and consumers joining the network (provided that $\mathbb{E}\left[V_R + U_R\right] \geq 0$):

$$N_H = \frac{\frac{N_H + N_L + N_F}{N_H}(p-k) + \frac{R_L + N_L}{S_p}\theta - \underline{c}}{\overline{c} - \underline{c}} \times \mu S_e, \tag{3.13a}$$

$$N_L = \frac{\frac{R_L + N_L}{S_p}\theta - \underline{c}}{\overline{c} - \underline{c}} \times (1-\mu) S_e, \tag{3.13b}$$

$$N_F = \frac{\frac{R_L + N_L}{S_p}\theta - \underline{c}}{\overline{c} - \underline{c}} \times F. \tag{3.13c}$$

A derivation of these (implicit) solutions is given in appendix 3.B.4. Proposition 3.4 follows from these equations:

Proposition 3.4. *The share of competent firms joining the network in the total population of competent firms is higher than the share of inept firms joining the network in the total population of inept firms, so:*

$$\frac{N_H}{\mu} > \frac{N_L}{1-\mu} \quad \textit{or, equivalently:} \quad \frac{N_L}{N_H} < \frac{1-\mu}{\mu}.$$

[10] To illustrate this, consider the joining costs of competent firms: of the total number of firms that have entered in the first stage (S_e), a share μ is competent, so there are μS_e competent firms in total. The competent firm with the lowest index $z_H = 0$ has cost of joining $c_H(0) = \underline{c}$, while the competent firm with the highest index $z_H = \mu S_e$ has a cost of joining equal to $c_H(\mu S_e) = \overline{c}$.

Decision to enter

Finally, we consider the decision to enter the game. At this point firms still do not know whether they are competent or inept and they also do not know what it will cost for them to join the network, so their expected profits will depend on their probability of being competent or inept and the associated expected profits:

$$\begin{aligned}
\mathbb{E}[\pi] &= \mu \times \left[\frac{N_H}{\mu S_e}\left(\mathbb{E}[V_R + V_N + U_N] - \hat{c}_H\right) + \frac{R_H}{\mu S_e}\left(\mathbb{E}[V_R + U_R]\right)\right] \\
&\quad + (1-\mu) \times \left[\frac{N_L}{(1-\mu)S_e}\left(\mathbb{E}[V_R + U_N] - \hat{c}_L\right) + \frac{R_L}{(1-\mu)S_e}\left(\mathbb{E}[V_R + U_R]\right)\right] \\
&\quad - f_e \\
&= \frac{N_H}{S_e}\left(\mathbb{E}[V_R + V_N + U_N] - \hat{c}_H\right) + \frac{N_L}{S_e}\left(\mathbb{E}[V_R + U_N] - \hat{c}_L\right) \\
&\quad + \frac{R_H + R_L}{S_e}\left(\mathbb{E}[V_R + U_R]\right) - f_e.
\end{aligned} \tag{3.14}$$

Here $\hat{c}_H$ and $\hat{c}_L$ are the expected joining costs for respectively a competent and an inept firm, given that it is going to join the network.[11]

In the equilibrium, expected profits must equal zero, otherwise a firm that has chosen not to enter would be better off by choosing to enter instead. This leads to the following free entry condition:

$$\mathbb{E}[\pi] = 0. \tag{3.15}$$

3.3.4 Solving the model

In the model, there are eight endogenous variables: the number of entering firms S_e, the number of producing firms S_p, the number of competent firms, inept firms and final consumers inside the network N_H, N_L and N_F, and the number of competent firms, inept firms and final consumers matching outside the network R_H, R_L and R_F. The number of entering firms is determined by equation (3.15). The number of firms and consumers that join the network is given by (3.13a–c). The number of firms that produce outside the

[11] So for a competent firm $\hat{c}_H = \mathbb{E}[c_H \mid c_H < c_H^*]$, where $c_H^* = c_H(N_H)$, the cost of the threshold competent firm that is indifferent between joining and not joining. Because $c_H(z_H)$ is linear, so:

$$\hat{c}_H = \frac{1}{N_H}\int_0^{N_H} c_H(z_H)\, dz_H = \underline{c} + \frac{\overline{c} - \underline{c}}{\mu S_e}\frac{N_H}{2}$$

Similarly, for an inept firm $\hat{c}_L = \mathbb{E}[c_L \mid c_L < c_L^*]$, where $c_L^* = c_L(N_L)$. Using the functional form:

$$\hat{c}_L = \frac{1}{N_L}\int_0^{N_L} c_L(z_L)\, dz_L = \underline{c} + \frac{\overline{c} - \underline{c}}{(1-\mu)S_e}\frac{N_L}{2}$$

network and their quality composition is determined by equations (3.11a–b) and (3.10). These equations together imply the number of producing firms S_p. Finally, the number of final consumers outside the network is determined by equation (3.2). By solving these equations, the outcome of the model can be determined.

For some parameter values it is possible to find an analytical solution, but for most parameter values numerical methods have been used to find equilibrium number of firms. This solution method relies on making the number of firms discrete instead of continuous and checking for every configuration whether a unit of firms has an incentive to deviate. Appendix 3.D.2 describes this way of finding solution in greater detail, as well as its limitations.

3.3.5 Interpretation of the model

The model incorporates elements of how businesses find trading partners. In reality, firms do not just consume a service from a random service provider, but try to get an idea of its quality in advance by considering the firms' reputation or asking around for past experiences. This can partly be done in professional organized settings, such as local business groups or chambers of commerce, but more often these interactions take place in informal settings, where social contacts play an important role. In some countries a more formal organization of networks has been documented, such as the *keiretsu* in Japan and the *chaebol* in South Korea (Rauch, 2001).

The network could be interpreted as the structure representing these informal interactions. Even though this is only a one period model, this network can be interpreted as an *arena of reputation sharing*: other members of the network have consumed services from this particular firm before and share this information with other firms inside the network, without any strategic considerations (i.e. no calumniation of high quality producers). This concept of a network follows the definition given by Rauch (2001), based on Granovetter (1973): the network is a "set of actors who know each other's relevant characteristics or can learn them through referral." Networks like these provide a bidirectional benefit for competent firms: from the buyer's perspective, the network allows the firm to consume high quality services, but also from a seller's perspective it provides a larger consumer base.

The model can be applied to a wide range of services which can be characterized by a consumer-provider relation, for example legal services, accountancy or management consultancy, for which the quality of the delivered service is often crucial to the consumer.

The heterogeneous cost of joining the network can be given multiple interpretations. The first interpretation is that some firms differ in their "networking skills" and therefore joining the network is more costly for some firms than for others. Firms consist of people and several studies have showed that networking patterns differ between individuals.[12] The second interpretation is one of increasing costs as the network grows: it is less costly to join a small network than a large network. So the first firms that join incur less costs of joining the network than the last firms that join. However, this is not a model of subsequent entry, but of simultaneous entry, and therefore this interpretation might fit the model less well.

The final consumers can be regarded as other actors in the economy who are also consuming services, such as other firms (e.g. producers of manufacturing), the government or households. Just like inept firms, who will not sell through the network, their main interest of joining the network is to match with higher quality suppliers.

The cost θ of receiving a bad quality service is the loss in payoff or (monetary) utility, which could just represent the dislike of people to receive bad quality services. However, in case of consuming services firms it can be seen as the additional cost that the firm needs to incur to make up for the bad quality, so that it is still able to produce its own good.

3.3.6 Relation to other models

The model is based on ideas that can be found in other matching models. However, the model differs from other models in several ways in order to incorporate some specific features of services exchange.

First of all, the model allows firms to *provide multiple services* and not just one. In the labour market referral model of Montgomery (1991) workers can only work for one employer. Current employees can refer new potential employees to their employer and because of the assumed inbreeding bias, employers will prefer to hire a new employee through referral if the current employee is of high ability. Potential workers with a higher ability are therefore more likely to be found through referral and therefore the workers that are left over (i.e. not referred) will consist of more than proportionally lower able workers, making non-referral hiring even less attractive. This *lemons effect* is not present in this model: if a producing firm is matched through network matching, it can still be matched through random matching. The introduction of a network does not lower the benefits of

[12]This feature has been documented in different types of networks, for example co-authorship data from economics journals (Goyal et al., 2006) or prison data on friendships between inmates (see also Jackson & Rogers, 2007; Jackson, 2008).

random matching, as was the case in the model of Montgomery (1991): in that model, when hiring through referral becomes more attractive, this will negatively affect those not hiring through referral. However, here there are no capacity constraints: in principle a firm can provide as many services as it can find consumers for.[13]

Second, the model acknowledges the *particular nature of the customer-producer relation*: most of the time only one party benefits from increased quality and the other party will only benefit if incentives are set correctly. Some models assume a symmetric relation between agents in the economy (e.g. Storper & Venables, 2004), where agents jointly undertake a project and benefit equally from the achieved quality. This setting may work well in an intrafirm interaction, but is less appropriate in this setting of service provision.

Furthermore, other models allow for a possibility for low quality service providers to provide high quality by increasing their investment or effort level (e.g. Grossman & Horn, 1988; Storper & Venables, 2004), which creates a richer environment, but also makes the network structure more complicated.

In this model the term *network* is used to describe the information sharing structure. The set-up of the network is limited compared to other models: a firm is either inside the network or outside, while other models emphasize the internal network structure (e.g. Kranton & Minehart, 2001). These models describe the exact structure of how information is shared, specifying every relationship between nodes and considering strategic interactions between them. This model does not focus on how the network is organized on the level of nodes: it does not describe the relation between two nodes inside the network or how links between nodes form. While in other models some degree in connectedness is allowed, in this model connectedness is *binary*: a firm is (completely) connected to the network or not connected at all.

Jackson & Rogers (2007) provide some justification for this binary structure in their discussion of the structure of social networks. They show that most social networks can be characterized by having a substantial group of highly-connected nodes as well as a substantial group of low-connected nodes, leading to "fat tails" in the distribution of the degree of connectedness of the nodes. Nodes with a medium degree of connectedness are relatively underrepresented. The binary structure of the model can be seen as a further simplification of this.

[13]Of course, in the case firms face capacity constraints, they have to make choices who they will serve and who not. However, in this model there is no increased benefit from a producer perspective to provide to someone inside the network compared to providing to a customer outside the network: profits from a sale are the same.

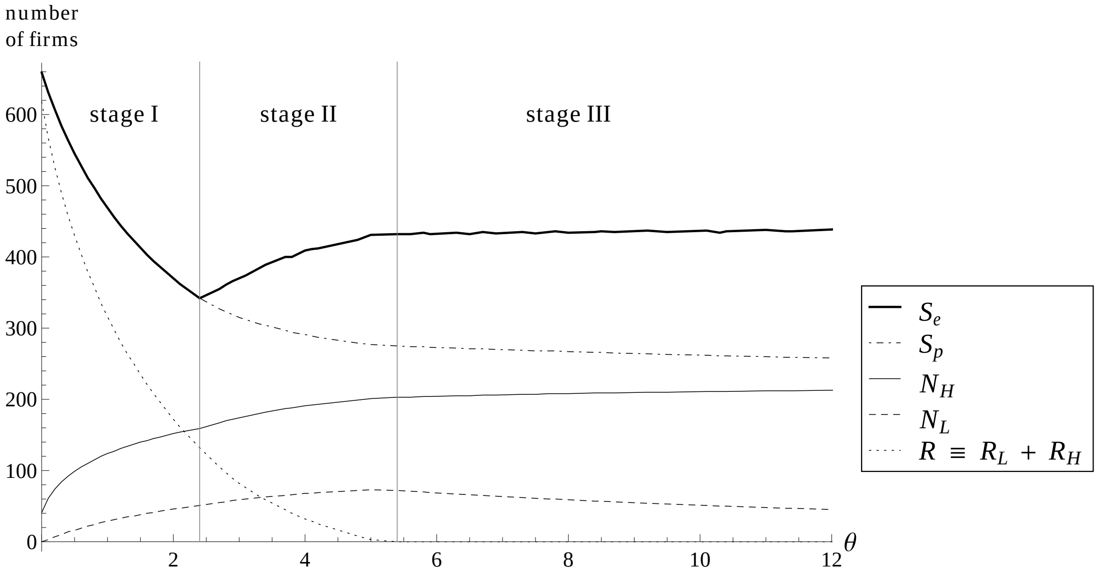

Figure 3.3: The effect of a rise in the cost of receiving a bad service θ on the number of entering firms S_e, the number of producing firms S_p, the number of competent firms inside the network N_H, the number of inept firms inside the network N_L and the number of firms producing through random matching $R \equiv R_H + R_L$.

3.4 Filtering and locking out

In this section the effects of changes in parameters on the outcome of the model are analyzed. Because the model is set-up for only one-period, these changes are not caused by dynamic effects: there is no process of adjustment from one outcome to another; it only shows how outcomes differ if the initial parameters are adjusted. This section focuses especially on the effect of a rise in the cost of receiving a bad quality service θ on the network structure and the consequences for welfare.

This network structure leads to two effects when θ increases. First, the network causes sorting in quality: if receiving bad quality services becomes too costly, fewer inept firms will find it profitable to produce, improving the quality composition of the economy, filtering out the lower quality firms. This is the *filtering effect*. Second, firms use networks not only as a way of improving the quality of their consumption, but also as a way of coordinating supply and demand: it provides an implicit mechanism for competent firms inside the network to agree to only consume each other's services, while still providing to firms outside the network. This means that firms outside the network receive fewer customers than they would have received without a network. This *lock-out effect* bears a cost, for low values of θ fewer firms will enter and produce than in the case of no networking.

Exogenous variables		Standard value
F	Number of final consumers	2000
μ	Share of competent (high quality type) firms in pool of firms	0.5
θ	Cost of consuming a bad quality service	1.0
p	Price of consuming a service	1.0
k	Marginal cost of providing a service	0.5
$\underline{c}$	Lowest cost of joining the network (i.e. if $z_j = 0$)	0.0
$\overline{c}$	Highest cost of joining the network	4.0
f_e	Entry cost for firm at first stage	1.0

Table 3.3: The exogenous variables of the model and the standard values used in the simulations.

3.4.1 The filtering effect of the network

One of the crucial parameters in the model is the cost of consuming a bad service θ. As θ increases, firms consuming through random matching will incur higher costs in expectation and therefore joining the network will become more attractive. This in turn affects expected profits and therefore also the number of firms entering S_e and the number of firms producing S_p. For increasing values of θ three "stages" can be identified in the outcome of the model: the first stage where all firms that have entered are producing, the second stage where some firms outside the network are not producing and the final stage where no firms outside the network are producing. Figure 3.3 shows the impact of a rise of θ on the number of firms entering, producing and joining the network.

In the first stage, all firms that have entered are producing, so $S_e = S_p$. The rise in θ means that consuming a service outside the network becomes more costly and therefore joining the network becomes more attractive. This leads to an increase in the number of firms joining the network. If the number of entering firms S_e is kept constant, the increase in θ will have a negative impact on expected profits in the first stage $\mathbb{E}[\pi]$, because firms buying outside the network incur higher costs when receiving a bad quality service and the average cost of joining the network $\hat{c}_j$ increases, because a larger proportion of the firm joins the network. As a result of the free entry condition (3.15), fewer firms will choose to enter and this will mean a decline in S_e as θ rises, as can be seen in figure 3.3. Figure 3.4 shows the effect of a rise in θ on profits of the firms. Since all firms outside the network are producing, it must be the case that expected profits outside the network excluding sunk costs, $\mathbb{E}[U_R + V_R]$, are positive.

The quality composition of the economy remains constant in this stage, because all firms that have entered are producing: a share μ of the firms that have entered are competent and so a share μ of the producing firms are competent as well. Because of the decrease in S_e as θ increases, the absolute number of competent and inept firms decreases as well.

As can be seen in figure 3.4b, at a sufficiently high value of θ expected profits for a firm outside the network excluding sunk costs, $\mathbb{E}[U_R + V_R]$, become zero. For higher values of θ this means that now some firms that have chosen to enter will not produce, as is stated in the following proposition:

Proposition 3.5. *For sufficiently high values of θ, it must be the case that the expected profits for a firm outside the network excluding sunk costs are nonpositive ($\mathbb{E}[U_R + V_R] \leq 0$) and that some firms that have entered are not producing ($S_p < S_e$).*

In the second stage, there are still some firms producing and condition (3.10) must be binding: $\mathbb{E}[U_R + V_R] = 0$. The firms outside the network consist disproportionately of inept firms, because a competent firm is more likely to have joined the network than an inept firm (see proposition 3.3). If fewer firms outside the network choose to produce, this also means that inept firms are more likely to not produce. In this situation the presence of this network structure improves the quality composition of the economy: proportionally there will be *fewer* inept firms producing than competent firms. This is the *filtering effect* of the network:

Definition. *The **filtering effect** occurs when the ratio of inept firms to competent firms producing is lower than the ratio of inept firms to competent firms entering, so:*

$$\frac{N_L + R_L}{N_H + R_H} < \frac{1-\mu}{\mu}.$$

Proposition 3.6. *If not all firms that have entered are producing ($S_p < S_e$), then the filtering effect occurs.*

A proof is given in the appendix. From proposition 3.5 and 3.6 we can conclude:

Corollary 3.7. *For sufficiently high values of θ the filtering effect occurs.*

In figure 3.5b it can be seen that from the point where $\mathbb{E}[U_R + V_R] = 0$, the share of producing competent firms increases, while the number of producing inept firms decreases. This filtering effect means that the negative effect of a rise of θ on expected buyer's profit $\mathbb{E}[U_R]$ is counterbalanced: fewer inept firms producing means that from a seller's point of view it is less likely to match an inept firm, even though θ is higher. This can be seen in figure 3.4a, where the decrease in expected buyer's profit outside the network $\mathbb{E}[U_R]$ in the second stage ($\theta \in [2.5, 5.0]$) is small.

The number of firms entering increases as well, even though the cost of receiving a bad service θ increases. To see why, fix the number of entering firms at $S_e{}'$ such that expected

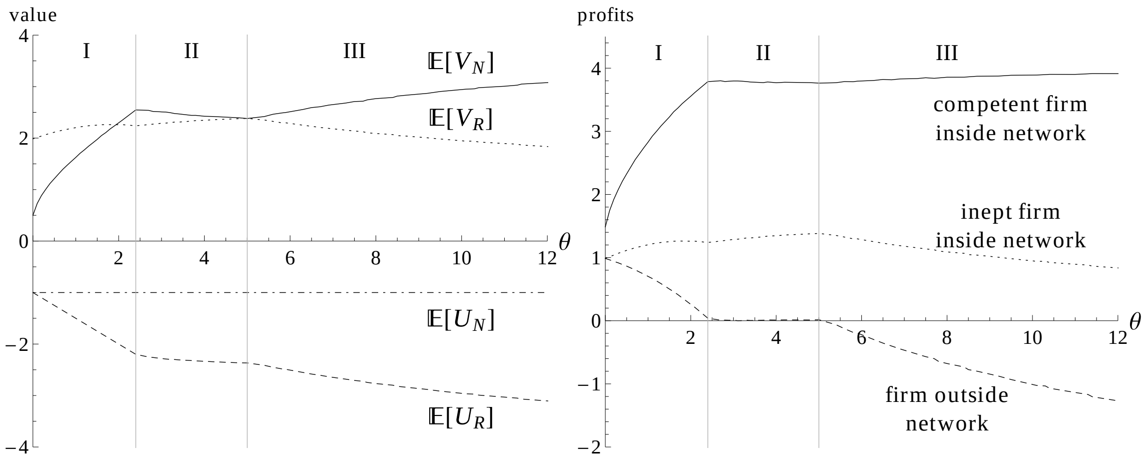

Figure 3.4: (a) *Left:* The expected utility from selling through network matching $\mathbb{E}[V_N]$ and through random matching $\mathbb{E}[V_R]$ and the expected utility from buying inside the network $\mathbb{E}[U_N]$ (constant) and outside the network $\mathbb{E}[U_R]$. (b) *Right:* Expected profits for a producing competent firm inside the network (equal to $\mathbb{E}[V_N + V_R + U_N]$), a producing inept firm inside the network (equal to $\mathbb{E}[V_R + U_N]$) and a producing firm outside the network (equal to $\mathbb{E}[V_R + U_R]$). Excludes entry cost f_e and cost of joining the network $c_j(z_j)$.

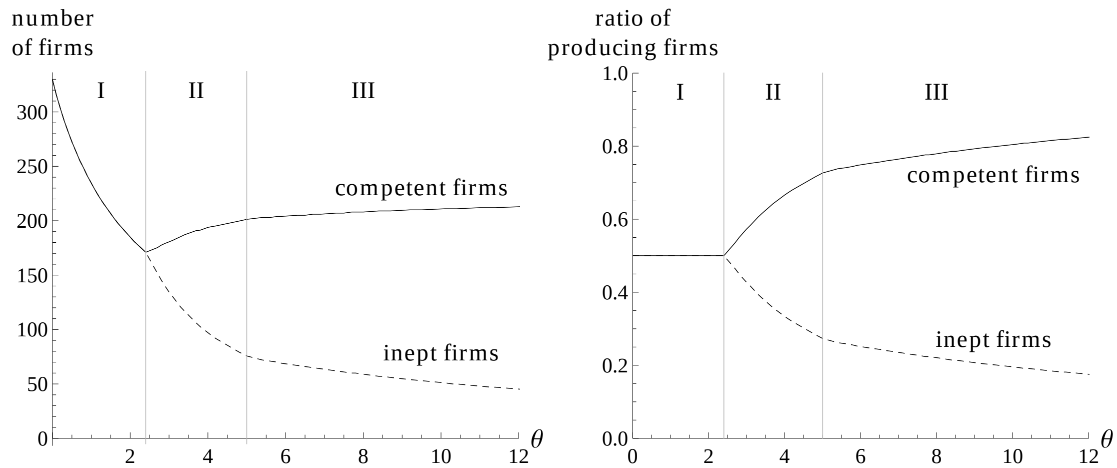

Figure 3.5: (a) *Left:* The number of competent firms producing ($R_H + N_H$) and the number of inept firms producing ($R_L + N_L$). (b) *Right:* The share of producing competent firms in the group of producing firms and the share of producing inept firms in this group.

profits are zero ($\mathbb{E}[\pi] = 0$) for a certain value $\theta = \theta'$ and let θ increase by a little bit (i.e. $\theta = \theta' + \delta$, with $\delta > 0$). As derived in equation (3.14), expected profits are determined by the expected profits of competent and inept firms inside the network as well as the expected profits of firms outside the network:

$$\mathbb{E}[\pi] = \frac{N_H}{S_e}\mathbb{E}\left[\pi_N^H\right] + \frac{N_L}{S_e}\mathbb{E}\left[\pi_N^L\right] + \frac{R_L + R_H}{S_e}\mathbb{E}[\pi_R].$$

In the first stage, as θ increases, the profits of a competent firm inside the network $\mathbb{E}\left[\pi_N^H\right]$ increase (more clientele inside the network), while $\mathbb{E}[U_R + V_R]$ decreases and so the profits of firms outside the network $\mathbb{E}[\pi_R]$ decrease as well. This decrease has a stronger effect on expected profits than the increase of expected profits for a competent firm inside the network. For fixed values of ${S_e}'$, this leads to lower expected profits , so $\mathbb{E}[\pi] < 0$, which means fewer firms will enter if we allow S_e to vary.

However, in the second stage $\mathbb{E}[U_R + V_R]$ is equal to zero and so $\mathbb{E}[\pi_R]$ does no longer decrease as θ increases (keeping S_e fixed again). Because $\mathbb{E}\left[\pi_N^H\right]$ still increases, expected profits $\mathbb{E}[\pi]$ will be larger than zero at the fixed value ${S_e}'$. If we allow S_e to vary again, than the number of firms entering S_e will increase such that $\mathbb{E}[\pi] = 0$. So more firms will find it beneficial to incur the entry costs and will enter, even though not all firms that have entered will produce. However, the prospect of possibly being a profitable competent firm inside the network will encourage them to enter, causing a rise in S_e as θ increases.

In the third stage, *no* firms outside the network are producing, only the firms inside the network are producing. If a firm outside the network would choose to produce, it would earn negative profits ($\mathbb{E}[U_R + V_R] < 0$) and therefore no firm outside the network will choose to produce. However, inept firms *inside* the network will still provide services to the final consumers outside the network, but as θ rises, the number of final consumer outside the network is dwindling as well, causing fewer inept firms to join the network and produce (so N_L decreases).[14] This means that the filtering effect of the network continues as θ increases, as can be seen in figure 3.5.

3.4.2 Strategic market protection: the lock-out effect

An interesting case to analyze is when the cost of consuming a bad service θ approaches zero. In this case very few final consumers and inept firms will join the network. However, under certain conditions competent firms will still join the network, even though the informational benefit of joining the network is small:

Proposition 3.8. *Even when the cost of receiving a bad service θ approaches zero, competent firms will join the network if $(p - k) > \underline{c}$. The limit from above of the number of competent firms inside the network is positive in this case:*

$$\lim_{\theta \downarrow 0} N_H > 0 \qquad \textit{if } (p - k) > \underline{c}.$$

As θ approaches zero, receiving a bad quality service is not very costly and it is clear that the network yields very little informational benefits. However, for competent firms it is still beneficial to join the network if $(p - k) > \underline{c}$. The reason for this is as follows: all firms are equally likely to provide a service to a consumer or firm outside the network. However, competent firms inside the network have an additional consumer base: the firms inside the

[14]Depending on the cost structure of entering the network, even at high values of θ there will be consumers outside the network.

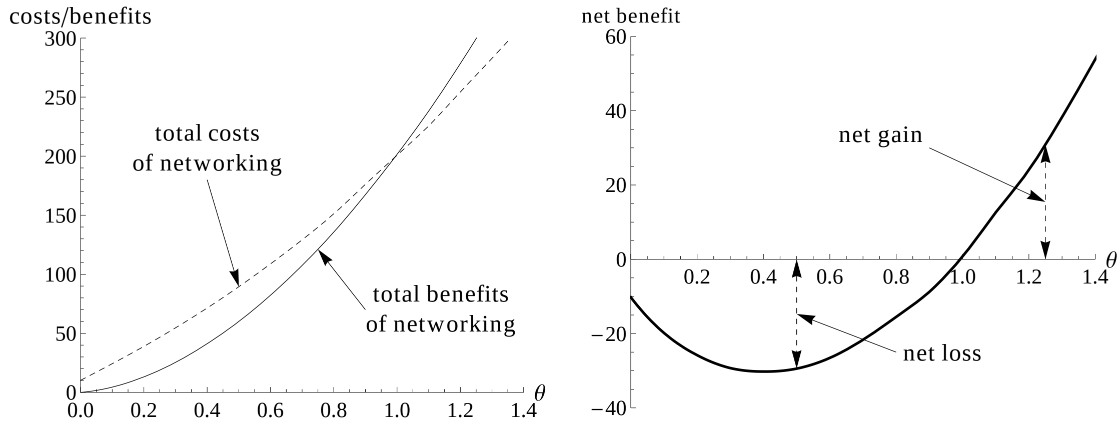

Figure 3.6: (a) *Left:* Total costs incurred to join the network and the total informational benefits of networking for low θ values. (b) *Right:* Difference in costs and benefits.

network. In this setting the competent firms that join the network are implicitly agreeing to consume each other's services, while still providing services to firms and consumers outside the network, just like every other firm. By jointly doing this they increase their number of customers at the expense of the firms outside the network, who now supply to fewer customers. In this way, networks "lock out" a part of the market, even though there is little informational benefit in doing so. This is the lock-out effect:

Definition. *The **lock-out effect** occurs when the total costs incurred to join the network are higher than the total benefits of networking.*

Total costs of joining the network are the sum of the costs that final consumers and firms incur to join the network, so:

$$\text{total costs} = \int_0^{N_H} c_H(z_H)\, dz_H + \int_0^{N_L} c_L(z_L)\, dz_L + \int_0^{N_F} c_F(z_F)\, dz_F.$$

The total benefits of joining the network are equal to the sum of benefits that final consumers and firms enjoy from networking (i.e. the expected costs of receiving a bad service that firms inside the network are now avoid by joining) and the benefits from the filtering effect for firms and consumers outside the network. So, when filtering does not occur total benefits are equal to

$$\text{total benefits} = \left(\mathbb{E}\left[U_N - U_R\right]\right) \times \left(N_H + N_L + N_F\right) = \frac{R_L + N_L}{S_p}\theta \times \left(N_H + N_L + N_F\right).$$

The lock-out effect only involves a *redistribution* of customers: the gain in clientele and profits by competent firms inside the network is equal to the loss in clientele and profits by

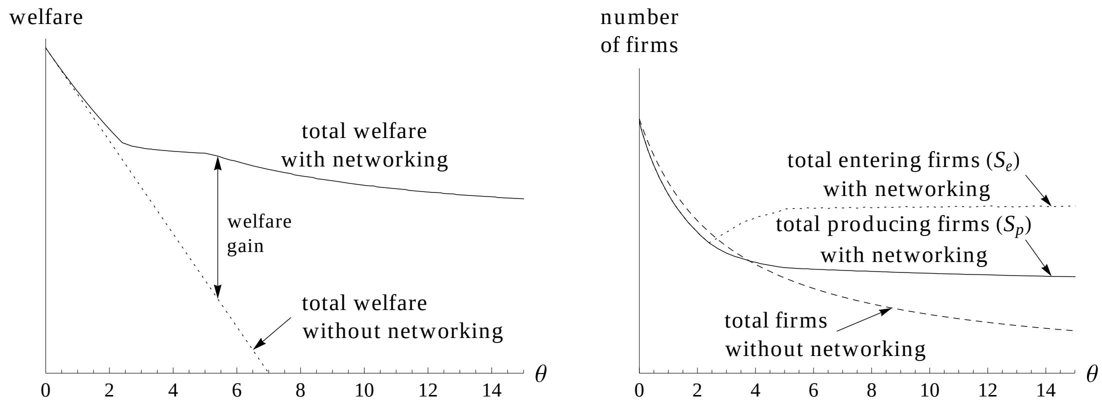

Figure 3.7: (a) *Left:* Welfare in a situation without network formation and the situation with network formation. (b) *Right:* The number of firms entering S_e and producing S_p with network formation and the number of firms without network formation (in the latter case all firms that have entered are also producing).

firms outside the network. For this reason this redistribution does *not* affect total benefits of networking.

Figure 3.6 shows the total costs and benefits of networking for low values of θ. It shows that for low values of θ the lock-out effect occurs, because the costs of joining the network outweigh the informational benefits.[15]

This lock-out effect has not been present in previous models that relied on the assumption that agents can only provide one service.[16] In that case it would not make sense for firms to join the network at low values of θ, since this would not have increased their payoffs. By making the assumption that firms can supply multiple services, this means that firms will also care about the number of clients they serve and will use the network structure to their benefit to achieve this.

3.4.3 Networks and welfare

Networks improve the information available to firms and consumers and prohibit its members from consuming bad quality services. In this way networks can improve welfare. However, the network itself is not costless: firms must incur a cost to join the network (e.g., a social cost), which on itself decreases welfare. Welfare is determined by adding up

[15] At $\theta = 0$ multiple equilibria can occur, because firms and consumers inside the network are indifferent between random matching and network matching. An equilibrium can occur where a group of firms joins the network and then only chooses to procure services through the network. Another equilibrium is one where no firms join the network and subsequently all services are procured through random matching.

[16] In previous models this assumption of providing only service was often the basis of which sorting occurs. For example in Montgomery (1991), if finding workers through referrals becomes more attractive, more high quality agents will be referred, meaning that fewer high quality agents are left over in the labour market outside of referrals. This is the so-called *lemons effect*, referring to the adverse selection model of Akerlof (1970).

all the profits of the firms and the buyer's side utilities of the final consumers:

$$W \equiv \int_0^{N_H} \pi_N^H(z_H)\,dz_H + \int_{N_H}^{\mu S_e} \pi_R(z_H)\,dz_H + \int_0^{N_L} \pi_N^L(z_L)\,dz_L + \int_{N_L}^{(1-\mu)S_e} \pi_R(z_L)\,dz_L$$
$$+ \int_0^{N_F} U_N - c_F(z_F)\,dz_F + \int_{N_F}^{F} U_R(z_F)\,dz_F, \tag{3.16}$$

where the profit terms $\pi_N^H(z_H)$, $\pi_N^L(z_L)$ and $\pi_R(z_j)$ take the sunk costs into account.

Expected welfare is the sum of all profits made by the firms and the expected buyer's side utilities of the final consumers:

$$\mathbb{E}[W] = S_e \times \mathbb{E}[\pi] + N_F \times (\mathbb{E}[U_N] - \hat{c}_F) + (F - N_F) \times \mathbb{E}[U_R]. \tag{3.17}$$

In this equation $\hat{c}_F$ is the average joining cost of final consumers. In the model the first term in the above equation is zero, because $\mathbb{E}[\pi] = 0$ by the free entry condition for firms (3.15). Welfare is therefore completely determined by the buyer's side utilities of the final consumer.

To answer to what extent the availability of a network actually improves welfare, the case where all firms and consumers engage in random matching (i.e. it is not possible to join the network) is used as a benchmark. Expected welfare in the benchmark model is determined by the expected profits of firms and final consumers. All firms in the benchmark model that have entered, are also producing ($S_e = S_p$). For a final consumer in the benchmark model, the expected payoff is determined by the price and the expected quality they get:

$$\mathbb{E}\left[\pi_F^B\right] = \mathbb{E}[U_R] = -p - (1-\mu)\theta.$$

Expected total profits for a firm (regardless of their type) also include the expected payoff from sales and the entry cost that has to be incurred:

$$\mathbb{E}\left[\pi_j^B\right] = -p - (1-\mu)\theta + \frac{S_p + F}{S_p}(p - k) - f_e \quad \text{(with } j = H, L\text{)}.$$

Expected welfare in the benchmark model is determined by adding up these profits for both the services producers and the final consumers:

$$\mathbb{E}\left[W^B\right] = \mathbb{E}\left[\pi_F^B\right] \times F + \mathbb{E}\left[\pi_{H,L}^B\right] \times S_e.$$

In the case of free entry, expected profits of the service producers $\mathbb{E}\left[\pi_{H,L}^B\right]$ must be zero. By comparing this benchmark model with the networking model, it can be shown that networking increases expected welfare:

Proposition 3.9. *If the free entry condition (3.15) holds, then*

$$\mathbb{E}[W] = \mathbb{E}\left[W^B\right] \qquad \text{if } \theta \leq \tfrac{c}{1-\mu},$$
$$\mathbb{E}[W] > \mathbb{E}\left[W^B\right] \qquad \text{if } \theta > \tfrac{c}{1-\mu}.$$

A proof is given in the appendix. Final consumers and inept firms will join the network if $\theta > \frac{c}{1-\mu}$ and as is shown in this proposition, in this case networks are welfare improving compared to the benchmark model. This can also be seen in figure 3.7a. This effect is enhanced for higher values of θ, when the filtering effect is present. Welfare decreases as θ increases, but the filtering effect mitigates this decrease: it causes "bad apples" to be taken out, which implies that the quality of matches improves.

The lock-out effect only affects firms and therefore does not affect welfare in the free entry case. However, the lock-out effect does have an effect on the number of producing firms: in figure 3.7b it is shown that the number of firms for low θ values is actually *lower* with the presence of networks than without networking. The number of producing firms is directly related to the *volume* of transactions, because every firm consumes one service and therefore requires one transaction. This is a consequence of the lock-out effect of the network: because a group of firms have (implicitly) agreed to consume each other's services, this redistribution of customers means that for fewer firms outside the network it is feasible to produce. This leads to a lower number of producing and entering firms. For high values of θ the lock-out effect is outweighed by the informational benefits and the resulting filtering effect, which means that more firms will enter and produce.

If the free entry condition is dropped and instead the number of entering firms S_e is kept *fixed*, welfare must now incorporate the profits of firms and the effects of networking. The lock-out effect causes in this case an inefficiency for low values of θ: compared to a situation where the same number of firms enters and networking is not possible, welfare is actually *lower* with networking.[17] However, compared to the welfare gains caused by the filtering effect, these welfare losses are small.

Of the two effects discussed in this paragraph, the filtering effect is the most influential on the welfare gains caused by networking. Networking not only improves information, which on its own is welfare improving, but it also causes a selection effect, which has a large positive impact on welfare.

[17] See appendix section 3.C for a graph showing this result.

3.5 Conclusion and discussion

This chapter analyzed the consequences of reputational business networks, such as business or trading associations, on the provision of low and high quality products or services. From the viewpoint of the final consumers, networks like these always improve welfare compared to a benchmark model where no networking can take place. This effect is further enhanced when networking pushes some incompetent firms out of business: an increase in networking means that fewer customers will consume through the regular market, reducing its profitability. This *filtering effect* of the network is important when identifying the benefits of networking.

However, networking also creates side effects: even when the informational benefits of joining the network are low, competent firms will have an incentive to join the network, as this provides an implicit agreement to consume each other's services, while still providing services outside the network. Other firms outside the network will now have fewer customers. This redistribution of customers results in only very few improvements in the quality of matching and is costly: the costs incurred to join the network are lower than the benefits of the improved quality of matching. This *lock-out effect* is a direct consequence of the reciprocal structure of the network, where firms are both producers and consumers. Without the double role of firms as both producers and consumers, this effect would not have been present.

The model relied on several assumptions, which do not necessarily hold for all transactions. First of all, the quality of the product or service is not contractible. This is often the case for services, which by their nature are intangible and therefore it is often difficult to draw up complete contracts that can fully specify the requirements that a specific service must fulfil. Second of all, firms in the network are both producers and consumers. This is often the case for service firms, who often have to rely on the consumption of other service providers to be able to produce their own services. For example, a financial institution needs to rely on law firms for legal advice and law firms need to rely on financial institutions for their banking needs. This creates a *reciprocal element* and reinforces the importance of networks: networks are not only used to improve the quality of consumption, but also as a way to increase one's clientele. If a firm joins the network to be able to consume higher quality services, this firm will also use the network to sell its services, given that its services are of good quality.

Face-to-face contacts are often seen as crucial in the transmission of such information (Storper & Venables, 2004) and the need for physical proximity means that such networks are often bound geographically. Given the more local nature of these networks, it is also more likely that service provision will be more locally oriented. Furthermore, the model predicts that in places where the cost of networking is low, the average quality that firms provide will increase.

Cities provide one such environment. Given the smaller distances, it is easier and less expensive to form networks. Empirical studies have indeed shown that cities have a relatively higher proportion of workers of high ability (Glaeser & Resseger, 2010). Assuming that better skilled workers produce better quality goods and services, cities will have relatively more high quality service providers. However, too tight trading networks can also be detrimental for outsider firms. The formal organized networks as seen in some Asian countries, such as the *keiretsu*, where a lot of trade is contained within a business agglomerate, make it difficult for new firms to enter, essentially locking them out.

A question is how valuable the information on quality is and whether a firm or consumer could sell this off to another firm. If a firm or consumer could credibly sell this information, there would be scope for this, because this reduces the expected cost of receiving a bad quality service. However, the credible selling of information could also signal whether the firm or consumer is inside the network. Given that firms inside the network tend to be of a higher quality, it might be beneficial for a consumer to not buy the information from such firm, but to directly procure the service from the firm selling information. This could result in interesting dynamics, especially if consumers can sell information on particular suppliers to each other as well. The model in this Chapter assumes that no credible information on quality can flow outside the network and therefore precludes such dynamics.

The model is limited in its set-up: it does not, for example, describe how the actual information sharing takes place or the strategic interactions behind it. Moreover, the model does not provide a full specification of the internal structure of the network itself: it is assumed that everyone in the network shares information with everyone else. The network is a binary design: firms and consumers are either a member or not, while it is more likely that in reality almost everyone is connected to someone else in one way or another, but that the *degree* of connectivity will vary across individuals. Theories of distributed networks, where some agents in a network collect information and spread this around, and theories of information dissemination in networks might provide starting points that could be used to consider these aspects in further detail.

Another topic not covered is the geographical scope of networks. There is no reason in the model why networks are solely attached to a specific geographic unit. In practice, some networks might cross geographical units, as is the case, for example, with ethnic networks between countries (discussed in further detail in Rauch, 2001; Rauch & Casella, 2003) or multinational firms. Reputations are nowadays not only locally oriented, but can also spread globally, implying that the role of local networking might decline and a more global form of network might result from this. However, even though new communication technologies might facilitate the formation of these networks, networking on a global scale will remain more costly to firms than networking on a local scale, which often relies heavily on face-to-face interactions. This would predict that local networks will remain important in the future and that the "proximity burden" should be seen as a characteristic of trade in services that will continue to exist.

Acknowledgements This chapter was part of my MPhil thesis, which was supervised by Tony Venables. I am grateful for his advice and guidance. I also would like to thank the participants of the Trade Workshop as well as the Firms & Development Group for helpful feedback and comments. I would like to thank Michael Koelle for further comments on this paper. Furthermore, I would like to thank two anonymous examiners for their comments and suggestions. All remaining errors are my own.

Appendices of Chapter 3

3.A List of variables

Exogenous variables	
F	Number of final consumers
μ	Share of competent (high quality type) firms in pool of firms
$1-\mu$	Share of inept (low quality type) firms in pool of firms
θ	Cost of consuming a bad quality service
p	Price of consuming a service
k	Marginal cost of providing a service
$c_j(z_j)$	Cost of joining the network for firm or consumer with index z_j with $j = H, L, F$ H competent (high quality) firms L inept (low quality) firms F final consumers
$\underline{c}$	Lowest cost of joining the network (i.e. if $z_j = 0$)
$\overline{c}$	Highest cost of joining the network
f_e	Entry cost for firm at first stage (before learning type or joining cost $c_j(z_j)$)
Endogenous variables	
S_e	Number of firms entered after the first stage
S_p	Number of firms actually producing
N_H	Number of competent (H-type) producing firms inside the network
N_L	Number of inept (L-type) producing firms inside the network
N_F	Number of final consumers inside the network
R_H	Number of competent (H-type) producing firms outside the network
R_L	Number of inept (L-type) producing firms outside the network
R_F	Number of final consumers outside the network
Other	
U_R	Cost to buyer when consuming through random matching
U_N	Cost to buyer when consuming through network matching
V_R	Benefit to seller when providing through random matching
V_N	Benefit to seller when providing through network matching
$u(p)$	Cost to buyer from consuming a good
$v(p)$	Benefit to seller from selling a good
ξ	Renegotiation cost
$\pi_R(z_j)$	Total profits of firm z_j ($j = H, L$) outside the network (including sunk costs)
$\pi_N^H(z_j)$	Total profits of competent firm z_H inside the network (including sunk costs)
$\pi_N^L(z_j)$	Total profits of inept firm z_L inside the network (including sunk costs)
$\pi(z_j)$	Total profits of firm z_j ($j = H, L$)
W	Welfare (of the networking model)
W^B	Welfare in the random matching benchmark model (without networks)

3.B Proofs and derivations

3.B.1 Proposition 3.1: Firms inside the network always produce

Proof (by contradiction). Consider the case where there is an inept firm that has chosen to join the network, but not to produce. Assume that this is a Nash equilibrium. Furthermore, assume that $c_j(z_j) > 0$. This means that the payoff of not producing must be larger or equal to the payoff when it would produce, so

$$\mathbb{E}\left[V_R + U_N\right] \leq 0.$$

Noting that the behaviour of other firms and consumers can be fully determined, because the cost of joining is determinate ($c_j(z_j)$ is exactly known and not drawn from a distribution), this can be anticipated in advance. This means that if this firm has chosen to enter, it must be that

$$\mathbb{E}\left[V_R + U_N\right] - c_L(z_L) \geq \max\left\{0, \mathbb{E}\left[V_R + U_N\right]\right\}$$

for the firm to make the decision to join the network. Since $\mathbb{E}\left[V_R + U_R\right] \leq \mathbb{E}\left[V_R + U_N\right] \leq 0$, it must be that the right hand side is zero. This is a contradiction, since the left hand side is smaller than zero. □

The reasoning for competent firms is similar. The assumption that $c_j(z_j) > 0$ does not hold in cases where $\underline{c} = 0$, since $c_j(0) = 0$ in these cases. However, this only occurs for this particular firm of infinitesimal size and is therefore ignored in the exposition.

3.B.2 Proposition 3.2: Expected profits outside network decreasing in S_p

This section shows that $\mathbb{E}\left[V_R + U_R\right]$ (given that it is positive) is decreasing in S_p, taking all the other parameters (except for R_L and R_H) as given. A change in S_p means a change in the number of firms selling through random matching. This includes both competent and inept firms selling through random matching, their number being respectively R_H and R_L. It is assumed that a competent and an inept firm are affected in the same way when S_p increases (i.e. an inept firm and a competent firm outside the network have an equal chance to be one of the producers, see proposition 3.3). However, since the number of competent and inept firms selling through random matching might not be equal, the effect of the change on R_H can differ from the effect of the change on R_L.

Proof. $\mathbb{E}[V_R + U_R]$ is given by:

$$\begin{aligned}
\mathbb{E}[V_R + U_R] &= \frac{S_p + F - (N_L + N_H + N_F)}{S_p}(p-k) - p - \frac{R_L + N_L}{S_p}\theta \\
&= (p-k) + \frac{F - (N_L + N_H + N_F)}{S_p}(p-k) - p - \frac{R_L + N_L}{S_p}\theta \\
&= \frac{F - (N_L + N_H + N_F)}{S_p}(p-k) - \frac{R_L + N_L}{S_p} - k. \qquad (3.18)
\end{aligned}$$

It is given that this is (weakly) positive, $\mathbb{E}[V_R + U_R] \geq 0$. We proceed by showing that (i) the first term is strictly decreasing in S_p and (ii) the second term is weakly decreasing in S_p. These two conditions imply that $\mathbb{E}[V_R + U_R]$ is strictly decreasing in S_p, which is what we want to prove.

Condition (i) is trivially satisfied if $F - (N_L + N_H + N_F) > 0$. Because S_p is in the denominator of the first term in (3.18), and the terms in the numerator are positive (note that $p > k$), an increase in S_p means a decrease of the whole fraction. We show that $F - (N_L + N_H + N_F) > 0$ by assuming the opposite and showing that this lead to a contradiction. So assume $F - (N_L + N_H + N_F) \leq 0$. In this case it is clear from equation (3.18) that $\mathbb{E}[V_R + U_R]$ must be strictly negative. This is a contradiction, so $F - (N_L + N_H + N_F) > 0$ must hold and condition (i) is satisfied.

The proof that condition (ii) is satisfied is less trivial and relies on the fact that competent firms are more likely to join the network than inept firms, and for this reason $N_H > \frac{\mu}{1-\mu}N_L$.[18] By proposition 3.3, the number of competent and inept firms is given by

$$\begin{aligned}
R_H &= (\mu S_e - N_H)\frac{S_p - N_H - N_L}{S_e - N_H - N_L}, \\
R_L &= ((1-\mu)S_e - N_L)\frac{S_p - N_H - N_L}{S_e - N_H - N_L}.^{19}
\end{aligned}$$

So:

$$\begin{aligned}
\frac{R_L + N_L}{S_p}\theta &= \frac{\theta}{S_p}\left[((1-\mu)S_e - N_L)\frac{S_p - N_H - N_L}{S_e - N_H - N_L} + N_L\right] \\
&= \frac{(1-\mu)S_e - N_L}{S_e - N_H - N_L}\theta - \frac{\theta}{S_p}\left[\frac{(1-\mu)S_e - N_L}{S_e - N_H - N_L} \times (N_H + N_L) - N_L\right]. \qquad (3.19)
\end{aligned}$$

This term is increasing in S_p (and therefore the second term in equation (3.18) is decreasing) if the term between brackets in the last equation is positive. So we need to show that

$$\frac{(1-\mu)S_e - N_L}{S_e - N_H - N_L}(N_H + N_L) - N_L > 0.$$

[18] This can be seen by first considering the situation that that competent firms are equally likely to join the network as inept firms. In this case $N_H = \alpha \times \mu S_e$ and $N_L = \alpha \times (1-\mu)S_e$, where α is the share of competent and inept firms joining the network. It follows that $N_H = \frac{\mu}{1-\mu}N_L$. If competent firms are strictly more likely to join the network, it must be that $N_H > \frac{\mu}{1-\mu}N_L$.

This holds if

$$\begin{aligned}
\left[\frac{(1-\mu)S_e - N_L}{S_e - N_H - N_L}\right] N_H + \left[\frac{(1-\mu)S_e - N_L}{S_e - N_H - N_L} - 1\right] N_L > 0 \\
\Leftrightarrow \left[\frac{(1-\mu)S_e - N_L}{S_e - N_H - N_L}\right] N_H + \left[\frac{-\mu S_e + N_H}{S_e - N_H - N_L}\right] N_L > 0 \\
\Leftrightarrow \left[(1-\mu)S_e - N_L\right] N_H + \left[-\mu S_e + N_H\right] N_L > 0 \\
\Leftrightarrow (1-\mu)S_e N_H - \mu S_e N_L > 0 \\
\Leftrightarrow (1-\mu)N_H > \mu N_L \\
\Leftrightarrow N_H > \frac{\mu}{1-\mu} N_L.
\end{aligned}$$

It is assumed that $N_H > \frac{\mu}{1-\mu} N_L$ indeed holds, which implies that the last term in brackets in equation (3.19) is indeed positive.

This proves that given that when $\mathbb{E}\left[V_R + U_R\right]$ is positive (for which $F > N_H + N_L + N_F$ is a necessary condition) and that competent firms are more likely to join the network (so $N_H > \frac{\mu}{1-\mu} N_L$), expected profits $\mathbb{E}\left[V_R + U_R\right]$ are decreasing in S_p. □

3.B.3 Proposition 3.3: Quality composition of firms outside the network

Each firm outside the network, whether competent or inept, is equally likely to be one of the producing firms. There are in total $S_e - N_H - N_L$ firms that have entered outside the network. The firms that have joined the network will produce (see proposition 3.1), so the total number of firms outside the network that produce is equal to $R = S_p - N_H - N_L$, where $R \equiv R_H + R_L$.

So the probability that a firm is going to produce is equal to

$$\text{Probability of producing} = \frac{S_p - N_H - N_L}{S_e - N_H - N_L}.$$

The number of competent firms outside the network is equal to the number of competent firms that have entered (μS_e), minus the number of competent firms that have joined the network (N_H). So the expected number of competent firms producing outside the network $\mathbb{E}\left[R_H\right]$ is equal to:

$$\mathbb{E}\left[R_H\right] = (\mu S_e - N_H)\frac{S_p - N_H - N_L}{S_e - N_H - N_L}. \tag{3.20a}$$

Similarly, the expected number of inept firms producing outside the network $\mathbb{E}\left[R_L\right]$ is equal to:

$$\mathbb{E}\left[R_L\right] = ((1-\mu)S_e - N_L)\frac{S_p - N_H - N_L}{S_e - N_H - N_L}. \tag{3.20b}$$

3.B.4 Derivation of network size

Each firm (either competent, with type H, or inept, with type L) and consumer (designated with F) is heterogeneous in its cost of joining the network. As given by equations (3.12a), (3.12b) and (3.12c), the costs of joining for a specific firm or consumer (indexed by z_j, where $j = H, L, F$) are

$$\begin{aligned} c_H(z_H) &= \underline{c} + \frac{\overline{c} - \underline{c}}{\mu S_e} \times z_H, \\ c_L(z_L) &= \underline{c} + \frac{\overline{c} - \underline{c}}{(1-\mu) S_e} \times z_L \\ \text{and} \quad c_F(z_F) &= \underline{c} + \frac{\overline{c} - \underline{c}}{F} \times z_F. \end{aligned}$$

The number of firms and final consumers joining the network is determined by the marginal firm z_j^* that is indifferent between joining and not joining the network. Given the number of firms and consumers in the network, a firm or consumer with $z_j > z_j^*$ will not join the network, since in this case the cost of joining outweigh the expected benefits when it would join the network.

Note that the values of N_j are bounded above and below: it value must be larger or equal to zero ($N_j \geq 0$) and smaller or equal to the total group that could potentially join the network (so $N_H \leq \mu Se$, $N_L \leq (1-\mu)Se$ and $N_F \leq F$). In case one of these conditions are binding, the indifference conditions derived here will not hold necessarily, since the marginal firm or consumer can now be strictly better off in network matching.

In the case of final consumers, consider the consumer that is indifferent between being inside the network and being outside and solve for its index z_F:

$$\begin{aligned} \mathbb{E}\left[U_N\right] - c_F(z_F) &= \mathbb{E}\left[U_R\right] \\ -p - c_F(z) &= -p - \frac{R_L + N_L}{S_p}\theta \\ c_F(z_F) = \underline{c} + \frac{\overline{c} - \underline{c}}{F} z_F &= \frac{R_L + N_L}{S_p}\theta \\ \Rightarrow \quad N_F = z_F &= \frac{\frac{R_L + N_L}{S_p}\theta - \underline{c}}{\overline{c} - \underline{c}} \times F. \end{aligned} \tag{3.21}$$

Suppose that $\mathbb{E}\left[U_R + V_R\right] \geq 0$. The derivation of the number of inept firms inside the network is similar:

$$\begin{aligned} \mathbb{E}\left[U_N\right] + \mathbb{E}\left[V_R\right] - c_L(z_L) &= \mathbb{E}\left[U_R\right] + \mathbb{E}\left[V_R\right] \\ \Rightarrow \quad N_L = z_L &= \frac{\frac{R_L + N_L}{S_p}\theta - \underline{c}}{\overline{c} - \underline{c}} \times (1-\mu) S_e. \end{aligned} \tag{3.22}$$

This condition can be solved explicitly (provided that $N_L < (1-\mu)S_e$, i.e. there are still some inept firms not inside the network):

$$N_L = \frac{\underline{c}S_p - R_L\theta}{(1-\mu)\theta S_e - S_p(\overline{c}-\underline{c})}(1-\mu)S_e. \tag{3.23}$$

For competent firms the expression is a bit more complicated, since the value of being inside the network will also depend on the number of firms inside the network. In the case that $\mathbb{E}\left[U_R + V_R\right] \geq 0$:

$$\begin{aligned}
\mathbb{E}\left[U_N\right] + \mathbb{E}\left[V_N\right] + \mathbb{E}\left[V_R\right] - c_H(z_H) &= \mathbb{E}\left[U_R\right] + \mathbb{E}\left[V_R\right] \\
-p + \frac{N_H + N_L + N_F}{N_H}(p-k) - c_H(z_H) &= -p - \frac{R_L + N_L}{S_p}\theta \\
c_H(z_H) = \underline{c} + \frac{\overline{c}-\underline{c}}{\mu S_e} z_H &= \frac{R_L + N_L}{S_p}\theta + \frac{N_H + N_L + N_F}{N_H}(p-k)
\end{aligned}$$

$$\Rightarrow \quad N_H = z_H = \frac{\frac{N_H + N_L + N_F}{N_H}(p-k) + \frac{R_L + N_L}{S_p}\theta - \underline{c}}{\overline{c}-\underline{c}} \times \mu S_e. \tag{3.24}$$

A closed form solution of this condition for N_H can be calculated, yielding two solutions, of which one is negative. This calculation and the final result are omitted for its length.

In the case that $\mathbb{E}\left[U_R + V_R\right] < 0$, the value of not joining the network is zero for competent and inept firms. The expressions are in this case determined by analyzing the firm that is indifferent between joining the network and producing or not joining the network and earning a payoff of zero. The expression for the final consumer remains the same.

3.B.5 Proposition 3.4: Number of firms inside the network

A simple comparison between (3.22) and (3.24) shows that

$$\frac{N_H}{\mu} > \frac{N_L}{(1-\mu)},$$

which is the result of proposition 3.4.

This implies that the ratio of competent firms inside the network to the general "population" of competent firms that have entered is higher than the ratio of inept firms inside the network to the "population" of inept firms that have entered. The networking rate of competent firms is higher than the networking rate of inept firms. In other words, a competent firm is more likely to join the network than an inept firm.

3.B.6 Proposition 3.5: Not all firms are producing for certain θ

Proof (by contradiction). Assume that all firms that have entered are also producing (i.e. $S_e = S_p$) for all values of θ. This implies that $\mathbb{E}\left[U_R + V_R\right] \geq 0$ or $R_H = R_L = 0$ (the number of producing firms outside the network is zero). If some firms outside the network are still producing (i.e. $R_H + R_L > 0$) and expected profits are negative ($\mathbb{E}\left[U_R + V_R\right] < 0$), a firm would have been better off by not producing and earning zero profits instead of the negative profits it is earning right now. It is shown that this assumption leads to a contradiction and should therefore be rejected.

Since all firms that have entered are also producing, this means that the share of inept firms in the population of producing firms, $\frac{R_L+N_L}{S_p}$, equals the share of inept firms in the population of entering firms, $1-\mu$. This means that $\frac{R_L+N_L}{S_p}\theta = (1-\mu)\theta$ and the ratio of firms and final consumers joining the network is equal to:

$$\frac{N_H}{\mu S_e} = \frac{\frac{N_H+N_L+N_F}{N_H}(p-k)}{\overline{c}-\underline{c}} + \frac{(1-\mu)\theta-\underline{c}}{\overline{c}-\underline{c}}, \tag{3.13a}$$

$$\frac{N_L}{(1-\mu)S_e} = \frac{(1-\mu)\theta-\underline{c}}{\overline{c}-\underline{c}}, \tag{3.13b}$$

$$\frac{N_F}{F} = \frac{(1-\mu)\theta-\underline{c}}{\overline{c}-\underline{c}}. \tag{3.13c}$$

As θ increases, these ratios increase as well, until for a certain value of θ they will equal their maximum possible value of one. Define θ^* as the lowest value of θ such that all firms and consumers are inside a network: so $N_H = \mu S_e$, $N_L = (1-\mu)S_e$ and $N_F = F$. This is the case when:

$$\frac{(1-\mu)\theta-\underline{c}}{\overline{c}-\underline{c}} = 1 \quad \Rightarrow \quad \theta^* = \frac{\overline{c}}{1-\mu}.$$

At $\theta \geq \theta^*$, expected profits from matching outside the network $\mathbb{E}\left[V_R + U_R\right]$ are negative:

$$\begin{aligned}\mathbb{E}\left[V_R + U_R\right] &= \frac{(S_p+F)-(N_L+N_H+N_F)}{S_p}(p-k) - p - (1-\mu)\theta \\ &= -p-(1-\mu)\theta < 0.\end{aligned}$$

This is a contradiction. Clearly, a firm that is now producing and earning negative profits can do better by not producing and earning a payoff of zero. So we reject the assumption that all firms that have entered are also producing and conclude that $S_p < S_e$ for sufficiently high θ. □

Using continuity of the solutions, we can also prove that there exist values $\theta < \theta^*$ where not all firms are producing. Note that the endogenous variables are continuous in

θ: taking limits from the left ($\lim_{\theta \uparrow \hat{\theta}}$) and right ($\lim_{\theta \downarrow \hat{\theta}}$) near a point $\hat{\theta}$ does not alter the equations or the outcome, on the domain $\theta \in (0, \infty)$.

Now assume again that all firms that have entered are producing (i.e. $S_p = S_e$) and consider a value of θ that is just smaller than θ^* , so take $\theta = \theta^* - \epsilon$, with $\epsilon > 0$. It must hold that $N_L < (1-\mu)S_e$ at this value of θ and therefore $R_L > 0$. This means that expected profits of an inept firm outside the network must at least as large as the outside option of earning zero profits, so $\mathbb{E}\left[V_R + U_R\right] \geq 0$. Continuity ensures that there is a range of values of $\epsilon > 0$ such that $\mathbb{E}\left[V_R + U_R\right] < 0$, because when $\epsilon = 0$, $\mathbb{E}\left[V_R + U_R\right] < 0$.

This means that there is a range values with lower bound $\theta < \theta^*$ for which $R_L > 0$ and $\mathbb{E}\left[V_R + U_R\right] < 0$, which is a contradiction: an inept firm that is producing is better off by not producing and this can therefore not be an equilibrium.

3.B.7 Proposition 3.6: Filtering effect

Proof. If not all firms that have entered are producing, then $S_p < S_e$. By proposition 3.4, we have:

$$\frac{N_H}{\mu} > \frac{N_L}{1-\mu}.$$

By proposition 3.3, the number of competent and inept firms outside the network is given by:

$$\begin{aligned} R_H &= (\mu S_e - N_H)\frac{S_p - N_H - N_L}{S_e - N_H - N_L}, \\ R_L &= ((1-\mu)S_e - N_L)\frac{S_p - N_H - N_L}{S_e - N_H - N_L}. \end{aligned}$$

This means that $N_H + R_H$ is equal to:

$$\begin{aligned} N_H + R_H &= N_H \times \left[1 - \frac{S_p - N_H - N_L}{S_e - N_H - N_L}\right] + \mu S_e \\ &= N_H \times \left[\frac{S_e - N_H - N_L - (S_p - N_H - N_L)}{S_e - N_H - N_L}\right] + \mu S_e \\ &= N_H \times \underbrace{\left[\frac{S_e - S_p}{S_e - N_H - N_L}\right]}_{>0} + \mu S_e. \end{aligned}$$

The calculation of $N_L + R_L$ is similar:

$$N_L + R_L = N_L \times \left[\frac{S_e - S_p}{S_e - N_H - N_L}\right] + (1-\mu)S_e.$$

Dividing $N_H + R_H$ by μ, and using that $\frac{N_H}{\mu} > \frac{N_L}{1-\mu}$:

$$\begin{aligned} \frac{N_H + R_H}{\mu} &= \frac{N_H}{\mu} \times \left[\frac{S_e - S_p}{S_e - N_H - N_L}\right] + S_e \\ &> \frac{N_L}{1-\mu} \times \left[\frac{S_e - S_p}{S_e - N_H - N_L}\right] + S_e = \frac{N_L + R_L}{1-\mu}. \end{aligned}$$

Rearranging this gives:

$$\frac{N_H + R_H}{\mu} > \frac{N_L + R_L}{1-\mu} \quad \Leftrightarrow \quad \frac{N_L + R_L}{N_H + R_H} < \frac{1-\mu}{\mu}.$$

□

3.B.8 Corollary 3.7: Filtering effect

This corollary follows directly from proposition 3.5 and 3.6.

3.B.9 Proposition 3.8: Size of the network for low θ values

When the cost of receiving a bad quality service approaches zero ($\theta \downarrow 0$), then there will be less of an incentive to join the network, and fewer firms will do so. However, competent firms will still have an incentive to join the network, even if θ is very close to zero.

When θ is very close to zero, both final consumers and inept firms will not join the network (in the limit, the number of firms joining the network is zero):

$$\lim_{\theta\downarrow 0} N_F = \lim_{\theta\downarrow 0} \max\left\{0, \frac{\frac{R_L+N_L}{S_p}\theta - \underline{c}}{\overline{c} - \underline{c}}\right\} = \max\left\{0, \underbrace{\frac{-\underline{c}}{\overline{c} - \underline{c}}}_{\leq 0} F\right\} = 0.$$

Using the explicit form for N_L (3.23), its limit value can be calculated:

$$\begin{aligned}\lim_{\theta\downarrow 0} N_L &= \lim_{\theta\downarrow 0} \max\left\{0, \frac{\underline{c}S_p - R_L\theta}{(1-\mu)\theta S_e - S_p(\overline{c} - \underline{c})}(1-\mu)S_e\right\} \\ &= \max\left\{0, \underbrace{\frac{-\underline{c}}{\overline{c} - \underline{c}}}_{\leq 0}(1-\mu)S_e\right\} = 0.\end{aligned}$$

Provided that its limit value $\in [0, \mu S_e]$, the number of competent firms inside the network N_H if $\theta \downarrow 0$ can be calculated.

$$\begin{aligned}\lim_{\theta\downarrow 0} N_H &= \lim_{\theta\downarrow 0} \frac{\frac{N_H+N_L+N_F}{N_H}(p-k) + \frac{R_L+N_L}{S_p}\theta - \underline{c}}{\overline{c} - \underline{c}}\mu S_e \\ &= \frac{(1 + \frac{\lim_{\theta\downarrow 0} N_L + \lim_{\theta\downarrow 0} N_F}{\lim_{\theta\downarrow 0} N_H}(p-k) - \underline{c}}{\overline{c} - \underline{c}}\mu S_e \\ &= \frac{(p-k) - \underline{c}}{\overline{c} - \underline{c}}\mu S_e.\end{aligned}$$

This value is not necessarily equal to zero. This implies that if $p - k > \underline{c}$ competent firms will still join the network for very small θ. For $(p-k) \geq \overline{c}$ all competent firms will join the network.

3.B.10 Proposition 3.9: Networks improve welfare

Proof. Expected welfare in the benchmark model is given by

$$\mathbb{E}\left[W^B\right] = \mathbb{E}\left[\pi_F^B\right] \times F + \underbrace{\mathbb{E}\left[\pi_{H,L}^B\right]}_{=0} \times S_e = \mathbb{E}\left[\pi_F^B\right] \times F = \int_0^F \mathbb{E}\left[U_R\right]\, dz_F,$$

since $\mathbb{E}\left[\pi_{H,L}^B\right] = 0$ by the zero profit condition. Expected welfare in the networking model is given by

$$\mathbb{E}\left[W\right] = \int_0^{N_F} U_N - c_F(z_F)\, dz_F + \int_{N_F}^F \mathbb{E}\left[U_R\right]\, dz_F.$$

If $\theta \leq \frac{c}{1-\mu}$ no final consumers are joining the network (see equation (3.21)) and welfare remains the same, so $\mathbb{E}\left[W\right] = \mathbb{E}\left[W^B\right]$. If $\theta > \frac{c}{1-\mu}$, final consumers will join the network. The consumer with index $z_F = N_F$ is indifferent between joining and not joining (i.e. $\mathbb{E}\left[U_N\right] - c_F(N_H) = \mathbb{E}\left[U_R\right]$). Consumers with index $z_F < N_F$ are strictly better off in expectation by joining the network, so for these values of z_H it must hold that

$$U_N - c_F(z_F) > \mathbb{E}\left[U_R\right] \quad (\text{if } z_F < N_F).$$

This implies:

$$\begin{aligned}
& \int_0^{N_F} U_N - c_F(z_F)\, dz_F > \int_0^{N_F} \mathbb{E}\left[U_R\right] dz_F \\
\Rightarrow \quad & \int_0^{N_F} U_N - c_F(z_F)\, dz_F + \int_{N_F}^F \mathbb{E}\left[U_R\right] dz_F > \int_0^{N_F} \mathbb{E}\left[U_R\right] dz_F + \int_{N_F}^F \mathbb{E}\left[U_R\right] dz_F \\
\Rightarrow \quad & \int_0^{N_F} U_N - c_F(z_F)\, dz_F + \int_{N_F}^F \mathbb{E}\left[U_R\right] dz_F > \int_0^F \mathbb{E}\left[U_R\right] dz_F.
\end{aligned}$$

The left hand term is equal to $\mathbb{E}\left[W\right]$, the right hand side is equal to $\mathbb{E}\left[W^B\right]$, so:

$$\mathbb{E}\left[W\right] > \mathbb{E}\left[W^B\right].$$

□

3.C Welfare with fixed entry

Without free entry condition (3.15), but instead a fixed number of entrants, expected profits of firms will no longer be zero, but vary, depending on the value of θ. Welfare, as defined by (3.16) will in this case incorporate profits. The figure below shows the difference in

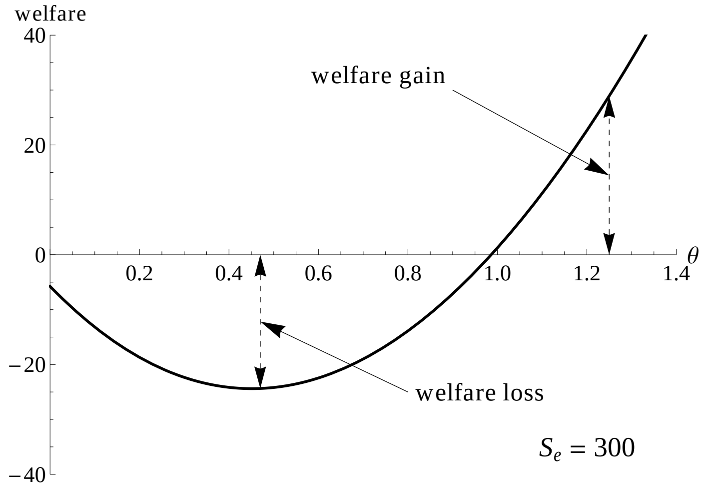

Figure 3.8: Difference in welfare between networking model and benchmark model with entry fixed at $S_e = 300$.

welfare between a model with only random matching and the networking model, both with the number of entrants S_e fixed at 300 firms:

(Note the similarity with figure 3.6b. However, these figures depict different situations, in the above figure S_e is held constant, while in figure 3.6b S_e will vary.)

3.D Solving the model

3.D.1 Overview of equations

The number of entering firms S_e is determined by the free entry condition (3.15):

$$\begin{aligned}\mathbb{E}[\pi] =& \frac{N_H}{S_e}\left(\mathbb{E}[V_R + V_N + U_N] - \hat{c}_H\right) + \frac{N_L}{S_e}\left(\mathbb{E}[V_R + U_N] - \hat{c}_L\right) \\ &+ \frac{R_H + R_L}{S_e}\left(\mathbb{E}[V_R + U_R]\right) - f_e = 0. \end{aligned} \tag{3.15}$$

If $\mathbb{E}[U_R + V_R] \geq 0$, the number of firms and consumers inside the network is given by:

$$N_H = \frac{\frac{N_H + N_L + N_F}{N_H}(p - k) + \frac{R_L + N_L}{S_p}\theta - \underline{c}}{\overline{c} - \underline{c}} \times \mu S_e, \tag{3.13a}$$

$$N_L = \frac{\frac{R_L + N_L}{S_p}\theta - \underline{c}}{\overline{c} - \underline{c}} \times (1 - \mu) S_e, \tag{3.13b}$$

$$N_F = \frac{\frac{R_L + N_L}{S_p}\theta - \underline{c}}{\overline{c} - \underline{c}} \times F. \tag{3.13c}$$

These values are bounded above and below (i.e. $N_H \in [0, \mu S_e]$, $N_L \in [0, (1-\mu)S_e]$ and $N_F \in [0, F]$). If $\mathbb{E}[U_R + V_R] < 0$ different equations will apply for N_H and N_L, such that N_H and N_L represent the firm that is indifferent between joining the network and not joining the network (and not producing).

All final consumers will consume a good, either through random matching or through network matching. The number of consumers engaging in random matching equals R_F, the number of consumers engaging in network matching equals N_F and their sum should equal the number of final consumers F:

$$F = R_H + N_H$$

$$\mathbb{E}\left[R_H\right] = (\mu S_e - N_H)\frac{S_p - N_H - N_L}{S_e - N_H - N_L} \tag{3.20a}$$

$$\mathbb{E}\left[R_L\right] = ((1-\mu)S_e - N_L)\frac{S_p - N_H - N_L}{S_e - N_H - N_L} \tag{3.20b}$$

$$\mathbb{E}\left[U_R + V_R\right] = -p - \frac{R_L + N_L}{S_p}\theta + \frac{(S_p + F) - (N_L + N_H + N_F)}{S_p}(p-k) \geq 0 \tag{3.10}$$

For certain parameter values, especially in the case where all the firms that have entered are also producing, this method can be used to solve analytically for a solution. The analytical solutions used in this research have been calculated either by hand or by applying the Solve-function of the Wolfram Mathematica software package. The Solve-function automatically reduces a set of equations and present a set of possible solutions. Solutions with negative values (i.e. a negative number of firms) or unrealistic values (e.g. the number of competent firms inside the network N_H being larger than the number of competent firms in total) have been discarded.

3.D.2 Numerical methods

For certain parameter values solving this model by analytical methods becomes very tricky or impossible. In these cases a program script written in Python has been used to find a numerical solution to the problem. This script checks a set of possible solutions against the equilibrium conditions and checks whether any firm has a reason to deviate from the equilibrium or not. If no firms have an incentive to deviate from the equilibrium, then this set of values constitute a solution.

In these cases the analysis is discrete instead of continuous, as is the case in the analytical: the unit of analysis is now a firm of size δ, instead of a firm of infinitesimal small size, as was the case in the analytical model. This results of the numerical simulation have been verified with the results of the analytical solutions (or if available, an approximation of this solution).

To simplify notation, define the expected profits in the last stage as (excluding sunk costs):

$$\hat{\pi}_N^H = \mathbb{E}\left[V_N + V_R + U_N\right],$$
$$\hat{\pi}_N^L = \mathbb{E}\left[V_R + U_N\right],$$
$$\hat{\pi}_R = \mathbb{E}\left[V_R + U_R\right].$$

As discussed in section 3.3.1, there are three stages in the game: first, the decision to enter, second, the decision to join the network and third, the decision to produce. By backwards induction, this leads to the following conditions (in reverse order):

- The decision to produce

  1. Inside the network: is it beneficial for a *competent* firm to produce? Or would a firm be better off not to produce and earn zero profits? The solution should satisfy
  $$\hat{\pi}_N^H(N_H, N_L, N_F, S_e, S_p) \geq 0.$$
  2. Inside the network: is it beneficial for an *inept* firm to produce? Or would a firm be better off not to produce? The solution should satisfy
  $$\hat{\pi}_N^L(N_H, N_L, N_F, S_e, S_p) \geq 0.$$
  3. Outside the network: is it beneficial for a firm (competent or inept) to produce? Or would it be better off not to produce? The solution should satisfy
  $$\hat{\pi}_R(N_H, N_L, N_F, S_e, S_p) \geq 0.$$
  4. For a firm not producing: would it be beneficial to that firm to have chosen to produce instead? If a firm that has entered, but is not producing would produce, profits must be lower than zero, otherwise this firm would be better off to produce:
  $$\hat{\pi}_R(N_H, N_L, N_F, S_e, S_p + \delta) < 0.$$

- The decision to join the network

  1. The marginal competent firm inside the network must not be better off by being outside the network:
  $$\hat{\pi}_N^H(N_H, N_L, N_F, S_e, S_p) - c_H(N_H) \geq \max\left\{0, \hat{\pi}_R(N_H - \delta, N_L, N_F, S_e, S_p)\right\}.$$

2. The marginal competent firm outside the network must not be better off by being inside the network:

$$\hat{\pi}_N^H(N_H+\delta, N_L, N_F, S_e, S_p) - c_H(N_H+\delta) \leq \max\{0, \hat{\pi}_R(N_H, N_L, N_F, S_e, S_p)\}.$$

3. Similarly for inept firms:

$$\hat{\pi}_N^L(N_H, N_L, N_F, S_e, S_p) - c_L(N_L) \geq \max\{0, \hat{\pi}_R(N_H, N_L-\delta, N_F, S_e, S_p)\}.$$

$$\hat{\pi}_N^L(N_H, N_L+\delta, N_F, S_e, S_p) - c_L(N_L+\delta) \leq \max\{0, \hat{\pi}_R(N_H, N_L, N_F, S_e, S_p)\}.$$

- The decision to enter

1. Expected profits must be larger than the free entry cost, but if S_p increases by δ it must be below the free entry cost. The solution should satisfy

$$\mathbb{E}[\pi](S_p) \geq f_e,$$

$$\mathbb{E}[\pi](S_p+\delta) < f_e.$$

Due to rounding errors, the solution for a certain value of S_e might be indeterminate. In that case a solution with a lower value of S_e has been taking (leading to higher expected profits). However, this might lead to small fluctuations in the solution.

Chapter 4

Aiming for the Moving Target: Learning and Strategic Experimentation

Abstract This chapter presents a model and experimental evidence on learning and strategic experimentation in the context of a repeated ultimatum game with imperfect information on the Responder's cost of acceptance. The model predicts that Proposers underbid in early periods, hoping that learning will benefit them in later periods. However, when Responders also have strategic considerations, the Responders will strategically reject low offers. If Proposers anticipate these higher reservation prices, they will respond by offering higher offer. I conduct a series of online experiments in which I vary the scope for strategic considerations, by varying the number of periods, whether the Responder is a computer or a human player and the information available to the Proposer. I find that Proposers make higher offers to Responders who are human and that Responders set the minimum required offer above their true cost. For most participants we cannot rule out that this is driven by distributional preferences or reciprocity concerns. There is a large degree of heterogeneity and both the analysis of play as well as self-reported descriptions of strategies provide evidence that some Responders indeed play strategically.

4.1 Introduction

Imagine being on an archery range, where you have to aim for a fixed target. Through trial and error you learn the best technique and where to aim to hit the target. Now imagine that the target is no longer fixed, but can move, like a clay pigeon. Aiming for the target now becomes harder, as you need to anticipate where the target is when your arrow has crossed the distance of the field.

Many economic interactions are characterized by moving targets. For example, a monopolist tries to set the optimal price given demand. An employer needs to decide what wage to offer in order to induce a prospective employee to accept their offer. However, if the preferences of the consumer or the employee are not known, experimentation is crucial. At the same time, if the consumer or employee realizes that the monopolist or the employer is experimenting, this gives rise to strategic considerations. By rejecting high prices or low wages, the consumer or employee can pretend to be of a more demanding type: they will move the target. The monopolist or employer now needs to take this movement into account while taking their aim.

In this chapter I present a model of learning and experimentation in the context of a repeated ultimatum game.[1] In this game a Proposer can make a proposed split of an endowment to a Responder and only if the Responder agrees, the proposed split is implemented. I introduce one change to the standard ultimatum game: the Responder has a cost of acceptance which is unknown to the Proposer, but drawn from a known distribution. First, consider the case where the Responder has no strategic considerations: the Responder will accept all offers at or above their true cost (their cost of acceptance). Through experimentation, the Proposer can gain information on the potential values of this true cost. However, experimentation could lead to a short-term loss if the offer is rejected and the Proposer has to weigh this short-term loss against the long-term benefit of having better information on the Responder's true cost.

So, the Proposer's optimal offer is determined by the trade-off between the short-run earnings and the long-run benefits from experimentation to increase information. The fewer number of periods are ahead, the lower this benefit from experimentation. The optimal series of offers is characterized by a pattern of *no downward experimentation*: the Proposer will start with a low offer and sticks to this offer when it is accepted by the employee. The opening offer itself will depend on the length of the game and the Proposer's discount rate.

This changes when the Responder takes into account how the acceptance and rejection decision will influence offers in future periods. Responders will now strategically reject some offers to elicit higher offers in the future. In a Perfect Bayesian Equilibrium, Proposers take this strategic rejection into account, both while updating the information on the true cost of acceptance and when making the offers. In all but the last period, the

[1] See Güth & Kocher (2014) for a recent literature overview of ultimatum bargaining.

Proposer will increase offers to reduce the risk of rejection, but decrease offers following acceptance, because only Responders with a low true cost of acceptance would have accepted the offer in the first period. Given that the Responder's benefit from strategic rejection declines with the number of remaining periods, Responders can lower the offer to ensure acceptance by these types.

Next, I conduct a series of online experiments. In this experiment participants play a repeated ultimatum game with either a human player or a computer player. The computer player is instructed to accept any offer above their randomly assigned true cost, while with a human player there is scope for strategic considerations. Furthermore, I change the scope for strategic considerations by varying the number of periods of the game and by revealing the Responder's true cost to the Proposer in certain treatments.

The experiments show that Proposers make higher offers in the human treatments, but that most Proposers are not responsive to other changes in the design. Furthermore, Responders set their reservation prices above their true cost, but do not change their reservation prices over time or when in a five or one period game. This could suggest that most Responders are not strategically rejecting, but that their behaviour is likely driven by distributional preferences or reciprocity concerns. Nevertheless, heterogeneity analysis shows that there is a subgroup of Responders whose behaviour can be described as strategic.

The contributions of this chapter are twofold. First, I develop a model of experimentation predicting patterns of learning, experimentation and strategic rejection. While previous models have often emphasized the patterns of learning (see e.g. Aghion et al., 1991), few papers have considered the consequences of strategic rejection for learning (e.g., Loginova & Taylor, 2008). Second, I conduct a series of online experiments to entangle the patterns of experimentation and strategic play. Previous experiments have shown how people learn social norms by playing many different players (Avrahami et al., 2013) or how people bargain in a repeated ultimatum game with a fixed partner (Slembeck, 1999). While some of these studies have suggested that strategic rejection could happen, they do not provide direct evidence for this. My experiment directly elicits reservation prices by using the strategy method and therefore can provide direct evidence for strategic rejection. Furthermore, my experiment allows to distinguish how learning differs when there is scope for strategic rejection.

This chapter is organized as follows. Section 2 discusses related literature. Section 3 presents the model. Section 4 presents the experimental design, the implementation and the hypotheses. Section 6 presents the results. Section 7 concludes.

4.2 Related literature

Most experimentation models imply a trade-off between the opportunity cost of experimenting and reaping the benefits from experimentation (see e.g. Bolton & Harris, 1999; Keller et al., 2005; Keller & Rady, 1999). Aghion et al. (1991) provide an example of a typical model of experimentation. A monopolist who is facing the same consumer for multiple periods, can learn more about the consumer's demand curve by experimenting. Each period the monopolist has to weigh the potential loss of custom when overcharging for the product against the informational benefit of learning more about the consumer's demand curve. At some points the potential loss might outweigh these benefits.

Aghion et al. show that in such a situation a monopolist can find the true optimal price if the payoff function is analytic, smooth and quasi-concave and there is no discounting. However, often consumer demands are far from satisfying these conditions. Models of ultimatum bargaining often assume that the consumer has a discontinuous demand: above her reservation price all prices are rejected, while prices below her reservation price are accepted. In such a situation the monopolist is facing a discontinuous payoff function. Aghion et al. show that in such a model with a discontinuous payoff learning is inadequate and the monopolist will never learn the true value of the reservation price. Unless the reservation price is zero, at some point the monopolist will stop increasing the price to avoid the risk of not being able to sell the good.

Aghion et al.'s paper does not discuss the potential for the Responder to reject offers strategically. In their paper, it is assumed that the consumer is acting myopically and accept all prices below their reservation price. However, many earlier ultimatum game experiments have shown that, when parties in bargaining have private information, for example about the "size of the pie" (Mitzkewitz & Nagel, 1993; Straub & Murnighan, 1995; Croson, 1996; Rapoport & Sundali, 1996), the value of the experimental currency units (Kagel et al., 1996) or the outside option for the responder (Knez & Camerer, 1994; Kahn & Murnighan, 1993), they will use this information strategically and use their informational advantage. For example, in Boles et al. (2000) the size of the pie is unknown to the Responder, but known to the Proposer. In their experiments, many Proposers take advantage of having this private information, by both deceiving the Responder in their communication, pretending that the size of the pie is smaller, and by making lower offers, anticipating that the Responder can not rule out that the size of the pie was actually smaller. Nevertheless, the Responders do not look favourably upon Proposers deceiving them this way and when they find out later in the game that the Proposer deceived them, they were more likely to reject future offers at their own cost, potentially as a means of retribution.

Much of the literature on learning in the ultimatum game focuses on learning the prevailing social norm. Many of the ultimatum game studies conducted repeated ultimatum games (e.g., Roth et al., 1991; Slonim & Roth, 1998; Cooper & Dutcher, 2011), with practically all of them using stranger matching or perfect stranger matching after each period. Many of these papers show that by repeated play with different partners a norm of equal sharing emerges: offers propose to split the surplus equally (Avrahami et al., 2013). However, few of them repeat bargaining with the *same* proposer or responder. To my knowledge, there is one paper studying the ultimatum game with fixed opponents. Slembeck (1999) conducted a repeated ultimatum game with both fixed and randomly rematched (i.e., stranger) opponents. Proposers in the fixed opponent treatment tend to start their offers at a lower rate than Proposers in the stranger treatment and subsequently increase their offers over time. Rejection rates decrease over time, indicating that Proposers are learning the optimal offer that is accepted.[2]

In a repeated ultimatum game there is scope for strategic rejection. However, to what extent does the Proposer anticipate the opportunity for strategic rejection in their offers? Higher-order rationality or level-k thinking matters here (Costa-Gomes et al., 2001). In the ultimatum game with strategic experimentation level-0 would correspond to a naïve Proposer making offers without taking the possibility of strategic rejection into account.[3] Level-1 then corresponds to the optimal Responder's choice of reservation prices given the Proposer's level-0 choices. Level-2 then corresponds to the Proposer's best choice of offers given the Responder's level-1 reservation prices, and so on. A Proposer adjusting their behaviour anticipating strategic rejection must think at least at level-2. Generally speaking, experiments on level-k thinking find considerable diversity in the degree of strategic thinking. For example, Kneeland (2015) estimates that 93% of the participants are acting rationally at at least level-1, 44% are rational at least level-2 and 22% are rational at least at level-3.

[2]It should be noted that the sample size of this study was small (58 students, of whom 38 were in the fixed opponent treatment) and many of the results are insignificant, e.g. the increasing offers in the fixed opponent treatment. This study did not use the strategy method and is therefore not informative about reservation prices.

[3]One topic of debate is how to define level-0 thinking (Hargreaves Heap et al., 2014). Generally level-0 is defined as the nonstrategic case, based on "nonstrategic, behaviorally plausible general principles" (Crawford et al., 2008).

4.3 Model of experimentation

The section develops the model and shows the pattern of experimentation without strategic considerations and then with strategic considerations.

4.3.1 An ultimatum game with private information

The basis of the model is a finitely repeated Ultimatum Game, where the Responder has an unknown cost of acceptance c_A, which I will call the true cost. Each period, the Proposer makes a proposal to split an amount to the Responder. If the Responder rejects the proposed split, both the Responder and the Proposer do not earn anything. If the Responder accepts the proposed split, the proposal gets implemented and the true cost c_A is deducted from the Responder's payoff. I normalize the total surplus to 1 and denote the amount that the Proposer offers to the Responder with x. If the Responder accepts, the payoffs of the Proposer π_P and the Responder π_R are:

$$\pi_P = 1 - x$$

$$\pi_R = x - c_A$$

This payoff gives utility $u(\pi_P)$ to the Proposer and utility $v(\pi_R)$ to the Responder. The true cost c_A represents potential outside opportunities or, in case of a buyer and a seller, the cost of producing.

Let $F(y) = P(c_A \leq y)$ be the cumulative distribution function of the true cost. I assume that both $F(y)$ and $u(y)$ are twice differentiable. Furthermore, I assume that the offer x can take any value between zero and one, so the support of c_A is $[0, 1]$ (including 0 and 1) and $F'(y) > 0$ for all $y \in [0, 1]$. Without loss of generality, I set the utility value of not receiving any amount to zero, so $u(0) = 0$. I assume that the players are expected utility maximizers. The game is played repeatedly for T periods and the discount factor of the Proposer is δ and discount factor of the Responder is β.

4.3.2 Without strategic considerations

First, I look at the case where the Responder has no strategic considerations, but where the Proposer can act strategically. I assume that the Responder is completely myopic ($\beta = 0$) and purely self-interested. This implies that the Responder will only accept any offer that equals or is larger than c_A.[4] The Responder's acceptance or rejection choice can now be used by the Proposer to put bounds on the potential values of c_A.

Each period, the Proposer faces a trade-off between trying to lower the offer to see if it is accepted and running the risk that the lower offer leads to rejection. As will be shown, under quite general assumptions on $u(y)$ and $F(y)$ this leads to a pattern of experimentation where the Proposer starts offering a low amount and increases this amount when it is rejected, but stops experimenting when the offer is accepted.

The Proposer's optimization problem can be formulated as a finite dynamic programming problem, with the lower and upper bound of c_A as the state variables. Let the bounds on c_A be $[a, b]$. In the initial period $a = 0$ and $b = 1$, which means that c_A can take any value in the domain $[0, 1]$. If the Responder accepts an offer x_t, then $c_A \leq x_t$ and the new bounds on c_A become $[a, \min\{b, x_t\}]$. If the Responder rejects, then $c_A > x_t$, the new bounds become $[\max\{a, x_t\}, b]$.

The value function for the Proposer is

$$\begin{aligned} V_t(a,b) &= \max_{x_t} u(1-x_t)P(c_A \leq x_t | a < c_A \leq b) \\ &\quad + \delta \left[P(c_A \leq x_t | a < c_A \leq b)\, V_{t+1}(a, \min\{b, x_t\}) \right. \\ &\qquad \left. + (1 - P(c_A \leq x_t | a < c_A \leq b))\, V_{t+1}(\max\{a, x_t\}, b) \right]. \end{aligned}$$

The first term in the maximization represents the expected payoff in the current period, which equals the utility $u(1 - x_t)$ times the probability that the offer x_t is accepted. The second term represents the next period's value when the offer is accepted times the probability that the offer is accepted. The third term represents the next period's value when the offer is rejected times the probability that the offer is rejected. The probability of acceptance, i.e. the probability that $c_A \leq x_t$ conditional on $a < c_A \leq b$, is

$$P(c_A \leq x_t | a < c_A \leq b) = \frac{F(x_t) - F(b)}{F(b) - F(a)}. \tag{4.1}$$

The game is finite with T periods and therefore $V_{T+1} = 0$.

The value function can be further simplified by restricting the choice x_t to the domain $[a, b]$. Offers outside this domain are always weakly dominated by offers inside this domain. An offer $x_t < a$ will lead to a certain rejection and therefore a zero payoff and is weakly dominated by an offer $x_t \geq a$. Similarly, an offer $x_t > b$ is weakly dominated by an

[4] I assume that when the Responder is indifferent between accepting and rejecting, the Responder will accept the offer. So, the Responder will accept offers equal to the true cost c_A.

offer of b (which is accepted for sure), since $u(1-x_t) \leq u(1-b)$. This means that the value function can be simplified to

$$\begin{aligned} V_t(a,b) &= \max_{x_t \in [a,b]} u(1-x_t)\frac{F(x_t)-F(a)}{F(b)-F(a)} \\ &\quad + \delta \left[\frac{F(x_t)-F(a)}{F(b)-F(a)} V_{t+1}(a,x_t) + \left(1 - \frac{F(x_t)-F(a)}{F(b)-F(a)}\right) V_{t+1}(x_t,b) \right]. \end{aligned}$$

We can solve for the Proposer's optimal choices of x_t^* by using backwards induction. In the final period, the optimal choice x_T^* is given by the solution of

$$-u'(1-x_T)(F(x_T)-F(a)) + u(1-x_T)F'(x_T) = 0, \tag{4.2}$$

provided that the solution falls within the domain $[a,b]$. The Proposer's optimal offer is a trade-off between receiving a higher "prize" (if a lower offer is accepted) and the probability that the offer is accepted. A lower offer gives the Proposer a higher payoff if accepted, but reduces the probability of acceptance. This is reflected in the above first order condition: the first term represents the marginal loss in utility of a higher offer, while the second term represents the increase in the probability of acceptance of the offer.

The optimal choice x_T^* is bounded between a and b. When there is no solution of equation (4.2) for x_T in the domain $[a,b]$, the optimal choice x_T^* equals b.[5] This means that the optimal offer will be $a < x_T \leq b$. Equation (4.2) does not feature b, which mean that *any interior solution* does not depend on b.

We can now use backward induction to derive the optimal choices in the periods before the last period, by solving the value function for each period. The solution is given by the following proposition:

Proposition 4.1. *The optimal choice in period t, x_t^*, is the minimum of b and the solution of the first-order condition*

$$-u'(1-x_t)(F(x_t)-F(a)) + u(1-x_t)F'(x_t) - \frac{\delta \sum_{k=t+1}^{T} \delta^{k-(t+1)} u(1-x_{k,f}^*(x_t))F'(a)}{\left(1 + \delta \sum_{k=t+1}^{T} \delta^{k-(t+1)} I(x_{k,s}^*(x_t) = x_t)\right)} = 0, \tag{4.3}$$

where $x_{k,s}^*(x_t)$ is the optimal offer in period k following acceptance ("success") in period t, $x_{k,f}^*(x_t)$ is the optimal offer in period k following rejection ("failure") in period t and $I(x_{k,s}^*(x_t) = x_t)$ is an indicator function equal to one if $x_{k,s}^*$ equals x_t and zero otherwise.

[5] The optimal solution cannot be a (provided $a < b$), because the first order condition (4.2) will then be positive, which means any offer $x_T > a$ yields a higher level of utility.

The derivation can be found in the Appendix. Compared to the final period first-order condition in equation (4.2), the first-order condition (4.3) now contains a third term which represents the benefit from experimentation for future periods. A lower offer means that the Proposer learns more about c_A and can therefore increase their earnings in the future period. We can show that under certain conditions the presence of a benefit from experimentation leads to lower initial offers.

Proposition 4.2. *Let* $H(x_t; a) \equiv -u'(1-x_t)(F(x_t) - F(a)) + u(1-x_t)F'(x_t)$. *If* $\frac{\partial H(x_t;a)}{\partial x_t} < 0$ *for all* x_t *and* $\delta > 0$ *and given beliefs* a *and* b*, offers in all but the last period will be lower than when no future periods are present.*

Proof. The first-order condition (4.3) can be rewritten as

$$H(x_t; a) = \frac{\delta \sum_{k=t+1}^{T} \delta^{k-(t+1)} u(1 - x^*_{k,f}(x_t))F'(a)}{\left(1 + \delta \sum_{k=t+1}^{T} \delta^{k-(t+1)}\right)}. \tag{4.4}$$

When $\delta > 0$ and $t < T$, the term on the right is positive, provided $u(1 - x^*_{k,f}(x_t)) > 0$.

Let $x^*_{1\cdot 1}$ be the optimal offer in the one shot game, so $H(x^*_{1\cdot 1}; a) = 0$.[6] Now, consider a game with multiple periods and consider a period $t < T$. In the first order condition (4.4) the last term is now strictly positive and the optimal choice in this period must satisfy $H(x^*_{t\cdot T}; a) > 0$. If $\frac{\partial H(x_t;a)}{\partial x_t} < 0$, then the only value that satisfies (4.4) is a value of $x_t < x^*_{1\cdot 1}$. □

The condition $\frac{\partial H(x_t;a)}{\partial x_t} < 0$ ensures that the solution of the first-order condition is indeed a maximum. Whether the condition holds depends on the utility function $u(x)$ and density functions $F(x)$. The condition holds, for example, when x is uniformly distributed and $u(x)$ is linear or logarithmic.[7]

Patterns of experimentation

We can use this result from the Proposition to compare the predicted play in the two-period game with the one-period game. Two patterns emerge:

First of all, the first-period offer in the two-period game $x^*_{1\cdot 2}$ is *lower* than the offer of the one-period game $x^*_{1\cdot 1}$. The Proposer will make a lower offer when there is another period ahead of her.

[6] I use the following notation: the optimal offer in period t of a game with in total T periods is denoted as $x^*_{t\cdot T}$.

[7] In the uniform distribution case $F(x) = x$. A suitable logarithimic specification is $u(x) = \ln(1 + x)$. Proposition 4.2 also holds under less stringent conditions, for example when $\frac{\partial H(x_t;a)}{\partial x_t} < 0$ for all values that $H(x_t; a) \geq 0$ (but can be positive otherwise). The isoelastic constant relative risk aversion (CRRA) utility function meets this condition.

Second, when the offer $x^*_{1\cdot 2}$ is accepted in the first period, it will be *re-offered* in the second period. Following acceptance in period 1, the Proposer's bound on c_A are now $[0, x^*_{1\cdot 1}]$, and the upper bound $b = x^*_{1\cdot 1}$ in the second period. However, the first-order condition for the last period does not depend on b, but only on a. The value of a has not changed, which means that the optimal choice is determined by (4.4), but now with the right hand term equal to zero, as no future periods are coming up. In contrast, the optimal choice in the first period $x^*_{1\cdot 2}$ was determined by (4.4), with a positive term on the right hand side (i.e., $H(x^*_{1\cdot 2}; a) > 0$). The final period first order condition is $H(x^*_{2\cdot 2}; a) = 0$, and if $\frac{\partial H(x_t;a)}{\partial x_t} < 0$, the solution to this equation must be larger than $x^*_{1\cdot 2}$ and therefore larger than the upper bound b. As discussed earlier, the optimal offer will never exceed b, so the optimal offer is now to re-offer $x^*_{1\cdot 2}$ in the final period as well.

Similar patterns also holds for longer games, because the benefit from experimentation increases with more periods to come.[8] The longer the game, the lower the initial offer, which will be re-offered in case of success. The crucial insight from this is that there is *no* downwards experimentation. If an offer was successful in the first period, it will be re-offered in the second period, but the offer will *not* be lowered. This result is driven by the fact that the first-order condition in equation (4.4) does not depend on b and therefore acceptance in the past will not influence the solution, apart from providing an upper bound for offers in future periods.

Intuitively, an offer is determined by the trade-off between the size of the "prize", the probability of losing it and the expected future gains from experimentation. After success, lowering an offer increases the size of the prize and increases the expected future benefit of experimentation, but decreases the probability of acceptance. The risk of losing the prize outweighs the increase of the size of the prize and the value of experimentation.[9] Learning therefore ceases after the first success and no further attempts to learn anything about c_A are made, simply because learning has become too costly.

Linear utility and uniform distribution

The described pattern of experimentation holds for a wide range of utility functions and distributions. The rest of this chapter will focus on the case where utility is linear, so

[8]This mostly depends on how the right-hand side term in equation (4.4) evolves when the number of periods that are still coming up. In case of the linear utility model with a uniform distribution the right-hand term indeed increases when there are more periods coming up.

[9]For example, in a two-period game, if utility is linear and the true cost is drawn from a uniform distribution, if you have offered 0.50 in the first period and this is accepted, lowering the offer in the second period to 0.49 increases the size of the earnings by 2%, but the probability that it is accepted has decreased by 2%. The expected earnings are $98\% \times (1 - 0.49) = 0.4998$, which is lower than 0.50. Offering 0.50 again will give you 0.50 for sure and is therefore preferred.

$u(y) = y$, and the true cost c_A is drawn from a uniform distribution, $c_A \sim U(0,1)$, and $F(c_A) = c_A$, with $c_A \in [0,1]$. The value function can now be written as

$$\begin{aligned} V_t(a,b) &= \max_{x_t \in [a,b]} (1 - x_t)\frac{x_t - a}{b - a} \\ &\quad + \delta \left[\frac{x_t - a}{b - a} V_{t+1}(a, x_t) + \left(1 - \frac{x_t - a}{b - a}\right) V_{t+1}(x_t, b) \right]. \end{aligned}$$

Again, we can use backward induction and the first order condition (4.4) to find the optimal choices x_t^*. The optimal offer in the final period is the solution of

$$1 + a - 2x = 0,$$

and so

$$x_T^* = \min \left\{ \frac{1+a}{2}, b \right\}. \tag{4.5}$$

The optimal offer is either the average of 1, representing the maximum value of x, and a, representing the highest unsuccessful offer, or b, the lowest successful offer, whichever values is smaller.

We can use this solution in the final period, combined with the first-order condition (4.4) to find the optimal offers in the periods before the last one. For example, in period $T-1$ the first order condition (4.4) equals

$$(1 + a - 2x)(1 + \delta) - (1 - x_T^*(x_{T-1}) = 0,$$

to which the solution is

$$x_{T-1}^* = \min \left\{ \frac{1 + a + (\frac{1}{2} + a)\delta}{2 + \frac{3}{2}\delta}, b \right\}. \tag{4.6}$$

Figure 4.1 shows the sequence of offers following rejection in the one to five period game, assuming $\delta = 1$. One can see that the opening offer is lower for longer games, because the benefit of experimentation increases with the number of periods.

The solutions (4.5) and (4.6) also show that the Proposer will not lower her offer after success in period $T-1$. Consider the two-period game. In the first period $a = 0$ and $b = 1$ and the Proposer will offer $\frac{1+\frac{1}{2}\delta}{2+\frac{3}{2}\delta}$. If the Responder accepts this offer, in the last period $a = 0$ and $b = \frac{1+\frac{1}{2}\delta}{2+\frac{3}{2}\delta}$. The optimal offer according to equation (4.5) is now equal to b, because

$$b = \frac{1 + \frac{1}{2}\delta}{2 + \frac{3}{2}\delta} < \frac{1+a}{2} = \frac{1}{2},$$

provided $\delta > 0$. So the Proposer will re-offer the offer from the first period in the second period. We can show that a similar pattern holds for longer period games as well.

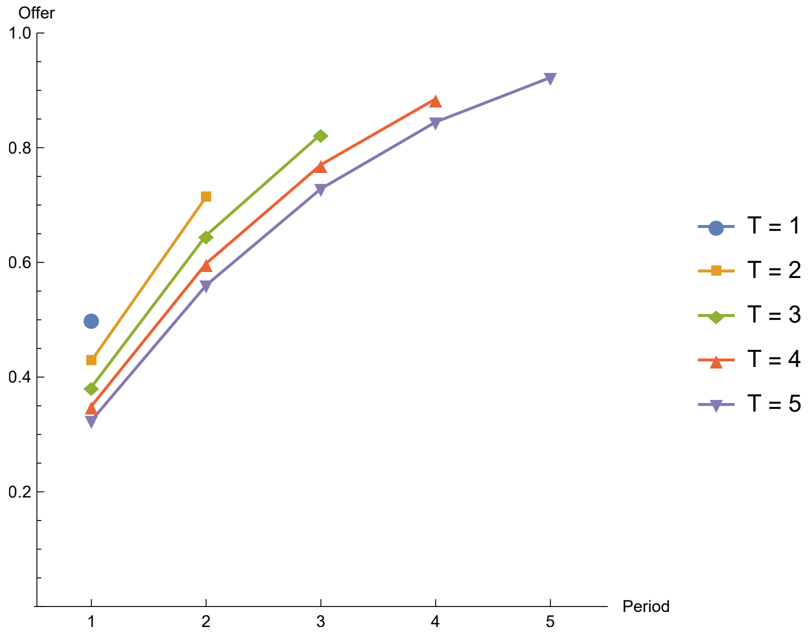

Figure 4.1: Offers following rejection in games with one to five periods.

4.3.3 With strategic considerations

Until so far I have assumed that the Responder is myopic. This means that the Responder is not acting strategically: the Responder accepts any offer above his true cost c_A. However, when the Responder in his decision to accept or reject also takes the future behaviour of the Proposer into account, the Responder might reject some offers above the true cost in order to elicit higher offers in the next period. In a Perfect Bayesian equilibrium the Proposer should take this strategic rejection into account and she will adjust her offers accordingly.

For the following part, I again assume that the utility functions of both the Proposer and the Responder are linear, so $u(1 - x_t) = 1 - x_t$ and $v(x_t - c_A) = x_t - c_A$, and that c_A is drawn from a uniform distribution $U(0, 1)$ (including 0 and 1), as before. The discount factor of the Proposer is δ and the discount factor of the Responder is $\beta > 0$. I illustrate this case by focusing on the two-period game.

Strategic rejection by the Responder

The main difference with the case discussed in the previous section is that Responders reject offers to signal a higher true cost and to receive higher offers in the future. In the previous case the reservation price equalled the cost of production. With strategic interactions this is no longer the case, apart from the final period, when no future interactions are present and the Responder will accept any offer above c_A.

Let x_{2s} be the second period offer when the offer in period 1 is accepted (success) and let x_{2f} be the offer when the offer in period 1 is rejected (failure). Let $\hat{x}_{2f}$ and $\hat{x}_{2s}$ be the Responder's beliefs of x_{2f} and x_{2s}, in period 1. In the first period the Responder faces a trade-off between accepting the offer now and receiving a low offer in the future period or rejecting the offer now and receiving a higher offer in the next period. The Responder will accept an offer in the first period if

$$v(x_1, c_A) + \beta v(\hat{x}_{2s}, c_A) \geq \beta v(\hat{x}_{2f}, c_A).$$

So,

$$x_1 \geq c_A + \beta \left[\max(\hat{x}_{2f} - c_A, 0) - \max(\hat{x}_{2s} - c_A, 0)\right].$$

Next, provided that $\hat{x}_{2f} \geq \hat{x}_{2s}$, the offer made after a rejection is higher than the offer made after an acceptance, the Responder's reservation price (the lowest offer the Responder accepts) is given by

$$x_{1,\min}(c_A) = \begin{cases} c_A + \beta\left(\hat{x}_{2f} - \hat{x}_{2s}\right) & \text{if } \hat{x}_{2s} \geq c_A \\ c_A + \beta\left(\hat{x}_{2f} - c_A\right) & \text{if } \hat{x}_{2f} \geq c_A > \hat{x}_{2s} \\ c_A & \text{if } c_A > \hat{x}_{2f}. \end{cases} \tag{4.7}$$

Proposition 4.3. *In a two-period game, if $\beta > 0$ the Responder rejects some offers above their true cost c_A in the first period.*

This follows from equation (4.7). When $c_A < \hat{x}_{2f}$ and $\beta > 0$, the reservation price $x_{1,\min}$ is larger than c_A. This means that the Responder rejects some offers that he would have accepted in a myopic game. Figure 4.2 shows the relation between the offer made and the true cost c_A for fixed $\hat{x}_{2f}$ and $\hat{x}_{2s}$. Given $\hat{x}_{2f}$ and $\hat{x}_{2s}$, we see that a Responder with a zero true cost will reject offers close to zero, while a Responder with a true cost close to one will accept offers equalling his true cost.

Anticipation by the Proposer

In a Perfect Bayesian equilbrium the Proposer will take this strategic rejection into account while making offers, both in the initial and in the second period.

Period 2 In the second period, the Proposer needs to take into account that some c_A-types strategically rejected the first period offer x_1. Let $\tilde{x}_{2s}$ and $\tilde{x}_{2f}$ be the Proposer's conjecture of $\hat{x}_{2s}$ and $\hat{x}_{2f}$, which are the Responder's conjecture of the final period offers following acceptance and rejection.

Based on the Responder's reservation price $x_{1,\min}(c_A)$ we can now find an expression for the highest c_A-type that would have accepted an offer of x_1 in the first period. Let us

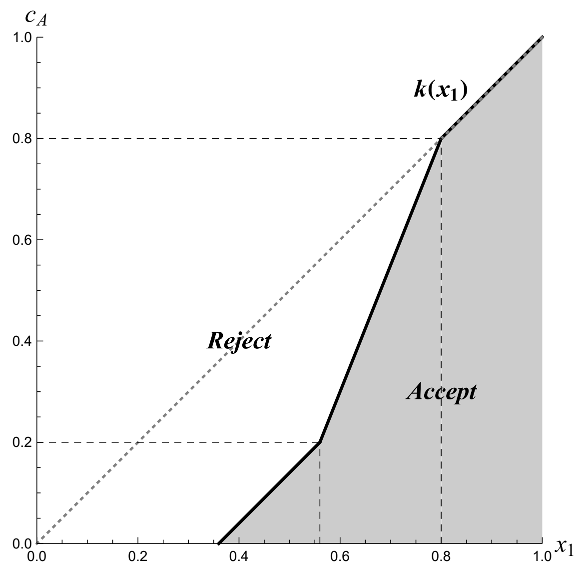

Figure 4.2: The cutoff point $k(x_1)$ for fixed $\hat{x}_{2f}$ and $\hat{x}_{2s}$, for $\beta = 0.6$.

denote the highest c_A-type that accepts the offer x_1 by $k(x_1)$. All types above this cutoff reject the offer in the first period, while types at or below this cutoff reject this offer. We can rewrite (4.7) to find an expression for $k(x_1)$,

$$k(x_1) = \begin{cases} x_1 - \beta\left(\tilde{x}_{2f} - \tilde{x}_{2s}\right) & \text{if } x_1 < (1-\beta)\tilde{x}_{2s} + \beta\tilde{x}_{2f}, \\ \frac{x_1 - \beta\,\tilde{x}_{2f}}{1-\beta} & \text{if } (1-\beta)\tilde{x}_{2s} + \beta\tilde{x}_{2f} \leq x_1 < \tilde{x}_{2f}, \\ x_1 & \text{if } x_1 \geq \tilde{x}_{2f}. \end{cases} \tag{4.8}$$

In case of acceptance, it must be that $c_A \leq k(x_1)$. Therefore acceptance will set the upper bound $b = k(x_1)$. In case of rejection, $k(x_1)$ becomes the new lower bound a. Given that no strategic rejection will happen in the last period, we can use (4.5) to derive the Proposer's final period offer:

$$x_{2f} = \frac{1 + k(x_1)}{2} \tag{4.9}$$

$$x_{2s} = \min\left\{\frac{1}{2}, k(x_1)\right\} \tag{4.10}$$

Using the condition that in a Perfect Baysian Equilibrium beliefs correspond to the actual actions, we can use these expressions for x_{2f} and x_{2s} to find the cutoff c_A-type accepting offer x_1, which equals

$$k(x_1) = \frac{2x_1 - \beta}{2 - \beta}. \tag{4.11}$$

The derivation can be found in the Appendix. Figure 4.3 graphs this relationship. All types with $c_A \leq k(x_1)$ will accept an offer x_1, while types with $c_A > k(x_1)$ will reject this offer. We can see that all c_A-types will reject offers when $x_1 < \frac{\beta}{2}$, because $k(x_1) = 0$.

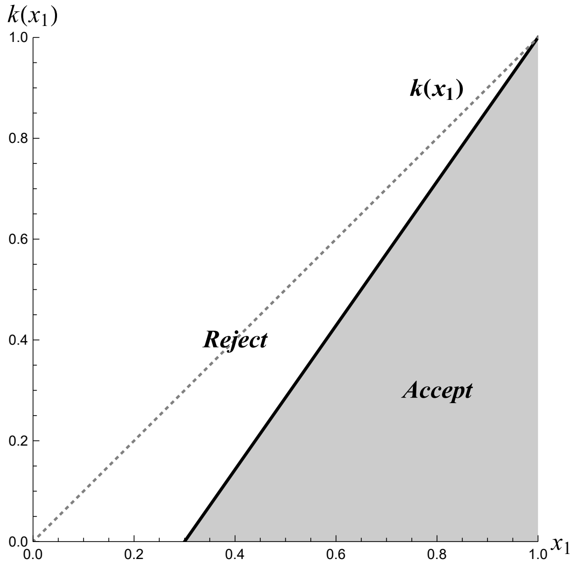

Figure 4.3: The cutoff point $k(x_1)$ in the Perfect Bayesian equilibrium, as a function of x_1, for $\beta = 0.6$. Compared to Figure 4.2 $\hat{x}_{2f}$ and $\hat{x}_{2s}$ are no longer fixed in this graph.

When $\beta = 0$ there are no strategic considerations and $k(x_1) = x_1$, so a Responder will accept any offer at or above their cost of acceptance. However, in the Perfect Bayesian equilibrium, $k(x_1)$ is smaller than x_1, provided that $x_1 < 1$ and $\beta > 0$. If a Responder accepts an offer x_1, it must be that the Responder's c_A-type is *lower* than when there are no strategic considerations (only the lower c_A-types accept offers when $\beta > 0$). This means that acceptance in the strategic considerations case has become a lot more informative of the Responder's c_A-type: if an offer was accepted, it must have been from a quite low c_A-type.

We can use this to formulate the *reservation price* $x_{1,\min}(c_A)$ as a function of β and c_A. The reservation price represents the minimum offer that the Responder accepts. We can rewrite equation (4.11) and show that the reservation price equals

$$x_{1,\min}(c_A) = \left(1 - \frac{\beta}{2}\right) c_A + \frac{\beta}{2}. \tag{4.12}$$

The *markup* of the reservation price, $\mu_1(c_A)$, which is defined as the difference of the reservation price and the true cost c_A, equals

$$\mu(c_A) = x_{1,\min}(c_A) - c_A = \frac{\beta}{2}\left(1 - c_A\right). \tag{4.13}$$

Period 1 Now let us turn to the period 1 offer x_1^*. Just like in the first game, the Proposer faces a trade-off between the size of the "prize" if the offer is accepted, the probability it is accepted and the benefit from experimentation.

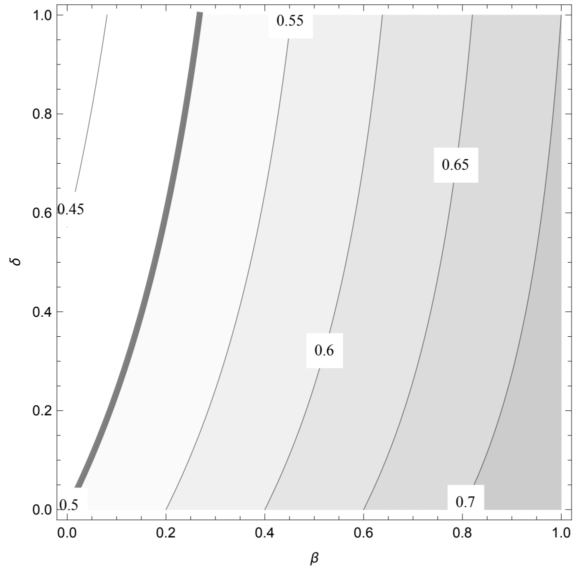

Figure 4.4: The optimal offer x_1 as a function of β and δ.

However, as is discussed previously, the Responder will now reject some of the offers. The probability of acceptance of an offer x_1 equals $k(x_1)$, since $k(x_1)$ represents the highest c_A-type accepting the offer.

The maximization problem in the first period, assuming linear utility, therefore becomes

$$\begin{aligned}\max_{x_1} \quad & (1-x_1)P(c_A \leq k(x_1)) + \delta\left[P(c_A \leq k(x_1))(1-x_{2s})P(c_A \leq x_{2s}^*|c_A \leq k(x_1))\right. \\ & \left. + (1-P(c_A \leq k(x_1)))(1-x_{2f}^*)P(c_A \leq x_{2f}^*|c_A > k(x_1))\right].\end{aligned}$$

Given that $P(c_A \leq k(x_1)) = F(k(x_1)) = k(x_1)$, this simplifies to

$$\max_{x_1} \quad (1-x_1)k(x_1) + \delta\left[k(x_1)(1-x_{2s}^*)\frac{x_{2s}}{k(x_1)} + (1-k(x_1))(1-x_{2f}^*)\frac{x_{2f}-k(x_1)}{1-k(x_1)}\right].$$

Using equations (4.9), (4.10) and (4.11) we can now derive the first-order condition and calculate the optimal first-period offer. This offer x_1^* is given by

$$x_1^* = \frac{1+\frac{1}{2}\delta+\frac{1}{2}\beta\delta-\frac{1}{4}\beta^2}{2+\frac{3}{2}\delta-\beta}. \tag{4.14}$$

Figure 4.4 shows the optimal first period offer x_1^* as a function of the discount factors for the Responder and the Proposer, β and δ. The higher the value of β is, the higher the first period offer will be. The reason for this is that a lower offer will have a higher probability of being rejected if β is higher. High values of δ lead to lower first period offers, because for a high value of δ the Proposer will put more weight on the benefits from experimentation and therefore offers a lower offer.

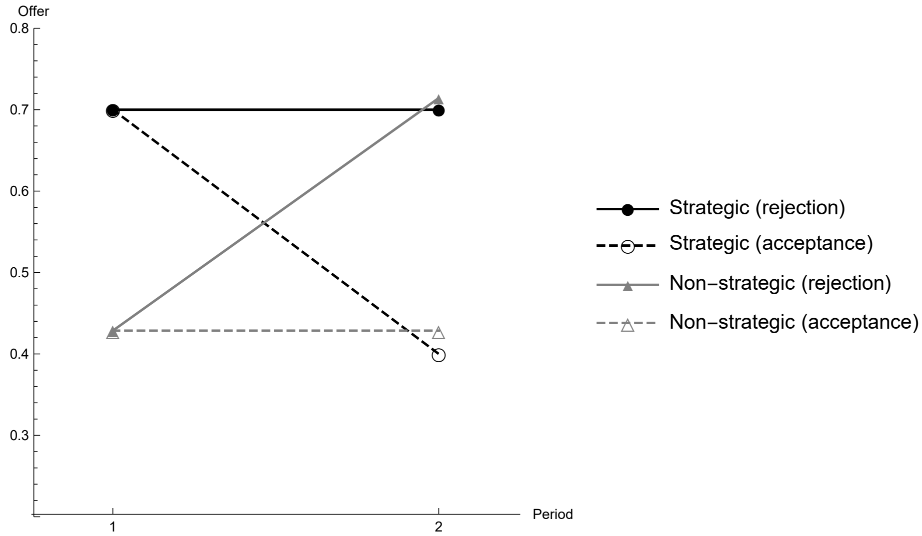

Figure 4.5: Optimal play in the non-strategic case ($\beta = 0$ and $\delta = 1$) and the strategic case ($\beta = \delta = 1$), in a two-period game, assuming linear utility and c_A uniformly distributed.

4.3.4 Comparing strategic with non-strategic play

Figure 4.5 shows the predicted play for the non-strategic game (assuming $\delta = 1$) and the strategic game (assuming $\beta = \delta = 1$). In the non-strategic game, the first period offer is lower than in the strategic game. In the strategic game the opening offer is inflated, due to the c_A-types rejecting low offers, making lower offers unviable. Following acceptance in the non-strategic game, the Proposer maintains the same offer in the next period. Rejection leads to a higher offer.

In the case of the strategic game, we see an opposite pattern. Now rejection leads to maintaining more or less the same offer in the next period (depending on the value of β). The reason for this is that in the first period, some c_A-types strategically rejected their offers. In the last period they do not have a reason for strategic rejection anymore and their reservation prices will now be equal to c_A. Acceptance of the first-period offer now leads to a *lower* offer in the second period. Only the lowest c_A-types will accept the first-period offer, even if it is inflated, and by accepting the offer this low c_A-type is revealed. Given that no further strategic considerations are present in the last period, the Proposer can now lower the offer in the second period.

4.4 Experiments

This section presents the experiment used to test the predictions of the model and to describe the patterns of learning and experimentation.

4.4.1 Design

The design of the experiment follows the setup of the model. A participant is assigned a role of either a Proposer or a Responder and then matched with either a human participant or a computer participant, for either one or five periods. The Responder is randomly assigned the true cost c_A, which can take any value between 0 and 40, and is informed about this. Each period, the Proposer can make an offer, between 0 and 40 points, and the Responder has to indicate their reservation price, the minimum amount at which they accept the offer. If the offer is accepted, the Proposer receives a payoff of 40 points minus the amount offered and the Responder receives this amount minus their true cost. If the offer is rejected, both payoffs are zero.

The computer player always takes the role of Responder and is by construction *non-strategic*. The computer is programmed to accept any offer at or above their true cost c_A. The Proposer is informed that the computer follows such strategy. The Proposer is also informed whether she faces a computer player or a human player. The human players are not constrained in their strategies and can potentially play strategically.

The experiment is framed in market language: the Proposer's role is described as being a *buyer* and the Responder's role is described as being a *seller*. The true cost c_A is described as a *cost of production* in order to sell the good. There are five treatments, which vary whether the Responder is a computer or a human player, the number of periods and whether the true cost c_A is revealed:

- C1. Matched with a computer player for one period.
- C5. Matched with a computer player for five periods.
- H1. Matched with a human player for one period.
- H5. Matched with a human player for five periods.
- H5Rev. Matched with human player for five periods, and the true cost c_A is revealed to the Proposer.

In treatments C1 and C5 the Responder is always a computer player and therefore there is no scope for strategic considerations on the Responder's side. In treatment C5 opportunities for the Proposer to experiment are present, while these are absent in treatment C1, which only consists of one period. Furthermore, these treatments control for social preferences, since the choices by the human Proposer will not influence any other partcipant's payoff. By comparing these treatments with H1, where there is no scope for experimentation but scope for social preferences on behalf of the Proposer, we can isolate the impact of social preferences.

Many ultimatum game experiments have shown that human Responders reject low offers, even if accepting would have given them a positive payoff (Güth & Tietz, 1990). This has often been motivated by other-regarding preferences (Fehr & Schmidt, 1999, e.g.). Another reason for differences could be that agency matters. Charness (2004) showed in a gift-exchange game experiment that workers were more likely to reciprocate high wage offers when this offer was set by a human (who had full agency over their offer) and not a computer.

There are two human treatments with five periods. In H5 there is scope for experimentation by the Proposer as well as for strategic rejection by the Responder. In H5Rev the cost of production is revealed, which limits the scope for both of these. Nevertheless, earlier gift-exchange experiments exhibited a large degree of heterogeneity in Responder's reservation prices, so the minimum offer that the Responder needs might still not be obvious to the Proposer and there is therefore still some scope for experimenting. Both treatments H1 and H5Rev function as control treatments for treatment H5, to identify the influence of experimenting. Treatment C5 functions as a control treatment to identify the influence of strategic rejection. Treatments C1 and H1, together with H5Rev, can be used to identify the impact of playing with a human Responder.

I conduct two versions of the experiment: a non-dichotomous and a dichotomous version. In the non-dichotomous version the true cost c_A can take any integer value between 0 and 40 (with each number equally likely) and in the dichotomous version the true cost c_A is either 5 or 25. The reason for the second variant is that both learning and strategic considerations become more salient: for the Proposer, there are only two values of c_A that she has to learn about and for the Responder choosing to reject strategically boils down to the choice whether to set the reservation price above 25, the high cost level, or not.[10] The parameters have been set to be able to identify strategic rejection separately from demanding an equal split of surplus when the cost of production is 5.[11] The dichotomous treatment

[10]One indication for the increase in salience of the dichotomous treatment can be found in the self-reported measures collected at the end of the experiment. As part of the questionnaire I asked candidates whether they played differently when facing a human or a computer Responder. In the non-dichotomous treatment 40.4% of the participants indicated they played differently, while in the dichotomous treatment 51.2% of the participants indicated that they played differently.

[11]When $c_A = 5$, an equal sharing of surplus is achieved by an offer of $x_t = 22.5$. An offer of $x_t \geq 25$ either corresponds to a Responder requiring more than half of the surplus or strategically rejecting an offer.

has the additional advantage of increased power due to Responders only having two possible values for costs of acceptances, which means that for each of these values there are more observations compared to the non-dichotomous treatments, where the number of observations for each true cost (by construction) is lower. This makes the dichotomous treatment the preferred specification of the experiment.

The participants play up to eight games and an algorithm was set up to make sure that each participant would play at least two computer treatments (C1 and C5). The advantage of such an approach is that it allows us for a within-subject analysis, which increases power (Charness et al., 2012), given the low number of participants (65 for the non-dichotomous game and 110 for the dichotomous game).[12]

4.4.2 Implementation

The experiment was conducted online, with participants recruited through Prolific Academic, an online recruitment platform.[13] In total 175 participants participated in the study. Each participant receives a written instruction and then completes a test of understanding consisting of five questions. After this the participant plays two practice games with a computer, one as a Proposer and one as a Responder. No points can be lost or earned in these practice games. After this, each participant plays up to eight games. If possible, each participant plays each human treatment and each computer treatment at least once. In the human treatments, the role of Proposer and Responder is randomized for each game.

At the end, the participant answers a questionnaire and receives their payment through the online system of Prolific Academic. For the payment a minimum of two games, covering at least six periods, are randomly selected. 100 points are worth 1.50 GBP and the show-up fee is 0.85 GBP. The average total earnings are around 2.00 GBP.

4.4.3 Hypotheses

This section presents the hypotheses, based on the predictions of the model as well as earlier experimental results.

In the computer treatments C1 and C5, no strategic considerations are present. These treatments allow me to investigate how Proposers make offers and how they experiment with finding the optimal offer. We can use this to test two predictions of the model:

[12] As discussed in the conclusion, we are potentially underpowered in some of our analyses.

[13] Unlike Amazon MTurk, Prolific Academic is not restricted to the United States and India. Most of our participant are from the United Kingdom (27.5%) and the United States (13.8%). 68.26% of the participants are male. The average age is 29.9 years. 38.0% of the participants are current students. 70.5% are in part-time or full-time employment.

Hypothesis 4.1 (Opening offer). *In the computer treatments, the opening offer is lower when the game is longer.*

Hypothesis 4.2 (Offer changes). *Rejection leads to a higher offer in the next period, while acceptance leads to a repeat of the successful offer.*

To test for the first hypothesis, we can compare the opening offer in treatment C1 (the one period game) with treatment C5 (the five period game). From the discussion in Section 4.3.2 follows that when there are more periods ahead, the scope for experimentation is larger. This drives down the opening offer. For the second hypothesis we can test how offer are revised in treatment C5. The model predicts that, following acceptance, no downward experimentation will occur in the periods after this.

In the human treatments, the potential for strategic considerations on the side of the Responder differs in the treatments. In treatment H1 no strategic considerations are present, while in treatment H5 there is a potential for strategic rejection. In treatment H5Rev, the true cost is revealed to the Proposer and therefore the potential for strategic rejection is more limited.

Hypothesis 4.3 (Reservation prices). *Reservation prices are higher when there is scope for strategic rejection, and in particular when (a) the game is longer, (b) there are more periods still coming up and (c) the true cost c_A is not revealed.*

To test this, we can compare (a) reservation prices in treatment H5 with prices in treatment H1, (b) the change in reservation prices over time in treatment H5 and (c) compare reservation prices between treatment H5 and H5Rev.

The model of experimentation does not incorporate distributional preferences and reciprocity concerns. A very robust result from many ultimatum game experiments is that Responders reject low offers, even if this offer would have given them a positive payoff and in one period games without any future interactions.

Hypothesis 4.4 (Distributional preferences I). *Humans will set reservation prices strictly above zero, even in single period games or when the true cost c_A is revealed.*

We can test for this by comparing the reservation prices in the human treatments with the true cost c_A in treatment H1 as well as the last period reservation price in treatment H5.

The next question is to what extent the Proposers anticipate the Responders' strategic considerations as well as to what extent Proposer take distributional preferences and reciprocity concerns into account. This leads to the following hypotheses:

Hypothesis 4.5 (Opening offer with strategic considerations)**.** *In the human treatments, the opening offer is higher when the game is longer.*

We can test for this by comparing the opening offer in treatment H5 with the opening offer in treatment H1. The model predicts that offers are higher when there is scope for the Responder to strategically reject an offer.

Hypothesis 4.6 (Distributional preferences II)**.** *Proposers will set higher offers to human Responders than to computer Responders.*

There are two channels through which distributional preferences influence the Proposer's offer: first, Proposers can have distributional preferences of their own and therefore increase their offers, and second, Proposers might anticipate the distributional preferences of the Responder and to avoid rejection, will increase the offer. We can test this by comparing the offers in treatment C1 with the offer in treatment H1. In both treatments, no strategic considerations for either the Proposer or Responder are present. The computer, by construction, does not earn anything and the Proposer's offers will only affect the Proposer's earnings and not the earnings of anyone else.[14]

The final hypothesis is related to learning:

Hypothesis 4.7 (Learning under strategic considerations)**.** *Learning is less effective when strategic considerations are present.*

We can test for this by comparing how offers are adjusted following acceptance and rejection and by comparing the last period offers with the true cost c_A and the Responder's reservation price. With strategic considerations, Responders with a low c_A-type might strategically reject low offers. A rejection is therefore less informative about the true cost, but an acceptance is more informative than in the non-strategic game.

4.5 Results

4.5.1 Opening offers

First, I discuss the opening offers. Hypotheses 1, 5 and 6 relate to these. Hypothesis 1 states that in the computer treatments, the opening offer is lower when the game is longer, so we expect a lower opening offer in treatment C5 than in treatment C1. Hypothesis 5 states that in the human treatments opening offers are higher when the game is longer, so the

[14]An interesting intermediary case would be to consider a treatment where play is conducted by the computer, but the earnings affect a human participant. Studies like Charness (2004) have shown that volition, i.e. the degree of control that the person receiving the money has over the choices, matters for the behaviour in the game.

Table 4.1: Summary of average offers, reservation prices and acceptance shares

	Non-dichotomous ($N = 65$)					**Dichotomous** ($N = 110$)				
	C1	C5	H1	H5	H5Rev	C1	C5	H1	H5	H5Rev
Offer x_t	20.4	21.0	22.5	24.3	25.4	18.5	19.2	20.7	22.3	22.6
- period 1 only	20.4	20.8	22.5	21.4	22.7	18.5	17.6	20.7	19.4	19.7
- period 5 only		21.4		25.8	26.9		19.6		23.3	24.2
Reservation price $x^R_{\min,t}$	20.8	18.3	23.0	24.4	24.3	16.0	15.0	23.4	22.2	22.0
- period 1 only	20.8	18.3	23.0	24.0	25.2	16.0	15.0	23.4	22.7	23.8
- period 5 only		18.3		24.8	23.3		15.0		22.1	21.1
Res. price markup μ_t	0.0	0.0	4.0	6.2	2.6	0.0	0.0	7.4	8.3	7.9
- period 1 only	0.0	0.0	4.0	5.8	3.6	0.0	0.0	7.4	8.8	9.6
- period 5 only		0.0		6.5	1.7		0.0		8.1	6.9
Δ Offer and Res. price	-0.3	2.7	-0.5	-0.1	1.1	2.5	4.2	-2.7	0.2	0.6
- period 1 only	-0.3	2.5	-0.5	-2.6	-2.6	2.5	2.6	-2.7	-3.3	-4.1
- period 5 only		3.0		1.0	3.6		4.6		1.3	3.1
Δ Offer and true cost	-0.3	2.7	3.5	6.1	3.8	2.5	4.2	4.7	8.4	8.5
- period 1 only	-0.3	2.5	3.5	3.2	1.0	2.5	2.6	4.7	5.4	5.5
- period 5 only		3.0		7.6	5.3		4.6		9.4	10.0
% Offers accepted	51.4%	69.2%	45.5%	55.9%	61.6%	62.6%	80.0%	38.7%	61.8%	64.2%
- period 1 only	51.4%	56.6%	45.5%	47.8%	39.5%	62.6%	60.9%	38.7%	43.8%	34.7%
- period 5 only		84.2%		66.7%	84.2%		90.6%		70.5%	86.1%
Proposer's earnings	9.6	12.6	6.3	7.8	6.3	11.6	16.2	5.5	9.7	9.8
- period 1 only	9.6	10.0	6.3	7.4	3.9	11.6	12.4	5.5	7.8	5.5
- period 5 only		15.0		8.6	9.0		18.7		10.4	13.3
Responder's earnings	5.6	5.8	5.2	6.4	4.5	6.0	5.8	5.1	7.5	6.9
- period 1 only	5.6	7.4	5.2	5.0	1.8	6.0	6.6	5.4	5.4	3.6
- period 5 only		4.8		8.1	7.8		4.9		8.6	9.5

Note: The reservation price refers to the minimum price the Responder specified to accept the Proposer's offer. The markup refers to the difference between the reservation and the true cost c_A. c_A is the randomly chosen true cost (cost of acceptance) and is either between 0 and 40 (with each integer equally likely) in the non-dichotomous experiment or chosen from 5 and 25 (with equal probability in the dichotomous experiment.

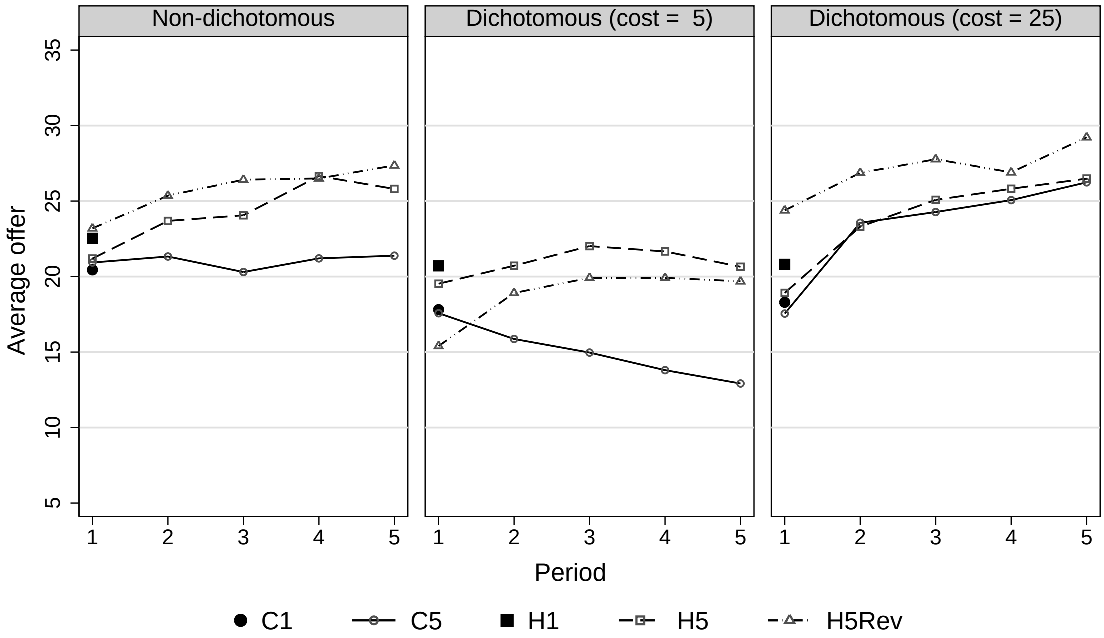

Figure 4.6: Average offers in the various treatments.

Table 4.2: Offers in the first period (treatments C1, C5, H1 and H5)

Dependent variable:	**Non-dichotomous**		**Dichotomous**	
Offer x_1 (period 1 only)	(1)	(2)	(3)	(4)
Computer treatment? (C1, C5)	-2.246^{**}	-2.179^{**}	-2.688^{***}	-1.718^{*}
	(1.102)	(1.009)	(0.875)	(0.901)
5 period game? (C5, H5)	-1.180	-1.599	-1.494	0.111
	(0.943)	(0.963)	(0.955)	(0.947)
Computer treatment? × 5 period game? (C5)	1.701	2.224^{*}	1.043	-0.499
	(1.436)	(1.326)	(1.153)	(1.073)
Constant	22.32^{***}	22.34^{***}	21.53^{***}	20.71^{***}
	(1.153)	(1.023)	(1.118)	(1.211)
Observations	272	272	460	460
R^2	0.0285	0.576	0.0338	0.538
Adjusted R^2	-0.00875	0.422	0.0101	0.387
Proposer fixed effects	No	Yes	No	Yes
Treatment index dummies	Yes	Yes	Yes	Yes

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: This is a linear regression of first period offer on the characteristics of the treatment. The observations from treatment H5Rev are excluded. Standard errors are clustered on the Proposer level. To control for order effects, indicator variables for the order of the treatment are included.

opening offer should be higher in treatment H5 than in treatment H1. Hypothesis 6 states that Proposers will set higher offers to human Responders than to computer Responders.

The second row of Table 4.1 shows the opening offer in the various treatments, while Figure 4.6 shows the average offers by period. For both the non-dichotomous and the dichotomous treatments, the opening offers in the human treatments are higher than the computer treatments. For example, in one-period computer treatment C1 the offer is 20.4 and 18.5, and in the one-period human treatment H1 the offer is higher at 22.5 and 20.7, where the first number refers to the non-dichotomous experiment and the second to the dichotomous experiment. The effect of having multiple periods is smaller and more ambiguous: the average offer in the first-period of treatment C5 is 20.8 and 17.6, which slightly higher than the offer in treatment C1 in the non-dichotomous experiment (20.4), but lower in the dichotomous experiment (18.5). For the human treatments the opening offers are lower in the longer period games, at 21.4 and 19.4 in treatment H5, compared to 22.5 and 20.7 in treatment H1.

In treatment H5Rev effort is revealed and Proposers base their offers on this information. In the non-dichotomous game, the opening offer is on average 1.0 higher than the true cost. In the dichotomous game, the average opening offer is 15.3 when the true cost is 5, and 24.1 when the true cost is 25.

Table 4.2 shows a regression of the opening offer on the characteristics of the treatment, without and with fixed effects.[15] Indicator variables for the game number are included to control for order effects. The regression shows that the offers in the computer treatments are on average 1.7 to 2.7 points lower than in the human treatments. This is significant at a 5% or 10% level, depending on the regression. The impact of being in a five period game is negative, but these estimates are not significant at a 10% level. The interaction between a five period game and a computer treatment is mainly positive, but mostly not significant. Only the treatment effect estimated by the fixed effect regression of the non-dichotomous game, in column (2), is significant at the 10% level.

The results seem to confirm Hypothesis 6 which predicts higher offers when the Responder is human. However, I do not find evidence for Hypotheses 1 and 5. The opening offers in treatment H5 are *lower* than in treatment H1, while the model predicts *higher* opening offers when there are stronger strategic considerations for the Responder. The opening offers in treatment C5 are *higher* than in treatment C1, while the model predicts *lower* opening offers when there are more periods, because the future benefits from experimentation

[15]Including fixed effects means that the coefficients can be interpreted as purely within subject estimates.

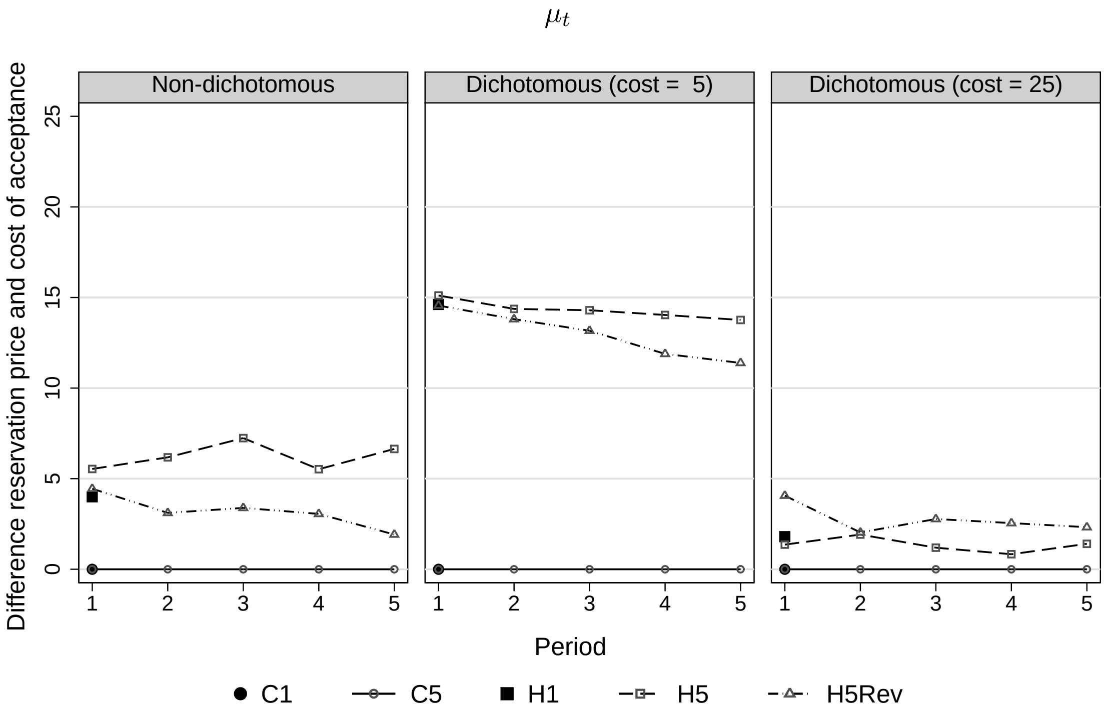

Figure 4.7: The average markup μ_t (the difference between the Responder's reservation price and the Responder's true cost).

are higher.

4.5.2 Reservation prices

Hypotheses 3 and 4 relate to the Responder's reservation price, which is the minimum offer that the Responder has indicated he would accept. Hypothesis 3 states that reservation prices are higher when there is scope for strategic rejection. Hypothesis 4 states that humans will set their reservation prices strictly above zero, regardless of the length of the game or whether the true cost c_A is revealed.

Table 4.1 shows the average reservation price that the Responder set as well as the Responder's markup, which is defined as the difference between the reservation price and the true cost c_A. Figure 4.7 shows the average Responder's markup by period. The computer's reservation price is by definition equal to the true cost c_A, and is therefore close to 20 in the non-dichotomous treatment and equal to 5 or 25 in the dichotomous treatment. As can be seen in Table 4.1, human Responders set their reservation price above their true cost, on average 4.0, 5.8 and 3.6 points higher in treatments H1, H5 and H5Rev in the non-dichotomous experiment. The markup, the difference between the reservation price and the true cost, is higher when the true cost is lower. In the dichotomous treatment the difference is the largest when the true cost is 5, at 14.6, 14.3 and 13.0 in treatments H1, H5

Table 4.3: Markups in the first period (treatments H1, H5 and H5Rev)

Dependent variable:	**Non-dichotomous**		**Dichotomous**	
Markup $\mu_t = x^R_{\min,t} - c_A$ (period 1 only)	(1)	(2)	(3)	(4)
5 period game? (H5, H5Rev)	1.996	2.082	0.968	1.516
	(1.498)	(2.029)	(1.171)	(1.658)
Treatment H5Rev?	-1.575	-2.593	0.821	-1.299
	(2.327)	(3.421)	(1.448)	(2.417)
Constant	5.104^{***}	5.529^{**}	7.967^{***}	8.495^{***}
	(1.841)	(2.296)	(1.394)	(1.773)
Observations	168	168	270	270
R-squared	0.0396	0.532	0.0247	0.431
Adjusted R-sq	-0.0151	0.219	-0.00902	0.0613
Responder fixed effects	No	Yes	No	Yes
Treatment index dummies	Yes	Yes	Yes	Yes

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: This is a linear regression with the first-period markup μ_t as the dependent variable. The markup μ_t is defined as the difference between the reservation price $x^R_{\min,t}$ and the true cost c_A. Only observations from the first period of treatments H1, H5 and H5Rev are included. Standard errors are clustered on the Responder level. To control for order effects, indicator variables for the order of the treatment are included.

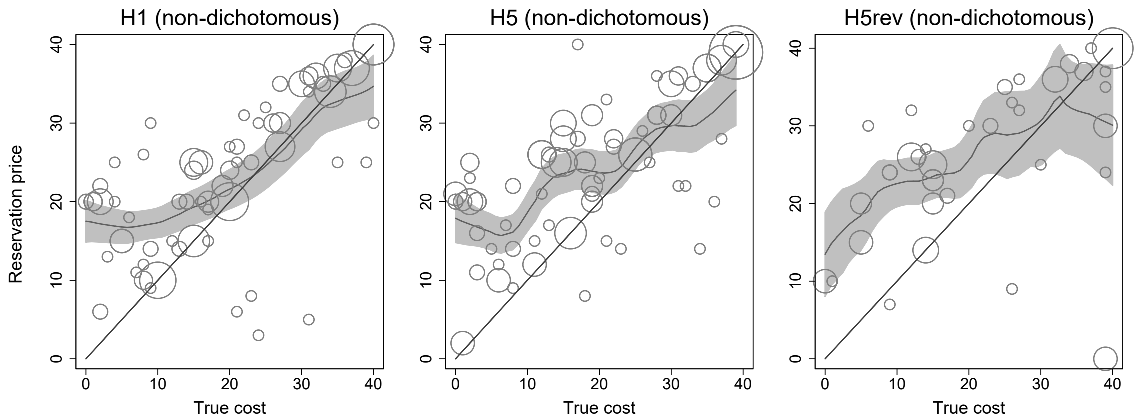

Figure 4.8: The relation between the true cost c_A and the reservation price $x^R_{\min,t}$ in treatments H1, H5 and H5Rev of the non-dichotomous experiment (period 1 only). This is a local polynomial regression with a 90% confidence interval. The size of the bubbles indicate the number of observations.

Table 4.4: Difference reservation price and true cost across time (treatments H5 and H5Rev)

Dependent variable:	**Non-dichotomous**		**Dichotomous**	
Markup $\mu_t = x^R_{\min,t} - c_A$	(1)	(2)	(3)	(4)
Period 2	0.642	0.642	-0.139	-0.139
	(0.724)	(0.765)	(0.444)	(0.470)
Period 3	1.701^{**}	1.701^{**}	-0.515	-0.515
	(0.781)	(0.825)	(0.513)	(0.543)
Period 4	-0.0149	-0.0149	-0.822	-0.822
	(0.748)	(0.791)	(0.532)	(0.563)
Period 5	1.104	1.104	-0.703	-0.703
	(0.899)	(0.950)	(0.683)	(0.723)
Treatment H5Rev?	-1.528	-3.540	0.586	-1.629
	(2.280)	(3.138)	(1.543)	(2.214)
× Period 2	-1.975	-1.975	-1.205	-1.205
	(1.274)	(1.347)	(0.734)	(0.777)
× Period 3	-2.757^{**}	-2.757^{**}	-0.828	-0.828
	(1.210)	(1.278)	(0.942)	(0.996)
× Period 4	-1.374	-1.374	-1.313	-1.313
	(1.354)	(1.431)	(1.103)	(1.167)
× Period 5	-3.632^{**}	-3.632^{**}	-1.804^{*}	-1.804
	(1.415)	(1.495)	(1.053)	(1.114)
Constant	10.09^{***}	8.535^{***}	9.358^{***}	10.39^{***}
	(2.525)	(2.966)	(2.111)	(2.529)
Observations	515	515	840	840
R^2	0.0799	0.616	0.0322	0.501
Adjusted R^2	0.0503	0.557	0.0134	0.430
Responder fixed effects	No	Yes	No	Yes
Treatment index dummies	Yes	Yes	Yes	Yes

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: This is a linear regression with the markup μ_t as the dependent variable. The markup μ_t is defined as the difference between the reservation price $x^R_{\min,t}$ and the true cost c_A. Only observations from treatments H5 and H5Rev are included (all periods). Standard errors are clustered on the Responder level. To control for order effects, indicator variables for the order of the treatment are included (not reported).

and H5Rev. When the true cost is 25, the corresponding values are 1.8, 1.3 and 2.7. Figure 4.8 shows the relation between the true cost and the first-period reservation price. It shows that there is a positive relation, but that at low true costs Responders set a higher markup than at high true costs.

Table 4.3 shows a linear regression of the first period difference of the Responder's reservation price and their cost acceptance on the characteristics of the treatment. The coefficient on being in a five period game is positive, but not significant. The model predicts that the reservation price will be set higher in games with multiple periods, because of the scope for strategic rejection, and that this coefficient should therefore be positive. When the cost of production is revealed, the reservation price is generally set lower, but again, this difference is not significant.

Table 4.4 shows how this difference changes over the several periods in the game, for treatments H5 and H5Rev. The regressions suggest that in the non-dichotomous case, the reservation prices increase over time in treatment H5 (even though only the coefficient for period 3 is significant) and that in treatment H5Rev reservation prices decrease over time (the coefficients for period 3 and 5 are significant at the 5% level). In the dichotomous experiments there is weak evidence of a similar pattern in treatment H5Rev.

These results provide evidence for Hypothesis 4. Even in treatment H1 and H5Rev human Responders set their reservation prices above their true costs. However I find only limited evidence for Hypothesis 3, which states that strategic rejection is more likely to happen in longer games and when there are more periods left. The regressions show that in longer games, the opening offer is higher, but this difference is not significant. I do not find evidence for decreasing reservation prices: if anything, in treatment H5, reservation prices are increasing with time. There is some evidence for decreasing reservation prices in the H5Rev treatment, but in this treatment the true cost is known to the Proposer.

Heterogeneity The previous analysis focused on aggregate behaviour. However, there is a large degree of heterogeneity in the sample, which suggests that there are some participants that strategically reject. Table 4.5 shows the share of workers in the dichotomous treatment who are assigned true cost $c_A = 5$ who set their reservation prices above 20 or 25 points. These are interesting cutoffs: offers between 20 and 25 points correspond to a more or less equal sharing of surplus (an offer of 22.5 points implies an equal sharing; however, offers have to be whole numbers), while 25 points corresponds to the higher possibility of the true cost. Reservation prices above 25 points could therefore indicate a strategic motive

Table 4.5: Share of Responders with $c_A = 5$ setting their reservation prices above 20 or 25 points (dichotomous experiment)

Period	1	2	3	4	5	Average
Reservation price $\geq$ 20 points						
Treatment H1	54.0%					54.0%
Treatment H5	62.9%	53.2%	53.2%	51.6%	50.0%	54.2%
Treatment H5Rev	61.5%	48.7%	46.2%	38.5%	30.8%	45.1%
Reservation price $\geq$ 25 points						
Treatment H1	32.0%					32.0%
Treatment H5	48.4%	38.7%	38.7%	35.5%	35.5%	43.4%
Treatment H5Rev	33.3%	28.2%	20.5%	15.4%	17.9%	28.7%

for the Responder, who could now pretend to be the high cost type, at least when the true cost c_A is not revealed.

Table 4.5 shows indeed that in treatment H5, where the true cost c_A is not revealed, the share of workers setting their reservation prices above 25 points is higher than in treatment H5Rev, where c_A is revealed. On average 43.4% of the workers set their reservation prices above 25 in treatment H5, while only 28.7% does that in treatment H5Rev. This difference is significant at the 5% level ($t = -2.32$, $p = 0.023$, with standard errors clustered on the Responder level).[16] In the final questionnaire 10 out of 110 participants stated explicitly that they strategically rejected to make the Proposer believe their true cost was 25, despite it being 5 points.[17]

4.5.3 Change of offers

The next question is how employers update their offer in response to acceptance or rejection in the five period games. Hypothesis 2 states that rejection leads to a higher offer in the next period, while acceptance leads to a repeat of the successful offer. Figure 4.9 shows the pattern of acceptance and rejection as well as the offers made in the period of acceptance or rejection and the revised offer afterwards for treatments C5 and H5. Generally, the offer is increased when the offer is rejected and the offer is lowered following an acceptance. For example, in the non-dichotomous treatment C5, the average first-period offer that was rejected was 19 and in the following period after this rejection the average offer

[16] A within subject comparison finds a larger difference between H5 and H5Rev (28.1%, with $p = 0.020$).

[17] Several participants reported motives like "I always set it at 30 so they don't think my producing cost was 5 (even if it was).", "Always play as if production was 25.", "Tried to maximise profit by attempting to trick the other player into assuming my costs were 25 if they were actually only 5.", "I set the price at or just above 25", "Pretended unit cost was 25 not 5", "I assumed they would test my rate for a 5 costs so I set at 26 regardless".

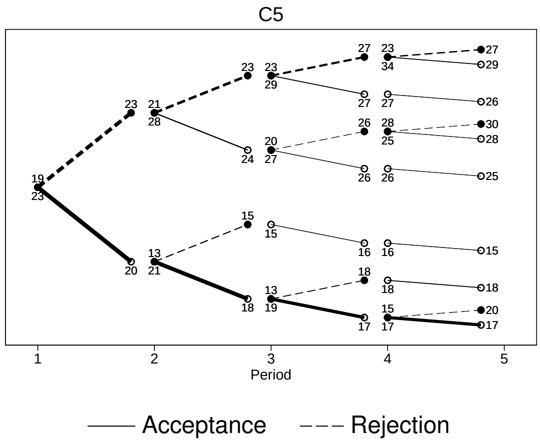

(a) Treatment C5 (non-dichotomous)

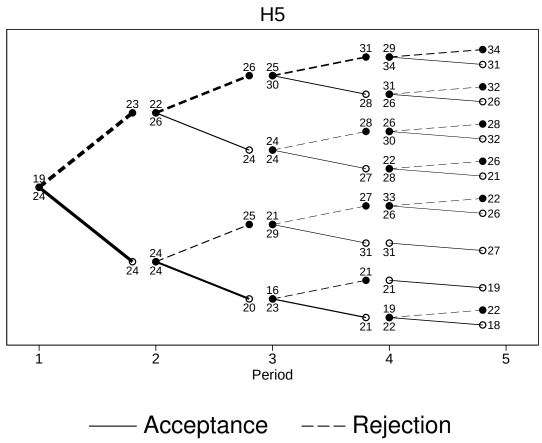

(b) Treatment H5 (non-dichotomous)

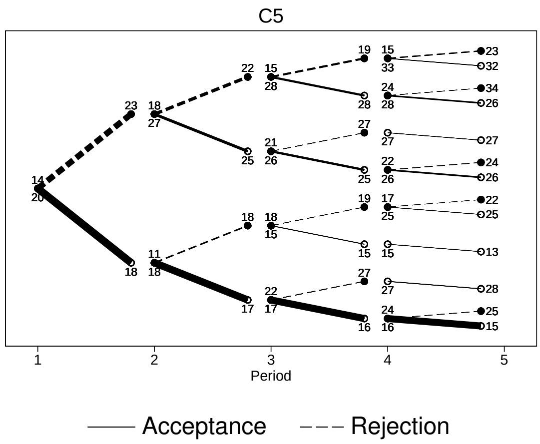

(c) Treatment C5 (dichotomous)

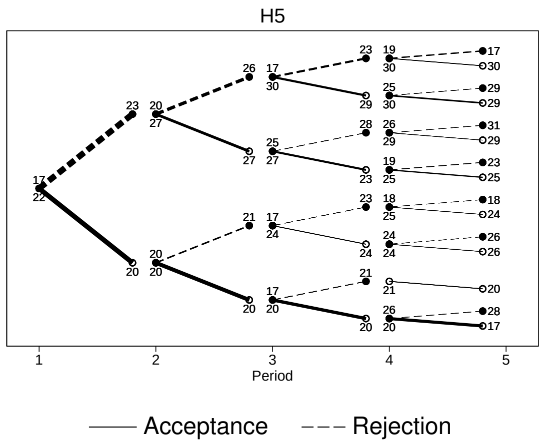

(d) Treatment H5 (dichotomous)

Figure 4.9: Offers and the pattern of acceptance and rejection. The thickness of the line indicates the frequency of this pattern. The numbers indicate the average offer in the period of the acceptance or rejection and the revised offer in the next period.

is 23. The average first-period offer that was accepted was 23 and in the second period following acceptance the average offer was 20.

Table 4.6 shows a linear regression of the change in offer Δx_t on whether the offer was rejected and treatment. Rejection leads to a higher offer in the next period. The coefficient suggest that the offer increases by 5.9 to 6.7 points. The coefficient on the interaction between treatment C5 and rejection in the previous period is small in the non-dichotomous experiment and positive in the dichotomous experiment. This suggests that rejection in the computer treatment in the dichotomous experiment leads to higher wage increases than in the human treatments, though most coefficients are not significant.

The constant is negative and significant at a 1% level. This suggest that when the offer is accepted, the next period's offer is 1.6 to 2.7 points lower. This means that employers are adjusting their offer *downwards* following success, contradicting Hypothesis 2. Nevertheless, the absolute change in offer following success (1.6 to 2.7 points) is lower than the (absolute) change in offer following failure (5.9 to 6.7), suggesting that less downward experimentation than upward experimentation happens. We do not find a significant difference between H5 and C5 (the coefficient for treatment C5 is close to zero).

However, the coefficient for treatment H5Rev, where the true cost is revealed, is positive, while the coefficient of the interaction with rejection in the previous period is negative. This suggests that there is *less* downwards adjustment following success and *less* upwards adjustment following failure. The explanation for this might be that in treatment H5Rev there is no need for experimentation to figure out the true cost, but only to figure out the value of the markup.

4.5.4 Success of learning

Hypothesis 7 relates to learning and states that learning is the least effective when there is more scope for strategic considerations. One way of measuring the success of learning is analyse the share of offers that are accepted in the final period, the difference between the final offer and the true cost or reservation price and the last period earnings of the Proposer. Figure 4.10 shows the share of offers accepted each treatment. This share is increasing each period. Interestingly, in treatment H5Rev, where the true cost is revealed, rejection is high in the first period, but then decreases drastically in later periods. This pattern is the clearest in the dichotomous case when the true cost is equal to 5. This could point to an underestimation by the Proposers of the Responder's reservation price. Even though they know the true cost, the availability of this information might lead to overconfidence

Table 4.6: Change of offer (treatments C5, H5 and H5Rev)

Dependent variable:	**Non-dichotomous**		**Dichotomous**	
Change in offer ($\Delta x_t = x_t - x_{t-1}$)	(1)	(2)	(3)	(4)
Rejected?$_{t-1}$	5.925^{***}	5.975^{***}	6.299^{***}	6.671^{***}
	(0.598)	(0.628)	(0.760)	(0.665)
Treatment C5?	−0.299	−0.290	0.244	−0.0494
	(0.389)	(0.528)	(0.301)	(0.332)
× Rejected?$_{t-1}$	0.194	0.538	1.273	1.973^{**}
	(0.818)	(0.922)	(0.919)	(0.926)
Treatment H5Rev?	1.122^{**}	1.657^{**}	0.551^{*}	0.296
	(0.530)	(0.709)	(0.326)	(0.422)
× Rejected?$_{t-1}$	-2.130^{***}	-1.967^{*}	−1.328	−0.950
	(0.793)	(1.018)	(0.993)	(1.060)
Constant	-2.464^{***}	-2.721^{***}	-1.630^{***}	-1.965^{***}
	(0.553)	(0.730)	(0.415)	(0.479)
Observations	704	704	1180	1180
R^2	0.191	0.238	0.264	0.344
Adjusted R^2	0.177	0.149	0.256	0.275
Proposer fixed effects	No	Yes	No	Yes
Treatment index dummies	Yes	Yes	Yes	Yes

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: This is a linear regression of the change in offer $\Delta x_t = x_t - x_{t-1}$. Rejected?$_{t-1}$ indicates whether the offer in the previous period was rejected. Only observations from periods 2 to 5 from treatments C5, H5 and H5Rev are included. Standard errors are clustered on the Proposer level. To control for order effects, indicator variables for the order of the treatment are included (not reported).

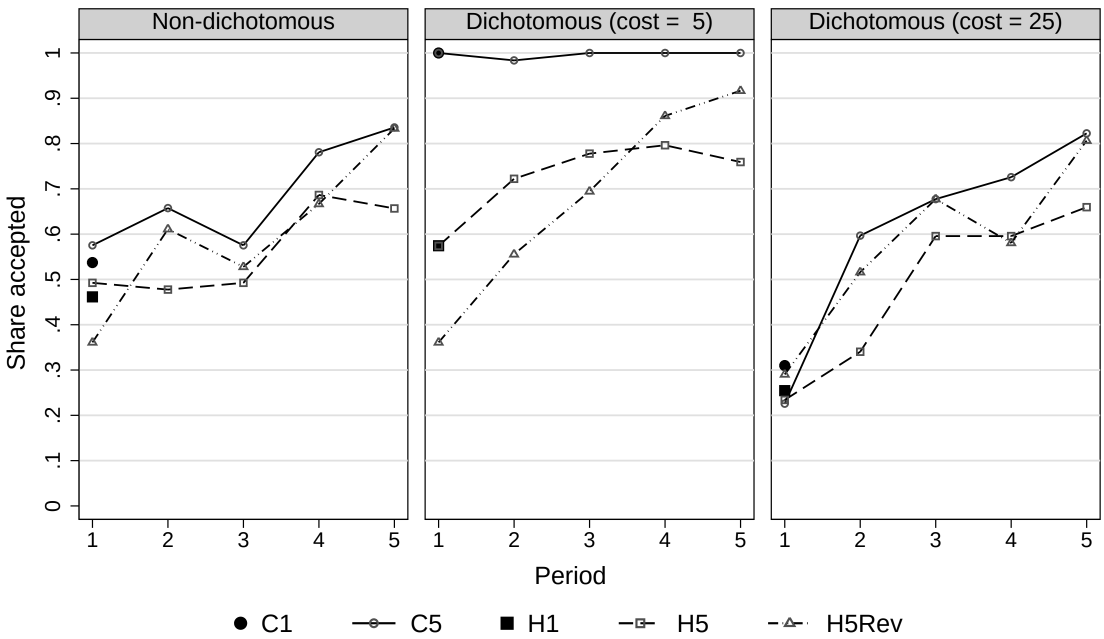

Figure 4.10: Share of offers accepted in the various treatments.

in bidding, not taking the high markup between the reservation price and the true cost into account.

In Table 4.1 we can see that in the final period the highest number of offers are accepted in treatments C5 and H5Rev. In the non-dichotomous experiment, 84.2% and 84.2% of offers are accepted in the final period of treatments C5 and H5Rev, compared to 66.7% of offers in the final period of treatment H5. The dichotomous experiment shows similar figures (90.6% and 86.1% against 70.5%). The difference between the final offer and the true cost for the computer treatments is on average 3.0 and 4.6 points for the non-dichotomous games. This is lower than for treatments H5 and H5Rev (7.6 and 5.3; 9.4 and 10.0), indicating that the Proposer's offers are closer to the true cost in the computer treatments. The final period Proposer's earnings are the highest in the C5 treatment (15.0 and 18.7 in the non-dichotomous and dichotomous game), and slightly higher in treatment H5Rev compared to treatment H5. This provides some supportive evidence for Hypothesis 7, which states that in the treatments with the largest scope for strategic considerations, learning is the least effective.

4.6 Discussion and conclusion

In this paper I presented a model of experimentation in a repeated ultimatum game and considered the role for both strategic experimentation on the Proposer's side and strategic rejection on the Responder's side. The model predicted that, without strategic considerations for the Responder, Proposers make lower opening offers in longer games, but that following acceptance experimentation would stop and Proposer would continue offering the same amount. However, while the model implied that no downward experimentation would occur, the experiments show that many Proposers decrease their offers following acceptance. Even though the downward adjustment following acceptance is smaller than the upwards adjustment following rejection, this suggests that many Proposers weigh the benefits of learning the reservation price of the proposers more highly than the opportunity cost of potentially losing the prize following unsuccessful experimentation.

When the Responder has strategic considerations, the model predicts that the Responder will reject low offers hoping for higher offers in the future. The Proposer will anticipate this and make higher offers. The model also predicts that if there are fewer periods left, the scope for strategic rejection will decline. As a result, the Proposer will decrease their offers following acceptance.

The experimental results show that in the human treatments, Responders set their reservation prices above their true cost. However, while the model predicts that the markup they set should decline over time, the experiments don't show evidence for this, apart from the H5Rev treatment where the true cost is revealed. Also, the markup does not seem to differ between the H1 (without strategic considerations) and H5 treatment (with strategic considerations). This suggests that, when considering the averages, we do not find evidence for higher reservation prices when there is more scope for strategic rejection.

The relatively constant markup between the reservation price and the true cost seems to suggest that the behaviour could be driven by distributional preferences and/or reciprocity concerns. We see at low true costs that many Responders set their reservation price such that the outcome would correspond to an equal sharing of surplus. Nevertheless, there could also be other motivations. One such motivation is a desire to be consistent or to not let the Proposer realize that they were deceived. This could explain why the markup does not decrease over time in the H5 treatment. If it would, then at some point the Proposer could realize that Responder strategically rejected in one of the earlier periods. As Boles et al. (2000) show, Proposers do not look favourably upon Responders who deceived them and the Responder might have internalized that norm. This would explain the lack of decline in markup in the H5 treatment, but not the difference between H5 and H1. Another reason could be group reputation concerns. The game is played multiple times, and a Responder in a one-shot game setting a high markup could provide a "public good" to the other Responders, potentially shifting the beliefs of the Proposer (Huck & Lünser, 2010). However, if this was the case, we would see a decline in markups in treatments towards later games in the experiment, which is not the case.[18]

The lack of significance is partially due to the low number of observations (there were 65 participants in the non-dichotomous game and 110 in the dichotomous game), despite the use of within-subject estimates. Nevertheless, the low point estimate of most coefficients suggest that even with a higher number of observations, a significant result is most likely only representing a small effect.[19]

The design of the experiment matters. The self-reported descriptions suggests that more Responders thought about strategic rejection in the dichotomous game than in the non-dichotomous game. Just having two options therefore increased the salience of such strategic considerations. A different design, for example one where the opportunities for

[18]It should be noted that participants were not made aware of the number of games they were playing. However, participants were aware of the average time the experiment took.

[19]There are a couple of point estimates that are sizeable, but where no significance is found, e.g. the coefficients in Table 4.3. The experiment is potentially underpowered to detect an effect here.

deceit are more salient (e.g., one where the Responder can send a non-binding message about their true cost) could potentially increase strategic considerations.

Another potential treatment that could be introduced to separate between the effect of social preferences towards a human player and the effect of (anticipating) strategic rejection is a stranger matching design, where the Proposer plays a new Responder each period, but all Responders have the same cost of acceptance c_A. In such a treatment, there is no scope for strategic rejection on behalf of the Responder. However, social preferences will still influence offers and reservation prices. I hope to further investigate this matter in future research.

Appendices of Chapter 4

4.A Proofs and derivations

4.A.1 Proposition 1: Optimal choice in period t

First, we start with redefining the value function to obtain a more tractable version of the Bellman equation. Let $\tilde{V}_t(a,b) = (F(b)-F(a))V_t(a,b)$. The dynamic programming problem now simplifies to

$$\tilde{V}_t(a,b) \quad = \quad \max_{x_t \in [a,b]} \left\{ u(1-x_t)(F(x_t)-F(a)) + \, \delta \, \left[\tilde{V}_{t+1}(a,x_t) + \tilde{V}_{t+1}(x_t,b) \right] \right\}.$$

In period t the first order conditions are

$$-\,u'(1-x_t)(F(x_t)-F(a)) + u(1-x_t)F'(x_t) + \delta \, \left[\frac{\partial}{\partial x_t} \tilde{V}_{t+1}(a,x_t) + \frac{\partial}{\partial x_t} \tilde{V}_{t+1}(x_t,b) \right] = 0. \tag{4.15}$$

Proposition 4.4. *The derivative of $\tilde{V}$ with respect to a is*

$$\frac{\partial \tilde{V}_t(a,b)}{\partial a} = -\sum_{k=t}^{T} \delta^{k-t} u(1-x_k^*)F'(a)$$

.

Proof. We can use the envelope theorem to differentiate $\tilde{V}(a,b)$ with respect to a:

$$\frac{\partial}{\partial a}\tilde{V}_t(a,b) = -u(1-x_t^*)F'(a) + \delta \frac{\partial}{\partial a}\tilde{V}_{t+1}(a,x_t^*).$$

The first term is zero or negative, given that $u \geq 0$ and $F'(a) \geq 0$. We can expand the second term using the envelope theorem again,

$$\frac{\partial}{\partial a}\tilde{V}_{t+1}(a,x_t^*) = -u(1-x_{t+1}^*)F'(a) + \delta \frac{\partial}{\partial a}\tilde{V}_{t+2}(a,x_{t+1}^*).$$

We can repeat this exercise until the last period, where $\tilde{V}_{T+1} = 0$, so $\frac{\partial}{\partial a}\tilde{V}_{T+1} = 0$. Each of the terms is zero or negative, so

$$\frac{\partial}{\partial a}\tilde{V}_t(a,b) = -\sum_{k=t}^{T} \delta^{k-t} u(1-x_k^*)F'(a) \leq 0.$$

□

Proposition 4.5. *The derivative of $\frac{\partial \tilde{V}_t(a,b)}{\partial b}$ is a multiple of $H(b;a) = -u'(1-b)(F(b)-F(a)) + u(1-b)F'(b)$.*

Proof. We can show using the envelop theorem, if $x_t^* < b$.

$$\frac{\partial}{\partial b}\tilde{V}_t(a,b) = \delta\frac{\partial}{\partial b}\tilde{V}_{t+1}(x_t^*, b).$$

When the constraint on x_t^* is binding, $x_t^* = b$ and the offer is accepted for sure (the optimal offer is repeating the lowest successful offer). In this case the lower b the better and the valuation $\tilde{V}$ equals

$$\tilde{V}_t(a,b) = u(1-b)(F(b)-F(a)) + \delta\left[\tilde{V}_{t+1}(a,b)\right]$$

So

$$\begin{aligned}
\frac{\partial}{\partial b}\tilde{V}_t(a,b) &= -u'(1-b)(F(b)-F(a)) + u(1-b)F'(b) + \delta\left[\frac{\partial}{\partial b}\tilde{V}_{t+1}(a,b)\right] \\
&= H(b;a) + \delta\left[\frac{\partial}{\partial b}\tilde{V}_{t+1}(a,b)\right] \\
&= \sum_{k=t}^{T}\delta^{k-t}I(x_k^* = b)H(b;a) = H(b;a)\sum_{k=t}^{T}\delta^{k-t}I(x_k^* = b),
\end{aligned}$$

where $I(x_k^* = b) = 1$ if the optimal choice for that period equals b and otherwise zero. □

This means that we can write down the first-order condition (4.15) as

$$\frac{dV_t(x;a,b)}{dx} = H(x_t,a)\left(1+\delta\sum_{k=t+1}^{T}\delta^{k-(t+1)}I(x_k^* = x_t)\right) - \delta\sum_{k=t+1}^{T}\delta^{k-(t+1)}u(1-x_k^*)F'(a) = 0 \tag{4.16}$$

4.A.2 Derivation of equation (4.11)

Let $\hat{k}(x_1)$ be the Responder's belief of $k(x_1)$ and $\tilde{k}(x_1)$ the Proposer's belief of $\hat{k}(x_1)$. From Equation (4.5) follows that the conjectures $\tilde{x}_{2s}$ and $\tilde{x}_{2f}$ equal

$$\tilde{x}_{2f} = \frac{1+\tilde{k}(x_1)}{2}, \tag{4.17}$$

$$\tilde{x}_{2s} = \min\left\{\frac{1}{2}, \tilde{k}(x_1)\right\}. \tag{4.18}$$

The expression for $k(x_1)$ then becomes:

$$k(x_1) = \begin{cases} x_1 - \beta\left(\frac{1+\tilde{k}(x_1)}{2} - \min\left\{\frac{1}{2}, \tilde{k}(x_1)\right\}\right) & \text{if } x_1 < (1-\beta)\min\left\{\frac{1}{2}, \tilde{k}(x_1)\right\} + \beta\frac{1+\tilde{k}(x_1)}{2}, \\ \frac{x_1 - \beta\frac{1+\tilde{k}(x_1)}{2}}{1-\beta} & \text{if } (1-\beta)\min\left\{\frac{1}{2}, \tilde{k}(x_1)\right\} + \beta\frac{1+\tilde{k}(x_1)}{2} \le x_1 < \frac{1+\tilde{k}(x_1)}{2}, \\ x_1 & \text{if } x_1 \ge \frac{1+\tilde{k}(x_1)}{2}. \end{cases} \tag{4.19}$$

In equilibrium beliefs must consistent, and therefore $\tilde{x}_{2s} = \hat{x}_{2s} = x_{2s}$, $\tilde{x}_{2f} = \hat{x}_{2f} = x_{2f}$ and $\tilde{k}(x_1) = \hat{k}(x_1) = \tilde{k}(x_1)$. We can now use (4.19) to find a solution for x_1. Only for the second condition we can find a solution:

- For the first condition, there is no $k(x_1)$that sets left and right equal. If $x_1 < (1-\beta)\tilde{x}_{2s} + \beta\tilde{x}_{2f}$, and $k(x_1) = x_1 - \beta(\tilde{x}_{2f} - \tilde{x}_{2s})$, then $k(x_1) < x_{2s} = \min\left\{\frac{1}{2}, k(x_1)\right\}$. This means that either $k(x_1) < k(x_1)$ if $k(x_1) \leq \frac{1}{2}$ or $k(x_1) < \frac{1}{2}$ if $k(x_1) > \frac{1}{2}$. Both statements lead to a contradiction.

- For the third condition, there is no $k(x_1)$that sets left and right equal to each other. If $x_1 > x_{2f}$ and $k(x_1) = x_1$, then $k(x_1) > \frac{1+k(x_1)}{2}$. This is only possible if $k(x_1) > 1$, which is outside the domain $[0,1]$, the domain of c_A, and is therefore not possible as solution.

If a solution exists, it must be under the second condition. Setting left and right equal leads Equation (4.11).

$$k(x_1) = \frac{2x_1 - \beta}{2-\beta} \tag{4.20}$$

4.B Additional graphs and tables

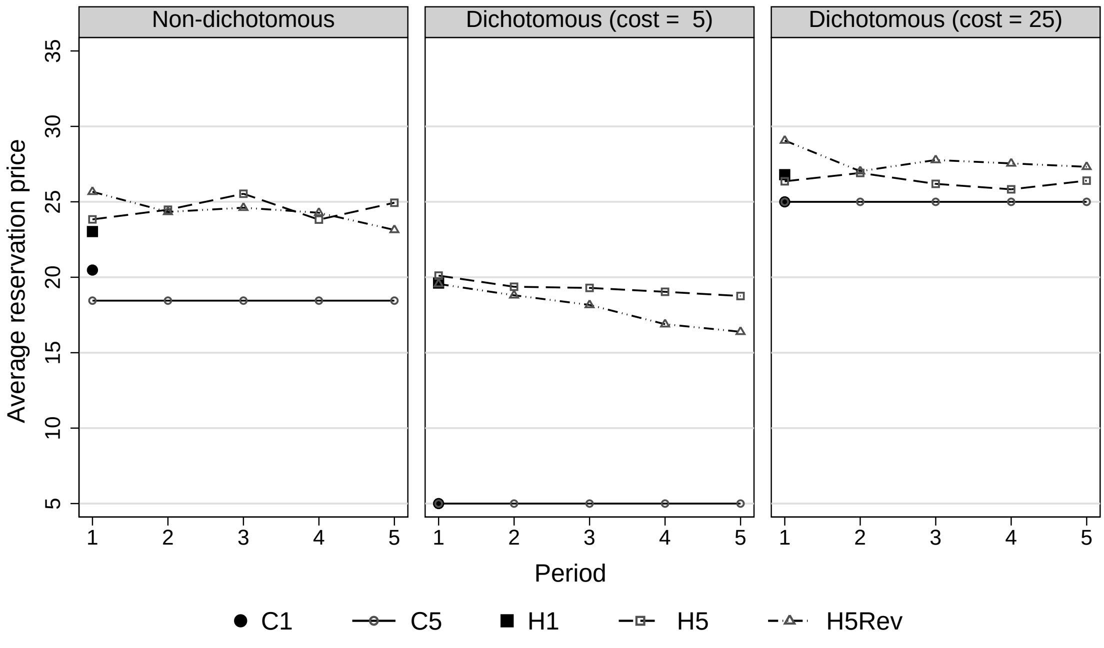

Figure 4.11: Reservation prices in the various treatments.

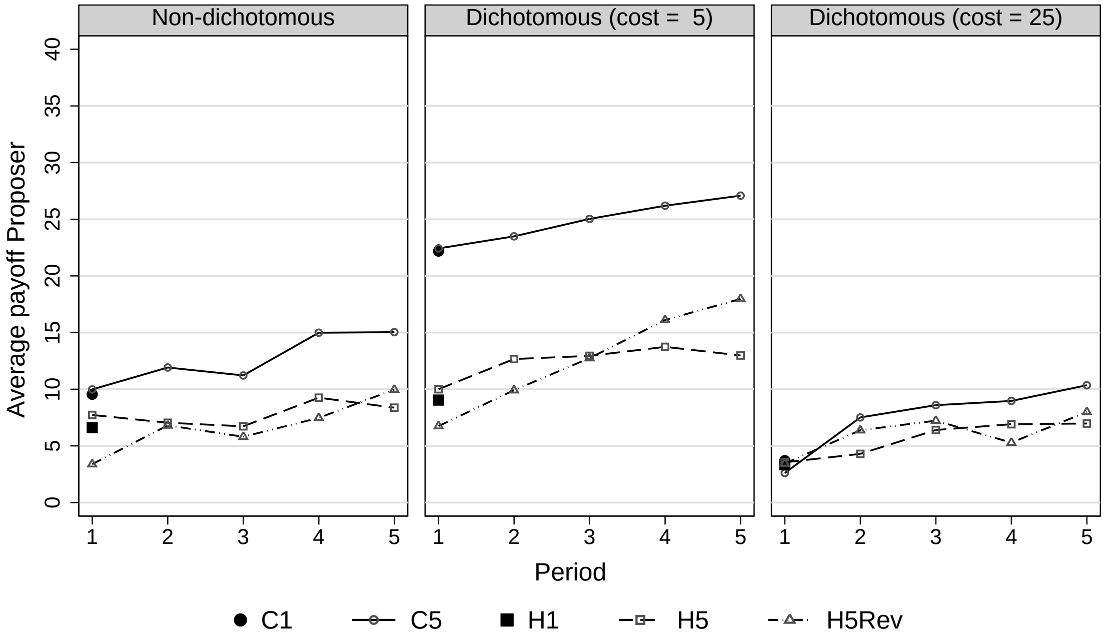

Figure 4.12: Proposer's earnings in the various treatments.

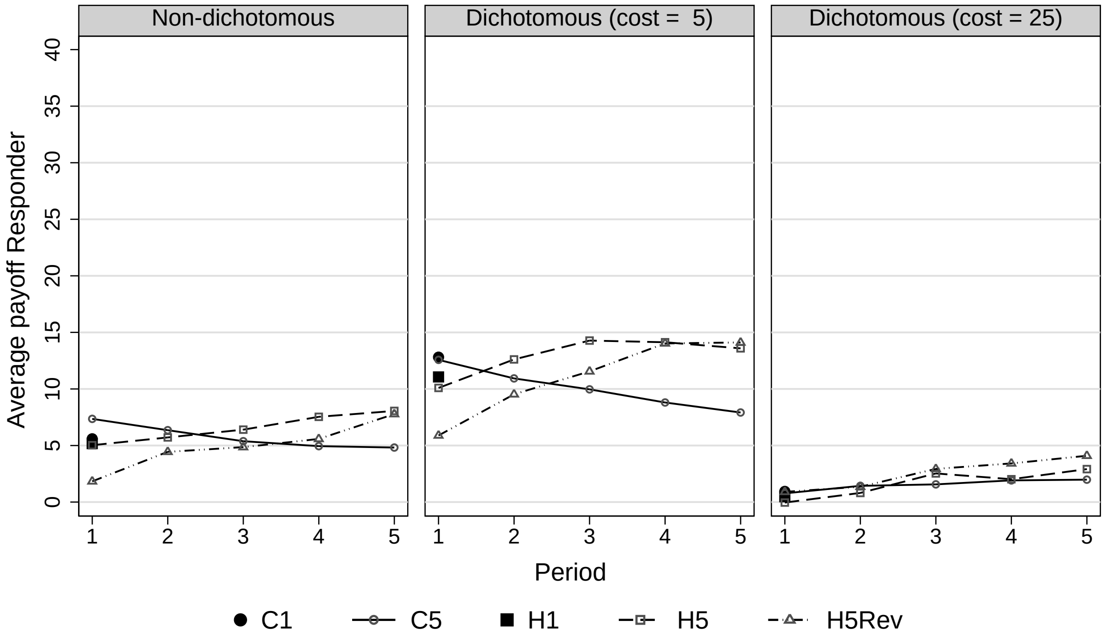

Figure 4.13: Responders's earnings in the various treatments.

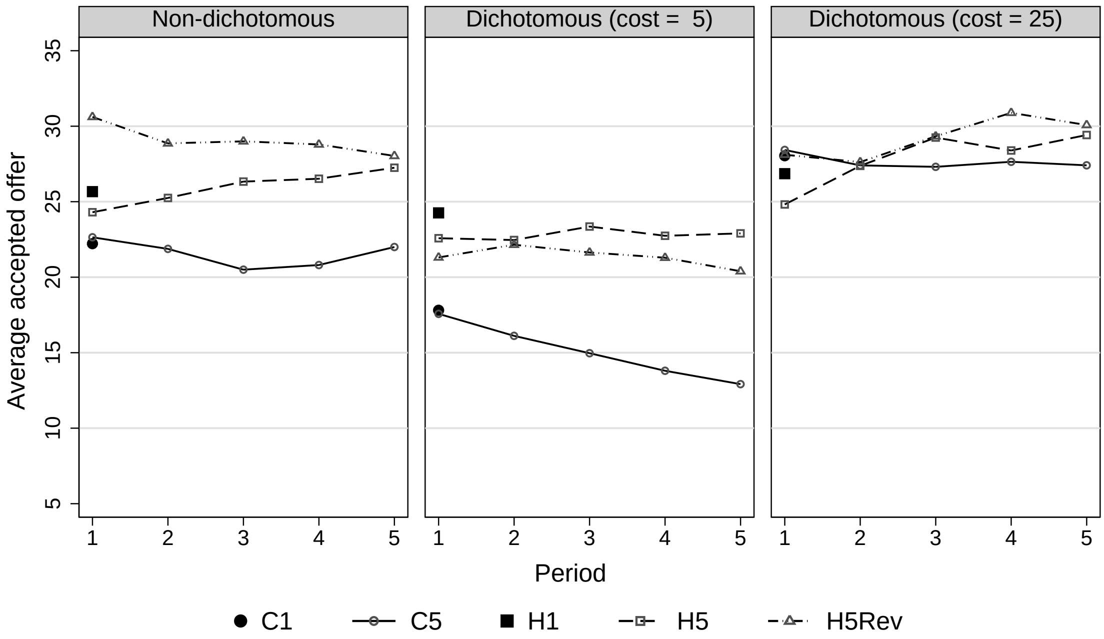

Figure 4.14: Average accepted offers in the various treatments.

4.C Instructions

4.C.1 Introduction screen

In this experiment you are given the role of either the buyer or a seller of a virtual product.

You will be matched up with another participant or a computer player. We will tell you in advance whether you are playing with another human or a computer participant.

You will be playing with this other participant or computer player for one or more periods. After this we will match you with a new player.

You will earn points which will be converted into pounds at the end of the session.

In this game, you will either have the role of a buyer or a seller:

- Each period the buyer offers a price to the seller. (between 0 and 40 points)
- The seller can choose to accept or reject this offer.

  We will ask the seller for the minimum price for which they are willing to sell. If the offer is at or above this minimum price, the offer is accepted. The buyer does not get to know the minimum price.
- If the offer is accepted, the buyer pays the seller the price offered and gets 40 points.
- So the buyer's earnings are: 40 points minus the price.

- Production is costly for the seller. This cost is, at the beginning of each game, randomly chosen by a computer. Each number between 0 and 40 is equally likely. The buyer does not know this cost.
- If the offer is accepted, the seller receives the price offered, and pays the cost of production.
- So the seller's earnings are: the price minus the cost of production.
- If the offer is rejected, both the buyer and the seller don't earn anything.

An example Suppose the seller has a cost of production of 25 points and the buyer offers a price of 35 points.

If the seller accepts the offer, the earnings are:

- **Buyer:** Gets 40 points. Pays the price of 35 points. So the buyer's earnings are: 5 points.
- **Seller:** Receives the price of 35 points. Pays the cost of production of 25 points. So the seller's earnings are: 10 points.

If the seller rejects the offer, both will earn 0 points this period.

About the cost of production You will play multiple games, with different participants or a computer player. Some games will consist of just one period, while others consist of **multiple periods**.

In games with multiple periods, the cost of production will **stay the same** across all periods.

In a new game, the cost of production can change again. It is randomly chosen between 0 and 40 points.

The buyer does not know this cost of production (unless we tell you).

How do we determine your earnings? We will pick at least two games, covering at least six periods, for which we will pay you all the points earned in this game. 100 points are worth 1.50 GBP.

- *Practice game.* We will start with two practice games. The points for these games do not count.

- *Time limit.* Make sure to make your choices within the time limit specified. If we do not hear from you for a longer period of time, we might exclude you from the experiment.

4.C.2 Quiz

The participants are asked to answer the following questions:

1. Suppose the buyer offers a price of 24 points. The seller selects 35 points as their minimum price. What will happen?

    - The seller accepts the offer. The buyer gets 24 points.
    - The seller accepts the offer. The buyer gets 16 points.
    - The seller rejects the offer. The buyer gets 24 points.
    - The seller rejects the offer. Both the buyer and the seller get 0 points. *(Correct answer)*

2. Suppose the seller accepts a price of 24 points. How many points will the buyer earn?

    - 24 points
    - 16 points *(Correct answer)*
    - 0 points

3. Suppose the seller accepts a price of 24 points. Their cost of producing is 5 points. How many points will the seller earn?

    - 24 points
    - 19 points *(Correct answer)*
    - 16 points
    - 5 points

4. The seller's cost of production is between 0 and 40 points, with each number equally likely.

    If a game consists of multiple periods, will this cost of production change every period?

    - Yes, the cost of production changes every period.

- No, the cost of production stays the same in all periods. *(Correct answer)*

5. The seller's cost of production is deducted from the price the seller receives.

   What are the possible values for the cost of production?

   - All values between 0 and 40.
   - 5 points or 25 points
   - 10 points or 30 points
   - 20 points or 40 points

   (Correct answer depends on experiment.)

4.C.3 Screens

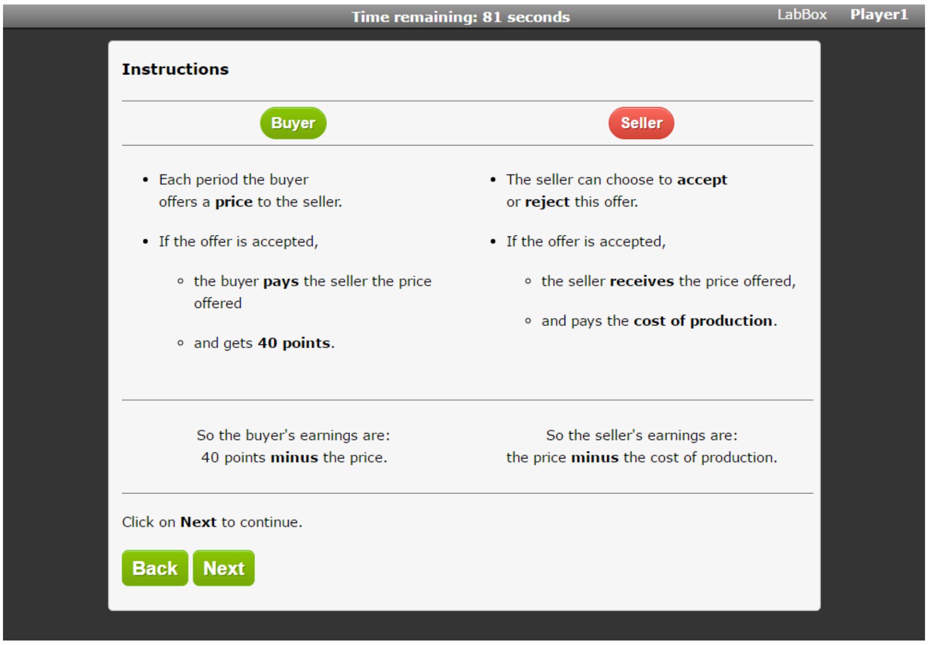

Figure 4.15: Instructions screen at the beginning of the experiment.

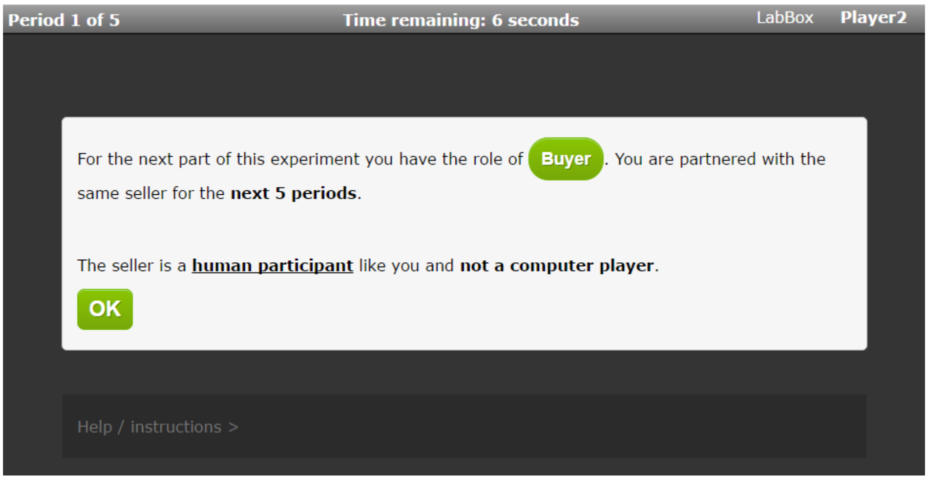

Figure 4.16: The first game screen for the buyer (Proposer).

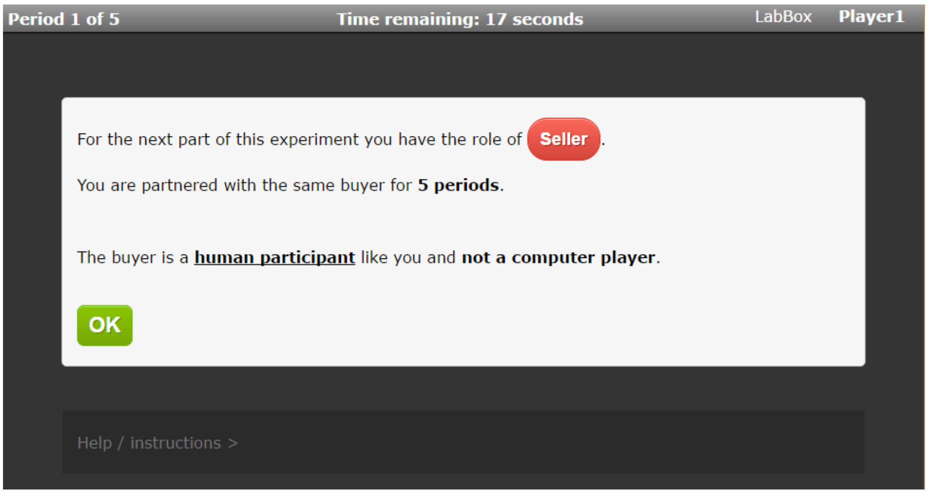

Figure 4.17: The first game screen for the seller (Responder).

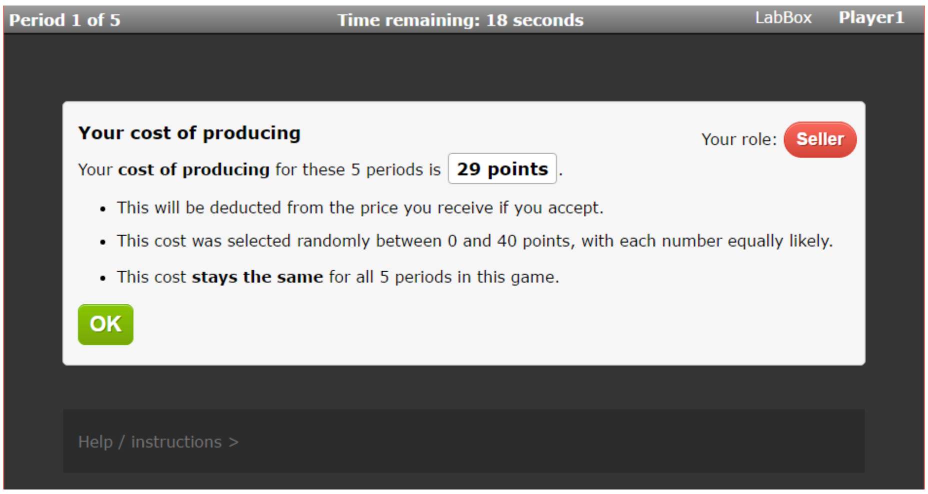

Figure 4.18: Informing the seller (Responder) about the cost of production (the true cost c_A).

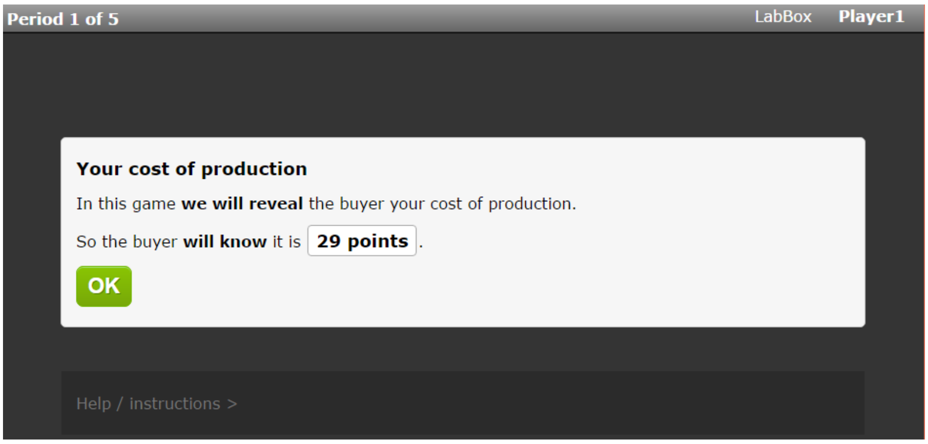

Figure 4.19: Informing the seller (Responder) that the cost of production (the true cost c_A) will be revealed to the buyer (Proposer).

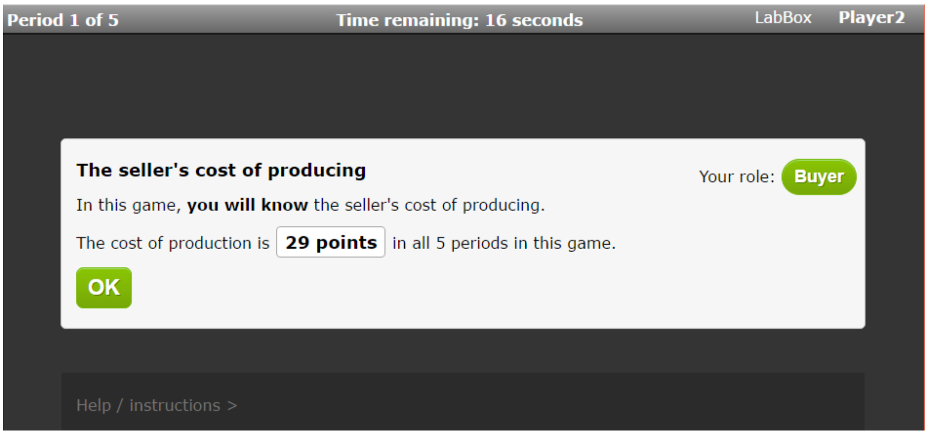

Figure 4.20: Revealing the cost of production (the true cost c_A) to the buyer (Proposer)

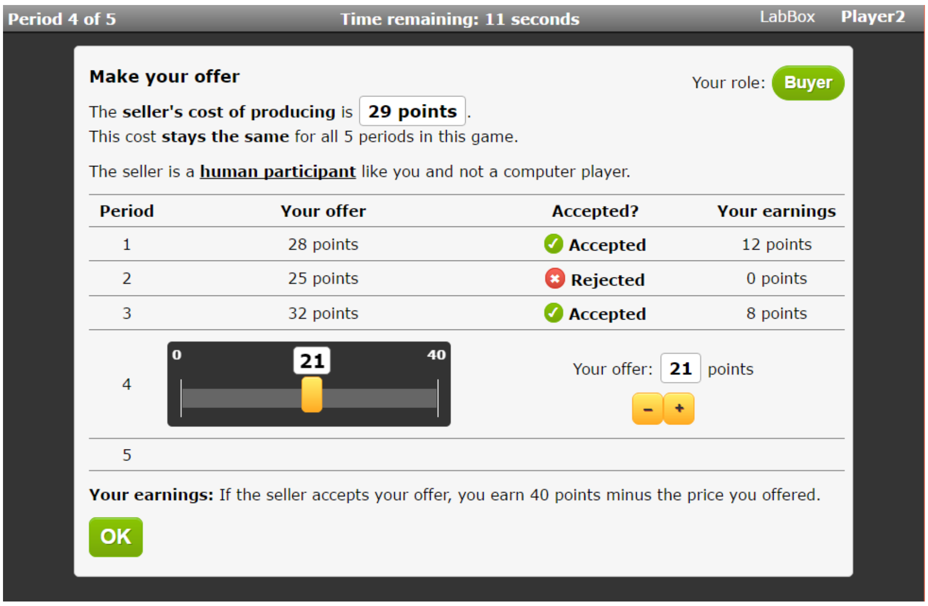

Figure 4.21: The buyer's (Proposer's) offer screen in treatment H5Rev.

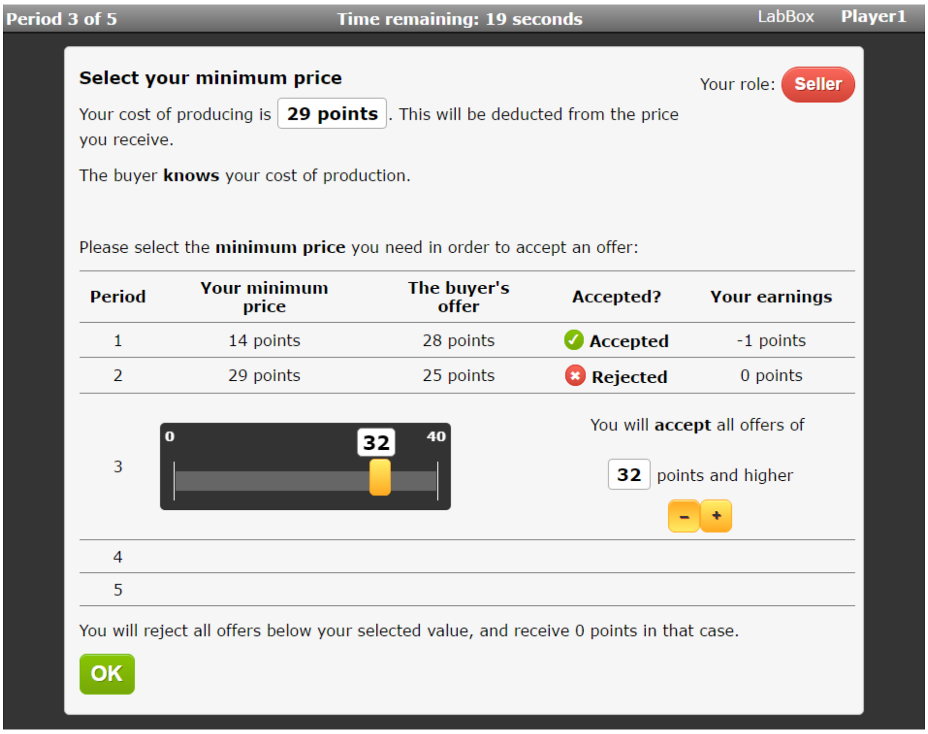

Figure 4.22: The seller's (Responder's) screen where the reservation price can be selected.

Acknowledgements I would like to thank the participants of the Firms & Development Group as well as the Oxford Behavioural Retreat for helpful feedback and comments. Furthermore, I would like to thank Pavel Luengas Sierra for an insightful discussion of this chapter. All remaining errors are my own.

Conclusion

This thesis studied how firms cope with imperfect information and imperfect enforcement and in particular, the role of incentives, reputation and learning. In Chapter 1 I showed that in economic experiments mimicking labour relations employers in Ghana were less likely to use wage incentives than in the United Kingdom. The Ghanaian participants who were assigned a role as employers were less likely to reward high effort or to punish low effort by workers, which resulted in low effort provision. Chapter 2 focused on whether nonmonetary incentives could provide an alternative enforcement mechanism in Ghana, by introducing the possibility for feedback and promises. I showed that feedback can indeed lead to the worker increasing effort, but only when the wage offered is high enough. Chapter 3 presented a theoretical model of the implications of business networks for the overall quality composition of firms in the economy. I showed that business networks, by providing information on its members, can lead to both a filtering effect, meaning that fewer low quality firms produce, and a lock out effect, where members of the network are redistributing demand from outside the network into the network. This means that under certain conditions the overall cost of networking outweighs the overall benefits of networks. Finally, Chapter 4 discussed patterns of learning in a repeated ultimatum game and how bargaining is affected by the presence of strategic considerations. I conducted an experiment, which showed that participants made higher offers to potentially strategic human players than to a myopic computer player. The evidence on strategic considerations is mixed: overall, there is no significant evidence that strategic considerations affect offers and reservation prices, but analyzing the heterogeneity in the sample shows that there is a group of Responders whose behaviour is consistent with strategic rejection.

Future directions The results from this thesis raise important questions for further research, especially the result from Chapter 1 that Ghanaian employers do not use monetary punishment or rewards as an incentive. The limited adoption of punishment and reward incentives could have several reasons, some of which have been discussed in the chapters. For example, I showed that the behaviour was not driven by using student recruits, as actual entrepreneurs acted very similarly. I also tested the effectiveness of nonmonetary incentives in Chapter 2. However, many questions remain that warrant further research.

Employers could refrain from using wage incentives because they do not see such an incentive as legitimate or they believe it is ineffective. Earlier research has shown that incentives are not effective in every context. In particular, incentives that create inequality among workers (see e.g. Breza et al., 2015) or put workers in direct competition with each other (Bandiera et al., 2010) can reduce productivity. Cultural differences as well as differences in the exposure to formal wage employment could also play a role. Together with Marcel Fafchamps and Edoardo Gallo, I have acquired a grant to study this further. Our plan is to conduct both internet as well as lab experiments with representative subject pools to measure beliefs on the legitimacy and the perceived effectiveness of incentives and relate this to experimental behaviour.

One other potential explanation for this behaviour is that employers instead of providing incentives to workers to improve their behaviour, prefer to fire the worker. In our experiment firing was possible, but not a very salient or attractive option: in principle firing was possible by not making an offer to a worker, but this also meant that the employer would not be allocated a worker for the remainder of the game. To study this further, we are planning to conduct additional experiments in which firing is both a more salient and attractive option. We want to test is to what extent the employers and workers see firing as an incentive mechanism or merely as the outcome of screening, by varying the scope for rehiring.

These results also call for more research on actual incentive usage in firms in developing countries to assess the external validity of our results. Current studies draw a mixed picture. There is evidence that bonuses are used in Ghana and these bonuses actually vary by worker (Barr & Serneels, 2009). However, there is some evidence that variations in remuneration in Ghana are often related to being of the same ethnicity of the employer and being related to the employer (Barr & Oduro, 2002). There is some anecdotal evidence suggesting that performance-based bonuses can be controversial (Debrah & Mmieh, 2009). Hence, it is still an unresolved question to what extent these incentives are actually effective in the workplace in Ghana.

The experiments on learning also warrant further research. The current results suggest that the higher-order rationality considerations were not salient to many participants or that they seem to be overpowered by other motivations, such as distributional preferences or reciprocity concerns. One way to proceed further is to run additional experiments where strategic rejection is more salient. One way of doing is this is asking Proposers at the end of the game what they thought the true cost for production of the worker was.

Other ways could be to increase opportunities for deception, for example by allowing the Responder to send a message about their true cost, or allowing participants to observe the play of other participants, to see if strategic thinking can be learnt. Another question is whether people who strategically reject also exhibit strategic anticipation: are participants who strategically reject as Responders also more strategic in their offers as proposers and do their offers take the possibility of strategic rejection into account?

* * *

Understanding how firms use incentives as well as cope with imperfect information is crucial from a policy perspective. As I noted in the introduction, differences in productivity between countries are very large and firm surveys have shown large differences in the way firms are managed. Low worker productivity has contributed to the relative underdevelopment of the manufacturing and agricultural industry (see e.g., Baptist & Teal, 2015). This is a missed opportunity: the experience from Southeast Asia has shown that the manufacturing industry could function as a source of jobs for many in sub-Saharan Africa. In this thesis I have shown that there are differences in the adoption of incentives in an experimental setting and that these differences lead to lower productivity in my experiments in Ghana. Understanding the way incentives work and how they are adopted across different settings is important for industrial policy, but also for firms designing the optimal contract for their employees. Gaining a deeper understanding is therefore not only an important topic for research but also an important question for policy.

References

Abeler, J., Altmann, S., Kube, S., & Wibral, M. (2010). Gift-exchange and worker's fairness concerns: when equality is unfair. *Journal of the European Economic Association*, 8(6), 1299–1324.

Aghion, P., Bolton, P., Harris, C., & Jullien, B. (1991). Optimal learning by experimentation. 58(4), 621–654.

Akerlof, G. (1970). The market for "lemons": Quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3), 488–500.

Akerlof, G. A. (1982). Labor contracts as partial gift exchange. *The Quarterly Journal of Economics*, 97(4), 543–569.

Akerlof, G. A. & Kranton, R. E. (2000). Economics and Identity. *Quarterly Journal of Economics*, 115(3), 715–753.

Altmann, S., Falk, A., Grunewald, A., & Huffman, D. (2014). Contractual Incompleteness, Unemployment, and Labour Market Segmentation. *Review of Economic Studies*, 81(1), 30–56.

Andreoni, J. & Miller, J. H. (1993). Rational Cooperation in the Finitely Repeated Prisoner's Dilemma: Experimental Evidence. *Economic Journal*, 103(418), 570–585.

Avrahami, J., Güth, W., Hertwig, R., Kareev, Y., & Otsubo, H. (2013). Learning (not) to yield: An experimental study of evolving ultimatum game behavior. *Journal of Socio-Economics*, 47, 47–54.

Axelrod, R. (1984). *The Evolution of Co-operation*. New York, NY: Basic Books.

Bandiera, O., Barankay, I., & Rasul, I. (2005). Social preferences and the response to incentives: Evidence from personnel data. *Quarterly Journal of Economics*, 120(3), 917–962.

Bandiera, O., Barankay, I., & Rasul, I. (2009). Social Connections and Incentives in the Workplace: Evidence From Personnel Data. *Econometrica*, 77(4), 1047–1094.

Bandiera, O., Barankay, I., & Rasul, I. (2010). Social incentives in the workplace. *Review of Economic Studies*, 77(2), 417–458.

Baptist, S. & Teal, F. (2015). Why do South Korean firms produce so much more output per worker than Ghanaian ones? *IZA Discussion Paper*, 9157.

Barr, A. (2001). Social Dilemmas and Shame-based Sanctions: Experimental results from rural Zimbabwe. *CSAE Working Paper*, (11).

Barr, A. & Oduro, A. (2002). Ethnic fractionalization in an African labour market. *Journal of Development Economics*, 68(2), 355–379.

Barr, A. & Serneels, P. (2009). Reciprocity in the workplace. *Experimental Economics*, 12(1), 99–112.

Berg, J., Dickhaut, J., & McCabe, K. (1994). Trust, reciprocity and social history. *Games and Economic Behavior*, 10, 122–142.

Bloom, N., Lemos, R., Sadun, R., Scur, D., & Van Reenen, J. (2014). JEEA-FBBVA Lecture 2013: The new empirical economics of management. *Journal of the European Economic Association*, 12(4), 835–876.

Bloom, N. & Reenen, J. V. (2010). Why Do Management Practices Differ across Firms and Countries ? *Journal of Economic Perspectives*, 24(1), 203–224.

Bloom, N. & Van Reenen, J. (2007). Measuring and Explaining Management Practices Across Firms and Countries. *Quarterly Journal of Economics*, 122(4), 1351–1408.

Bohnet, I., Herrmann, B., & Zeckhauser, R. J. (2010). Trust and the reference points for trustworthiness in Gulf and Western countries. *Quarterly Journal of Economics*, 125(2), 811–828.

Boles, T. L., Croson, R. T., & Murnighan, J. K. (2000). Deception and Retribution in Repeated Ultimatum Bargaining. *Organizational behavior and human decision processes*, 83(2), 235–259.

Bolton, P. & Harris, C. (1999). Strategic Experimentation. *Econometrica*, 67(2), 349–374.

Bowles, S. (1985). The Production Process in a Competitive Economy: Walrasian, Marxian, and Neo-Hobbesian Models. *American Economic Review*, 76(1), 16–36.

Brandts, J. & Charness, G. (2004). Do Labour Market Conditions Affect Gift Exchange? Some Experimental Evidence. *The Economic Journal*, 114(July), 684–708.

Brandts, J., Ellman, M., & Charness, G. (2015). Let's Talk: How Communication Affects Contract Design. *Experimental Economics*, 14(4), 943–974.

Breza, E., Kaur, S., & Shamdasani, Y. (2015). The Morale Effects of Pay Inequality. *Mimeo*.

Brown, M., Falk, A., & Fehr, E. (2004). Relational Contracts and the Nature of Market Interactions. *Econometrica*, 72(3), 747–780.

Brown, M., Falk, A., & Fehr, E. (2012). Competition and Relational Contracts: the Role of Unemployment As a Disciplinary Device. *Journal of the European Economic Association*, 10(4), 887–907.

Calvó-Armengol, A. & Jackson, M. O. (2004). The effects of social networks on employment and inequality. *American Economic Review*, 94(3), 426–454.

Camerer, C. & Weigelt, K. (1988). Experimental Tests of a Sequential Equilibrium Reputation Model. *Econometrica*, 56(1), 1–36.

Cardenas, J. C. & Carpenter, J. (2008). Behavioural Development Economics: Lessons from Field Labs in the Developing World. *Journal of Development Studies*, 44(3), 311–338.

Caselli, F. & Coleman, W. J. I. (2006). The World Technology Frontier. *American Economic Review*, 96(3), 499–522.

Charness, G. (2004). Attribution and Reciprocity in an Experimental Labor Market. *Journal of Labor Economics*, 22(3), 665–688.

Charness, G. & Dufwenberg, M. (2006). Promises and partnership. *Econometrica*, 74(6), 1579–1601.

Charness, G. & Dufwenberg, M. (2010). Bare promises: An experiment. *Economics Letters*, 107(2), 281–283.

Charness, G., Gneezy, U., & Kuhn, M. A. (2012). Experimental methods: Between-subject and within-subject design. *Journal of Economic Behavior and Organization*, 81(1), 1–8.

Charness, G. & Haruvy, E. (2002). Altruism, equity, and reciprocity in a gift-exchange experiment: an encompassing approach. *Games and Economic Behavior*, 40(2), 203–231.

Charness, G. & Kuhn, P. (2011). Lab Labor: What Can Labor Economists Learn from the Lab? In *Handbook of Labor Economics*, volume 4 (pp. 229–330). Elsevier Inc.

Chen, D. L., Schonger, M., & Wickens, C. (2016). oTree - An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9, 88–97.

Christen, E. & Francois, J. (2010). Modes of delivery in services. *CEPR Discussion Paper*, (7912).

Coles, M. G. & Smith, E. (1998). Marketplaces and matching. *International Economic Review*, 39(1), 239–254.

Cooper, D. J. & Dutcher, E. G. (2011). The dynamics of responder behavior in ultimatum games: A meta-study. *Experimental Economics*, 14(4), 519–546.

Cooper, D. J. & Lightle, J. P. (2013). The gift of advice: Communication in a bilateral gift exchange game. *Experimental Economics*, 16(4), 443–477.

Costa-Gomes, M., Crawford, V. P., & Broseta, B. (2001). Cognition and Behavior in Normal-Form Games: An Experimental Study. *Econometrica*, 69(5), 1193–1235.

Crawford, V. P., Gneezy, U., & Rottenstreich, Y. (2008). The power of focal points is limited: Even minute payoff asymmetry may yield large coordination failures. *American Economic Review*, 98(4), 1443–1458.

Croson, R. T. (1996). Information in ultimatum games: An experimental study. *Journal of Economic Behavior & Organization*, 30(2), 197–212.

Davies, E. & Fafchamps, M. (2016). Who gains from competition? The ultimatum game in a labour market setting in Ghana. In *Social Economics (CESifo Conference Volume)* chapter 3. MIT Press.

Davies, E. & Fafchamps, M. (2017). When No Bad Deed Goes Punished: Relational Contracting in Ghana versus the UK. *NBER Working Paper*.

Debrah, Y. A. & Mmieh, F. (2009). Employment relations in small- and medium-sized enterprises: insights from Ghana. *The International Journal of Human Resource Management*, 20(7), 1554–1575.

DellaVigna, S., List, J. A., Malmendier, U., & Rao, G. (2016). Estimating Social Preferences and Gift Exchange at Work. *Working paper*.

Duranton, G. & Puga, D. (2004). Micro-foundations of urban agglomeration economies. In V. Henderson & J.-F. Thisse (Eds.), *Handbook of Regional and Urban Economics*, volume 4 (pp. 2063–2117). Amsterdam: North-Holland.

Fafchamps, M. (2002). Spontaneous market emergence. *Topics in Theoretical Economics*, 2(1).

Fafchamps, M. (2004). *Market Institutions in Sub-Saharan Africa: Theory and Evidence*. Cambridge, Massachusetts: MIT Press.

Fafchamps, M. & Söderbom, M. (2006). Wages and Labor Management. *Journal of Human Resources*, 41(2), 346–379.

Falco, P. & Teal, F. (2012). Unemployment, Not Working and Jobs for the Young in Urban Africa: Evidence from Ghana. *Mimeo*.

Falk, A. (2007). Gift exchange in the field. *Econometrica*, 75(5), 1501–1511.

Falk, A., Becker, A., Dohmen, T., Enke, B., Huffman, D., & Sunde, U. (2015). The Nature and Predictive Power of Preferences: Global Evidence. *IZA Discussion Paper*, 9504.

Falk, A., Fehr, E., & Zehnder, C. (2005). Reputation and Performance. *Mimeo*.

Falk, A., Gächter, S., & Kovács, J. (1999). Intrinsic motivation and extrinsic incentives in a repeated game with incomplete contracts. *Journal of Economic Psychology*, 20(3), 251–284.

Falk, A., Huffman, D., & MacLeod, W. B. (2008). Institutions and Contract Enforcement. *IZA Discussion Paper*, 3435.

Fehr, E., Brown, M., & Zehnder, C. (2009). On reputation: A microfoundation of contract enforcement and price rigidity. *Economic Journal*, 119(536), 333–353.

Fehr, E. & Falk, A. (1999). Wage rigidity in a competitive incomplete contract market. *Journal of Political Economy*, 107(1), 106–134.

Fehr, E., Gächter, S., & Kirchsteiger, G. (1997). Reciprocity as a Contract Enforcement Device: Experimental Evidence. *Econometrica*, 65(4), 833–860.

Fehr, E., Kirchler, E., Weichbold, A., & Gächter, S. (1998). When Social Norms Overpower Competition: Gift Exchange in Experimental Labor Markets. *Journal of Labor Economics*, 16(2), 324–351.

Fehr, E., Kirchsteiger, G., & Riedl, A. (1993). Does Fairness Prevent Market Clearing? An Experimental Investigation. *The Quarterly Journal of Economics*, 108(2), 437–459.

Fehr, E., Klein, A., & Schmidt, K. M. (2007). Fairness and Contract Design. *Econometrica*, 75(1), 121–154.

Fehr, E. & List, J. A. (2004). The hidden costs and returns of incentivestrust and trustworthiness among CEOs. *Journal of the European Economic Association*, 2(5), 743–771.

Fehr, E. & Schmidt, K. (2002). Theories of Fairness and Reciprocity - Evidence and Economic Applications. In M. Dewatripont, L. Hansen, & S. Turnovsky (Eds.), *Advances in Economics and Econometrics*, number 75. Cambridge: Cambridge University Press.

Fehr, E. & Schmidt, K. M. (1999). A Theory of Fairness, Competition, and Cooperation. *Quarterly Journal of Economics*, 114(3), 817–868.

Fehr, E., Tougareva, E., & Fischbacher, U. (2014). Do high stakes and competition undermine fair behaviour? Evidence from Russia. *Journal of Economic Behavior and Organization*, 108, 354–363.

Fischbacher, U. (2007). Z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2), 171–178.

Fudenberg, D. & Maskin, E. (1986). The Folk Theorem in repeated games with discounting or with incomplete information. *Econometrica*, 54(3), 533–554.

Gächter, S. & Falk, A. (2002). Reputation and Reciprocity: Consequences for the Labour Relation. *Scandinavian Journal of Economics*, 104(19), 1–27.

Gächter, S. & Fehr, E. (1999). Collective action as a social exchange. *Journal of Economic Behavior & Organization*, 39(4), 341–369.

Gächter, S. & Fehr, E. (2001). Fairness in the Labour Market A Survey of Experimental Results. *Surveys in Experimental Economics. Bargaining, Cooperation and Election Stock Markets*, (114).

Gächter, S., Huang, L., & Sefton, M. (2016). Combining "Real Effort" with induced effort costs: the ball-catching task. *Experimental Economics*, 19(4), 687–712.

Georg, S. J. & Kaiser, J. (2009). Nonparametric testing of distributions - the Epps-Singleton two-sample test using the empirical characteristic function. *Stata Journal*, 9(3), 454–465.

Ghana Statistical Service (2014). *Ghana Living Standards Survey Round 6 (GLSS 6). Labour Force Report*. Technical report.

Ghosh, P. & Ray, D. (1996). Cooperation in Community Interaction without Information Flows. *The Review of Economic Studies*, 63, 491–519.

Glaeser, E. L., Laibson, D. I., Scheinkman, J. A., & Soutter, C. L. (2000). Measuring trust. *Quarterly Journal of Economics*, 115(3), 811–846.

Glaeser, E. L. & Resseger, M. G. (2010). The complementarity between cities and skills. *Journal of Regional Science*, 50(1), 221–244.

Gneezy, U. & List, J. A. (2006). Putting behavioral economics to work: testing for gift exchange in labor markets using field experiments. *Econometrica*, 74(5), 1365–1384.

Gneezy, U. & Rustichini, A. (2000). Pay Enough or Don't Pay at All. *Quarterly Journal of Economics*, 115(3), 791–810.

Goyal, S., van der Leij, M. J., & Moraga-González, J. L. (2006). Economics: An emerging small world. *Journal of Political Economy*, 114(2), 403–412.

Granovetter, M. (1973). The strength of weak ties. *American Journal of Sociology*, 78(6), 1360–1380.

Granovetter, M. (2005). The impact of social structure on economic outcomes. *Journal of Economic Perspectives*, 19(1), 33–50.

Green, F. & Weisskopf, T. E. (1990). The worker discipline effect: A disaggregative analysis. *The Review of Economics and Statistics*, 72(2), 241–249.

Greif, A. (2000). The fundamental problem of exchange: A research agenda in Historical Institutional Analysis. *European Review of Economic History*, 4(3), 251–284.

Greif, A., Milgrom, P., & Weingast, B. (1994). Coordination, Commitment, and Enforcement: The Case of the Merchant Guild. *Journal of Political Economy*, 102(4), 745–776.

Grossman, G. & Horn, H. (1988). Infant-industry protection reconsidered: the case of informational barriers to entry. *Quarterly Journal of Economics*, 103(4), 767–787.

Güth, W. & Kocher, M. G. (2014). More than thirty years of ultimatum bargaining experiments: Motives, variations, and a survey of the recent literature. *Journal of Economic Behavior and Organization*, 108, 396–409.

Güth, W. & Tietz, R. (1990). Ultimatum Bargaining Behaviour: A survey and comparison of experimental results. *Journal of Economic Psychology*, 11, 417–449.

Hall, R. E. & Jones, C. I. (1999). Why do some countries produce so much more output per worker than others? *Quarterly Journal of Economics*, 114(1), 83–116.

Hargreaves Heap, S., Rojo Arjona, D., & Sugden, R. (2014). How Portable is Level-0 Behaviour? A Test of Level-K Theory in Games with Non-neutral Frames. *Econometrica*, 82(3), 1–27.

Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *The Behavioral and Brain Sciences*, 33, 61–83; discussion 83–135.

Henrich, J., McElreath, R., Barr, A., Ensminger, J., Barrett, C., Bolyanatz, A., Cardenas, J. C., Gurven, M., Gwako, E., Henrich, N., Lesorogol, C., Marlowe, F., Tracer, D., & Ziker, J. (2006). Costly punishment across human societies. *Science*, 312(5781), 1767–70.

Huck, S. & Lünser, G. K. (2010). Group reputations: An experimental foray. *Journal of Economic Behavior and Organization*, 73(2), 153–157.

Ismayilov, H. & Potters, J. (2015). Why do promises affect trustworthiness, or do they? *Experimental Economics*, 19(2), 382–393.

Jackson, M. O. (2008). *Social and Economic Networks*. Princeton, NJ and Oxford, United Kingdom: Princeton University Press.

Jackson, M. O. & Rogers, B. W. (2007). Meeting strangers and friends of friends: how random are social networks? *American Economic Review*, 97(3), 890–915.

Jones, R. & Ruane, F. (1990). Appraising the options for international trade in services. *Oxford Economic Papers*, 42(4), 672–687.

Kagel, J. H., Kim, C., & Moser, D. (1996). Fairness in Ultimatum Games with Asymmetric Information and Asymmetric Payoffs. *Games and Economic Behavior*, 13(1), 100–110.

Kahn, L. M. & Murnighan, J. K. (1993). A General Experiment on Bargaining in Demand Games with Outside Options. *American Economic Review*, 83(5), 1260–1280.

Karlan, D. S. (2005). Using Experimental Economics to Measure Social Capital and Predict Financial Decisions. *American Economic Review*, 95(5), 1688–1699.

Keller, G. & Rady, S. (1999). Optimal Experimentation in a Changing Environment. *Review of Economic Studies*, 66(3), 475–507.

Keller, G., Rady, S., & Cripps, M. (2005). Strategic Experimentation with Exponential Bandits. *Econometrica*, 73(1), 39–68.

Kirchler, E., Fehr, E., & Evans, R. (1996). Social exchange in the labor market: Reciprocity and trust versus egoistic money maximization. *Journal of Economic Psychology*, 17(3), 313–341.

Kneeland, T. (2015). Identifying Higher-Order Rationality. *Econometrica*, 83(5), 2065–2079.

Knez, M. J. & Camerer, C. F. (1994). Outside Options and Social Comparison in Three-Player Ultimatum Game Experiments.pdf. *Games and Economic Behavior*, 10, 65–94.

Kranton, R. & Minehart, D. (2001). A theory of buyer-seller networks. *American Economic Review*, 91(3), 485–508.

Kranton, R. E. (1996). Reciprocal Exchange: A Self-Sustaining System. *American Economic Review*, 86(4), 830–851.

Kreps, D. M., Milgrom, P., Roberts, J., & Wilson, R. (1982). Rational Cooperation in the Finitely Dilemma Repeated Prisoners' Dilemma. *Journal of Economic Theory*, 252(27), 245–252.

Kube, S., Maréchal, M. A., & Puppe, C. (2006). Putting Reciprocity to Work: Positive versus Negative Responses in the Field. *Universität St. Gallen Discussion Paper*, (2006-27).

Loginova, O. & Taylor, C. R. (2008). Price experimentation with strategic buyers. *Review of Economic Design*, 12(3), 165–187.

Macchiavello, R. & Morjaria, A. (2015). The value of relational contracts : evidence from a supply shock to Kenyan flower exporters. *American Economic Review*, 105(9), 2911–2945.

MacLeod, W. B. (2007). Reputations, Relationships, and Contract Enforcement. *Journal of Economic Literature*, 45(3), 595–628.

Masclet, D., Noussair, C. N., Tucker, S., & Villeval, M.-C. (2003). Monetary and Nonmonetary Punishment in the Voluntary Contributions Mechanism. *American Economic Review*, 93(1), 366–380.

McMillan, J. & Woodruff, C. (1999). Interfirm Relationships and Informal Credit in Vietnam. *Quarterly Journal of Economics*, 114(4), 1285–1320.

Milgrom, P. R., North, D. C., & Weingast, B. R. (1990). The Role of Institutions in the Revival of Trade: the Law Merchant, Private Judges, and the Champagne Fairs. *Economics and Politics*, 2(1), 1–23.

Mitzkewitz, M. & Nagel, R. (1993). Experimental results on ultimatum games with incomplete information. *International Journal of Game Theory*, 22, 171–198.

Montgomery, J. (1991). Social networks and labor-market outcomes: Toward an economic analysis. *American Economic Review*, 81(5), 1408–1418.

Morck, R., Wolfenzon, D., & Yeung, B. (2005). Corporate Governance, Economic Entrenchment, and Growth. *Journal of Economic Literature*, 43, 655–720.

Pereira, P. T., Silva, N., & Andrade e Silva, J. (2006). Positive and negative reciprocity in the labor market. *Journal of Economic Behavior and Organization*, 59(3), 406–422.

Prendergast, C. (1999). The Provision of Incentives in Firms. *Journal of Economic Literature*, 97(1), 7–63.

Rapoport, A. & Sundali, J. A. (1996). Ultimatums in Two-Person Bargaining with One-Sided Undertainty: Offer Games. *International Journal of Game Theory*, 30, 475–494.

Rauch, J. (2001). Business and social networks in international trade. *Journal of Economic Literature*, 39(4), 1177–1203.

Rauch, J. E. (1999). Networks versus markets in international trade. *Journal of International Economics*, 48(1), 7–35.

Rauch, J. E. & Casella, A. (2003). Overcoming Informational Barriers to International Resource Allocation: Prices and Ties. *Economic Journal*, 113(484), 21–42.

Rege, M. & Telle, K. (2004). The impact of social approval and framing on cooperation in public good situations. *Journal of Public Economics*, 88(7-8), 1625–1644.

Rogerson, R., Shimer, R., & Wright, R. (2005). Search-Theoretic Models of the Labor Market: A Survey. *Journal of Economic Literature*, 43(4), 959–988.

Roth, A. E., Prasnikar, V., Okuno-Fujiwara, M., & Zamir, S. (1991). Bargaining and Market Behavior in Jerusalem, Ljubljana , Pittsburgh, and Tokyo: An Experimental Study. *American Economic Review*, 81(5), 1068–1095.

Selten, R. & Stoecker, R. (1986). End behavior in sequences of finite Prisoner's Dilemma supergames. A learning theory approach. *Journal of Economic Behavior & Organization*, 7(1), 47–70.

Shapiro, C. (1983). Premiums for high quality products as returns to reputations. *Quarterly Journal of Economics*, 98(4), 659–679.

Shapiro, C. & Stiglitz, J. E. (1984). Equilibrium as a Worker Unemployment Discipline Device. *American Economic Review*, 74(3), 433–444.

Siang, C. K., Requate, T., & Waichman, I. (2011). On the role of social wage comparisons in gift-exchange experiments. *Economics Letters*, 112(1), 75–78.

Slembeck, T. (1999). Reputations and Fairness in Bargaining - Experimental Evidence from a Repeated Ultimatum Game With Fixed Opponents. *Universität St. Gallen Discussion Paper*, (9904).

Slonim, R. & Roth, A. E. (1998). Learning in high stakes ultimatum games: an experiment in the Slovak Republic. *Econometrica*, 66(3), 569–596.

Spence, M. (1973). Job market signaling. *Quarterly Journal of Economics*, 87(3), 355–374.

Storper, M. & Venables, A. J. (2004). Buzz: face-to-face contact and the urban economy. *Journal of Economic Geography*, 4(4), 351–370.

Straub, P. G. & Murnighan, J. K. (1995). An experimental investigation of ultimatum games: information, fairness, expectations, and lowest acceptable offers. *Journal of Economic Behavior and Organization*, 27(3), 345–364.

Uzzi, B. (1996). The sources and consequences of embeddedness for the economic performance of organizations: the network effect. *American Sociological Review*, 61(4), 674–698.

Van Der Heijden, E. C. M., Nelissen, J. H. M., Potters, J. J. M., & Verbon, H. A. A. (2001). Simple and complex gift exchange in the laboratory. *Economic Inquiry*, 39(2), 280–297.

Vanberg, C. (2008). Why Do People Keep Their Promises? An Experimental Test of Two Explanations. *Econometrica*, 76(6), 1467–1480.

Venables, A. J. (2011). Productivity in cities: self-selection and sorting. *Journal of Economic Geography*, 11(2), 241–251.

Wadhwani, S. B. & Wall, M. (1991). A Direct Test of the Efficiency Wage Model Using UK Micro-Data. *Oxford Economic Papers*, 43(4), 529–548.

Weisbuch, G., Kirman, A., & Herreiner, D. (2000). Market organisation and trading relationships. *The Economic Journal*, 110(463), 411–436.

Wu, S. Y. & Roe, B. (2007a). Contract enforcement, social efficiency, and distribution: Some experimental evidence. *American Journal of Agricultural Economics*, 89(2), 243–258.

Wu, S. Y. & Roe, B. E. (2007b). Discretionary Latitude and Relational Contracting. *IZA Discussion Paper*, 2879.

Xiao, E. & Houser, D. (2005). Emotion expression in human punishment behavior. *Proceedings of the National Academy of Sciences of the United States of America*, 102(20), 7398–401.