Experimental Markets with Frictions

by
Knut-Eric Neset Joslin

Knut-Eric Neset Joslin
Experimental Markets with Frictions:

Contents

Acknowledgments

Introduction

Summary

This dissertation investigates how individuals behave in markets with frictions. Search and information frictions are present in many markets and it is critical to develop useful representations of behavior in such settings.

The four studies included in this dissertation combine theory and laboratory experiments. These studies test prevailing models of markets with frictions. Overall, we find that standard theories often perform well, especially in search markets. This suggests that these settings exhibit market forces that create tendencies for participants to engage in specific types of behavior. This is a deep result that implies that these models have a high degree of internal validity. Moreover, this result has value for further theorizing and, in general, for thinking systematically about these types of markets.

The reader may ask what in particular is achieved by combining economic theory with laboratory experiments. Economic theory provides insights into how individuals interact with each other and institutions, and what consequences this has for society. Although pure theory helps organize our thinking, the value of a model is strengthened if it is corroborated by evidence from actual decision-making. Traditionally, such evidence has come from econometric studies. Econometric studies have the strength that the conclusions are derived from real-world decisions. However, econometric studies also encounter a variety of limitations. At a basic level, econometric studies are limited by the data that are available. Furthermore, even when rich data are available, many of the objects of interest may be inherently difficult or impossible to observe.

Experimental tests using laboratory techniques provide an alternative and complementary type of evidence. Experimental testing enables researchers to perform a close test of the basic mechanisms at work in a model. This provides evidence that may be difficult to

obtain by other means. This is especially useful in the study of frictional markets in which vital aspects of decisions—such as the available information—tend to be un-observable to the researcher. Moreover, because the researcher exerts considerable control over the experimental environment, it enables a close alignment between the theory and the test of the theory. This is important for assessing the internal validity of the model. The purpose is to test how well a model explains the observed results at both the individual level and the aggregate level. A fascinating finding from experimental work—that is reproduced in this dissertation—is that a model can perform well at an aggregate level even if individual behavior deviates from predictions.

A challenge with experimental testing is how to re-present models in the lab. It is often necessary to pare away features of theoretical models in order to make them appropriate for lab testing. For example, it is impossible to directly implement an infinite horizon model in a lab setting. In each experiment included in this dissertation, the model taken to the lab is a simplified version of a more complex model. However, in each case, key trade-offs are preserved. The analogy between the situation of interest and the environment that is tested is preserved—even if some details of models are omitted.

Papers

This dissertation comprises four papers, each presented in a Chapter: (1) *Market Entry with Frictional Matching and Bargaining: Labor Search in the Lab* (2) *Discrimination in Small Markets with Directed Search: Part I Theory* (3) *Discrimination in Small Markets with Directed Search: Part II Experiment* and (4) *Should I Stay or Should I Go: Bandwagons in the Lab*.

The first paper, *Market Entry with Frictional Matching and Bargaining: Labor Search in the Lab*, tests elements from the standard labor search model using a modified market entry game that includes labor market features. We reproduce key comparative statics in the lab and conclude that the no-profit condition at the heart of the standard labor search model is nearly satisfied. This is an important result for the credibility of the model. The one deviation that we observe is that vacancy creation does not respond quite as strongly to the level of productivity as predicted.

The second two papers are part of a single project, *Discrimination in Small Markets with Directed Search*. This project investigates how discriminatory hiring impacts the black-white wage gap in markets with directed search. The first paper develops the theory while the second paper takes a simple version of the model to the lab. Although we find that

search behavior in this setting aligns well with theory, the anticipated discrimination effect is not as strong as predicted by the model. Nevertheless, the study does substantiate a segregation effect. Relative to the case without discrimination, more firms offer low wages and discriminated workers earn lower income.

The last paper included in the dissertation, *Should I Stay or Should I Go: Bandwagons in the Lab*, is co-authored with Leif Helland and Tom-Reiel Heggedal. This project tests a seminal model of platform coordination. We find that the equilibrium of the model effectively predicts behavior. When we in addition allow for some noise in the equilibrium concept, the model matches observation closely.

Background: Markets with Frictions

The classical theory of markets assumes that agents can costlessly meet each other and that there are no information problems. The properties of such markets are well understood: When there are many agents active in the market, equilibrium between supply and demand leads to a single "competitive" price. Because agents can immediately and costlessly transact at this price, these models are described as "frictionless." Even though the assumptions of the frictionless model are rarely—if ever—satisfied exactly, the frictionless model often performs well.

Despite its successes, however, there are many phenomenon of central economic importance that can not be accommodated, or only accommodated with great difficultly, in the frictionless framework. This includes such basic phenomenon as the coexistence of unemployed workers and vacant jobs, price heterogeneity, the use of money, and even the existence of firms. This has necessitated other models, with conclusions derived from other premises.

This dissertation investigates such models, which (in contradistinction to the frictionless model) are referred to as models of markets with frictions. In particular, the dissertation addresses models that deviate from the frictionless paradigm by relaxing the assumption that trading partners instantly meet each other and that they have access to all relevant information.

Search Frictions A basic observation is that most markets operate in a decentralized fashion. Individuals and firms are not instantaneously matched with appropriate trading partners. Rather, matching usually involves both pecuniary and time costs. That these costs can be significant should be apparent to anybody who has participated in the labor

market or the housing market. In the economic parlance, we refer to these costs as *search frictions*.

The presence of search frictions re-configures how we think about markets relative to the classical model. Rather than a fluid and impersonal exchange of goods, goods are traded when suitable individuals encounter each other and agree on terms of trade. By design, this accommodates the co-existence of searching agents. It also emphasizes that trade has a personal dimension. Individuals do not simply receive a "market price." Rather, prices arise at the junctures between agents. Search models can thereby account for phenomenon such as price heterogeneity. It also highlights the fact that there is surplus associated with the meeting of *specific* agents and that this surplus is the opportunity cost of further search.

Two common ways of representing search frictions are *random search* and *directed search*. Chapter 1 (*Market Entry with Frictional Matching and Bargaining*) investigates a model of the former while Chapter 2 and Chapter 3 (*Discrimination in Small Markets with Directed Search*) are representative of the latter approach. While both types of models address search frictions, they have a different conceptual basis and are implemented differently in a modelling framework.

Random search models are characterized by the assumption that in decentralized, anonymous markets, appropriate trading partners encounter each other randomly. It can be reasonable to represent job matching in this fashion: It is a chance outcome that a particular worker applies to a particular job and is selected from among a particular set of applicants.

An implication of random search is that the terms of trade can only be agreed after matching: The parties in the transaction must agree how the idiosyncratic match surplus should be divided. Although the ability to negotiate outcomes is circumscribed by market conditions—the ease of matching, discounting, etc.—the division of the surplus ultimately entails some type of bargaining. There is thus a non-market process at the core of random search models. This provides intuition for why models of random search need not be efficient, even in a constrained sense. Bargaining outcomes need not be aligned with socially optimal participation in the market.

The effects of bargaining institution on market participation is a focus of the first study in this dissertation. This study demonstrates in a simple experiment that the allocation of bargaining power has substantial consequences for market outcomes exactly because the terms of trade can only be agreed after search. Test subjects face a hold up problem that is exacerbated when the bargaining institution favors one side of the market.

In contrast to models of random search, in directed search models, agents observe the

terms of trade (i.e. often a price) prior to search. Frictions arise in this setting because of congestion and the possibility of rationing. Directed search models thus have the attractive feature that prices play an important role. Agents trade off prices and the probability of transacting. The experimental work in Chapter 3 demonstrates that test subjects do in fact adjust their behavior in a fashion consistent with this principle. Optimal search equalizes the expected income available among all sellers conditional on the seller actually receiving applications.

Information Frictions Another source of frictions are so-called information frictions.[1] When the quality of a good or the quality of a relationship are uncertain, this introduces costs (or potential costs) that are not present in the frictionless model. These costs may arise because of the presence of private information or because the degree of complementarity between trading partners is revealed over time. For example, when purchasing an automobile it may be costly to verify that it is not a "lemon." Similarly, when hiring an employee, firms often use considerable time and resources to figure out if the employee is a good match.

In many cases, it is useful to explicitly incorporate information frictions into a model using the tools of game theory. This is the approach taken in Chapter 4 (*Should I Stay or Should I Go?*). In this study, agents have private information about payoffs but these payoffs are also determined by a positive network externality. This model captures in a reduced form way the tension between competition and cooperation that is present in many market situations. The findings from this study suggest that test subjects can indeed be modelled as using cut-off strategies, with only minor deviations from the predictions of the theoretical model.

Theory, Representation, and Economic Experiments

The stated goal of the dissertation is to test models of markets with frictions. But it may be unclear what is meant by "test." What is being tested? And what is learned from such an exercise?

Economic Theory As alluded to in the previous subsection, there is rarely a one-to-one mapping between reality and the assumptions that underpin economic models. The strict

---

[1]Random search models can also be thought of as representing information frictions, albeit in a reduced form fashion.

assumptions of models, which extend to both the environment and to the preferences of agents in the model, are rarely consistent with the underlying reality. Likewise, there are many aspects of real markets not accounted for by economic theory. Economic models are at best simplified representations or idealizations of economic interactions. For instance, in his monograph on the standard search and matching model, Pissarides (2000) explicitly acknowledges that his model is unrealistic in key respects.

Economic theorizing is therefore about *useful representation*. Of course, usefulness may be judged by various standards. At the most abstract level, economic models provide a disciplined way of reasoning. Even if a model is unrealistic, it may be useful because it clarifies our thinking. Another standard of usefulness is the quality of predictions. For example, the Nash equilibria identified in search models often predict experimental outcomes even though test participants almost certainly do not have the ability to compute the equilibrium (See Helland et al. 2017, for a remarkable example).

A pragmatic position—and the position taken in this study—is that an economic model is useful if it captures something substantial about how individuals behave when confronted by a particular economic institution. In particular, if the model has comparative statics implications that are borne out, then this is useful information about economic behavior. A "test" of theory can therefore entail an assessment about how test subjects adjust their behavior under different sets of incentives, not just a comparison with exact equilibrium predictions.

Experimental Economics   Economic experiments have the ability to create reproducible behavior. Typically, they also facilitate a closer approximation of the environment specified in economic theory than occurs in real world markets. The ability to manipulate the environment, moreover, enables clear comparative statics tests. Economic experiments are thereby able to produce some of the most direct and dramatic tests of economic theory.

Experimental evidence is complementary to both theory and econometric studies. Results in theoretical or econometric work are far more credible if they are consistent with the behavioral results found in properly designed experimental studies. As illustrated by policy-oriented economists such as Alvin Roth and Charles Holt, experimental studies help us understand real-world markets and can be used to test policies before they are implemented "in the real world."

The studies in this dissertation employ economic experiments to generate insights into how frictional markets behave and how individuals in such markets make decisions. This

includes direct tests of equilibrium predictions as well as comparative statics analyses. This contributes to a small but growing experimental literature that has focused on frictional markets (See Helland et al. 2017; Kloosterman 2016, for prime examples.). Overall, there is (1) a high degree of consistency between the findings in this dissertation and theoretical predictions, and (2) a high degree of consistency with earlier experimental studies. A conclusion seems to be that the theoretical models that we investigate do usefully represent economic behavior in the domains of interest.

Methodology

Over the past few decades, a set of norms regarding the conduct of economic experiments has grown up. The norms include lack of deception, anonymity of test subjects, and financial incentives. As pointed out by Camerer et al. (2016) in a survey study, a shared set of norms may be a reason why experimental economics results have proven to be reproducible to a greater extent than econometric studies and experimental studies in other social sciences. There has also been widespread collaboration among researchers in the form of sharing instructions and program files. The use of *zTree* in particular has made it possible for experimental economists to easily develop and exchange standard game modules (See Fischbacher 2007, for a description of *zTree.*).

In recent years, greater attention has also been devoted to the design of experimental studies. For example, there has been considerable attention to issues related to the power of statistical tests and to the relative merits of between versus within subject designs (List et al. 2011; Charness et al. 2012). The codification of standards for experimental design represents a maturation of experimental economics as a scientific enterprise.

The experimental portions of this project have followed these "state of the art" principles as closely as possible. For example, many of the key statistical tests are performed on independent block-level observations using non-parametric approaches, taking account of the requirements of statistical power. The expectation is that the findings in this dissertation are credible and would be robust to reproduction by other researchers.

References

Camerer, C. F., Dreber, A., Forsell, E., Ho, T.-H., Huber, J., Johannesson, M., . . . Wu, H. (2016). Evaluating replicability of laboratory experiments in economics. *Science.*

Charness, G., Gneezy, U., & Kuhn, M. A. (2012). Experimental methods: Between-subject and within-subject design. *Journal of Economic Behavior & Organization, 81*(1), 1–8.

Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics, 10*(2), 171–178.

Helland, L., Moen, E. R., & Preugschat, E. (2017). Information and coordination frictions in experimental posted offer markets. *Journal of Economic Theory, 167*, 53–74.

Kloosterman, A. (2016). Directed search with heterogeneous firms: An experimental study. *Experimental Economics, 19*(1), 51–66.

List, J. A., Sadoff, S., & Wagner, M. (2011). So you want to run an experiment, now what? some simple rules of thumb for optimal experimental design. *Experimental Economics, 14*(4), 439–457.

Pissarides, C. A. (2000). Equilibrium unemployment theory. *MIT Press Books, 1*.

Chapter 1.

Market Entry with Frictional Matching and Bargaining: Labor Search in the Lab

Abstract

This paper studies an experimental labor market that incorporates elements from the standard theory of equilibrium unemployment. Specifically, we test in a controlled lab setting a novel market entry game that includes matching frictions and wage bargaining. The model predicts that firms will enter up to the point at which stochastic rationing of workers equalizes the value of a vacancy with its costs. Between treatments, we vary productivity and bargaining strength. Consistent with theory, we find that increases in productivity increase job creation and thereby reduce unemployment. We also reproduce the expected outcomes associated with different forms of wage negotiation. When wages are determined by bargaining after match, firms face a hold-up problem. As a consequence, job creation collapses when workers have excessive bargaining power. In contrast, when wages are determined prior to entry, workers moderate their wage claims to induce vacancy creation. Although our findings tend to align with theory, we observe some deviations. In particular, there is a systematic bias in the aggregate level of entry: There is too much vacancy creation when productivity is low and too little vacancy creation when productivity is high. To explain this bias and to account for heterogeneity at the individual level, we estimate a quantal response equilibrium in which we allow for idiosyncratic preferences for entry.

Author: Knut-Eric N. Joslin [1]

Keywords: market entry, bargaining, search, labor economics, experiment.

JEL Classification: C78, J64.

1.1 Introduction

In this study, we investigate an experimental labor market that incorporates firm entry, frictional matching, and wage bargaining. The purpose of this study is to scrutinize in a controlled laboratory environment some of the main assumptions that underpin standard labor market search models.[2] In particular, we test components of the Diamond-Mortensen-Pissarides (DMP) model of equilibrium employment. Our experiment thus has a macroeconomic motivation.[3] Although the model we take to the lab is simple, it enables us to confront participants with an environment that closely matches theory. The principle research question we address is how vacancy creation varies in response to changes in productivity and the structure of wage bargaining. We also investigate a set of subsidiary research questions that relate to individual behavior when trade is mediated by a frictional process and negotiated via bargaining.

A key feature of trade in the labor market is that it is decentralized and uncoordinated. Because jobs and workers are heterogeneous, firms must invest resources to identify and recruit suitable candidates. A tractable way of representing these search frictions is via the use of a *matching function*. This modelling device is at the core of the Diamond-Mortensen-Pissarides model of equilibrium unemployment (Pissarides 2000). The matching function gives a number jobs formed as a function of the number of job vacancies and the number of unemployed workers. Matching is frictional because the jobs are allocated randomly among the searching agents. This may be thought of as an urn-ball process in which applications correspond to balls and vacancies as jobs. Agents of a given type thus impose congestion externalities on each other and labor market participants face risk due to stochastic rationing that depends on market tightness. In particular, when the number of firms increases relative to the number of workers, matching probabilities for firms decrease.

Another characteristic of trade in the labor market is the presence of employment contracting. For production to take place, employers and potential employees must negotiate a wage. Because it is time-consuming and costly to recruit workers, good matches will be associated with a productive surplus. Trade in the labor market can thus be conceptualized as a two-stage process in which jobs and workers are randomly matched together in

---

[2]The use of the lab to study the labor market is well established. See Charness and Kuhn (2011) and Falk and Fehr (2003) for overviews of this approach.

[3]For discussion of how experimental economics can be used to understand macroeconomic markets see Noussair et al. (1995). Duffy (2008) provides a comprehensive review of the experimental macroeconomics literatures.

the first stage and then make a wage agreement that divides a match surplus in the second stage.[4] Search equilibrium of this type will typically be inefficient as the costs of market participation are sunk prior to matching.[5]

We incorporate these features of the labor market—matching frictions and wage contracting—into a market entry game that can be implemented in the lab. Market entry games are an established class of binary choice games that have received considerable attention in the theoretical, econometric, and experimental literature.[6] Our game extends this literature to the labor market context. Specifically, the model that we take to the lab has the following structure: In the first stage, firms make a decision about whether to invest in vacancy creation and thereby participate in the labor market. In the second stage of the game, vacancies are randomly matched with workers according to a constant returns to scale matching function.[7] And, in the the third stage, matched firm-worker pairs divide a match surplus via a wage agreement. This game captures the notion that when considering whether to open a job vacancy, firms must anticipate the ease with which workers can be recruited and at what wage costs.

The essential prediction of the model is that firms will enter up to the point at which stochastic rationing of workers equalizes the value of a vacancy with its costs. Just as in the Diamond-Mortensen-Pissarides model, a zero-profit condition determines the degree of vacancy creation. In terms of comparative statics, the model predicts that when hiring becomes more valuable, more firms can profitably compete for workers. Other factors equal, job creation will increase when productivity increases or when wages are lower. Because of matching frictions, there will be unemployment in equilibrium even when the equilibrium is efficient.[8]

The tension in this game is the coordination problem associated with entry. Although

---

[4]Random matching should be contrasted with *directed* (or competitive) search in which firms commit to wages and and workers observe the wages prior to sending applications.

[5]Efficiency in random search models follows from the *Hosios condition*. Essentially, this condition relates bargaining power of firms to the sensitivity of the matching function to the presence of more firms. In contrast to models of random search, models of directed/competitive search tend to be efficient (Moen 1997).

[6]An early game-theoretic description of market entry is provided by Selten and Güth (1982). Rapoport and Seale (2008) summarize some of the main experimental tests of market entry games in a handbook chapter.

[7]Most of the experimental literature on labor market search focuses on directed search (See Helland et al. 2017, for a study that organizes the experimental results in this area).

[8]There are only a handful of studies that produce unemployment in the lab. one example is Fehr, Kirchsteiger, et al. (1996). In this study, they test the shirking model of unemployment in which unemployment is a by-product of efficiency wages. A motivation for this study is the fact that the search and matching framework has supplanted the shirking model as the basic way to understand unemployment.

the zero-profit condition pins down a level of vacancy creation, it does not identify which subset of firms should enter the labor market. Firms thus face strategic uncertainty that is not resolved by theory and beliefs are critical.[9] This highlights a crucial difference between the small market setting we investigate and the assumption of atomistic agents made in most macro models. In the large market setting, the decision of an individual firm will not matter for the matching probability.[10] Although matching is stochastic at the individual level, a fixed proportion of firms are matched. In contrast, in the small market setting there is a coordination problem. Because the small market setting is an empirically relevant case, we are interested to know if the zero-profit assumption still holds.

Our treatment variables are the size of the match surplus and the wage contracting institution. The design of the experiment aims to cleanly identify the effect of each. At the beginning of each session, subjects were assigned either the firm role or the worker role. Subjects remained in this role throughout the duration of the experiment. Every session consisted of 30 repetitions of the game. Between each play of the game, subjects were re-matched within a block into new markets comprised of 6 firm players and 4 worker players. The goal of re-matching was to represent in a lab setting an anonymous macro labor market. This contrasts with most experimental tests of market entry games in which tacit coordination and equilibrium selection is the primary interest.

Our first four treatments comprise a 2×2 design in which we vary the level of productivity and the presence or absence of wage bargaining after matching. In the absence of bargaining, the entry decision is a choice between a fixed payment and a binary lottery with a prize equal to half of a match surplus. In the presence of bargaining, firms are randomly matched with workers and participate in an ultimatum bargaining stage in which the proposer is determined by a fair coin. These four treatments enable us to identify the effect of productivity and bargaining on entry. In the fifth treatment, we examine ultimatum bargaining after matching but let the worker propose. This tests the sub-game structure of the model. The prediction for this treatment is that vacancy creation collapses because firms get expropriated whenever they negotiate with a worker.

In the sixth treatment, we make a substantial change to the bargaining institution. At the beginning of each round, a wage proposal is elicited from each worker and a group wage is computed as the average of these independent proposals. The group wage is then advertised

---

[9]Strategic uncertainty may be characterized as "uncertainty concerning the actions and beliefs (and beliefs about the beliefs) of others" (Morris and Shin 2002).

[10]Notice that if all firms enter with probability p then the standard deviation of the matching probability goes to 0 at a rate of $1/\sqrt{N}$ as the number of firms N increases to infinity.

to the firms prior to the entry decision. This treatment gives workers the opportunity to induce entry by moderating their wage demand. To make this treatment as clean as possible, the group wage binds for all wage negotiations.

We make contributions to three distinct literatures. The first is the literature on labor market search. Our study is the first to take elements of the Diamond-Mortensen-Pissarides model to the laboratory. This tests the behavioral assumptions of the model. This is difficult to achieve by other means. The second contribution is to the literature on market entry games. We test whether frictional matching and bargaining alter the basic findings in the literature. Our results tend to reinforce the existing stylized facts. The third contribution is to the experimental literature on bargaining. We are the first to test whether the presence of entry prior to bargaining affects bargaining outcomes.

Our foremost finding is that the aggregate outcomes respond to changes in the environment in the fashion anticipated by theory. When productivity increases, test subjects create more vacancies. This generates an inverse relationship between vacancies and unemployment—an experimental equivalent of the Beveridge curve. We also find that the allocation of bargaining power affects vacancy creation via its effect on the expected value of a hire. Notably, when workers have more bargaining power, it generates high unemployment because firms cannot recoup the resources they invest in vacancy creation.

Nevertheless, we do not reproduce the exact predictions of the model. There is too little entry when productivity is high and too much entry when productivity is low. This appears to be a systematic bias. As a consequence, the zero-profit condition does not hold—even on average—and persistent arbitrage opportunities exist. To give this a macroeconomic interpretation, the experimentally observed elasticity of unemployment with respect to productivity is smaller than anticipated. This is notable because it is contrary to the pattern observed in macroeconomic data but in line with results from other market entry games.[11]

Given the additional complications and dynamics that we introduce in our labor market entry game, the consistency of our results with the general findings in the market entry literature is remarkable. Although one might have expected the risk introduced by frictional matching and bargaining to reduce entry, this effect is small if it even exists. A possible explanation is that the reduction in the entry frequency of an individual player creates an opportunity for another player to profitably enter. Our findings support the conclusion from the literature that market entry games are a robust environment that create strong incentives for entry.

---

[11]This behavioral finding makes the Shimer critique perhaps even more puzzling.

Our bargaining results also deliver new findings. In the treatments with ultimatum bargaining, we find that offers tend to be lower and much more tightly distributed than in most of the literature. Bargaining offers also do not vary (in absolute terms) across treatments despite changes in the level of productivity and bargaining strength. Contrary to our expectation, workers do not appear to reward firms for entry. The modal offer is just slightly in excess of the direct costs of vacancy creation. We also show in a dramatic fashion that the offers are consistent with individual payoff maximization. Because workers have no way of internalizing the negative effect of high wage claims, vacancy creation collapses when workers have proposal power.

When we provide an institution by which workers commit to a group wage claim, the results are starkly different. In this environment, workers moderate their wage demands to induce entry by firms. Efficiency is restored because workers internalize the effect that their wage claims have on entry. A fascinating finding is that the distribution of individual wage proposals is roughly tri-modal with peaks at zero, half, and the entire surplus. Individuals appear to make offers strategically to move the group wage claim in the direction they prefer. This treatment demonstrates that labor organizations and centralized bargaining can have an efficiency enhancing effect by creating the preconditions for job creation. This treatment also suggests that the low offers associated with bargaining after match are a by-product of individual incentives rather than the outcome of a heuristic sharing rule.

In an attempt to reconcile our data to a model of behavior, we estimate a quantal response equilibrium (QRE) for our entry game. This approach is motivated by the observation that in a symmetric QRE the entry probabilities are closer to 0.5 than the Nash prediction. This can help explain the bias in the entry frequencies. However, a symmetric QRE cannot account for the heterogeneity that we observe at the individual level. In particular, there are a substantial number of individuals who enter in all periods. To account for both the aggregate bias and the individual pattern of entry, we therefore estimate a heterogeneous QRE in which we allow idiosyncratic subject-level preferences for entry. This helps account for some—though not all—of the variation in the data. This exercise suggests that aggregate biases can survive in environments with noise. This may be important beyond the lab.

The paper is organized as follows. In the next section, we provide context for this study and situate this study in the literature. After this background, we present the model that we test in the lab. This section includes the equilibrium analysis of each treatment and the associated predictions. In the third section, we go through the design and review the procedures. The fourth section presents the results and analysis of the results. The last

section concludes.

1.2   Background and Related Literature

We make contributions to three distinct literatures. The first is the literature on labor market search. Our study is the first to take elements of the Diamond-Mortensen-Pissarides model to the laboratory. The second is the literature on market entry games. Despite the addition of frictional matching and bargaining, our results tend to reinforce the existing stylized facts from this literature. The third literature to which we contribute is the experimental literature on bargaining. We are the first to test whether the presence of entry prior to bargaining affects bargaining outcomes. We also test a novel multilateral bargaining game.

1.2.1   The Standard Search and Matching Model of the Labor Market

Our study is motivated by the Diamond-Mortensen-Pissarides (DMP) model (Pissarides 2000). The DMP model is the workhorse model of the aggregate labor market because it is theoretically appealing and useful in empirical applications. Crucially, the model accounts for how fluctuations in productivity affect vacancy creation and thereby determine the level of unemployment. The DMP thus predicts movements in labor market variables over the business cycle.

Despite its successes, the DMP model exhibits some limitations. In his famous critique, Shimer (2005) shows that the standard calibration of the DMP model under-predicts volatility in the vacancy-unemployment ratio by more than an order of magnitude. The DMP model also struggles to account for certain empirical patterns. For instance, observed shifts in the Beveridge curve seem to imply adverse developments in matching efficiency (Elsby et al. 2015).

These shortcomings have stimulated research in a number of directions. One response to the quantitative limitations of the DMP models has been to propose alternative calibrations (Hagedorn and Manovskii 2008). Another approach has been to reexamine the theory, including reassessment of the free entry condition, the nature of wage determination, and the microfoundations of the search process. See, for example, Moen and Rosen (2006) who show how private information can increase the response of unemployment to changes in productivity.

We make a modest contribution to this literature. Although the DMP model addresses market level outcomes, an understanding of individual decision-making can help explain patterns in the aggregate data. The lab enables us to perform a clear test of how individuals behave when faced with the incentives from the model.[12] Specifically, we expose test subjects to reduced form matching frictions and directly test the no-profit condition. Our study thus lends credibility to the behavioral premises of the model. Comparable evidence is difficult to collect by other means. One issue is the availability of relevant data.[13] Another issue is the difficulty associated with identification.[14]

1.2.2   Market Entry

The model that we take to the lab is a version of a market entry game. Market entry games are a class of N-player binary choice games in which symmetric players simultaneously decide to either "enter" or to "stay out" (Selten and Güth 1982; Gary-Bobo 1990).[15] In this class of games, the payoff $\pi(v)$ is non-increasing in the number of entrants v and the payoff X from staying out is fixed. This environment is animated by the assumption that there exists some *market capacity* C such that $\pi(C) - X \geq 0$ but $\pi(C+1) - X < 0$.[16] Because any configuration of C total entrants is an equilibrium, these games are characterized by a large number of pure strategy Nash. There is also a symmetric mixed strategy equilibrium.[17] This creates a coordination problem. In the absence of a coordinating institution, agents face strategic uncertainty: The decision to enter is predicated on beliefs about the entry behavior of other players.

Market entry games are well known from industrial organization. For example, $\pi(v)$ could represent profits associated with Cournot competition between v firms when X is a sunk cost associated with market participation. If few firms choose to operate, profits are above the competitive level. However, if many firms produce, the market is oversupplied and *ex post* firms would have preferred to stay out.[18]

---

[12]Given the model is only an approximation to reality, what we are ultimately interested in is how well the assumptions represent real outcomes. An assumption is that findings inherit credibility from consistency with how individuals actually behave.

[13]Even data on aggregates such as vacancies and unemployment pose challenges (Elsby et al. 2015).

[14]For a discussion of identification challenges and other issues in the econometric studies of market entry and market structure see Toivanen and Waterson (2000) and Berry and Reiss (2007).

[15]This literature is part of the broader literature on coordination games (Ochs 1990; Cooper et al. 1990; Van Huyck et al. 1991; Cooper 1999).

[16]Minimally, we require that $\pi(F) - X < 0 < \pi(1) - X$, where F is the total number of participants.

[17]Other, "asymmetric equilibria," are also possible.

[18]There exist some econometric studies that attempt to structurally estimate discrete choice market entry

Market entry games have attracted substantial attention from experimental economists and been examined in many variations.[19] One variation has been with respect to the payoff structure. Another has been whether the market capacity is constant or fluctuating. A third variation has been with respect to the matching protocols.

The basic finding in this literature is that test subjects manage to coordinate in such a way that profits from entry are nearly equalized with the outside option (Ochs 1990; Rapoport, Seale, Erev, et al. 1998; Sundali et al. 1995; Morgan, Orzen, Sefton, and Sisak 2012). This is despite the absence of any organizing institution or possibility of communication.[20] The high degree of coordination has even been described as "magic" (Kahneman 1988; Erev and Rapoport 1998).

Although the stylized fact of a high degree of coordination is well-established, at least one systematic bias has been identified. When the market capacity is low, there tends to be excessive entry while the opposite tends to holds when market capacity is high (Rapoport, Seale, and Ordóñez 2002; Goeree and Holt 2000; Morgan, Orzen, and Sefton 2012). We find evidence of the same bias in our study.[21] In addition, individual behavior is heterogeneous and inconsistent with mixed strategy play at the individual level (Duffy and Hopkins 2005; Erev and Rapoport 1998). Zwick and Rapoport (2002) identify four "clusters" of subjects that employ distinct strategies. Of the four groups, the largest is a group of players exhibiting "sequential dependencies" (i.e. play that depends on the experienced history of "successes" or "failures") that is inconsistent with any model of randomization. Our data also mirror this finding.

The study most closely aligned with the present work is Rapoport, Seale, and Ordóñez (2002). Rapoport, Seale, and Ordóñez (2002) also investigate market entry under uncertainty. In their study, players who enter the market participate in a lottery for which the probability of winning depends on the number of entrants. This is comparable to the matching stage in our game. Although our studies differ in most other respects, our study corroborates their main finding that coordination is good on the aggregate level but not necessarily at the individual level. In terms of theory, Anderson and Engers (2007) develop results that are useful for understanding strategic uncertainty in market entry games. Although they study a specific and extreme game—the "blonde game" in which payoffs are

---

games using field data (Bresnahan and Reiss 1990; Ciliberto and Tamer 2009).

[19]As background, section 1.A.1 presents a classic experimental implementation.

[20]See Andersson and Holm (2010) for a study that incorporates communication.

[21]In appendix section , we present the findings from Sundali et al. (1995). This study illustrates in a clear way the main stylized facts frm the literature.

zero for entrants if there is more than one—their results generalize in a natural way to our setting.

A second strand in the literature that is relevant to the present study is the set of studies in which participants participate in a second stage after market entry. Examples include Camerer and Lovallo (1999) and Morgan, Orzen, and Sefton (2012). In the study by Camerer and Lovallo (1999), there is a skill-based tournament after entry, while in Morgan, Orzen, and Sefton (2012) test subjects make an investment decision. As in our study, this introduces a subgame dimension that is important for the entry decision. Players must anticipate outcomes in the second stage when considering an entry decision in the first stage. Notably, the findings in these studies are similar to our own. This includes the finding of an aggregate bias in the entry frequencies.

1.2.3   Bargaining in the Lab

Our main treatments embed ultimatum bargaining in an entry context. Relative to the existing literature, the entry margin is new. The presence of an entry margin is important because it could plausibly affect bargaining via a number of channels. For instance, in the presence of reciprocity, one might expect workers to reward firms for vacancy creation. Workers depend on jobs for income and can only get hired if firms invest in vacancies. The labor market entry game thus bears some resemblance to the binary trust games.[22] Reciprocity of the kind normally observed in the trust game would predict that workers reward firms and allow them to appropriate a large part of the surplus in bargaining. Notably, we do not find evidence of reciprocity effects. We suspect that this may be a consequence of the multilateral nature of matching in our setting.

Another aspect of our model that is not present in most other studies is that participants must take an active decision to participate in bargaining. The selection effect might explain why offers in our study tend to be lower than in most of the literature.

The study also contributes a new bargaining structure. The treatment in which workers negotiate a group wage claim prior to firm entry does not have a close analog in the literature. This intra-worker bargaining dramatically tests the ability of test participants to trade-

---

[22]The basic trust game was proposed by Berg et al. (1995). This game is sometimes as referred to as the investment game. In this game, a "sender" can choose to invest some portion of an endowment with a "receiver." This investment is then scaled up by a factor larger than 1. In the second stage, the receiver can return some portion of the scaled-up investment to the sender. For a comprehensive overview of findings in such games see that meta-analysis by Johnson and Mislin (2011). The main finding in such games is a relatively high level of investment by the sender and a correspondingly high level of return by the receiver.

off the benefit of a low wage with the benefit from a higher matching probability. An important aspect of this bargaining institution is that only a few sophisticated types are required to generate an optimal outcome. If naive workers propose super-optimally high wages, sophisticated workers can propose sub-optimal wages such that the average proposal maximizes expected earnings.

1.3 Model and Treatments

1.3.1 Model

We refer to the game that we investigate as a *labor market entry game*. The version we consider is populated by F identical firm players and U identical worker players. The labor market entry game has the following basic structure:

1. **Vacancy Creation**: In the first stage, the N firms independently decide whether or not to create a vacancy. Firms that create a vacancy participate in the matching stage for workers. Firms that do not participate receive a fixed payment X and take no further actions. The fixed payment X may be interpreted as the cost of opening a vacancy.[23] We refer to the decision to open a vacancy as a decision to *enter* the labor market and we refer to firms in the labor market as *entrants*. We denote the total number of vacancies by V.

   The U unemployed workers automatically participate in the labor market and are passive in the first stage. These U workers represent an equilibrium level of unemployment.

2. **Matching**: In the second stage, a subset of vacancies and unemployed workers are matched into pairs via a constant returns to scale matching function: $M(U, V)$. The matching function is a reduced form way of representing labor market frictions and the constant returns to scale specification is the empirically relevant specification in the labor market context (Petrongolo and Pissarides 2001). In our experimental implementation, we use a Cobb-Douglas form

$$M(U, V) = AU^\alpha V^{1-\alpha} \tag{1.1}$$

---

[23]The cost of vacancy creation may include direct costs as well as the opportunity cost of unused capital. For a discussion of a magnitude of these costs see Hagedorn and Manovskii (2008). In the absence of opening a vacancy, the cost X could be invested in some other opportunity.

where $\alpha \in (0,1)$.[24] The constant returns to scale assumption implies that that the matching probability for firms is declining and convex in the *market tightness*, $\theta = V/U$.

$$\frac{M(U,v)}{v} = A\frac{U^\alpha V^{1-\alpha}}{V} = A\theta^{-\alpha} = q(\theta).\text{[25]}$$

The function $q(\theta)$ captures the notion that when additional firms compete for workers, it becomes more difficult for firms to hire.[26] Because workers do not make an entry decision in our experiment, the degree of matching friction is determined by the number of vacancies created. To emphasize this, we write the matching probability as a function of V rather than θ: $q(V)$.[27]

3. **Wage Bargaining**: After matching, matched firms and workers negotiate a wage w to divide a *production (match) surplus* Y. We refer to Y as the *productivity* of a match. If a firm and worker reach a wage agreement, a job is created and production takes place. The firm earns $Y - w$, and the worker earns w. If a matched firm and worker fail to reach an agreement, both parties earn zero.

Throughout, we use η to summarize the (expected) share of the production surplus that the firm can appropriate. We interpret η as the bargaining strength of the firm. The expected wage payment is therefore $w = (1 - \eta)Y$.

The *ex ante* expected payoff for firm n when $V - 1$ other firms create vacancies is

$$\mathbb{E}[\pi_n] = \begin{cases} X & \text{if } n \text{ doesn't enter} \\ q(V)\,\eta Y & \text{if } n \text{ and } V - 1 \text{ other firms enter.} \end{cases} \tag{1.2}$$

---

[24]Section 1.B.1 in the appendix shows the shape of the matching function and the associated matching probabilities for firms.

[25]The equivalent probability for workers is computed in analogous fashion and denoted by $\mu(\theta)$. This function increases in V such that the matching probability for workers increases when more vacancies are created.

[26]Rapoport, Seale, and Ordóñez (2002) study a setting with *endogenously determined lotteries* in which the probability of winning a lottery is determined by the number of entrants. This is analogous to the matching probability in our study. However, in contrast to our study, Rapoport, Seale, and Ordóñez (2002) use a linear specification for the probability of winning, $p = \frac{N-V}{N}$. Although this formulation has the advantage of tractability, it does not correspond to any reasonable matching function. For instance, a linear specification of the matching function implies that the expected number of matches declines when the number of entrants gets sufficiently large. To see this, note that the expected number of matches is pV. This implies a number of matches $V - \frac{V^2}{N}$ for the linear matching case.

[27]Future work may extend the experiment to include an entry margin for workers.

Notice that firm n faces stochastic risk because of the matching function even in the absence of uncertainty about the number of (other) entrants: With a probability $1 - q(V)$ the firm remains unmatched and earns zero.

Equilibrium   We limit the equilibrium analysis of the model to the case of expected payoff maximization. Earlier work in a similar setting has found that the Nash equilibrium for risk-neutral players effectively organizes the aggregate results and this is corroborated in the present study (see Rapoport, Seale, and Ordóñez 2002). The impact of risk aversion is, however, simple to characterize in qualitative terms. Because hiring and wage costs are uncertain, entrants would need to be compensated for this risk. In the presence of risk aversion, we would therefore expect relatively fewer vacancies to be created but for those vacancies to have an expected value in excess of X.

The labor market entry game has both pure and mixed strategy equilibria.[28] Consistent with the notion of a labor market with random matching, we focus on the symmetric mixed strategy equilibrium. As we discuss in the following section, we make a number of design decisions to enforce play of this equilibrium. In particular, we randomly re-match anonymous players into new markets before the start of each repetition of the game. We also limit the feedback to information about market level outcomes. This circumscribes the ability of players to coordinate on a particular equilibrium and makes the symmetric mixed strategy equilibrium seem most plausible as a model of behavior.

As in the standard Diamond-Mortensen-Pissarides model, the level of vacancy creation in the symmetric mixed strategy equilibrium is determined by a *zero-profit* condition. This zero-profit condition is characterized by a probability p^* such that

$$\sum_{V=0}^{N-1} \binom{N-1}{V} p^{*V} (1-p^*)^{N-V-1} q(V) \eta Y = X. \tag{1.4}$$

This condition says that when $N-1$ identical firms randomize with a probability p^*, the N^{th}

---

[28]The pure strategy equilibria are characterized by the *market capacity* C. C is the maximum number of firms that can profitably create vacancies. For a given level of Y and w, C is identified by

$$q(C)(Y-w) \geq X \quad \text{and } q(C+1)(Y-w) < X. \tag{1.3}$$

These conditions imply that no firm has a profitable deviation. Each of the C firms that create a vacancy expect to earn at least X and, simultaneously, none of the $F - C$ firms that take the fixed payment could expect to earn more than X by creating a vacancy. Because each possible way of picking C firms out of F is a valid equilibrium, there are $\binom{F}{C}$ pure strategy equilibria in total. Much of the experimental literature on market entry games has been interested in coordination issue associated with multiple equilibria.

firm is indifferent between creating a vacancy and taking the guaranteed payment. Since no individual firm has an incentive to deviate from p^*, this is an equilibrium. [29]

Although the assumption of symmetric randomization implies that all firms randomize with the same probability, what is crucial from the perspective of an individual firm is not whether each firm randomizes with the same probability but whether the average probability of entry is above or below p^*. In particular, if the average probability of entry is below p^*, then the firm should enter more often as there are excess profits associated with a vacancy. In practice, we may therefore expect that competition for profits will induce a mixing probability close to p^* even though there is latitude for individual firms to randomize with different probabilities. As long as the average probability of entry in the population of firms is p^*, an individual firm will be indifferent between the fixed payment or entry. If p^* arises because a portion p^* of the population of firms enters with certainty (and a portion $1 - p^*$ never enters), this has an identical implication for individual behavior as if all firms were randomizing with probability p^*.

While it is not possible to solve for p^* analytically due to the non-linearity introduced by the matching function, it is straightforward to solve for p^* numerically (from condition 1.4). In addition, condition 1.4 has obvious comparative statics. Holding other factors fixed, if Y, the match productivity, or η, the amount of the surplus the firm can appropriate in bargaining, increase, then p^* must also increase in order to maintain the equality. Increases in the value of hiring are thus compensated by increases in vacancy creation.[30]

1.3.2   Parameters and Treatments

In all treatments, a market is comprised of $N = 6$ firms and $U = 4$ workers. Throughout, we use a simple parametrization of the matching function: $A = \frac{1}{2}$ and $\alpha = \frac{1}{2}$. Table 1.1 summarizes the matching outcomes associated with each level of entry for the specification of the matching function used in our study. The guaranteed payment X was fixed equal to 95. The level of X was chosen in combination with the two different levels of Y to ensure clear treatment differences. The reason that X was perturbed below 100 was to disrupt a focus on round numbers. 100 provides a natural focal point for bargaining offers and might also give the impression that the ratio X/Y has some special importance.

---

[29]Existence and uniqueness (in the class of symmetric strategies) of the mixed strategy equilibrium, as well as efficiency results follow, with minor modifications, from Anderson and Engers 2007.

[30]Although not the focus of our experiment, improvements in the matching efficiency would have the same implication for vacancies.

Table 1.1: Matches and Matching Probabilities

v	U	$M(U,V)$	$q(\theta)$
0	4	0	.
1	4	1	1.00
2	4	1.41	0.71
3	4	1.73	0.58
4	4	2	0.50
5	4	2.24	0.44
6	4	2.45	0.41

For this specification of matching technology, an equal split of the production surplus is approximately efficient in the sense of maximizing total surplus in the market, $M(U,V)Y + (N - V)X$. This analogous to the *Hosios condition* in the random matching setting. Efficiency arises when firms are compensated to the degree that they contribute to job creation. Since firms and workers contribute equally to the matching function, this implies efficiency of an equal split. Accordingly, this result would change if the matching function had a different specification. The basic efficiency results are located in the appendix, subsection 1.B.2.

We conduct six treatments in which we vary the productivity Y and the bargaining strength η. We label the treatments by T^Y where $T \in \{M, B, W, G\}$ denotes the form of wage determination (what we also refer to as the bargaining institution) and $Y \in \{h, l\}$ denotes the surplus—h for the high surplus level $Y = 400$ high and l for the low surplus level $Y = 300$. The first two treatments M^h and M^l focus on matching frictions. The second two treatments N^h and N^l add a bargaining stage implemented via an ultimatum bargaining protocol after matching. To approximate equal bargaining strength between the workers and firms (that is, $\eta = 0.50$) we have the proposer in the bargaining stage be determined by the flip of a fair coin. In the fifth treatment W^h, we use the same set-up as in N^h and N^l but have the worker always propose at the ultimatum stage. In the last treatment G^h, we make a more substantial change to the bargaining institution. In this treatment, workers as a group commit to a wage level that is advertised to firms before they make the vacancy creation decision.

Treatments M^h and M^l : Matching Frictions (M)

M^h and M^l examine the impact of frictional matching on market entry in the absence of bargaining. In both treatments, the wage w is fixed equal to half of the production surplus. This exogenously imposes the bargaining weight $\eta = 0.5$. These two treatments provide a baseline to which the other treatments can be compared.
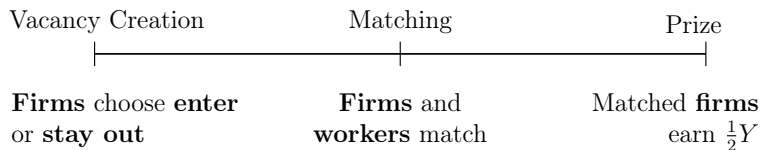
Figure 1.1: Treatments M^h and M^l

Figure 1.1 shows the structure of treatments M^h and M^l . In the first stage, firms choose to either take the outside payment X or to participate in a binary matching lottery. Firms that participate earn ηY with probability $q(\theta)$ and 0 with a complementary probability. This corresponds to a situation in which the firm pays the wage $w = (1 - \eta)Y$ conditional on hiring. In M^h the value of hiring is $\eta Y = \frac{1}{2}400 = 200$ and in M^l the value of hiring is $\eta Y = \frac{1}{2}300 = 150$. Although the value of hiring is fixed, the probability of hiring declines as additional firms enter because of matching frictions. For the two levels of productivity, the symmetric Nash equilibrium identified by equation 1.4 predicts an entry frequency of $p_{T1}^* = 0.73$ in treatment M^h and $p_{T2}^* = 0.37$ in treatment M^l . This corresponds to 4.38 vacancies and 2.22 vacancies respectively.

Treatments M^h and M^l are the treatments most directly comparable to the canonical market entry game. The difference from the basic implementation is that the payoff from entry is a stochastic rather than deterministic function of the number of other entrants. The most closely related study is Rapoport, Seale, and Winter (2002). This study examines lotteries in which the probability of winning is determined by an entry protocol, and is a linear function of the number of entrants. Treatments M^h and M^l in effect re-examine this setting but with a lottery that is a non-linear rather than linear function of the number of entrants.

Treatments N^h and N^l : Fair Bargaining after Match (B)

In treatments N^h and N^l , we extend the basic treatments by including a bargaining stage in which the firm and worker have *ex ante* equal bargaining power. These treatments represent

a situation in which firms open vacancies knowing that once a suitable worker is identified, a wage must be negotiated. The timing of this version of the game is shown in figure 1.2.

As in M^h and M^l, firms choose to either take the outside payment X or to participate in a matching market in which the probability of meeting a suitable worker has probability $q(V)$. However, in contrast to the first two treatments, the wage is determined by bargaining over the production surplus Y. As in the first two treatments, $Y = 400$ in the h treatment and $Y = 300$ in the l treatment.

Vacancy Creation	Matching	Bargaining State	Proposal	Contracting
Firms choose **enter** or **stay out**	**Firms** and **workers** match	**Proposer** assigned by a fair coin	**Proposer** offers $x_r \in [0, Y]$	**Responder** accepts or rejects x_r
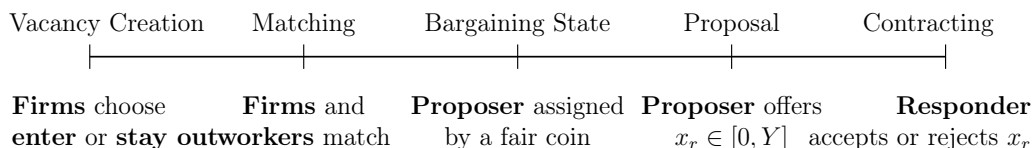
Figure 1.2: Treatments N^h and N^l

We implement the bargaining stage as ultimatum bargaining in which the proposer is determined by a fair coin flip. In half of instances, the firms get to propose the wage and, in the other half of instances, the workers get to propose the wage. Because X is a sunk cost, firms and workers are symmetric at the bargaining stage, both with an outside option of 0. This represents a situation in which the firm and the worker have roughly equal bargaining power. This protocol is an attempt to endogenously represent Nash Bargaining from the Diamond-Mortensen-Pissarides model with bargaining weight $\eta = 0.50$.

Regardless of the proposer's type (firm or worker), let $x_r \in [0, Y]$ denote the offer extended to the responder. If the responder accepts the offer, the proposer earns $Y - x_r$ and the responder gets x_r.[31] If the responder rejects the offer, both the proposer and the responder earn zero.

The standard prediction from economic theory is that the proposer in an ultimatum games will appropriate the entire surplus. That is, the proposer will offer $x_r = 0$ (or $x_r = \epsilon$) and keep Y. The expected payoff in the bargaining stage is therefore equal to earning the entire surplus half of the time, in expectation $\frac{1}{2}Y$. This means $\eta = \frac{1}{2}$, just as in treatments M^h and M^l. Moreover, that η equals 0.5 does not depend on an equal split. $\eta = 0$ will hold in practice as long as the sharing norm is consistent: If the proposer always gets a share $\gamma \in [0, 1]$ of the surplus and the responder always gets the complementary share $1 - \gamma$, then the expected earnings will be $\frac{1}{2}\gamma Y + \frac{1}{2}(1 - \gamma)Y = \frac{1}{2}Y$. In terms of expected

---

[31]If the firm proposes, the wage is $w = x_r$ while if the worker proposes, the wage is $w = Y - x_r$.

earnings, treatments N^h and N^l are therefore strategically equivalent in expected earnings to treatments M^h and M^l, with the same predictions for entry probabilities and profits.

Note that in these treatments, the bargaining stage is "decoupled" from the entry decision because the firm gets to propose half of the time. Post-match bargaining is thus introduced into an entry game in a fashion that should have minimal consequences for entry relative to M^h and M^l.

Treatment W^h : Worker Wage Demand after Match (W)

Treatment W^h is identical to treatment N^h but with a single change to the bargaining stage. Instead of a random draw determining the proposer, the worker always get to make the wage claim. The structure of the treatment is shown in figure 1.3.

Vacancy Creation	Matching	Proposal	Contracting
Firms choose **enter** or **stay out**	**Firms** and **workers** match	Matched **workers** propose $w \in [0, 400]$	Matched **Firms** accept or reject offer $400 - w$
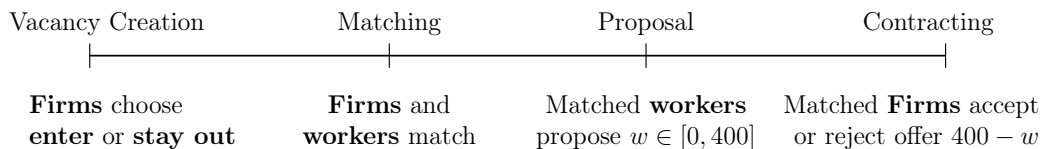
Figure 1.3: Treatment W^h

In the absence of a commitment device, the standard model anticipates that the proposer will appropriate all of the surplus at the bargaining stage. And since workers always propose in treatment W^h, firms will anticipate $\eta = 0$. This yields a stark prediction: Knowing that they will be expropriated, firms should choose to take the fixed payment X rather than open vacancies. When bargaining takes place after matching and workers have all the bargaining power, theory thus predicts complete unemployment. In effect, firms face a hold up problem in which they cannot recoup the investment in vacancy creation.

Even if test subjects in the worker role do not in practice demand the entire surplus, we expect that the level of unemployment will be higher and the level of efficiency will be lower than in N^h. The problem for workers is that they do not have a mechanism by which they can limit individual wage demands yet it is expectations about bargaining outcomes that will determine the availability of vacancies.

Treatments G^h : Group Wage Demand prior to Entry (G)

In contrast to W^h, in G^h we investigate an environment in which workers commit to a group wage prior to firm entry. This treatment incentivizes workers to set a wage that maximizes

the expected *ex ante* payoffs and thereby creates the maximum surplus for workers. We can think of this as centralized wage formation with similarities to a corporatist arrangement in which a labor organization negotiates a binding wage based on the negotiation amongst the members.[32]

In the first stage of G^h , each worker i makes a wage proposal w_i. These wages are then averaged into a group wage claim $\bar{w} = \frac{1}{U} \sum_{i=1}^{U} w_i$. Bargaining thus occurs among the workers. The group wage claim $\bar{w}$ is then advertised to firms prior to the entry decision. From the perspective of firms, G^h is comparable to M^h in the sense that the decision to open a vacancy amounts to the choice between a fixed payment and a lottery with prize $Y - \bar{w}$. The structure of G^h is shown in figure 1.4.

Proposal	$\bar{w}$ advertisement	Vacancy Creation	Matching	Contracting
Workers	**Firms** observe	**Firms** choose	**Firms** and	Matched **firms**
propose	$\bar{w} = \frac{1}{U} \sum_{i=1}^{U} w_i$	**enter** or	**workers** match	accept or reject
$w_i \in [0, 400]$		**stay out**		offer $400 - \bar{w}$
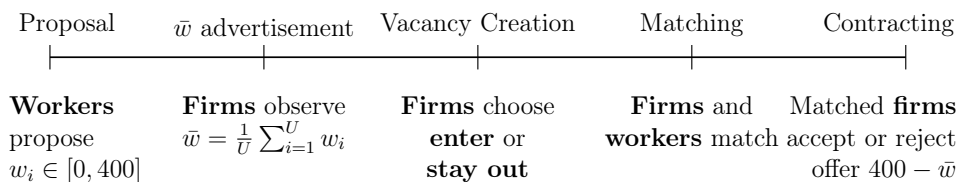
Figure 1.4: Treatment G^h

Because workers commit to a wage prior to vacancy creation, workers can influence the matching probability by moderating their wage claim. Specifically, we expect workers to set a wage $\bar{w}$ to induce the entry probability $\tilde{p}$ that maximizes the expected payoff

$$\sum_{v=0}^{F} \binom{F}{v} (\tilde{p})^v (1 - \tilde{p})^{F-v} \mu(\theta) \bar{w} \tag{1.5}$$

where $\tilde{p}$ is determined from the no-profit condition (equation 1.4) and $\mu(\theta) = M(U, v)/U$ is the matching probability for workers. Optimally, workers set wages to induce vacancy creation up to the point at which the increase in the probability of getting hired is equalized with the cost of lower wages. Solving this problem must be done numerically. For $Y = 400$, the payoff maximizing group claim is $\bar{w} = 207$ with associated entry probability of $\tilde{p} = 0.78$. This means that $\eta = 0.52$. Because $\eta = 0.52$ is just slightly in excess of of η in M^h and N^h , theory predicts similar results in these three treatments.

Although workers in G^h have all the bargaining power—just as in treatment W^h — treatment G^h predicts that wages will be lower, vacancy creation will be higher, and un-

---

[32]See Layard et al. (2005) for a more comprehensive discussion of the merits of centralized versus firm-level wage bargaining.

employment will be lower. This treatment anticipates that institutions such as unions can be efficiency enhancing because they enable workers to internalize the negative externalities associated with bargaining after match (See Layard et al. 2005, chapter 2).

1.4   Design

Our treatment variables are the size of the match surplus and wage contracting institution. The design of the experiment aims to cleanly assess the effect of each of these factors. We employ a conservative approach and the identification of the main treatment effects rely on non-parametric comparisons of independent block level outcomes. Credibility of our findings is enhanced because the tests are sufficiently powerful (List et al. 2011).[33]

In the framework of the previous section, the key parameters can be thought of as Y and η, where η is exogenous in M^h and M^l but is endogenous in the treatments with bargaining, N^h –G^h .[34] Table 1.2 gathers the theoretical predictions for p^* and η worked out in the previous section. The main comparisons are between the level of productivity (reading across table 1.2) and how wages are determined (reading down table 1.2).

Table 1.2: Summary of Treatments

		$Y = 400$			$Y = 300$	
		p	η		p	η
Matching	M^h	0.73	0.50	M^l	0.38	0.50
Fair	N^h	0.73	0.50	N^l	0.38	0.50
Worker	W^h	0.00	0.00			
Group	G^h	0.78	0.52			

The predictions in table 1.2 are derived from the standard model in which agents maximize payoffs, do not make mistakes, hold correct beliefs etc. As we know from other existing evidence, many of these assumptions are questionable. The goal of the experiment is to investigate the extent to which the environment reproduces the equilibrium predictions even though reality does not coincide exactly with the model.

---

[33]We plan to do a final round of data collection in the spring of 2017 in which we add observations to strengthen our bargaining results.

[34]Y and η together determine the wage level.

Our first four treatments, M^h-N^l, comprise a 2×2 design in which we examine productivity and post-match bargaining. The 2×2 format facilitates a comparative statics assessment of the effect of productivity changes and if there are effects related to bargaining. There is also the possibility of interactions between these two features.

Differences in the level of productivity translate into predictions for the distribution of entrants. Given symmetric randomization according to the mixed strategy Nash, the distribution of the total number of entrants should be binomially distributed as shown in figure 1.5. This plot may be interpreted in the following fashion: If players are playing according to the symmetric mixed strategy equilibrium then in any given period there is a some probability of $v = 0, \ldots, 6$ entrants. The probability of observing v entrants is equal to the height of the associated bar. Our choice of productivity levels predicts distinct distributions in which the low productivity treatment has a right skew and the high productivity treatment has a left skew. Moreover, the modal number of entrants is different by two entrants.

Figure 1.5: Predicted Probability of Number of Entrants by Surplus
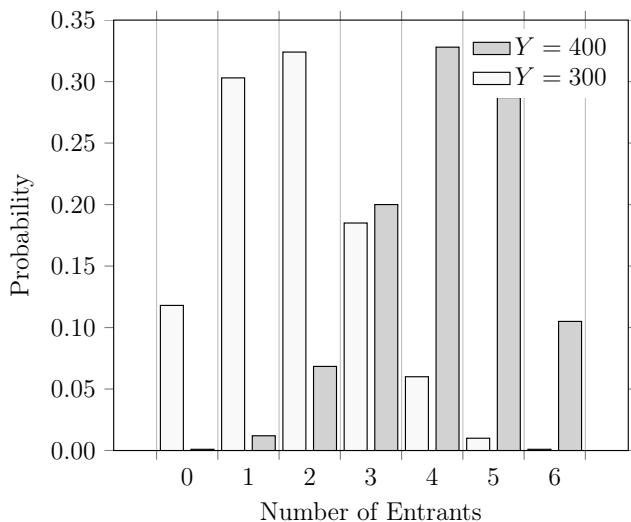
As discussed above, the ultimatum bargaining stage in N^h and N^l is specified in a neutral fashion such that all players can expect to earn half of the surplus given a symmetric sharing norm. We chose this bargaining protocol so that N^h and N^l are the mildest possible extension of M^h and M^l. This facilitates inference about the independent effect of bargaining.

There a number of reasons to expect bargaining to have an impact on entry. One reason

is that bargaining exposes test subjects to an additional source of risk. This would predict lower entry. For example, risk averse subjects would prefer a guaranteed payment of half the surplus (as in M^h and M^l) rather than earning the entire surplus half of the the time (which is the prediction in N^h and N^l). Another way in which bargaining can be important is if η deviates from 0.5 due to behavioral factors. Average bargaining outcomes could vary depending on the identity of the proposer (firm or worker) or due to differences in the level of unemployment. For instance, workers may reciprocate entry by firms in similar fashion as in the trust game and allow firms to appropriate a large portion η of the surplus as a reward. High unemployment might also increase η because matched workers have less leverage in bargaining. If workers endure a longer unemployment spell before again getting matched they may not be able to afford to reject low offers. This could justify differences between treatments N^h and N^l .

In treatments W^h and G^h , we introduce variations in how wages are determined. We compare these treatments with M^h and N^h which have the same level of surplus. The comparison between N^h and W^h accounts for how changes in bargaining strength impact the expected payoff that firms can expect in the last stage. Specifically, we expect lower entry in W^h . Unless workers in W^h propose an equal 1/2 split of the productive surplus, the expected value of entry will be lower for firms in W^h than in N^h .[35] Anticipating this, firms will enter less frequently.

In comparison to W^h , G^h creates incentives for workers to moderate wage demands. Theory predicts "hyper fair" offers of more than 50% of the surplus. Although workers have all the bargaining power, the outcomes should be less like W^h and more similar to N^h and M^h . The parallel with M^h is particularly strong because firms know prior to entry what they will earn if they get matched.

To try to discern the impact of bargaining outcomes, we compare the offers across types (firms compared with workers) and treatments. One hypothesis that can be checked is whether workers adjust their offers between N^h and W^h , the two treatments in which they have all the bargaining power. Whereas in N^h firms get to propose half of the time, in W^h firms must be rewarded by workers. If there are trust game like effects, we would expect these to be larger in treatment W^h because the firms only enter if they are rewarded by workers. Treatment G^h illuminates the analysis of bargaining because workers have the chance to impact the degree of entry via the wages that they set. This tests the ability of workers to

---

[35]This conclusion is predicated on the plausible assumption that the sharing norm in N^h is such that the proposer is able to appropriate more of the surplus than the responder.

think through the incentives that they create via wages. Furthermore, wage formation in G^h is such that individual workers can affect the wage by setting extreme offers. This means that even a few sophisticated individuals can counteract inappropriately low or inappropriately high offers by making extreme wage proposals in the other direction. Difference in the wage proposals in W^h and G^h thus give some measure of how individuals response to the incentives created by bargaining after match.

Because of the entry margin, individuals who choose to participate in the bargaining stage may not be representative of the general population. This is a facet of our study that can result in differences relative to other studies of ultimatum bargaining.

1.4.1 Experimental Implementation

At the beginning of each session, subjects were assigned either the firm role or the worker role. Subjects remained in this role throughout the duration of the experiment. Each session implemented a single set of treatment parameters and results are therefore based on between subject (session) comparisons. For a discussion of the strengths and weaknesses of between-subject designs see Charness and Kuhn (2011). Given the novelty and complexity of our environment, we use a between-subject design to avoid confounds and establish a clear set of results. Moreover, between-subject comparisons are congruous with the realities of a marketplace in which different firms recruit different workers at different times.[36]

With respect to practical aspects of implementing the model in the laboratory, the use of matching in the lab requires that the matching technology generate an integer number of matches. This introduces a problem relative to the CRS matching function which can yield fractional matches (see table 1.1). For instance, if five firms enter we expect 2.24 matches. This could be dealt with in various ways. We elect to enforce a "psuedo-law of large numbers": When there is a fractional match, an additional match is generated with a probability equal to the remainder. For example, if there are 2.24 matches, then there will be 2 matches in 76% of cases and 3 matches in 24% of cases. This means that given the number of entrants, there will be *at least* $\lfloor M(U,V) \rfloor$ matches but *no more than* $\lceil M(U,V) \rceil$ matches. This means that, on average, expected matching probabilities are correct and that the number of matches is almost deterministically determined by the number of entrants.

A major distinction between our implementation of the market entry game and most

---

[36]If we instead interpret the model as one in which an individual firm repeatedly returns to the marketplace, a within-subject design could be appropriate to understand how firms adapt to changes in the environment over time. This could, however, introduce order effects that we wish to avoid at this initial stage.

previous work is our choice of matching protocol. Prior to each round, subjects were re-matched within a block into new markets of 6 firm players and 4 worker players. The goal of re-matching was to represent in a lab setting an anonymous macro labor market. This contrasts with most experimental tests of market entry games in which tacit coordination and equilibrium selection are the primary interest (Kahneman 1988; Sundali et al. 1995; Erev and Rapoport 1998; Rapoport, Seale, Erev, et al. 1998). Our aim is to simulate a series of independent market entry decisions and we therefore want to disrupt tacit collusion of the form that most studies hope to induce. Moreover, although we cannot induce play of a specific equilibrium, by re-matching we introduce strategic uncertainty such that test subjects must rely on beliefs about the likely level entry. In terms of implementation, the present study thus bears some similarity to studies in which the market capacity is varied from period to period. [37] Because re-matching circumscribes opportunities for coordination, we expect subjects to play as if the remaining players were randomizing symmetrically.[38] Furthermore, the only way for all players to be using best responses is for entry frequency in the pool of potential entrants to adjust to the probability dictated by the unique mixed strategy equilibrium. If this were not the case, then an individual would have an incentive to either enter more often (arbitrage) or less often.

Given this set-up, the model and market parameters described in 1.4, and the levels of productivity denoted in table 1.2, our *ex ante* prediction is aggregate entry behavior consistent with the symmetric mixed strategy Nash equilibrium. For the treatments in which the surplus is shared equally (M and N), this implies a binomial distribution over the number of entrants with a binomial probability $p^* = 0.37$ for the low productivity treatments and $p^* = 0.73$ for the high productivity treatments. The expected distribution of entrants are shown in figure 1.5.

---

[37]Although studies such as Sundali et al. 1995 (shown in appendix section 1.A.1) use fixed groups, when the market capacity changes from period to period, equilibrium requires at least some players modify their strategies. This means that, by design, there is persistent strategic uncertainty. In contrast, in studies in which neither the market capacity nor the market composition are altered, the uncertainty is considerably reduced. In a notable contribution to the study of learning in games, Duffy and Hopkins 2005 show that in such an environment players should in the limit of many repetitions of the game converge to a pure strategy Nash equilibrium.

[38]If there are some players in the population who always enter while others never enter, and every period these individuals are selected randomly, then the experienced level of entry is consistent with individual randomization equal to the relative proportion of individuals who always enter. For example, if half of individuals always enter and the other half never enter, then a random draw of entrants implies that half will enter. This is equivalent to the case in which all players randomize with one half probability.

1.4.2 Procedures

All sessions were conducted in the research lab of the BI Norwegian Business school using participants recruited from the general student population at the BI Norwegian Business School and the University of Oslo, both located in Oslo, Norway. Recruitment and session management were handled via the *ORSEE* system (Greiner 2004). *z-Tree* was used to program and conduct the experiment (Fischbacher 2007). Anonymity of subjects was preserved throughout.

Table 1.3: Blocks per Treatment

Treatment	Blocks
M^h	5
M^l	4
N^h	4
N^l	4
W^h	2
G^h	2

We collected 17 block level observations, distributed among the treatments as shown in table 1.3. Each block consisted of 2-3 markets. In total 292 individuals participated in the experiment.

On arrival, subjects were randomly allocated to cubicles in the lab in order to break up social ties. After being seated, instructions were distributed and read aloud in order to achieve public knowledge of the rules. All instructions were phrased in neutral language. Although interested in capturing features of the labor market, we wished to avoid any direct connotations with this setting. For example, we did not wish to frame the bargaining portion of the experiment in such a way that it seemed important to pay a "fair" wage. We therefore referred to firms and workers players as A types and B types. For similar reasons, we described the labor market as simply a "matching market." This term conveys the essential information that A and B players will be randomly paired without implying anything about the relative status of the two types (such "boss" vs. "employee"). Another important detail of the language that we used was with respect to the description of the bargaining situation. Rather than ask the proposer to make an "offer" or to "divide the surplus," we instructed the proposer to state a *claim*: A claim was defined as an amount that the proposer gets if his or her claim is accepted. Our motivation for using this language was to avoid establishing a "property right" for the proposer. This is appropriate for the labor market context because

the surplus is only realized after the firm and worker together produce.

Each session of the experiment began with two test periods in which players could get acquainted with the software. This was immediately followed by 30 periods in which the players earned payoffs. At the beginning of every round, subjects within a block were randomly matched into markets with the composition described above. Each round consisted of a single repetition of the market entry game.

Gameplay was formulated in the following fashion: In the first stage of the game, subjects in the firm role chose to either participate in the market or to take a guaranteed payment. A player who chose to take the guaranteed payment took no further actions in the game but did receive market level feedback at the same junctures as other players. After the entry stage, players received information on how many players of each type chose to enter, the number of realized matches, the matching probabilities, and whether or not they were matched. In the treatments without bargaining, players that matched earned the payoff ηY. In the other treatments, matched players were assigned either the proposer or responder role at this stage. In the ultimatum bargaining portion of the game, a screen was devoted to both the proposal decision and the responder decision. At the conclusion of the game, all players were provided with information on their own profits and a historical statistic describing previous decisions and market outcomes, along with accumulated profits.

The one exception to the basic structure was treatment G^h which featured two additional stages at the beginning of the experiment. In the first stage, wage proposals were collected from the worker types. The average of these proposals was then presented to all participants in a given market on a dedicated screen. Otherwise, the structure of the game was the same. In a similar way as the other treatments with bargaining, the proposer (workers) and responder (firms) were identified separately. In the final stage before payoffs and feedback, the responder had the right to reject the (group) wage offer. This stage was in some sense unnecessary (since the firms know the wage before the entry decision) but preserved as close as possible the experience in the treatments with ultimatum bargaining.[39]

Subjects earned experimental currency units (ECU) based on the realized outcome in each period. After the final game, accumulated earnings in ECU were converted to Norwegian Kroner (NOK) using a fixed and publicly announced exchange rate. Subjects were paid in cash privately as they left the lab. On average subjects, earned between 220 and 250 NOK (about 36 USD at the time). A session typically lasted 60 minutes.

---

[39]No firm that entered in G^h chose to reject the wage offer.

Sample instructions and screen shots are available upon request.[40]

1.5 Results

We present the main findings as a series of *Results*. Unless otherwise noted, statistics and tabulations are based on data from the last half of periods. We focus on the last half of the data because there is a negative trend in the level of entry during the initial periods of some of the treatments. By comparison, behavior is more or less stable in the second half of the experiment. With the exception of W^h , none of our conclusions depend substantively on the choice of data.
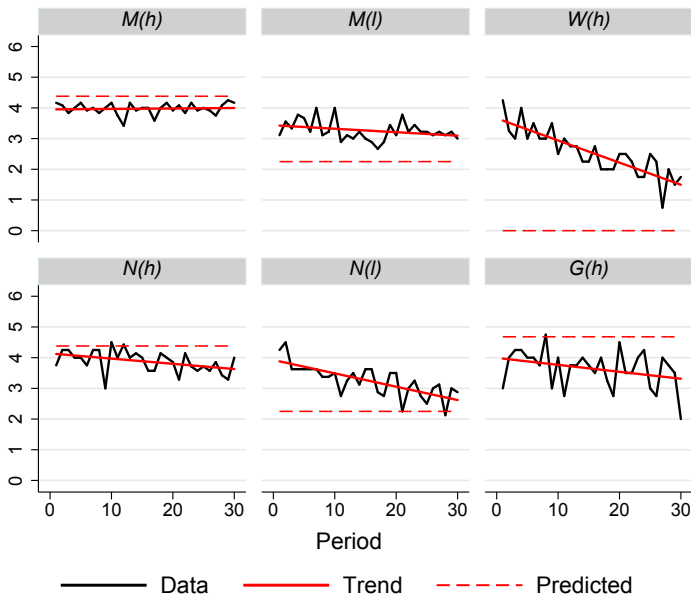
Figure 1.6: Average vacancy creation by period

Figure 1.6 presents for each treatment the average number of vacancies created per period (solid black). Also included on this plot is a linear trend line (solid red) and line indicating

---

[40]Instructions and screen shots will made available when final data collection concludes.

the level of entry predicted by the mixed strategy equilibrium (dashed red). The negative trend mostly disappears in the second half of the data.[41]

We summarize the data on entry in table 1.4. The first row presents the aggregate entry frequency for each treatment while the second row shows the predicted entry frequencies associated with the symmetric mixed strategy equilibrium. A key observation is that all the observed entry frequencies are closer to 0.5 than predicted in theory. The third row presents the payoffs associated with entry. These payoffs are computed as the total earnings from entry divided by the total number of entry decisions.

Table 1.4: Aggregate Entry Frequencies and Associated Payoffs

Treatment	M^h	M^l	N^h	N^l	W^h	G^h
Observed	0.66	0.53	0.62	0.49	0.34	0.58
Predicted	0.73	0.37	0.73	0.37	0.00	0.78
Payoffs	99.5	81.1	98.0	80.1	69.5	92.4

Result 1.5.1. *Vacancy creation increases when the productivity of a job increases.*

The average number of vacancies varies in the expected direction with productivity. As can be seen in table 1.5, which presents summary statistics for block level observations, the high surplus treatments generate almost one additional vacancy per period. This holds in the treatments without bargaining (3.99 in M^h and 3.23 in M^l) and those with bargaining (3.81 in N^h and 2.94 in N^l). Moreover, the maximum observed level of entry in the low surplus treatments is below the minimum observation in the high surplus treatment.

Table 1.5: Summary Statistics Vacancies (Block Level Observations)

		$Y = 400$					$Y = 300$			
		Mean	Std.Dev.	Min.	Max.		Mean	Std.Dev.	Min.	Max.
Matching	M^h	4.00	0.33	3.64	4.33	M^l	3.23	0.29	2.98	3.60
Bargaining	N^h	3.81	0.29	3.54	4.13	N^l	2.94	0.31	2.63	3.27

To formally test the impact of productivity on vacancy creation, we employ Wilcoxson Rank-Sum (WSR) tests.[42] These tests are carried out on the block level data presented

---

[41]This can be seen in figure 1.C.2 which presents the same plots for the last half of the data. Out of treatments M^h - N^l , only N^l exhibits a slight negative trend.

[42]The Wilcoxson Rank-Sum test is a non-parametric test that evaluates the hypothesis that two independent samples come from the same distribution.

in table 1.C.1. The WSR test between treatments M^h and M^l identifies a statistically significant difference between the treatments (W=2.45, p=0.01 with a power of 0.8).[43] The analogous comparison between treatments N^h and N^l yields a similar conclusion (W=2.31, p=0.02 with a power of 0.92). When we pool the high productivity treatments M^h and N^h and compare these with the low productivity treatments M^l and N^l , the results are even stronger (W=3.32, p=0.00; power of 1). These findings indicate that there is a systematically higher degree of entry in the high productivity treatments.

Result 1.5.2. *The presence of bargaining has a small and statistically insignificant impact on vacancy creation.*

The difference between treatments M^h and N^h and between M^l and N^l is the presence of a bargaining stage. Casual inspection of the level of vacancy creation suggests that the presence of bargaining after matching may have some negative impact on firm entry. The average level of entry in the treatments with bargaining is depressed by about 0.25 vacancies per period relative to the case without bargaining (0.19 and 0.29 for the high and low surplus treatments respectively). This difference could be consistent with a small amount of risk aversion related to bargaining uncertainty.[44] We can not, however, substantiate this based on block-level comparisons. Wilcoxon rank-sum tests between M^h and N^h (W=1.23, p=0.22) and between M^l and N^l (W=1.16, p=0.25) do not identify a statistically significant difference between the treatments. This is not surprising since these tests are not sufficiently powered (for the observed average difference, both tests have a power of less than 0.2). As can be seen in table 1.C.1, there is too much overlap between the distributions of observations to reject equality. To achieve a power of 0.80 would require about a ten-fold increase in the number of observations.

If we relax the standard of comparing block level observations and instead treat each block-period observation as independent, we find weak statistical significance (W=1.76, p=0.08 for the high surplus treatments; W=1.8, p=0.07 for the low surplus treatments). However, also these tests are underpowered (power of about 0.55). We conclude that bar-

---

[43]To compute power we use the Stata package developed by Bellemare et al. (2016). The power computation is based on the assumption of a significance level of 0.05.

[44]The ultimatum bargaining stage can be thought of as causing a mean preserving spread of the payoff in the treatments with a fixed wage. In the extreme case, the proposer appropriates the entire surplus half of the time. As a crude exercise based on the aggregate data, we consider CRRA preferences and estimate a risk aversion parameter between 0.05 and 0.1, depending on whether we use the assumption that the proposer appropriates the entire surplus or use the empirically observed proposer share of $(Y - 100)/100$. As we discuss, however, the role of risk aversion is difficult to justify when it does not have a significant impact on entry. An alternative explanation might therefore be that individuals display some aversion to bargaining.

gaining may have some small impact on the level of entry but that this effect is difficult to identify in the data.

Result 1.5.3. *Vacancy creation is less sensitive to productivity changes than predicted by theory.*

Although the number of vacancies created varies in the expected direction, it does not respond as strongly to changes in productivity as predicted by theory. Relative to the theoretical benchmark, the the degree of vacancy creation is too great when the productivity is low while the degree of vacancy creation is too low when the productivity is high (see table 1.4). In the high productivity treatments M^l and N^l , the aggregate entry frequency is more than 10 percentage points higher than the equilibrium prediction (0.53 and 0.49 vs. 0.37). In the low productivity treatments M^h and N^h , the aggregate entry frequency is lower than the equilibrium prediction by a similar amount (0.67 and 0.62 vs. 0.73). The significance of these observations is confirmed by a one-sample Kolmogorov-Smirnov test that compares the block level observations to the anticipated distributions and finds significant differences (p=0.031 for high surplus and p=0.021 for the low surplus). As a consequence of the bias toward 0.5, the distribution of entrants is shifted right in the low surplus treatments and left in the high productivity treatments. This is evident in figure 1.7 which shows the relative frequency of each number of entrants aggregated for a given level of productivity.[45] For example, in both the high and low surplus cases, four vacancies were created about 30% of the time. The observed distributions thus overlap substantially more than anticipated by the theoretical prediction (see figure 1.5 for the theoretical prediction).

The finding that vacancy creation responds more weakly that anticipated to changes in productivity is notable because it goes in the opposite direction of the stylized facts from the labor market (i.e the Shimer critique) but is consistent with the findings from other market entry games.

The pattern of over-entry in M^l and N^l but under-entry in M^h and N^h implies that the no-profit condition is not satisfied, even on average. Test subjects that enter in treatments M^l and N^l earn less than the outside option while test subjects that enter in treatments M^h and N^l earn more than the outside option. The quantitative effect on payoffs is shown in the bottom row of table 1.4. In the high surplus treatments entrants earn excess profits from entry while in the low surplus treatments entrants earn more than 10 ECUs less than

---

[45]The same figures for just the M treatments and just the N treatments are presented in appendix figure 1.C.3.

Figure 1.7: Average Number of Vacancies by Surplus
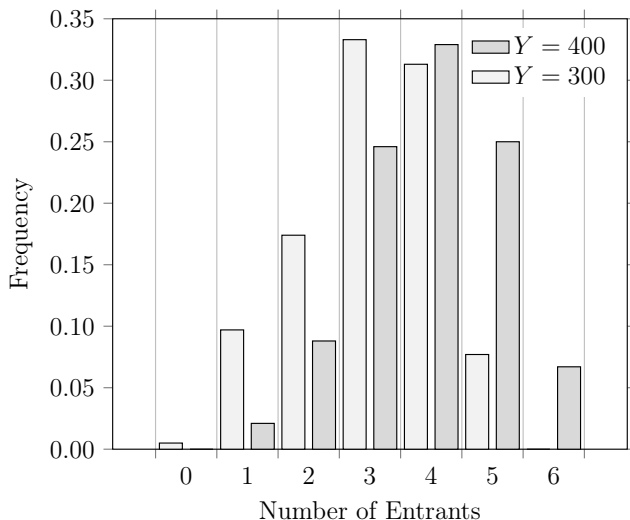
what they could have earned by taking the guaranteed payment.[46]

In the presence of risk aversion, we would expect a lower level of vacancy creation in both the high and low surplus treatments. The finding that entry is too high in the low surplus treatment thus suggests that risk aversion does not play a major role in determining aggregate entry. Over-entry in the low surplus treatments requires risk-seeking preferences while under-entry in the high surplus treatments requires the opposite. While risk aversion may play some role at the individual level in determining who enters, it does not have a discernible effect on the average level of entry.[47] This is consistent with the conclusions from previous studies in this area (Rapoport, Seale, and Ordóñez 2002).

Result 1.5.4. *Offers in ultimatum bargaining are lower and more tightly distributed than in other studies.*

   1. There is substantial coordination on a modal ultimatum bargaining offer and this offer

---

[46]In the treatments with bargaining, payoffs are also impacted by the acceptance and rejection decision. Although this changes the quantitative conclusion about the payoffs from entry (average payoff of 106.4 in N^h and 87.0 in N^l), it does not change the conclusion about the degree of entry.

[47]Although we collected individual level data on risk aversion using a Holt-Laury test conducted at the end of some sessions in treatments N^h and N^l , these data did not have any explanatory power in a regression predicting entry. We question, however, the value of this risk aversion measure because only 65% of respondents demonstrated consistent preferences that did not exhibit preference reversals.

does not vary between treatments.

2. *The average ultimatum bargaining offer varies slightly with the level of productivity but does not depend on whether the proposer is a firm or a worker.*

3. *Workers offer slightly more generous shares to firms in W^h compared with N^h .*

In treatments N^h , N^l , and W^h there is ultimatum bargaining after matching. In N^h and N^l , the proposer is determined by a fair coin while in W^h the workers always get to propose. We present the bargaining results in terms of *offers*. An offer is the the amount that the proposer—who may be either a worker or a firm in treatments N^h and N^l —offers as a take-it-or-leave-it offer to the responder.

Offers fall from a roughly split in the early periods but rapidly stabilize in the vicinity of 100. Figure 1.C.4 in the appendix shows the evolution of bargaining offers over time. This pattern is consistent with test subjects in the proposer role learning the level of acceptable offers. A handful of super equal offers in the early periods suggests that there may have been some confusion about whether a "claim" accrues to the proposer or responder. However, these types of offers quickly disappear. There appears to be rapid convergence on a sharing convention in which offers of 100 are accepted but offers below this level are rejected with high probability.

Figure 1.8 shows the distribution of offers by size. The defining feature of these plots is the large fraction of offers at near 100. In N^h , 47.4% of offers were equal to 100 with nearly 61% of offers 100 ± 10; in N^l the percentage was 39.7% with 58% 100 ± 10; and in W^h the percentage was 39% with 56% 100 ± 10.

In terms of the average offer, we identify a small differences between the two levels of productivity. The average offer is 98 in the high surplus treatment and 107 in the low surplus treatment. This difference is significant at the 5% level. This can be accounted for by the more common occurrence of offers of 50 in N^h , as can be seen at the left hand side of the left panel in figure 1.8.

Whether using the modal offer or the mean offer as the main statistic, the level of offers is more or less constant in terms of absolute levels of ECUs. In terms of shares of the surplus, however, this implies that individuals in the high surplus treatment share about a quarter of the surplus while individuals in the low surplus treatment share a bit more than a third of the surplus. This contrasts with the findings from most studies of ultimatum bargaining. An initial hypotheses as to why bargaining shares vary with the size of the surplus is that players might be willing to accept smaller *shares* when the surplus is large. While Andersen

Figure 1.8: Offers in Ultimatum Bargaining
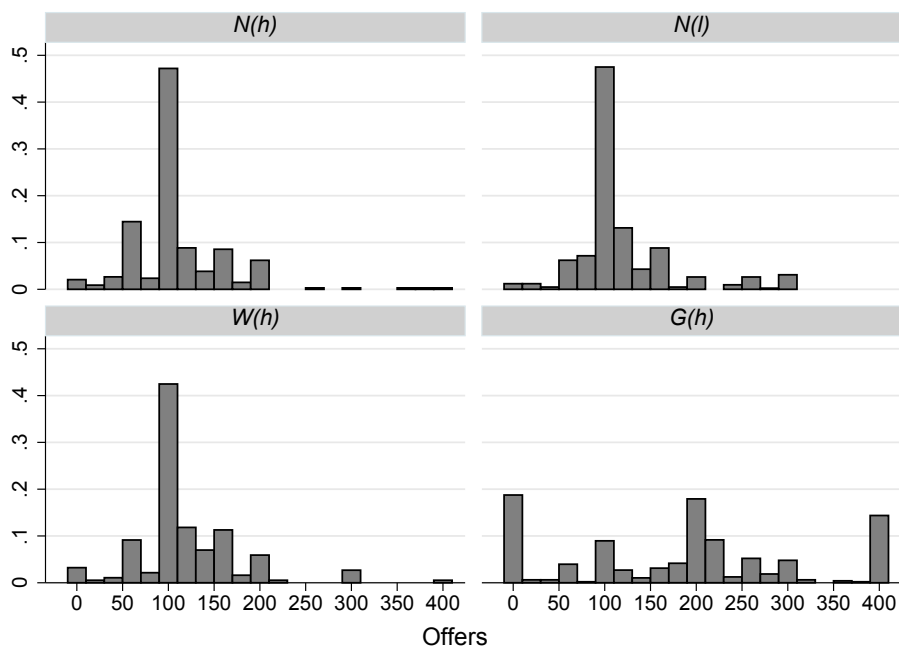
et al. (2011) finds some support for this conjecture, it only holds for enormous sums. The stakes would therefore need to be much larger than those seen in this study. We therefore do not expect this effect to be relevant in our results. We consider some alternative explanations in the next section.

With respect to type, we do not find a significant difference between offers: Overall, 97 by firms compared with 104 by workers.

With respect offers that get rejected, we find that the probability that offers get rejected increases steeply as the offers fall below 100.[48] Histograms showing the offers that were rejected are presented in figure 1.C.5 in the appendix. There do not appear to be any

---

[48]The stylized facts from ultimatum bargaining is that there are few low offers and offers less than 20 percent are rejected about half the time (Camerer 2003). These outcomes are also surprisingly insensitive to the size of the stakes; the size of stakes must be increased dramatically before low offers are accepted (Andersen et al. 2011). Although generosity may result in part from fear of rejection, positive offers in dictator games suggest some role of altruism. A notable *caveat* is that combination of design features such as anonymity and earned endowments can reduce dictator offers close to zero (Cherry et al. 2002)

difference between types in terms of which offers were rejected.

Overall, the level of bargaining offers are lower than in most studies of ultimatum bargaining but not implausible. Standard results from studies of ultimatum bargaining are modal and median offers of 40-50 percent and mean offers of 30-40 percent—but with considerable variability across studies (See Camerer 2003, Tables 2.2 and 2.3). The main difference from earlier studies seems to be the tight distribution of offers in the vicinity of offers of 100. Notably. this is close to 95, the direct search cost borne by firms. The high degree of coordination on offers around 100 also implies an average η close to 0.5 since a firm is equally likely to be in the proposer and responder position.

The overall lower level of offers we find may be due to the presence of frictional matching. Individuals realize that they have limited opportunities for bargaining and therefore seek to earn as much as possible when they have the chance. We also speculate that the overall lower level of offers may in part be driven by the selection of firm players that choose to enter the bargaining portion of the game.

Result 1.5.5. *Wage offers in ultimatum bargaining maximize individual payoffs.*

It is natural to ask why there is coordination on offers near 100. From an empirical perspective, this question has a clear answer: Offers of 100 maximize expected payoffs. Although the standard model predicts that any non-negative offers should be accepted in ultimatum bargaining, in practice lower offers are rejected with a higher probability. The optimal bargaining offer therefore trades off the likelihood of being accepted with the value of agreement. To assess this formally, we estimate a logistic regression to yield the probability that an offer of a given size is accepted.[49] This exercise shows that subjects almost never reject offers above 100 but reject offers with an increasing probability as offers fall below this level. When we weight the probability that an offer is accepted by the value that the offer is accepted, this gives the expected payoff from a given offer.

In figure 1.9, we present a plot of the expected payoff associated with each ultimatum offer along with a simple polynomial fit. This plot illustrates that expected payoffs are maximized in the vicinity of offers of 100. We conclude from this exercise that test subjects maximize their individual payoffs, taking into account that offers less than 100 are often rejected.

Result 1.5.6. *Increases in the bargaining strength of workers in ultimatum bargaining increases unemployment.*

---

[49]See figure 1.C.6 in the appendix.

Figure 1.9: Expected Payoff by Offer
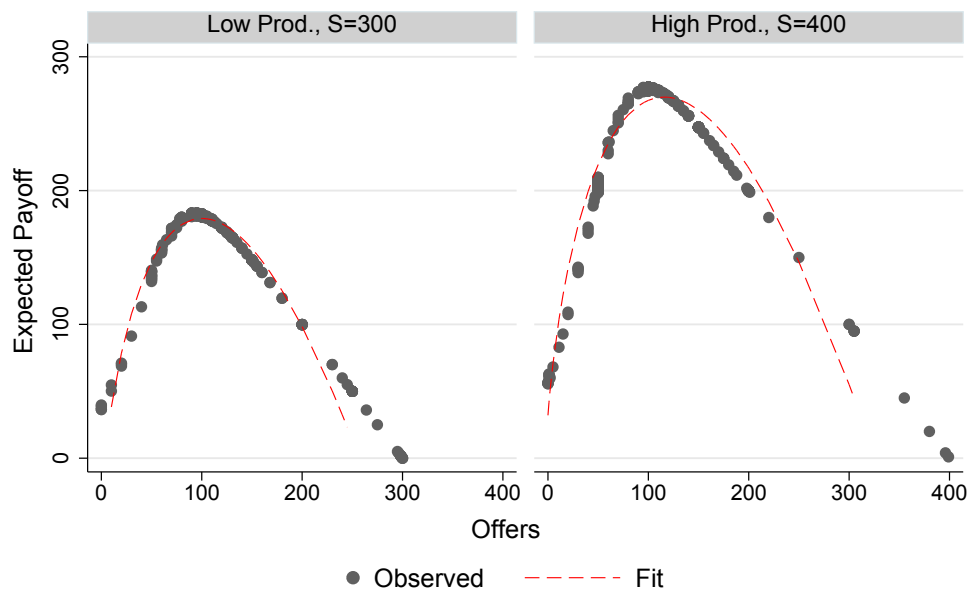
In treatment W^h, the workers always get to propose in the bargaining stage. In figure 1.6, we see that the average number of vacancies falls from around 4 in the first period to about 1.5 in the final period. This is due to the fact that the workers are not able to coordinate on a higher wage level and firms gradual cease to enter. Given that the expected wage in W^h is only slightly greater than 100, the market capacity is only a single vacancy. It is in some sense remarkable that firms persist in entering as long as they do given that they sacrifice almost 25 ECUs per entry decision. Notably, workers in W^h earn relatively high payoffs because workers exploit the firms who do persist in entering.

Result 1.5.7. *When workers can commit to a common wage offer, this restores efficiency.*

Treatment G^h changes the bargaining institution and allows workers as a group to commit to a wage level. In contrast to the ultimatum bargaining treatments in which offers tended

to be in the vicinity 100, in this treatment the average offer was 174. The expected share that workers could expect was therefore $\eta = 0.44$. This translates to a market capacity of about three firms. The substantial moderation of wage demands leads to a much higher level of entry than in the W^h.

Although the wage demands are substantially higher than the level that is predicted to maximize expected profits, workers in this treatment also do better than expected in theory because of somewhat too much entry by firms.

Result 1.5.8. *When workers decide a group wage, there is considerable variation in beliefs about the appropriate level of the wage.*

Perhaps the most fascinating element of treatment G^h is the wage agreement stage that occurs between workers. In this treatment, each worker submitted a wage proposal at the beginning of the stage. This was then advertised to the firms before the entry decision. Workers thus have an incentive to set a wage that trades off increases in the probability of matching with the size of the wage.

Figure 1.10: Wage Proposals
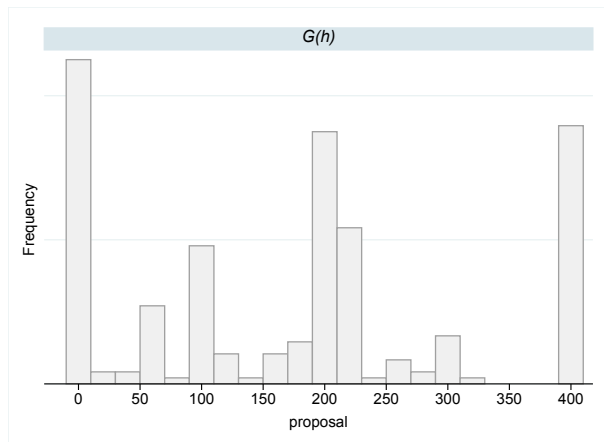
Figure 1.10 shows the wage proposals made by workers in G^h. In contrast to the offer distributions in the treatments with ultimatum bargaining which were tightly distributed around 100, in G^h we observe a tri-modal distribution. The distribution of wage offers has peaks at 0, 200, 400. These offers correspond to all, half, and none of the surplus. This testifies to substantial differences in beliefs about the optimal group wage level. The

extreme offers suggest that some individuals are strategically setting wage offers to try to influence the average wage proposal. It is evident from this finding that (at least some) test subjects understand the subgame structure of the experiment and understand that different wage levels translate into payoffs not just through their impact on the wage but also via the vacancy creation channel. The result that test subjects dramatically adjust their offers relative to the treatments with ultimatum bargaining after match—only a few make proposals of 100—shows that the wage formation institution has a strong impact on what offers can be sustained. In particular, workers are not able to restrain their wage demands in the absence of a centralized institution.

Result 1.5.9. *There is more heterogeneity at the individual level than can be accounted for by the symmetric mixed strategy Nash equilibrium*

The aggregate entry probabilities in the experiment are somewhat biased but reasonably close to the level predicted by symmetric randomization. However, a given aggregate level of entry is consistent with a wide range of patterns of entry at the individual level. In the simplest case, consider two test subjects. In terms of the expected number of entrants, the case when both randomize with a probability 0.5 and the case when one test subject always enters and one subject never enters is identical.

At the individual level, we find that there is substantial heterogeneity in excess of that predicted by symmetric randomization in how often subjects choose to enter. For each individual, we compute the proportion of the total number of periods in which an individual enters. Figure 1.11 shows this distribution of entry frequencies at the individual level. Some test subjects enter often and some rarely do. The strategy of always entering and nearly never entering are both relatively common.

In addition, consistent with expectations, in the treatments with higher productivity the distribution of entry frequencies stochastically dominates the equivalent treatment with lower productivity.[50]

If all test subjects in a given treatment randomize with the same probability p, this predicts a binomial distribution of entry frequencies across the 30 periods with mean $30p$ and variance $30p(1-p)$ and a probability of a given number of entry decisions k equal to $P(X = k) = \binom{30}{k} p^k (1-p)^{30-k}$. This is the outcome we expect if we take the symmetric Nash equilibrium as the solution.

In figure 1.11, we present histograms showing the frequency of each level of entry. A first observation is that the distributions are much flatter than those predicted from the mixed

---

[50]We present the CDF in figure 1.C.7 in the appendix.

Figure 1.11: Individual Entry Frequencies
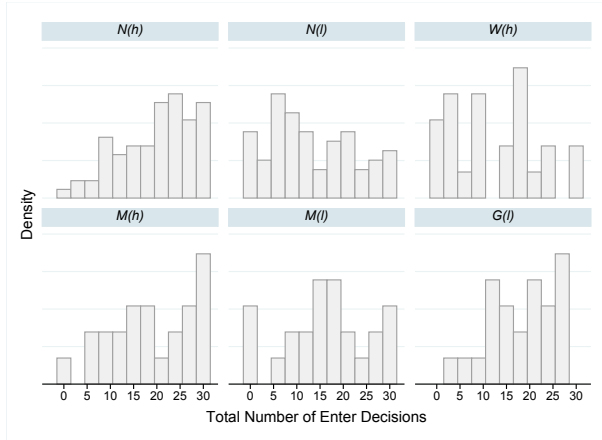
strategy Nash prediction with larger standard deviations. A second crucial observation is that there is a substantial mass of test subjects in each treatment that choose to enter in every period or nearly every period. In treatment T1, 8 individuals choose to enter in every period (about 10% of all participants).

The question is whether these data are consistent with randomization. If subjects randomized with a symmetric probability p we would expect a binomial distribution over the $N = 30$ trials. Relative to this benchmark, the observed distributions are too flat and wide (with, again, the caveat that the results are based on only 24 individuals in each treatment. The distributions might become more peaked with additional observations.). This suggests that entry can not be well accounted for by symmetric randomization.

To further investigate the role of randomization, we investigate the *pattern* of entry decisions. Specifically, we investigate how persistent subjects are in their decision to enter. In figure 1.C.8, we present the distribution of "streaks," where a streak is defined as consecutive decisions to enter. As is evident from the figure, most of the mass is concentrated on streaks of length less than four. In fact, for short streaks, the frequency falls close to geometrically. This implies that players switch fairly often between strategies.

Notably, the main differences across distributions is observed in the tails. For instance, the geometric pattern for short streaks is quite similar across treatments. In *both* high and low treatments, the geometric probability is close to 0.5. The main difference between treatments instead arises because in the high surplus treatment there are more players who enter very

many times in a row. This implies that the aggregate differences between treatments can be attributed to a subset of players who use extended streaks.

From the individual analysis, we find arrive at a conclusion similar to that in Zwick and Rapoport (2002). There exists a subset of subjects who play as if randomizing—although not necessarily with symmetric probabilities and perhaps as consequence of sequential dependencies—as well as a population of subjects who do not adjust their strategies over time.

1.5.1 Analysis

Quantal Response Equilibrium

The environment we investigate is complex. Even the baseline treatments feature both strategic and stochastic uncertainty. It is reasonable to expect participants to make mistakes. The pertinent question in this regard is how errors affect aggregate outcomes. In a naive framework, we would expect idiosyncratic errors to cancel out on average. However, if players are sophisticated and recognize that other players may make errors, and that these other players also have beliefs about error (and beliefs about beliefs etc.), then mistakes in decision-making can have equilibrium consequences.

A rigorous way of modelling this is via the estimation of a quantal response equilibrium (QRE). In this framework, players decisions are perturbed by some random noise that is introduced into payoffs. In addition, this framework requires that beliefs be correct. That is, beliefs correspond to equilibrium choice probabilities.

The QRE provides an intuitive explanation for the bias we observe in aggregate entry (See Goeree, Holt, and Palfrey 2016, chapter 8, for a discussion of QRE in the context of participation games). In the market entry context, the expected payoff from entry depends negatively on the probability that other individuals enter. We denote this by $\mathbb{E}[\pi(p)]$ where p is the symmetric entry probability. Letting payoffs from entry and staying out be affected by some independently and identically distributed noise ϵ_e and ϵ_o respectively, then the choice to enter will be chosen in the case when $\mathbb{E}[\pi(p)] + \lambda\epsilon_e > X + \lambda\epsilon_o$, where λ is a parameter that can be interpreted as magnitude of noise in the environment. Notice that as noise increases behavior becomes more random, with choice probabilities going to $1/2$ as $\lambda \to \infty$. This characterization implies a choice probability $p = Pr\left(\epsilon_o - \epsilon_e < \frac{\mathbb{E}\pi_e(p) - X}{\mu}\right) = F(\mathbb{E}[\pi(p)] - X)$.

Manipulating this further yields

$$\mu F^{-1}(p) = \mathbb{E}[\pi(p)] - X \qquad (1.6)$$

The two sides of this equation are shown in figure 1.12.
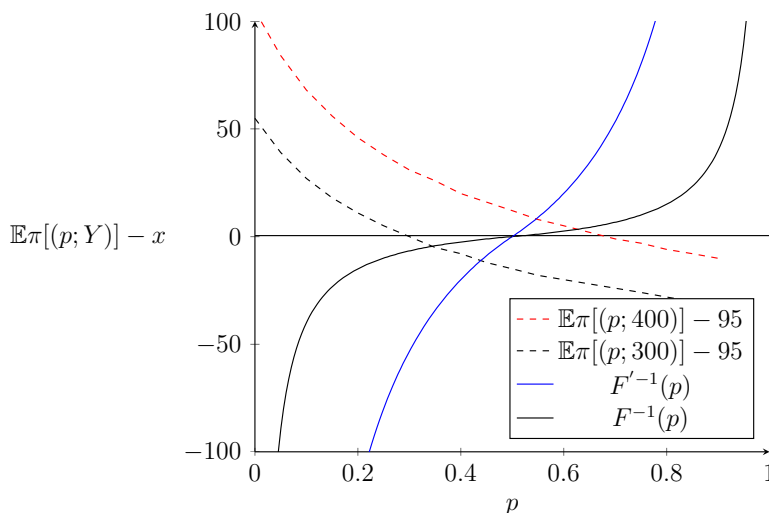
Figure 1.12: Quantal Response Equilibrium

The left hand side is illustrated by the upward sloped curves while the right hand side is illustrated by the downward sloping curves; the intersection of these curves thus gives the QRE entry probabilities (on the horizontal axis). With respect to the Nash equilibrium—which is given by the intersection of the expected payoff (downward sloping curve)—it's clear that the QRE is always biased toward 0.5. Furthermore, as the noise parameter λ increases, the upward sloping curves become steeper and lead to probabilities that are close to 0.5. Qualitatively, the QRE can justify the bias observed in the aggregate data.

At the individual level, however, we observe substantially more heterogeneity in entry than predicted by a symmetric randomization. We observe a substantial number of individuals who always choose to enter, even when they persistently earn less than the fixed payment. In addition, in our low surplus treatment the entry frequency is close to 0.50.[51] This can only be reconciled to a symmetric QRE if behavior is totally random. It is therefore

---

[51]In fact, if we consider all periods, the entry frequency is a few percentage points in excess of 0.50.

impossible to simultaneously fit both the low and high surplus treatments.

To try to reproduce our data in a model framework, we therefore estimate a heterogeneous QRE (HQRE) in which we allow idiosyncratic subject-level preferences for entry. We assume that every individual draws a preference parameter δ at the beginning of the experiment such that the payoff from entry is $\mathbb{E}[\pi(p)] + \delta$. Our approach is similar to that employed by Goeree, Holt, and Moore (2015). We model individuals as drawing a type δ_f from a discrete distribution $\tilde{\phi}(\mu, \sigma)$ that is generated from a truncated normal distribution. Specifically, we truncate a normal distribution to the region $\mu \pm \sigma$ and consider 11 points equal spaced in this interval $(\mu - \sigma, \mu - 0.8\sigma, \ldots, \mu + \sigma)$. To discrete-ize the distribution, the density at each point is divided by the sum of the densities at all points. Maintaining the assumption that each individual experiences some idiosyncratic extreme value distributed noise each period, the quantal response probabilities can still be written down in the logistic form—for each individual f, δ_f is simply a constant that adjusts the expected payoff from entry.[52]

Formally, the estimation has two stages. In the first stage, a QRE is fit for a vector of parameters that include the level of noise λ and the parameters of the normal distribution μ and σ. This fitting was done by minimizing the difference between beliefs and the quantal response probabilities based on the logit specification. After estimating the theoretical QRE entry probabilities, we then compute the likelihood of observing data given the choice of parameters. The likelihood function is based on the observed number of entry decisions for each individual. For each individual f in each treatment T, we tabulate the number of vacancy creation decisions $v(f, T)$ and then compute the likelihood function

$$\Pi_{f=1}^{F(T)} \left[\sum_{z=1}^{Z} p\left(\lambda, \mu, \sigma, T\right)^{v(f,T)} \left(1 - p\left(\lambda, \mu, \sigma, T\right)\right)^{30-v(f,T)} \tilde{\phi}(z, \mu, \sigma) \right]. \qquad (1.7)$$

Relative to a binary QRE without heterogeneity, the difference is the presence of the $\tilde{\phi}(z, \mu, \sigma)$ which accounts for the distribution of possible types.

Despite adding some explanatory power to our estimation and helping to account for the excess heterogeneity at the individual level, it remains difficult to jointly fit both the high and low surplus treatments. One issues appears to be the data are skewed in a fashion that makes it difficult to fit with a symmetric distribution. We are therefore investigating the use

---

[52]We also estimated a simple model in which we estimated a simple additive constant for an estimated proportion of the population q. The idea was to try to account for the significant portion of players who always enter. However, this approach gave quite dramatic results. The difficult is that there is a left skew in entry in most the treatments. As a consequence, q must be relatively large. The outcome was therefore that a large proportion of individuals never entered while another large proportion always entered.

of an asymmetric distribution such as the Gamma.

The HQRE helps account for some—though not all—of the variation in the data. This exercise suggests that aggregate biases can survive in environments with noise. This may be important beyond the lab.

Learning Rules

As an alternative way of accounting for the observed patter of entry we estimate the predictive ability of two simple learning rules, *payoff reinforcement* and *fictitious play*. Payoff reinforcement is based on the notion that if an action had a positive outcome, then an agent becomes more likely to repeat the action (Roth and Erev 1995; Erev and Roth 1998). By comparison, fictitious play is a type of belief learning in which the probability of an action increases if it *would have been* a best response.[53]

For both types of learning, we compute the propensities that a subject will choose to enter given their history of payoffs and market information. Although both rules require "initialization" (i.e. initial choice propensities and the "stickiness" of initial propensities), the outcomes are mostly insensitive to these choices. After estimating the propensities, we use the estimated entry probabilities to predict actual entry[54]. In a regression based on payoff reinforcement probabilities, the estimated coefficient is highly significant and of a large size. In contrast, the same coefficient for fictitious play is insignificant.

Payoff Reinforcement		
$\mathbb{P}$	Enter	Stay Out
Enter	76.2	23.8
Stay Out	18.8	81.2

Figure 1.13: $\mathbb{P}_{\text{Enter}}$ reinforcement learning predictions

Fictitious Play		
$\mathbb{P}$	Enter	Stay Out
Enter	52.5	48.5
Stay Out	38.7	61.3

Figure 1.14: $\mathbb{P}_{\text{Enter}}$ belief learning predictions

---

[53]Note that players do not need to know their opponents' payoffs: Belief learning is simply about expectations about other players' actions.

[54]We estimate a logistic regression with fixed/random subject and time effects

To give a gross intuition for the predictive ability of these learning rules, we categorize the learning rule as predicting "enter" whenever the estimated propensity is greater than 0.5. Next, we cross-tabulate the prediction with the actual entry decision. Figure 1.13 shows the results for payoff reinforcement: In about 80% of cases, the prediction is correct. In contrast, results for reinforcement learning, shown in 1.14, demonstrate that this learning rule has almost not predictive ability.

Fairness Preferences

With respect to the bargaining portion of the experiment, we find rapid and consistent coordination on a sharing norm. This sharing norm is the same in nominal (ECUs) terms across treatments. This implies that the share that the proposer keeps increases with the size of the pie. This last finding is is unusual.

A common way of explaining outcomes in ultimatum and dictator games is via the use of fairness preferences. In the conceptualization of Fehr and Schmidt (1999), other players' payoffs enter the utility function explicitly, with penalties associated with unequal payoffs. In the version proposed by Bolton and Ockenfels (2000), players are penalized when their earnings deviate from a sharing norm. Both approaches have been successful in reproducing a variety of stylized facts not well accounted for by basic rational behavior.

Because this study incorporates a standard ultimatum bargaining game, it is plausible that fairness preferences could help explain the results. Moreover, because fairness preferences concern *utility*, outcomes of the bargaining stage will also affect the entry decision. This means that fairness preferences have implications for both bargaining and entry.

However, theoretical consideration of these preferences suggests that neither the inequity preferences of Fehr and Schmidt (1999) nor the equality-reciprocity-competition preferences of Bolton and Ockenfels (2000) can explain the results from this experiment. Even under favorable assumptions on preference parameters, solving for an equilibrium with fairness preferences implies sharing of the surplus and a pattern of entry inconsistent with what we observe. Fairness preferences in this study would predict substantially *lower* offers than what we observe. The reason is that unmatched players earn 0 and this actually loosens the constraint on the offers that proposers can make. Specifically, predicted offers tend to be in the range of 20-30 ECUs, that is, a third of what we find (allowing for some differences given various parametrizations). Details of the computation are found in appendix 1.C. It is also the case that bargaining outcomes predicted by these fairness preferences do not reproduce the pattern across treatments.

Although the two main varieties of fairness preferences do not explain the results, the notion that fairness concerns are important has appeal. An observation that is compatible with this intuition is that the modal offer of a 100 is close to the *average* earnings in the market. In tables 1.15 and 1.16, we show the average payoffs for the case of all subjects in a market ("All") and the case of just the players in the matching stage ("In").

Average Payoffs, Y High		
v	All	In
1	87.5	80.0
2	94.6	94.3
3	97.8	98.0
4	99.0	100.0
5	98.9	99.4
6	98.0	98.0

Figure 1.15: Average payoffs by comparison group, Y High

Average Payoffs, Y Low		
v	All	In
1	77.5	60.0
2	80.4	70.7
3	80.5	74.2
4	79.0	75.0
5	79.6	74.5
6	73.5	73.5

Figure 1.16: Average payoffs by comparison group, Y Low

Under this norm, an offer is accepted if it is at least as large the average earning. This explanation would be consistent with the observations from this study. Once offers fall below the market average, then begin to get rejected.

1.6 Conclusion

This study examines entry in a market entry game with labor market features. We find a high degree of coordination even though the level of entry is less sensitive to changes in productivity than predicted by the Nash equilibrium. This mirrors the results in the existing literature. The coordination success of agents in the model, though not perfect,

does not seem to be perturbed by the additional complications we introduce. With respect to the market entry literature, we conclude that the complications that we introduce do not strongly alter aggregate outcomes. Indeed, the main bias that we observe is a general finding in the literature and not a peculiar aspect of our environment.[55] The finding that individual behavior is chaotic also is repeated in our findings. Overall, the "magic" of coordination appears to be a robust phenomenon even if some puzzles remain.

With respect to the bargaining results, we find that that offers are "optimal" in the sense that proposers makes offers that maximize their payoffs. We believe that the overwhelming reason for modal sharing of 100 relates to the direct cost of vacancy creation. Although it is fair to compensate firms for this cost, workers do not reward firms for the matching risk they bear. This leads to negative consequences for efficiency in treatment W^h . This illustrates why bargaining after matching is often inefficient. Notably, however, when we create an institution in which workers can coordinate their wage setting, they are able to rectify the situation and improve market outcomes. This suggests that in certain instances, workers should be willing to engage an organization to negotiate on their behalf. This result provides a new finding that explains the value of labor organizations.

Relative to the patterns observed in actual labor markets (and, indeed, our *ex ante* expectation, especially given the matching technology), we would perhaps have expected the opposite of our results: More volatility in entry and employment than in productivity. Despite the deviations from predictions, the results should still be interpreted as good news for the validity of the market entry assumption and for the basic assumptions that underpin labor search models. In the average across treatments, the level of entry is about correct. This is due to the fact the aggregate biases cancel out across high and low productivity settings. In the absence of more specific information about a market—and when considering industries aggregated across an entire economy—neither over nor under entry is more likely. In addition, our study examines an extremely small market yet subjects are still relatively close to satisfying the zero-profit condition. Even though the assumption of a continuum of agents and free-entry in macrolabor models is strong, the assumptions still hold—more or less—in the small market setting. That the assumptions holds both in small markets as well as large markets, for which it seems likely that the effects of strategic uncertainty are smaller, is good news for the DMP model. The one caveat to the positive conclusion for the model is that the presence of noise may lead to persistent over or under entry. This

---

[55]It is, however, difficult to tell if this effect is more or less strong in our environment, given that we make a number of alternate design decisions.

could help explain patterns of entry in new industries in which the potential profits are less certain.

1.7 References

References

Andersen, S., Ertaç, S., Gneezy, U., Hoffman, M., & List, J. A. (2011). Stakes matter in ultimatum games. *The American Economic Review*, 3427–3439.

Anderson, S. P. & Engers, M. (2007). Participation games: Market entry, coordination, and the beautiful blonde. *Journal of Economic Behavior & Organization*, *63*(1), 120–137.

Andersson, O. & Holm, H. J. (2010). Endogenous communication and tacit coordination in market entry games: An explorative experimental study. *International Journal of Industrial Organization*, *28*(5), 477–495.

Bellemare, C., Bissonnette, L., & Kröger, S. (2016). Simulating power of economic experiments: The powerbbk package. *Journal of the Economic Science Association*, *2*(2), 157–168.

Berg, J., Dickhaut, J., & McCabe, K. (1995). Trust, reciprocity, and social history. *Games and economic behavior*, *10*(1), 122–142.

Berry, S. & Reiss, P. C. (2007). Empirical models of entry and market structure. *Handbook of industrial organization*, *3*, 1845–1886.

Bolton, G. E. & Ockenfels, A. (2000). Erc: A theory of equity, reciprocity, and competition. *American economic review*, 166–193.

Bresnahan, T. F. & Reiss, P. C. (1990). Entry in monopoly market. *The Review of Economic Studies*, *57*(4), 531–553.

Camerer, C. (2003). *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press.

Camerer, C. & Lovallo, D. (1999). Overconfidence and excess entry: An experimental approach. *American economic review*, 306–318.

Charness, G. & Kuhn, P. (2011). Lab labor: What can labor economists learn from the lab? *Handbook of labor economics*, *4*, 229–330.

Cherry, T. L., Frykblom, P., & Shogren, J. F. (2002). Hardnose the dictator. *American Economic Review*, 1218–1221.

Ciliberto, F. & Tamer, E. (2009). Market structure and multiple equilibria in airline markets. *Econometrica*, *77*(6), 1791–1828.

Cooper, R. W. (1999). *Coordination games*. Cambridge University Press.

Cooper, R. W., DeJong, D. V., Forsythe, R., & Ross, T. W. (1990). Selection criteria in coordination games: Some experimental results. *The American Economic Review*, 218–233.

Duffy, J. (2008). Macroeconomics: A survey of laboratory research. *Handbook of experimental economics*, *2*.

Duffy, J. & Hopkins, E. (2005). Learning, information, and sorting in market entry games: Theory and evidence. *Games and Economic behavior*, *51*(1), 31–62.

Elsby, M. W., Michaels, R., & Ratner, D. (2015). The beveridge curve: A survey. *Journal of Economic Literature*, *53*(3), 571–630.

Erev, I. & Rapoport, A. (1998). Coordination, "magic," and reinforcement learning in a market entry game. *Games and Economic Behavior*, *23*(2), 146–175.

Erev, I. & Roth, A. E. (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American economic review*, 848–881.

Falk, A. & Fehr, E. (2003). Why labour market experiments? *Labour Economics*, *10*(4), 399–406.

Fehr, E., Kirchsteiger, G., & Riedl, A. (1996). Involuntary unemployment and non-compensating wage differentials in an experimental labour market. *The Economic Journal*, 106–121.

Fehr, E. & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *Quarterly journal of Economics*, 817–868.

Fischbacher, U. (2007). Z-tree: Zurich toolbox for ready-made economic experiments. *Experimental economics*, *10*(2), 171–178.

Gary-Bobo, R. J. (1990). On the existence of equilibrium points in a class of asymmetric market entry games. *Games and Economic Behavior*, *2*(3), 239–246.

Goeree, J. K. & Holt, C. A. (2000). An explanation of anomalous behavior in binary-choice games: Entry, voting, public goods, and the volunteers' dilemma. *Unpublished paper, University of Virginia*.

Goeree, J. K., Holt, C. A., & Moore, A. K. (2015). An experimental examination of the volunteer's dilemma. *Unpublished working paper*.

Goeree, J. K., Holt, C. A., & Palfrey, T. R. (2016). *Quantal response equilibrium: A stochastic theory of games*. Princeton University Press.

Greiner, B. (2004). The online recruitment system orsee 2.0-a guide for the organization of experiments in economics. *University of Cologne, Working paper series in economics*, *10*(23), 63–104.

Hagedorn, M. & Manovskii, I. (2008). The cyclical behavior of equilibrium unemployment and vacancies revisited. *The American Economic Review*, *98*(4), 1692–1706.

Helland, L., Moen, E. R., & Preugschat, E. (2017). Information and coordination frictions in experimental posted offer markets. *Journal of Economic Theory*, *167*, 53–74.

Johnson, N. D. & Mislin, A. A. (2011). Trust games: A meta-analysis. *Journal of Economic Psychology*, *32*(5), 865–889.

Kahneman, D. (1988). Experimental economics: A psychological perspective. In *Bounded rational behavior in experimental games and markets* (pp. 11–18). Springer.

Layard, P. R. G., Nickell, S. J., & Jackman, R. (2005). *Unemployment: Macroeconomic performance and the labour market*. Oxford University Press on Demand.

List, J. A., Sadoff, S., & Wagner, M. (2011). So you want to run an experiment, now what? some simple rules of thumb for optimal experimental design. *Experimental Economics*, *14*(4), 439–457.

Moen, E. R. (1997). Competitive search equilibrium. *Journal of political Economy*, *105*(2), 385–411.

Moen, E. R. & Rosen, A. (2006). Incentives in competitive search equilibrium and wage rigidity.

Morgan, J., Orzen, H., & Sefton, M. (2012). Endogenous entry in contests. *Economic Theory*, *51*(2), 435–463.

Morgan, J., Orzen, H., Sefton, M., & Sisak, D. (2012). *Strategic and natural risk in entrepreneurship: An experimental study*.

Morris, S. & Shin, H. S. (2002). Measuring strategic uncertainty.

Noussair, C. N., Plott, C. R., & Riezman, R. G. (1995). An experimental investigation of the patterns of international trade. *The American Economic Review*, 462–491.

Ochs, J. (1990). The coordination problem in decentralized markets: An experiment. *The Quarterly Journal of Economics*, 545–559.

Petrongolo, B. & Pissarides, C. A. (2001). Looking into the black box: A survey of the matching function. *Journal of Economic literature*, 390–431.

Pissarides, C. A. (2000). Equilibrium unemployment theory. *MIT Press Books*, *1*.

Rapoport, A. & Seale, D. A. (2008). Coordination success in non-cooperative large group market entry games. *Handbook of Experimental Economics Results*, *1*, 273–295.

Rapoport, A., Seale, D. A., Erev, I., & Sundali, J. A. (1998). Equilibrium play in large group market entry games. *Management Science*, *44*(1), 119–141.

Rapoport, A., Seale, D. A., & Ordóñez, L. (2002). Tacit coordination in choice between certain outcomes in endogenously determined lotteries. *Journal of Risk and Uncertainty*, *25*(1), 21–45.

Rapoport, A., Seale, D. A., & Winter, E. (2002). Coordination and learning behavior in large groups with asymmetric players. *Games and economic behavior*, *39*(1), 111–136.

Roth, A. E. & Erev, I. (1995). Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games and economic behavior*, *8*(1), 164–212.

Selten, R. & Güth, W. (1982). Equilibrium point selection in a class of market entry games. In *Games, economic dynamics, and time series analysis* (pp. 101–116). Springer.

Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. *American economic review*, 25–49.

Sundali, J. A., Rapoport, A., & Seale, D. A. (1995). Coordination in market entry games with symmetric players. *Organizational behavior and Human decision processes*, *64*(2), 203–218.

Toivanen, O. & Waterson, M. (2000). Empirical research on discrete choice game theory models of entry: An illustration. *European Economic Review*, *44*(4), 985–992.

Van Huyck, J. B., Battalio, R. C., & Beil, R. O. (1991). Strategic uncertainty, equilibrium selection, and coordination failure in average opinion games. *The Quarterly Journal of Economics*, 885–910.

Zwick, R. & Rapoport, A. (2002). Tacit coordination in a decentralized market entry game with fixed capacity. *Experimental economics*, *5*(3), 253–272.

Appendix

1.A   Background

1.A.1   Market Entry Games

This study builds on the experimental literature on market entry games. In the version of the market entry game that has attracted the most attention, payoffs decline linearly in the number of entrants while the opportunity cost is constant (Kahneman 1988; Rapoport et al. 1998; Sundali et al. 1995). In this set-up, the payoff for player i when $v - 1$ other players enter is

$$\pi_i = \begin{cases} X & \text{if } i \text{ doesn't enter} \\ X + \beta(C - v) & \text{if } i \text{ enters} \end{cases} \tag{1.8}$$

where β is a scaling factor that dictates how sensitive payoffs are to entry and C is the market capacity.

The basic result in the literature on market entry games is that there is a high degree of coordination. Consider the study by Sundali et al. (1995). In this study, players in groups of 20 played a market entry game in which the market capacity, C, was varied between each period. The payoffs are a linear function of entrants such that in equilibrium the marginal entrant is indifferent between entering and staying out. Both C and $C - 1$ entrants are therefore equilibria. In one treatment, players were only provided information on market capacity ("no info") while in another treatment players were also given feedback on payoffs ("info"). Figure 1.A.1 presents results. In this figure, the market capacity is given on the horizontal axis and the number of entrants is shown on the vertical axis. This figure illustrates the remarkable degree of coordination, even in the treatment with limited information. The level of entry is in almost all cases close to the level that equalizes expected payoffs from entry with the outside option.
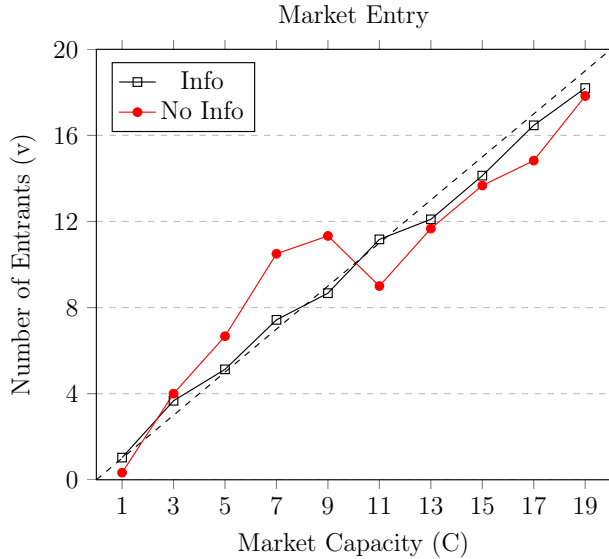
Figure 1.A.1: Entry in Sundali et al. (1995)

The second stylized fact on market entry games that is relevant to our study is the observation that there is more moderate entry—in the sense that participants enter with closer to a 0.5 probability—than predicted by the Nash equilibrium. The study by Sundali et al. (1995) also illustrates this bias. They find too much entry when the capacity is low and too little entry when the capacity is high. This can be seen in figure 1.A.1. When the market capacity exceeds 50% there is a switch from over to underentry. This feature is especially prominent in the "no info" treatment.

1.B   Model

1.B.1   Matching Function

The matching technology is shown in figure 1.B.1. The total number of matches is a function of the number of firm entrants and has the shape shown for the specification of the matching function used in this study. This function increases monotonically in the number of entrants.
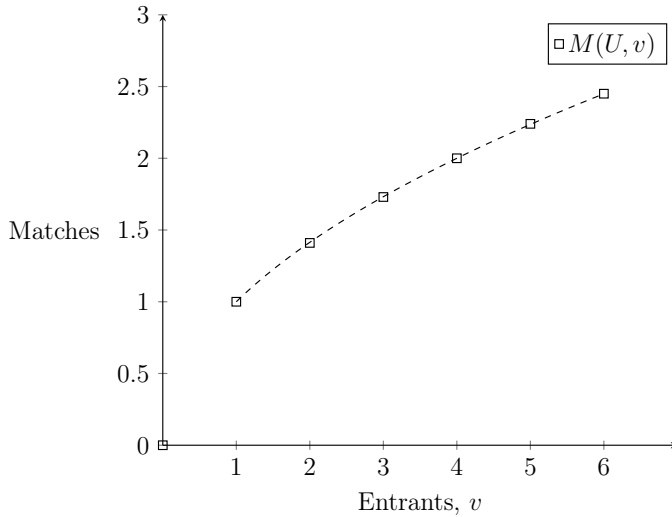
Figure 1.B.1: Total number of matches, given entry v

Although the number of matches increases with the number of firm entrants, the matching probabilities decline. Figure 1.B.2 is a plot of the associated matching probabilities for firms, $q(\theta)$. The effect of an additional entrant is to reduce the matching probability for all other players. Moreover, the negative externality imposed on other firms is most severe when entry is low: For instance, if a single firm enters, then that firm is guaranteed to match. However, if a second firm were to enter the matching probability drops to 71%.

1.B.2   Efficiency

We define expected total *social surplus* as the sum of payoffs to firms and workers:

$$\Sigma = M(U, v)Y + (F - v)X \tag{1.9}$$

$$1$$

$$0.75$$

Match prob. 0.5

$$0.25$$

1   2   3   4   5   6

Entrants, V

$\square q(\theta)$
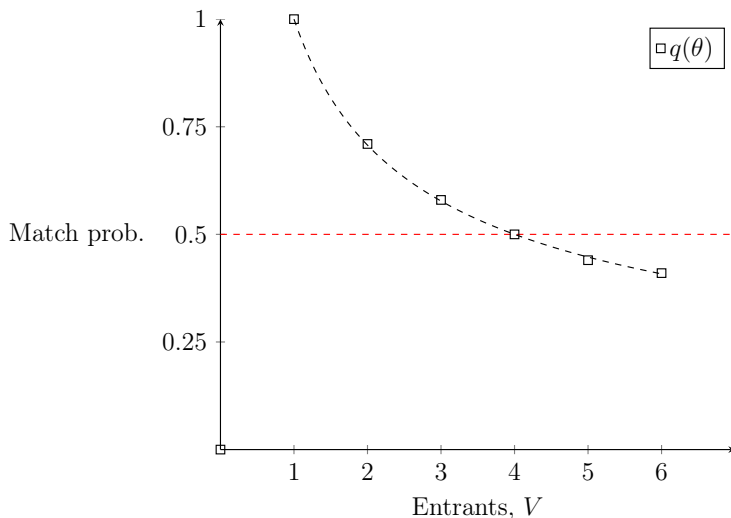
Figure 1.B.2: Matching probabilities for firms

The first component is the expected payoff from production. Each match M produces Y. The second component is the value of (avoided) search costs. A risk-neutral social planner who maximizes social surplus should therefore increase vacancy creation up to the point at which stochastic rationing of workers drives down the expected productivity of a match to equal the vacancy creation cost X. We see this from the first-order condition:

$$\frac{d\Sigma}{dv} = \frac{dM}{dv}Y - X$$
$$\Rightarrow \frac{dM}{dv}Y = X$$

Efficiency requires the number of vacancies to satisfy $\frac{dM}{dv} = \frac{X}{Y}$.

Given our parametrization of the matching function:

$$v^\Sigma = \left(\frac{95}{Y/2}\right)^2.$$

When $Y = 400$, efficiency requires $v^\Sigma = 4.4$. When $Y = 300$, efficiency requires $v^\Sigma = 2.5$. Hence, if the planner were able to perfectly coordinate entry, the planner would choose 4 entrants and 2 entrants respectively. If the planner must choose a symmetric entry proba-

bility, the problem is somewhat more complex. In such a case, the planner chooses a p^Σ to maximize

$$\mathbb{E}\left[\Sigma\right] = \sum_{v=0}^{F} \binom{F}{v} \left(p^\Sigma\right)^v \left(1 - p^\Sigma\right)^{F-v} \left[M(U, v)Y - (F - v)X\right].$$

This probability can be computed from the first order condition

$$\sum_{v=0}^{F} \binom{F}{v} \left(p^\Sigma\right)^v \left(1 - p^\Sigma\right)^{F-v} \left(\frac{v - p^\Sigma F}{p^\Sigma(1 - p^\Sigma)}\right) \left[M(U, v)Y - (F - v)X\right] = 0.$$

Closely related to the expected social surplus is the expected total profits of firms. If firms could collude, they would like to restrict entry relative to the non-cooperative Nash equilibrium. The reason is that entry imposes a congestion externality on other firms. Given our parametrization of the matching function, the level of vacancy creation that maximizes joint firm payoffs is reduced by a factor of η^2 relative to the socially efficient level of entry. For $\eta = 0.5$ this implies that if the efficiency level of entry is 4 firms, then the joint payoffs for firms are maximized if only a single firm (out of 6) enters.

1.C   Results

Figure 1.C.1 presents the analogous results to those shown in figure 1.6 for vacancy creation but at the individual block level. Each block consisted of 2-3 markets so the data are already somewhat aggregated at the level of a block observation. Although there is some heterogeneity at the block level, the overall level of vacancy creation is quite similar to the results at the higher level of aggregation. Notably, there is no perceivable time trend in most of the blocks.

Figure 1.C.1: Average Vacancy Creation $M^h - N^l$
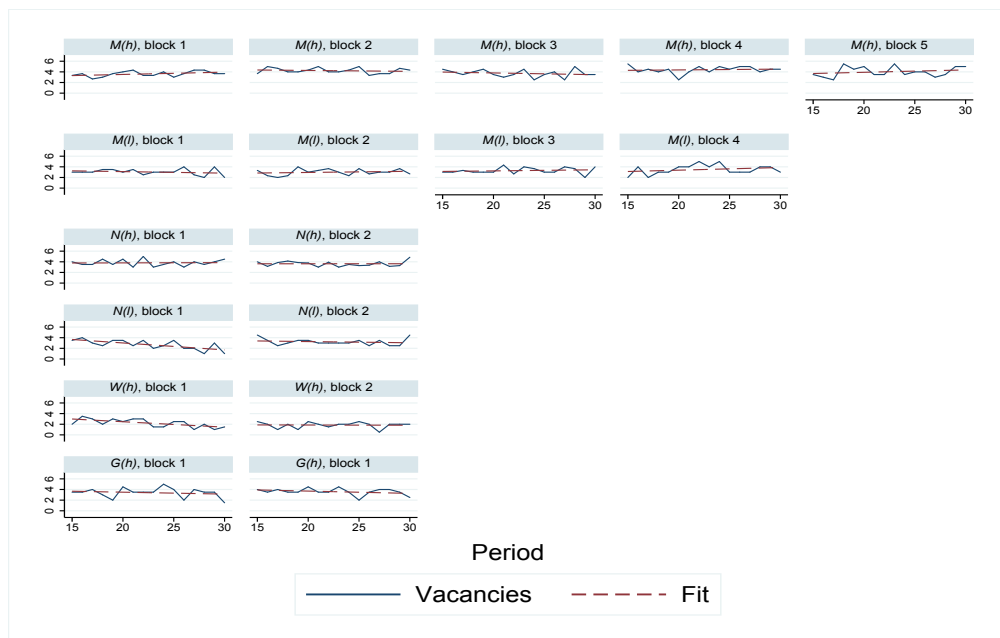
Figure 1.C.2 presents the entry in the second half of periods for the M and N treatments. For this portion of the data, only N^l exhibits any trend in the degree of entry. We conclude that it takes some time for behavior to converge to a stable level but that once it reaches that level it tends to stay there. This is perhaps surprising given that there continues to be considerable randomization at the individual level throughout the duration of the experiment.

Figure 1.C.2: Average Vacancy Creation M^h $-N^l$
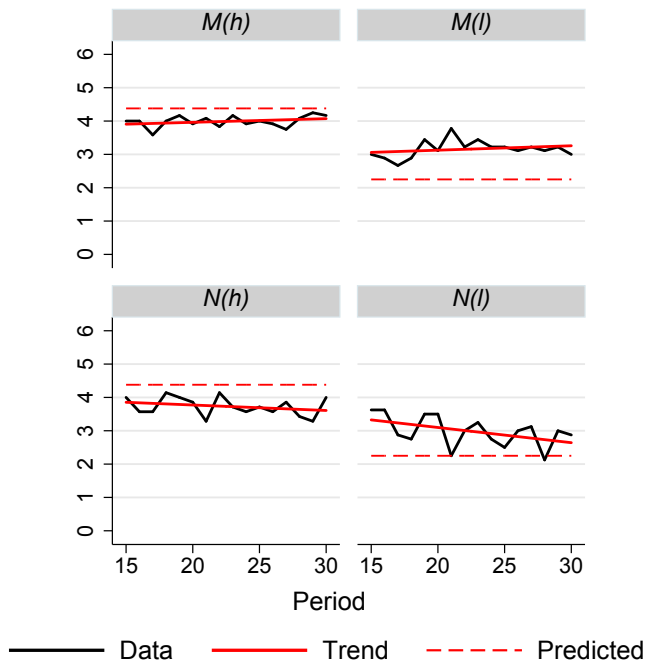
Result 1

Table 1.C.1 presents the block level averages for number of vacancies. These are computed as the average of the market level observations for each block for all 30 periods. These are the data used in the WSR tests presented in the main text.

Table 1.C.1: Block Level Averages

Blocks		$Y = 400$						$Y = 300$			
		1	2	3	4	5		1	2	3	4
Matching	M^h	3.64	3.67	4.07	4.27	4.33	M^l	2.98	3.03	3.31	3.60
Neutral	N^h	3.54	3.60	3.98	4.13		N^l	2.63	2.73	3.13	3.27
Worker	W^h	1.80	2.23								
Group	G^h	3.40	3.57								

Result 3

In 1.C.3 we show the distributions for entry for treatments M^h and M^l together (just matching) and treatments N^h and N^l together (neutral bargaining). In both case we see much more overlap than is expected from the Nash randomization probabilities.

Figure 1.C.3: Distribution of Entrants N and M treatments
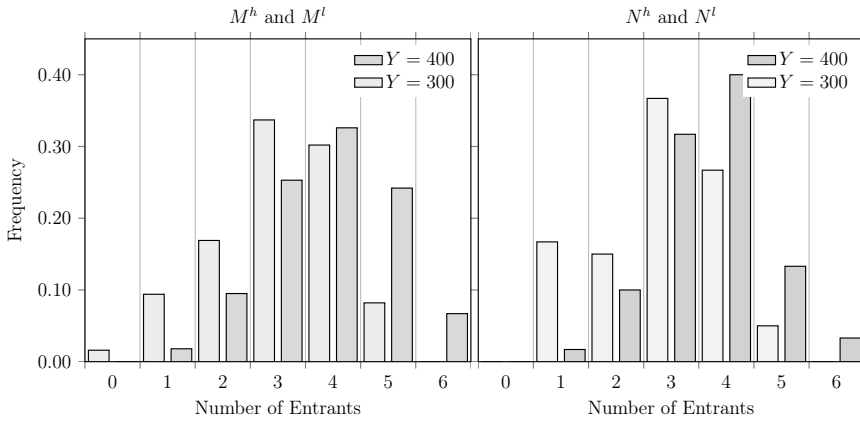
Result 4

Figure 1.C.4 shows the bargaining offers over time aggregated by market (group) level. The offers are stable near 100 during most periods of the experiment with, perhaps, a slight downward trend overall:

Figure 1.C.4: Individuals Offers, N^h , N^l , and W^h
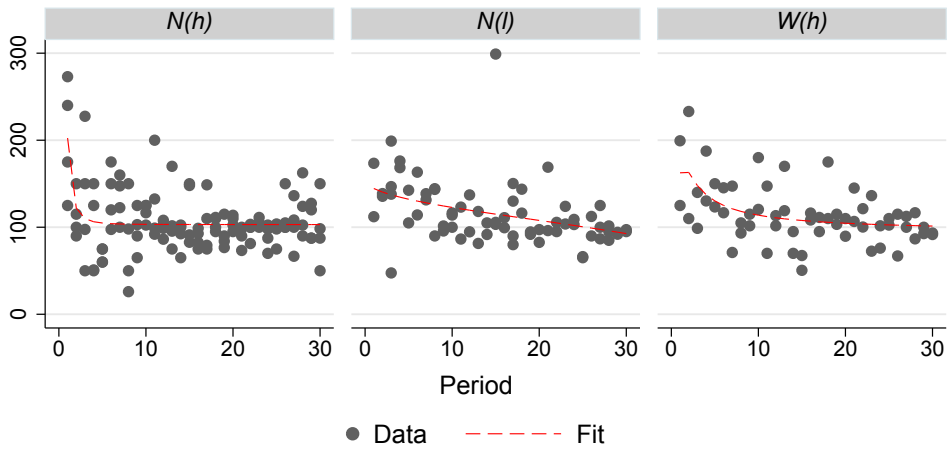
In figure 1.C.5, we show histograms showing all of the rejected offers in the treatments with ultimatum bargaining. The remarkable finding shown in this plot is that no offers in excess of 100 were rejected (in the second half of the data). This gives intuition as to why the offers cluster near 100. Below this level (that is, to the left of 100), the offers are rejected relatively often but offers above this level were never rejected.

Figure 1.C.5: Rejected Offers, N^h , N^l , and W^h
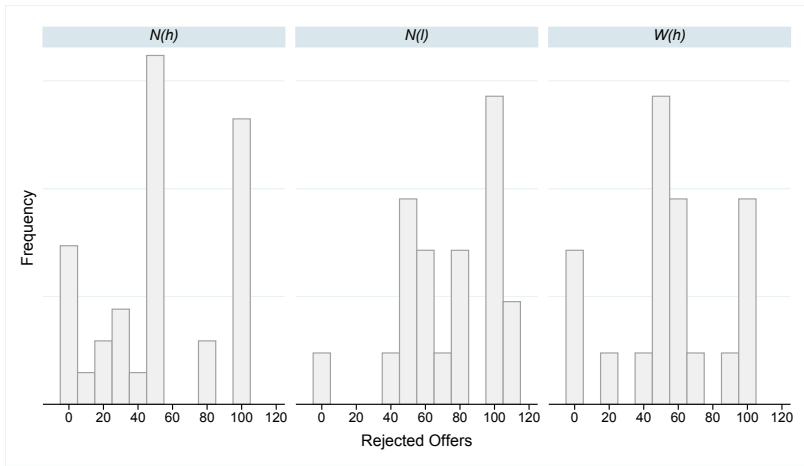
Result 5

A revealing analysis is based on a logistic regression in which we estimate the probability that a given offer will be accepted. These estimated acceptance probabilities are shown in figure 1.C.6. As is apparent, the probabilities fall steeply below the modal offer. Moreover, using these probabilities, we can compute expected payoffs from the offers: That is, the surplus net of the offer, weighted by the acceptance probability. This is shown in figure 1.9 (in the main text). We see that the expected payoffs are maximized in the vicinity of the modal offer and, indeed, fall steeply on either side of this offer. This seems to imply consistency of beliefs about sharing among players.

Figure 1.C.6: Probability of Acceptance
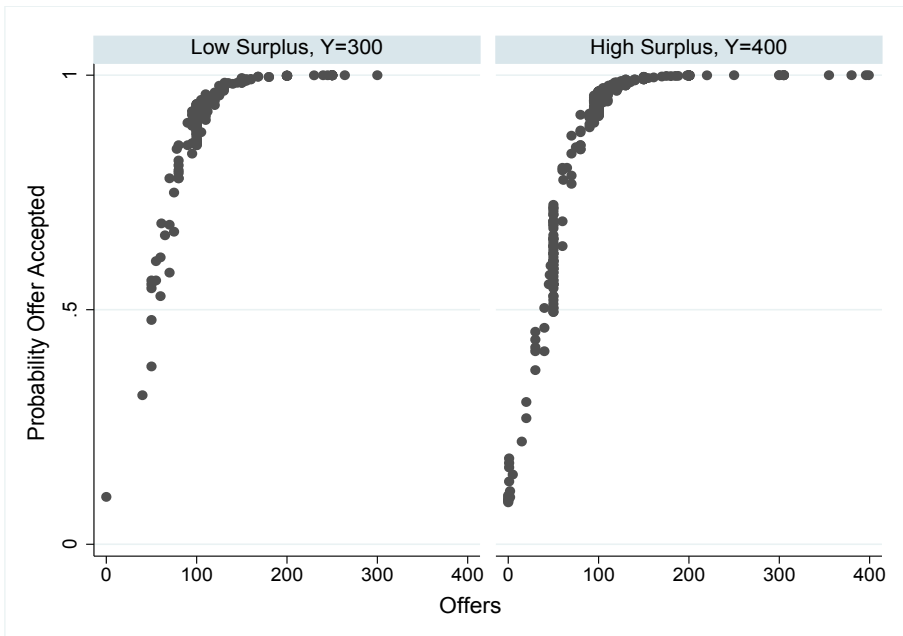
Result 9

In figure 1.C.7, we present the empirical distributions showing the proportion of entrants that enter a given number of times or less. That is, if we randomly sampled an individual, these graphs show how likely it is that they enter x times or less. The main result here is that the high productivity treatments stochastically dominate the low productivity treatments.

Figure 1.C.7: Empirical Distributions of the Number of Enter Decisions
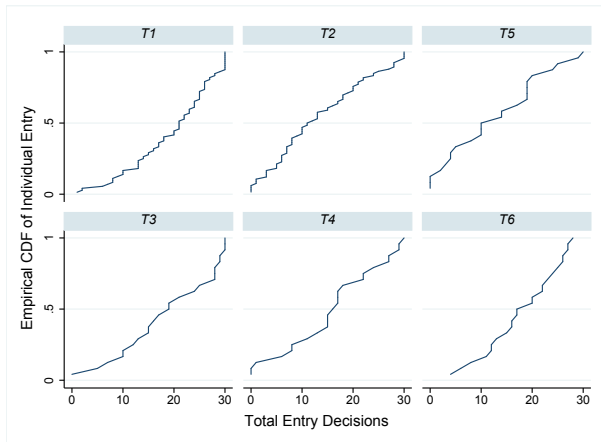
The distribution of streaks—that is, consecutive enter decisions—is shown in figure 1.C.8. Compared with the theoretical (geometric) distribution associated with symmetric randomization, there are many streaks in which the same individual enters many times in a row. This suggests that some individuals persistently enter.

Figure 1.C.8: Distribution of entry streaks (number of consecutive decision to enter
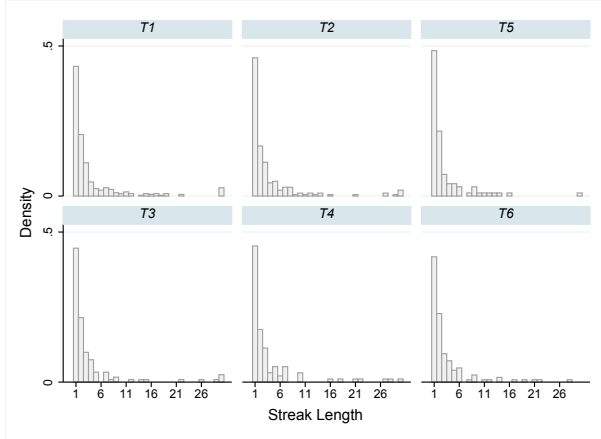
Additional Analyses

Inequity Aversion

We begin by considering inequity preferences as proposed by Fehr and Schmidt (1999) with unfavorable inequity parameter α and favorable inequity parameter β:

$$U_i(x_i, x_j) = x_i - \alpha \max\{0, x_j - x_i\} - \beta \max\{0, x_i - x_j\}. \qquad (1.10)$$

Inequity Preferences with Match Reference

We first consider the case when only payoffs to a player and their match in bilateral bargaining enter the utility function. We will refer to this as inequity preference with match reference. Such preferences have been extensively explored and predict offers of a constant share of the surplus, independent of the size of the surplus. This is inconsistent with the empirical results which predict a negative relationship between the size of the surplus and the offer. We show this below.

Let p denote the proposer state and r denote the responder state. We denote material payoffs by x with the relevant subscript. In the bargaining stage, it does not matter whether a player is a firm or worker type. By subgame perfection, we first solve for the minimal acceptable offer of the responder and then back out the remaining outcomes. Utility for the

74

responder is

$$U_r(x_p, x_r) = x_r - \alpha(x_p - x_r)$$
$$U_r(x_p, x_r) = x_r - \alpha(S - 2x_r).$$

and the the utility for the proposer is

$$U_p(x_p, x_r) = x_p - \beta(x_p - x_r)$$
$$U_p(x_p, x_r) = x_p - \beta(2x_p - S)$$

where we use the fact that when bargaining over a fixed surplus $x_r = S - x_p$. Given this specification, the responder will accept any offers that yield positive utility because rejection gives zero utility to both players. That is, offers satisfying $x_r \geq \frac{\alpha}{1+2\alpha}S$. Next, given the value of β, the proposer will either propose an offer of half the surplus or the lowest offer acceptable to the responder, depending on whether $\beta \lessgtr \frac{1}{2}$. Only the case with $\beta < 1/2$ is empirically relevant since proposers nearly always make offers that are much smaller than $1/2$ of the surplus. As an example, consider that $\beta < 1/2$ and $\alpha = 1$. Then the proposer will offer $1/3$ of the surplus and this is the minimum that the responder will accept. We conclude that shares offered in bilateral bargaining are independent of the size of the surplus. To explain our observations, α (or, more generally, the distribution of α) would need to be substantially different between treatments.

Inequity Preferences with Market Reference

Next we consider the case when payoffs to all other agents in the game also enter the utility function, that is a specification of equation 1.10 for all agents in the game. We refer to this as inequity preference with market reference. We assume initially that agents share the same values for α and β and that this is common knowledge. Again, we denote material payoffs by x with a subscript. x_o is a constant and $x_\varnothing = 0$. We denote the *states* that players can be in by o, p, r, and $\varnothing$. r and p denote the same types as above, o denotes firms that take the outside option, and $\varnothing$ denotes players in the bargaining stage who do not get matched. There is, in addition, a fifth type of agent that we denote by X. These are players who got matched but for whom bargaining ended with a rejection. This type of agent determines which offers are accepted, but should not be observed in the (common knowledge) equilibrium.

We begin by writing down the utilities for players in the bargaining stage of the game. These are *ex post* utilities that depend on final realizations in the game. When v players enter and M matches are *realized*, then there are $N - v$ firm players who do not enter, M proposers and M responders, and $U + v - 2M$ agents in the matching market who end up unmatched.[56] Next, we show that for plausible values of β proposers would like to maximize their own payoffs. Thereafter, we solve for the minimal acceptable offer for the responder. Finally, we back out v.

We show that proposers maximize their utility by maximizing their payoffs. We *assume* that proposers offer responders the minimal amount that will not be rejected. Then we derive a condition that ensures that this is an optimal decision. Moreover, given complete information, all proposers will offer this amount. This justifies the assumption.

The proposer's utility is:

$$U_p(x_p, \mathbf{x_r}, \mathbf{x_o}, \mathbf{x_\varnothing}) = x_p + \frac{\beta}{N+U-1}M\left(x_r - x_p\right) + \frac{\beta}{N+U-1}\left[(v-M)+(U-M)\right](x_\varnothing - x_p)$$
$$+ \frac{\beta}{N+U-1}(N-v)(x_o - x_p)$$

The second term results from the M matches for which the proposer earns more than the responder, the third term results from a comparison of earnings with the $v + N - 2M$ agents who did not get matched, and the last term results from the $N - v$ firm players who chose to stay out of the market. Simplifying further:

$$U_p(x_p, \mathbf{x_r}, \mathbf{x_o}, \mathbf{x_\varnothing}) = x_p + \frac{\beta}{N+U-1}M\left(S - 2x_p\right) + \frac{\beta}{N+U-1}\left[(v-M)+(U-M)\right](x_\varnothing - x_p)$$
$$+ \frac{\beta}{N+U-1}(N-v)\left(x_o - x_p\right)$$
$$= x_p - \frac{\beta}{N+U-1}(U+N)x_p + \frac{\beta}{N+U-1}(MS + (N-v)x_o)$$

From the last line, we see that when $\beta < \frac{N+U-1}{N+U}$ the proposer's utility increases with x_p since own payoff increases at a faster rate than the penalty from favorable inequity.[57] It is

---

[56]We refer to *realized* matches because v does not map deterministically to a number of matches: There is a probability $q = (M(v,U) - \lfloor M(V,U) \rfloor)$ that $M = \lceil M(V,U) \rceil$ and a complimentary probability that $M = \lfloor M(V,U) \rfloor$.

[57]If proposers happen to offer different amounts, this condition will be more complicated and could be either more or less strict. However, a similar intuition will hold.

very likely that this requirement is satisfied [58] If so, the proposer will offer a responder the minimal offer that the responder will accept. Moreover, it must be so that $x_p \geq x_r$.[59] In the empirical results, the relevant ranking is $x_p > x_r > x_o > x_\varnothing$. However, theoretically, nothing guarantees that $x_r > x_o$.

Next we compare the utility for a responder who accepts an offer with the utility for a responder who rejects an offer. On the equilibrium path, proposers will make offers that responders will accept. This means that a decision to reject leads to $M-1$ matches in which the offer is accepted, given that the other responders accept the equilibrium offer. For the responder, write the utility for the case when $x_r > x_o$:

$$
\begin{aligned}
U_r(x_r, \mathbf{x_p}, \mathbf{x_o}, \mathbf{x_\varnothing}) = x_r &- \frac{\alpha}{N+U-1}M(x_p - x_r) + \frac{\beta}{N+U-1}\left[v + U - 2M\right](x_\varnothing - x_r) \\
&+ \frac{\beta}{N+U-1}(N-v)(x_o - x_r) \\
= x_r &+ \left[\frac{\alpha}{N+U-1}2M - \frac{\beta}{N+U-1}(N+U-2M)\right]x_r \\
&- \left[\frac{\alpha}{N+U-1}MS - \frac{\beta}{N+U-1}(N-v)x_o\right].
\end{aligned}
$$

and for the case when $x_r < x_o$:

$$
\begin{aligned}
U_r(x_r, \mathbf{x_p}, \mathbf{x_o}, \mathbf{x_\varnothing}) = x_r &- \frac{\alpha}{N+U-1}M(x_p - x_r) - \frac{\alpha}{N+U-1}(N-v)(x_o - x_r) \\
&+ \frac{\beta}{N+U-1}(v + U - 2M)(x_\varnothing - x_r) \\
= x_r &+ \left[\frac{\alpha}{N+U-1}(2M + N - v) - \frac{\beta}{N+U-1}(v + U - 2M)\right]x_r \\
&- \frac{\alpha}{N+U-1}(MS + (N-v)x_o)
\end{aligned}
$$

---

[58]In their original study Fehr and Schmidt (1999) suggest an upperbound on β of 0.6 (See Fehr and Schmidt 1999, table III).

[59]It is a dominant strategy to accept offers of half the surplus (that is $x_p = x_r$) whenever $\beta < \frac{N+U-1}{N+U-2M}$, a condition that is slack (to solve for this condition, set $x_p = x_r$, and solve for a condition on β such that U_r is positive. The condition derived here is for the case when $x_r > x_o$). Moreover, a proposer would like to maximize their payoffs which are increasing in x_p. Hence, proposers will only make offers equal to less than half the surplus. An analogous computation can be carried out for the case when $x_r < x_o$.

Utility for a responder who rejects an offer such that payoffs are $x_p = x_r = 0$ is:

$$U_X(x_X, \mathbf{x_r}, \mathbf{x_p}, \mathbf{x_o}, \mathbf{x_\varnothing}) = x_x - \frac{\alpha}{N+U-1}(M-1)(x_p - x_x) - \frac{\alpha}{N+U-1}(M-1)(x_r - x_x)$$
$$- \frac{\alpha}{N+U-1}\left[v + U - 2M\right](x_\varnothing - x_x) - \frac{\alpha}{N+U-1}(N-v)(x_o - x_x)$$
$$= -\frac{\alpha}{N+U-1}\left[(M-1)S + (N-v)x_o\right]$$

Then we solve for the minimum x_r such that $U_r \geq U_x$. In the case when $x_r > x_o$,

$$x_r > \frac{\frac{\alpha}{N+U-1}S - \frac{\alpha+\beta}{N+U-1}(N-v)x_o}{1 + \frac{\alpha}{N+U-1}2M - \frac{\beta}{N+U-1}(N+U-2M)} \tag{1.11}$$

and in the case when $x_r < x_o$:

$$x_r > \frac{\frac{\alpha}{N+U-1}S}{1 + \frac{\alpha}{N+U-1}(2M+N-v) - \frac{\beta}{N+U-1}(v+U-2M)} \tag{1.12}$$

Let us denote these minimal acceptable offers by $\underline{x_r}$.

To assess whether $x_r \lessgtr x_o$ we compare $\underline{x}$ with x_o and derive conditions on α and β. For $\underline{x_r} > x_o$ we consider equation 1.11 relative to x_o:

$$\frac{\alpha S - (\alpha+\beta)(N-v)x_o}{N+U-1+\alpha 2M - \beta(N+U-2M)} > x_o$$
$$\alpha S - \alpha(N-v)x_o - \beta(N-v)x_o > (N+U-1)x_o + \alpha 2Mx_o - \beta(N+U-2M)x_o$$
$$\alpha\left[S/x_o - (N-v) - 2M\right] > (N+U-1) - \beta\left[v+U-2M\right]$$

Conclude that $\underline{x_r} > x_o$ only for the case when

$$\beta > \frac{(N+U-1) - \alpha\left[S/x_o - (N-v) - 2M\right]}{v+U-2M} \tag{1.13}$$

The numerator on the right hand side will be larger than the denominator. For our parameter values, $S/x_o < (N-v)-2M$. The numerator will therefore be *larger* than $N+U-1$ while the denominator is smaller.[60] But we concluded above that $\beta < \frac{N+U-1}{N+U}$. So inequity preferences predict the opposite ranking $(x_r < x_o)$ of what is observed empirically.

---

[60]The numerator is a quantity larger than the total number of players while the denominator is equal to the number of unmatched players.

Similarly, beginning from equation 1.12 for the case when $x_r < x_o$, we derive the complimentary condition:

$$\frac{\alpha S}{N + U - 1 + \alpha(2M + N - v) - \beta(v + U - 2M)} < x_o$$

$$\alpha \frac{S}{x_o} < N + U - 1 + \alpha(2M + N - v) - \beta(v + U - 2M)$$

$$\alpha \left[\frac{S}{x_o} - (N - v) - 2M \right] - (N + U - 1) < -\beta(v + U - 2M)$$

And conclude again that $x_r < x_o$ when:

$$\beta < \frac{(N + U - 1) - \alpha \left[S/x_o - (N - v) - 2M \right]}{v + U - 2M} \tag{1.14}$$

which, as argued above, will always be satisfied for this study.

This result predicts the incorrect ordering relative to our findings. In addition, not only is the ranking wrong, but the predicted difference between x_o and x_r is substantial. for reasonable values of α and β, offers are around 30 in the high surplus case and somewhat larger than 20 for the low surplus case. This is strongly counterfactual. Given these small shares, both less than a share of 0.1, the pattern across treatments is also unclear.

The details of this computation are included below. Note that we must consider the *expectation* because up to two different values of M can be realized for each level of entry v. To solve for the optimal entry v^*, we solve

$$U_o(v^*) = \mu(\theta) \left[\frac{1}{2} \mathbb{E}U_p(v^*) + \frac{1}{2} \mathbb{E}U_r(v^*) \right] + (1 - \mu(\theta))\mathbb{E}U_\varnothing \tag{1.15}$$

$$\text{s.t } U_r(v^*) \geq U_X(v^*). \tag{1.16}$$

The utility for players who take the outside option is

$$U_o(x_o, \mathbf{x_r}, \mathbf{x_p}, \mathbf{x_o}, \mathbf{x_\varnothing}) = x_o - \frac{\alpha}{N + U - 1}M(x_p - x_o) + \frac{\beta}{N + U - 1}M(x_r - x_o) + \frac{\beta}{N + U - 1}(v + U - 2M)(x_\varnothing - x_o)$$

$$= x_o \left(1 + \frac{\alpha - \beta}{N + U - 1}M - \frac{\beta}{N + U - 1}(v + U - 2M) \right) - \frac{\alpha}{N + U - 1}Mx_p + \frac{\beta}{N + U - 1}Mx_r$$

given that $x_r < x_o$.[61] Using our result on the offer $\underline{x_r}$, we can back out values of v^*.

Given that the theoretically predicted offers to responders are much lower than observed in the data, it may not be surprising that theoretically predicted entry also does not conform to the data. Moreover, the qualitative pattern of entry across treatments is ambiguous: It need not be the case that higher entry is predicted in the treatment with higher surplus. Finally, we note that the degree of entry is quite sensitive to changes in α and β.

ERC Model

Next consider *equity, reciprocity, and competition* (ERC) preferences (Bolton and Ockenfels 2000). ERC preferences specify a motivation function $U_i(y_i, \sigma_i)$ where y_i is the payoff to player i, σ is the share of the surplus i receives, and c is the total size of the surplus. The motivation function is continuous and differentiable in both arguments, increasing and concave in own payoff, and—for a given level of y_i—maximized by an equal division of c. De Bruyn and Bolton (2008) suggest a formulation similar to

$$
U_i = \begin{cases} c\left(\sigma_i - \frac{b}{2}\left(\sigma_i - \frac{1}{2}\right)^2\right) & \text{if } \sigma \leq 1/2 \\ c\sigma_i & \text{if } \sigma > 1/2 \end{cases}
$$

Narrow ERC Preferences

Specifying this for narrow preferences including just the proposer and the responder, the responder gets share $\sigma_r = (c - \sigma_p c)/c$, and has preferences

$$
U_i = c\left(\sigma_r - \frac{b}{2}\left(\sigma_i - \frac{1}{2}\right)^2\right) \tag{1.17}
$$

---

[61]Although only the case where $x_r < x_o$ is empirically relevant, the opposite case gives utility

$$
\begin{aligned}
U_o(x_o, \mathbf{x_r}, \mathbf{x_p}, \mathbf{x_o}, \mathbf{x_\varnothing}) =& x_o - \frac{\alpha}{N+U-1}M(x_p - x_o) - \frac{\alpha}{N+U-1}M(x_r - x_o) + \frac{\beta}{N+U-1}(N+U-2M)(x_\varnothing - x_o) \\
=& x_o\left(1 + \frac{\alpha}{N+U-1}2M - \frac{\beta}{N+U-1}(v+U-2M)\right) - \frac{\alpha}{N+U-1}MS
\end{aligned}
$$

. And the utility for players that do not get matched is

$$
\begin{aligned}
U_\varnothing(x_\varnothing, \mathbf{x_p}, \mathbf{x_r}, \mathbf{x_o}) =& -\frac{\alpha}{N+U-1}M(x_p - x_\varnothing) - \frac{\alpha}{N+U-1}M(x_r - x_\varnothing) - \frac{\alpha}{N+U-1}(N-\nu)(x_o - x_\varnothing) \\
=& -\frac{\alpha}{N+U-1}(MS + (N-v)x_o).
\end{aligned}
$$

The responder will therefore accept offers satisfying $\sigma_r > \frac{b}{2}\left(\sigma_i - \frac{1}{2}\right)^2$. Minimal offers therefore lie to the left of $1/2$ and, moreover, be described by a *share* of the surplus independent of the size of the pie. This can not be reconciled to our results.

Broad ERC Preferences

To attempt to generate sensitivity to the size of the pie, we consider broad preferences for which c is the *sum of all payoffs* in the market and σ_i individual i's share of this total. Moreover, we also adjust $1/2$ to $1/(N+U)$ (that is, equal division among all agents in the game is the social fairness benchmark): $U_i = c\left(\sigma_i - \frac{b}{2}\left(\sigma_i - \frac{1}{N+U}\right)^2\right)$. The share for player i will then be relative to total payoffs

$$\sigma_i = \frac{x_i}{MS + (N-v)x_o} \tag{1.18}$$

where the denominator is the total payoff from matches plus the value of payoffs to agents who do not enter the market. Although again the preference structure implies a constant *share*, this is of *overall* payoffs. Offers x_r will therefore vary depending on S (and v).

However, given this specification of the preferences, minimal acceptable offer for responders will be below $1/(N+U)$ (imagine the intersection between the linear component of the preferences and the quadratic portion: Minimal offers must lie to the left of $1/(N+U)$). Indeed, offers much be substantially below the social fairness level if b is specified to yield reasonable values of x_p. This specification can therefore not fit the data. More complicated ERC preferences can be engineered, but this requires additional parameters or assumptions of functional form. Another possibility is that the preferences include a stochastic element.

As an exercise, consider an empirically relevant value of σ_r in the treatment with $S = 400$. Since we expect 4 firms to enter, there will be 2 matches and 2 firms that take the outside option. Given offers of about 100 in this treatment, $\sigma_{S=400} \sim 0.1$. Similar computations for the low surplus treatment translate to $\sigma_{S=300} \sim 0.125$. Using a value of $\sigma = 0.11$ we can solve for the offers x_r using equation 1.18: In this case, offers of share size $1/(N+U)$ are *never rejected* and, perhaps, agents do not offer less than this in order to avoid rejections.

Empirically consistent minimal acceptable offers		
v	$S = 300$	$S = 400$
1	85.3	96.3
	(0.284)	(0.241)
2	88.5	104.0
	(0.295)	(0.260)
3	88.5	107.6
	(0.295)	(0.269)
4	86.9	108.9
	(0.290)	(0.272)
5	84.2	108.8
	(0.281)	(0.272)
6	80.8	107.8
	(0.269)	(0.269)

Figure 1.C.9: Minimal offers by surplus Y

Chapter 2.

Discrimination in Small Markets with Directed Search: Part I Theory

Abstract
To understand the consequences of discrimination for labor market outcomes, it is critical to account for how discrimination interacts with labor market frictions. In an important contribution in this regard, Lang, Manove, and Dickens (2005) show how discriminatory hiring leads to labor market segregation in a large (continuum) market environment with posted wages. The essential reason for this outcome is that discriminated workers will apply to low wage jobs to avoid competition from workers of the preferred type. In this paper, we investigate the theoretical properties of a small market analog of the model of Lang, Manove, and Dickens (2005). The model we consider is populated by a finite set of homogeneous firms and two types of unemployed worker, white and black. Discrimination takes the form of a preference at the hiring stage and white workers are always selected ahead of black workers if they compete for the same vacancy. Although the setting we analyze reproduces the incentives from Lang, Manove, and Dickens (2005), we find that a pure strategy segregation outcome is not a general finding in the small market.

Author: Knut-Eric N. Joslin [1]

2.1 Introduction

The black-white wage gap is a persistent feature of the American labor market. In fact, the wage gap appears to have widened in recent years (Wilson and Rodgers III 2016). While some of the difference in wages is explicable in terms of differences in productivity, evidence suggests that this is not a complete explanation (See Lang, Lehmann, et al. 2012, for an introduction to the extensive literature on the black-white wage gap). Efforts to understand how and why a wage gap can persist are therefore worthwhile.

A feature of labor markets that has been neglected in most models of discrimination is the presence of labor market frictions. Recent empirical and theoretical studies suggest that these frictions may be important for explaining the operation of discrimination and persistence of discriminatory outcomes in labor market settings (Lang and Manove 2011; Lang, Lehmann, et al. 2012; Holden and Rosén 2014).

One of the most striking models of discrimination is the directed search model of Lang, Manove, and Dickens (2005). In this model, if a white and black worker apply for the same wage, the white worker is always hired. This creates a powerful incentive for black workers to avoid applying to the same wages as white workers. In response, some firms can offer wages that are so low that only black workers apply. The model illustrates how discrimination creates a tendency toward segregation. The equilibrium in the model is characterized by a high and a low wage such that white workers only apply to the high wage and black workers only apply to the low wage. Crucially, the low wage is depressed to a level exactly equal to the white expected income.

This project reexamines the model of Lang, Manove, and Dickens (2005) in a small market setting. We represent discrimination in the same fashion as they do and the agents in our model face similar incentives. However, in the small market setting we examine, firms have market power and wage-setting is strategic. These features are absent from the original model. The motivation for this study is to determine whether the segregation result is robust to these additional effects. Many real world labor markets have a small scale and it is useful to understand the general applicability of the segregation result.

Our main finding is that a segregated equilibrium need not exist. Although the candidate for a segregated equilibrium associated with Lang, Manove, and Dickens (2005) characterizes a local profit maximum for firms, firms typically have an incentive to deviate either to a much higher or much lower wage. Notably, this opportunity disappears as the market becomes large. Our negative finding is consistent with the absence of an existence theorem for pure

strategy equilibria when payoff functions fail to be quasi-concave.

Formally, the small market model that we investigate builds on the literature on directed search with a finite number of agents (Burdett et al. 2001; Galenianos and Kircher 2012; Peters 1984). From a theoretical standpoint, what is new is the presence of ranked types. This complicates the characterization of the worker application subgame. To handle the presence of multiple types, we extend the approach developed by Galenianos and Kircher (2012). Although tedious to compute, we derive analytic expressions for how worker behavior responds to changes in wage announcements. This enables us to assess analytically certain candidate equilibria in pure strategies.

The ability to develop general results is, however, limited by the complex relationship between the set of wages offered and the profits available from deviations. Our main results thus rely on case-by-case computations. For a large number of cases, we compute the optimal deviation from the segregated equilibrium candidate and typically—though not in all cases—find that a profitable deviation exists. These deviations are profitable because workers respond strongly to changes in the available wages.

In the next section, we summarize the literature with direct relevance to this project. This provides an orientation to the analysis of directed search models in the small market setting. It also serves to clarify the differences between this study and previous work. With this background, we present the model in the third section. Thereafter, in the fourth section, we turn to the results. We show first that there is no equilibrium in which all firms set the same wage. Second, we demonstrate that a segregated equilibrium does not exist in the general case. And third, we discuss what other types of equilibria are plausible. In the fifth section, we briefly conclude.

2.2 Background and Theory

At a conceptual level, our study draws inspiration from the discrimination model of Lang, Manove, and Dickens (2005). We review their key results below. The fundamental question is whether the insights from the large market situation carry over to a strategic setting.

In terms of theoretical analysis, the most closely related studies are Burdett et al. (2001) and Galenianos and Kircher (2012). These studies develop results for directed search models in finite settings. The analysis of the finite setting is non-trivial because worker behavior will depend on the exact set of wages offered in the market. From a theoretical standpoint, this means that the full announcement of wages needs to be known in order to deduce behavior.

Because the techniques developed in Burdett et al. (2001) and Galenianos and Kircher (2012) form the basis for our analysis, we present them in considerable detail.

2.2.1 Directed Search Models

In directed search models, firms first post wages and, after observing these prices, workers then decide where to apply. There are thus two subgames, the *application subgame* and the *wage-posting subgame*. Analysis of directed search models depends on solving for an equilibrium among workers, taking wages as given, and then solving for an equilibrium in the full wage-posting game. An equilibrium in the whole game is a set of wages such that no firm wishes to deviate given the optimal search strategies of workers.

The standard assumptions, which we also maintain, are that worker application behavior is uncoordinated and that workers apply to the same wages with the same probability. For these assumptions, the matching technology is of the urn-ball type. Workers randomize their applications and an outcome of the application stage of is a set of worker queues at each firm.

The crucial trade-off for agents in this type of environment is that agents trade off the probability of a transaction with the price.[1] For example, a firm in this environment can increase their probability of hiring at the cost of setting a higher wage. Similarly, a worker can increase their probability of getting hired by applying to a low wage firm. Models of directed search thus have the desirable feature that prices play a role in determining the equilibrium.

In all directed search models, the object that dictates worker behavior is the expected market income. Optimal search demands that workers use application strategies such that all wages to which they apply offer this level of income. If workers did not adjust their behavior in this way, it would imply that they could increase their expected income by applying more to some firms and less to others. This is referred to as the *market income* or *market utility property*. An implication of the market income property is that a firm that increases its wage is rewarded with more applicants—at least in a setting with identical workers.[2]

---

[1] Directed search models thus share some similarities with oligopoly models in which firms can trade off price and quantity.

[2] There are some additional details that we omit. For instance, if a firm is setting a wage that is too low, in the sense of being discretely below the market expected income, then a marginal increase in the wage will have no impact on the number of applicants.

2.2.2 Directed Search in a Large Market

The large market case with identical firms and identical workers is the most straightforward to analyze and serves to illustrate the key concepts that dictate behavior. This case was first studied by Montgomery (1991) and has formed the basis for much of the subsequent work on directed and competitive search, including the discrimination model of Lang, Manove, and Dickens (2005).

The critical assumption in the large market setting is that individual market participants can treat the income available to workers as a parameter $\bar{U}$. This means that the wage set by a single firm does not impact workers' expected income.[3] This simplifies the analysis of optimal wage-setting because it pins down the relationship between the wage set by an individual firm and the rate of applications. From a modelling perspective, each firm independently solves a maximization problem taking as given that workers must expect to receive at least $\bar{U}$.

Using notation that we build on in the rest of the paper, let θ_w denote the probability of applying to a firm offering the wage w and let Ω_w denote the probability that at worker is hired at such a firm when all other workers apply with probability θ_w. Moreover, let M and N denote a number of identical firms and workers respectively. The firm thus solves

$$\max_w \pi(w) = (1 - w)\left(1 - (1 - \theta_w)^N\right)$$
$$\text{subject to } w\Omega_w = \bar{U}.$$

This problem has two components. The objective is the profit from hiring weighted by the probability of hiring. The constraint is the requirement that the firm offer the market income. In the case of urn-ball matching, this problem is straightforward to solve. Treating the market tightness $k = \frac{M}{N}$ as a fixed parameter, and allowing the market size to increase then delivers expressions for the wage, the matching probability, and profits (these computations are shown in appendix 2.A).

2.2.3 Directed Search in a Small Market

In the context of a small market, the market income property manifests itself in a richer form. The wage set by an individual firm impacts the expected income available to workers and the application probabilities must be worked out from a set of indifference conditions.

---

[3]This is sometimes referred to as the assumption of a subgame competitive equilibrium.

Individual firms have market power and the application behavior of workers is dictated by the exact vector of wages offered. The optimal wage to set depends on the wages at other firms and there is an interdependency between wages that is absent in the large market setting.

Consider a directed search market with M identical firms and N identical workers. We assume that workers employ symmetric and anonymous application strategies such that all workers employ the same mixed strategy and all firms that set the same wage expect to attract the same number of applicants. Let firms be indexed by i, such that workers apply to firm i that offers wage w_i with probability θ_i and get hired with a probability Ω_i. Then an equilibrium in the worker application subgame is a set of application probabilities such that all firms that receive applications offer the maximum expected income. Concretely, this means that if workers apply to the T highest wages, then the application probabilities are dictated by the indifference conditions

$$w_2 \Omega_2 = w_1 \Omega_1 \tag{2.1}$$

$$\vdots$$

$$w_T \Omega_T = w_1 \Omega_1 \tag{2.2}$$

where we index the firms in such a way that $i = 1$ indicates the highest wage firm and $i = T$ indicates the lowest wage to which workers direct applications. This means that firms that offer $w_i \leq w_T$ do not receive any applicants.

With urn-ball matching, the probability that an additional applicant is hired, conditional on applying to a firm offering wage w_i is

$$\Omega_i = \sum_{k=0}^{N-1} \binom{N-1}{k} \frac{1}{k+1} \theta_i^k (1 - \theta_i)^{N-1-k}. \tag{2.3}$$

This expression results from the fact that there are $\binom{N-1}{k}$ outcomes in which k other applicants apply, each of which occurs with a probability $\theta_i^k (1 - \theta_i)^{N-1-k}$, and the probability of getting hired conditional on k other applicants is $\frac{1}{k+1}$. Manipulating this expression, it is straightforward to show that it can be re-written as

$$\Omega_i = \frac{1 - (1 - \theta_i)^N}{N \theta_i}. \tag{2.4}$$

We show this result in appendix section 2.24. The expression is the ratio of the probability that the firm hires divided by the expected number of applicants.[4] To see this clearly, when all M firms set the same wage w, the symmetric application probability is $\theta_w = \frac{1}{M}$ and the expression arises naturally from the following identity

$$N\Omega_w = M\left(1 - \left(1 - \frac{1}{M}\right)^N\right)$$

which states that the number of workers that get hired must equal the number of firms that hire. Hence,

$$\Omega_i = \frac{1 - \left(1 - \frac{1}{M}\right)^N}{N/M}$$

where the numerator is the probability of hiring and the denominator is the expected number of applicants.

For the case of urn-ball matching, the worker application subgame has a unique (symmetric) equilibrium for any announcement of wages (See Peters 1984). This guarantees that applications vary in a predictable fashion with the wages that are offered. Via implicit differentiation of system 2.1 it is possible to solve for how worker behavior is affected when a firm makes a marginal change in the wage (see appendix 2.A). A difficulty that arises in practice is, however, that the system of indifference equations 2.1 must be solved by numerical methods. As is clear from equation 2.3, this amounts to solving a system of higher order polynomials.

An equilibrium in the full game is an equilibrium in the wage-setting subgame. Specifically, an equilibrium is a set of wages such that no firm can profitably deviate to another wage given the optimal application behavior of workers. Characterizing an equilibrium depends on finding a set of mutually consistent wages such that all wages are optimal given the wage announcement.

The expected profit for a firm offering wage w is the value of hiring weighted by the probability of hiring, where the probability of hiring equals the probability of receiving at least one application (which is the complement of receiving no applications).

$$\pi_w = (1 - w)\left(1 - (1 - \theta_w)^N\right). \tag{2.5}$$

---

[4]Note that the probability of hiring is equal to the probability of getting at least one applicant.

Correspondingly, the derivative of profits is

$$\frac{\partial \pi_w}{\partial w} = -\left(1 - (1-\theta)^N\right) + (1-w)N\left(1-\theta_w\right)^{N-1}\frac{\partial \theta_w}{\partial w}. \tag{2.6}$$

The first term is the loss in profits associated with hiring at a higher wage while the second term is the increase in the number of applicants. Because we are able to solve for $\frac{\partial \theta_w}{\partial w}$ from differentiation of system 2.1, it is possible to show that profits for firm i are quasi-concave in the wage w_i. A fixed point argument then establishes that an equilibrium exists (See Galenianos and Kircher 2012, who derive the result in setting with more general matching functions and heterogeneity on the firm side.).

Concretely, the optimal wage will balance the two effects captured in expression 2.6. After plugging in for $\frac{\partial \theta_w}{\partial w}$ it is possible to show that for M identical firms and N identical workers, an equilibrium is for all firms to set the wage

$$w(M, N) = \frac{1}{1 - \frac{M-1}{M}\left(1 - \frac{\Omega}{\left(1-\frac{1}{M}\right)^N}\right)}. \tag{2.7}$$

By algebraic manipulation it is possible to reconcile this expression to that derived in Burdett et al. (2001). Table 2.A.1 in the appendix shows the level of wage associated with different levels of M and N. Table 2.A.2 shows the associated profits.

2.2.4 Directed Search in a Large Market with Discrimination

The main result in Lang, Manove, and Dickens (2005) is that there is a segregated equilibrium. This equilibrium is characterized by a high wage that attracts exclusively white applicants, a low wage that attracts exclusively black applicants, and an equal profit condition that ensures that firms are indifferent between the two wages. In addition, the low wage is exactly equal to the expected income of white workers. This means that the presence of discrimination presses the black wage down exactly to the level at which white workers cease to apply. The reason for this result is that the arrival of white workers is so costly for blacks that it can never be optimal to set a wage that attracts both worker types. Whenever a firm is attracting both worker types, a reduction in the wage is profitable because the loss of white workers is more than compensated by an increase in black applications.

The consequences of the segregation are noteworthy. Because the black wage is equal to the white expected income, black workers earn substantially less than whites. Perhaps even

more remarkable, firms earn higher profits than in the absence of discrimination. However, the result is predicated on the assumption that individual firms can not impact the expected income available to workers. It is this assumption that we relax in this study.

2.3 Model

We begin with a formal description of the game that we investigate. We argue that for any set of wages, there is an equilibrium in the worker application subgame. We characterize worker application behavior analytically and derive expressions for how workers respond to changes in the set of offered wages. After that, we examine the firms' problem and describe how profits depend on worker application behavior. Finally, we define an equilibrium in the full game.

2.3.1 The Environment

The labor market we consider is populated by a set of M identical firms $m \in \{1, 2, \ldots, M\}$, a set of N identical "white" workers $n \in \{1, 2, \ldots, N\}$, and a set of B identical "black" workers $b \in \{1, 2, \ldots, B\}$. Each firm has a production opportunity that generates a surplus of 1 if they employ a worker.

In the first stage of the game, each firm m posts a wage w_m that they are committed to paying if a worker is hired. This wage divides the production surplus. The wage can not be conditioned on worker type.

In the second stage, workers make a single application. A strategy for a worker specifies for any wage announcement $\mathbf{w} = (w_1, w_2, \ldots, w_M)$ a vector of application probabilities. We let $\theta_m^n(\mathbf{w})$ denote the probability that a white worker n applies to firm m when the wages $\mathbf{w}$ are offered and $\gamma_m^b(\mathbf{w})$ denote the same for a black worker b. Workers may only send a single application so that $\sum_{m=1}^{M} \theta_m^n = 1$ and $\sum_{m=1}^{M} \gamma_m^b = 1$ for all n and b. An outcome of the worker subgame is thus a set of queues of white and black workers at each firm. It is possible for some firms to receive multiple applications and for other firms to not receive any applications.

Any firm that receives at least one application hires and produces an output with a (normalized) value of 1. If a firm receives multiple applications, a worker is selected randomly subject to the requirement that the firm always select a white worker ahead of a black worker. Discrimination thus takes the form of an ordinal selection rule in hiring. Conditional on hiring

at the wage w, a firm earns $1 - w$ and the worker earns w. There this is no productivity difference associated with worker type.

Equilibrium in the overall game is an equilibrium in the wage-setting game among firms which in turn depends on an equilibrium in the application subgame among the workers.

2.3.2 Worker Subgame

As is standard in the literature, we focus on the symmetric mixed strategy equilibrium in the worker subgame. This is a situation in which all workers employ the same application strategy. We also assume anonymity of firms in the sense that if two or more firms set the same wage, all white workers apply to these firms with equal probability and all black workers apply to these firms with equal probability. In other words, if $w_i = w_j$ then $\theta_i^n = \theta_j^n$ and $\gamma_i^b = \gamma_j^b$.[5] This means that we can drop the worker specific superscripts and for a given $\mathbf{w}$ all white workers apply to the M firms with probabilities $\theta = (\theta_1, \ldots, \theta_M)$ and all black workers apply with probabilities $\gamma = (\gamma_1, \ldots, \gamma_M)$. In addition, we assume anonymity of firms in the sense that workers of a given type apply to firms that set the same wage with the same probability.

The focus on the symmetric mixed strategy equilibrium is justified on the basis that workers generally do not have a mechanism by which to coordinate their applications. It also seems reasonable to assume that workers apply to identical firms with an equal probability if the firms set the same wage. The selection of the mixed strategy equilibrium also delivers tractability because it implies that application probabilities will vary (hemi-)continuously in the set of wage announcements.

We let $\Omega(\theta_m)$ denote the probability that a white worker will get hired when applying to firm m when white workers apply to firm m with probability θ_m, where Ω has the same definition as shown in the previous section. The probability of getting hired for black workers, $\Gamma(\theta_m, \gamma_m)$, is analogous but depends on the application probabilities of both white and black workers:

$$\Gamma(\theta_m, \gamma_m) = w_m (1 - \theta_m)^N \frac{1 - (1 - \gamma_m)^B}{B \gamma_m}.$$

This is the expected payoff from applying to a firm offering a wage w_m when other black workers apply to the firm with probability γ_m, weighted by the probability that no white workers apply to the given firm, $(1 - \theta_m)^N$.

[5]Although white workers and black workers will typically not apply to the same wage level with the same probability.

A symmetric equilibrium in the worker application subgame is a set of application probabilities θ and γ such that for all firms m

$$\theta_m > 0 \Rightarrow w_m\Omega(\theta_m) = \max_{k \in M} w_k\Omega(\theta_k) \tag{2.8}$$

$$\gamma_m > 0 \Rightarrow w_m\Gamma(\theta_m, \gamma_m) = \max_{k \in M} w_k\Gamma(\theta_k, \gamma_k). \tag{2.9}$$

This is a statement of the market utility property: All firms to which workers of a given type apply must offer the market utility for their type. In the symmetric equilibrium, the application probabilities will therefore be governed by a set of indifference relationships, as presented below. Given the application probabilities of the other white workers, each white worker must be indifferent among the set of firms that attract white applicants. Similarly, given the application probabilities of black workers, each black worker must be indifferent among the set of firms that attract black applications. Since some firms may receive applications from both types of workers, some firms may receive applications from only one type, and some firms may not receive any applications at all, we let the first $1, \ldots, Q$ firms denote the firms to which white workers and black workers apply, the $Q + 1$ to $Q + T$ firms denote firms which only black workers apply, and the last $Q + T + 1$ to L firms denote firms to which only white workers apply.[6] The set of relationships that dictate white worker behavior are therefore

$$w_2\Omega(\theta_2) - w_1\Omega(\theta_1) = 0 \tag{2.10}$$
$$\vdots$$
$$w_Q\Omega(\theta_Q) - w_1\Omega(\theta_1) = 0$$
$$w_{Q+T+1}\Omega(\theta_{Q+T+1}) - w_1\Omega(\theta_1) = 0$$
$$\vdots$$
$$w_L\Omega(\theta_L) - w_1\Omega(\theta_1) = 0$$
$$\sum_{i=1}^{Q} \theta_i + \sum_{i=Q+T+1}^{L} \theta_i = 1. \tag{2.11}$$

---

[6]There can therefore exist firms $L + 1$ to M that do not receive any applications.

Similarly, letting $B = Q + T$, the set of relationships that govern black worker behavior are

$$w_2 \Gamma(\theta_2, \gamma_2) - w_1 \Gamma(\theta_1, \gamma_1) = 0 \tag{2.12}$$

$$\vdots$$

$$w_B \Omega(\theta_B) - w_1 \Omega(\theta_1) = 0$$

$$\sum_{i=1}^{B} \gamma_i = 1. \tag{2.13}$$

A corollary of the proof of Peters (1984) guarantees that these systems are have a solution and that the worker subgame with discrimination has a unique equilibrium.

Corollary 2.3.0.1. *A symmetric equilibrium to the workers' subgame consisting of vectors θ and γ exists and is unique.*

Proof. The proof in Peters (1984) applies directly for white workers because white workers are not affected by the presence of black workers. To see that this proof also holds for black workers, consider the set of *residual wages* $\tilde{W} = \left((1 - \theta_1)^N w_1, \ldots, (1 - \theta_M)^N w_M\right)$. This vector describes the expected portion of each wage available to black workers after taking into account the white applications. The same logic that holds for the white workers given an arbitrary announcement of wages then holds for the black workers given an arbitrary set of *residual wages*. $\qquad \square$

The existence of an equilibrium in the worker subgame has the important implication that the application probabilities vary continuously on the set of wages with the only complication associated with the case when all firms offer a wage of zero. The existence of an equilibrium also means that we can apply the implicit function theorem to solve for the set of probabilities as functions of the wages (See Sydsæter et al. 2008, theorem 2.7.1). For white workers, this is analogous to the case from Burdett et al. (2001) shown in the appendix. For black workers, the approach is similar but the with the added complication that the application behavior of blacks must also account for the application behavior of whites (an example is shown in appendix 2.B).

In practice, the application probabilities have to be solved for numerically.[7] The basic strategy that we use is iterative. We begin by ordering the wages W from highest to lowest. Starting with the highest wage, we consider at each stage of the iteration whether the *expected*

---

[7]Since Ω and Γ are binomial sums, the system is comprised of high order polynomials.

income for whites from applying to the n highest wages is greater than the $n + 1^{th}$ highest wage. If not, then the application probabilities are recomputed for the $n + 1$ highest wages and the process is repeated until it stops or it reaches $n = M$ (in which case all firms receive applications). When the iteration stops, the set of firms to which whites apply and the associated application probabilities are identified. After solving for the white application probabilities, we perform the same exercise for black workers except we use the *residual wages* rather than the original set of wages. The residual wages are the wages adjusted for the presence of white worker. For example, if a wage w_n attracts white workers with a probability θ_n, then the residual wage is $w_n(1 - \theta_n)^N$.

Example   In figure 2.1 we show how the white (left panel) and black application probabilities (right panel) in a market with $M = 6$ firms, $N = 5$ white workers, and $B = 2$ black workers vary when there are $x = 2$ firms setting the wage $w_h = 0.289$, $y = 3$ firms setting the wage $w_l = 0.151$, and a single firm setting the wage w_d. We use a similar style of plot throughout the paper, so it is worthwhile to elaborate on the presentation in some detail. In each plot, there are three wages that are considered. Two of the wages, w_h and w_l, are fixed and indicated by vertical lines. The third wage, w_d, is varied between 0 and 1, as indicated by the horizontal axis. At each point, the three curves show the probability that workers of a given type apply to each of the wages, given the level of the wages $\mathbf{w} = (w_d, w_h = 0.289, w_l = 0.151)$ and the number of firms setting each wage (2, 1, and 3, respectively): The blue curve shows the probability that workers apply to the two firms offering the wage w_h, the red curve shows the probability that workers apply to w_d, and the yellow curve shows the probability that workers apply to w_l. For example, in the left panel, we can see that when $w_d = w_h$ that $\theta_h = \theta_d = 1/3$ (the point where the horizontal line at 0.289 intersects the blue and red curves). This means that in this market when three firms set the wage w_h that the white workers randomize exclusively among the firms setting the high wage.

The white worker response to an increase in w_d, shown in figure 2.1a, is simple to characterize. As the wage w_d is increased, white workers reallocate their applications to w_d, away from w_l and w_h, as indicated by the red line in the left panel. Consistent with the theoretical results, θ_d is a nondecreasing, (quasi-) concave function of w_d. For the particular market presented here, it is also worthwhile to observe that when $w_d = w_h$, no white workers apply to the low wage firm, $\theta_l = 0$. This indicates that when 3 firms are setting the wage 0.289 that w_l is sufficiently low that it receives no white applicants. Notice moreover that if w_d

(a) $\theta_h, \theta_d, \theta_l$ (b) $\gamma_h, \gamma_d, \gamma_l$
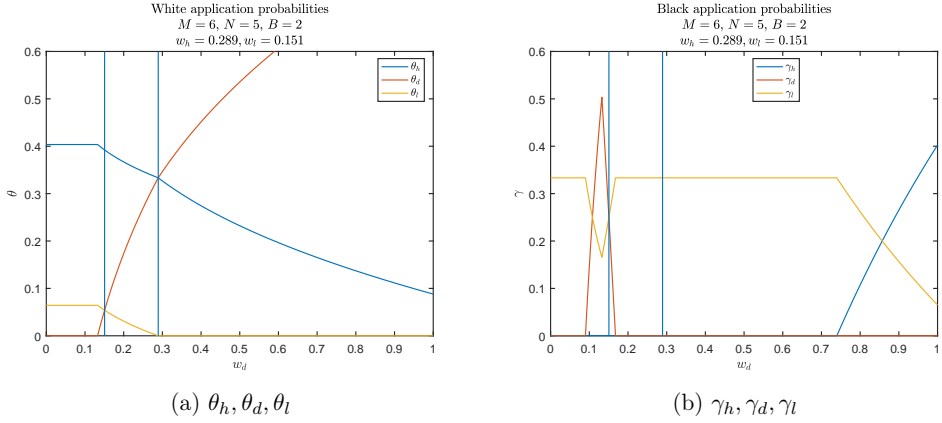
Figure 2.1: Worker Application Behavior

is reduced even marginally, the expected income available to white drops and θ_l begins to increase. This indicates that w_l must be exactly equal to the white income when three firms set w_h.

The response of black workers to an increase in w_d is more difficult to characterize. As indicated by the red "spike" in the right panel, application behavior is not monotonic in the wage. As w_d is increased from zero, there is a point at which black workers begin to apply to w_d. In an interval above this point but below w_l, only black workers are applying to w_d and, as a consequence, the rate of black applications increases as w_d is increased. θ_d increases above θ_l because the presence of white workers at w_l makes it more attractive for black workers to apply to w_d. However, once white workers begin to apply to w_d, the rate of black applications begins to fall until a point above w_l at which the presence of white workers is so costly for blacks that blacks cease applying to w_d altogether.

2.3.3   Firm Subgame

The expected profit for a firm that sets the wage w that attracts white workers with probability θ and black workers with probability γ is

$$\pi_w = (1 - w) \left(1 - (1 - \theta)^N \right) + (1 - w - \delta) (1 - \theta)^N \left(1 - (1 - \gamma)^B \right) \qquad (2.14)$$

where θ and γ are determined by the set of indifference conditions 2.10 and 2.12, and $\delta > 0$ is the difference in productivity between white and black workers. Observe that the first

term in 2.14 is the expected profit associated with white applicants while the second term is the expected profit associated with black applicants. Since we maintain the assumption that black and white workers are equally productive, equation 2.14 simplifies to

$$\pi_w = (1 - w) \left(1 - (1 - \gamma)^B (1 - \theta)^N \right). \tag{2.15}$$

For convenience, we denote the profits from white and black applicants by π_w^N and π_w^B respectively,

$$\pi_w^N = (1 - w) \left(1 - (1 - \theta)^N \right) \tag{2.16}$$

$$\pi_w^B = (1 - w) (1 - \theta)^N \left(1 - (1 - \gamma)^B \right). \tag{2.17}$$

π_w^N and π_w^B have an identical interpretation as profits in the case with homogeneous workers (equation 2.5) except that black workers are only hired if no white worker applies, a event that occurs with probability $(1 - \theta)^N$.

The question for a firm is how to adjust the wage, taking as given the wages of other firms. In general, this is informed by the derivative of profits

$$
\frac{\partial \pi}{\partial w} = - \left(1 - (1 - \theta)^N (1 - \gamma)^B \right)
$$
$$
+ (1 - w) N (1 - \theta)^{N-1} (1 - \gamma)^B \frac{\partial \theta}{\partial w}
$$
$$
+ (1 - w) B (1 - \theta)^N (1 - \gamma)^{B-1} \frac{\partial \gamma}{\partial w}. \tag{2.18}
$$

The first term is the marginal cost of paying a higher wage. Since an increase in the wage reduces the profit from hiring one-to-one, this term equals the negative of the probability of hiring. The second and third terms in expression 2.18 relate to the change in the number of applicants. The second term is the marginal increase in profits associated with an increase in the number of white applicants while the third term is the marginal change in profits associated with changes in the number of black applicants. Because a higher wage always increases the probability that whites apply (as long as the wage is equal to or above the white expected income), $\frac{\partial \theta}{\partial w} \geq 0$ and the second term is non-negative. In contrast, the third term has an ambiguous sign because $\frac{\partial \gamma}{\partial w}$ can take both positive and negative values (see the behavior in the right panel of figure 2.1).

Gathering together the terms associated with the change in the number of applications,

we see that the change in profits from increasing the wage is negative if the change in the number of applicants weighted by the profit of hiring is less than the odds of hiring:

$$(1 - w) \left(\frac{N}{1 - \theta} \frac{\partial \theta}{\partial w} + \frac{B}{1 - \gamma} \frac{\partial \gamma}{\partial w} \right) < \frac{1 - (1 - \theta)^N (1 - \gamma)^B}{(1 - \theta)^N (1 - \gamma)^B}.$$

A sufficient (but not necessary) condition for profits to decrease is therefore that $\frac{N}{1 - \theta} \frac{\partial \theta}{\partial w} + \frac{B}{1 - \gamma} \frac{\partial \gamma}{\partial w} < 0$. If an increase in the wage is not more than compensated by increases in applications, then a increase in the wage is not worthwhile.

A pure strategy equilibrium in the game is a set of wages $\mathbf{w}^*$ such that no firm can deviate to some other wage and increase their profits given the optimal application behavior of workers.

2.4 Results

We begin by showing that there is no equilibrium in which all firms set the same wage. Although this result is intuitive it demonstrates the basic approach that we utilize in this section. In the main result of the paper, we consider the possibility of a segregated pure-strategy equilibrium of the type identified by Lang, Manove, and Dickens (2005). This is an equilibrium in which the labor market divides into a white high wage submarket and a black low wage submarket. We demonstrate that such equilibria tend not to exist in small markets.

From a technical vantage point, analysis of this setting is complicated because there is no guarantee that there is a pure strategy equilibrium. Black workers avoid competing with white workers and the profit function can have multiple peaks. This means that the profit function fails to be quasi-concave.[8] To the best of our knowledge, there is no theorem that guarantees existence of a pure strategy equilibrium in the absence of quasi-concavity. The form of the profit functions thus makes an analytic characterization of firm behavior difficult.

An additional challenge arises from points of non-differentiability in the profit function. Because black workers respond strongly to the presence of white workers, black worker application behavior changes discretely when they encounter white workers. If the low wage is equal to the white expected income, a low wage firm that increases its wage will experience

---

[8]To see that the profit function fails the requirement of quasi-concavity, it sufficient to observe that the profit function can have multiple peaks. In any such case, there must necessarily be an upper level set of the function that fails to be convex.

a discrete loss of black applications. Moreover, the loss of black applicants will outweigh the increase in white applicants. Hence, profit maxima can exist at points of non-differentiability. This corresponds to the segregated equilibrium suggested by Lang, Manove, and Dickens (2005). Because standard optimization techniques fail at such points, it is necessary to demonstrate that the derivative changes sign by *ad hoc* techniques.

2.4.1   Single Wage Equilibrium

Given that all firms are identical, it is natural to ask whether there can exist an equilibrium in which all M firms set the same wage and workers of both types randomize uniformly among the posted wages. Because this is a situation in which firms attract workers of both types, we refer to this as a *pooling* equilibrium.

Result 2.4.1. *There do not exist equilibria in which all firms set the same wage.*

A pooling equilibrium in which all firms set the same wage w_p can not occur. Regardless of the level of the wage, it is profitable to deviate to a wage w_d that is less than w_p.[9] This is a consequence of the fact that black workers have a strong incentive to avoid competing with white workers. When a firm reduces the wage, the loss of white workers is more than offset by the gain in black workers.

To show this formally, we demonstrate that

$$\frac{N}{1-\theta}\frac{\partial\theta}{\partial w} + \frac{B}{1-\gamma}\frac{\partial\gamma}{\partial w} < 0.$$

Substantively, this amounts to showing that the net change in the number of applicants associated with a change in the wage is less than zero. We plug in for $\frac{\partial\theta}{\partial w}$ and $\frac{\partial\gamma}{\partial w}$ and take advantage of the fact that $\theta = \gamma = 1/M$ at the point $w_d = w_p$. We relegate the details of the derivation to appendix 2.C.

The intuition for this result is illustrated in figure 2.1 for a market with $M = 6$ firms, $N = 5$ white workers, and $B = 2$ black workers and the candidate pooling wage $w_p = 0.25$. The vertical line is the point at which $w_d = w_p = 0.25$ and the heavy blue line shows the level of profits π_d associated with the wage w_d when $M - 1$ firms are setting the wage $w_p = 0.25$. As is evident from the plot, there exists a profitable downward deviation when $w_d = w_p$ since the heavy blue line is downward sloping where it crosses the vertical line:

---

[9]Unless $w_p = 0$, in which case a firm can deviate marginally and attract applicants with certainty and thereby discretely increase their profits.

Decomposing profits into the part arising from black applicants, π_d^B (blue dash), and white
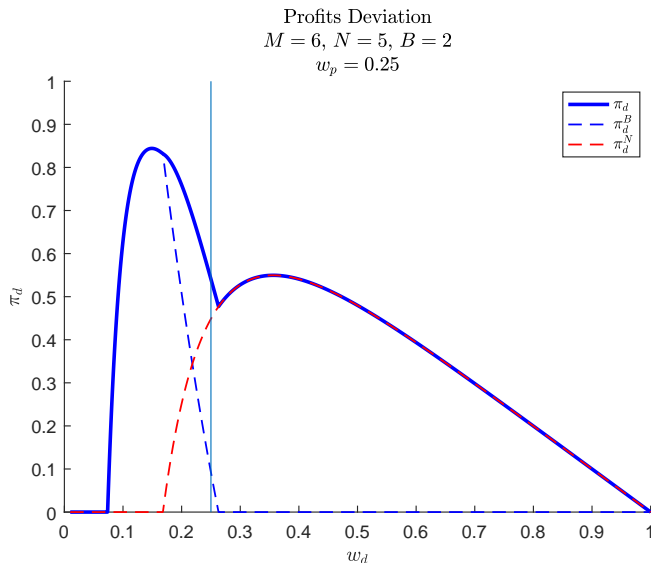
Figure 2.1: No Pooling Equilibrium

applicants, π_d^N (red dash), we see that the opportunity for a profitable downward deviation arises because the black workers respond more strongly to a marginal change in the wage than white workers. Moreover, it is optimal to deviate discretely to a point substantially below w_p, to a point at which white workers cease applying to w_d.

In the appendix in figure 2.C.1, we present analogous figures to figure 2.1 but for $w_p = 0.10$ and $w_p = 0.50$. These plots deliver the same intuition as figure 2.C.1: There always exists a profitable downward deviation. However, the plot for $w_p = 0.5$ also makes it clear that there need not exist a profitable deviation upward, something that is ambiguous from figure 2.1 given that there is a region above w_p for which profits are increasing.

2.4.2 Two Wage Segregated Equilibrium

Our primary interest is the existence of segregated equilibria in which a set of x firms post a high wage w_h to which only white workers apply and another subset of $y = M - x$ firms set a low wage w_l to which only black workers apply. We refer to such a constellation of firms and wages as a *candidate* for a segregated equilibrium. This is a situation in which there

are two sub-markets, a high wage white sub-market and a low wage black sub-market. The essential question is whether there exists a sorting of x of the total M firms into the high wage market such that no firm has a profitable deviation given wages w_h and w_l.

Our basic finding is that segregated equilibria of this kind do not necessarily exist, and in fact tend not to exist in the small market setting. The main part of this section is dedicated to showing why this is the case. We structure the argument in several steps. We begin by establishing results that structure and simplify the analysis. In particular, we argue that for any market constellation only a subset of wages are possibly consistent with an equilibrium and that there must be sufficient competition—in the sense of number of agents—on each side of the market. Next, we provide an example of a representative market with $M = 6, N = 5$, and $B = 2$ in which a segregated equilibrium does not exist.[10] The reason that the segregated equilibrium does not exist is that there is either the opportunity to undercut the lowest wage or overbid the highest wage. The example thus establishes that a segregated equilibrium need not always exist. But how general is this finding? As it turns out, the example market is representative of a large number of different markets and the deviations that are present in the example market are reproduced in other markets. Based on the insights from this example, it is possible to systematically search through a large number of different markets and unambiguously establish that there a large number of cases in which segregated equilibria do not exist.

Before turning to the analysis, it is worthwhile to elaborate on the differences between the continuum market and the discrete market that we investigate. In contrast to the continuum market, in the discrete case the profits in the two sub-markets need not be equal. Consider, for example, a situation in which firms in the high market earn higher profits than firms in the low market. To show that there is a segregated equilibrium, it is necessary to demonstrate that despite the higher level of profits, if an additional firm tries to compete for white workers, that the increased competition for white workers drives down profits in the white market sufficiently that the firms in the low market prefer not to deviate. Perhaps not surprisingly, it is not always possible to do so.

Preliminary Assertions

Wages We begin by observing that the high wage w_h in any segregated equilibrium must be given by $w(x, N)$, expression 2.7. This is the same wage as would arise in a market

---

[10]The market is representative in the sense that there are sufficiently many participants of each type that a segregated equilibrium could plausibly exist.

constituted by N homogeneous workers and x firms. Because the high wage market is by definition segregated from the low market, the high wage must be determined by the number of firms that compete for white workers in the same fashion as if the low wage market did not exist. If the wage were above or below $w(x, N)$, there would be an incentive for a firm in the high market to deviate. The wages and profits associated with the high market are therefore limited to the levels shown in tables 2.A.1 and 2.A.2, where x replaces M.

For a segregated equilibrium to exist, the low wage w_l must also be sufficiently low that it does not attract any white workers. Specifically, w_l cannot be higher than the white expected income:

$$w_l \leq \Omega(x)w_h. \tag{2.19}$$

To the contrary, if w_l were greater than $\Omega(x)w_h$, then white workers would occasionally apply to the low wage. Notice that this condition can hold with equality because workers employ symmetric application strategies. If white workers were to begin to apply to firms offering the wage $\Omega(x)w(x)$, the expected income would be reduced strictly below that available at the high wage firms. In combination with the first assertion, this implies that $w_l \leq \Omega(x)w(x, N)$. In the remainder, we omit N where it does not cause confusion.

In the segregated equilibrium identified by Lang, Manove, and Dickens (2005), the low wage is exactly equal to the white expected income. Translating this to our setting, this implies that $w_l = \Omega(x)w_h$, where $\Omega(x)$ denotes the probability that a white worker is hired when there are x firms offering the wage w_h. These low wages are shown in table 2.C.1. Thus, in these equilibrium candidates, the high and low wage are pinned down by the number of firms that participate in the white sub-market, x. For a given number of firms M, there are therefore only a limited number of possible candidates that need to be considered.

Moreover, note that if the low wage is equal to the white expected income then this implies that the wage in the low market is constrained by the presence of the white sub-market in the sense that $w_l < w(y, B)$, where $w(y, B)$ is the wage that would arise if y firms competed for B workers in the absence of the constraint imposed by the white market. In other words, the presence of white workers depresses the wage in the black submarket.

Market Size In order for a segregated equilibrium to exist, there must be sufficient competition for each worker type in the sense that no firm can hire with certainty. Intuitively, this follows from the fact that a firm that doesn't face competition has no incentive to maintain a high wage. A consequence of this result is that there must be at least two firms in the white submarket and two firms in the black submarket.

Consider a market in which there is a single firm that sets a wage w and another subset of firms that offer the (non-zero) wage w_l. Assume in addition that $w_l \leq \frac{w}{N}$ such that white workers apply to w with probability $\theta_h = 1$. In this case, black workers randomize among the firms offering w_l since they are never hired at w because of the presence of white workers.

In general, the derivative of profits for the firm that sets a wage w to which white workers apply with probability θ_w is

$$\frac{\partial \pi}{\partial w} = -(1 - (1 - \theta_w)^N) + (1 - w)N(1 - \theta_w)^{N-1}\frac{\partial \theta_w}{\partial w}.$$

As long as $w_l < \frac{w}{N}$ this derivative is clearly negative. Because the low wage is lower than what whites can expect to earn at the high wage, $\theta_w = 1$, the first term equals -1, and the second term equals zero. The second term equals zero because $\theta_w = 1$ does not change and as a consequence, both the $(1 - \theta_w)^{N-1}$ and $\frac{\partial \theta_w}{\partial w}$ terms are zero. When $w_l = \frac{w}{N}$, θ_l is exactly equal to zero and any additional reduction in w will cause white workers to begin to apply to the low wage firm. However, the derivative remains negative at this point because $\frac{\partial \theta_w}{\partial w}$ approaches a fixed positive number when $\theta_l \to 0$ while $(1 - \theta_w)^{N-1}$ becomes arbitrarily small because $\theta_w = 1 - y\theta_l$ (shown in appendix 2.C). It must therefore be optimal for a firm to continue to reduce the wage. The outcome is that the firm reduces the wage sufficiently that w attracts both worker types. There must therefore be at least two firms competing for white workers.

A single firm in the black market is impossible for the same reason. For any wages that satisfy $w_l > B(1 - 1/x)^N w_h$, the firm can reduce the wage without any loss of applicants because black workers strictly prefer the low wage. Moreover, by a similar computation as for white workers, it is possible to show that the derivative remains negative for w_l close to w_l. This means that it is optimal for a firm that receives all the black applications to reduce the wage sufficiently that the black workers begin to apply to the high wage.

Can there be a situation with only a single white worker? Since the single white worker always gets hired, the white worker would always apply to the highest wage. But since there must be at least two firms competing for the white worker (or else the wage would fall), a firm would either wish to deviate upwards and hire the white worker with certainty or, if the wage is sufficiently high, deviate to the low market and compete instead for black workers. As a consequence, there is no wage compatible with a pure strategy equilibrium if there is only one white worker. The only case with a single participant on one side of the market which we cannot exclude is a single black worker. Analogous to white workers,

firms in the black sub-market gain discretely from a marginal increase in the wage if it is below the white expected income. However, when w_l equals $\Omega(x)w(x)$ a marginal increase begins to attract white workers. The black worker may therefore strictly prefer w_l to wages above this level because of competition from white workers. This constrains the incentives of firms in the black sub-market to bid the wage above the white expected income. In such a situation, the y firms setting the low wage earn $\frac{1-\Omega(x)w(x)}{y}$ while the x high wage firms earn $\left(1 - \left(1 - \frac{1}{x}\right)^N\right)(1 - w(x))$.

Critical Points   The intuition from Lang, Manove, and Dickens (2005) largely carries over to the small market setting. As is the case in the large market analog, we find that $w_h = w(x)$ and $w_l = \Omega(x)w_h$ are local maxima in the sense that firms setting either of the wages cannot deviate to some nearby wage and increase their profits.
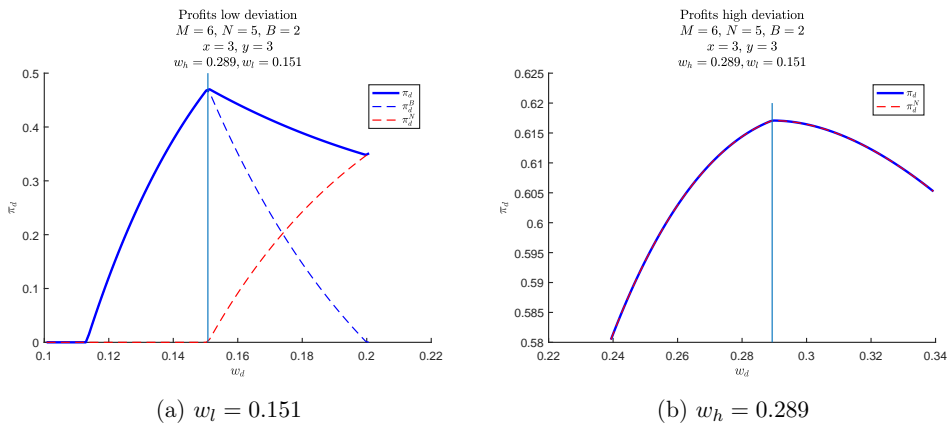
Figure 2.2: Local Maxima at w_h and w_l

w_h is a local maximum because it is determined as the optimal response to $x - 1$ other firms setting the high wage (the same as in Burdett et al. (2001)). Figure 2.2b shows this for the same market that was used as an example in the previous section ($M = 6, N = 5, B = 2$). That w_l should also be a local maximum follows from a similar intuition as why there are not pooling equilibrium. An increase in the wage increases competition from whites and this drives away black applicants faster than profits increase from additional white applicants. An illustration of this is shown in figure 2.2a.[11]

---

[11]To see that this also holds for other numbers of high firms x, see figure 2.C.2 in the appendix.

To demonstrate that this intuition holds in general is, however, tricky because the profit function (for a firm that deviates from w_l) will be kinked at w_l whenever the low wage is constrained by the presence of the white workers. This kink exists because white workers begin to apply to wages above w_l (since they are above the white expected income and therefore get applied to with positive probability) and this has a discrete negative impact on how often black workers apply. Demonstrating that w_l is a critical point therefore depends on showing that the derivative of profits is increasing at w_l but decreasing for an arbitrarily small increase in the wage beyond w_l.

That profits are increasing for $w_d < w_l$ follows from the assumption that w_l is constrained.[12] This means that the increase in profits from attracting more black applicants outweighs the increase in wage costs for any wage at or below w_l. To show that profits are decreasing for an infinitesimal increase beyond w_l, we manipulate the derivative of profits and show that as $\theta_d \to 0$ this derivative is strictly negative. This proof relies on the fact that $w_l = \Omega(x)w_h$ each worker prefers their own market such that $\theta_h = 1/x$ and $\gamma_h = 1/y$. Critically, the level of the wage, w_h, cancels from the expression.[13] We may therefore conclude that w_l is a critical point for all equilibrium candidates in which the low wage equals the white expected income. The derivation is shown in appendix 2.C.

Example, $M = 6, N = 5$, and $B = 2$

Although $w_h = w(x, N)$ and $w_l = \Omega(x)w_h$ are local maximum, the outstanding question is whether there exists a discrete (rather than marginal) deviation that will increase profits. An example serves to demonstrate that such a deviation can exist. This shows that the pure strategy equilibrium from Lang, Manove, and Dickens (2005) need not carry over to the small market setting.

Result 2.4.2. *A segregated equilibrium does not always exist.*

Consider a market with $M = 6, N = 5$, and $B = 2$. One candidate for a segregated equilibrium is a configuration with three firms setting the high wage ($x = 3$) and 3 firms setting the low wage ($y = 3$). In this candidate, the firms in the high wage market earn higher profits than firms in the low wage market. However, this can not be an equilibrium because a deviant firm in the low wage market setting w_d can deviate from $w_d = w_l$ to $w_d = w_h$ and

---

[12]We argue below that there are no segregated equilibrium candidates in which the low wage is not constrained.

[13]Which means that also w_l cancels, as w_l can be written as a function of w_h.

increase their profits, as illustrated in figure 2.3a: At the far left in this figure are the profits available from setting the low wage while the far right is the profit available from deviating to w_h. The level of profits from setting w_l is indicated by the horizontal line. As is evident, profits are higher at the far right. This suggests that the tightness in the white market is too low. One might therefore expect that the problem is simply that the firms have not sorted properly. We therefore consider a different equilibrium candidate in which the tightness in the high market is greater, with four high wage firms ($x = 4$) and two low wage firms ($y = 2$) and a corresponding increase in the wages since there is more competition in the high wage market. In this candidate, the low wage firms earn higher profits. However, this too cannot be an equilibrium. As is illustrated in figure 2.3b, a high wage firm can deviate to w_l and increase their profits. Since neither $x = 3$ nor $x = 4$ are deviation proof, this illustrates that an segregated equilibrium need not exist.[14] Tables 2.C.2 and 2.C.3 in appendix section 2.C summarize the market outcomes for this example.

(a) $x = 3$, low deviation possible   (b) $x = 4$, high deviation possible
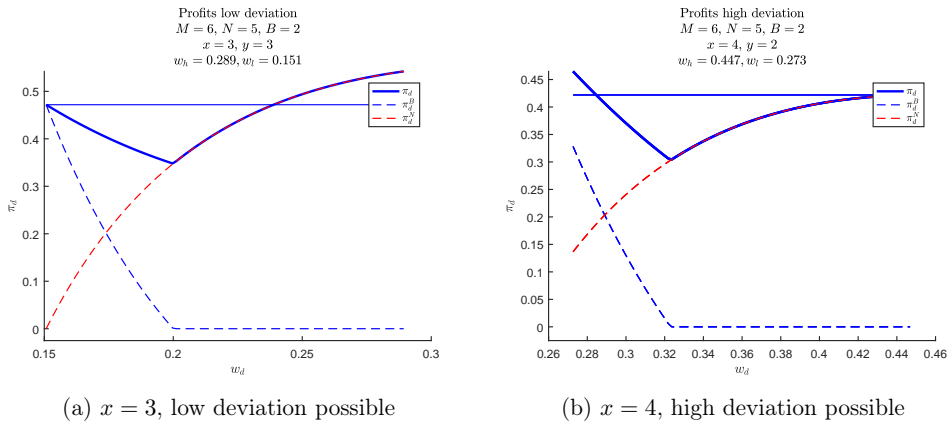
Figure 2.3: Profitable Deviations for $x = 3$ and $x = 4$

Optimal Deviation Analysis

The example yields two complementary observations. Both observations relate to deviations in which a firm jumps discretely and begins to compete for workers of the "other type." The

---

[14]To be complete, we would also need to show that $x = 2$ exhibits a profitable deviation. We omit this case, but it turns out to be consistent with intuition. If there is a deviation to the high wage when $x = 3$ then it is also the case that there exists a deviation to the high wage when $x = 2$. The case with $x = 2$ is similar to that in figure 2.3a but with an even greater incentive for a deviation.

first observation is that for a low wage firm that deviates to the high wage, the derivative of profits is increasing when $w_d = w_h$ as seen the left panel of figure 2.3. This means that it is even better for such a firm to deviate to a wage strictly above w_h. The second observation is that for a high wage firm that deviates to the low wage, the derivative of profits is negative when $w_d = w_l$ as illustrated by the right panel of figure 2.3. This means that it is even better for such a firm to deviate to a wage strictly below w_l.

The reason for the first observation is that a firm that unexpectedly deviates to the high wage market can take advantage of the fact the wage is now too low relative to the tightness. If there is a segregated equilibrium with x firms setting the high wage, then the wage in the white sub-market must be $w_h = w(x)$. But if an additional firm participates in the high wage market, then this wage is lower than what would arise in a market with $x + 1$ firms. This means that for a low wage firm that deviates to the vicinity of w_h that the increase in profits from attracting additional applicants must exceed the loss from offering a higher wage and the optimal "upward" deviation is above w_h.

The reason for the second observation is that firms in the low wage market were initially setting the wage $w_l = \Omega(x)w(x)$, which is exactly low enough that no white workers apply. When instead a high wage firm deviates and only $x - 1$ firms compete in the high market, this induces white workers to apply to w_l. In turn, this creates an incentive for black workers to now apply to even lower wages. As a result, if a high wage firm wishes to deviate to a wage $w_d \leq w_l$, then the high wage firm should deviate to strictly below w_l and benefit both from hiring a black worker with high probability and from paying a lower wage.

Together these observations imply that there may be profitable deviations above w_h and/or below w_l. Even if there are no profitable deviations for low firms in the vicinity of w_l, there may be an opportunity to deviate and compete for white workers by overbidding the wage in the high market. Analogously, a firm in the high market may be able to deviate and undercut the low wage and attract black workers who wish to avoid competition from white workers.

To systematically check whether whether segregated equilibria exist, we compute the optimal deviations of the kind described above for a large number of different markets and market constellations. Specifically, for each combination of x, $y = M - x$, N, and B, we compute the following quantities:

1. The wages and profits π_h and π_l associated with the candidate equilibrium.

2. The wages and profits associated with (1) the optimal "low" deviation for a high wage

firm that deviates and begins to compete for black workers and (2) the optimal "high" deviation for a low wage firm that deviates and begins to compete for white workers. In all cases, the best deviation is associated with either undercutting the low wage (the "low" deviation) or overbidding the high wage (the "high" deviation). We let $\overline{\pi_d}$ denote the profit associated with the optimal deviation by a low firm to a wage $\overline{w_d} \geq w_h$ and let $\underline{\pi_d}$ denote the profit associated with the optimal deviation by a high firm to a wage $\underline{w_d} \leq w_l$.

3. Based on computations (1.) and (2.), the difference in payoffs

$$\Delta \pi_h = \underline{\pi_d} - \pi_h$$
$$\Delta \pi_l = \overline{\pi_d} - \pi_l.$$

If either $\Delta \pi_h > 0$ or $\Delta \pi_l > 0$, then there is a profitable deviation from the candidate equilibrium and some other combination of x and y must be checked. If there are no profitable deviations for any combination of x and y, then we may conclude that there is no segregated equilibrium for the given M, N and B. If, however, $\Delta \pi_h < 0$ and $\Delta \pi_l < 0$ for some x and y then a segregated equilibrium is possible.

As an example, table 2.1 presents the optimal deviations associated with the market presented in the example above. As can be seen in the columns indicated by $\Delta \pi_h$ and $\Delta \pi_l$ there is always a profitable deviation available (one or both are positive for every combination of x and y):

Table 2.1

			$M = 6$,	$N = 5$,	$B = 2$				
x	y	w_h	w_l	π_h	π_l	$\underline{w_d}$	$\Delta \pi_h$	$\overline{w_d}$	$\Delta \pi_l$
2	4	0.088	0.034	0.884	0.423	0.028	-0.064	0.130	0.413
3	3	0.289	0.151	0.617	0.472	0.132	0.034	0.340	0.087
4	2	0.447	0.273	0.422	0.546	0.248	0.139	0.487	-0.167

To illustrate that this holds for most small markets, figure 2.4 illustrates how $\Delta \pi_h$ (in blue) and $\Delta \pi_l$ (in red) behave for a large number of different markets. For the all markets in the figure, $M = 6$ firms, $B = 2$ black workers, and a number of N white workers is given by the title of each panel. For instance, panel 5 located at the left in the middle row summarizes the results from table 2.1.
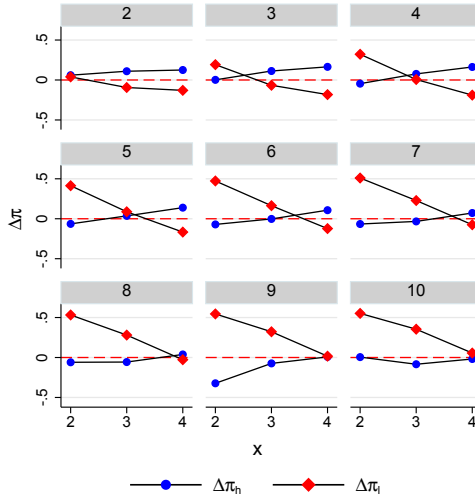
Figure 2.4: No segregated equilibrium, $M = 6$, $B = 2$

In order for a segregated equilibrium to exist, both $\Delta\pi_h$ and $\Delta\pi_l$ would need to be below zero, as indicated be the dashed red line. This would indicate the absence of profitable deviations. There are no such cases in figure 2.4. This is a general finding: In the majority of cases, there does not appear to exist a segregated equilibrium of the kind predicted by Lang, Manove, and Dickens (2005).

Existence

The results presented above demonstrate that a segregated equilibrium is not guaranteed to exist in the small market setting. In fact, this is the typical case. However, the result can not be strengthened to conclude that segregated equilibrium never exist. Using a similar method as above to investigate a large number of markets, it is possible to find constellations for which both $\underline{w_d}$ and $\overline{w_d}$ are negative. We examined all combinations of x and y for markets with $M \in \{4, \ldots, 10\}$, $N \in \{2, \ldots, 20\}$, and $B \in \{2, 3\}$ and found a handful of instances in which a segregated equilibrium does appear to exist. This implies that there are cases in which no low firm can overbid the high wage and no high wage firm can undercut the low wage.

One such case is a market with $M = 7$, $N = 14$, and $B = 3$. Table 2.2 presents the

results for this case. The bottom row presents the case for $x = 5$ and $y = 2$. For this market composition and this constellation of firms, $\Delta \pi_h = -0.033$ and $\Delta \pi_l = -0.031$ (also illustrated in figure 2.C.3 in the appendix). This means that neither the high or low deviation are profitable. The conclusion is that this is an instance in which a segregated equilibrium may be possible.

Table 2.2

$M = 7$, $N = 14$, $B = 3$									
x	y	w_h	w_l	π_h	π_l	$\underline{w_d}$	$\Delta \pi_h$	$\overline{w_d}$	$\Delta \pi_l$
2	5	0.000	0.000	0.999	0.488	0.000	0.000	0.000	0.508
3	4	0.016	0.003	0.980	0.576	0.003	-0.189	0.022	0.397
4	3	0.065	0.018	0.918	0.591	0.015	-0.140	0.079	0.210
5	2	0.133	0.045	0.829	0.835	0.039	-0.033	0.151	-0.031

Figure 2.5: Segregated Equilibrium, No Profitable Deviations

To confirm that this is indeed the case, we plot the profits associated with a deviation w_d from the candidate $x = 5$ and $y = 2$. Figure 2.5a shows the case when a high wage firm deviates to some other wage w_d and figure 2.5b shows the case when a low wage firm deviates to some other wage w_d. In both cases, it is clear that no firm can unilaterally deviate. The highest profit available to the firms setting the high wage is at w_h and the highest profit available to firms setting the low wage is at w_l.

110

Segregated equilibrium candidates, $w_l < \Omega(x)w(x)$

The other possibility for a segregated equilibrium is the case in which the low wage is strictly below the white expected income, $w_l < \Omega(x)w_h$. The existence of such an equilibrium is limited by two countervailing requirements. The first requirement is that in equilibrium, firms in the white submarket and the black submarket earn about the same level of profits. The equal profit requirement implies that the tightness in the black submarket should be higher than in the white submarket in order to compensate for the fact that $w_l < w_h$. If this were not the case, then the low wage market would be more profitable both because of the lower wage and because of a greater chance of hiring. However, if the tightness is greater in the low wage market then this implies more competition for black workers and in turn a higher wage. This is inconsistent with the requirement that the low wage be lower than the white expected income, $w_l < \Omega(x)w_h$. These two requirements appear to be incompatible and in practice we do not find situations in $w_l < \Omega(x)w_h$ can arise as an equilibrium.

The challenge in formalizing this argument is related to the small market setting. In particular, the scale of the market matters. Even with the same tightness, the wage can vary meaningfully. For example, for a tightness equal to $1/2$, the wage grows from 0.154 when there are two firms up to 0.313 when the market becomes very large, a more than doubling of the wage. In addition, the small market differs from the large market setting because the profits earned by firms need not be exactly equalized.

Although we are not able to establish a completely general proof, in the vast majority of cases, this type of segregated equilibrium can be ruled out. To see this, assume that $N > B$ and that the tightness is the same in both markets such that $y/B = x/N = k$. Assume in addition that the wage in the high wage market is larger simply because of scale effects.[15] For any such market, the profits earned by firms in the black market are greater than those in the white market for two reasons: The wage is lower and the matching probability is higher.[16] However, regardless of the scale of the high wage market, the low wage tends to be greater than the expected income of whites, even in the case when $N \to \infty$.[17] This means that for the low wage to be set freely, the tightness in the low wage market must be at least somewhat lower than that in the high wage market. But this means that there must be at

---

[15]Whether the wage is higher or lower because of scale effects depends on the relative tightness of firms to workers (see the table of wages in the appendix). We focus on cases in which the wage is higher because we are interested in cases in which w_l falls below the white expected income.

[16]the matching probability is higher because the market has a smaller scale but the same tightness.

[17]The exception is when the tightness is so low that the high wage is close to zero. The result therefore depends on a sufficient level of competition.

least one less firm competing for black workers (and one more competing for white workers) relative to the case when the tightness was equal. If this is the case, then a firm in the high market could deviate and achieve at least the same profits as a low wage firm when the tightnesses were equal. But this means that there is profitable deviation.

2.4.3 Other Equilibrium Candidates

The game we investigate does not have a segregated equilibrium in most cases. The question is therefore what other equilibria, whether pure or mixed, may be possible.

Pooling Equilibria

One possibility for a pure strategy equilibrium not explicitly considered earlier is a pooling equilibrium in which there are multiple wages, at least some of which attract both types of workers. We are skeptical that such equilibria exist. As in the model of Lang, Manove, and Dickens (2005), it appears that it is rarely if ever optimal to attract both white and black workers. While we were not able to prove this result in general for the small market setting, several pieces of evidence support the conjecture. First, in all analyses of segregated candidates, the profits associated with deviations to a wage w_d between w_l and w_h characteristically first fall and then climb, with the change in sign occurring at the wage at which black workers cease to apply to the deviant firm. This indicates that the derivative of profits is negative whenever w_d is attracting both worker types. Second, if a high wage firm deviates to a point below w_l it appears that the profit maximizing point is the point at which white workers cease to apply to the deviant. The implication is again that the firm should continue adjusting the wage until only one of the worker types is applying.

Mixed Strategy equilibrium

Although existence results for pure strategy equilibria do not apply for our model, the existence results for mixed strategy equilibria depend on more forgiving requirements. In particular, the existence of mixed strategy equilibria depend on continuity of the payoff function rather than quasi-concavity (See Fudenberg and Tirole 1989, Chapters 1 and 12, in particular the Maskin and Dagupta result for discontinuous games, theorem 12.4.). With the exception of a discontinuity when all firms set the wage w, the profit functions of firms are continuous in our setting (albeit not necessarily differentiable). Mixed strategy equilibria

will therefore exist in general. The challenge is that characterizing the mixed strategies turns out to be difficult.

Given the motivation to study segregated equilibria, a natural candidate is a mixed strategy equilibrium in which firms randomize among two wages, a high wage w_h and a low wage w_l, with a symmetric probability p. If firms randomize between the high and low wages, then each firm must be indifferent between setting either of the wages given the probability with which the other firms set the high and low wages. Let $\theta_h(k)$ and $\theta_l(k)$ denote the probability that white workers apply to the high and low wage firms when k firms set the high wage and let $\gamma_h(k)$ and $\gamma_l(k)$ denote the same quantities for the black workers. Then the randomization probability p must be such that

$$(1 - w_h) \sum_{k=0}^{M-1} \binom{M-1}{k} p^k (1-p)^{M-1-k} \left(1 - (1 - \theta_h(k+1))^N (1 - \gamma_h(k+1))^B\right) =$$
$$(1 - w_l) \sum_{k=0}^{M-1} \binom{M-1}{k} p^k (1-p)^{M-1-k} \left(1 - (1 - \theta_l(k))^N (1 - \gamma_l(k))^B\right).$$

In addition, it must be such that firms would not like to deviate to any other wage. This latter condition is enormously stringent.

To get a sense of whether randomization of this kind can mitigate the possibility of profitable deviations, we use a simple routine to search through various possibilities. While computationally intensive, the routine has three basic steps:

1. Fix a low wage w_l and high wage w_h.

2. Solve for a mixing probability p^* such that firms are indifferent between w_l and w_h.

3. Compute whether there exists a profitable deviation to some other w_d in the range $(0, 1)$.

   - If there is a profitable deviation, adjust w_h and begin from step 1.

   - If there is no profitable deviation, then (w_h, w_l, p^*) constitutes a two-point mixed strategy equilibrium.

To facilitate the interpretation of the results, the routine produced a plot showing the profits as a function of $w_d \in (0, 1)$ at the completion of every iteration. The main output of the procedure was thus a set of plots of π_d for every combination of w_l and w_h. An example for $M = 6$, $N = 5$, and $B = 2$ is shown in figure 2.6. The value of w_l is fixed to 0.15 while the
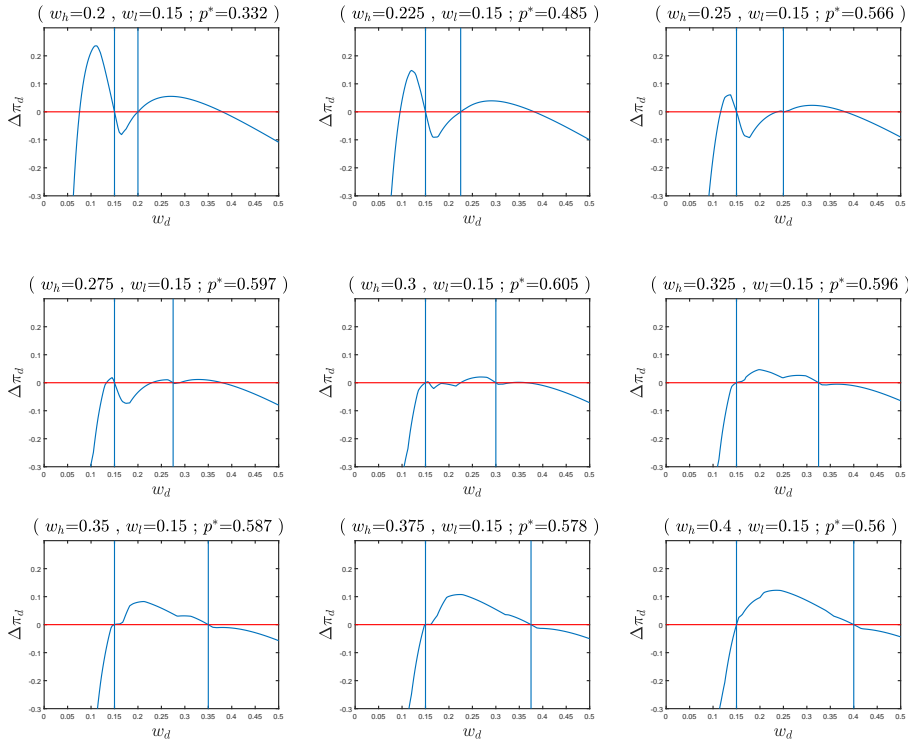
Figure 2.6: Two-point distribution candidates, $M = 6$, $N = 5$, and $B = 2$

high wage is adjusted in each panel. The wages are indicated by the vertical lines while the profit associated with a deviation is indicated by the blue curve.

Although our procedure is relatively coarse, our preliminary findings suggest that a two-point distribution is sensitive to similar types of deviations as the segregated pure strategy case—and is perhaps even more delicate. As can be seen in the top row of figure 2.6, when the wages are close together, there are both high and low deviations that are profitable. Although the gain from deviations become smaller as the wages approach the level predicted by the pure strategy segregated equilibrium (row two), we do not observe a movement toward a bimodal form as would be necessary to get mixing over two wages.

Another candidate for a mixed strategy equilibrium is a mixed strategy in which firms randomize their wages according to a continuous distribution similar to standard models of wage dispersion. This type of equilibrium is plausible exactly because it eliminates the ability of firms undercut or overbid each other. It is, however, challenging to say how such

a a distribution might look or even what the support would be. One approach would be to discretize the wages and use numerical techniques to find a fixed point for the (discrete) distribution. Then, as the discretization is made more and more fine, the hope would be that the result converges to some limit. We leave this for future work.

2.5   Conclusion

The model of Lang, Manove, and Dickens (2005) makes a stark prediction: In a posted wage environment, discriminatory hiring leads to labor market segregation. This is an intuitive result, but it does not carry over to the small market setting. We find that when there are strategic interactions between firms that the postulated segregated equilibrium tends not to exist.

2.6   References

References

Burdett, K., Shi, S., & Wright, R. (2001). Pricing and matching with frictions. *Journal of Political Economy*, *109*(5), 1060–1085.

Fudenberg, D. & Tirole, J. (1989). Noncooperative game theory for industrial organization: An introduction and overview. *Handbook of industrial Organization*, *1*, 259–327.

Galenianos, M. & Kircher, P. (2012). On the game-theoretic foundations of competitive search equilibrium. *International economic review*, *53*(1), 1–21.

Holden, S. & Rosén, Å. (2014). Discrimination and employment protection. *Journal of the European Economic Association*, *12*(6), 1676–1699.

Lang, K., Lehmann, J., & Yeon, K. (2012). Racial discrimination in the labor market: Theory and empirics. *Journal of Economic Literature*, *50*(4), 959–1006.

Lang, K. & Manove, M. (2011). Education and labor market discrimination. *The American Economic Review*, *101*(4), 1467–1496.

Lang, K., Manove, M., & Dickens, W. T. (2005). Racial discrimination in labor markets with posted wage offers. *The American economic review*, *95*(4), 1327–1340.

Montgomery, J. D. (1991). Equilibrium wage dispersion and interindustry wage differentials. *The Quarterly Journal of Economics*, *106*(1), 163–179.

Peters, M. (1984). Bertrand equilibrium with capacity constraints and restricted mobility. *Econometrica: Journal of the Econometric Society*, 1117–1127.

Sydsæter, K., Hammond, P., & Seierstad, A. (2008). *Further mathematics for economic analysis.* Edinburgh Gate, harlow, Essex CM20 2JE, England: Financial Times Prentice Hall.

Wilson, V. & Rodgers III, W. M. (2016). *Black-white wage gaps expand with rising wage inequality.* Economic Policy Institute.

Appendix

2.A Background

Large Market

To derive outcomes in the large market case with homogeneous firms and homogeneous workers with a market tightness $k = \frac{M}{N}$ each firm solves

$$\max_{w} \pi(w) = (1 - w_i)\left(1 - (1 - \theta)^N\right)$$

$$\text{subject to } w\frac{1 - (1 - \theta)^N}{N\theta} = \bar{U}.$$

where the hiring probability is shown in 2.24. Solving for w from the constraint, $w = \frac{N\theta\bar{U}}{1 - (1-\theta)^N}$ and inserting this into the objective, we maximize

$$\max_{\theta} 1 - (1 - \theta)^N - N\theta\bar{U}.$$

From the first order condition we find that $\bar{U} = (1 - \theta)^{N-1}$ and that $\theta = 1 - \bar{U}^{1/(N-1)}$.

When all the firms are identical and take the market expected utility $\bar{U}$ as given, all choose to induce $\theta = 1 - \bar{U}^{1/(N-1)}$ by setting the wage w. Since all firms set the same wage, workers randomize uniformly among the set of firms and $\theta = 1/M$. This has the consequence that in equilibrium the market utility for workers equals

$$\bar{U} = \left(1 - \frac{1}{M}\right)^{N-1}.$$

For a fixed market tightness $k = \frac{M}{N}$

$$\bar{U} = \left(1 - \frac{1}{M}\right)^{M/k-1}.$$

Allowing the market size to become large

$$\lim_{M \to \infty} \bar{U} = e^{-1/k}. \tag{2.20}$$

Observe that $1/k$ is equal to the queue length of workers per job opening N/M. Similarly, the wage $w = \frac{N\theta\bar{U}}{1-(1-\theta)^N}$ approaches

$$\lim_{M \to \infty} w = \frac{1/k}{e^{1/k} - 1}, \tag{2.21}$$

the expected profits for a firm $\pi = (1 - (1 - \theta)^N) - \frac{N}{M}\bar{U}$ approaches

$$\lim_{M \to \infty} \pi = 1 - \left(1 + \frac{1}{k}\right)e^{-1/k}, \tag{2.22}$$

and the hiring probability for workers approaches

$$\lim_{M \to \infty} \Omega = \frac{1 - e^{-1/k}}{1/k}. \tag{2.23}$$

Notice that equation 2.23 is the ratio of the probability that a firm gets at least one applicant relative to the expected queue length.

Worker Behavior

Hiring Probability

The expression for the hiring probability Ω

$$\Omega(\theta) = \frac{1 - (1 - \theta)^N}{N\theta} \tag{2.24}$$

118

arises from the expected probability of getting hired in all the different possible configurations of applicants:

$$
\begin{aligned}
\Omega(\theta) &= \sum_{k=0}^{N-1} \binom{N-1}{k} \frac{1}{k+1} \theta^k (1-\theta)^{N-k-1} \\
&= \sum_{k=0}^{N-1} \frac{N-1!}{(N-k-1)!k!} \frac{1}{k+1} \theta^k (1-\theta)^{N-k-1} \\
&= \sum_{k=0}^{N-1} \frac{N-1!}{(N-k-1)!k+1!} \theta^k (1-\theta)^{N-k-1} \\
&= \frac{1}{N\theta} \sum_{k=0}^{N-1} \frac{N!}{(N-(k+1))!k+1!} \theta^{k+1} (1-\theta)^{N-k-1} \\
&= \frac{1}{N\theta} \sum_{l=1}^{N} \frac{N!}{(N-l)!l!} \theta^l (1-\theta)^{N-l} \\
&= \frac{1}{N\theta} \left(\sum_{l=0}^{n} \frac{N!}{(N-l)!l!} \theta^l (1-\theta)^{N-l} - (1-\theta)^N \right) \\
&= \frac{1-(1-\theta)^N}{N\theta}.
\end{aligned}
$$

Response to changes in the wage

We begin by considering a situation with M firms and N identical workers in which a single firm sets the wage w_1 and the remaining firms set the wage w_2. In this case, the application behavior of workers is summarized by two equations:

$$
\begin{aligned}
w_1 \Omega_1 - w_2 \Omega_2 &= 0 \\
\theta_1 + (M-1)\theta_2 &= 0
\end{aligned}
$$

where Ω_i is the probability of a worker being hired at a firm offering the wage w_i (and to which workers apply with probability θ_i). We refer to this system of equations as $\mathbf{F}$.

The main result from Peters (1984) guarantees that the application probabilities are well defined for any (non-zero) announcement of wages. We can therefore implicitly differentiate

with respect to the wages. For example, differentiating this system with respect to w_1 yields

$$\Omega_1 + w_1 \frac{\partial \Omega_1}{\partial \theta_1} \frac{\partial \theta_1}{\partial w_1} - w_2 \frac{\partial \Omega_2}{\partial \theta_2} \frac{\partial \theta_2}{\partial w_1} = 0$$

$$\frac{\partial \theta_1}{\partial w_1} + (M-1) \frac{\partial \theta_2}{\partial w_1} = 0.$$

Notice that the terms of the form $w_i \frac{\partial \Omega_i}{\partial \theta_i}$ describe the payoff consequences for a worker of applying to firm i when the other workers increase their application to this firm. Since an increase in θ_i means that there is more congestion, this derivative will be negative. Going forward, we denote $\xi_i = w_i \frac{\partial \Omega_i}{\partial \theta_i}$.

Summarizing the implicit differentiation with respect to w_1 and w_2 is the equation

$$\mathbf{D_w F} + \mathbf{D_\theta F}\, \mathbf{D_w}\theta = \mathbf{0}$$

where

$$\mathbf{D_w F} = \begin{bmatrix} \Omega_1 & -\Omega_2 \\ 0 & 0 \end{bmatrix}, \quad \mathbf{D_\theta F} = \begin{bmatrix} \xi_1 & -\xi_2 \\ 1 & M-1 \end{bmatrix}, \quad \text{and} \quad \mathbf{D_w}\theta = \begin{bmatrix} \frac{\partial \theta_1}{\partial w_1} & \frac{\partial \theta_1}{\partial w_2} \\ \frac{\partial \theta_2}{\partial w_1} & \frac{\partial \theta_2}{\partial w_2} \end{bmatrix}.$$

From this equation it is then possible to solve for the $\frac{\partial \theta_i}{\partial w_j}$ terms from

$$\mathbf{D_w}\theta = -\mathbf{D_\theta F}^{-1}\, \mathbf{D_w F}$$

where

$$\mathbf{D_\theta F}^{-1} = \frac{1}{(M-1)\xi_1 + \xi_2} \begin{bmatrix} M-1 & \xi_2 \\ -1 & \xi_1 \end{bmatrix}.$$

such that

$$\begin{bmatrix} \frac{\partial \theta_1}{\partial w_1} & \frac{\partial \theta_1}{\partial w_2} \\ \frac{\partial \theta_2}{\partial w_1} & \frac{\partial \theta_2}{\partial w_2} \end{bmatrix} = -\frac{1}{(M-1)\xi_1 + \xi_2} \begin{bmatrix} M-1 & \xi_2 \\ -1 & \xi_1 \end{bmatrix} \begin{bmatrix} \Omega_1 & -\Omega_2 \\ 0 & 0 \end{bmatrix}.$$

Since we are typically interested in how applications change when a single "deviant" firm adjusts its wage, the two derivatives $\frac{\partial \theta_1}{\partial w_1}$ and $\frac{\partial \theta_2}{\partial w_1}$ are of prime importance:

$$\frac{\partial \theta_1}{\partial w_1} = -\frac{M-1}{(M-1)\xi_1 + \xi_2}\Omega_1$$

$$\frac{\partial \theta_2}{\partial w_1} = \frac{1}{(M-1)\xi_1 + \xi_2}\Omega_1.$$

Lemma 3 from Galenianos and Kircher (2012) ensures furthermore that $\frac{\partial \theta_i}{\partial w_i}$ is quasi-concave for any set of wage announcements.

Wage

When $M - 1$ firms are setting the wage w, the optimal response for a firm d is the wage w_d dictated by the first-order condition

$$\frac{\partial \pi_d}{\partial w_d} = -(1 - (1 - \theta_d)^N) + (1 - w_d)N(1 - \theta_d)^{N-1}\frac{\partial \theta_d}{\partial w_d} = 0.$$

Plugging in from above for $\frac{\partial \theta_d}{\partial w_d}$,

$$(1 - (1 - \theta_d)^N) = -(1 - w_d)N(1 - \theta_d)^{N-1}\frac{M - 1}{(M - 1)\xi_d + \xi_w}\Omega_d.$$

In a symmetric equilibrium, all firms set the same wage w^* and $\theta_d = \theta_w = 1/M$. This means that the wage can be solved from

$$(1 - (1 - 1/M)^N) = -(1 - w^*)N(1 - 1/M)^{N-1}\frac{M - 1}{Mw^*\xi}\Omega.$$

Rearranging and simplifying, the wage when there are M identical firms and N identical workers is the function

$$w^* = w(M, N) = \frac{1}{1 - \frac{M}{M-1}\left(1 - \frac{\Omega}{\left(1 - \frac{1}{M}\right)^{N-1}}\right)} \tag{2.25}$$

where $\Omega = \frac{1 - \left(1 - \frac{1}{M}\right)^N}{\frac{N}{M}}$. Algebraic manipulations reconcile expression 2.25 to that found in Burdett et al. (2001).

Table 2.A.1 presents wages for a large number of combinations of N and M. table 2.A.2 presents for the same combinations of N and M the profits associated with $w(N, M)$.

Table 2.A.1: Wages, $w(N, M)$

N	\multicolumn{9}{c}{M}								
	2	3	4	5	6	7	8	9	10
2	0.500	0.727	0.818	0.865	0.893	0.911	0.925	0.934	0.942
3	0.273	0.533	0.669	0.747	0.796	0.830	0.854	0.872	0.887
4	0.154	0.393	0.547	0.644	0.709	0.755	0.788	0.814	0.834
5	0.088	0.289	0.447	0.555	0.631	0.685	0.726	0.758	0.784
6	0.050	0.213	0.364	0.477	0.560	0.622	0.669	0.706	0.736
7	0.028	0.156	0.296	0.410	0.497	0.563	0.616	0.657	0.691
8	0.016	0.115	0.240	0.351	0.440	0.510	0.566	0.611	0.648
9	0.009	0.084	0.195	0.300	0.389	0.461	0.520	0.568	0.608
10	0.005	0.061	0.157	0.256	0.343	0.416	0.477	0.527	0.569

Table 2.A.2: Profits $w(N, M)$

N	\multicolumn{9}{c}{M}								
2	0.375	0.152	0.080	0.049	0.033	0.024	0.018	0.014	0.011
3	0.636	0.328	0.191	0.123	0.086	0.063	0.048	0.038	0.031
4	0.793	0.487	0.310	0.210	0.151	0.113	0.088	0.070	0.057
5	0.884	0.617	0.422	0.299	0.221	0.169	0.133	0.108	0.089
6	0.935	0.718	0.523	0.386	0.293	0.228	0.182	0.149	0.124
7	0.964	0.794	0.610	0.467	0.363	0.288	0.234	0.193	0.161
8	0.980	0.851	0.684	0.540	0.430	0.347	0.285	0.237	0.200
9	0.989	0.893	0.745	0.606	0.493	0.405	0.336	0.283	0.240
10	0.994	0.923	0.796	0.664	0.551	0.459	0.386	0.327	0.281

2.B   Model

Black Worker Behavior

When there are two wages that attract black applications, w_1 and w_2 set by l and k number of firms respectively, γ_1 and γ_2 are determined by the equations

$$w_1 \Gamma_1 - w_2 \Gamma_2 = 0$$
$$l\gamma_1 + k\gamma_2 = 0$$

We refer to this as system $\mathbf{Q}$. Because we are most interested in the case with a single deviant firm, we illustrate the case when $l = 1$ and $M = l + k - 1$.

Implicitly differentiating with respect to w_1 yields terms of the form

$$\Gamma_1 + \eta_1 \frac{\partial \theta_1}{\partial w_1} + \zeta_1 \frac{\partial \gamma_1}{\partial w_1} - \eta_2 \frac{\partial \theta_2}{\partial w_1} - \zeta_2 \frac{\partial \gamma_2}{\partial w_1} = 0$$

$$\frac{\partial \gamma_1}{\partial w_1} + (M - 1)\frac{\partial \gamma_2}{\partial w_1} = 0$$

where

$$\zeta_i = w_i \frac{\partial \Gamma_i}{\partial \gamma_i} \tag{2.26}$$

and

$$\eta_i = w_i \frac{\partial \Gamma_i}{\partial \theta_i}$$
$$= -w_i N (1 - \theta_i)^{N-1} \frac{1 - (1 - \gamma_i)^B}{\gamma_i B}$$
$$= -w_i \frac{N}{1 - \theta_i} \Gamma_i \tag{2.27}$$

The η terms denote the payoff impact of white workers on black workers and the ζ terms denote the payoff impact of black workers on other black workers. Notice that the Γ and ζ terms are analogous to Ω and ξ for white workers. What is new relative to the case with white workers are the η terms that are weighted by the $\frac{\partial \theta}{\partial w}$. These terms indicate the additional (negative) effect that white workers have on black payoffs. Organizing the results as a matrix equation

$$\begin{bmatrix} \Gamma_1 & -\Gamma_2 \\ 0 & 0 \end{bmatrix} + \begin{bmatrix} \eta_1 & -\eta_2 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \frac{\partial \theta_1}{\partial w_1} & \frac{\partial \theta_1}{\partial w_2} \\ \frac{\partial \theta_2}{\partial w_1} & \frac{\partial \theta_2}{\partial w_2} \end{bmatrix} + \begin{bmatrix} \zeta_1 & -\zeta_2 \\ 1 & M-1 \end{bmatrix} \begin{bmatrix} \frac{\partial \gamma_1}{\partial w_1} & \frac{\partial \gamma_1}{\partial w_2} \\ \frac{\partial \gamma_2}{\partial w_1} & \frac{\partial \gamma_2}{\partial w_2} \end{bmatrix} = \mathbf{0}$$

we can solve for the derivative of interest from

$$\begin{bmatrix} \frac{\partial \gamma_1}{\partial w_1} & \frac{\partial \gamma_1}{\partial w_2} \\ \frac{\partial \gamma_2}{\partial w_1} & \frac{\partial \gamma_2}{\partial w_2} \end{bmatrix} = -\frac{1}{(M-1)\zeta_1 + \zeta_2} \begin{bmatrix} M-1 & \zeta_2 \\ -1 & \zeta_1 \end{bmatrix} \left(\begin{bmatrix} \Gamma_1 & -\Gamma_2 \\ 0 & 0 \end{bmatrix} + \begin{bmatrix} \eta_1 & -\eta_2 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \frac{\partial \theta_1}{\partial w_1} & \frac{\partial \theta_1}{\partial w_2} \\ \frac{\partial \theta_2}{\partial w_1} & \frac{\partial \theta_2}{\partial w_2} \end{bmatrix} \right)$$
$$= -\frac{1}{(M-1)\zeta_1 + \zeta_2} \begin{bmatrix} M-1 & \zeta_2 \\ -1 & \zeta_1 \end{bmatrix} \begin{bmatrix} \Gamma_1 + \eta_1 \frac{\partial \theta_1}{\partial w_1} - \eta_2 \frac{\partial \theta_2}{\partial w_1} & -\Gamma_2 + \eta_1 \frac{\partial \theta_1}{\partial w_2} - \eta_2 \frac{\partial \theta_2}{\partial w_2} \\ 0 & 0 \end{bmatrix}.$$

We can conclude that

$$\frac{\partial \gamma_1}{\partial w_1} = -\frac{M-1}{(M-1)\zeta_1 + \zeta_2}\left(\Gamma_1 + \eta_1 \frac{\partial \theta_1}{\partial w_1} - \eta_2 \frac{\partial \theta_2}{\partial w_1}\right).$$

The first term in the parenthesis describes how black workers respond to a change in the wage, not taking into account the change in the application behavior of whites. This is analogous to the Ω component of $\frac{\partial \theta}{\partial w}$. We refer to this part as the direct wage effect since it implies that a higher wage increases black applications. Moderating the direct wage effect are the two η terms. These two terms represent the reduction in payoffs at w_1 associated with white workers.

2.C Results

No Single Wage Pooling Equilibrium

Figure 2.C.1 shows two examples of pooling wages for which the derivative of profits is negative. The two different wages are picked to illustrate that the intuition holds both when the wage is low ($w_p = 0.10$) and when the wage is high ($w_p = 0.50$). In addition, we see from the right panel that when the wage is sufficiently high, it is never be optimal to increase the wage.

(a) $w_p = 0.10$ (b) $w_p = 0.50$
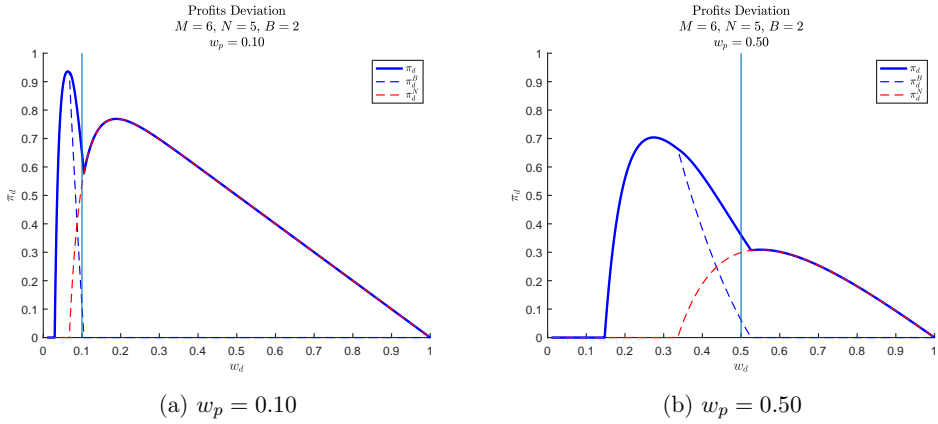
Figure 2.C.1: No Pooling Equilibrium

To demonstrate formally that there exists a profitable downward deviation, we consider a case when all M firms are setting a pooling wage w_p and then a single firm deviates to a wage w_d close to w_p. We aim to show that an increase in the wage is always associated with a loss of applicants when $w_d = w_P$. Recalling that

$$\frac{\partial \gamma_d}{\partial w_d} = -\frac{M-1}{(M-1)\zeta_d + \zeta_p}\left(\Gamma_d + \eta_d \frac{\partial \theta_d}{\partial w_d} - \eta_p \frac{\partial \theta_p}{\partial w_d}\right)$$

and plugging this into the profit expression, it is sufficient to show that

$$(1 - w_d)\left(\frac{1}{1-\theta_d}N\frac{\partial \theta_d}{\partial w_d} - \frac{1}{1-\gamma_d}B\frac{M-1}{(M-1)\zeta_d + \zeta_w}\left(\Gamma_d + \eta_d\frac{\partial \theta_d}{\partial w_d} - \eta_p\frac{\partial \theta_p}{\partial w_d}\right)\right) < 0, \quad (2.28)$$

where

$$\frac{\partial \theta_d}{\partial w_d} = -\frac{M-1}{(M-1)\xi_d + \xi_p}\Omega_d \quad \text{and} \quad \frac{\partial \theta_p}{\partial w_d} = \frac{1}{(M-1)\xi_d + \xi_p}\Omega_d.$$

Because all firms set the same wage w in a pooling equilibrium, both worker types randomize uniformly and the various components simplify. In particular,

$$\frac{\partial \theta_d}{\partial w_d} = -\frac{M-1}{M}\frac{\Omega}{\xi}$$

$$\frac{\partial \theta_p}{\partial w_d} = \frac{1}{M}\frac{\Omega}{\xi}$$

and

$$\begin{aligned}
\frac{\partial \gamma_d}{\partial w_d} &= -\frac{M-1}{M}\frac{1}{\zeta}\left(\Gamma + \eta\frac{\partial \theta_d}{\partial w_d} - \eta\frac{\partial \theta_p}{\partial w_d}\right)\\
&= -\frac{M-1}{M}\frac{1}{\zeta}\left(\Gamma - \eta\frac{M-1}{M}\frac{\Omega}{\xi} - \eta\frac{1}{M}\frac{\Omega}{\xi}\right)\\
&= -\frac{M-1}{M}\left(\frac{\Gamma}{\zeta} - \frac{\eta}{\zeta}\frac{\Omega}{\xi}\right).
\end{aligned}$$

Plugging these into expression 2.28

$$-\frac{1}{1-\frac{1}{M}}N\frac{M-1}{M}\frac{\Omega}{\xi} - \frac{1}{1-\frac{1}{M}}B\frac{M-1}{M}\frac{1}{\zeta}\left(\Gamma - \eta\frac{\Omega}{\xi}\right) < 0$$

$$-N\frac{\Omega}{\xi} - B\frac{1}{\zeta}\left(\Gamma - \eta\frac{\Omega}{\xi}\right) < 0.$$

This yields the condition for no pooling

$$-N\frac{\Omega}{\xi} - B\frac{\Gamma}{\zeta} < -B\frac{\eta}{\zeta}\frac{\Omega}{\xi}.$$

The left hand side is the direct wage effect on white and black workers. The right hand side is the negative impact of white workers on black workers. If the loss of black workers due to the presence of whites is greater than the gain in (both) types of workers because of the increased wage, then a wage increase is unprofitable. Manipulating the condition and

plugging in for η,

$$N\zeta\Omega + B\xi\Gamma > B\eta\Omega$$

$$N\zeta\Omega + B\xi\Gamma > -Bw\frac{N}{1 - \frac{1}{M}}\Gamma\Omega$$

$$N\frac{\zeta}{w\Gamma} + B\frac{\xi}{w\Omega} > -BN\frac{M}{M - 1}.$$

Notice that the wage cancels out of this expression as there is a w term in both ξ and ζ. Next, plugging in for $\frac{\zeta}{w\Gamma}$ and $\frac{\xi}{w\theta}$ yields

$$N\left(\frac{(1 - 1/M)^{B-1}}{\frac{1-(1-1/M)^B}{B/M}} - 1\right) + B\left(\frac{(1 - 1/M)^{N-1}}{\frac{1-(1-1/M)^N}{N/M}} - 1\right) > -\frac{BN}{M - 1}$$

After some more algebra,

$$\left(\frac{(1 - 1/M)^{B-1}}{1 - (1 - 1/M)^B} - \frac{M}{B}\right) + \left(\frac{(1 - 1/M)^{N-1}}{1 - (1 - 1/M)^N} - \frac{M}{N}\right) > -\frac{M}{M - 1}$$

$$1 + \frac{(1 - 1/M)^B}{1 - (1 - 1/M)^B} + \frac{(1 - 1/M)^N}{1 - (1 - 1/M)^N} > \frac{M - 1}{B} + \frac{M - 1}{N}.$$

The last condition holds because $\frac{1}{2} > \frac{M-1}{N} - \frac{(1-1/M)^N}{1-(1-1/M)^N}$ for any N and M.

Two Wage Segregated Equilibrium

White Expected Income and Low Wages, $w_l = \Omega w_h$

Table 2.C.1 shows the low wages associated with the white expected income when there M firms and N white workers:

Table 2.C.1: Low Wages

| N | \multicolumn{9}{c}{M} |
	2	3	4	5	6	7	8	9	10
2	0.375	0.606	0.716	0.778	0.819	0.846	0.867	0.882	0.895
3	0.159	0.375	0.516	0.608	0.671	0.717	0.752	0.779	0.801
4	0.072	0.236	0.374	0.476	0.551	0.608	0.652	0.688	0.717
5	0.034	0.151	0.273	0.373	0.453	0.516	0.566	0.608	0.642
6	0.016	0.097	0.200	0.293	0.372	0.438	0.492	0.537	0.575
7	0.008	0.063	0.147	0.231	0.307	0.372	0.427	0.475	0.515
8	0.004	0.041	0.108	0.183	0.253	0.316	0.371	0.420	0.462
9	0.002	0.027	0.080	0.144	0.209	0.269	0.323	0.371	0.414
10	0.001	0.018	0.059	0.114	0.173	0.229	0.281	0.328	0.371

Sufficient Competition

We show that $\frac{\partial \theta_w}{\partial w}$ is bounded when a single firm sets the wage w, $M-1$ firms set the wage w_l, and $\theta_w \to 1$ and $\theta_l \to 0$. In this case,

$$\frac{\partial \theta_w}{\partial w} = -\frac{M-1}{(M-1)\xi_w + \xi_l}\Omega_w.$$

When $\theta_w \to 1$ and $\theta_l \to 0$, each component has a simple form:

- $\Omega_w \to \frac{1}{N}$.

- $\xi_w = -w\frac{1}{\theta_w}\left((1-\theta_w)^{N-1} - \Omega_w\right) \to \frac{w}{N}$.

- $\xi_l = -w_l\frac{1}{\theta_l}\left((1-\theta_l)^{N-1} - \Omega_l\right) \to -w_l\frac{N-1}{2} = \frac{w}{N}\frac{N-1}{2}$.

Plugging these into the derivative yields

$$\frac{\partial \theta_w}{\partial w} = \frac{M-1}{\frac{w}{N}\left((M-1) + \frac{N-1}{2}\right)} \frac{1}{2}$$
$$= \frac{1}{w} \frac{M-1}{M + \frac{N}{2} - \frac{3}{2}}.$$

For a given w, the last expression is a bounded positive number.

Analogous computations for black workers deliver the same result.

Critical Point at w_l

Figure 2.C.2 illustrates that there is a critical point at w_l when w_l is equal to the white expected income. The left panel indicates a situation in which there are relatively few firms competing ($x = 2$ firms out of $M = 6$ total) for white workers and the right panel indicates a situation in which relatively many firms are competing for white workers ($x = 4$ firms out of $M = 6$ total). This highlights the fact that what is important for there to be a critical point is that the wage to which black workers apply is exactly at the point at which white workers cease to apply.
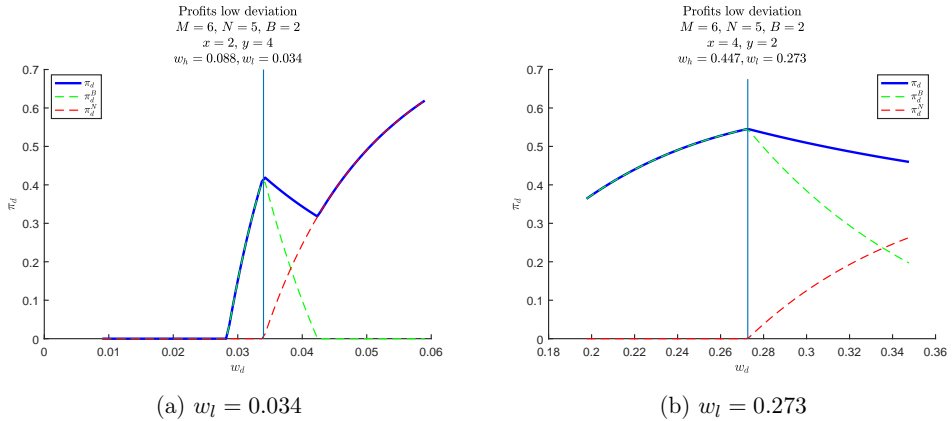

(a) $w_l = 0.034$ (b) $w_l = 0.273$

Figure 2.C.2: Critical Point at w_l

To complete the argument from the main text that there is a critical point at w_l, we demonstrate that

$$(1 - \gamma_d)N\frac{\partial \theta_d}{\partial w_d} + (1 - \theta_d)B\frac{\partial \gamma_d}{\partial w_d} < 0.$$

129

For w_d close enough to w_l, θ_d is close to 0. We can exploit this to simplify the problem. To begin with, observe that no white workers apply to the low wage w_l. As a consequence, $\frac{\partial \theta_l}{\partial w_d} = 0$ and

$$\frac{\partial \gamma_d}{\partial w_d} = -\beta \left(\Gamma_d + \eta_d \frac{\partial \theta_d}{\partial w_d} \right)$$
$$= -\beta \left(\Gamma_d - w_d \frac{N}{1 - \theta_d} \Gamma_d \frac{\partial \theta_d}{\partial w_d} \right),$$

where $\beta = \frac{(y-1)}{(y-1)\zeta_d + \zeta_l}$. Inserting this into the expression for change in the number of applicants, and letting $\theta_d \to 0$, we get

$$(1 - \gamma_d) N \frac{\partial \theta_d}{\partial w_d} - B\beta \Gamma_d + w_d B\beta \Gamma_d N \frac{\partial \theta_d}{\partial w_d} < 0$$

This can rearranged to the condition:

$$-\frac{1 - \gamma_d}{B\beta \Gamma_d} + \frac{1}{N \frac{\partial \theta_d}{\partial w_d}} < w_d. \tag{2.29}$$

Note the direction of the inequality doesn't change because $-\beta > 0$. Observe in addition that the denominator in each term is a function of the change in the expected number of applicants associated with the direct effect of the wage.

Next, we plug in for $\frac{\partial \theta_d}{\partial w_d}$, and use limit results for $\theta_d \to 0$. Consider when $\theta_d = 0$. Then we have the results

$$\xi_d = -w_d \frac{N - 1}{2}$$
$$\xi_h = -w_h \frac{1}{\theta_h} \left((1 - \theta_h)^{N-1} - \Omega_h \right)$$
$$= -x \frac{w_d}{\Omega_h} \left((1 - \theta_h)^{N-1} - \Omega_h \right)$$

and

$$\frac{\partial \theta_d}{\partial w_d} = -\frac{x}{x\xi_d + \xi_h}\Omega_d$$

$$= -\frac{x}{x\xi_d + \xi_h}$$

$$= \frac{1}{w_d}\frac{1}{\frac{N-1}{2} + \left(\frac{(1-\theta_h)^{N-1}}{\Omega_h} - 1\right)}$$

$$= \frac{1}{w_d}\frac{1}{\frac{N-3}{2} + \frac{(1-\theta_h)^{N-1}}{\Omega_h}}.$$

The first follows from the limit of ξ as θ goes to zero; the second follows from $w_d = w_l = \Omega_h w_h$ (so that $w_h = \frac{w_d}{\Omega_h}$); the third follows from the fact that $\Omega_d = 1$ if $\theta_d = 0$ combined with the two previous results. Plugging these in gives for the first term in equation 2.29:

$$-\frac{1}{N\frac{\partial \theta_d}{\partial w_d}} = w_d\frac{\left(\frac{N-3}{2} + \frac{(1-\theta_h)^{N-1}}{\Omega_h}\right)}{N}.$$

Addressing the second term in equation 2.29 in a similar fashion, we have the results

$$\zeta = -w_d\frac{1}{\gamma_d}\left((1-\gamma_d)^{B-1} - \Gamma_d\right)$$

$$= -w_d y\left((1-\gamma_d)^{B-1} - \Gamma_d\right)$$

and

$$\beta = \frac{y-1}{y}\frac{1}{\zeta}$$

$$= -\frac{y-1}{y}\frac{1}{w_d y\left((1-\gamma_d)^{B-1} - \Gamma_d\right)}$$

which gives

$$-\frac{1-\gamma_d}{B\beta\Gamma_d} = w_d\frac{y-1}{y}\frac{y^2}{y-1}\frac{\left((1-\gamma_d)^{B-1} - \Gamma_d\right)}{B\Gamma_d}$$

$$= w_d y\frac{\left(\frac{(1-\gamma_d)^{B-1}}{\Gamma_d} - 1\right)}{B}$$

Plugging in for both terms, 2.29 is

$$w_d \frac{\left(\frac{N-3}{2} + \frac{(1-\theta_h)^{N-1}}{\Omega_h} \right)}{N} + (1-\gamma_d) w_d y \frac{\left(\frac{(1-\gamma_d)^{B-1}}{\Gamma_d} - 1 \right)}{B} < w_d$$

$$w_d \frac{\left(\frac{N-3}{2} + \frac{(1-\theta_h)^{N-1}}{\Omega_h} \right)}{N} + (1-\gamma_d) w_d y \frac{\left(\frac{(1-\gamma_d)^{B-1}}{\Gamma_d} - 1 \right)}{B} < w_d$$

$$\frac{1}{N} \left(\frac{N-3}{2} + \frac{N\theta_h(1-\theta_h)^{N-1}}{1-(1-\theta_h)^N} \right) + (1-\gamma_d) \frac{y}{B} \left(\frac{B\gamma_d(1-\gamma_d)^{B-1}}{1-(1-\gamma_d)^B} - 1 \right) < 1$$

$$\frac{\theta_h(1-\theta_h)^{N-1}}{1-(1-\theta_h)^N} + (1-\gamma_d) \frac{\gamma_d(1-\gamma_d)^{B-1}}{1-(1-\gamma_d)^B} < 1 + B(1-\gamma_d) - \frac{N-3}{2N}$$

$$\frac{1}{x} \frac{(1-\frac{1}{x})^{N-1}}{1-(1-\frac{1}{x})^N} + \frac{y-1}{y} \frac{1}{y} \frac{(1-\frac{1}{y})^{B-1}}{1-\left(1-\frac{1}{y}\right)^B} < 1 + B\frac{y-1}{y} - \frac{N-3}{2N}$$

This condition is always satisfied because the left hand side cannot be larger than $\frac{1}{N} + \frac{1}{B}$ while the right hand side cannot be smaller than $\frac{N+3}{2N}$ which is the same as $\frac{1}{B} < \frac{1}{2} + \frac{1}{N}$.

Example $M = 6$, $N = 5$, $B = 2$

Table 2.C.2

| | | $M = 6, N = 5, B = 2$ | | | | |
| | | $w_h = 0.289, w_l = 0.151$ | | | | |
x	π_h	π_l	θ_h	θ_l	γ_h	γ_l
1	0.676	0.544	0.454	0.109	0.000	0.200
2	0.652	0.487	0.392	0.054	0.000	0.250
3	0.617	0.472	0.333	0.000	0.000	0.333
4	0.542	0.637	0.250	0.000	0.000	0.500
5	0.499	0.805	0.200	0.000	0.046	0.771

Table 2.C.3

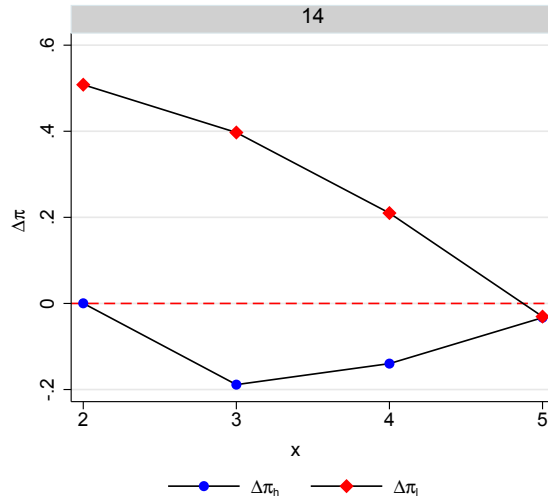| | | $M = 6, N = 5, B = 2$ | | | | |
| | | $w_h = 0.447, w_l = 0.273$ | | | | |
x	π_h	π_l	θ_h	θ_l	γ_h	γ_l
1	0.503	0.487	0.382	0.124	0.000	0.200
2	0.482	0.460	0.336	0.082	0.000	0.250
3	0.455	0.464	0.293	0.041	0.000	0.333
4	0.422	0.546	0.250	0.000	0.000	0.500
5	0.377	0.724	0.200	0.000	0.013	0.933

Existence of Segregated Equilibrium

Figure 2.C.3: Existence of Segregated Equilibrium

Chapter 3.

Discrimination in Small Markets with Directed Search: Part II Experiment

Abstract

Lang et al. (2005) develop a model of discrimination in labor markets with posted wages that predicts a segregated outcome in which white workers apply exclusively to a high wage and black workers apply exclusively to a low wage. We develop a simplified version of this game that predicts the same outcome. The key simplification is that we exogenously fix the wages that firms are allowed to set. We investigate this game using a controlled laboratory experiment and reproduce some of the postulated segregation effect. Firms learn to take advantage of discrimination and the experimental markets become more segregated over time. Preferred worker types apply almost exclusively to high wage firms and the income of discriminated workers is reduced by about 30% relative to the income of prioritized workers. Relative to treatments without discrimination, firms also increase their payoffs by about 25%. Although we do not reproduce the complete segregation predicted by the theory, we show that the presence of discrimination affects both worker search and firm wage-setting.

Author: Knut-Eric N. Joslin [1]

Keywords: Discrimination, Search, Coordination, Experimental Economics.
JEL Classification: C91, J64, J71 .

3.1 Introduction

Contemporary models of the labor market emphasize the role of frictions in the hiring process. In such a setting, discriminatory hiring practices will impact not just who gets hired at a given job but also the pattern of applications across jobs. In particular, we can expect that workers will avoid applying to jobs at which they are systematically discriminated. For example, if employers find white workers more desirable than black workers, then black workers have an incentive to avoid applying for the same positions as whites. Furthermore, firms may anticipate this response and exploit it to offer black workers low wages. Discriminatory hiring thus creates a tendency for segregation.

Lang et al. (2005) formalize this intuition in a posted wage model of discrimination. They represent discrimination in a simple fashion: If a white and a black worker apply to the same job vacancy, then the white worker is hired. The result is an equilibrium in which the market self-segregates. A subset of firms offers a high wage to which only white workers apply while another subset of firms offers a low wage to which only black workers apply. Blacks are so sensitive to the presence of white workers that the black wage can be pushed down to the level at which it is unattractive for white workers to apply. This generates a substantial wage gap between the worker types that is unrelated to productivity differences. It also leads to higher profits for firms. The fact that firms earn higher profits is remarkable because it may help explain why discrimination is a durable phenomenon.

In this study, we investigate a small market analog of the model of Lang et al. (2005) using a laboratory experiment. The motivation is to test whether a segregated outcome of the type predicted in their model emerges as a stable outcome. The value of an experiment is that it enables us to directly observe how discriminatory hiring affects how workers search for jobs and how firms set wages. The results from this study provide insights into how discrimination functions and can complement the interpretation of real-world data. This may be of particular value given that self-segregation would imply that discrimination is impossible to observe at the level of an individual firm.

The key difference between the model we take to the lab and the model of Lang et al. (2005) is that we restrict firms to pick between two wages, a high wage and a low wage. Although this simplifies the environment it does not change the basic implications for behavior. Discriminated workers should favor low wages at which they have a high probability of getting hired and this creates an incentive for firms to sort into distinct submarkets.

With respect to the experimental literature, this study most closely resembles Helland

et al. (2017) and Kloosterman (2016) who test a variety of small market directed search models. However, because firms make a choice between just two wages, the wage-setting portion of the model we take to the lab also has similarities with binary choice games. In the same fashion as in market entry games, the model predicts a sorting of firms into the high wage and low wage submarkets such that the expected profits are equalized. Hence, our model also exhibits the same multiplicity of equilibria.

The study is organized around two treatments, one in which all the workers are homogeneous and another in which there are two worker types, a prioritized white worker type and a discriminated black worker type. In both treatments, firms choose to either set a high wage or a low wage. The wage announcements are then observed by workers before they decide where to apply. The only difference between the treatments is that there is discriminatory hiring in the treatment with two worker types. In that treatment, firms that receive applications from both white and black workers must select a white worker.

For the parameters in our experimental implementation, the model predicts substantial treatment differences. The prediction in the treatment with homogeneous workers is for all the firms to set the high wage. In contrast, the prediction in the treatment with discrimination is for the majority of firms to set the low wage. Although the analysis of the treatment with discrimination is complicated by the presence of multiple equilibria, the anticipated treatment differences are large. The pure strategy equilibria in the discrimination treatment are characterized by two high wage firms and three low wage firms while the mixed strategy equilibrium predicts slightly more than three low wage firms on average.[1]

With respect to worker behavior, we find that theory effectively predicts worker behavior for most wage announcements. Our results thus corroborate the basic finding from other experimental studies of directed search models (Helland et al. 2017; Kloosterman 2016).

Overall, we find some support for the predictions of the model. In the presence of discrimination, firms set low wages and discriminated workers apply to these wages while the prioritized workers do not. Relative to the case without discrimination, firms earn higher profits and discriminated workers have lower income.

However, the treatment differences are more moderate than anticipated and insignificant when compared using non-parametric tests. Discriminated workers persistently apply to the high wages and firms in the treatment with discrimination do not set the low wage often enough. With respect to payoffs, the excess number of high wage firms reduces the payoffs of

---

[1]With respect to the mixed strategy equilibrium, the model predicts four or more high wage firms in only about 1.5% of instances.

firms below the predicted level with the consequence that both worker types end up better off. Of the equilibrium solutions, behavior in about half of the sessions coincides with the pure strategy equilibrium. However, the remaining sessions do not appear to coincide with any equilibrium play or even collusive play. We conclude that the presence of discrimination affects both worker search and firm wage-setting but that the overall impact is more moderate than anticipated.

The remainder of the paper is organized as follows: We begin with a general presentation of the model. This includes a discussion of both pure and mixed strategy equilibria. In the third section, we present the experimental parameters and the associated predictions. The fourth section goes over the design and the details of the laboratory implementation. The fifth section presents the results, which we divide into a section on the overall treatment differences, a section on worker behavior, and a section on firm behavior.

3.2   Model

The game that we take to the lab is a posted-wage market populated by M firms, N **prioritized** workers, and B **discriminated** workers. As is standard in the directed search setting, firms post wages and workers observe the available wages before workers make their application decision. The timing of the game is summarized below.

1. **Wage posting**: M firms independently choose to post the high wage w_h or a low wage w_l. Let x denote the number of firms that set the high wage and y denote the number of firms that set the low wage, where $y = M - x$.

2. **Application**: Workers observe the full set of wages and independently apply to *one* job.

3. **Hiring**: Firms hire a single worker if one or more workers applies to their firm.

   - **Discrimination**: A firm can only hire a discriminated worker if no prioritized worker applies.

4. **Profits** Firms earn $100 - w$ if they hire; workers who get hired earn w. Firms that do not hire and workers that do not get hired earn zero.

This game adapts the labor market discrimination model studied by Lang et al. (2005) to the experimental setting. As in their setting, discrimination takes the form of an ordinal

ranking at the hiring stage. Discriminated workers are only hired if no prioritized worker applies to the same vacancy. To aid the exposition and for consistency with Lang et al. (2005), we refer to the prioritized workers as *white* workers and the discriminated workers as *black* workers.

Relative to the full model of Lang et al. (2005), there are two crucial differences. The first difference is that we exogenously fix the wages to two levels. The second difference is that we consider a setting in which there are only a small number of market participants.

In the model of Lang et al. (2005), the wages w_h and w_l are determined endogenously as a function of the ratio of white and black workers relative to firms. This feature is absent in the model we take to the lab. However, from the perspective of an individual firm, there is not much loss of consistency with the model. In the full model, an individual firm takes the market outcomes as given (specifically, the expected income for each worker type is treated parametrically) and best responds by also setting w_h or w_l—in effect making a decision between two wages. Restricting the wages clarifies the sorting mechanism by which discrimination occurs and makes the model appropriate for laboratory testing. What is crucial is that we preserve the incentives of black workers to avoid competition from white workers and that this enables some firms to offer low wages. By fixing the wages, we sharpen the identification of the sorting effect in which we are interested. From a technical standpoint, restricting the wages to two levels also makes the equilibrium analysis of the small market setting tractable and, for the parameters that we investigate, ensures that there is a pure strategy segregated equilibrium.

The more substantive difference relative to Lang et al. (2005) is that there is an equilibrium selection problem in the small market case. In the full model, firms sort into the two submarkets in such a fashion that firms are exactly indifferent between setting the high wage and low wage. Moreover, with a continuum of agents, whether an individual firm picks the high or low wage has no impact on the outcomes. In contrast, with a finite number of agents, the wage decision is strategic. Test subjects in the firm role must anticipate which wages the other firms will choose and there are a variety of possible equilibria in both pure and mixed strategies. We present these equilibria next. In a certain sense, what we address in the lab is whether the forces in the model induce the test subjects to coordinate on a particular segregated equilibrium. In the phrasing used earlier, we are interested to know whether segregation can arise as a persistent phenomenon.

3.2.1 Equilibrium

Because an equilibrium in the full (wage-posting) game depends on the worker subgame, we present the worker behavior first.

Worker application

A strategy for a worker is a set of application probabilities to the M firms. We maintain that firms are "anonymous" and can only be identified by the wage that they post. This means that the strategy for a worker can be summarized by the application probabilities to the two different wage levels, h and l. We also focus on symmetric worker strategies in which all workers of a given type use the same application strategy. This is a natural assumption given that workers do not have an institution by which they can coordinate their application behavior. As is standard in the literature, we limit our presentation to the case of payoff maximization.

Given the assumptions of anonymity and symmetry, the application behavior of white workers is described by θ_h and θ_l. θ_h denotes the probability that white workers apply to the high wage and θ_l denotes the probability that white workers apply to the low wage. The same objects for the black workers are denoted by γ_h and γ_l.

An equilibrium in the worker application subgame is a set of probabilities such that workers of a given type are indifferent among the set of wages to which they apply. We refer to this as the *market income property*. If this property were not satisfied, then it would be possible for a firm to increase their profits by adjusting their application behavior. But this type of adjustment would then serve to re-establish the indifference relationship. The strong implication of the *market income property* is that all firms that receive applications offer the maximum expected payoff.

Throughout, we use Ω_w to denote the probability that a white worker is hired at a wage w when white workers apply with a probability θ_w:

$$\Omega_w = \frac{1 - (1 - \theta_w)^N}{N\theta_w}.$$

Intuitively, this expression may be understood as the probability that a firm hires (in the numerator) divided by the expected number of applicants (in the denominator).[2] The market income property thus dictates that θ_h and θ_l equalize the expected payoff from applying to

---

[2] The derivation of the hiring probability may be found in appendix 2.A.

w_h and w_l in the sense that

$$w_h \frac{1 - (1 - \theta_h)^N}{N\theta_h} = w_l \frac{1 - (1 - \theta_l)^N}{N\theta_l} \tag{3.1}$$

subject to $x\theta_h + y\theta_l = 1$, unless $w_h \frac{1 - \left(1 - \frac{1}{x}\right)^N}{\frac{N}{x}} \geq w_l$ and it is better for white workers to only apply to the x firms that offer the high wage in which case $\theta_h = 1/x$, $\theta_l = 0$.

The black application probabilities are derived from a similar indifference condition, except that black workers adjust their application behavior to account for the presence of white workers. The presence of white workers introduces an additional term that adjusts for the fact that blacks are only hired if no whites show up, an event that occurs at wage w with probability $(1 - \theta_w)^N$. We denote the black hiring probability at a wage w by

$$\Gamma_w = (1 - \theta_w)^N \frac{1 - (1 - \gamma_w)^B}{B\gamma_w}.$$

Hence, γ_h and γ_l are such that

$$w_h (1 - \theta_h)^N \frac{1 - (1 - \gamma_h)^B}{B\gamma_h} = w_l (1 - \theta_l)^N \frac{1 - (1 - \gamma_l)^B}{B\gamma_l} \tag{3.2}$$

subject to $x\gamma_h + y\gamma_l = 1$. Again this holds unless it is strictly better for black workers to only apply to one of the wages. For example, if $w_h (1 - \theta_h)^N \leq w_l (1 - \theta_l)^N \frac{1 - (1 - \gamma_l)^B}{B\gamma_l}$ then black workers only apply to the low wage and randomize among those y low wage firms with probability $\gamma_y = 1/y$.

The term $(1 - \theta_h)^N$ that adjusts the hiring probability of black workers is perhaps the crucial term in the model. If θ_h is large enough, then equality 3.2 fails to hold because the black workers strictly prefer the low wages. This creates a segregating tendency in the market. We refer to this as the *segregation effect*. Anticipating that black workers will seek to avoid competition from white workers, a subset of firms can profitably offer lower wages. Black workers are willing to apply to the low wages because they can avoid competition from white workers who target high wage firms. This increase firm profits but reduces the expected income available to black workers.

Firm wage-setting

For any constellation of x firms setting the high wage and y firms setting the low wage, there are well-defined application probabilities for workers such that the profit associated

with each wage can be written

$$\pi_x = (1 - w_h) \left(1 - (1 - \theta_h)^N (1 - \gamma_h)^B\right)$$
$$\pi_l = (1 - w_l) \left(1 - (1 - \theta_l)^N (1 - \gamma_l)^B\right)$$

where θ_h, θ_l, γ_h and γ_l are functions of w_h, w_l and x. These profit functions are composed of two factors. The first is the part of the surplus earned if the firm hires. The second factor is the probability of hiring. Notice that the probability of hiring is the complementary probability of not receiving any applicants of either type. To simplify the experiment and to emphasize the pure discrimination effect, we assume that both worker types are equally productive. However, none of the results depend in a critical way on this assumption.

Pure Strategy Equilibria A pure strategy equilibrium in this game is a number of firms setting the high wage x and a number of firms setting the low wage y such that no firm would prefer to set the other wage. Of particular interest are pure strategy *segregated* equilibria in which the high wage firms attract only prioritized workers while the low wage firms attract only the discriminated workers. For the treatment that we conduct in which a segregated equilibrium is possible, this is the only possible type of pure strategy equilibrium. We present the details of the experimental implementation in the section on treatments.

For a combination of x and y firms to be a segregated equilibrium, it must be impossible for a firm to profitably deviate to the other wage. Let $\pi_h(x^*)$ and $\pi_l(x^*)$ denote the profits associated with the high and low wage when x^* firms set the high wage. x^* is a pure strategy equilibrium if there no profitable deviations:

$$\pi_h(x) > \pi_l(x - 1)$$
$$\pi_l(x) > \pi_h(x + 1).$$

These two conditions say respectively that no high wage firm would like to deviate and set the low wage, in which case there would be $x - 1$ which wage firms and the deviant earns $\pi_l(x - 1)$, and that no low wage firm would like to deviate and set the high wage, in which case there would be $x + 1$ high wage firms and the deviant earns $\pi_h(x + 1)$. A benefit of limiting the wages to two levels is that the pure strategy equilibria can be identified by direct comparison of the expected profits for various combinations of x and $y = M - x$.

From a theoretical standpoint, a critical feature of the pure strategy segregated equilibria is that any combination of x out of M firms is a legitimate equilibrium. Firms thus face

142

a coordination issue associated with which firms should set the high wage and which firms should set the low wage. We discuss this further in the context of the design.

Mixed Strategy Equilibrium   Mixed strategy equilibria are also possible in this setting. We limit the discussion to the symmetric case in which all firms randomize with the same probability. A mixed strategy equilibrium is characterized by a probability p such that every firm is *ex ante* indifferent between setting the high wage and the low wage:

$$(1 - w_h) \sum_{k=0}^{M-1} \binom{M-1}{k} p^k (1-p)^{M-1-k} \left(1 - (1 - \theta_h(k+1))^N (1 - \gamma_h(k+1))^B \right) = \quad (3.3)$$

$$(1 - w_l) \sum_{k=0}^{M-1} \binom{M-1}{k} p^k (1-p)^{M-1-k} \left(1 - (1 - \theta_l(k))^N (1 - \gamma_l(k))^B \right) \quad\quad (3.4)$$

Notice that such an equilibrium depends on the application behavior for every possible realization of k other firms setting the high wage. Although it is difficult to provide intuition for when p is high and when p is low, our choice of parameters ensures that there is a unique p.[3]

3.3   Treatments and Predictions

The experiment is organized around two treatments, a discrimination treatment and a control treatment. In both treatments, the game described in the previous section is played by 5 firms and 4 workers, with wages fixed to $w_h = 50$ and $w_l = 32.5$. The difference between the treatments is that the discrimination treatment features 2 prioritized workers and 2 discriminated workers whereas the control treatment makes no distinction between the worker types. Because the workers are undifferentiated in the control treatment, we refer to the control treatment as the *homogenous* treatment.

The levels of the wages were chosen to generate substantial differences in the number of firms setting each wage and the overall outcomes in the market. In the homogeneous worker treatment, the equilibrium of the model is such that all firms set the high wage, that is $x = 5$ and $y = 0$, and the four workers randomize uniformly. In contrast, in the

---

[3]The mixing probability depends in a complex way on the features of the market. For example, for a fixed w_l, p can vary non-monotonically as w_h is increased, first rising and then falling. Moreover, it appears in some cases that it is possible for there to exist two different p values, one associated with a high level of profits and one associated with a low level of profits.

anticipated segregated (pure strategy) equilibrium of the discrimination treatment, only two firms should set the high wage, such that $x = 2$ and $y = 3$. Prioritized workers should apply strictly to the two firms setting the high wage and the discriminated workers should apply strictly to the three firms setting the low wage. The difference between the number of firms setting the high wage is therefore an overall measure of the treatment difference. Only in the treatment with discrimination should it be possible for the test subjects in the firm role to offer the low wage. Hence, a lower wage level in the discrimination treatment provides evidence of the postulated segregation effect.

Although a difference in the aggregate level of the wages provides evidence of a segregation effect, a more difficult question is whether and to what extent the test subjects coordinate on a segregated equilibrium. This poses challenges from a design standpoint as well as in terms of analysis. To get the pure segregation that is possible in theory requires that test subjects in the firm role coordinate on a particular segregated equilibrium and that test subjects in the worker role randomize their applications in the fashion dictated by the market income property. Although it is unrealistic to expect these requirements to be fulfilled exactly, the question is whether they are satisfied to a sufficient degree that the pure strategy segregated Nash equilibrium is a useful representation of the market.

The pertinent issue is whether the test subjects are able to coordinate on one of the pure strategy segregated equilibria.[4] Mis-coordination is a possibility and there is an equilibrium selection challenge. To facilitate *tacit coordination*, we therefore had test subjects interact in the same markets for all 30 periods of the experiment, with each period consisting of a single play of the game described in detail below. Evidence from market entry games suggests that repeated play of this kind is conducive to coordination, even when players are anonymous and have minimal feedback (for a brief review, see Rapoport and Seale 2008).

The danger associated with repeated interaction is that the participants collude and, more generally, that the repeated game differs in some un-modelled respect from the game we aim to test. One type of collusion is coordination between firms and workers in which a specific worker always applies to a specific firm. This is inconsistent with the use of symmetric application strategies by workers. Another type of collusion is among firms. For example, the firms would benefit as a group if they only offered the low wage.[5]

Both types of collusion are made difficult by the design choices and the structure of the game. Collusion between firms and workers is made difficult by anonymity and the structure

---

[4]Since all firms should set the same wage, equilibrium selection is only relevant in the treatment with discrimination.

[5]Other types of collusion among the firms would involve rotating which firms set the high wage.

of wage-posting and hiring. When firms post wages, the wages are presented to workers in a random order and are not associated to a particular firm. Workers are therefore only able to target a wage level, not a specific firm. Similarly, when a firm receives applicants, the applicants are presented in a random order, sorted by type. This circumscribes opportunities for repeated interactions.

Collusion on wages among firms is perhaps more plausible. However, also this type of collusion would be difficult to sustain. Foremost, there is a significant gain from unilateral deviations. For example, if all the firms are setting the low wage, then a deviation to the high wage is rewarded by a substantial increase in the number of applicants. The increase in payoffs associated with such a deviation is sufficiently large that it outweighs many periods of sustained collusion. Because of anonymity and randomness, the scope for punishment is also limited. Furthermore, collusion only improves payoffs slightly relative to the segregated equilibrium. Collusion would need to be sustained for more than 10 periods to outweigh the benefit of a unilateral deviation (assuming a grim trigger punishment strategy). The forces driving test subjects toward the pure strategy equilibrium are thus strong whereas the those favoring collusion are weak.

Homogenous Treatment Table 3.1 shows the expected market outcomes for the homogeneous treatment. For instance, the bottom row shows the case when all firms set the high wage such that $y = 0$: The workers randomize equally among the high wage firms $\theta_h = 0.200$, and firms can expect to earn $\pi_h = 0.295$. As is apparent from the table, the ability of firms

Table 3.1: Homogeneous Treatment

y	θ_h	θ_l	π_h	π_l
5	–	0.200	–	0.399
4	0.426	0.143	0.446	0.312
3	0.369	0.087	0.421	0.206
2	0.313	0.031	0.389	0.079
1	0.250	0.000	0.342	0.000
0	0.200	–	0.295	–

to offer low wages in the treatment with homogeneous workers is limited by the application behavior of workers. Say that 4 firms offer the high wage while a single firm offers the low wage ($y = 1$). This is the case is summarized in the second row from the bottom. In this case, the low wage firm does not attract any applicants. The reason is that the workers

strictly prefer to apply to high wages even though there is congestion at those firms:

$$w_h \frac{1 - (1 - \theta_h)^N}{N\theta_h} = 0.5 \frac{1 - \left(1 - \frac{1}{4}\right)^4}{4\frac{1}{4}} = 0.342 > w_l = 0.325.$$

As a consequence, if the other firms set the high wage, then the best response of a firm is also to set the high wage and earn 0.295. This is therefore an equilibrium. What about other constellations? As we can see by comparing the expected profits, π_h and π_l in columns 4 and 5 of table 3.1, any firm that sets a low wage could always increase their profit by deviating to the high wage. For instance, if all 5 firms were setting the low wage then a firm could increase their profit from 0.399 to 0.446 by deviating to w_h. The only equilibrium is therefore one in which all 5 firms set the high wage. For the same reason, there are no mixed strategy equilibrium in this treatment.

Discrimination Treatment The analysis of the treatment with discrimination is somewhat more complex. Table 3.2 shows the worker application behavior when 0–5 firms set the low wage y: These application probabilities are derived from the optimal application

Table 3.2: Discrimination Treatment, Application Behavior

y	θ_h	θ_l	γ_h	γ_l
5	–	0.200	–	0.200
4	0.742	0.065	0.000	0.250
3	0.500	0.000	0.000	0.333
2	0.333	0.000	0.000	0.500
1	0.250	0.000	0.150	0.399
0	0.200	–	0.200	–

behavior of workers. To see that this is the case, consider the case when 1 firm is offering the high wage and 4 firms are offering the low wage. The expected payoff for a white worker that applies to the high wage w_h is

$$\mathbb{E}U_l^N = w_h \frac{1 - (1 - \theta_h)^N}{N\theta_h} = 0.5 \frac{1 - (1 - 0.742)^2}{2 * 0.742} = 0.315$$

and the same payoff when applying to the low wage w_l is

$$\mathbb{E}U_l^N = w_l \frac{1 - (1 - \theta_l)^N}{N\theta_l} = 0.325 \frac{1 - (1 - 0.065)^2}{2 * 0.065} = 0.315.$$

Notice that the expected payoff for whites are equal at both wages. This illustrates the market income property that governs worker behavior. For this constellation of firms, both wages thus attract applications from white workers. The expected payoff for black workers is computed in a similar fashion, except that black workers do not apply to the high wage and their payoffs from the low wage are reduced by the presence of white workers:

$$\mathbb{E}U_l^B = w_l(1 - \theta_h)^N \frac{1 - (1 - \gamma_l)^B}{B\gamma_h} = 0.325(1 - 0.065)^2 \frac{1 - (1 - 0.25)^2}{2 * 0.25} = 249.$$

Taking the application behavior of workers as given by table 3.2, what is critical for firms are the expected payoffs. Consider the case when 1 firm is offering the high wage and 4 firms are offering the low wage. The application probabilities are given by the second row in table 3.1: $\theta_h = 0.742$, $\theta_l = 0.065$, $\gamma_h = 0.000$, and $\gamma_h = 0.250$. Then a high wage firm can expect to earn

$$(1 - w_h)\left(1 - (1 - \theta_h)^N(1 - \gamma_h)^B)\right) = (1 - 0.5)\left(1 - (1 - 0.742)^2\right) = 0.467$$

and a low wage firm can expect to earn

$$(1 - w_l)\left(1 - (1 - \theta_l)^N(1 - \gamma_l)^B\right) = (1 - 0.325)\left(1 - (1 - 0.065)^2(1 - 0.25)^2\right) = 0.343.$$

We summarize the expected payoffs for the various levels of low wage firms y in table 3.2

Table 3.3

y	π_h	π_l	$\mathbb{E}U_h^N$	$\mathbb{E}U_l^N$	$\mathbb{E}U_h^B$	$\mathbb{E}U_l^B$
5	–	0.399	–	0.292	–	0.187
4	0.467	0.343	0.315	0.315	0.033	0.249
3	0.375	0.375	0.375	0.325	0.125	0.271
2	0.278	0.506	0.417	0.325	0.222	0.244
1	0.297	0.431	0.438	0.325	0.260	0.260
0	0.295	–	0.450	–	0.288	–

The number of high wage firms y that is consistent with a pure strategy equilibrium can be identified in this table. In the row corresponding to $y = 3$, we see that firms earn 0.375 from setting the high wage and 0.375 from setting the low wage. To see that there are no profitable deviations from this equilibrium, consider both possibilities. If a low wage firm were to deviate to the high wage, then there would be $y = 2$ low wage firms and the deviant

firm could expect to earn $\pi_h = 0.278$. No low wage firm would therefore like to adjust their wage. Similarly, if a high wage firm were to deviate to a low wage then there would be $y = 4$ low wage firms and the deviant firm could expect to earn $\pi_l = 0.343$. It is therefore also the case that no high wage firm would like to deviate to the low. This means that $x = 2$, $y = 3$ is an equilibrium.

Observe that the segregated outcome in the discrimination treatment predicts higher profits for firms than in the homogeneous treatment (0.375 vs. 0.295), and that the white expected income is substantially above the black expected income (0.375 vs. 0.271). In addition, the treatment is designed so that in the segregated equilibrium, there should be no benefit in terms of expected payoffs from setting the high or the low wage.

Besides the pure strategy segregated equilibrium, we also consider the mixed strategy equilibrium as governed by expression 3.3. For the market parameters given in this experiment, $p = 0.263$ and the expected profits are 0.376. This means that if all the test subjects in the firm role set the high wage with a probability 0.263, then every player will be indifferent between the high and low wage and can expect to earn 0.376. This translates to an expected number of high wage firms $x = 1.32$ and an expected number of low wage firms $y = 3.68$. That is, fewer high wage firms and more low wage firms than in the pure strategy equilibrium.

3.4 Design and Implementation

All participants in the experiment played in fixed groups of 9 individuals for the duration of the 30 periods of the experiment. It was made explicitly clear in the instructions that the groups were fixed and that there was no re-matching at any point. Players were randomly assigned to one of the roles at the beginning of the session—firm, prioritized worker, or discriminated worker—and these roles were maintained throughout the experiment. In total, 5 sessions of the homogeneous treatment were conducted and 6 sessions of the discrimination treatment were conducted. In total, 45 individuals participated in the homogeneous treatment and 54 individuals participated in the discrimination treatment. Prior to the experiment, 1 session of each treatment was conducted as a pilot study. Results in the pilot were largely consistent with segregated equilibrium. In addition, we conducted 2 pilot sessions in which workers faced exogenously given wages (that is, there were no players in the firm role). Results from those experiments were similar to the results for worker behavior presented in the next section.

In terms of implementation, we worried that a presentation of the experiment using a labor market framing might affect how test subjects chose to set wages. Given the strong negative connotations associated with discrimination, it seems plausible that test subjects might adjust their behavior to avoid sustaining discriminatory outcomes or unfair wage offers. To avoid this, we chose to instead re-present the experiment using a market context in which firms posted prices and buyers with unit demand chose a single seller to visit. In the treatment with discrimination, prioritized workers were assigned the role of "Buyer A" and discriminated workers were assigned the role of "Buyer B." Each buyer was provided with a fixed endowment of 100 that expired at the end of each period. It was made clear that there was no "roll-over" and that any part of the endowment not spent on a good was lost. In addition to these differences in the description of the environment, the wages were converted to prices such that the low wage 32.5 was instead a high price of 67.5.

Both treatments were conducted at the BI Norwegian Business School in a dedicated research lab. Test subjects were recruited from the population of university students in Oslo. Students from the business school represented the majority of the participants although a substantial number of students from other fields of study at the University of Oslo also attended. Recruitment and session management were handled via the ORSEE system and test subjects were only provided the opportunity to participate in a single session of the experiment (Greiner 2015). The experiment was programmed and administered via *z-Tree*.

Subjects were assigned to cubicles via lottery tickets as they arrived at the lab. Apart from signing a payment form, anonymity of the subjects was preserved throughout the experiment. In the case of over-recruitment, randomly selected students were excluded and paid a show-up fee of 100 Norwegian kroner (about 13 USD). Before starting the computerized portion of the experiment, printed instructions were distributed and read aloud. The intent was to achieve public knowledge of the rules and to establish that all participants were playing by the same rules. Participants were thus exposed to the instructions prior to knowing the role that they would have in the experiment. After allowing for brief factual clarification with regard to the instructions, the students were then asked to start their *z-Tree* client.

After a short instruction screen that summarized the key features of the instructions, test subjects were assigned their role. Although they kept this role throughout the experiment, they were reminded of their type at the beginning of each round of the game. The first two rounds were trial periods in which actions did not have payoff consequences. This provided an opportunity for test subjects to familiarize themselves with the interface. Each group then participated in 30 repetitions of the labor market game in which they could

earn experimental currency called *points*. In total, each session took about an hour. At the end of the experiment, the accumulated points for each player were converted to Norwegian kroner according to an exchange rate that was published in the instructions. On average, test subject earned a bit less than 250 NOK, about 30 USD. As test subjects exited the lab they signed a confirmation of payment and were paid in cash.

In each round of play, participants moved through three screens which varied somewhat depending on their type. The first screen that firms (in the language of the experiment, *sellers*) encountered was a price posting screen. This consisted of two vertically aligned buttons, one for the high wage and one for the low wage (implemented as the prices 50 and 67.5 in the experiment). The first screen encountered by workers (in the language of the experiment, *buyers*) was an application screen in which the wages were presented vertically in a randomly sorted order. There was as a result no way to target a particular firm. At best, test subjects could target a wage level.

The second screen encountered by test subjects represented the hiring stage. If a firm received applications, they selected which worker to hire from a randomly sorted vertical list labelled by worker type (buyer A and buyer B in the experiment). If white and black workers showed up, the firm player was only shown and allowed to pick from the list of white workers. The corresponding screen for workers informed them about how many of each type of worker applied to the same firm.

The final screen in each round provided information about play in the present round and rich feedback on the market level outcomes including how many of each type of buyer tried to buy at each wage. Also provided as feedback was a historical summary of which wages the firm (worker) had posted (applied to) along with whether they had hired (been hired) and how many of each type of worker also applied to the wage.

3.5 Results

We begin with a summary of market level outcomes. We identify some evidence of segregation, although this is not statistically significant.

3.5.1 Summary

Figure 3.1 shows the average number of firms setting the low wage by period. The initial level is similar in both treatments. However, as indicated by the vertical red dashed line, in

the second half of the periods the number of low wage firms is systematically higher in the treatment with discrimination (see figure 3.A.1 in the appendix).
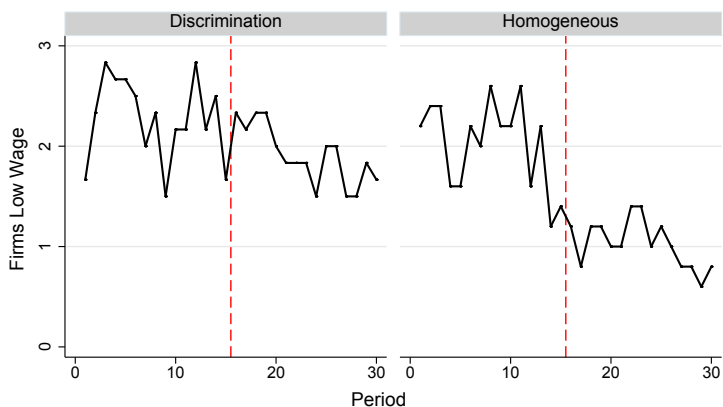

Figure 3.1: Number of Firms setting the Low Wage

The behavior in the early part of the experiment is consistent with an initial period of learning in which the test subjects familiarize themselves with the mechanics of the game and the behavior of other players. Play in both treatments is close to random and quite variable during the early periods. For this reason, we base our analysis on the second half of the data unless otherwise noted. Our interest is whether discrimination can emerge as firms coordinate their wage-setting and it would be surprising if the behavior early on was representative of the long run outcome.

Tables 3.1 and 3.2 summarize the market outcomes over the last 15 periods of the experiment. Table 3.1 shows the average number of firms that set the low wage (y) along with summary statistics. Table 3.2 shows the payoffs in the market, computed as the average for a given type of participant when setting/applying to a given wage. For example, white workers earned on average 44.1 from application to to the high wage h in the D treatment while firms in the same treatment earned on average 35.4 from setting the high wage. Also included in parentheses is how often workers of a given type applied to a given wage. For instance, white workers applied to the high wage with a relative frequency of 0.94.

Table 3.1: Firms Setting the Low Wage

	Mean y	Std. Dev.	Min	Max	Obs
H	1.03	1.25	0	4	75
D	1.91	1.25	0	4	90

Table 3.2: Payoffs

	π_h	π_l	$\mathbb{E}U_h^N$	$\mathbb{E}U_l^N$	$\mathbb{E}U_h^B$	$\mathbb{E}U_l^B$
H	33.1	26.3	37.3 (0.88)	27.1 (0.12)	–	–
D	35.4	33.0	44.1 (0.94)	32.5 (0.06)	26.4 (0.49)	26.4 (0.51)

Result 3.5.1 (Segregation Effect). *There is weak evidence for the segregation effect.*

Our key statistic for overall market behavior is the number of firms setting the low wage. About one more firm sets the low wage in the treatment with discrimination relative to the treatment without discrimination, 1.91 in D and 1.03 in H. This suggests that discrimination enables a subset of firms to offer lower wages. Also consistent with the expected segregated equilibrium is the finding that firms in the D treatment earn higher profits and that white workers apply much more often to the high wage firm than black workers (a relative frequency of 0.94 compared with 0.49).

Table 3.3

			y		
	0	1	2	3	4
H	46.7	28.0	6.7	13.3	5.3
D	16.7	24.4	17.8	33.3	7.8

However, the substantial variability in outcomes at the block level compromises an unambiguous interpretation. The variability is evident from the standard deviations in table 3.1 and the relative frequency of numbers of firms setting the low wage shown in table 3.3 (we show the block level results in the section on firm behavior). In particular, the frequency of $y = 0$ and $y = 1$ is far too high in the discrimination treatment. This is indicative of a large number of cases in which there is little or no segregation. As we discuss below, this is inconsistent with any equilibrium.

As a consequence of the variability in the observations, we cannot reject the null hypothe-

sis of no treatment difference using a non-parametric Wilcoxon rank-sum test (p-value=0.361 based on the block level averages presented in table 3.A.1).[6] Moreover, to get sufficient power would require a more than ten-fold increase in the number of observations.[7]

3.5.2 Worker Behavior

Figures 3.2 and 3.3 present the worker application behavior observed in the data (grey diamond) along with the worker application behavior anticipated by the model (black circle). The first figure shows the application behavior to low wage firms (θ_l, γ_l) plotted by the number of low wage firms y. The second figure shows the application behavior to high wage firms (θ_h, γ_l) plotted by the number of high wage firms. Omitted from these figures are the results when all firms set the same wage ($y = 0$ or $y = 5$). In those cases, the average application probabilities are mechanically in line with predictions and therefore do not reveal anything about search.
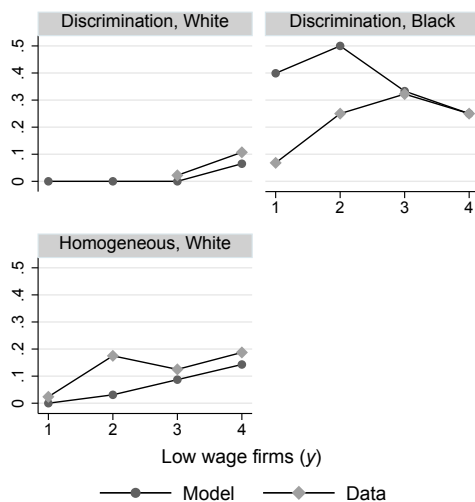

Figure 3.2

[6] Power computations were carried out using the *Stata* package developed by Bellemare et al. (2016).

[7] In a pilot study comprised of two observations from each treatment, the outcomes were essentially in line with the model. From a design and identification standpoint, the extreme outliers were therefore an unexpected and problematic finding.
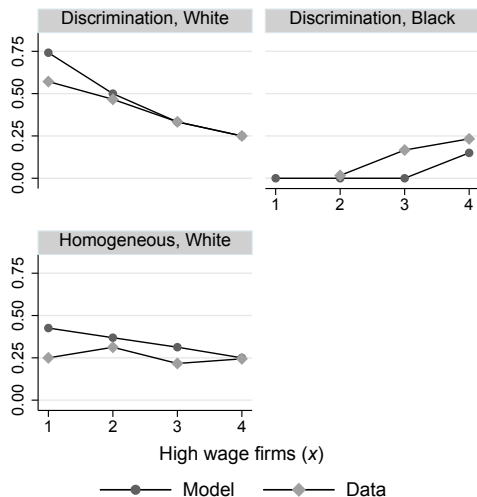
Figure 3.3

Result 3.5.2 (Worker Behavior). *(1) White worker behavior is consistent with the market income property. (2) Black workers apply to the high wage too often when there are few low wage firms.*

Overall, there is a high degree of congruence between the search behavior in the lab and that anticipated by theory. The application behavior for white workers in the discrimination treatment nearly overlaps with that predicted by theory. In the top left panel of 3.2, we see that white workers rarely apply to low wage firms. Exactly as predicted by theory, when $y = 1$ and $y = 2$ there are no instances in which white workers choose the low wage firm. When $y = 3$ and $y = 4$ the probabilities are also close to the theoretically predicted levels.

For black workers, there is also a tight fit with theory when there are relatively many low wage firms, as shown the top right panel of 3.2 ($y = 3$ and $y = 4$).[8] In these instances, blacks avoid applying to the high wage firms because the high wage firms are congested by the presence of white workers. When the number of firms is equal to or less than the number of white workers, black workers never apply to the high wages. Notably, this means that in the anticipated segregated equilibrium with $x = 2$ and $y = 3$, white and black workers behave nearly as theory would predict.

---

[8]The complementary presentation in terms of high wage firms is shown in figure 3.3.

The differences from predicted behavior arise when there are few firms of a given type. In the case of white workers, the white workers shift too much of their search to low wage firms when $x = 1$ (and $y = 4$). This may represent an effort to avoid competition. [9] The most prominent deviation is, however, for black workers when there are few low wage firms. As is shown in the top right panel of figure 3.2, blacks do not apply often enough to the low wage firms when $y = 1$ and $y = 2$. This means that blacks are instead applying to high wages. Black workers appear to mis-perceive the degree of congestion in the high market. When there are more high wage firms than white workers, this tempts black workers to begin to apply to the high wage. Digging into the data, the black workers appear to have benefited from some "luck of the draw" and were not punished as strongly as might be expected. Overall, blacks were hired at a slightly higher rate than anticipated when they applied to the high wage. In particular, out of 16 observations of $x = 3$, blacks were hired at the high wage 65 percent of the time relative to a prediction of 44.[10]

Also worth remarking on is the difference between treatments. The application behavior for white workers hews closer to predictions in the discrimination treatment than the corresponding behavior in the homogeneous treatment. An explanation is that white workers in the discrimination treatment have a more transparent decision problem. Because white workers only compete with a the single other white worker, it is relatively straightforward to compute expected payoffs in the discrimination treatment. Say, for example, that there are two high wage firms. If a white worker believes that white workers always apply to high wage firms conditional on $x = 2$, then it is not too difficult for the worker to realize that they will encounter another worker half the time at the high wage. This translates to an expected payoff from the high wage of $\frac{3}{4}50 = 37.5$. In contrast, in the treatment with homogeneous workers, each worker competes with 3 others. Even if a worker has correct beliefs about the behavior of the other workers, it can still be challenging to translate this into an expected payoff. For instance, if there are $x = 4$ high wage firms, then the probability of getting hired at a high wage firm when workers only apply to the high wages is $(1 - (1 - 0.25)^4)/(4 * 0.25) = 0.684$.

---

[9]There were only a few observations of $x = 1$ which makes the effect difficult to judge.

[10]Using data from all periods, these numbers are more in line with the predictions, with blacks getting hired in a proportion 0.5 of the time out of 36 observations

3.5.3 Firm Behavior

With respect to firm behavior, our interest is in whether test subjects coordinate on an equilibrium. Recall that the theoretical prediction associated with the pure strategy equilibrium in the discrimination treatment is for 3 firms to set the low wage in the treatment with discrimination and for 0 firms to set the low wage in the treatment without discrimination. In contrast, the mixed strategy equilibrium is associated with a somewhat higher number of low wage firms but also with more variability. Since the aggregated results obscure important differences at the market level outcome, we present the block level results in figures 3.4a and 3.4b (block level averages can be found in table 3.A.1).
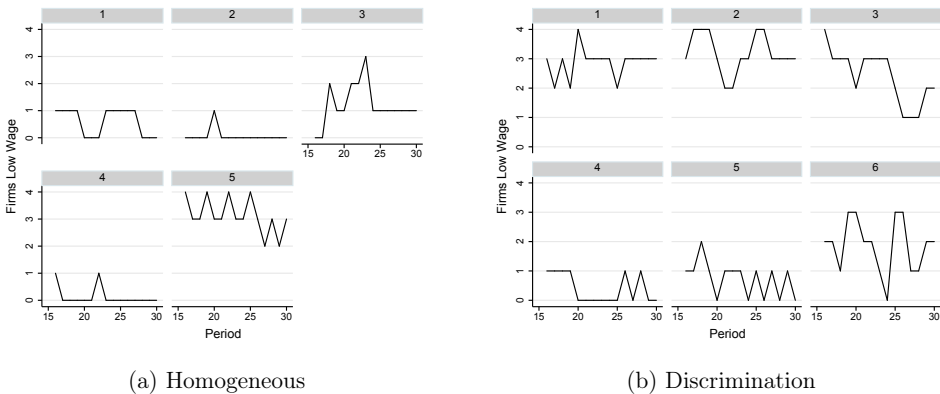

(a) Homogeneous   (b) Discrimination

Figure 3.4: Number of Low Wage Firms, Block Level

Result 3.5.3 (Firm Behavior). *(1) There is some coordination on the equilibrium in the homogeneous treatment. (2) There is some coordination on the pure strategy segregated equilibrium in the discrimination treatment. (3) At the individual level there is considerable randomization that is inconsistent with any equilibrium.*

There is considerable heterogeneity across sessions in both treatments. In the homogeneous treatment, blocks 1, 2, and 4 are consistent with equilibrium play, block 3 is ambiguous, and block 5 an outlier. In block 5, the majority of firms set the low wage even though none should do so in equilibrium. Similarly in the discrimination treatment, treatments 1 and 2 are consistent with a pure strategy segregated equilibrium, blocks 3 and 6 are somewhat ambiguous, and blocks 4 and 5 are dramatic outliers. In blocks 4 and 5 less than one firm

156

sets the low wage on average, even though the anticipated outcome is for 3 or more firms to do so on average.

In the treatment with homogeneous workers, the question is why some firms offer the low wage. In the equilibrium, workers should not apply to such firms. A possible explanation is that test subjects are risk averse and this makes it possible for some firms to offer the low wage. The difficulty in sustaining this argument is that such firms earn lower payoffs.

In the treatment with discrimination, four of the six treatments display evidence of coordination on a segregated outcome with more than two firms setting the low wage on average. Overall, the number of firms setting the low wage was somewhat closer to the level predicted by the pure strategy equilibrium than the mixed strategy equilibrium. However, behavior at the individual level is not compatible with pure coordination on a pure strategy outcome.

At the individual level, about three-fifths of the test subjects in the firm role adjusted their wage-setting behavior one or fewer times during the last 15 periods of play. This suggests some degree of coordination. However, this was accompanied by many individuals who randomized their wage-setting throughout the duration of the experiment.

3.6 References

References

Bellemare, C., Bissonnette, L., & Kröger, S. (2016). Simulating power of economic experiments: The powerbbk package. *Journal of the Economic Science Association*, *2*(2), 157–168.

Greiner, B. (2015). Subject pool recruitment procedures: Organizing experiments with orsee. *Journal of the Economic Science Association*, *1*(1), 114–125.

Helland, L., Moen, E. R., & Preugschat, E. (2017). Information and coordination frictions in experimental posted offer markets. *Journal of Economic Theory*, *167*, 53–74.

Kloosterman, A. (2016). Directed search with heterogeneous firms: An experimental study. *Experimental Economics*, *19*(1), 51–66.

Lang, K., Manove, M., & Dickens, W. T. (2005). Racial discrimination in labor markets with posted wage offers. *The American economic review*, *95*(4), 1327–1340.

Rapoport, A. & Seale, D. A. (2008). Coordination success in non-cooperative large group market entry games. *Handbook of Experimental Economics Results*, *1*, 273–295.
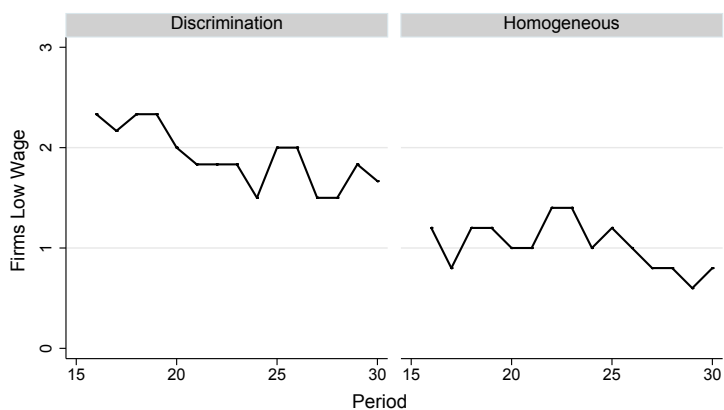
Appendix

3.A   Results

Figure 3.A.1: Number of Firms setting the Low Wage

Table 3.A.1: Firms Setting the Low Wage, by block

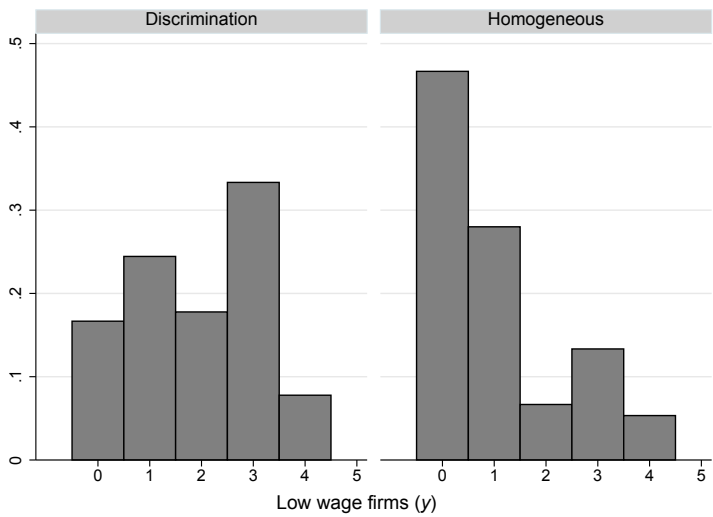|	1	2	3	4	5	6
H	0.60	0.07	1.20	0.13	3.13	
D	2.87	3.20	2.40	0.40	0.73	1.87

Figure 3.A.2: Histogram y

Chapter 4.

Should I Stay or Should I Go? Bandwagons in the Lab

Abstract

We experimentally investigate the impact of strategic uncertainty and complementarity on leader and follower behavior using the model of Farrell and Saloner (1985). At the core of the model are endogenous timing, irreversible actions and private valuations. We find that strategic complementarity strongly determines follower behavior. However, there is a reluctance to lead when leading is a conditional best response. We explain this deviation from the neo-classical equilibrium by injecting some noise in the equilibrium concept. We also find that cheap talk improves efficiency.

Author: Tom-Reiel Heggedal[*], Leif Helland[†], and Knut-Eric N. Joslin[‡]

[*]BI Norwegian Business School and Centre for Experimental Studies and Research (CESAR)

[†]BI Norwegian Business School and CESAR, e-mail: leif.helland@bi.no (corresponding author)

[‡]BI Norwegian Business School and CESAR

4.1 Introduction

Many economic environments are characterized by the presence of asymmetric information and strategic complementarity. Examples include bank runs (Garratt and Keister 2009; Goldstein and Pauzner 2005); speculative currency attacks (Morris and Shin 1998); setting of industry standards (Farrell and Saloner 1985; Farrell and Klemperer 2007); technology adoption (Katz and Shapiro 1985; Katz and Shapiro 1986); political revolts (Edmond 2013; Egorov and Sonin 2011); and foreign direct investment (Rodrik 1991; Goldberg and Kolstad 1995). In such environments there is a potential for joint welfare improvements through coordination of actions. When players' moves are endogenous, the timing of moves may in itself serve as an important coordinating device. For players with conditional best responses, strategic uncertainty enters the picture and may impact on the ability to coordinate actions.[1]

We investigate the seminal model of Farrell and Saloner (1985) (FS) in a controlled laboratory experiment.[2] In the model, players endogenously time their actions in the presence of strategic complementarity and incomplete information about types. In stage one, players simultaneously decide whether to Stay with the status quo or Go to the alternative, where Go is an irreversible action.[3] In stage two, players that are not committed to Go choose between Stay or Go again. If no player committed in stage one, second stage decisions are again simultaneous. All payoffs are obtained after the second stage. The key decision in the model is whether to Lead or Follow. A leader is defined as a player that chooses to Go in the first stage. A follower is defined as a player that Stays in the first stage and matches the first stage decision of her opponent in the second stage. Due to strategic complementarity, when a player leads, this may create incentives for the opponent to "jump on the bandwagon." The strength of the incentive depends on the private valuations of the opponent with respect to the status quo and its alternative. Thus, a player may regret the decision to Lead if the opponent fails to Follow.

In FS, the combination of a specific information structure and the endogeneity of moves produces a unique equilibrium.[4] This provides an unequivocal benchmark for our analysis

---

[1]We follow Morris and Shin (2002) in defining strategic uncertainty as "uncertainty concerning the actions and beliefs (and beliefs about the beliefs) of others." In neo-classical theory, while a player with a dominant best reply may be strategically uncertain, this has no bearing on her choice of action.

[2]For textbook treatments see Shy (2001) and Belleflamme and Peitz (2015).

[3]E.g. the action Go could—depending on the application—be "switch to the new technology platform"; "rise against the ruler"; or "make an investment". The action Stay would have the prefix "do not" attached.

[4]Coordination problems are defined by the presence of multiple, Pareto-ranked equilibria. Coordination failure results if players beliefs lead them to play a payoff dominated equilibrium. Thus, in a strict sense, there are no coordination problems in the game we use.

and facilitates separate assessment of the role of strategic uncertainty and complementarity. Our two main treatments explore variations in strategic uncertainty with respect to leadership decisions. This treatment variation turns out to be consequential also for follower decisions, through strategic complementarities.

We make two main contributions. Foremost, we find that the effect of strategic complementarity is strong. If a subject takes the lead, subjects who should follow in equilibrium do so with high probability. This contrasts with recent findings in a similar environment where players have incomplete information about fundamentals rather than types. We comment further on this below. Second, we find that subjects often do not lead when this is a conditional best response. This effect of strategic uncertainty is unaccounted for by the model. Leading carries the risk of failure; the leader might end up alone. We find that it is the variation in the cost of failed leadership, rather than the sharp cut-off between dominant and non-dominant equilibrium strategies, that appears to cause the reluctance to lead. We clarify this argument by introducing some noise in the decision making process. Such noise makes beliefs relevant everywhere, eroding the sharp divide between dominant and non-dominant equilibrium strategies. In particular, we show that an agent quantal response equilibrium (AQRE) organizes our data well.

In addition we investigate a simple extension of the model which permits cheap talk. We find that cheap talk improves players' ability to coordinate actions on mutually beneficial actions and increases efficiency.

To the best of our knowledge ours is the first experiment to address the FS-model. The paper closest to ours is Brindisi, Boğaçhan Çelen, et al. (2014).[5] While they use the same sequence of moves as we do, type uncertainty is replaced by uncertainty about fundamentals. Agents get a private signal about the true state of fundamentals, resembling the global games set-up. In contrast to us, they find that strategic complementarity does not strongly determine outcomes, as it should do in equilibrium. This indicates that the information structure is crucial in determining the strength of bandwagon behavior in the presence of complementarities and irreversible choices. While strategic complementarity is a strong force in environments with private information about types, it appears not to be so under private information about fundamentals.

More generally, most, if not all, economic situations of interest embody a mix of type uncertainty and uncertainty about fundamentals. Usually, it is not evident what the crucial source of uncertainty is in a particular situation. Accordingly, the choice of information

---

[5]Brindisi, Bogaçhan Çelen, et al. (2009) provides a thorough exposition of the theory.

structure should be determined with a view to the context.[6] For these reasons, we believe that models such as the one analyzed in this paper have the potential to shed further light on situations in which the current practice is to rely on a global games approach.

There is an experimental literature on leadership effects in weak-link games. In contrast to our setting multiple Pareto ranked equilibria coexist in these games. Like in our setting strategic complementarities are strong in these games. Several instruments of leadership have been been found to increase efficiency in weak-link games. This holds for leadership by example (Cartwright et al. 2013); leadership by communication (Brandts, Cooper, and Weber 2015; Brandts and Cooper 2007; Chaudhuri and Paichayontvijit 2010); and leaders committing to help (low ability) followers (Brandts, Cooper, Fatas, et al. 2016). There is also an experimental literature on leadership in public goods provision in which there are no strategic complementarities (see Helland et al. (2015) for a review).

The remainder of the paper is organized as follows. In the next section, we describe the model. For concreteness, we present the model using the parameters of the experiment. Thereafter, in the third section, we review our design and the experimental procedures. In section four, we present the experimental results. The fifth section considers how noisy behavior impacts the equilibrium. The final section concludes.

4.2   Model

There are two players, $i \in \{1, 2\}$, and two stages, $t \in \{1, 2\}$.[7] Players choose between the actions S_i^t and as G_i^t, that refer to Stay and Go respectively. Payoffs in the game only depend on the outcome at $t = 2$. The payoff matrix is as follows:

	S_2	G_2
S_1	7 , 7	5 , $\alpha\,\theta_2$
G_1	$\alpha\,\theta_1$, 5	$\theta_1{+}2$, $\theta_2{+}2$

Table 4.1: Payoff Matrix

---

[6]This is also the view taken in the seminal work on global games (see the discussion in Carlsson and Van Damme (1993) pp.251-2).

[7]The model can be generalized to the case with n players and n stages. The essential conclusions translate to that setting.

In the experiment we use two main treatments, D (Dominant) and N (Non-dominant). Details are provided in the next section. The essential difference is that in the D treatment Leading is sometimes a dominant strategy while in the N treatment this is never the case. The treatment difference is captured in the value of α which is $\alpha = 1$ in the D treatment and $\alpha = 1/2$ in the N treatment. For later use we define π as the payoff function conditional on type with outcomes as arguments. For instance, if subjects are in treatment D and the outcome at the second stage is Go for the subject and and Stay for her match, we write the payoff of the subject as $\pi(G_i, S_{-i}; \theta_i) = \theta_i$.

The game has the following timeline:

0. Prior to the first stage, nature draws a type θ_i for each player. Type draws are i.i.d from a uniform distribution: $\theta_i \sim U[0, 10]$. Each player's type is private information (i.e. is revealed to the player but not to the player's match). θ parameterizes preferences, with higher realizations associated with higher payoffs from action Go.

1. In the first stage, players simultaneously select action Go or Stay. The choice of Go in the first stage commits the player to Go in the second stage. Players observe the first stage action of their match at the conclusion of the stage.

2. In the second stage, players who chose Stay in the first stage choose between Stay and Go. If both players chose to Stay in the first stage, second stage actions are taken simultaneously.

We limit our self to considering symmetric equilibria. It can be shown that a unique symmetric (Bayesian perfect) equilibrium exists. In this equilibrium players use bandwagon strategies (monotone threshold strategies). In a bandwagon strategy there are two strategic thresholds $\underline{\theta}$ and θ^*. These thresholds divide players into three strategic regions: First, a range $[0, \underline{\theta})$ in which players Stay, i.e. they use the strategy $a_1 = (S_i^1, S_i^2)$. Second, a range $[\theta^*, 10]$ in which players Lead, i.e. they use the strategy $a_2 = G_i^1$. Third, a range $[\underline{\theta}, \theta^*)$ in which players Follow, i.e. they use the conditional strategy $a_3 = \left(S_i^1, (S_i^2|S_{-i}^1; G_i^2|G_{-i}^1)\right)$.[8]

The two strategically relevant thresholds, $\underline{\theta}$ and θ^* are defined by indifference conditions. First, $\underline{\theta}$ is defined as the point of indifference between strategies a_1 and a_3. If a player using

---

[8]Note that only the strategies a_1, a_2, and a_3 needs to be considered. First, we may disregard strategy $\left(S_i^1, G_i^2\right)$ since it is dominated by a_2. The reason is that a_2 induces abandonment of the status quo if the match has a type in the range $[\underline{\theta}, \theta^*)$ while by choosing strategy $\left(S_i^1, G_i^2\right)$ the player forgoes such inducements. Second, we may disregard strategy $\left(S_i^1, (G_i^2|S_{-i}^1; S_i^2|G_{-i}^1)\right)$ since it guarantees that players choose different actions.

a_1 is matched with a type above θ^* the player ends up staying alone, otherwise both players end up staying and the status quo prevails. If a player using a_3 is matched with a type above θ^* both players end up going and the status quo is abandoned, otherwise the status quo prevails. The indifference condition is then

$$\mathbb{P}(\theta_{-i} > \theta^*)\pi(S_i, G_{-i}; \underline{\theta}) + (1 - \mathbb{P}(\theta_{-i} > \theta^*))\,\pi(S_i, S_{-i}; \underline{\theta}) =$$
$$\mathbb{P}(\theta_{-i} > \theta^*)\pi(G_i, G_{-i}; \underline{\theta}) + (1 - \mathbb{P}(\theta_{-i} > \theta^*))\,\pi(S_i, S_{-i}; \underline{\theta}).$$

Note that the payoffs $\pi(S_i, G_{-i}; \theta_i)$, $\pi(S_i, S_{-i}; \theta_i)$, and $\pi(G_i, G_{-i}; \theta_i)$ are the same for both treatments. Using the payoffs, the equation can be written as

$$\frac{(10 - \theta^*)}{10}5 + \frac{\theta^*}{10}7 = \frac{(10 - \theta^*)}{10}(\theta_i + 2) + \frac{\theta^*}{10}7,$$

which gives $\underline{\theta} = 3$. Thus for players with types below $\underline{\theta} = 3$ strategy a_1 is preferred to strategy a_3. Since a_1 is an unconditional strategy, and since strategies are monotone threshold strategies, players with types below $\underline{\theta} = 3$ has Stay as a dominant choice of action in the game.

Second, the threshold θ^* is defined as the point of indifference between strategies a_2 and a_3. If a player using a_2 is matched with a type above $\underline{\theta}$ the status quo is abandoned, otherwise the player ends up going alone. If a player using a_3 is matched with a type above θ^* the status quo is abandoned, otherwise the status quo prevails. The indifference condition is then

$$\mathbb{P}(\theta_{-i} > \underline{\theta})\pi(G_i, G_{-i}; \theta^*) + (1 - \mathbb{P}(\theta_{-i} > \underline{\theta}))\,\pi(G_i, S_{-i}; \theta^*) =$$
$$\mathbb{P}(\theta_{-i} > \theta^*)\pi(G_i, G_{-i}; \theta^*) + (1 - \mathbb{P}(\theta_{-i} > \theta^*))\,\pi(S_i, S_{-i}; \theta^*).$$

Using the payoff functions from the D treatment (i.e., $\alpha = 1$), this equation can be written as

$$\frac{(10 - \underline{\theta})}{10}(\theta^* + 2) + \frac{\underline{\theta}}{10}\theta^* = \frac{(10 - \theta^*)}{10}(\theta^* + 2) + \frac{\theta^*}{10}7,$$

which reduces to

$$\theta^{*2} - 5\theta^* - 2\underline{\theta} = 0.$$

Using the fact that $\underline{\theta} = 3$, the only positive root of this equation is $\theta^* = 6$. The bandwagon strategy $(\underline{\theta} = 3, \theta^* = 6)$ is thus a unique best response to itself. Moreover, any equilibrium strategy must have the threshold form: Regardless of a player's beliefs, the benefits of leading

are non-decreasing in the player's type θ. As a consequence, if it is optimal for a player of type θ' to Go in the first stage, then it is optimal for any types $\theta > \theta'$ to also Go in the first stage. Types above θ^* will therefore Go in the first stage while types below $\underline{\theta}$ will always Stay, and the bandwagon strategy $(\underline{\theta} = 3, \theta^* = 6)$ completely characterizes the equilibrium. An analogous calculation establishes that $\theta^* = 7.3$ for treatment N (i.e., $\alpha = 1/2$).[9]

For convenience we also identify the thresholds $\bar{\theta}$ and θ°. A player with a type below θ° prefers the status quo prevailing while a player with a type above θ° prefers abandonment of the status quo. This threshold is $\theta^\circ = 5$ for both treatments. The threshold $\bar{\theta}$ defines the region in which Lead is a dominant strategy; a player with a type greater than $\bar{\theta}$ prefers to end up going alone rather than preserving the status quo. This threshold is $\bar{\theta} = 7$ in the D treatment but $\bar{\theta} = 10$ in the N treatment. This means that no players in the N treatment have a dominant strategy to Lead.

Players with types in the interval $(\underline{\theta}, \bar{\theta})$ face genuine trade-offs. Such players have conditional best responses. In particular, players with types in the interval $\theta_i \in (\theta^\circ, \bar{\theta})$ prefer the status quo to be abandoned but would choose to Stay if they knew that their match would end up staying.[10] These are players that *ex post* regret a decision to Lead if their match chooses to Stay in the final stage. These players therefore assess the expected benefits of a_2 against their best alternative strategy a_3. Such a player balances the benefits of leading, in the hope of promoting abandonment of the status quo, with the cost of possibly ending up being the only one to Go. In the unique equilibrium of the bandwagon game, θ^* denotes the point at which a player is indifferent between leading and following. Players with types greater than θ^* should therefore Lead.

Signaling We also investigate a version of the game with communication (a formal analysis of the signaling equilibrium is provided in the supplementary material 4.C). This game is identical to the basic model but with the addition of a cheap talk stage after the agents have observed their type θ_i but prior to the first stage. The message is either Stay or Go. This allows players to announce their preference for one of the outcomes. Importantly, this message is non-binding and this is common knowledge.

There exists a truth-telling equilibrium. That is, players with types below θ° prefer the

---

[9]In the N treatment we violate parts of assumption A3 in FS. This assumption assures that the threshold θ^* is in the support of θ for all monotonically increasing payoff functions. As our calculations show, the violation is inconsequential for the parameters used in our N treatment.

[10]Symmetrically, players in the interval $\theta_i \in (\underline{\theta}, \theta^\circ)$ prefer to Stay jointly but will Go if their match chooses to Go.

status quo to prevail and send the message Stay, while types above $\theta°$ prefer the abandonment of the status quo and send the message Go. The truth telling equilibrium co exists with uninformative equilibria. In such babbling equilibria, the thresholds are the same as in a game without cheap talk. In the data analysis we focus on the truth-telling equilibrium, where players can partially update their beliefs about their match's type. In the truth telling equilibrium Pareto inefficiency is eliminated as players that send the same message will choose the same action. In the case in which agents send conflicting signals, the agents use a threshold strategy analogous to in the game without communication, and Pareto inefficiency may still occur. Thus, our reason for focusing on the truth telling equilibrium is that it payoff dominates babbling equilibria. Note that the probability that a match will Follow, conditional on giving signal Stay, is lower than the unconditional probability in the absence of communication.[11] As a consequence, the threshold θ^* is higher in the game with communication.

4.3   Design and procedures

Design   The design is organized around the comparison of three treatments: Two treatments without pre-play communication, D and N, and a single treatment with pre-play communication, S. Tables 4.1 and 4.2 summarize the payoffs predictions associated with the various treatments. As shown in 4.1, the payoff functions are identical in the D, S, and N treatments except that the payoff of unilaterally switching is reduced by half in the N treatment. The fourth and fifth columns in 4.2 indicate the regions in which the best

<div align="center">Payoffs</div>

	$\pi_{\theta_i}(x_i^2, x_j^2)$	$\pi_{\theta_i}(x_i^2, y_j^t)$	$\pi_{\theta_i}(y_i^t, x_j^2)$	$\pi_{\theta_i}(y_i^t, y_j^t)$
D	7	5	θ_i	$\theta_i + 2$
N	7	5	$\theta_i/2$	$\theta_i + 2$
S	7	5	θ_i	$\theta_i + 2$

Table 4.1: Payoffs for D, N, and S $(i = \{1,2\}; t = \{1,2\})$

response to *Lead* is, respectively, not dominant and dominant.

Our key comparison is between behavior in the D and N treatments. The strategic difference between these treatments is that $\theta^* = 6.0$ in the D treatment and $\theta^* = 7.3$ in the

---

[11]The computations are analogous to those presented for the model without communication, but take into account the partial updating that results from observing the message of the match.

Predictions

	Thresholds			Best Response = y_i^1	
	$\underline{\theta}$	θ°	θ^*	Conditional	Dominant
D	3.0	5.0	6.0	$\theta_i \in [6.0, 7.0)$	$\theta_i \in [7.0, 10.0]$
N	3.0	5.0	7.3	$\theta_i \in [7.3, 10.0]$	$\theta_i \in \emptyset$
S	3.0	5.0	6.2	$\theta_i \in [6.2, 7.0)$	$\theta_i \in [7.0, 10.0]$

Table 4.2: Predictions 4.2 for D, N, and S ($i = \{1, 2\}$; $t = \{1, 2\}$)

N treatment. The other bandwagon threshold, $\underline{\theta}$, is the same in both treatments. Because the only strategically relevant threshold that changes is θ^*, a comparison of D and N provides a clean test of the model. The basic analysis compares the behavior of players in the strategic ranges *Stay*, *Follow*, and *Lead*. The model predicts that behavior in the same range will be identical across treatments.

In the first stage, the crucial decision is whether to *Lead*. However, although the model predicts that all players with types greater than θ^* will lead, the relevance of beliefs is distinctly different in the D and N treatments. In the D-treatment, the decision to *Lead* is dominant for players with types greater than $\bar{\theta} = 7.0$ (see column five in table 4.2). Such players face no strategic uncertainty. In contrast, the decision to *Lead* in the N treatment is always predicated on beliefs, and strategic uncertainty enters the picture. Comparison of subjects in the D and N treatment thus facilitates a test of the behavioral impact of beliefs on leadership.

In the second stage, the equilibrium predicts that players in the *Follow* range will choose the same action as their match chose in the first stage. We refer to this as a strategic complementarity effect. In particular, due to irreversibility, when a player's match chooses y_j^1, this resolves all strategic uncertainty in the second stage of the game. We therefore expect that leadership will powerfully determine behavior through rendering beliefs irrelevant for actions.

Our second stage analysis compares behavior of subjects conditioned on the first stage choice of their match. Of special interest is the difference between the behavior of the subjects in the *Stay* and *Follow* ranges: When a player's match chooses to *Stay*, players in the two ranges should behave identically. However, when a player's match chooses to *Lead*, players in the two ranges should make opposite choices. Hence, a direct measure of the complementarity effect may be computed by comparing the difference of the *Go* frequency in the presence and absence of a leader in the *Stay* and *Follow* ranges.

In the S treatment, players send a cost-free signal simultaneously, prior to taking their first stage action. According to theory, access to a cost-free signal should eliminate Pareto inefficiency. We implement this treatment with the same parameters as the D treatment. This allows a direct assessment of differences in efficiency, including Pareto inefficiency, due to communication. We compute overall measures of efficiency and also report the prevalence of specific varieties of inefficiency.

Moreover, in contrast to most other studies, we examine the role of strategic complementarity in the presence of conflicts of interest. For instance, players who prefer the status quo to prevail will send the signal that they intend to *Stay*. However, if they are in the *Follow* range, they will *Go* if their match leads. This preference flipping highlights the strength of strategic complementarity in our setting.

Experimental procedures All sessions were conducted in the research lab of BI Norwegian Business school using participants recruited from the general student population at the BI Norwegian Business School and the University of Oslo, both located in Oslo, Norway. Recruitment and session management were handled via the ORSEE system (Greiner 2004). In the D and N treatments we ran five sessions per treatment with between 16 and 20 subjects per session. These data are supplemented with one session of 20 subjects for the S treatment. No subject participated in more than one session. z-Tree was used to program and conduct the experiment (Fischbacher 2007). Anonymity of subjects was preserved throughout.

On arrival, subjects were randomly allocated to cubicles in the lab in order to break up social ties. After being seated, instructions were distributed and read aloud in order to achieve public knowledge of the rules. All instructions were phrased in neutral language. Subjects were asked to choose either shape *Circle* (that is, *Stay*) or shape *Square* (that is, *Go*). A goal was to avoid prioritizing one of the actions as a default option. Sample instructions and screen shots are provided in the supplementary materials.

Each session of the experiment began with two non-paying test games for subjects to get acquainted with the software. This was immediately followed by $n-1$ games in which the subjects earned payoffs, where n is the total number of participants in the session.[12] In each game, subjects were matched with one other subject, their "match", according to a highway protocol. Every subject thus met every other subject once and only once.[13] In total our data consists of 2673 unique games (excluding test games). Each game consisted

---

[12]Hence, in a treatment with 20 participants, each participants played 19 rounds with payoffs.

[13]This protocol eliminates certain dynamic problems, such as strategic teaching and reciprocity (see Fréchette (2012) for a discussion).

of a single repetition of the two-player, two-stage game with the rules and payoff functions outlined above. Subjects earned experimental currency units (ECUs). After the final game, accumulated earnings in ECU were converted to NOK, using a fixed and publicly announced exchange rate. Subjects were paid in cash privately as they left the lab. On average subjects earned 250 Norwegian Kroner (about 36 USD at the time). A session lasted on average 50 minutes.

Gameplay was formulated in the following fashion: At the beginning of each each new game, each subject received a private number drawn from a uniform distribution on the interval $(0, 10)$ with two decimal points of precision. This number corresponded to the subject's type θ_i. A dedicated screen was used to display this information. Thereafter, subjects observed a 2×2 matrix with their own payoffs from the four possible combinations of outcomes and a button to choose a first-stage action. The first stage concluded when both subjects in the match had made their decisions. If both subjects decided to *Go* in the first stage, they continued directly to the feedback and bypassed the second stage.

The second stage began with a screen that revealed the first stage actions of both subjects in the match. Subjects who decided to *Stay* in the first stage made a second stage decision. If a subject's match decided to *Go* in the first stage, then the subject observed a truncated 2×1 matrix in which the payoffs conditioned on the match choosing *Stay* was removed. This reflects the fact that the subject's match had committed to *Go*.

After all second stage decisions were resolved, the subjects move to a feedback screen. The feedback displayed payoffs from the current game as well as a history of type draws, choices, and outcomes in all previous periods in which the subject had participated.

The signal treatment S included an additional stage between the type draw stage and the first stage action choice. In this stage, subjects selected either the message "I choose circle" or the message "I choose square". Next, the message was revealed to their match on a dedicated screen. Apart from this additional stage, the screens and information are identical to those used in the two other treatments.

4.4   Results

First stage behavior   The first stage behavior of the subjects is consistent with the use of bandwagon strategies and the essential predictions of the model. Figure 4.1 shows first stage behavior of subjects in the D and N treatments. On the horizontal axis is a set of twenty bins, corresponding to 0.5 intervals over subject types: The first bin includes subjects with

types $\theta \in [0, 0.5)$, the second bin includes subjects with types $\theta \in [0.5, 1)$, etc. Each of the bins shows the proportion of subjects in the given range who chose to *Go* in the first period. We interpret this as a probability. The bubbles are scaled by the number of observations within a bin, relative to the total number of observations within a treatment.

Figure 4.1: First Stage Behavior
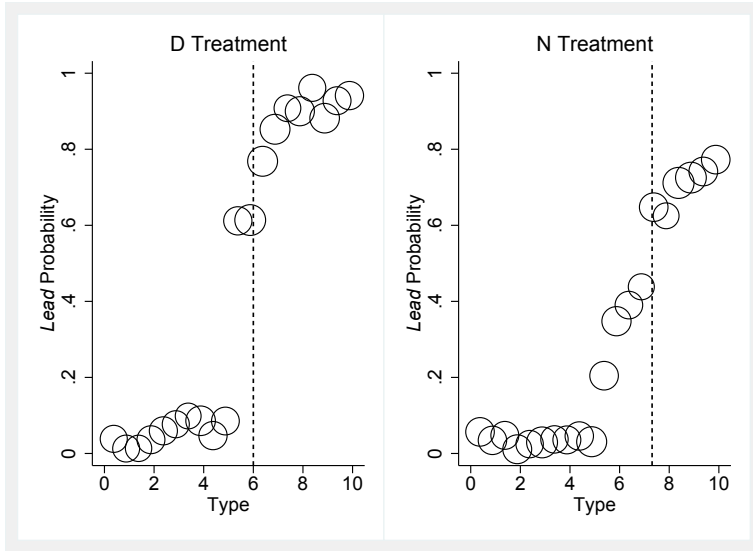
The theoretical prediction for the first stage behavior is a step function at θ^*: In the equilibrium of the model, players with types below θ^* *Stay* in the first stage while those above θ^* *Lead*. For each treatment, this threshold is indicated by a stapled line. The plots in figure 4.1 illustrate that subjects with low types tend to *Stay* in the first stage while subjects with high types tend to *Go*. Moreover, the frequency of switching increases steeply in the vicinity of θ^* in both treatments. This is consistent with the use of bandwagon strategies.

To formally assess the predictions of the model, we compare behavior across treatments using Wilcoxson Rank Sum (WRS) tests. Throughout, we denote tables and figures in the supplementary materials using a letter (A, B, etc.). Using session level data, between treatment comparisons find no significant differences in behavior over the D and N treatments in either the Stay range or the Follow range.[14] Likewise, in the region of the Lead range

---

[14]Details are provided in tables 4.A.1 and 4.A.2. Inspecting figure 4.A.2 there appears to be treatment differences in the error rate also to the left of θ^*. As noted these differences are not statistically significant at conventional levels, and we refrain from further discussion of them.

in which decisions are conditioned on beliefs (i.e. in the region $[\theta^*, \bar{\theta})$), we are unable to identify a significant treatment difference.[15] In short, when beliefs are relevant, we find no behavioral differences over treatments.

In contrast, when we compare behavior over the entire Lead range, $[\theta^*, 10]$, we reject equality of the treatments. In this region, the probability that a subject chooses to Lead is a full 18 percentage points higher in the D treatment than in the N treatment (89 percent compared to 71 percent). This difference is strongly significant.[16]

Given a 5 percent significance level and the observed variances of the D and N treatments, this test has a power of 99 percent.[17]

Result 4.4.1 (Leader Behavior)**.** *Subjects are relatively more reluctant to* Lead *when leading is a conditional best response.*

Relative to the model, subjects with types above θ^* do not choose to Go often enough. In doing so, subjects forgo an opportunity to induce their favored outcome if their match is in the Follow range.[18] This is costly. When we compare the payoff consequences of a decision to Stay for players with a best response to Lead*,* we find that on average players in the D treatment earn 3.6 ECU less when failing to Lead while the comparable number for players in the N treatment is 1.6 ECU.[19]

Below, we use the (agent form) quantal response equilibrium—in which beliefs are relevant everywhere—to rationalize observed deviations from the equilibrium of the model.

Second Stage Behavior Figure 4.2 presents the second-stage behavior by treatment, conditional on the match's action. The figure includes only those subjects who take a second stage decision. On the horizontal axis is subject type, grouped in half unit bins, and on the vertical axis is the proportion of subjects that Go in the second stage. The left panel presents the second stage behavior for subjects whose match chose to Stay in the first stage while the right panel presents the second stage behavior for subjects whose match

---

[15]We do not identify a difference regardless of whether we compare the range $\theta_i^D \in (6, 7)$ in the D treatment with the entire Lead range in the N treatment $\theta_i^N \in (7.3, 10)$ or if we base the test on a balanced set of data and use the restricted ten base point region just beyond the Go threshold, $\theta_i^N \in (7.3, 8.3)$ (See tables 4.A.3 and 4.A.4).

[16]See table 4.A.5.

[17]Calculated using the simulation routine of Bellemare et al. (2016).

[18]Because types are distributed uniformly, subjects are expected to be in the Follow range 30 percent of the time in the D treatment and 43 percent of the time in the N treatment. The actual rates realized in the treatments were 35 percent and 43 percent.

[19]These are subjects in the region $\theta \in [\theta^* = 6, \bar{\theta} = 7)$ in the D treatment and $\theta \in [\theta^* = 7.3, 10]$ in the N treatment.

chose to Lead in the first stage. Data from the N treatment are presented as hollow bubbles and data from the D treatment are presented as shaded bubbles. The size of the bubbles reflects the proportion of observations in a bin relative to the total number of observations within a treatment. The thresholds identified by the equilibrium of the model are marked by vertical lines: We denote $\underline{\theta}$ by a short dashed black line (this threshold is identical for both treatments) while we denote θ^* by a short dashed black line for the D treatment ($\theta^* = 6$) and a long dashed gray line for the N treatment ($\theta^* = 7.3$).

Figure 4.2: Second Stage Behavior by Match Action
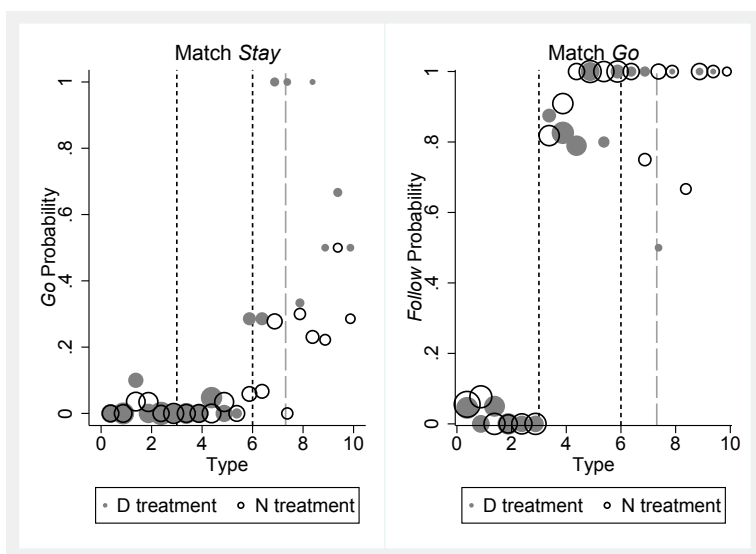
As is evident from figure 4.2, complementarity has a strong effect on the outcomes. When a subject in the Follow range has a match that stays in the first stage, the subject also stays with high probability: In 92 percent of cases in the D treatment and in 93 percent of cases in the N treatment. In contrast, when a subject in the Follow range has a match that Leads, the subject tends to Follow with a high probability: In 90 percent of cases in the D treatment and in 87 percent of cases in the N treatment. Subjects in the Follow range essentially mirror the first stage action of their match. Thus we find strong evidence of strategic complementarity. WRS tests do not identify differences in the behavior across treatments.[20] Second stage actions thus appear to be consistent with the use of bandwagon

---

[20]Details are provided in tables 4.B.1, 4.B.2, 4.B.3, and 4.B.4.

strategies.[21]

Result 4.4.2 (Follower Behavior)**.** *Strategic complementarity strongly determines outcomes.*

Although the model predicts that all subjects with types above θ^* will Lead, some fail to do so. These subjects make "mistakes" relative to the equilibrium. Nevertheless, the optimal second stage behavior for these subjects is easy to characterize: When a subject's match chooses to Lead, all subjects above $\underline{\theta} = 3$—including those above θ^*—should Follow. This is due to strategic complementarity. Otherwise, if a subject's match chooses to Stay in the first stage, only subjects with Go as a dominant strategy should Go in the second stage. Since only players in the D treatment have dominant strategies, we have distinct predictions for the D and N treatments: When a subject's match chooses to Lead, behavior should be identical in the treatments. However, when a subject's match chooses to Stay, only subjects in the D treatment with types above $\bar{\theta}$ should Go.

Consistent with these predictions, when a subject's match has chosen to Lead, subjects with types above $\underline{\theta}$ Follow with high probability. Also consistent with predictions, when a subject's match Stay, high types in the D treatment Go with higher probability than analogous types in the N treatment.[22]

Signal In the signal treatment, there are four possible outcomes from the communication stage: Two outcomes in which the subjects give the same signal, either Go or Stay, and the two outcomes in which the subjects give opposite signals. We present the first-period results in figure 4.3.

The opportunity for pre-play communication enables players to update their beliefs about their match's type (section 4.C shows the computation of the bandwagon thresholds in this case). If both subjects signal the same action, then both subjects should choose that action in the first stage. These results are shown on the diagonal. In the top left panel, we see that when subjects with types below $\theta < \theta^\circ = 5$ meet each other they nearly always signal Stay and then Stay in the first stage. Similarly, on the bottom right, we see that when subjects with types $\theta > \theta^\circ = 5$ meet each other they nearly always send the signal Go and then Go in the first stage. Relative to the D and N treatments, the coordination success of subjects in the signal treatment is substantially higher.

---

[21]These effects are confirmed in the supplementary materials, using a logistic regression (see the analysis in relation to figure 4.B.1).

[22]See the left panel figure 4.2 in which players with dominant strategies correct their mistake in the second stage even if their match choose to Stay.

Figure 4.3: First Stage GO by Signal
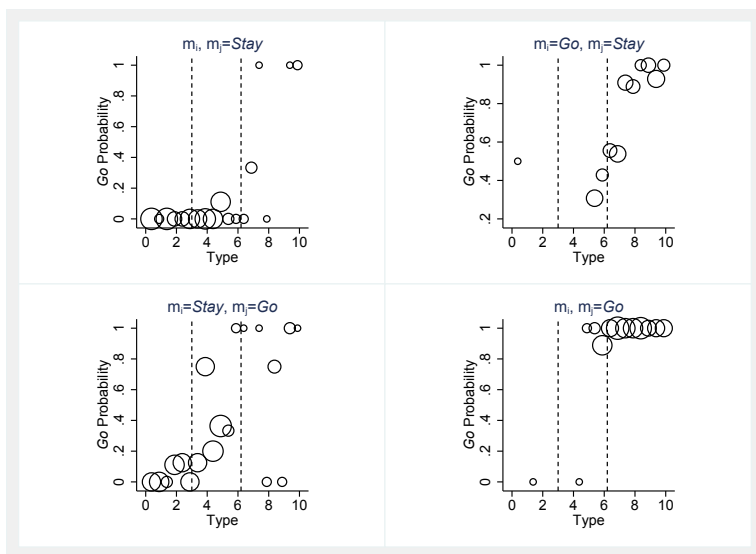
On the off diagonal, we see the instances in which subjects sent conflicting signals. In these cases, the subjects should play bandwagon strategies similar to the D treatment, with the exception that $\theta^* = 6.2$ in the S treatment. The pattern of behavior is similar to the D treatment: subjects with a dominant strategy (respectively $\theta < \underline{\theta}$ and $\theta > \overline{\theta}$) behave as predicted while there is an over eagerness for subjects just below θ^* to Lead. Comparison of the D and S treatments thus demonstrates that communication improves coordination of actions substantially whenever subjects have the same preferred outcome.

Result 4.4.3 (Communication). *Cheap talk promotes subjects ability to coordinate actions on mutually beneficial outcomes.*

Efficiency The different treatments affect the incentives and ability of subjects to achieve efficient outcomes. Figure 4.4 presents the theoretical and empirical efficiency of each treatment, computed as the fraction of the maximum total earnings.[23] The theoretical efficiency

---

[23]The *maximum* total earnings is computed as the payoff that would be realized if a social planner chose the subjects' actions to maximize total payoff. The *theoretical* payoff is computed as the payoff that would be realized if all subjects played the equilibrium bandwagon strategy. The *empirical* payoff is computed based on the actual payoffs of test subjects. Note that the maximum total surplus is *nearly* identical across treatments.

Figure 4.4: Efficiency
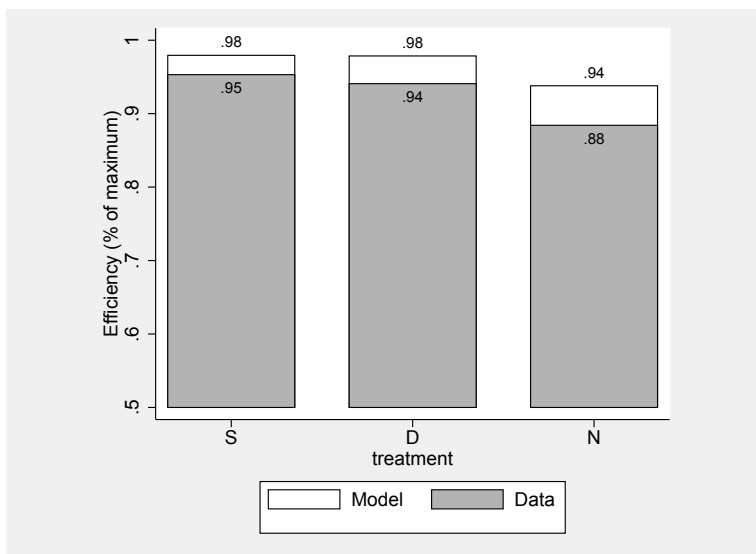
associated with equilibrium play of the model is marginally higher in the S treatment than in the D treatment, but substantially lower in the N treatment. In comparison, the efficiency realized from the data is 95.3 percent of the maximum in the S treatment, compared to 94.1 percent in the D treatment and 88.4 percent in the N treatment. Furthermore these differences are statistically significant at conventional levels using one-sided Wilcoxson Rank Sum tests.[24] Efficiency thus improves with access to cheap talk and with the redundancy of beliefs.

Result 4.4.4 (Efficiency). *Cheap talk and unconditional best responses to* Lead *improve efficiency.*

4.5   Agent Quantal Response Equilibrium

Although the predictions of the model tend to be supported by the data, behavior is not uniformly consistent with the equilibrium of the model. In particular, we observe an asym-

---

[24]See the appendix section 4.D for WSR statistics for pairwise comparisons of the treatments and the associated tabulations of realized efficiency by session.

metric pattern of errors in the vicinity of θ^*.[25] Although it is natural for subjects to make mistakes in the computation of θ^*, we would expect a symmetric pattern of error if mistakes were idiosyncratic. Asymmetry suggests instead a systematic deviation from the equilibrium. The overall level of errors is also higher in the N treatment than the D treatment.

A key observation is that the frequency of errors is inversely related to their costs. Subjects with types close to θ^*, who are nearly indifferent between Lead and Stay, often make mistakes while subjects with extreme types, who strongly prefer either abandonment or preservation of the status quo, rarely do. This is consistent with the core intuition for a quantal response equilibrium. We therefore estimate the agent quantal response equilibrium (AQRE) of the model (McKelvey and Palfrey 1998). This framework enables us to assess whether the observed pattern of behavior is consistent with an equilibrium in which decisions are noisy. Furthermore, the AQRE perspective emphasizes that beliefs are consequential everywhere. Since the model we study has a unique equilibrium, this allows us to gauge the impact of beliefs on behavior in a smooth way.

Employing the notation in Turocy (2010), let a, a' denote actions and $I(a)$ denote the information set that includes action a. In a game of perfect recall, like the bandwagon game, any node appears at most once along any path of play. Let ρ denote a behavior strategy profile. Such a profile denotes, for each action a, the probability ρ_a that action a is played if information set $I(a)$ is reached. Finally, let $\pi_a(\rho)$ denote the expected payoff to the player of taking action a on reaching information set $I(a)$, contingent on the behavior profile ρ being played at all other information sets. We say that the strategy profile is a logit AQRE if, for all players, for some $\lambda \geq 0$, and for all actions a and every information set:

$$\rho_a = \frac{e^{\lambda \pi_a(\rho)}}{\sum_{a' \in I(a)} e^{\lambda \pi_{a'}(\rho)}}$$

In an AQRE $\rho_a > 0$ for all actions a. Thus, beliefs are relevant everywhere. Equilibrium requires that beliefs are correct at each information set. The set of logit AQRE maps $\lambda \in [0, \infty]$ into the set of totally mixed behavior profiles. Letting $\lambda \to \infty$ identifies a subset of the set of sequential equilibria as limiting points (McKelvey and Palfrey 1998). Thus, when noise vanishes one is back in the neo-classical equilibrium theory. On the other hand, and for a given game, moderate noise can get amplified in an AQRE, resulting in substantial deviations from neo-classical equilibrium theory.

We estimate the logit AQRE on 20 equally sized bins (i.e. the empirically observed

---

[25]We document this asymmetry further in figure 4.A.2 which plots the distribution of errors along with a smoothed trend.

leading frequency in that range) for the three decision nodes: The first stage action and two second stage stage actions that relate to whether the match chose Stay or Go in the first stage. Our estimation performs a fixed point iteration in which we loop through the QREs for each stage, taking behavior in the other stages as given. We fit λ by minimizing the distance between the binned empirical data and the estimates. Figure 4.1 presents the best fit for each treatment individually.[26] We choose to present the individually estimated logit AQREs because the treatments are quite different, both in terms of the costs of unilaterally choosing Go (which are higher for the N treatment) and in terms of the complexity of the environment (in the D treatment the majority of the players have a dominant strategy whereas the majority of players in the N treatment have only a conditional best response).[27] [28]

Figure 4.1: Efficiency
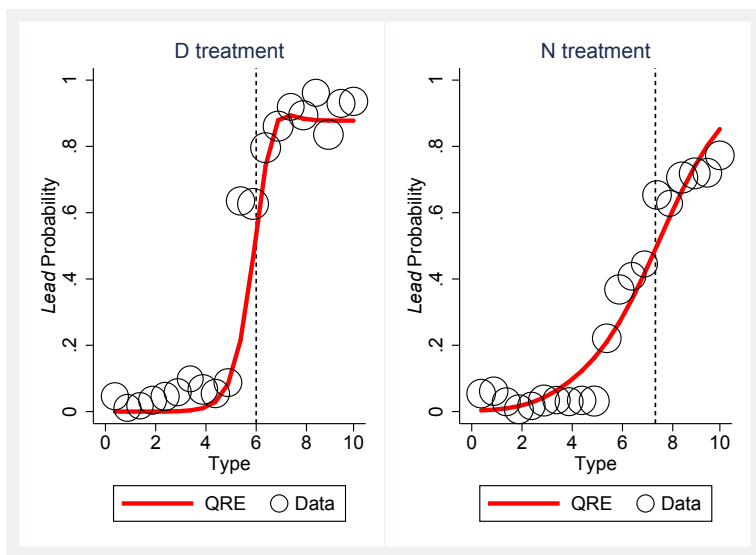
The AQRE reproduces key features of the data. Crucially, it captures the asymmetry

---

[26]Section 4.E in the supplementary materials provides a detailed discussion of the estimation procedure.

[27]Jointly estimated logit AQRE are presented in the supplementary materials figure 4.E.1. With joint fitting of the data, it is primarily the fit for high types in the N treatment that suffers. Qualitatively, however, the jointly estimated logit AQREs are consistent with the ones presented in the main text.

[28]Haile et al. (2008) demonstrate the lack of falsifiability of QRE when any error distribution is permitted. However, even a treatment by treatment estimation of the logit AQRE is disciplined by the extreme value distributional assumption necessary to arrive at the logit form of choice probabilities.

around θ^*: The AQRE correctly predicts that types just below θ^* deviate from the model to a greater extent than types just above this cut-off. The AQRE also identifies key features such as the stable level of leading for high types in the D treatment.[29] Between treatments, the AQRE correctly predicts that there should be a rapid change in behavior around the cut point $\theta^* = 6$ in the D treatment whereas behavior in the N treatment should change more gradually. The AQRE thus predicts the difference in the level of errors we observe in the data. Relative to our earlier discussion of the role of beliefs with regard to conditional and unconditional best responses, the AQRE provides a more nuanced perspective: It suggests that beliefs vary continuously and that this is an important feature for modeling actual behavior.

Result 4.5.1 (Noisy Leadership)**.** *The AQRE rationalizes observed behavior. In particular, it explains the reluctance to Lead when leading is a conditional best response in the neo-classical equilibrium.*

4.6   Conclusion

We have investigated the model of Farrell and Saloner (1985) in a controlled laboratory experiment. We find that subjects by and large respond to the incentives of the model as predicted. However, there is a reluctance to Lead not accounted for by the model. This reluctance is primarily present when leadership failure is costly. For our parameters leadership failure is more costly when leading is a conditional best response. We use a quantal response equilibrium to account for this phenomenon. In the quantal response equilibrium beliefs are relevant everywhere. We find that the observed deviations from neo-classical equilibrium is explained well by injecting some noise in the equilibrium concept.

Once a subject decides to Go he or she produces a strong incentive for moderate types to Follow. This is because the leader resolves all uncertainty on behalf of potential followers. We find that this complementarity in actions strongly determine follower behavior. Hence, the main driver of deviations from neo-classical equilibrium is weak leadership. As

---

[29]The flat (and even declining for high noise) Lead probability for players in the D treatment with high types is the outcome of the subgame structure. For players with high types, if they fail to Lead in the first period, there is still a high probability of Leading in the second (since they prefer y alone). The payoff consequence is therefore about the same for all players in this range: It is the size of the payoff externality from not inducing the preferred outcome. This predicts similar behavior for these players. In addition, when behavior is noisy, lower types are less likely to correct their mistakes in the second period than higher types. This can explain why for lower levels of noise it is actually types in the vicinity of $\bar{\theta}$ for whom a error to not Lead is most costly.

a consequence, efficiency losses are greater when potential leaders have non-dominant best responses. However, we find that cheap talk improves players' ability to coordinate on mutually beneficial actions and increases efficiency.

4.7   References

References

Belleflamme, P. & Peitz, M. (2015). *Industrial organization: Markets and strategies.* Cambridge, UK: Cambridge University Press.

Bellemare, C., Bissonnette, L., & Kröger, S. (2016). Simulating power of economic experiments: The powerbbk package. *Journal of the Economic Science Association, 2*(2), 157–168.

Brandts, J. & Cooper, D. J. (2007). It's what you say, not what you pay: An experimental study of manager-employee relationships in overcoming coordination failure. *Journal of the European Economic Association, 5*(6), 1223–1268.

Brandts, J., Cooper, D. J., Fatas, E., & Qi, S. (2016). Stand by me—experiments on help and commitment in coordination games. *Management Science, 62*(10), 2916–2936.

Brandts, J., Cooper, D. J., & Weber, R. A. (2015). Legitimacy, communication, and leadership in the turnaround game. *Management Science, 61*(11), 2627–2645.

Brindisi, F., Çelen, B. [Bogaçhan], & Hyndman, K. (2009). *On the role of information and strategic delay in coordination games: Theory and experiment.* Unpublished manuscript, https://www.mcgill.ca/economics/files/economics/hyndman.pdf.

Brindisi, F., Çelen, B. [Boğaçhan], & Hyndman, K. (2014). The effect of endogenous timing on coordination under asymmetric information: An experimental study. *Games and Economic Behavior, 86*(July), 264–281.

Carlsson, H. & Van Damme, E. (1993). Equilibrium selection in stag hunt games. In K. Binmore, A. Kirman, & P. Tani (Eds.), *Frontiers of game theory* (Chap. 12, pp. 237–253). Cambridge, MA: MIT Press.

Cartwright, E., Gillet, J., & Van Vugt, M. (2013). Leadership by example in the weak-link game. *Economic Inquiry, 51*(4), 2028–2043.

Chaudhuri, A. & Paichayontvijit, T. (2010). Recommended play and performance bonuses in the minimum effort coordination game. *Experimental Economics, 13*(3), 346–363.

Edmond, C. (2013). Information manipulation, coordination, and regime change. *Review of Economic Studies*, *80*(4), 1422–1458.

Egorov, G. & Sonin, K. (2011). Dictators and their viziers: Endogenizing the loyalty–competence trade-off. *Journal of the European Economic Association*, *9*(5), 903–930.

Farrell, J. & Klemperer, P. (2007). Coordination and lock-in: Competition with switching costs and network effects. In M. Armstrong & R. Porter (Eds.), *Handbook of industrial organization, volume 3* (1st ed., Chap. 31, pp. 1967–2072). Amsterdam: North Holland.

Farrell, J. & Saloner, G. (1985). Standardization, compatibility, and innovation. *RAND Journal of Economics*, *16*(1), 70–83.

Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, *10*(2), 171–178.

Fréchette, G. R. (2012). Session-effects in the laboratory. *Experimental Economics*, *15*(3), 485–498.

Garratt, R. & Keister, T. (2009). Bank runs as coordination failures: An experimental study. *Journal of Economic Behavior & Organization*, *71*(2), 300–317.

Goldberg, L. & Kolstad, C. (1995). Foreign direct investment, exchange rate variability and demand uncertainty. *International Economic Review*, *36*(4), 855–873.

Goldstein, I. & Pauzner, A. (2005). Demand–deposit contracts and the probability of bank runs. *Journal of Finance*, *60*(3), 1293–1327.

Greiner, B. (2004). The online recruitment system orsee 2.0-a guide for the organization of experiments in economics. *University of Cologne, Working paper series in economics*, *10*(23), 63–104.

Haile, P. A., Hortacsu, A., & Kosenok, G. (2008). On the empirical content of quantal response equilibrium. *American Economic Review*, *91*(1), 180–200.

Helland, L., Hovi, J., & Sælen, H. (2015). *Climate leadership by conditional commitments*. CESAR Working paper 1/15.

Katz, M. L. & Shapiro, C. (1985). Network externalities, competition, and compatibility. *American Economic Review*, *75*(3), 424–440.

Katz, M. L. & Shapiro, C. (1986). Technology adoption in the presence of network externalities. *Journal of Political Economy*, *94*(4), 822–841.

McKelvey, R. D. & Palfrey, T. R. (1998). Quantal response equilibria for extensive form games. *Experimental Economics*, *1*(1), 9–41.

Morris, S. & Shin, H. S. (1998). Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review*, *88*(3), 587–597.

Morris, S. & Shin, H. S. (2002). *Measuring strategic uncertainty*. Unpublished manuscript, http://www.nuff.ox.ac.uk/users/shin/pdf/barcelona.pdf.

Rodrik, D. (1991). Policy uncertainty and private investment in developing countries. *Journal of Development Economics*, *36*(2), 229–242.

Shy, O. (2001). *The economics of network industries*. Cambridge, UK: Cambridge University Press.

Turocy, T. (2010). Computing sequential equilibria using agent quantal response equilibria. *Economic Theory*, *42*(1), 255–69.
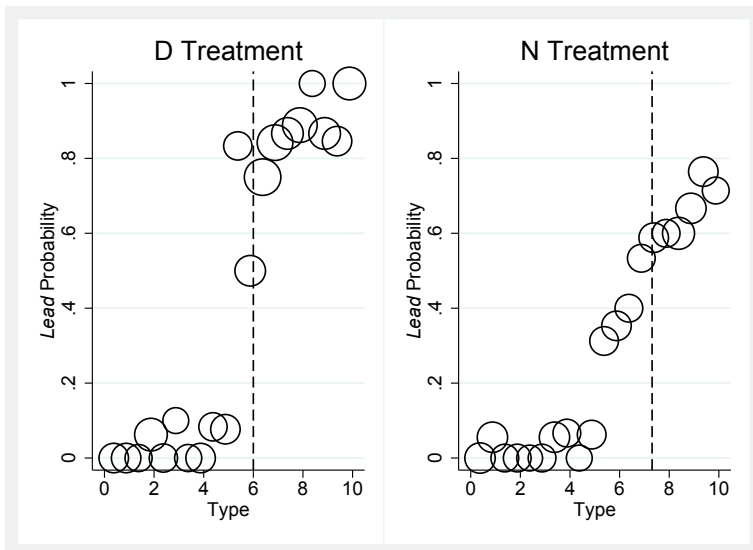
Appendix

4.A   First-Stage Behavior

Figures

Figure 4.A.1 presents the same plot as figure 4.1, but based on only the last half of periods in each treatment. The main difference from the plot in the main text is an increased frequency of switching for subjects in the D treatment with types in the range $[5, 5.5)$.

Figure 4.A.1: First Stage, Last Half of Periods

Between Treatment Comparisons

To assess whether behavior in the strategic ranges is the same in the D and N treatments, we use Wilcoxon rank-sum (WSR) tests to compare the frequency with which participants choose to *Go*. These tests are based on between treatment comparisons of session-level data. For each test we present the relevant session data and the associated means and standard deviations. We denote the WSR test statistic by W. The *p*-value indicates how likely it is that the given observations come from the same distribution.

Comparison of behavior in the *Stay* Range (Table 4.A.1).   A two-sample Wilcoxon rank-sum (Mann-Whitney) test can not reject equality of behavior in the *Stay* range, $\theta \in [0, \underline{\theta}]$, in the D and N treatments: $W = 0.1$, $p = 0.92$.

Session	D	N
1	0.00	0.01
2	0.02	0.03
3	0.04	0.04
4	0.08	0.05
5	0.08	0.07
Mean	0.04	0.04
Std	0.04	0.02

Table 4.A.1: $\theta_D \in [0, 3)$ vs. $\theta_N \in [0, 3)$

Comparison of behavior in the *Follow* Range (Table 4.A.2).   A two-sample Wilcoxon rank-sum (Mann-Whitney) test can not reject equality of behavior in the *Follow* range, $\theta \in [\bar{\theta}, \theta^*]$, in the D and N treatments: $W = 0.94, p = 0.35$.

Session	D	N
1	0.22	0.12
2	0.22	0.18
3	0.24	0.23
4	0.33	0.24
5	0.36	0.29
Mean	0.27	0.21
Std	0.07	0.06

Table 4.A.2: $\theta_D \in [3, 6)$ vs. $\theta_N \in [3, 7.3)$

Comparison of behavior when *Lead* is a conditional best response (Tables 4.A.3 and 4.A.4). A two-sample Wilcoxon rank-sum (Mann-Whitney) test can not reject equality of behavior when *Lead* is a conditional best response regardless of whether we compare the region $\theta \in [\theta^*, \overline{\theta}]$) in the D treatment with the entire *Lead* region in the N treatment or just the restricted region $\theta \in [\theta^*, \theta^* + 1]$: Both tests deliver identical results, $W = 1.36, p = 0.17$.

Session	D	N
1	0.69	0.63
2	0.73	0.70
3	0.76	0.71
4	0.85	0.73
5	0.92	0.77
Mean	0.79	0.71
Std	0.09	0.05

Session	D	N
1	0.69	0.44
2	0.73	0.57
3	0.76	0.71
4	0.85	0.74
5	0.92	0.79
Mean	0.79	0.65
Std	0.09	0.14

Table 4.A.3: $\theta_D \in [6, 7)$ vs. $\theta_N \in [7.3, 10]$

Table 4.A.4: $\theta_D \in [6, 7)$ vs. $\theta_N \in [7.3, 8.3)$

Comparison of behavior in the *Lead* range (Table 4.A.5). A two-sample Wilcoxon rank-sum (Mann-Whitney) test rejects equality of behavior for the *Lead* range, $\theta \in [\theta^*, 10]$: $W = 2.61, p = 0.01$.

Session	D	N
1	0.83	0.63
2	0.86	0.70
3	0.90	0.71
4	0.91	0.73
5	0.94	0.77
Mean	0.89	0.71
Std	0.04	0.05

Table 4.A.5: $\theta_D \in [6, 10]$ vs. $\theta_N \in [7.3, 10]$,

Comparison of behavior when *Lead* is an unconditional best response in D (Table 4.A.6). A two-sample Wilcoxon rank-sum (Mann-Whitney) test rejects equality of behavior for subjects with high type draws; when we compare test subjects in the D treatment in the region $\theta_D \in [\theta^* + 1, 10]$ with test subjects in the N treatment in the region $\theta_N \in [\theta^* + 1, 10]$, we strongly reject equality of behavior: $W = 2.61, p = 0.01$.

Session	D	N
1	0.87	0.68
2	0.89	0.69
3	0.92	0.74
4	0.94	0.78
5	0.97	0.79
Mean	0.92	0.74
Std	0.04	0.05

Table 4.A.6: $\theta_D \in [7, 10]$ vs. $\theta_N \in [8.3, 10]$

The Cost of Failing to Lead

In the Bandwagon game, the strategic decision to *Stay* or *Lead* is most difficult for players in the vicinity of the first stage switching threshold θ^* who have a conditional best response to *Lead*. These are players who prefer a joint choice of y but would stay on x if they knew that their match will choose x with certainty. In the D treatment this range is relatively small while in the N treatment it is relatively large. To get a measure of how costly it is for players in this region to forgo leading, we tabulate the frequency with which subjects in the relevant ranges encounter a subject in the follow range. Next, we tabulate the frequency with which subjects take the correct strategic timing decision and the matched subject does in fact follow.[30] Finally, as a crude measure of the importance of correctly taking the strategic timing decision, we list the average payoff from the (correct) decision to *Lead* relative to the average payoff from the choice to *Stay*:

| Data	Correct Go	Correct Go & Followed	Follow Error	ΔPayoff
Joint	0.79	0.72 (0.91)	6.51	3.28
D	0.79	0.79 (1.00)	9.17	1.59
N	0.78	0.71 (0.91)	5.24	3.61

Table 4.A.7: Strategic Timing

Errors

Figure 4.A.2 shows the empirical error frequency in 0.1 unit bins along with a think black black line which is a 0.5 unit moving average of the error frequencies. This figure has three main features: First, the pattern of errors is asymmetric around the cut point in both

---

[30]In about 9% of cases subjects in the *Follow* range fail to follow.

treatments. Second, there is sharper pattern of errors in the D treatment relative to the N treatment. This indicates that the range in which the decision to lead is uncertain is more narrow in the D treatment relative to the N treatment. Third, there is a generally higher level of errors in the N treatment relative to the D treatment for subjects with high types. This testifies to an overall higher level of uncertainty in the N treatment.

Figure 4.A.2: First Stage Errors
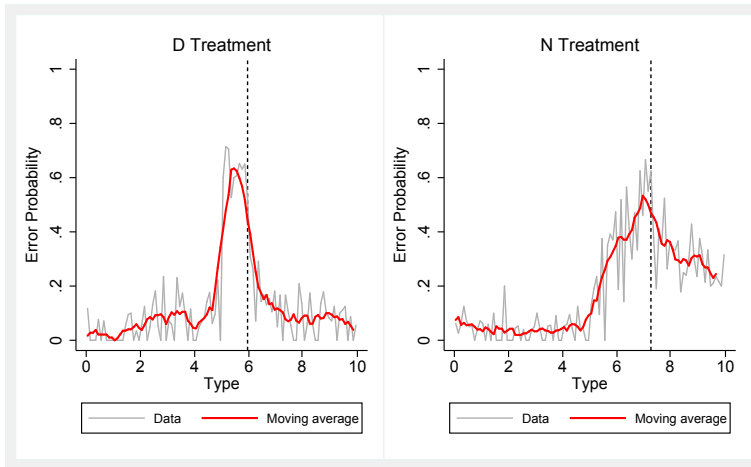
4.B   Second-Stage Behavior

Between Treatment Comparisons

Comparison of behavior in the Stay Range (Table 4.B.2 and 4.B.1).   Two-sample Wilcoxon rank-sum (Mann-Whitney) test can not reject equality of behavior when test subjects are in the Stay range. This holds when the match chooses to Stay ($W = 1.16, p = 0.24$) and when the match chooses to Go ($W = -0.63, p = 0.53$).

Comparison of behavior in the Follow Range (Table 4.B.3 and 4.B.4).   Two-sample Wilcoxon rank-sum (Mann-Whitney) test can not reject equality of behavior when test subjects are in the Follow range. This holds when the match chooses to Stay ($W = 0.52$, $p = 0.60$) and when the match chooses to Go ($W = 1.79, p = 0.07$).

Match Stay		
Session	D	N
1	0.00	0.00
2	0.03	0.00
3	0.04	0.02
4	0.05	0.04
5	0.08	0.04
Mean	0.04	0.02
Std	0.03	0.02

Table 4.B.1: $\theta_D \in [0,3)$ vs. $\theta_N \in [0,3)$

Match Go		
Session	D	N
1	0.00	0.00
2	0.03	0.03
3	0.03	0.07
4	0.05	0.08
5	0.10	0.13
Mean	0.04	0.06
Std	0.04	0.05

Table 4.B.2: $\theta_D \in [0,3)$ vs. $\theta_N \in [0,3)$

Match Stay		
Session	D	N
1	0.02	0.05
2	0.03	0.06
3	0.10	0.07
4	0.12	0.08
5	0.13	0.08
Mean	0.08	0.07
Std	0.05	0.01

Table 4.B.3: $\theta_D \in [3, \theta_D^* = 6)$ vs. $\theta_N \in [3, \theta_N^* = 7.3)$

Match Go		
Session	D	N
1	0.88	0.84
2	0.89	0.85
3	0.89	0.86
4	0.92	0.88
5	0.92	0.90
Mean	0.90	0.87
Std	0.02	0.03

Table 4.B.4: $\theta_D \in [3, \theta_D^* = 6)$ vs. $\theta_N \in [3, \theta_N^* = 7.3)$

Logistic Regression

To demonstrate the predictive ability of the bandwagon model and to illustrate the role of complementarity, we estimate the logistic regression for individual i in period t

$$
\begin{aligned}
\mathbb{P}\left(y_{it}^2 | a_{jt}^1\right) = & F\left(\beta_0 + \beta_1 \theta_{it} + \beta_2 \mathbf{range}_{it} + \beta_3 \mathbf{treatment}_{it} + \beta_4 a_{jt}^1 \right. \\
& \left. + \beta_5 \mathbf{range}_i a_{jt}^1 + \beta_6 \mathbf{range}_{it} \mathbf{treatment} + \beta_7 \mathbf{treatment} a_{jt}^1 + \epsilon_{it}\right)
\end{aligned}
$$

with clustered errors for each individual. We include three dummy variables: The **range** variable players with types in the follow range (relative to the stay range), **treatment**$_i$ the D (relative to the N treatment), and a_j^1 the first stage action of i's match (1 if Go, 0 otherwise). We include interactions between these last three variables.

y_{it}^2	Coeff. (p-value)	Std. Err.
θ_{it}	0.44 (0.00)	0.11
range	-0.44 (0.31)	0.43
treatment	0.40 (0.46)	0.53
a_{jt}^1	0.84 (0.06)	0.44
range x a_{jt}^1	3.94 (0.00)	0.41
treatment x a_{jt}^1	-0.39 (0.38)	0.44
treatment x range	0.33 (0.52)	0.52
constant	-4.50 (0.00)	0.45

Table 4.B.5: Logistic estimates of second stage *Go* probability

In this specification, only type θ_{it}, match action a_{jt}^1, and the interaction between range and a_{jt}^1 are significant.[31] Contrary to theory, type has an independent effect on the probability to Go and higher types are more likely to Go regardless of their match's action. However, this effect is relatively weak. The much stronger effect is the interaction between the match's action and range. In particular, when a subject is in the follow range and their match chooses to Go in the first period, the average probability that the test subject will switch increases by about 70 percentage points relative to the case when their match chooses to Stay. In addition, we see that the impact of the match's action is close to zero in the case when the test subject is in the Stay range. These effects are clearly illustrated by figure 8 which plots the predicted Go probability given the match's action in both the D and N treatments.[32]

4.C   Signalling Game

The equilibrium with signaling takes a form similar to the equilibrium without signals. First, observe that players will always send the message that promotes the outcome that they prefer. The message that players send therefore perfectly reveals whether a player has a type above or below θ^o. An implication is that players who send the same message will choose the same action. In addition, if players send conflicting messages, then the player who prefers outcome y can update her belief about the type of their match. Relative to the game without signaling, the player has more information since the range of possible types for the match is truncated from θ^* to θ^o. In addition, if the players do not make a joint

---

[31]Alternative specifications yielded nearly identical results with small differences in estimated coefficients and p-values.

[32]Recall that in both treatments, the dividing line between the Stay and Follow ranges is at $\underline{\theta} = 3$.

Figure 4.B.1: Predicted Switching Probabilities Conditional on Match Action
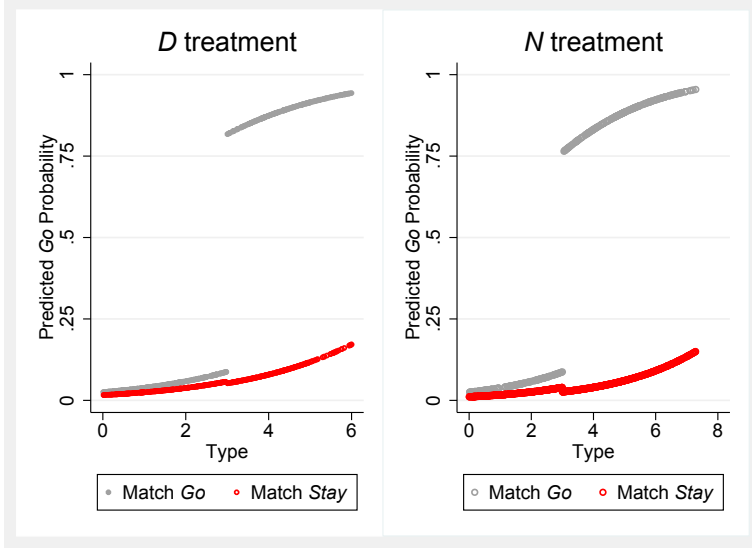
move to y, then they will remain on x. This means that the upper bandwagon threshold must satisfy:

$$\mathbb{P}(\theta_j > \underline{\theta}|\theta_j < \theta^o)\pi_{\theta^*}(y_i^1, y_j^2) + (1 - \mathbb{P}(\theta_j > \underline{\theta}|\theta_j < \theta^o))\pi_{\theta^*}(y_i^1, x_j^2) = \pi_{\theta^*}(x_i^2, x_j^2)$$

Given the uniform distribution of types on the interval $[0, 10]$ and our parameterization of the payoff functions, this reduces to

$$\frac{\theta^o - \underline{\theta}}{\theta^o}(\theta^* + 2) + \frac{\underline{\theta}}{\theta^o}\theta^* = 7$$

where $\bar{\theta} = 3$ and $\theta^o = 5$. Relative to the analogous D treatment, the upper bandwagon threshold increases slightly from 6 to 6.2 in the game in the game with signaling.

4.D   Efficiency

Wilcoxson rank-sum tests identify a statistically significant difference with respect to the *realized efficiency* (empirically observed payoff as a percent of the maximum possible payoff) between each pair of treatments, between S and D ($W = 1.4, p = 0.087$) , between S and N

$(W = 2.6, p = 0.005)$, and between D and N ($W = 2.6, p = 0.005$).

Session	S	D	N
1	0.97	0.95	0.89
2	0.96	0.95	0.88
3	0.95	0.94	0.88
4	0.95	0.93	0.88
5	0.94	0.93	0.88
Mean	0.95	0.94	0.88
Std	0.01	0.01	0.01

Table 4.D.1: Empirical efficiency as percent of maximum by treatment–session,

4.E   Equilibrium with Noise

Regarding the estimation of the AQRE, there are a few features on which it is worth commenting. Because the type space is continuous, in our estimation we discretize the type space into B equally spaced bins with types corresponding to the mid-point in the bin . Since all actions are played with some non-zero probability in an AQRE, the expected payoff for a player in bin i depends on how likely it is to get a match j in each of the $1...B$ bins and the probability that the match will choose to *Go* conditional on the actions taken so far. In particular, in the second stage, players update their beliefs about their match's type based on whether their match chose to *Stay* or *Go* in the first stage. Although we present estimates based on 20 bins in the paper, we estimated versions with up to 100 bins. Since increasing the number of bins did not change any conclusions—even delivering the same estimate of the noise parameter—we choose to present the simpler version.

Given the discretization, the estimation involves two stages. In the first stage, we estimate a fixed point for the vector of first stage *Go* probabilities taking the second stage *Go* probabilities as given. Since agents are forward looking, they anticipate how likely it is that their match will *Go* in the first period and how their own first period action will affect the second stage action of their match. In the second stage, agents that chose to *Stay* are in one of two possible situations: Either their match chose to *Go* or their match chose to *Stay*. In both cases, we must estimate a QRE for the second stage switching probability for each of the B types. Moreover, in the case when the match chose to *Stay*, the second stage estimation depends on the first-stage probability estimates because agents need to update their beliefs about how likely each type is because a high type will be more likely to *Go*

in the second stage. The first and second stage decisions are thus interlinked because the first-stage decisions depend on the anticipated second stage probabilities and the second stage decisions depend on the updated beliefs generated in the first stage.

The actual estimation proceeded by looping through the first and second stage, using the estimated probabilities from the previous iteration of the procedure as beliefs. To efficiently estimate the model, we vectorize the computations. For example, in the first stage we compute the payoff from choosing *Go* for all the types $i = 1, \ldots, B$ from the matrix multiplication

$$\frac{1}{B}\mathbf{1}\left(\mathbf{p_1^T}\pi_{\mathbf{i}}(\mathbf{y_i^1y_j^1}) + (\mathbf{1^T} - \mathbf{p_1^T})\mathbf{p_{2,Go}^T}\pi_{\mathbf{i}}(\mathbf{y_i^1y_j^2}) + (\mathbf{1^T} - \mathbf{p_1^T})(\mathbf{1^T} - \mathbf{p_{2,Go}^T})\pi_{\mathbf{i}}(\mathbf{y_i^1x_j^2})\right)$$

where all vectors are denoted in bold, are of length $1 \times B$, and transposes are indicated by a T. The vector $\mathbf{p_1}$ denotes the probability of match $j = \{1, \ldots, B\}$ going in the first stage, $\mathbf{p_{2,Go}}$ denotes the probability of match $j = \{1, \ldots, B\}$ going in the second stage conditional on i choosing to *Lead*, and $\pi_{\mathbf{i}}$ denotes payoffs to a type $i = \{1, \ldots, B\}$ that depends on the outcome realized in the second stage. The first term in the entire product, $\frac{1}{B}\mathbf{1}$, is the probability of meeting each of the B types while the second term (everything inside the outer parenthesis) is a $B \times B$ matrix that in each i, j-cell gives the expected payoff to a type i of meeting a type j. Notice that the second factor is composed of three outer products that repectively give the payoff from (1) both choosing to *Lead*, (2) i leading and j following in the second stage, and (3) i leading while j chooses to *Stay* in the second stage. The product thus produces a vector of length length $1 \times B$ that in each position gives the expected payoff to a type i that results from the sum of payoffs from meeting all the $j \in \{1, \ldots, B\}$ types. Analogous computations were carried out for the first stage payoff of choosing to *Stay*, as well as for the second stage payoffs that also depend via updated beliefs on the first stage action of the match. QRE probability estimates in each stage could then be computed as the ratio of payoffs from *Go* relative to the sum of the payoffs from *Go* plus the payoffs of *Stay* as shown in section 4.5. The QRE probabilities from the first and second stage for all B types then characterize the behavior profile. The estimation terminated when the estimates converged sufficiently that the maximum distance between the previous best estimates and the current best estimates fell below the tolerance.

How to select the level of noise and fit the AQRE is an open question. Foremost, the question is whether the noise should be jointly estimated across the two stages of the model and whether the noise should should be jointly estimated across treatments. To discipline

194

our analysis against over-fitting, we chose to constrain noise to be the same across both the first and second stage of the estimation. However, this does create some issues. Specifically, the second-stage decisions are more "simple" than the first stage decisions in the sense that the payoff consequences of making a mistake are more clear and more stark. Thus, even with low levels of noise, the second stage behavior is sharper. In turn, this has consequences for the first-stage estimates.

In figure 4.E.1, we present the same plots as in figure 4.1 but with the noise parameter estimated jointly for both treatments. In the top panel we present the data and fitted AQRE for all the periods while in the bottom panel the same information for the last ten periods.

Figure 4.E.1: AQRE, Joint estimate of lambda
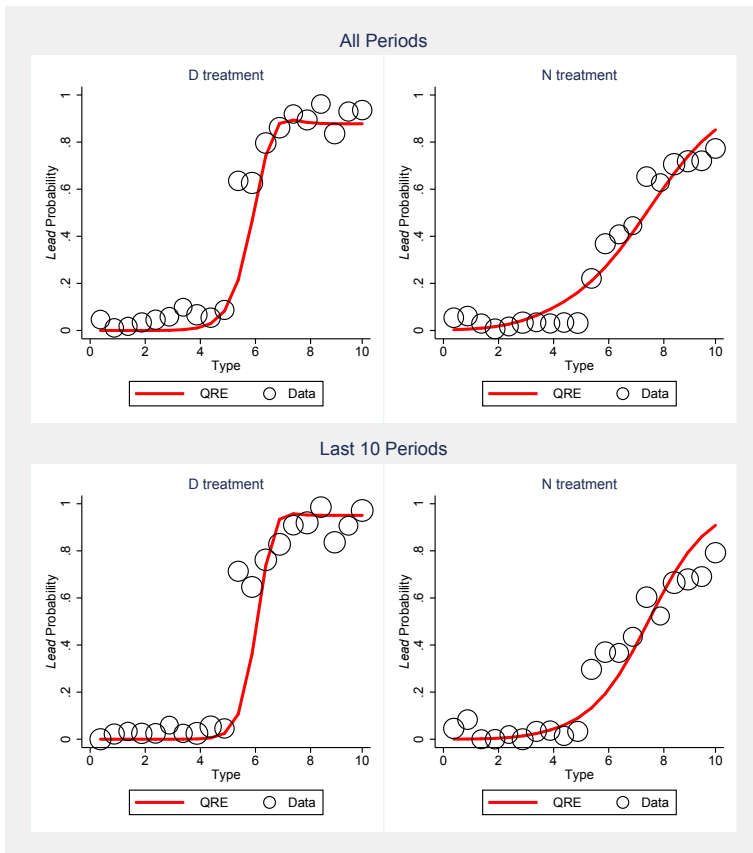
We investigated several procedures to select the level of noise: Maximum likelihood,

Euclidean distance, and absolute distance. In all cases, our fit was based on the closeness of the first-stage estimates to the data. All three procedures produced similar results, although maximum likelihood and Euclidean distance estimated somewhat higher levels of noise. The maximum likelihood estimates were strongly affected by the fact that for the lower half of the types, the predicted first-stage behavior is close to zero. Because actual behavior was somewhat greater than zero—even for lowest types—the maximum likelihood requires a high level noise. The second feature that created issues was that subjects in the bin $[5, 5.5)$ switched at a higher rate in this bin than in the bin from $[5.5, 6)$. This non-monotonic behavior is strongly punished by the quadratic distance and led to a higher noise for the Euclidean distance. Given these issues and our interest in using the model for prediction, we fit the AQRE based on the absolute distance.

Series of Dissertations

The dissertations may be ordered in printed copies from BI's website:
www.bi.no/en/Research/Research-Publications/

2017

9/2017 *Knut-Eric Neset Joslin*
Experimental Markets with Frictions

8/2017 *Sumaya AlBalooshi*
Switching between Resources: Psychosocial Resources as Moderators of the
Impact of Powerlessness on Cognition and Behavior

7/2017 *Zongwei Lu*
Three essays on auction theory and contest theory

6/2017 *Frode Martin Nordvik*
Oil Extraction and The Macroeconomy

5/2017 *Elizabeth Solberg*
Adapting to Changing Job Demands**:** A Broadcast Approach to Understanding Self-
Regulated Adaptive Performance and Cultivating it in Situated Work Settings

4/2017 *Natalia Bodrug*
Essays on economic choices and cultural values

3/2017 *Vegard Høghaug Larsen*
Drivers of the business cycle: Oil, news and uncertainty

2/2017 *Nikolay Georgiev*
Use of Word Categories with Psychological Relevance Predicts Prominence in Online
Social Networks

1/2017 *Sigmund Valaker*
Breakdown of Team Cognition and Team Performance: Examining the influence of media
and overconfidence on mutual understanding, shared situation awareness and
contextualization

2016

12/2016 *Xiaobei Wang*
Strategizing in Logistics Networks: The role of Logistics Service Providers as mediators
and network facilitators

11/2016 *Prosper Ameh Kwei-Narh*
A mid-range theory of monitoring behaviors, shared task mental models, and team
performance within dynamic settings

10/2016 *Anton Diachenko*
Ownership type, performance, and context: Study of institutional and industrial owners

9/2016 *Ranvir S. Rai*
Innovation in Practice: A Practice-Theoretical Exploration of Discontinuous Service
Innovations

8/2016   *Gordana Abramovic*
Effective Diversity Management on the Line - Who and How? On the role of line managers in organisations with a diverse workforce

7/2016   *Mohammad Ejaz*
Why do some service innovations succeed while others fail? A comparative case study of managing innovation processes of two internet-based aggregation financial services at Finn.no (Penger.no)

6/2016   *Asmund Rygh*
Corporate governance and international business: Essays on multinational enterprises, ownership, finance and institutions

5/2016   *Chen Qu*
Three Essays on Game Theory and Patent-Pool Formation

4/2016   *Arash Aloosh*
Three Essays in International Finance

3/2016   *John-Erik Mathisen*
Elaborating and Implemental Mindsets in Business-related Behavior: An Investigation of the Interface between Cognition and Action. How goals and plans emerge and relate to cognition and action in business.

2/2015   *Øyvind Nilsen Aas*
Essays in industrial organization and search theory

1/2016   *Dominque Kost*
Understanding Transactive Memory Systems in Virtual Teams: The role of integration and differentiation, task dependencies and routines

2015

8/2015   *Andreea Mitrache*
Macroeconomic Factors and Asset Prices – An Empirical Investigation

7/2015   *Lene Pettersen*
Working in Tandem: A Longitudinal Study of the Interplay of Working Practices and Social Enterprise Platforms in the Multinational Workplace

6/2015   *Di Cui*
Three Essays on Investor Recognition and Mergers & Acquisitions

5/2015   *Katja Maria Hydle*
Cross border practices: Transnational Practices in Professional Service Firms

4/2014   *Ieva Martinkenaitė-Pujanauskienė*
Evolutionary and power perspectives on headquarters-subsidiary knowledge transfer: The role of disseminative and absorptive capacities

3/2015   *Tarje Gaustad*
The Perils of Self-Brand Connections: Consumer Response to Changes in Brand Image

2/2015   *Jakob Utgård*
Organizational form, local market structure and corporate social performance in retail

1/2015   *Drago Bergholt*
Shocks and Transmission Channels in Multi-Sector Small Open Economies

2014

13/2014   *Nam Huong Dau*
Asset Pricing with Heterogeneous Beliefs and Portfolio Constraints

12/2014   *Vegard Kolbjørnsrud*
On governance in collaborative communities

11/2014   *Ignacio García de Olalla López*
Essays in Corporate Finance

10/2014   *Sebastiano Lombardo*
Client-consultant interaction practices: Sources of ingenuity, value creation and strategizing

9/2014   *Leif Anders Thorsrud*
International business cycles and oil market dynamics

8/2014   *Ide Katrine Birkeland*
Fire Walk with Me: Exploring the Role of Harmonious and Obsessive Passion in Well-being and Performance at Work

7/2014   *Sinem Acar-Burkay*
Essays on relational outcomes in mixed-motive situations

6/2014   *Susanne H.G. Poulsson*
On Experiences as Economic Offerings

5/2014   *Eric Lawrence Wiik*
Functional Method as a Heuristic and Research Perspective: A Study in Systems Theory

4/2014   *Christian Enger Gimsø*
Narcissus and Leadership Potential - The measurement and implications of narcissism in leadership selection processes

3/2014   *Mehrad Moeini-Jazani*
When Power Has Its Pants Down: Social Power Increases Sensitivity to Internal Desires

2/2014   *Yuriy Zhovtobryukh*
The role of technology, ownership and origin in M&A performance

1/2014   *Siv Staubo*
Regulation and Corporate Board Composition

2013

9/2013   *Bjørn Tallak Bakken*
Intuition and analysis in decision making: On the relationships between cognitive style, cognitive processing, decision behaviour, and task performance in a simulated crisis management context

8/2013   *Karl Joachim Breunig*
Realizing Reticulation: A Comparative Study of Capability Dynamics in two International Professional Service Firms over 10 years

7/2013   *Junhua Zhong*
Three Essays on Empirical Asset Pricing

6/2013   *Ren Lu*
Cluster Networks and Cluster Innovations: An Empirical Study of Norwegian Centres of Expertise

5/2013   *Therese Dille*
Inter-institutional projects in time: a conceptual framework and empirical investigation

4/2013   *Thai Binh Phan*
Network Service Innovations: Users' Benefits from Improving Network Structure

3/2013   *Terje Gaustad*
Creating the Image: A Transaction Cost Analysis of Joint Value Creation in the Motion Picture Industry

2/2013   *Anna Swärd*
Trust processes in fixed-duration alliances: A multi-level, multi-dimensional, and temporal view on trust

1/2013   *Sut I Wong Humborstad*
Congruence in empowerment expectations: On subordinates' responses to disconfirmed experiences and to leaders' unawareness of their empowerment expectations

2012
9/2012   *Robert Buch*
Interdependent Social Exchange Relationships: Exploring the socially embedded nature of social exchange relationships in organizations

8/2012   *Ali Faraji-Rad*
When the message feels right: Investigating how source similarity enhances message persuasiveness

7/2012   *Marit Anti*
Commercial friendship from a customer perspective: Exploring social norm of altruism in consumer relationships and self-interest-seeking behavior

6/2012   *Birgit Helene Jevnaker*
Vestiges of Design-Creation: An inquiry into the advent of designer and enterprise relations

5/2012   *Erik Aadland*
Status decoupling and signaling boundaries: Rival market category emergence in the Norwegian advertising field, 2000-2010

4/2012   *Ulaş Burkay*
The Rise of Mediating Firms: The Adoption of Digital Mediating Technologies and the Consequent Re-organization of Industries

3/2012   *Tale Skjølsvik*
Beyond the 'trusted advisor': The impact of client-professional relationships on the client's selection of professional service firms

2/2012   *Karoline Hofslett Kopperud*
Well-Being at Work: On concepts, measurement, and leadership influence

1/2012   *Christina G. L. Nerstad*
In Pursuit of Success at Work: An Empirical Examination of the Perceived Motivational Climate, Its Outcomes and Antecedents

2011

12/2011 *Kjell Jørgensen*
Three Articles on Market Microstructure at the Oslo Stock Exchange (OSE)

11/2011 *Siri Valseth*
Essays on the information content in bond market order flow

10/2011 *Elisabet Sørfjorddal Hauge*
How do metal musicians become entrepreneurial? A phenomenological investigation on opportunity recognition

9/2011 *Sturla Lyngnes Fjesme*
Initial Public Offering Allocations

8/2011 *Gard Paulsen*
Betwixt and between: Software in telecommunications and the programming language Chill, 1974-1999

7/2011 *Morten G. Josefsen*
Three essays on corporate control

6/2011 *Christopher Wales*
Demands, designs and decisions about evaluation: On the evaluation of postgraduate programmes for school leadership development in Norway and England

5/2011 *Limei Che*
Investors' performance and trading behavior on the Norwegian stock market

4/2011 *Caroline D Ditlev-Simonsen*
Five Perspectives on Corporate Social Responsibility (CSR): Aan empirical analysis

3/2011 *Atle Raa*
Fra instrumentell rasjonalitet til tvetydighet: En analyse av utviklingen av Statskonsults tilnærming til standarden Mål- og resultatstyring (MRS) 1987-2004

2/2011 *Anne Louise Koefoed*
Hydrogen in the making - how an energy company organises under uncertainty

1/2011 *Lars Erling Olsen*
Broad vs. Narrow Brand Strategies: The Effects of Association Accessibility on Brand Performance

2010

8/2010 *Anne Berit Swanberg*
Learning with Style: The relationships among approaches to learning, personality, group climate and academic performance

7/2010 *Asle Fagerstrøm*
Implications of motivating operations for understanding the point-of-online-purchase: Using functional analysis to predict and control consumer purchasing behavior

6/2010 *Carl J. Hatteland*
Ports as Actors in Industrial Networks

5/2010 *Radu-Mihai Dimitriu*
Extending where? How consumers' perception of the extension category affects brand extension evaluation

4/2010   *Svanhild E. Haugnes*
          Consumers in Industrial Networks: a study of the Norwegian-Portuguese bacalhau network

3/2010   *Stine Ludvigsen*
          State Ownership and Corporate Governance: Empirical Evidence from Norway and
          Sweden

2/2010   *Anders Dysvik*
          An inside story – is self-determination the key? Intrinsic motivation as mediator and
          moderator between work and individual motivational sources and employee outcomes.
          Four essays

1/2010   *Etty Ragnhild Nilsen*
          Opportunities for learning and knowledge creation in practice

2009
8/2009   *Erna Senkina Engebrethsen*
          Transportation Mode Selection in Supply Chain Planning

7/2009   *Stein Bjørnstad*
          Shipshaped: Kongsberg industry and innovations in deepwater technology, 1975-2007

6/2009   *Thomas Hoholm*
          The Contrary Forces of Innovation*: An Ethnography of Innovation Processes in the Food
          Industry.*

5/2009   *Christian Heyerdahl-Larsen*
          Asset Pricing with Multiple Assets and Goods

4/2009   *Leif-Magnus Jensen*
          The Role of Intermediaries in Evolving Distribution Contexts: A Study of Car Distribution

3/2009   *Andreas Brekke*
          A Bumper? An Empirical Investigation of the Relationship between the Economy and the
          Environment

2/2009   *Monica Skjøld Johansen*
          Mellom profesjon og reform: Om fremveksten og implementeringen av enhetlig ledelse i
          norsk sykehusvesen
1/2009
          Mona Kristin Solvoll
          Televised sport: Exploring the structuration of producing change and stability in a public
          service institution

2008
7/2008   *Helene Loe Colman*
          Organizational Identity and Value Creation in Post-Acquisition Integration: The Spiralling
          Interaction of the Target's Contributive and the Acquirer's Absorptive Capacities

6/2008   *Fahad Awaleh*
          Interacting Strategically within Dyadic Business Relationships: A case study from the
          Norwegian Electronics Industry

5/2008   *Dijana Tiplic*
          Managing Organizational Change during Institutional Upheaval: Bosnia-Herzegovina's
          Higher Education in Transition

4/2008   *Jan Merok Paulsen*
Managing Adaptive Learning from the Middle

3/2008   *Pingying Zhang Wenstøp*
Effective Board Task Performance. Searching for Understanding into Board Failure and Success

2/2008   *Gunhild J. Ecklund*
Creating a new role for an old central bank: The Bank of Norway 1945-1954

1/2008   *Øystein Strøm*
Three essays on corporate boards

2007
6/2007   *Martha Kold Bakkevig*
The Capability to Commercialize Network Products in Telecommunication

5/2007   *Siw Marita Fosstenløkken*
Enhancing Intangible Resources in Professional Service Firms. A Comparative Study of How Competence Development Takes Place in Four Firms

4/2007   *Gro Alteren*
Does Cultural Sensitivity Matter to the Maintaining of Business Relationships in the Export Markets? An empirical investigation in the Norwegian seafood industry

3/2007   *Lars C. Monkerud*
Organizing Local Democracy: The Norwegian Experience

2/2007   *Siv Marina Flø Karlsen*
The Born Global – Redefined. On the Determinants of SMEs Pace of Internationalization

1/2007   *Per Engelseth*
The Role of the Package as an Information Resource in the Supply Chain. A case study of distributing fresh foods to retailers in Norway

2006
10/2006   *Anne Live Vaagaasar*
From Tool to Actor - How a project came to orchestrate its own life and that of others

9/2006   *Kjell Brynjulf Hjertø*
The Relationship Between Intragroup Conflict, Group Size and Work Effectiveness

8/2006   *Taran Thune*
Formation of research collaborations between universities and firms: Towards an integrated framework of tie formation motives, processes and experiences

7/2006   *Lena E. Bygballe*
Learning Across Firm Boundaries. The Role of Organisational Routines

6/2006   *Hans Solli-Sæther*
Transplants' role stress and work performance in IT outsourcing relationships

5/2006   *Bjørn Hansen*
Facility based competition in telecommunications – Three essays on two-way access and one essay on three-way access

4/2006   *Knut Boge*

Votes Count but the Number of Seats Decides: A comparative historical case study of 20th century Danish, Swedish and Norwegian road policy

3/2006 *Birgitte Grøgaard*
Strategy, structure and the environment. Essays on international strategies and subsidiary roles

2/2006 *Sverre A. Christensen*
Switching Relations - The rise and fall of the Norwegian telecom industry

1/2006 *Nina Veflen Olsen*
Incremental Product Development. Four essays on activities, resources, and actors

2005
6/2005 *Jon Erland Bonde Lervik*
Managing Matters:Transferring Organizational Practices within Multinational Companies

5/2005 *Tore Mysen*
Balancing Controls When Governing Agents in Established Relationships: The Influence of Performance Ambiguity

4/2005 *Anne Flagstad*
How Reforms Influence Organisational Practices: The Cases of Public Roads and Electricity Supply Organisations in Norway

3/2005 *Erlend Kvaal*
Topics in accounting for impairment of fixed asset

2/2005 *Amir Sasson*
On Mediation and Affiliation. A Study of Information Mediated Network Effects in The Banking Industry

1/2005 *Elin Kubberød*
Not just a Matter of Taste – Disgust in the Food Domain

2004
10/2004 *Sverre Tomassen*
The Effects of Transaction Costs on the Performance of Foreign Direct Investments - An empirical investigation

9/2004 *Catherine Børve Monsen:*
Regulation, Ownership and Company Strategies. The Case of European Incumbent Telecommunications Operators

8/2004 *Johannes A. Skjeltorp*
Trading in Equity Markets: A study of Individual, Institutional and Corporate Trading Decision.

7/2004 *Frank Elter*
Strategizing in Complex Contexts

6/2004 *Qinglei Dai*
Essays on International Banking and Tax-Motivated Trading

5/2004 *Arne Morten Ulvnes*
Communication Strategies and the Costs of Adapting to Opportunism in an Interfirm

Marketing System

4/2004   *Gisle Henden*
Intuition and its Role in Strategic Thinking

3/2004   *Haakon O. Aa. Solheim*
Essays on volatility in the foreign exchange market

2/2004   *Xiaoling Yao*
From Village Election to National Democratisation. An Economic-Political Microcosm
Approach to Chinese Transformation

1/2004   *Ragnhild Silkoset*
Collective Market Orientation in Co-producing Networks

2003
2/2003   *Egil Marstein*
The influence of stakeholder groups on organizational decision-making in public hospitals

1/2003   *Joyce Hartog McHenry*
Management of Knowledge in Practice. Learning to visualise competence

2002
6/2002   *Gay Bjercke*
Business Landscapes and Mindscapes in People's Republic of China. A Study of a Sino-
Scandinavian Joint Venture

5/2002   *Thorvald Hærem*
Task Complexity and Expertise as Determinants of Task Perceptions and Performance:
Why Technology-Structure Research has been unreliable and inconclusive

4/2002   *Norman T. Sheehan*
Reputation as a Driver in Knowledge-Intensive Service Firms:
An exploratory study of the relationship between reputation and value creation in
petroleum exploration units

3/2002   *Line Lervik Olsen*
Modeling Equity, Satisfaction and Loyalty in Business-to-Consumer Markets

2/2002   *Fred H. Strønen*
Strategy Formation from a Loosely Coupled System Perspective. The Case of Fjordland

1/2002   *Terje I. Våland*
Emergence of conflicts in complex projects. The role of informal versus formal governance
mechanisms in understanding interorganizational conflicts in the oil industry

2001
6/2001   *Kenneth H. Wathne*
Relationship Governance in a Vertical Network Context

5/2001   *Ming Li*
Value-Focused Data Envelopment Analysis

4/2001   *Lin Jiang*
An Integrated Methodology for Environmental Policy Analysis

3/2001   *Geir Høidal Bjønnes*
Four Essays on the Market Microstructure of Financial Markets