

Boston University

OpenBU

<http://open.bu.edu>

Boston University Theses & Dissertations

Boston University Theses & Dissertations

2017

Three essays in information and its acquisition

<https://hdl.handle.net/2144/24104>

Downloaded from DSpace Repository, DSpace Institution's institutional repository

BOSTON UNIVERSITY
GRADUATE SCHOOL OF ARTS AND SCIENCES

Dissertation

**THREE ESSAYS IN INFORMATION AND ITS
ACQUISITION**

by

CONSTANTINE S. CAVOUNIDIS

B.A., Tufts University, 2009
M.A., Boston University, 2011

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

2017

© 2017 by
CONSTANTINE S. CAVOUNIDIS
All rights reserved

Approved by

First Reader

Kevin Lang, Ph.D.
Professor of Economics

Second Reader

Barton L. Lipman, Ph.D.
Professor of Economics

Third Reader

Sambuddha Ghosh, Ph.D.
Associate Professor of Economics
Shanghai University of Finance and Economics

Acknowledgments

I want to thank Kevin Lang for hundreds of discussions, insights and pieces of advice. Despite the fact he infected me with the Labor disease, I am grateful for the countless interesting questions he has introduced me to.

Without Bart Lipman's appetite for refinement, this thesis would be far more boring to read. Without Sambuddha Ghosh's eye for elegance, there would be little left once the chaff was scraped away.

I want to thank various other faculty, family members and friends as well as my partner for their support and above all patience (which, by this point, they're either selected for or trained in).

Costas Cavounidis

THREE ESSAYS IN INFORMATION AND ITS ACQUISITION

CONSTANTINE S. CAVOUNIDIS

Boston University, Graduate School of Arts and Sciences, 2017

Major Professors: Kevin Lang, Ph.D., Professor of Economics

Barton L. Lipman, Ph.D., Professor of Economics

Sambuddha Ghosh, Ph.D., Associate Professor of
Economics, Shanghai University of Finance and
Economics

ABSTRACT

This thesis consists of three essays in economic theory, two on search models with information acquisition and one on repeated games when precise information about discount factors is unavailable.

In the first essay, I develop a model in which optimal costly information acquisition by individual firms causes adverse selection in the market as a whole. Each firm's information acquisition policy determines which customers it serves, which in turn affects the distribution of remaining customers and hence other firms' incentives. I show that when information acquisition is 'smooth', the adverse selection externality due to each firm is dampened, and in equilibrium all firms make positive profits. By contrast, with lumpy information acquisition, only a limited number of firms are profitable. I establish that my results apply to a broad class of continuous-time information acquisition processes.

The second essay explores information acquisition in labor markets. Noting that

African-Americans face shorter employment durations than similar whites, we hypothesize that employers discriminate in acquiring ability-relevant information. We construct a model with a binary information generating process, ‘monitoring’, at the disposal of firms. Monitoring black but not white workers is self-sustaining. This ‘bad’ equilibrium is not merely a matter of coordination; rather, it is determined by history and not easily reversed. The model’s additional predictions, lower lifetime incomes and longer unemployment durations for blacks, are both strongly empirically supported.

In the third essay, we investigate the possibility of repeated games equilibria that are robust to the discount factors. We prove a negative result which shows that a sizable part of the set of feasible individually rational payoffs can never be produced by such equilibria. We find the cutoff defining this region and interpret it as a limit on the ability to punish deviations when future rewards for randomization cannot be finely calibrated. Furthermore, we present a robust folk theorem to support payoffs in the complementary region with strategies that remain Subgame Perfect Nash Equilibria at all greater discount factors.

Contents

Acknowledgements	iv
Abstract	v
1 Search, Information Acquisition and Adverse Selection	1
1.1 Introduction	1
1.2 Model	6
1.2.1 Timing	6
1.2.2 Beliefs	7
1.2.3 Information Acquisition	7
1.2.4 Strategies	8
1.2.5 Mutual Information	9
1.2.6 Payoffs	10
1.3 Equilibrium	11
1.3.1 Inference	11
1.3.2 Equilibrium with unrestricted information acquisition	12
1.3.3 Equilibrium with restricted information acquisition	14
1.4 Market Size	14
1.4.1 Market size with unrestricted information acquisition	14
1.4.2 Market size with restricted information acquisition	16
1.4.3 Discussion	18

1.5	A theorem with continuous-time learning	20
1.5.1	Setting	20
1.5.2	Setup	21
1.5.3	Result	22
1.6	Conclusion	24
2	Discrimination and Worker Evaluation (with Kevin Lang)	26
2.1	Introduction	26
2.2	A simple example of churning	31
2.3	The Model	33
2.3.1	Setup	34
2.3.2	Match Quality	34
2.3.3	Job Search	36
2.3.4	Bargaining	36
2.3.5	Steady State	43
2.3.6	Solution Concept	44
2.3.7	Parametric Assumptions	45
2.4	Solution	48
2.4.1	The Non-Monitored Market	48
2.4.2	The Monitored Market	49
2.5	Implications for Labor Markets	50
2.5.1	Employment Duration	50
2.5.2	Persistence of Discrimination	51
2.5.3	Wages	52
2.5.4	Unemployment Duration	55
2.5.5	Extensions	56
2.5.6	Additional Empirical Content	62

2.6	Example	62
2.7	Conclusions	63
3	An Impossibility Theorem in Repeated Games (with Sambuddha Ghosh)	66
3.1	Introduction	66
3.2	Preliminaries	69
3.3	Discount Robustness	70
3.3.1	The restricted minmax	73
3.4	An Impossibility Theorem	74
3.5	Blackwell-Nash Equilibrium Payoffs	75
3.6	Example	78
3.7	Conclusion	79
A	Proofs of results in Chapter 1	81
A.1	Proof of Theorem 2	81
A.2	Proof of Theorem 3	81
A.3	Proof of Theorem 4	83
A.4	Proof of Theorem 5	86
A.5	Proof of Lemma 1	88
A.6	Proof of Theorem 6	88
B	Proofs of results in Chapter 2	91
B.1	Proof of Lemma 2	91
B.2	Proof of Lemma 3	92
B.3	Proof of Proposition 1	93
B.4	Proof of Proposition 2	96
B.5	Proof of Proposition 3	98

B.6 Proof of Proposition 4	100
C Proofs of results in Chapter 3	101
C.1 Proof of Theorem 7	101
C.2 Proof of Theorem 8	103
D References	105
Curriculum Vitae	112

List of Tables

3.1 Example stage game G 78

List of Figures

1.1	Optimal acceptance rates in the unrestricted information acquisition market with $\Theta = \{-1, 1\}$, $k = .4$ as a function of the probability that $\theta = 1$	16
1.2	Restricted information acquisition model when $\Theta = \{-1, 1\}$. Payoff to each strategy as a function of probability that $\theta = 1$ (left); Optimal type-wise and average acceptance probability as a function of probability that $\theta = 1$ (right).	18
1.3	Equilibrium beliefs p_n for an inactive firm n are a convex combination of the prior p_0 and points on the marginal rejection line (bold).	23
2.1	White workers' perpetual employment, black workers' churning cycle	28
2.2	Equilibrium α and β payoffs as a function of average hire quality θ	53
3.1	The payoff space for game G . Starred region: payoffs that can be delivered by a BN equilibrium. Hatched region: SPNEs exist, but they are not robust.	79

List of Abbreviations

BN	Bayes-Nash
DRSP	Discount Robust Subgame Perfect
FSIR	Feasible and Strictly Individually Rational payoff set
PDF	Probability Distribution Function
PRD	Public Randomization Device
SPNE	Subgame Perfect Nash Equilibrium

Chapter 1

Search, Information Acquisition and Adverse Selection

1.1 Introduction

Consider a venture capitalist inspecting a start-up soliciting an investment. Should she invest in the start-up, she will get a payoff that depends on its profitability. Ideally, she invests if the return on investment clears some benchmark. However, start-ups vary in type, and this variation is not immediately observable to the venture capitalist. She may aid her decision by using an information acquisition technology to learn about the start-up's expected profitability, at a cost. The venture capitalist must, naturally, keep in mind the prior distribution over start-up types before deciding whether and how to use this technology.

However, she must also consider that she may not be the first potential investor solicited. Could the start-up have been rejected for investment by other venture capitalists? What would those other venture capitalists have learned about the start-up, and how would this information influence rejections? The answers to these questions affect the venture capitalist's beliefs about the start-up, and therefore the way the technology will be used. There are clearly spill-over effects in information acquisition

decisions. However, an intriguing possibility arises: could it be that the induced update on beliefs by others' decisions is so pessimistic that the start-up is immediately rejected, without consideration? The answer to this question will turn on the nature of the information acquisition technology.

I develop a framework that allows exploration of the link between information acquisition technology and the limitations it imposes on the number of its users - the 'market size'. The framework is simple yet flexible, with the emphasis on allowing for variation of the information acquisition technology. I obtain precise market size predictions in the two most natural instances of the model, which are then augmented with a generalization to a large class of continuous-time information acquisition processes.

This paper studies a model of sequential search with asymmetric information regarding the value of transacting. A 'firm', faced with uncertainty, may choose whether, and possibly additionally what, to learn (at a cost) about the 'applicant' before deciding to accept or reject her. Each firm will choose what to learn optimally, given other firms' behavior.

Each firm does not know, when visited, if the applicant has visited other firms. As the applicant leaves the market when accepted and the same information acquisition technology is available to every firm, in equilibrium there is a form of adverse selection in the unmatched that is directly related to the technology's characteristics. Therefore, learning is interactive not because a firm is directly affected by its opponents' actions, but rather because these actions determine the level of adverse selection in the market. This endogenous adverse selection in turn affects what information the firm must acquire to protect itself. I study how the information acquisition technology, through this channel, will affect equilibrium payoffs.

An applicant sequentially visits firms until either she is accepted, or she has vis-

ited every one of the finitely many firms. The applicant is characterized by one of finitely many types, the value common to all firms of a transaction with her. A firm, upon being visited, may choose to acquire a signal whose distribution varies with the applicant's type, bearing a cost that depends on the signal's informativeness. The choice set of such signals is a parameter of interest I will vary, but is common to all firms. Conditional on the applicant's type, signals acquired by different firms are independent. Once the cost is paid and the signal observed, the firm may reject the applicant, getting no further payoff, or accept her, getting a payoff based on her type. An accepted applicant exits the model.

If the information acquisition technology is *smooth*, it allows the firm to decide what to learn about the applicant flexibly. On the other hand, if the technology is *lumpy*, it allows a firm to either become fully informed or stay uninformed (or either with some probability, via mixing).

The main insight in this paper is that when firms possess a lumpy technology, the adverse selection externality they collectively impose on one another grows with the number of firms; whereas if firms use a smooth technology, they adjust to increased competition in a way that reduces their individual externalities, keeping adverse selection in check. As a result, with restricted information acquisition only a limited number of firms can attain a strictly positive payoff, but with flexible information acquisition all firms will be strictly profitable. Therefore, the number of firms in a market is affected by the available information acquisition technology.

The intuition for these results lies in the distribution of types rejected by each firm. When the technology is smooth, the marginal cost of additional information acquired about rejected applicants is equal to the marginal benefit. When a firm is uncertain whether a visiting applicant has been rejected at another firm or not, its beliefs about her are therefore at least a bit better than about the rejected; it

is therefore beneficial to acquire at least some information. This will in equilibrium imply strictly positive profits. When the technology is lumpy, with enough firms, beliefs about rejected applicants are not marginal but rather excessively bad, so that the number of firms that can profitably become informed is limited.

If the cost of the fully informative signal was fixed, one would be concerned that the result is driven by the fact that, as adverse selection grows worse with the number of firms, the information provided by the signal is diminishing but the cost is not. To avoid this trap, I allow the cost of the signal to depend continuously on the informational content, which in turn depends on the prior. With information costs based on Shannon's mutual information, firms only 'pay for what they get' - as the ex-ante uncertainty goes to zero, so does the signal cost.¹ The result is also not driven by discontinuity in the lumpy case either, as mixing allows a firm to produce a signal with any probability of being fully informative. The 'lumpiness' of the information comes not from an all-or-nothing choice of signal, then, but rather from the extreme nature of the informed posteriors.

In a sense, flexible learning, combined with continuous costs based on mutual information, 'dampens' adverse selection. The marginality of a rejected applicant ensures that the updating on being visited is not as pessimistic as in the lumpy case. As the number of firms expands in the smooth learning model, they each learn less, effectively always leaving surplus for the rest. However, the main results are not dependent on the mutual information cost function, and a generalization of the 'smooth' case to a very broad setting is considered in section 1.5.

The model in this paper can apply to a variety of contexts. Other than investment in start-ups, applications include hiring in labor markets with non-negotiable (e.g. minimum) wages, insurance markets, and even human mating. It can be used, for

¹This is going to require a somewhat non-standard formulation of the game, as beliefs do not typically enter payoffs directly.

instance, to think about the effects of the prohibition of health insurance rejections in the Affordable Care Act. While one may think the prevention of screening by insurers would cause adverse selection, the results in this paper show that screening creates a different, potentially worse, kind of adverse selection. The model does not address price-setting, but it can accommodate pricing based on variable observable flow outside options - as long as it is not determined after information acquisition. One can show the paper's main results continue to hold, *mutatis mutandis*, even when market size affects such prices.

The present paper builds heavily on work on search with adverse selection. Inderst (2005) introduces a search model where contracts are used to separate types in equilibrium. Lauer mann and Wolinsky's (2016) search model determines the level of information aggregation when exogenous, conditionally i.i.d. signals of a buyer's type are available to sellers. Zhu (2012) uses a decentralized search market with a similar signaling technology to model exploding offers in over-the-counter markets. These last two models strongly develop the intuition for the *solicitation effect*, the fact that being visited at all can be a negative signal, as it speaks of possible rejection from other potential transactions. However, the information acquisition process is taken to be both free and exogenous, with the focus being on prices.

The model in this paper has certain parallels with common values auctions, in particular in the existence of a probabilistic version of the Winner's Curse; in equilibrium, an accepted applicant will have had worse signal realizations at all firms visited previously. Persico (2000) examines optimal information acquisition in auctions, finding that different auction procedures induce different levels of information acquisition. Bergemann, Shi and Välimäki (2009) study auctions with interdependent values and a binary information choice. In that context they model the number of bidders who choose to become informed much in the same spirit as this paper addresses market

size.

Finally, the present model is heavily influenced by the literature on rational inattention. The model uses as a measure of informativeness mutual information, built on Shannon (1948)'s notion of entropy, used in economics as a cost of information measure since Sims (2003). Strategic interaction with rational inattention has been studied by both Yang (2015) and Denti (2016) in the context of coordination games. Gentzkow and Kamenica (2014) write a model in which the informed party must pay to disclose information, rather than to acquire it. Ravid (2016) examines rational inattention in a bargaining model with one-sided offers. Additionally, Matějka and McKay (2015) study discrete choices with rational inattention, and special attention is paid to binary choice by Woodford (2008).

The rest of this paper proceeds as follows. Section 1.2 introduces the main model. Section 1.3 contains existence theorems for the model's main cases, while section 1.4 provides the main market size results. Section 1.5 generalizes the main results to a broad class of settings with continuous-time learning, and section 1.6 concludes.

1.2 Model

There exist N identical firms and a single applicant. The applicant is characterized by a private type, θ , which is distributed according to prior $p_0(\cdot)$ with full support on finite, non-singleton $\Theta \subset \mathbb{R} \setminus \{0\}$. θ is the net benefit to a firm of accepting the applicant. Crucially, I assume $\min \Theta < 0 < \max \Theta$; firms want to accept some, but not all, applicants.

1.2.1 Timing

Nature chooses the type θ of the applicant according to p_0 , and a *visit order* σ from the set of permutations of the N firms equiprobably. Then, the game proceeds in up

to N stages, starting with stage 1.

In stage n , firm $\sigma^{-1}(n)$ is visited. The visited firm may then choose to acquire information about θ using the available technology. Once information is acquired, the firm may choose to accept or reject the applicant. If the applicant is accepted or $n = N$, the game ends. If the applicant is rejected and $n < N$, stage $n + 1$ follows.

That is, the applicant visits firms in a (uniform) random order, until either he is accepted, or he has been rejected by every firm. When the applicant visits a firm, the firm sees only that an applicant has arrived, not her history, the visit order or (equivalently) the date.

1.2.2 Beliefs

Each firm acts at a single information set. That is to say, firms are not aware of the applicant's history of visiting other firms. As a history with rejections is bad information for the applicant, the applicant would not disclose her history even if the model were augmented with a cheap-talk stage.

In equilibrium, each firm n has beliefs over the nodes in its information set whose marginal over θ is not, in general, $p_0(\cdot)$ but rather the *equilibrium belief* p_n . The equilibrium belief is computable from the prior, the applicant's random visit path, and the equilibrium acceptance probability of each type of applicant at each rival firm via Bayes' rule. These distributions are endogenous variables of the model, as they are produced by firms' strategies.

1.2.3 Information Acquisition

When the applicant visits a firm, that firm can then acquire information about the applicant's type. A *signal structure* is a set of conditional distributions $\{g(\cdot|\theta)\}_{\theta \in \Theta}$

for a signal $s \in S$, where S is a finite alphabet.² A collection of signal structures $\mathcal{G} = \{\{g_i(\cdot|\theta)\}_{\theta \in \Theta}\}_{i \in I} = \{g_i(\cdot|\cdot)\}_{i \in I}$ is called an *information menu*. Signal structures will vary in cost, so a more informative signal will not always be preferable.

I consider two different information menus. The first is the *unrestricted* information menu, $\mathcal{G}_U = (\Delta S)^\Theta$. As it is comprised of all conditional distributions for the signal, it is the largest possible set of signal structures. The second is the *restricted* information menu, $\mathcal{G}_R = \{g_{no}, g_{all}\}$, comprised of g_{no} with $g_{no}(s|\theta) = g(s)$, a completely uninformative signal structure, and the completely informative signal structure g_{all} , with $g_{all}(s|\theta) = \delta_\theta$, so that the signal is always the same as the applicant's type.

1.2.4 Strategies

A *single-structure* strategy for firm n is a pair (g, a) where

- $g \in \mathcal{G}$ is a signal structure; and
- $a : S \rightarrow [0, 1]$ maps the signal to a probability of accepting the applicant.

The information menu \mathcal{G} is common to all firms. As the firm's problem may have multiple solutions, the firm may pursue a strategy that mixes over signal structures. As the posterior distribution of types conditional on a signal realization will depend on the signal structure that produced the signal in question, the conditional acceptance probability must be free to vary with the chosen signal structure. Therefore the set of mixed strategies is defined as the set of distributions over single-structure strategies, $\Delta(\mathcal{G} \times [0, 1]^S)$, rather than $\Delta\mathcal{G} \times \Delta([0, 1]^S)$.

²In general, a signal alphabet of cardinality equal to that of the set of resulting actions is sufficient in unrestricted models of rational inattention. I assume $|S| \geq |\Theta|$.

1.2.5 Mutual Information

Shannon (1948) lays out a particular measure of uncertainty blind to economic consequences. For an arbitrary distribution P with finite support X , the Shannon entropy is defined as

$$H(P) \equiv - \sum_{x \in X} P(x) \ln P(x) \quad (1.1)$$

and can be interpreted, if $-\ln P(x)$ is thought of as the *surprisal*³ in observing the realization x , as the *average surprisal* in P .⁴

If a signal distributed by $g(\cdot|\cdot)$ provides a lot of information about θ , it reduces the uncertainty about θ by on average concentrating its posterior distribution $p_n(\cdot|\cdot)$. Our measure of the informativeness of a signal structure is the expected reduction in the average surprisal achieved by updating the equilibrium belief p_n using the signal. This expected reduction in the entropy of p_n by observing a draw from g is known as the *mutual information* between p_n and g and is computed as

$$M(p_n, g) = H(p_n) - E_s[H(p_n(\cdot|s))] \quad (1.2)$$

where the expectation is taken according to g .

Alternatively, mutual information can be defined as the expected Kullback-Leibler divergence between the prior and the posterior distributions of θ :

$$M(p_n, g) = \sum_{s \in S} \left(\sum_{\theta \in \Theta} p_n(\theta) g(s|\theta) \right) D_{KL}(p_n(\cdot|s) || p_n(\cdot)). \quad (1.3)$$

This defines mutual information in terms of the extent to which it will in expectation

³The notion of surprisal originates in the information theory literature. Intuitively, observing an ex-ante unlikely event conveys more information than observing one initially thought to be likely. The log functional form (uniquely) allows additivity over intersections of independent events as $\ln(P(A \cap B)) = \ln(P(A) \cdot P(B)) = \ln P(A) + \ln P(B)$.

⁴For this and all other purposes, this paper uses the common convention $0 \ln 0 = 0$ in accordance with the limit.

shift the posterior, where the notion of posterior-shifting is given by the divergence.

Starting with Sims (2003), it is common for work on rational inattention to use mutual information as the cost of information. The main advantages are the treatment of uncertainty as generic and separable from economic consequences, the agreement with Blackwell-informativeness, the logit-like discrete choice probabilities as highlighted by Matěka and McKay (2015) and the ability to price all signal structures.

1.2.6 Payoffs

A firm n pays for signal structure g a cost equal to a multiple k of the mutual information $M(\cdot, \cdot)$ between its equilibrium belief of the applicant's type p_n and that of the signal g .⁵

A firm that accepts an applicant with type θ gains utility θ for doing so, whereas rejecting any applicant gives a payoff of 0. Therefore, a firm n with beliefs p_n about a new applicant will choose what information to acquire via g and how to respond to it via the conditional acceptance probability a so as to maximize profits. I normalize payoffs so that they are ex-post of applicant arrival. Firm n 's payoff, given equilibrium belief p_n and single-structure strategy $(g, a) \in \mathcal{G} \times [0, 1]^S$

$$a(s)\theta - kM(p_n, g). \tag{1.4}$$

As this quantity depends directly on beliefs, the game defined above is a Psychological Game as described in Geanakoplos, Pearce and Stacchetti (1989). In effect, the requirement that firms only pay for the amount of information they acquire means that the signal's cost must depend on how much it shifts the firm's beliefs from the equilibrium beliefs.

⁵ An abstract generalization of the model's main results to other information acquisition technologies can be found in section 1.5.

1.3 Equilibrium

An equilibrium for the N -firm model is given by mixed strategies for all firms, with firm n playing $F_n \in \Delta(\mathcal{G} \times [0, 1]^S)$, that form a Nash Equilibrium, given that each firm's beliefs about arriving applicants follow Bayesian inference.

In equilibrium, each firm n will therefore choose a strategy that solves

$$\max_{F_n \in \Delta(\mathcal{G} \times [0, 1]^S)} \int_{\mathcal{G} \times [0, 1]^S} \left[\sum_{\theta \in \Theta} p_n(\theta) \sum_{s \in S} g(s|\theta) a(s) \theta - kM(p_n, g) \right] dF_n(g, a). \quad (1.5)$$

1.3.1 Inference

Consistency requires firms' beliefs $(p_n)_{n \leq N}$ are derived from firms' strategies $(F_n)_{n \leq N}$ and the prior p_0 via Bayes' Rule. A type θ applicant has a probability

$$A_m(\theta) \equiv \int_{\mathcal{G} \times [0, 1]^S} \sum_{s \in S} g(s|\theta) a(s) dF_m(g, a) \quad (1.6)$$

of being accepted by firm m . The applicant's visiting sequence is a uniform draw from the permutations of N , with each permutation getting probability $1/N!$. Therefore, the probability a type θ applicant ever visits firm n is $1/N!$ times the sum over these permutations of the probability that the applicant is rejected at every previously visited firm:

$$\frac{1}{N!} \sum_{\sigma \in \text{perm}(N)} \prod_{m: \sigma(m) < \sigma(n)} (1 - A_m(\theta)). \quad (1.7)$$

Thus the posterior probability that an applicant arriving at firm n is of type θ is the prior $p_0(\theta)$ times the sum over all visiting sequences of the probability that θ visits firm n , over the same for all types:

$$p_n(\theta) = \frac{p_0(\theta) \sum_{\sigma \in \text{perm}(N)} \prod_{m: \sigma(m) < \sigma(n)} (1 - A_m(\theta))}{\sum_{\theta' \in \Theta} p_0(\theta') \sum_{\sigma \in \text{perm}(N)} \prod_{m: \sigma(m) < \sigma(n)} (1 - A_m(\theta'))}. \quad (1.8)$$

I denote by $\hat{p}_n : [0, 1]^{(N-1)|\Theta|} \rightarrow \Delta\Theta$ the *belief function for firm n* taking all other firms' acceptance probability vectors $(A_m)_{m \neq n}$ into firm n 's posterior distribution for θ using Bayes' rule:

$$\hat{p}_n((A_m)_{m \neq n}) = \left(\frac{p_0(\theta) \sum_{\sigma \in \text{perm}(N)} \prod_{m: \sigma(m) < \sigma(n)} (1 - A_m(\theta))}{\sum_{\theta' \in \Theta} p_0(\theta') \sum_{\sigma \in \text{perm}(N)} \prod_{m: \sigma(m) < \sigma(n)} (1 - A_m(\theta'))} \right)_{\theta \in \Theta}. \quad (1.9)$$

As the expression in the numerator is a polynomial in the $A_m(\theta)$ s and the denominator is a sum of such polynomials, each of which is bounded below by $\frac{1}{N}p_0(\theta)$, \hat{p}_n is a continuous function. Continuity of the belief functions will be useful in showing existence of equilibria.

1.3.2 Equilibrium with unrestricted information acquisition

I first turn to the case when $\mathcal{G} = \mathcal{G}_U$. Let two signal structures g, g' be *equivalent* if they induce the same distribution over posteriors on θ . It is useful here to use the characterization of binary choice with mutual information costs given by Woodford(2008). This will provide both uniqueness and continuity of best responses in p_n .

Theorem 1 (Woodford 2008). *In the model with unrestricted information acquisition, for a firm holding beliefs p_n , there is a unique up to equivalence optimal signal structure. Each signal in the support of an optimal signal structure leads to a pure choice, that is $a(s) \in \{0, 1\}$. Also, the induced optimal acceptance probabilities $A_n(\theta)$ satisfy*

- if $\sum p_n(\theta) e^{-\frac{\theta}{k}} \leq 1, \forall \theta$ $A_n(\theta) = 1$

- if $\sum p_n(\theta)e^{\frac{\theta}{k}} \leq 1, \forall \theta$ $A_n(\theta) = 0$
- otherwise, $A_n(\theta)$ is given by

$$\frac{A_n(\theta)}{1 - A_n(\theta)} = \frac{\bar{A}}{1 - \bar{A}} e^{\frac{\theta}{k}}, \text{ where} \quad (1.10)$$

$$\bar{A} = \sum p_n(\theta)A_n(\theta). \quad (1.11)$$

Importantly, we can use the theorem to exclude mixtures over non-equivalent signal structures (as they would each be optimal but non-equivalent). Mixing over equivalent signal structures merely amounts to randomizing *how* information is acquired, but not *what* is learned; the induced distributions over posteriors are the same. Effectively, equivalent signal structures are mere relabellings.

Furthermore, Woodford (2008) shows that optimal actions can be summarized as $(A_n(\theta))_{\theta \in \Theta}$, the optimal acceptance probability for each type; and that those are unique. In effect, the theorem allows us to restrict to a much smaller strategy space when looking for best responses. As this space is compact, equilibrium existence will be easy to show.

It is useful at this point to define the function that takes firm n 's beliefs p_n into optimal acceptance probabilities as $\hat{A}_n : \Delta\Theta \rightarrow [0, 1]^{| \Theta |}$ using the rule given in (1.10) and (1.11).

Theorem 2. *For any $N \in \mathbb{N}$, there exists an equilibrium in the model with N firms and unrestricted information acquisition.*

All proofs are in the appendix. As I have shown that we only need to consider a compact and convex action space, and Theorem 1 gives a continuous best response, the proof of Theorem 2 is a straightforward application of Brouwer's Fixed Point Theorem.

1.3.3 Equilibrium with restricted information acquisition

When $\mathcal{G} = \mathcal{G}_R$, when the firm becomes informed, it is optimal to accept all $\theta \geq 0$. Thus the only relevant choice is whether to acquire full or no information, with each firm choosing a probability of playing each of these two actions. However, (1.9) implies that the cost of becoming informed depends non-linearly on opponent strategies, the Nash Existence Theorem does not apply. However, the restricted action space is compact and allows an application of Kakutani's Fixed Point Theorem once the best response correspondences are shown to be closed-graph.

Theorem 3. *For any $N \in \mathbb{N}$, there exists an equilibrium in the model with N firms and restricted information acquisition.*

Neither of the existence theorems presented requires that the sequence in which the applicant visits firms is uniformly distributed. What is required is that there is a strictly positive probability for each firm to be the first visited, so that Bayes' rule is well-defined for every strategy profile⁶

1.4 Market Size

So far, I have kept the number of firms in the market fixed. This paper's main results pertain to the number of firms that can profit in the market.

1.4.1 Market size with unrestricted information acquisition

For the market with unrestricted information acquisition and N firms, two possibilities arise. If the proportions and relative value of the types in the prior are too low, low enough that even at the prior a monopolist is unwilling to acquire any information or ever accept applicants, then the only equilibrium will be one in which no firms profit. This will occur iff $\sum_{\theta} p_0(\theta) e^{\frac{\theta}{k}} \leq 1$. Firms that have no profits acquire no

⁶That is, (1.9) must have a positive denominator.

information⁷. If a monopolist cannot profit, it is quite clear that neither can more firms, with the associated negative externality each imposes on the others. This rather trivial case aside, Theorem 4 will show that as long as a monopolist would be profitable, any number of firms are profitable, as equilibrium beliefs p_n for any firm n will also satisfy $\sum_{\theta} q_m(\theta) e^{\frac{\theta}{k}} > 1$.

This guaranteed profitability occurs because in any equilibrium, any profitable firm's *rejected* applicants will have a probability distribution q_m so that $\sum_{\theta} p_n(\theta) e^{\frac{\theta}{k}} = 1$ holds exactly. Then, assuming a firm n gets a payoff of 0 and thus rejects all applicants, the pool of potential applicants it receives is comprised of rejects from some other firm m with type distribution q_m as well as first-time applicants, with a type distribution p_0 ; thus, the average applicant p_n will be a convex combination of q_m 's and p_0 and as such will satisfy $\sum_{\theta} p_n(\theta) e^{\frac{\theta}{k}} > 1$. Therefore, the assumption that firm n rejects all applicants and gets a payoff of 0 is contradicted.

This surprising result occurs because rejected applicants have been *marginally* learned about. That is, optimal learning in the unrestricted information acquisition setting requires that at the time of rejection, the rejecting firm has beliefs about the applicant on the boundary of the set of profitable beliefs. If a firm is uncertain whether an applicant has just been rejected at some other firm or is a new draw, any interior probability weights on these two cases will produce beliefs about the applicant's type that allow the firm to make a profit.

Theorem 4. *If a monopolist would be profitable, for every $N \in \mathbb{N}$, in every equilibrium of the model with unrestricted information acquisition and N firms each firm gets a strictly positive payoff.*

The theorem implies that if there exist a finite number of potential firms, each of which can decide to enter at sufficiently low cost, in equilibrium, they all will. If the

⁷The converse is not true; it is possible that a firm acquires no information but profitably accepts the applicant.

choice of information structure is unrestricted, the information acquisition technology does not erect a barrier to entry.

Figure 1.1 shows how acceptance rates react to beliefs p_n in a two-type market. As the number of firms expands, the average acceptance probability for each firm \bar{A}_n will approach 0, but reach it only in the limit; for any finite number of firms, information will be worth acquiring for each firm. The figure illustrates how a lower prior forces a firm to be more cautious, decreasing the acceptance rates for both types, while increasing the ratio of the probability a good type is accepted to that a bad type is accepted. As the number of firms increases, each firm moderates its acceptance probabilities, reducing its individual contribution to the market-level adverse selection.




Figure 1.1: Optimal acceptance rates in the unrestricted information acquisition market with $\Theta = \{-1, 1\}$, $k = .4$ as a function of the probability that $\theta = 1$.

1.4.2 Market size with restricted information acquisition

I now turn to the case when $\mathcal{G} = \mathcal{G}_R$, when firms can choose to acquire either all or no information. The main assumption I make is $\Pi = \sum_{\theta > 0} p_0(\theta)\theta + k \sum_{\theta \in \Theta} p_0(\theta) \ln p_0(\theta)$,

which means that a monopolist would choose to become informed by purchasing the perfectly informative signal structure. As with more firms the equilibrium beliefs are worse than the monopolist's, a profitable firm in this market must choose to become informed.

As the number of profitable firms in the market grows, equilibrium beliefs get worse. However, this may be uncertainty-reducing, as it increasingly concentrates the prior on the negative types; therefore, the mutual information-based costs of acquiring information may be decreasing.

Theorem 5 shows that the expected net benefit to acquiring information (rather than rejecting all applicants) as a function of the number of firms that acquire information is decreasing and eventually becomes negative.

It therefore shows that in a market with few firms, all will be profitable. For a larger market, it provides an upper bound on the number of profitable firms, though equilibria exist where mixed strategies make even fewer firms profitable. However, if firms only enter if they expect to make strictly positive profits, Theorem 5 provides a precise prediction about the number of firms.

Theorem 5. *When $k \sum_{\theta \in \Theta} p_0(\theta) \ln p_0(\theta) > \sum_{\theta < 0} \theta p_0(\theta)$, there exists an $\bar{N} \in \mathbb{R}$ such that*

- (a) *in the unique equilibrium of the model with restricted information acquisition and $N < \bar{N}$ firms, every firm gets a strictly positive payoff and*
- (b) *in every equilibrium of the model with restricted information acquisition and $N > \bar{N}$ firms, at most \bar{N} firms get a strictly positive payoff.*

Figure 1-2 displays, on the left, the payoffs to the strategies of rejecting the applicant without information, accepting the applicant without information, and becoming informed in order to accept if $\theta > 0$. As the probability of the 'good' type varies, the optimal strategy changes. If the prior is concentrated, whether it is good or bad, it

is not worth becoming informed; it is in the middle region where gathering information is optimal. The right-hand panel shows how choosing the optimal strategy as displayed on the left results in acceptance probabilities for each type, and the average acceptance probability. As the equilibrium belief becomes worse with the addition of more firms, so long as firms are profitable, their behavior does not change; therefore, the effect on adverse selection is not dampened by changes in equilibrium information acquisition.




Figure 1-2: Restricted information acquisition model when $\Theta = \{-1, 1\}$. Payoff to each strategy as a function of probability that $\theta = 1$ (left); Optimal type-wise and average acceptance probability as a function of probability that $\theta = 1$ (right).

1.4.3 Discussion

Comparison

The theorems in this section have shown that the information acquisition technology limits the number of profitable firms in the restricted information market, but not in the unrestricted information market. This results from the fact the restricted information market forces firms to buy ‘chunky’ information, which produces a larger negative externality on other firms. By contrast, the unrestricted information acquisition market allows for such fine-tuning that rejected applicants are not excessively

bad draws - merely *marginally* so. This sort of fine-tuning is a property shared by a variety of continuous-time learning procedures, detailed in section 1.5, and is by no means an artefact of mutual information costs.

I assumed only a single applicant. It is useful to think of the applicant side of the market changing in size, and tracking the effect on the ‘market’ size. Crucially, both market size theorems address the profitability of firms *per applicant*; there are no fixed costs. Therefore, surprisingly, the market size prediction is invariant to the number of applicants when entry is free. If we consider costly entry, as the number of applicants grows, the restricted information market will see an increase in the number of active firms if initially less than $\lfloor \bar{N} \rfloor$; but once there, it remains there. The unrestricted information acquisition market with entry costs, on the other hand (considering the symmetric equilibrium at each N) will grow without bounds as the number of applicants increases.

Implicit Information Cascades

In the restricted information acquisition market, with too many (more than \bar{N}) firms, for some firms a visit leads to inference that the applicant is too likely to have been rejected elsewhere to be worth considering. Although the actions of other firms are not directly observable, information is carried by their strategies and Bayes’ Rule. Furthermore, as the inferred action is rejection, and that action is then copied by the inactive firm, the effect is similar to an information cascade, as in Bikhchandani, Hirshleifer and Welch (1992). The inference, rather than the direct observation, of the action being copied gives rise to what might be called an *implicit information cascade*. Firms other than the \bar{N} allowed by Theorem 5 ‘copy’ the action they believe has been taken with high probability by other firms, without even paying to view their signal, in the same way that agents ignore their signals in BHW.

By contrast, such a cascade does not occur in the market with unrestricted infor-

mation acquisition. Rejections are not as informative as in the restricted information market, and the additional type uncertainty generated by the visit order leaves some surplus for an additional firm.

Note that clearly, both technologies would lead to cascades if the visit order were known. In the restricted information case, an applicant known to have been rejected by another firm is known to have $\theta < 0$ and therefore not worth considering. Similarly, if a firm is observed to reject the applicant in the market with unrestricted information acquisition, the posterior $q(\cdot)$ for that applicant now has $\sum_{\theta \in \Theta} q(\theta)e^{\theta/k} = 1$ which means that the optimal action for all subsequent firms is to learn nothing and reject.

1.5 A theorem with continuous-time learning

Recent advances in the literature explore continuous-time learning for economic settings. Drift-diffusion models have been used to discuss two-alternative choice as in Fudenberg, Strack and Strzalecki (2015) and Hébert and Woodford (2016) consider a broad class of continuous time information acquisition technologies. I will present a theorem providing sufficient conditions for the outcome of Theorem 4 to apply in settings where firms have such a technology available. Moreover, by stating assumptions in terms of the properties of the solution to an unmodeled continuous-time information acquisition problem, I will arrive at a very general characterization, one that in fact implies Theorem 4.

1.5.1 Setting

In this section, I take a general statement of the outcome of a continuous time learning process without modeling it in detail. It may represent an exogenous process that the firm passively observes in preparation for a decision, or it may be driven by the firm optimally choosing how to learn. Details such as costs enter only implicitly. I

denote by X the resulting stochastic process that describes in continuous time the firm's beliefs about θ given this learning. I give three conditions: that X 's path is almost surely continuous with respect to time, that the firm makes its choice once X exits some convex continuation set that is not growing over time, and that a firm that engages in learning will in expectation make a profit. These conditions will be enough to ensure a general version of Theorem 4 holds.

The rather sparse assumptions on this continuous-time result allow for a variety of commonly used learning processes. Drift-diffusion models with constant (increasing) time costs, for instance, satisfy. The mutual information model described earlier in this paper can be sequentialized as well, in a way that fits these assumptions. These stipulations do not hold, however, if X is the result of a fully revealing Poisson signal - this latter case is the equivalent to the restricted information acquisition model.

1.5.2 Setup

Let $\{X_t : t \in \mathbb{R}_+\}$ be a continuous-time martingale process in the probability simplex over Θ denoting beliefs resulting from the firm's optimal acquisition process and the true value of θ . Its initial value is $x_0 = p_n$. The process results from some unmodeled policy that depends on the current value x_t , the time t , or both. Suppose furthermore that this optimal policy contains a stopping rule, given by a closed stopping region $\mathcal{S}(t) \subseteq \Delta\Theta$ that may vary with time. Assume that $\forall t, \mathcal{S}(0) \subseteq \mathcal{S}(t)$ and that \mathcal{S} is a closed-graph correspondence of t . The stopping rule requires that the stopping time τ satisfies $\tau = \inf\{t | x_t \in \mathcal{S}(t)\}$. Once the stopping rule is triggered at some τ , the firm is assumed to make a decision to accept or reject the applicant based on the current beliefs x_τ . Assume furthermore that if $x_0 \notin \mathcal{S}(0)$, the firm's ex-ante expected payoff is positive.

1.5.3 Result

This section's result will rely, as did Theorem 4, on the intuition that rejections are in some sense *marginal*. To generate this property, assume that the path of X_t is almost surely a continuous function of t . Then, almost surely the posterior for rejections will be marginal, and a similar theorem holds.

Lemma 1. *If $p_n \in \Delta\Theta \setminus \mathcal{S}(0)$, if X_t is almost surely continuous in t , then almost surely for terminal τ , $x_\tau \in \overline{\Delta\Theta \setminus \mathcal{S}(0)}$.*

The optimal learning rule therefore implies that when learning terminates, the rejecting firm's beliefs about the applicant lie in the closure of the continuation set $\Delta\Theta \setminus \mathcal{S}(0)$.

Theorem 6. *If X_t is almost surely continuous, almost surely terminates and $\Delta\Theta \setminus \mathcal{S}(0)$ is convex, if $p_0 \in \Delta\Theta \setminus \mathcal{S}(0)$, for every $N \in \mathbb{N}$, in every equilibrium of the model with N firms, every firm gets a positive expected payoff.*

The theorem is proven by contradiction. If a n firm is inactive and gets 0 payoff, it must be that its initial beliefs p_n about applicants are in the stopping set. However, each firm's posterior about its rejected applicants will (almost surely) lie in the closure of the initial continuation set $\Delta\Theta \setminus \mathcal{S}(0)$ due to the continuity of the stochastic process. Then, p_n is just an appropriately-weighted convex combination of the prior p_0 which is in the interior of the initial continuation set, and points in the continuation set's closure. Given a convex initial continuation set $\Delta\Theta \setminus \mathcal{S}(0)$, p_n must lie in its interior, and by hypothesis a belief in the initial continuation set corresponds to a positive expected payoff. Figure 1-3 illustrates the theorem for an arbitrary process when $|\Theta| = 3$.

As the unrestricted information acquisition one-shot model with mutual information costs in this paper can be written as the output of optimal learning in a continuous-time model with mutual information-restricted learning per unit time and

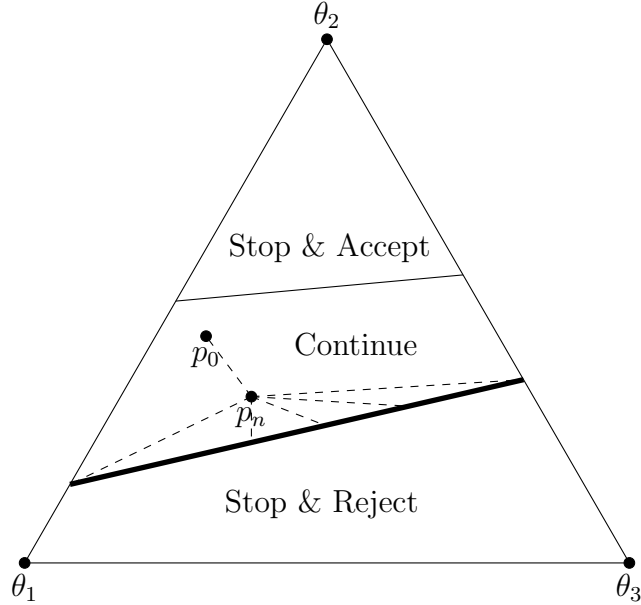


Figure 1-3: Equilibrium beliefs p_n for an inactive firm n are a convex combination of the prior p_0 and points on the marginal rejection line (bold).

a constant time cost, Theorem 6 implies Theorem 4. The theorem also applies to a wide class of learning processes described in Hébert and Woodford (2016), drift diffusion models with non-decreasing time costs and other settings. To attain such generality, however, Theorem 6 must be written in terms of the solution, not the primitives, of some unspecified model. Nevertheless, an almost surely continuous path for the learning process and a convex continuation set are both easy to verify for a particular model. The market size result is therefore generalizable to a fairly broad setting and not particular to the unrestricted information acquisition process or a particular cost function.

Transporting the restricted information case into continuous time is also feasible, as it can be modeled by a fully revealing variable with a Poisson arrival process and constant time costs. Firms will choose to either accept the applicant without information, reject the applicant without information, or wait until the arrival of the

fully informative signal⁸, equivalently to the behavior in the restricted information one-shot case. Too many firms using the waiting strategy would push equilibrium beliefs for other firms into the region where profits cannot be attained; this limits the number of firms.

1.6 Conclusion

In this paper I study a model of information acquisition in a search market. As each firm's actions affect the aggregate level of adverse selection, I explore how these externalities compound as the number of firms expands. Unexpectedly, the model shows that, with more firms, flexibility in information acquisition dampens the adverse selection externality each causes, as they will produce marginal posteriors. On the other hand, chunky information acquisition results in more pessimistic posteriors, so that with enough firms, adverse selection grows to the point of eradicating profits. I have argued that these results are relatively robust within a broad class of information acquisition technologies.

I have shown that as a result of this form of adverse selection, the information acquisition technology in a market can have strong implications about the number of firms that will be active. A rich information menu will allow all firms to be profitable, so with low enough entry costs many firms will enter. On the other hand, a chunky information acquisition technology will limit the number of potential entrants.

This research highlights the need to study the qualitative aspects of information acquisition as they appear in different contexts. It investigates what can loosely be considered a novel type of barrier to entry, and its determinants.

Finally, the model invites a broad class of applications in investment, labor, insurance, and other markets. As each market is ostensibly endowed with its own

⁸Naturally, an indifferent firm could also employ a mixed strategy, or wait a finite amount of time, which is also equivalent to mixing.

information acquisition technology, it is thus of paramount importance to consider the differential effects policies may have on otherwise similar markets.

Chapter 2

Discrimination and Worker Evaluation (with Kevin Lang)

2.1 Introduction

Many African-Americans believe black workers ‘don’t get second chances’¹ or that they face additional scrutiny in the workplace. Similarly, black workers are admonished to be ‘twice as good’² in order to succeed. If black workers are subject to higher standards or scrutinized more heavily, we expect this to be reflected in more separations.

Indeed, the data support the idea of shorter employment duration³ for black workers. Bowlus, Kiefer and Neumann (2001) detect and ponder the disparity in job destruction rates; Bowlus and Eckstein (2002) estimate⁴ that young black male high school graduates had roughly 2/3 the job spell duration of their white counterparts, despite more of their job spells ending in unemployment. Both papers assume an

¹This assertion can be found in a range of occupations including football coaching (Reid, 2015), music and films (*The Guardian*, 2014) as well as more generally (Spencer, 2014).

²Coates, Ta-Nehisi (2012) and Mabry, Marcus (2012)

³Throughout this paper we refer to employment duration by which we mean the length of an employment spell rather than job duration by which we mean the time a worker spends with a particular employer. Job duration depends on, among other factors, the arrival rate of outside offers. Our model abstracts from job-to-job transitions.

⁴Using the NLSY data for 1985 and 1988.

exogenously higher separation rate for black workers to fit their models to the data. Lang and Lehmann (2012) show that differences in unemployment duration alone are insufficient to account for the black/white unemployment gap and therefore that black workers' employment stints are shorter. This aspect of labor discrimination has thus far eluded theoretical explication.

In this paper, our proposed explanation for differential employment durations is, in its broadest sense and consistent with the aforementioned observations, that firms discriminate in acquisition or use of productivity-relevant information. That is, firms either learn differently about black workers or, when information regarding ability is received, they condition how they act on it on workers' race. Crucially, we establish that such discrimination can be self-perpetuating.

We develop a model in which differences in job duration arise naturally and their relation to skill is plausible. The essence of our model is that, because black workers are more closely scrutinized, a larger share of low-performance workers will separate into unemployment. As a result, since productivity is correlated across jobs, the black unemployment pool is 'churned' and therefore weaker than the white unemployment pool. Since workers can, at least to some extent, hide their employment histories, race serves as an indicator of expected worker productivity. This in turn makes monitoring newly hired black (but not white) workers optimal for firms. Figure 1 illustrates employment in the two labor markets.

There are multiple equilibria in our model, a property it shares with models of rational stereotyping or self-confirming expectations (Coate and Loury, 1993). However, in our model discrimination is not simply a product of coordination failure; instead, history matters. A group that begins with a low level of skills for which only the bad (monitoring) equilibrium exists will remain in that equilibrium even if its skill level rises to a level consistent with the existence of both the good and bad equilibria.

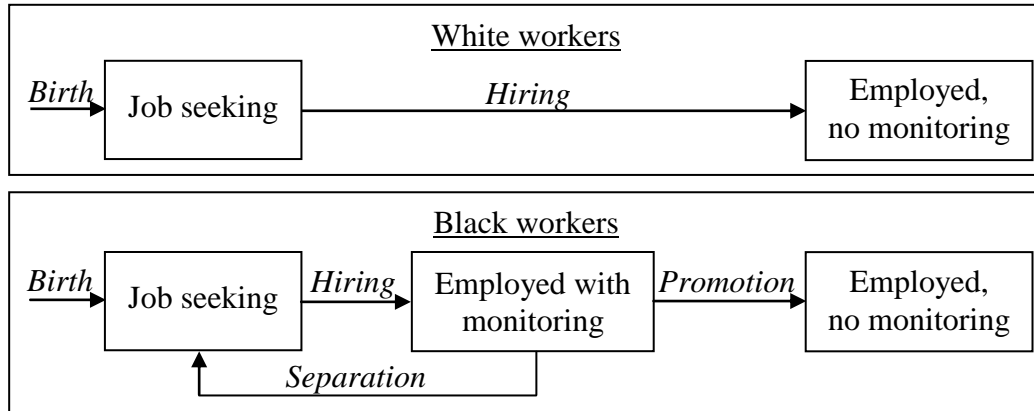


Figure 2.1: White workers' perpetual employment, black workers' churning cycle

Even if blacks are, on average, more skilled than whites, whites can be in the good steady-state and blacks in the bad steady-state because of a history of lower access to schooling and other human capital investments. Equalizing the human capital that blacks and whites bring to the labor market may be insufficient to equalize labor market outcomes. In contrast, in self-confirming expectations models, if we could just convince blacks to invest in themselves and employers that blacks have invested, we would immediately jump to the good equilibrium.

There is an abundance of evidence that black workers face lower wages and longer unemployment duration than white workers. Moreover, these disparities are less prevalent and, perhaps, in some cases nonexistent for the most skilled workers as measured by education or performance on the Armed Forces Qualifying Test. While there are a plethora of models intended to explain wage or unemployment differentials, none addresses both and their relation to skill.⁵ Since in our model newly hired black

⁵Many models (e.g. Aigner and Cain, 1977; Becker, 1971; Bjerck, 2008; Charles and Guryan, 2011; Coate and Loury, 1993; Fryer, 2007; Lang, 1986; Lang and Manove, 2011; Lundberg and Startz, 1983; Moro and Norman, 2004) assume market clearing and therefore cannot address unemployment patterns. Search models (e.g. Black, 1995; Bowlus and Eckstein, 2002; Lang and Manove, 2003; Lang, Manove and Dickens, 2005; Rosen, 1997) can explain unemployment differentials, but assume otherwise homogeneous workers and thus cannot address wage differentials at different skill levels. Peski and Szentes (2013) treat wages as exogenous. In general, discrimination models have not

workers are on average less productive than white ones, firms that expect to hire blacks anticipate less profit from a vacancy and therefore offer fewer jobs. Consequently, blacks have longer unemployment durations. Also, since they both spend more time searching for a job and are believed to be less productive on hiring, blacks earn less over their lifetimes. Additionally, the higher level of scrutiny increases the return to skill for blacks, consistent with evidence that blacks invest more in schooling compared with apparently equivalent whites.

We derive additional implications from informal extensions to the model. The higher level of scrutiny increases the return to skill for blacks, consistent with evidence that blacks invest more in schooling compared with apparently equivalent whites. In addition, if unemployment history is partially observable, black job seekers who have experienced enough turnover may be permanently relegated to low-skill, low-wage jobs. Although we do not wish to overstate the predictive power of the model, we note that until around 1940, blacks and whites had similar unemployment rates (Fairlie and Sundstrom, 1999), while blacks faced lower wages. This is consistent with a setting in which, due to low human capital investments, blacks were assumed to have low productivity at most jobs and therefore not monitored for quality. ‘Churning’ of the black labor market would not begin until human capital investments were sufficiently high.

We believe that the broad implications of our model can be derived through a variety of formalizations. The key elements common to these are:

- i. that a worker’s productivity at different firms is correlated,
- ii. that workers cannot or do not signal their ability and that they can, at least imperfectly, hide their employment histories,⁶

addressed employment duration. See the review in Lang and Lehmann (2012).

⁶In particular, they must sometimes be able to omit or mischaracterize prior bad matches.

- iii. that firms must therefore, to some degree, statistically infer worker ability,
- iv. that further information about match productivity is costly, imperfect, or both,
and
- v. that this information, if obtained, may affect hiring or retention, so that firm behavior affects the average unemployed worker's ability.

The details of our formal model are driven by our desire for a theoretically rigorous model of wage-setting in a dynamic framework with asymmetric information. Firms and workers bargain over wages and use a costly monitoring technology to assess the quality of the match, which is correlated with the worker's underlying type. We argue that separating worker-types is impossible without commitment to monitor, as the workers with the greatest incentive to be monitored for match quality while their type is unknown to the firm (those privately sure to be a good match) are also the ones for whom monitoring is most ex-post inefficient.

Therefore, use of the monitoring technology depends on the firm's prior: if the belief that a worker is well-matched is sufficiently high or sufficiently low, it will not be worth investing resources to determine match quality. However, if the cost of determining the match quality is not too high, there will be an intermediate range at which this investment is worthwhile. Firm beliefs about black, but not white, workers fall in this region. Consequently, they are subject to heightened scrutiny and are more likely to be found to be a poor match and fired. The increased scrutiny ensures that the pool of unemployed black workers has a higher proportion of workers who have been found to be a poor match at one or more prior jobs. And therefore employers' expectations that black workers are more likely to be poor matches is correct in equilibrium. This, nested in a search model, generates the empirical predictions discussed above.⁷

⁷Note that our model abstracts from moral hazard and that performance is observed objectively.

This churning equilibrium is hard to escape. This is disheartening since *policy succeeding at convergence of group characteristics may fail to equate labor market outcomes*. Only if the skill level of blacks is raised sufficiently above that of whites (technically the proportion of good workers is sufficiently high), does the bad equilibrium cease to exist and white and black workers receive similar treatment.

2.2 A simple example of churning

To provide some intuition, we first consider a simple discrete-time market in which we abstract from wage bargaining and vacancy creation decisions. These will play a central role in the full model.

A unit mass of worker is born every period. Suppose new workers have a probability $g = 2/3$ of being type α and producing $q_\alpha = 1$ unit per period and with the remaining probability are type β and produce $q_\beta = 0$. Each unemployed worker who has not been publicly revealed to be type β is matched to a firm at the beginning of the period. Wages are set to $w = 1/3$ exogenously,⁸ to be endogenized in the full model. Firms can either hire a worker indefinitely, or hire for a single period with monitoring costing $b = 4/3$, which reveals a β employee to the firm with a probability $1/2$, then firing those revealed and keeping the rest indefinitely following that. Matches do not dissolve naturally and the discount factor is $\delta = .95$.

To show churning can persist in environments where the market learns about worker ability rather quickly, we assume that the second revelation of a β worker is public. Such a worker is not hired again, and thus exits this labor market for one with lower wages and production that is less type-sensitive⁹ - unlike one revealed only

MacLeod (2003) develops an interesting model in which biased subjective assessments interact with moral hazard concerns.

⁸This is half the expected surplus of a new worker.

⁹In the context of our main model, we call this degenerating into ‘dead end jobs’ in the extension presented in Section 2.5.5.

once, who is re-matched next period.

Consider first a market with only first-time job seekers. A firm that hires but does not monitor a worker earns

$$g(q_\alpha - w)/r + (1 - g)(q_\beta - w)/r = \frac{80}{9} - \frac{20}{9} = \frac{20}{3}.$$

One that does monitor to fire revealed β s (recall they stay for one period) earns

$$g(q_\alpha - w)/r + (1 - g) \left[\frac{1}{2}(q_\beta - w)/r + \frac{1}{2}(q_\beta - w) \right] - b = \frac{80}{9} - \frac{21}{18} - \frac{4}{3} = \frac{115}{18} < \frac{20}{3}.$$

So, a firm will prefer not to monitor newly hired workers. Consequently, no β workers are revealed or fired, and all unemployed workers are first-time job seekers as assumed.

Now consider a market that has been churned by the monitoring technology. In each period, half the β workers who got their first job the previous period are fired and return to the job-seeking pool where they join a new batch of β workers of size $1 - g$ and α workers of size g ; thus the probability a newly hired worker is of type α is $g_c = g / (g + (1 - g) + .5(1 - g)) = 4/7$.

An employer in this market who does not monitor will get a payoff of

$$g_c(q_\alpha - w)/r + (1 - g_c)(q_\beta - w)/r = \frac{160}{21} - \frac{60}{21} = \frac{100}{21}$$

whereas one who does monitor expects a payoff of

$$g_c(q_\alpha - w)/r + (1 - g_c) \left[\frac{1}{2}(q_\beta - w)/r + \frac{1}{2}(q_\beta - w) \right] - b = \frac{160}{21} - \frac{21}{14} - \frac{4}{3} = \frac{201}{42} > \frac{100}{21}.$$

Thus workers in the second, or ‘*churned*’, market are monitored and can be fired.

This simplistic model demonstrates how two groups with the same underlying abilities can face very different treatment, and that this process can be self-enforcing. It captures churning-induced discrimination. Since only one group suffers separa-

tions, we interpret this as an employment duration differential. To address wage and unemployment duration differentials however, we will need the main model.

This simple example also helps demonstrate an important point: history matters. It is readily confirmed that the market switches from monitoring to not monitoring when the proportion of α s in the labor market surpasses $11/19$ and that firms will make a loss if this proportion is less than $29/101$. Consider a group for which historically the proportion α was less than $29/101$ and was therefore employed in some other type of job. Now let improvements in human capital lead new entrants in period t to have a proportion α equal to $.3$; also, let this proportion grow to $2/3$ in period $t + 1$ and remain at this level thereafter. The group never exits the churning equilibrium. Despite a legacy of only one generation in which the quality of the inflow favored churning, the group would remain stuck in the churning equilibrium until some time after the proportion α in the new generation exceeded $33/49$.

The example also shows that it is not essential that a worker's employment history be entirely opaque. Even though workers can only hide a single dismissal, the churning mechanism operates and induces a worse steady state.

2.3 The Model

We now present our model. As in the model in Section 2, employers statistically infer past employment based on race and may therefore discriminate in monitoring and retention; this can in turn churn the labor market and thus self-perpetuate. The main model's richer ontology will now enable us to address wage setting, unemployment duration and a host of other questions.

2.3.1 Setup

There are two worker groups, ‘blacks’ and ‘whites’. Race is observable by the worker and employers but does not have any direct impact on production.

At all times a steady flow of new workers is born into each population group.¹⁰ A proportion $g \in (0, 1)$ of new workers are type α , for whom every job is a good match.¹¹ The rest, referred to as type β , have probability $\beta \in (0, 1)$ of being a good match at any particular job. The probability of a worker being good at a job, conditional on her type, is independent across jobs. Worker type is private to the worker. Workers begin their lives unemployed. We define the average probability of being good at a particular job among new job seekers as

$$\theta_0 = g + (1 - g)\beta. \quad (2.1)$$

Employers cannot directly observe worker type or employment history,¹² but can instead draw statistical inferences from race.

2.3.2 Match Quality

Production, the payment of wages and the use of the monitoring technology occur in continuous time using a common discount rate r .

Workers can be either well-suited to a task (a ‘good’ match), producing q per unit time; or ill-suited (a ‘bad’ match), producing expected output $q - \lambda c$ per unit time. We can interpret the lower productivity of bad workers as errors or missed opportunities, each costing the firm c , that arrive at a constant rate λ . Under this

¹⁰We do not allow for death but could do so at the cost of a little added complexity.

¹¹Having type α workers perform well at every job does not appear to be essential to the argument but does appear to be essential to having comprehensible mathematics.

¹²At a more informal level, we believe that workers have some ability to hide their employment history and that they will not report information speaking to their own low ability. We show the model is robust to imperfect history revelation in Section 2.5.5.

interpretation, opportunities for error are also opportunities to learn the quality of the match as well-matched workers are observed to avoid errors.¹³

For monitoring to ever be useful, matches revealed to be bad must separate. To this end, we make the sufficient and simple assumption that such a match is unproductive:

$$(C1) \quad q - \lambda c \leq 0.$$

It is *much* stronger than necessary. In general, if the worker and firm know that the match is bad, it will be efficient for the worker to experience some unemployment in order to try a new match; this is a consequence of productivity conditional on worker type being match-specific. Assumption (C1) ensures that such separation in search of a better match is efficient regardless of the expected duration of unemployment.¹⁴

Neither the employer nor a type β worker can know the match quality without monitoring. The parties can agree to a costly regime of monitoring that may produce a fully informative, bilaterally observable signal about match quality. In keeping with the opportunities-for-errors interpretation, we assume the signal arrives at a constant hazard rate λ . The monitoring technology costs b per unit time, so that the expected cost of information is $\int_0^\infty b e^{-\lambda t} dt = b/\lambda$ and its expected discounted cost is $\int_0^\infty (e^{-rt} b) e^{-\lambda t} dt = b/(\lambda + r)$. The principal benefit of a signal whose arrival is exponentially distributed, rather than one that arrives deterministically, is that it makes the employment survival function more realistic. In addition, it allows for a certain stationarity in the model: so long as no signal has arrived, the underlying incentives do not change.

¹³Alternatively, we could assume that the flows are $q - d$ and q with $d \equiv \lambda c$ and that λ is the arrival rate of opportunities to measure the flows.

¹⁴Nothing of interest is ruled out here; if known bad matches don't end, then costly monitoring for separation is never worth paying for.

2.3.3 Job Search

When a worker is born or her match is terminated, she becomes unemployed. Unemployed workers are stochastically matched to firms, which occurs at a constant hazard μ . For the moment, we treat this rate as exogenous; it will be endogenized in Section 2.5.4 to address unemployment duration. When a match dissolves, transfers cease and the worker becomes unemployed. A firm does not recoup a vacancy and therefore receives a payoff of 0 on termination.¹⁵

In the unemployed state, workers merely search for new jobs; we normalize the flow utility from this state to 0. The value from unemployment is thus simply the appropriately discounted expected utility from job-finding and is invariant to history. The discount on job-finding is $\int_0^\infty e^{-rt} \mu e^{-\mu t} dt = \mu / (\mu + r)$; the value of a new job will depend on the equilibrium. We denote the value of the job-finding state as U_α for type α workers and U_β for type β workers.

2.3.4 Bargaining

Informal Description

In the interest of modeling wage determination, this section ends up being more technical than may be of interest to readers who are primarily interested in discrimination. We therefore begin with a brief intuitive discussion which we hope will be sufficient to permit such readers to skip the technical discussion.

We cannot use Nash bargaining because there is no accepted model of Nash bargaining with asymmetric information. Instead, we use a bargaining model in which workers and firms make alternating offers. We assume that the parties may unilaterally reopen bargaining at any time but with a delay. Offers take the form of a wage and monitoring regime. If the regime involves no monitoring, no new information

¹⁵This occurs naturally due to free entry when vacancy creation is endogenized; see Section 2.5.4.

arises. If the regime involves monitoring, bad matches will separate, and those shown to be good will renegotiate so as to not continue monitoring.

A critical question is whether the bargaining can reveal workers' private information about their type. Intuitively, firms might propose a monitoring offer that would attract one type and then a no monitoring offer that would attract the other. Alternatively, a worker could try to signal her type by bargaining tactics, effectively engaging in 'money burning'. The problem is that if separation occurs, α workers and firms should immediately renegotiate to a no-monitoring regime with a high wage reflecting the fact that the match is known to be good. But this is also the best possible outcome for a β worker. So, knowing that renegotiation will occur immediately, β workers will pretend to be α s, so such type-separating solutions fail to exist. Thus separation in this setting would require commitment, even in the face of Pareto-improving alternatives.

Since they do not wish to reveal themselves, in the solution β workers negotiate as if they were type α . The firm therefore evaluates offers as though the worker is an average of the two types. As in the Rubinstein (1982) model, there is no utility flow while bargaining.

When there is no monitoring and the firm believes the match is good with probability θ , as the bargaining delay disappears the worker receives w/r and the firm receives $(q - (1 - \theta) \lambda c - w) / r$. This is split by an average (over proposers) wage of

$$.5(q - (1 - \theta) \lambda c),$$

as in the Nash bargaining solution. Lemma 2 shows each party values a revealed good match at $.5q/r$.

When monitoring takes place, bargaining splits the cost equally on average. In addition to that however, a new term appears reflecting that firms and workers eval-

uate the probability of the match being good differently. Again abstracting from bargaining delays, the wage is

$$\frac{1}{2} (q - b - \lambda c(1 - \theta)) - \frac{(1 - \theta)}{2} \lambda \frac{q}{2r}.$$

This would also obtain as the equal-weights Nash bargaining outcome of an α worker bargaining with a firm with belief θ that the match is good, with 0 outside options. Since all workers bargain as α s who know the match is good but firms have belief θ , workers are more impatient to get to revelation and therefore bargain as if delays were more costly. This means that monitored workers bear not only their share of the monitoring cost but an additional “Pooling Penalty.” As in Nash bargaining, the monitoring policy is efficient from the standpoint of firms and α workers.

In the next subsections we impose conditions to ensure that monitoring is indeed optimal in the churned (black) labor market only and that type separation is infeasible. We furthermore make bargaining stationarity assumptions that empower off-path renegotiation as a way to exclude equilibria supported by either unreasonable off-path beliefs or repeated-games-style ‘punishments’. We then derive both steady-state solutions of the full alternating-offers bargaining model when the time between offers is small. Readers who are less interested in the technical details may wish to skim the material until Section 2.4.

The Formal Bargaining Model

Wage and monitoring contracts are determined by alternating-offers bargaining with a delay of Δ .¹⁶ An offer is a pair $(w, m) \in \mathbb{R} \times \{0, 1\}$ comprised of a wage w per unit

¹⁶Although the standard Mortensen and Pissarides (1994) model uses Nash bargaining, it requires symmetric information and therefore is unusable in our setting. Evidence from Hall and Milgrom (2008) suggests that their own Rubinstein variant with added pecuniary costs of delay is able to produce far more realistic unemployment predictions than Nash bargaining. Our model shares the feature that enables this prediction (workers’ outside option not dampening firm payoff fluctuations).

time paid continuously and a policy m of using or not the monitoring technology.¹⁷

When a match is first formed, a first proposer is chosen with equal probability on the firm and worker. Production, monitoring and wages cease during bargaining. We are interested in solutions when Δ is low.

Most importantly, either partner may *unilaterally* choose to re-open negotiations at any time by causing a single delay of length Δ during which production and wages are suspended. Once this delay expires, the party instigating renegotiation is placed in the role of proposer.¹⁸ The choice to reopen negotiation is logically simultaneous at each time, and if both partners wish to reopen negotiations at the same instant they each assume the role of proposer with probability $1/2$.

Thus, there is no commitment to any agreement. This is important. If the wage is independent of worker type, it will generally fall between the wages that would be negotiated by a known β and by a known α . Therefore, starting from a common wage, if a worker is revealed to be an α , she will renegotiate to raise her wage, while the firm will renegotiate a lower wage if the worker is revealed to be a β . This creates an environment hostile to separating equilibria.

With the impermanence of deals, however, we now open ourselves to repeated-games type equilibria where the acceptance of bad offers, and intransigence in insisting on them, is enforced by off-path punishment. To recover the uniqueness of Rubinstein bargaining from this, we make an assumption:

(S0) Stationarity: Consider histories where firm beliefs put probability 1 on a certain

¹⁷We assume that offers entail constant wages and monitoring, a limitation. Allowing time-varying wage profiles to be offered does not affect our findings but results in the loss of some elegance. We can show that our results hold for the average wage over a small interval that is nevertheless large relative to the bargaining delay but cannot rule out wages that, for example, alternate between a high and low wage with each wage maintained for a period equal to the bargaining delay. If we further assume that wages and monitoring can be contingent on the signal arriving, we require additional assumptions on the delay, Δ , to preserve our results; at the cost of considerable complexity, the equilibrium derived here is essentially unique as $\Delta \downarrow 0$.

¹⁸This delay on renegotiation ensures that disagreeable offers are rejected rather than accepted with the intent to renegotiate instantly.

worker type or match quality. There are no deviating offers at such histories that if not renegotiated (in the case of uncertain match quality, until revelation) improve¹⁹ the payoff of the proposer while giving the receiver more than the once-discounted expected value at their previous offer (or, if this is the first offer, the receiver's once-discounted value of offering first).

Stationarity allows for precisely the kind of argument present in standard Rubinstein bargaining. A party who makes offers it values at x should be willing to accept offers it values at $e^{-r\Delta}x$. This further allows us to dispense with repeated-games type inefficient behavior, such as strategies that waste most of the surplus under the threat of wasting even more of the surplus.

At this point we want to assume that the bargaining delay is not too large for the parties to renegotiate to shut off monitoring after match quality revelation.

$$(C2) \quad e^{-r\Delta} \cdot q > q - b.$$

In fact, we want to think of the bargaining delay as being vanishingly small and do our analysis in Sections 5 and 6 treating it as such.

Bearing this in mind, we additionally postulate that

$$(C3) \quad e^{-r\Delta} > \frac{\mu}{\mu + r}$$

to ensure that for a worker, rejecting an offer and making a counter offer is, in expectation, faster than separating in order to find a new match where the worker might be the first proposer (for simplicity, we formalize this as though he will be the first proposer). Counter-offering is quicker than finding a new employer to make an offer to. Again, this condition must always be satisfied for sufficiently small Δ .

¹⁹A delay caused by rejecting an equilibrium offer or reopening negotiations is of course factored in to deciding whether a deviating proposal is payoff-improving to the proposer.

Bargaining solution with symmetric information

First, let us find the subgame perfect Nash equilibrium (SPNE) solution where bargaining occurs under symmetric information. Using S0, the parties will make stationary offers and split the output according to the Rubinstein shares, $1/(1+e^{-r\Delta})$ for the first proposer and $e^{-r\Delta}/(1+e^{-r\Delta})$ for the responder. These shares are delivered via a constant wage, avoiding renegotiation in the absence of information from monitoring.

Known Bad Match. By (C1), the match would forever produce a negative average flow should it persist, so it instead separates.

Known Good Match. The total match surplus is the discounted value of producing q for all time, $\int_0^\infty qe^{-rt} dt = q/r$. The first proposer therefore earns

$$\frac{1}{1+e^{-r\Delta}} \cdot \frac{q}{r}.$$

We can now show that matches in which revelation of good quality occurred via a policy of monitoring will instantly renegotiate:

Lemma 2. *If monitoring reveals a match to be good, both parties request renegotiation; they each expect a payoff of $e^{-r\Delta}q/(2r)$ upon such revelation.*

Proof. See B.1 □

As the bargaining under symmetric information will be efficient, we can decouple the monitoring decision from wage setting and proceed to examine the latter.

Known β , no monitoring. Worker type is commonly known to be β , and the match is of unknown quality. The average over match qualities cost of errors per unit time is $(1-\beta)\lambda c$. Total match surplus is therefore

$$S_{N\beta} = \frac{q - (1-\beta)\lambda c}{r}. \tag{2.2}$$

The first proposer therefore receives

$$\frac{1}{1 + e^{-r\Delta}} \cdot \frac{q - (1 - \beta)\lambda c}{r}. \quad (2.3)$$

Known β , monitoring. If the match is revealed by the signal to be bad, separation occurs; the firm receives 0 and the worker U_β . By Lemma 2, when the match is revealed to be good, each player expects a payoff of $qe^{-r\Delta}/(2r)$. Expected discount on revelation is $\int_0^\infty e^{-rt} \cdot \lambda e^{-\lambda t} dt = \lambda/(\lambda + r)$. For expected discounted pre-revelation total wages W , the worker's total expected payoff is

$$W + \frac{\lambda}{\lambda + r} \left(\beta \frac{qe^{-r\Delta}}{2r} + (1 - \beta)U_\beta \right). \quad (2.4)$$

The firm's payoff, remembering a separation has a value of 0, is

$$\frac{q - (1 - \beta)\lambda c - b}{\lambda + r} - W + \frac{\lambda}{r + \lambda} \left(\beta \frac{qe^{-r\Delta}}{2r} + (1 - \beta) \cdot 0 \right). \quad (2.5)$$

The total surplus from the match is therefore

$$S_{M\beta} = \frac{q - b}{\lambda + r} - \frac{(1 - \beta)\lambda c}{\lambda + r} + \frac{\lambda\beta}{\lambda + r} \frac{qe^{-r\Delta}}{r} + \frac{\lambda(1 - \beta)}{\lambda + r} U_\beta. \quad (2.6)$$

We can solve for the instantaneous wage, which averages (over proposers) to

$$w_{M\beta} = .5(q - (1 - \beta)\lambda c - b) - .5(1 - \beta)\lambda U_\beta. \quad (2.7)$$

Continuation play from an off-path history in any of the above cases is for both players to request immediate renegotiation and propose the equilibrium shares, unless both players are receiving greater than the receiver's share by the current offer (in which

²⁰The solution has the somewhat disturbing property that the worker's value following separation *lowers* the wage. The worker is impatient for the opportunity to 'try again' if she turns out to be bad at the job. Similarly to most alternating offers models, the outside option does not directly affect the outcome here. However, as β workers will not be strategically revealed, we do not observe wages with this property.

case the status quo offer continues until revelation).

2.3.5 Steady State

A steady state of a labor market is a mass of α job seekers, a mass of β job seekers and a mass of monitored β workers along with equilibrium firm and worker strategies that make these populations constant over time. There are two kinds of stable steady states: those in which all employees are monitored until match quality is revealed, and those in which no monitoring occurs.²¹

Consider the case where no employees are monitored: the white labor market. Matches never deteriorate and therefore the only source of job seekers is newly born workers. In this scenario, a firm just matched with an employee infers his probability of being of type α is the population prevalence g ; the chance of a white job-seeker being good at a job to which he is matched is therefore

$$\theta_W = \theta_0 = g + (1 - g)\beta.$$

Now suppose that all newly hired black employees are monitored and all bad matches are terminated. Newly matched black workers will be worse than average.

Lemma 3. *The probability a newly hired black worker is in a good match is*

$$\theta_B = \frac{\beta}{\beta g + (1 - g)} < \theta_W. \quad (2.8)$$

Proof. See B.2 □

Therefore, although monitoring may be individually prudent for each matched pair, it creates a negative externality by feeding a stream of workers who are worse than the population average (i.e. containing more β types) back into the job-seeker

²¹A steady state in which only some workers are monitored until revelation is not stable as it implies indifference and a mixed strategy for the firm. A perturbation in θ will lead to either complete or no monitoring, causing movement away from the steady state.

pool. Surprisingly, the steady state θ_B of this process does not depend on the rate of information λ , the worker matching rate μ or the rate at which new workers enter the market.²²

2.3.6 Solution Concept

We are interested in solutions that fulfill the following criteria in addition to S0:

S1 Steady State: The labor market is in steady state.

S2 PBE: Firm and worker strategies form a perfect Bayesian equilibrium.

S3 Stationarity/No Dictatorial Beliefs: At no history with firm beliefs $\theta_h \in (\beta, 1)$ on the match being good can a deviating offer be made that, should it stay in place until revelation and beliefs be fixed at θ_h until revelation:

- a) strictly improves the payoff of a proposing firm or α worker, and either
- b1) if the first offer in a match, improves on the lesser of the receiving firm or α worker's once-discounted first-proposer payoff or equilibrium payoff at the current node, or
- b2) if a subsequent offer, gives a greater payoff to a receiving firm or α worker than their once-discounted expected payoff at their previous offer.

Restriction S3 requires some explanation. Its primary purpose is to provide uniqueness. It allows α workers to make off-equilibrium offers that are beneficial to them without having to worry about the offers' effect on beliefs. It furthermore allows firms to make offers that the best workers should accept without those workers worrying about a deleterious effect acceptance has on beliefs.

²²This is an artifact of the assumption that workers are infinitely lived.

This restriction ensures that the bargaining protocol will produce real bargaining and Rubinstein-like solutions rather than dogmatic offers backed by the threat of belief change or punishment off-path. Lacking S3, low wages could be maintained by the firm believing any deviation is due to the worker being type β . By providing for deviations from such situations without belief ramifications, we eliminate these distasteful equilibria.²³

2.3.7 Parametric Assumptions

Now we impose certain restrictions on the joint values of parameters sufficient to ensure the existence of both solutions.

For an equilibrium with no monitoring to exist for white workers, we want to assume that monitoring costs are not too low. Initially, we want to abstract from bargaining frictions; in the limit as Δ disappears, the first parameter restriction can be stated as saying that monitoring costs exceed the sum of the benefits to the firm and α worker.

$$\underbrace{\frac{b}{\lambda}}_{\text{Mon. cost}} > \underbrace{\underbrace{(1 - \theta_W) \frac{\lambda c}{2r}}_{\text{Revelation Gain to } \alpha}}_{\text{Wage Increase}} + \underbrace{\underbrace{(1 - \theta_W) \left(\frac{\lambda c}{r} - \frac{q}{2r} - (1 - \theta_W) \frac{\lambda c}{2r} \right)}_{\text{Revelation Gain to Firm}}}_{\text{Bad Match End Benefit}} - \underbrace{\theta_W (1 - \theta_W) \frac{\lambda c}{2r}}_{\text{Wage Increase}} \quad (2.9)$$

Restating this limiting condition to constrain β and g rather than the monitoring costs and recalling that $\lambda c - q/2$ is guaranteed to be positive due to (C1), we get

$$\theta_W = g + (1 - g)\beta \geq 1 - \frac{r}{\lambda} \cdot \frac{b}{\lambda c - \frac{q}{2}} \Rightarrow \underbrace{(1 - \beta)(1 - g)}_{\text{Probability Match is Bad}} < \frac{\frac{b}{\lambda}}{\frac{\lambda c}{r} - \frac{q}{2r}}. \quad (2.10)$$

However, as we wish to derive equilibria when bargaining delays are nonzero, things are a bit more complicated. Expression (2.9), when accounting for bargaining

²³Without S3, workers could be forced to accept far below half the average surplus even if match quality was unobservable in principle - despite types having identical incentives.

frictions, transforms into the rather less interpretable

$$(C4) \quad \frac{b}{\lambda} > (1 - \theta_W) \frac{\lambda c}{r} - \frac{q(2 - e^{-r\Delta}(1 + \theta_W))}{2r}.$$

Unsurprisingly, this condition is still of the form “*monitoring costs must not be too low.*” On the left hand side is the instantaneous cost of information, b/λ . On the right hand side, the losses to the employer that result from bad matches, which would end if monitoring is successful, minus a measure of production lost to both separations and delays due to renegotiation when match quality revelation occurs.²⁴ Strictness of the inequality ensures persistence and stability. Condition (C4) guarantees that a type α worker (who is more eager for monitoring) does not want to propose monitoring if, having done so, the employer will not update beliefs; firm beliefs θ_W are good enough.

Our second condition, antisymmetrically to (C4), posits that “monitoring costs must not be too high” and ensures the monitoring equilibrium also exists:

$$(C5) \quad \frac{b}{\lambda} < (1 - \theta_B) \frac{\lambda c}{r} - \frac{q(2 - e^{-r\Delta}(1 + \theta_B))}{2r}.$$

This condition establishes that, in the black (‘churned’) labor market with a pooling monitoring equilibrium, it is not profitable for α types to deviate to a no-monitoring offer if this would not change the employer’s belief. In other words θ_B , the belief about the average ability in the black unemployed pool, must be sufficiently low that α workers prefer to be monitored. Strictness of the inequality ensures that switching to an unchurned market is not simply a matter of switching equilibria (as the non-monitoring one will not exist here).

Combining conditions (C4) and (C5), in the limit as $\Delta \downarrow 0$, we get

$$(1 - \beta)(1 - g) < \frac{\frac{b}{\lambda}}{\frac{\lambda c}{r} - \frac{q}{2r}} < \frac{(1 - \beta)(1 - g)}{\beta g + (1 - g)}. \quad (2.11)$$

²⁴As an α worker and firm hold different beliefs about the probability of a bad match, the lost production is estimated using their average beliefs.

That is, the ratio of the average costs of monitoring to a measure of its benefits lies strictly between the rate of bad matches in the two markets.

The two remaining conditions rule out the complexities associated with partial pooling. We assume that the efficient outcome for a match of unknown quality in which the worker is revealed to be type β is to monitor. This is the case even when it takes a single delay to renegotiate:

$$(C6) \quad \max\{S_{N\beta}, 0\} < e^{-r\Delta} \cdot S_{M\beta}$$

where $S_{N\beta}$ is the surplus from a known β match without monitoring and $S_{M\beta}$ is the surplus from a known β match with monitoring, both derived in Section 2.3.4. This further allows us to say that there is positive surplus from these matches; it is not detrimental to social welfare that β workers are employed at all. Assuming (C6) allows us to say that since α s benefit more than β s from monitoring, if β s want to reveal themselves in order to be monitored, α s will want to pretend to be β s. Second, in the absence of (C6), it is difficult to rule out mixed strategy equilibria where some β workers reveal themselves to avoid monitoring costs.

Finally, we require that the black equilibrium pooling wage is no lower than the wage β workers could get by revealing their type:

$$(C7) \quad \frac{e^{-r\Delta}(q - b - \lambda(1 - \theta_B)) - (1 - e^{-r\Delta}\theta_B)qe^{-r\Delta}/(2r)}{(1 + e^{-r\Delta})} > w_{M\beta}.$$

Condition (C7) provides that β s must prefer pooling with monitoring to revealing their type and being monitored. As discussed, β s gain from pooling with α s because of believed higher productivity but lose because it is costly for them to bargain as if they were α s. This can be rewritten as a restriction on β . If β is near 0, θ_B will also be near 0, and the benefits of pooling will be diminished.

Note that if (C6) and/or (C7) is violated, it will still not be an equilibrium for

all β s to reveal themselves, as any worker not doing so would earn the revealed good payoff, $.5e^{-r\Delta}q/r$. Thus the benefit of these conditions is that they allow us to rule out equilibria in which some, but not all, β s reveal their type.

It is not self-evident that (C4)-(C7) can hold simultaneously. In Section 2.6 we provide an example with plausible parameters in which they do.

2.4 Solution

We present the main results of the paper: existence and essential uniqueness of equilibria in the two markets that perpetuate their associated steady states.

2.4.1 The Non-Monitored Market

Proposition 1. *Assuming (C1)-(C7), the white (non-churned) labor market has a solution where the monitoring technology is not used on-path. Employed workers, regardless of type, receive their Rubinstein share of the surplus. In the limit as $\Delta \downarrow 0$ their wage is*

$$w_{N\theta_w} = .5[q - (1 - \theta_W)\lambda c]. \quad (2.12)$$

Proposition 2. *The above equilibrium is unique.*

All proofs are in the appendix.

The main intuition here flows from (C4), the lack of commitment and S3. On the one hand, (C4) tells us that α workers would only really want to deviate to a monitoring offer if they could affect beliefs by doing so - beliefs are high enough that the wage is already good in equilibrium. On the other hand, if it were possible for α workers to reveal themselves by making a monitoring offer, then as soon as they made it, beliefs would change, and S3 would allow them to make a new, improved no-monitoring offer also preferable to β workers.

Interestingly, since the firm cannot learn the worker's type in this non-churned equilibrium, type has no effect on wages. The firm's prior, θ_W , is high enough that

even good workers do not wish to pay to reveal match quality.

2.4.2 The Monitored Market

Whether or not the equilibrium involves monitoring, α s are always willing to pay more than β s to be monitored since they know the match is good. In the monitored (black) labor market, average job-seeker quality θ_B is low and the firm's expected error costs, $\lambda c(1 - \theta_B)$, are high. These expected costs are shared with the workers, so both firms and α workers wish to reveal match quality. But then β s must follow suit lest they be revealed. All workers must therefore bargain as if they were type α and make offers with monitoring.

Proposition 3. *Assuming (C1)-(C7), the black (churned) labor market has a monitoring employment equilibrium with a wage limiting to*

$$w_{M\theta_B} = \frac{1}{2} [q - b - \lambda c(1 - \theta_B)] - \frac{(1 - \theta_B)}{2} \lambda \frac{q}{2r}. \quad (2.13)$$

Proposition 4. *This equilibrium is unique.*

Note that since in order not to reveal his type, a β -worker has to bargain as type α , he acts as though the probability of promotion is 1 even though the firm treats him as being of average type θ_B . If the worker were truly a “type θ_B ,” with probability of matching well of θ_B , known to be one and bargained as one, the Rubinstein bargaining solution would substitute the value of unemployment as a θ_B for q/r in the final term. Instead, the firm here extracts additional surplus over the baseline of a “ θ_B ” type; as $\Delta \downarrow 0$ this limits to $.5(1 - \theta_B)\lambda(.5q - U_{\theta_B})$, the ‘pooling penalty’.²⁵ Type α s are hurt by pooling with β s not only because of the pooling penalty but also because the firm underestimates their output by $\lambda c(1 - \theta_B)$.

As the equilibrium strategies induce full monitoring, employees who are revealed

²⁵The Pooling Penalty is always positive as the unemployment value of the worker is never higher than the payoff to matching well.

to be in bad matches separate from the firm. This sends only β workers back into the job-seeking pool, churning the market quality to θ_B .

2.5 Implications for Labor Markets

The previous sections establish conditions under which there are two distinct steady-states of the labor market. In this section, we compare labor market outcomes for workers in these steady states. Our comparative statics are performed in the limit as the bargaining delay goes to 0.

2.5.1 Employment Duration

Absent monitoring, there is no new information to dissolve the match. Therefore, taken literally, the model implies no turnover in the white equilibrium. In contrast, with monitoring, some workers prove ill-suited for the job and return to the unemployment pool. We interpret this as predicting that black workers will have lower average employment duration. Recall that workers who return to the unemployment pool are all type β . Therefore, turnover is even higher than if only new entrants were monitored.

The model, again taken literally, implies that the separation hazard for blacks is

$$h_t = \frac{(1 - \beta)(1 - g)\lambda e^{-\lambda t}}{1 - (1 - \beta)(1 - g)e^{-\lambda t}} \quad (2.14)$$

which is decreasing in t . We expect the prediction that the exit hazard into unemployment should decline more rapidly for blacks than for whites is robust to consideration of important real world elements not addressed by the model. Unfortunately, all the estimates of this hazard by race that we have been able to find assume a constant hazard. The closest result we know of is Bowlus and Seitz (2000) who find that this hazard is much higher for young blacks than for young whites but that this differ-

ence disappears among older workers, a finding consistent with our model but that nevertheless does not directly measure the relations between hazards and seniority.

As our model abstracts from firm-to-firm hiring, we have no prediction with regards to it. Although it may seem that firms would be out to poach black workers with high seniority (that are likely to have passed monitoring), adverse selection effects (with the worst workers more willing to leave) could unravel such effects, depending on the ability of outside employers to commit. Still, our predictions are in terms of employment, not job, duration.

2.5.2 Persistence of Discrimination

A key result of the churning mechanism in this paper is that deleterious steady states are persistent. In this section we show just how hard it is to transition to a good steady state. We regard this as illustrating the difficulty of addressing labor market discrimination in the context of policy, particularly policy aimed at improving the skills of black workers. The existence of a range of g values for which both equilibria exist allows us to talk about *persistence* of the deleterious equilibrium.

Heretofore we have assumed that average skill levels for the two population groups are identical. Suppose instead that skill levels are $g_B \neq g_W$ and the initial equilibrium has monitoring of blacks but not whites. Monitoring will persist as the equilibrium in the black labor market until g_B rises above some critical level while the no monitoring equilibrium will persist in the white market provided that g_W remains above a lower critical level. In principle, we can have the black workers in the bad equilibrium and the white workers in the good equilibrium provided that (C4) is satisfied and

$$g_B \leq \frac{g_W}{\beta + (1 - \beta) g_W}. \quad (2.15)$$

To set ideas, suppose that g_W and β both equal .5, then we could observe the black

workers in the bad equilibrium if g_B is as large as $2/3$. In short, not only may discriminatory markets persist when skill levels for whites and blacks are identical, but they may persist *even when black skill levels are significantly higher*. Policy aimed at accomplishing convergence of labor market outcomes via changes in population skill may fail to clear the hurdle of inertia.

2.5.3 Wages

Wages are lower for black workers at the point of hiring. Not only do they pay a share of the monitoring cost, they also pay what we dubbed the Pooling Penalty. In addition, each type expects lower lifetime earnings than its white counterpart. To see this, consider the following:

- (i) Rearranging (C4), we have that the payoff to α s is higher in the unchurned market for a no-monitoring strategy:

$$\frac{q - \lambda c(1 - \theta_W)}{2r} > \frac{q - \lambda c(1 - \theta_W) - b + \lambda(1 + \theta_W)\frac{q}{2r}}{2(r + \lambda)}. \quad (2.16)$$

But as the right-hand side of that inequality is increasing in θ , we further have

$$\frac{q - \lambda c(1 - \theta_W) - b + \lambda(1 + \theta_W)\frac{q}{2r}}{2(r + \lambda)} > \frac{q - \lambda c(1 - \theta_B) - b + \lambda(1 + \theta_B)\frac{q}{2r}}{2(r + \lambda)} \quad (2.17)$$

and therefore

$$\frac{q - \lambda c(1 - \theta_W)}{2r} > \frac{q - \lambda c(1 - \theta_B) - b + \lambda(1 + \theta_B)\frac{q}{2r}}{2(r + \lambda)}, \quad (2.18)$$

which implies that white α s have a higher ex-ante payoff compared to their black counterparts. As all worker payoff derives from wages, this means that *lifetime wages are lower for black α workers*.

- (ii) On the other hand, white α and β workers expect the same lifetime wages. Since

β workers value monitoring strictly less than α workers and black α workers are worse off than white ones, black β workers must expect lower lifetime wages than their white counterparts.

Significantly, the model predicts that the realized strategies produce payoffs that maximize the joint surplus of α s and θ -belief firms. This implies that as a function of market θ , payoff to newly matched firms and type α workers is continuous, being an upper envelope of linear functions. However, the strategy shift produces a jump discontinuity in the payoff to β types forced to follow suit. Figure 2 illustrates this jump.²⁶




Figure 2.2: Equilibrium α and β payoffs as a function of average hire quality θ .

²⁶Figure 2, and this discussion, only concern expected lifetime payoff starting at a new job. As time goes to infinity, any single black worker will eventually be revealed good at a match and will therefore receive a better wage than white workers. We don't dwell on this issue as it is an artefact of the irrelevance of the outside option (much lower for black workers) in our particular bargaining model and the perfectly revealing nature of the monitoring technology. An alternate ultimate fate for workers is discussed in the model variant of section 2.5.5, where unlucky black workers can get stuck in low-wage jobs.

In a sense, because of the sharp discontinuity in the earnings of β s, the model predicts that the return to skill is higher for blacks than for whites, consistent with the empirical findings in Neal and Johnson (1996) and Lang and Manove (2011). We are reluctant to push this point strongly because the evidence concerns either observable ability in the form of education or potentially observable ability in the form of performance on the Armed Forces Qualifying Test. In section 2.5.5 we consider the case of observable investments.

The model has no role for human capital acquisition although below we briefly discuss the possibility that workers invest in human capital before they enter the labor market. Since there is no post-employment investment in the model, conditional on seniority, there is no return to experience for either blacks or whites. Since blacks spend more time unemployed, we might expect that, once we allow for such investments, the return to potential experience would be higher for whites, at least conditional on seniority. On the other hand, since the probability that blacks, but not whites, are well-matched to their jobs increases with job duration (through the selection effect), if we do not condition on seniority, this will tend to give blacks a higher return to potential experience.

But even this ignores the potential complementarity between match quality and human capital investment. On the one hand, firms are less likely to invest in workers whom they believe may be badly matched. On the other hand, they may be more (or less) likely to invest in black workers who have been revealed to be a good match than in white workers whose match quality is unknown and will not be revealed.

As a consequence of these considerations, we do not view the predictions regarding the differing effects of seniority and experience on the wages of blacks and whites to be robust. This is less disturbing than it could have been since we interpret the

empirical literature on this issue as fairly mixed.²⁷

2.5.4 Unemployment Duration

We have so far treated the workers' matching rate, μ , as exogenous. Making the standard assumption of free entry, we now allow firms to post and maintain vacancies at a cost k per unit time. When a firm creates a vacancy, it can direct its search. This can take several forms, most notably locating production operations in an area with specific population characteristics or advertising the vacancy in different areas and through different media. In general, a firm can target markets indexed by i where a proportion ρ_i of unemployed workers are white. The open vacancy cost k is invariant to this target choice. We assume that in each market i the bargaining equilibria and population group steady states break down along the discriminatory lines described so far.

Define ϕ as market tightness and let the worker job-finding rate function follow the commonly assumed form

$$\mu(\phi) = m\phi^\gamma \tag{2.19}$$

for constants $m > 0$ and $0 < \gamma < 1$. Note that if firms expect a match to be worth V , the free-entry level of ϕ in such a market sets

$$\frac{\mu(\phi)}{\phi}V - k = 0. \tag{2.20}$$

So

$$\phi = \left(\frac{Vm}{k}\right)^{\frac{1}{1-\gamma}}. \tag{2.21}$$

²⁷Monk (2000) finds that the experience effect on wages for blacks exceeds that for whites until roughly fifteen years of experience while the seniority effect is larger for whites through thirteen years of seniority. Bronars and Famulari (1997) also find that the black-white wage differential tends to fall with experience. On the other hand, D'Amico and Maxwell (1994) find that the gap between blacks and whites widens with experience, a result that Altonji and Blank (1999) view as confirming earlier work.

Therefore, ϕ is an increasing function of V .

Assuming that (C6) and (C7) hold for the entire breadth of derived matching rates, we can now derive the free-entry equilibrium level of μ_{ρ_i} for each market i . The payoff to a firm for matching is the same as for an α worker, that is, when hiring from pool i , the firm expects a successful match to pay

$$V_i = \rho_i \frac{q - \lambda c(1 - \theta_W)}{2r} + (1 - \rho_i) \frac{q - \lambda c(1 - \theta_B) - b + \lambda(1 + \theta_B) \frac{q}{2r}}{2(r + \lambda)}. \quad (2.22)$$

Since the payoff to white α workers is higher than for blacks, the expression above is increasing in ρ_i . Therefore, markets with more black workers will have a lower expected payoff for a filled vacancy. Therefore, the free-entry $\phi(\rho_i)$ and $\mu(\phi(\rho_i))$ are increasing in ρ_i . As average unemployment duration is $\frac{1}{\mu}$, this implies that *markets with higher black concentration will experience higher average unemployment duration*. In the extreme case where markets are fully segregated, that is $\rho_i \in \{0, 1\}$, we can derive the ratio of the matching rates in the two markets:

$$\frac{\mu(\phi(0))}{\mu(\phi(1))} = \left(\frac{q - \lambda c(1 - \theta_B) - b + \lambda(1 + \theta_B) \frac{q}{2r}}{q - \lambda c(1 - \theta_W)} \frac{r}{r + \lambda} \right)^{\frac{\gamma}{1-\gamma}} < 1. \quad (2.23)$$

2.5.5 Extensions

Eventual revelation in all matches

We have assumed unrealistically that the match quality of workers who are not monitored is never revealed. More plausibly, heightened scrutiny speeds the rate at which match quality is revealed. In a model in which workers live forever, this change considered in isolation would eliminate our result because the composition of the jobless pool is independent of the rate at which bad matches are revealed. However, if workers do not live forever, then reducing the rate at which match quality is revealed does affect the quality of the unemployment pool, and our basic results go through.

Skill level and discrimination

Further, we can allow for observable heterogeneity among workers. If there are groups of workers for whom g is high, only the no-monitoring equilibrium will exist for these groups, regardless of race. This is also true at very low g and very low β (although we have assumed away this case to simplify the proofs). The first result is consistent with similar outcomes for blacks and whites with high levels of skill as measured by education or the Armed Forces Qualifying Test (Neal and Johnson, 1996; Lang and Manove, 2011). The latter is consistent with some evidence that the bottom of the labor market is similarly bad for blacks and whites. On the other hand, Lang and Manove find that the market learns the productivity of white but not black high school dropouts. This is consistent with an equilibrium in which white dropouts are, on average, more skilled than black dropouts and therefore in which white but not black dropouts are monitored. Nevertheless, without additional, largely *ad hoc* assumptions, this story cannot account for the very high unemployment rate among black dropouts.

Investment in unobservable skills

We have heretofore postulated that the proportion of α types is exogenous. Assume instead that some fraction of workers are innately of type α . Others can transform themselves from β s into α s at some cost ω with cdf $F(\omega)$. Provided that the fraction of natural α s satisfies (C4) and (C5), both equilibria will continue to exist. However, since in the no-monitoring equilibrium α s and β s receive the same wage, there is no incentive to invest in becoming an α . In contrast, in the monitoring equilibrium, lifetime earnings are strictly higher for α s than for β s. Thus, some individuals will have an incentive to make the investment.²⁸ This prediction contrasts with Coate

²⁸It might appear that the incentive to undertake such investments would unravel the monitored equilibrium. However, if this were the case, no worker would have an incentive to invest. This raises

and Loury (1993), where black workers are less willing to invest in skills.

Education

Suppose now that there exists a signal,²⁹ which we identify with education, that α workers can purchase at some personal cost $\kappa \sim F(\kappa)$. Assume doing so assures that any employer will be immediately aware that the worker is indeed type α . A worker of either population will then anticipate a lifetime utility of $V_{Educ}(\kappa) = \mu q / (2r(\mu + r)) - \kappa$. In Section 2.5.3 we showed that unrevealed white α workers receive a higher lifetime payoff than their black counterparts; therefore, the incentive for the latter to invest in education is greater. As this implies that $\kappa_W \equiv \max\{\kappa : V_{Educ}(\kappa) \geq V_W^\alpha\} < \max\{\kappa : V_{Educ}(\kappa) \geq V_B^\alpha\} \equiv \kappa_B$, we must have that $F(\kappa_W) < F(\kappa_B)$ and therefore more black workers will purchase education. In particular, there exists some range of idiosyncratic costs for which black workers will purchase education but white workers will not. This is consistent with the finding in Lang and Manove (2011) that, conditional on past test scores, blacks get more education than whites do. The intuition here is simple; if a worker of high skill is treated as if she has the average hire's skill for her group, she has a greater incentive to reveal her high skill if that average is lower.³⁰

Perhaps equally importantly, this extension suggests that blacks and whites with high observable skills will have similar outcomes as discussed in the previous subsection.

messy dynamic issues which we sidestep by assuming that the fraction of additional workers who would choose to invest is insufficient to overturn (C5).

²⁹We analyze the case of a pure signal. If education can also turn a β into an α , the analysis is a combination of the analysis in this and the prior subsection since productive investment increases the fraction of workers who are α but investment that reveals workers to be α reduces the fraction of unrevealed workers who are α .

³⁰Strictly speaking, this creates a feedback loop from lower wages for the uneducated to a greater measure of education. The right assumptions on F rule out associated complexities.

Imperfect monitoring

Reader may have noted that the intuitive example presented first is distinct from the main model where monitoring resolves all uncertainty about worker type. As the example demonstrates, one can write a very similar model in which β workers always match badly but monitoring can result in false positive good matches. Given a wage-determination mechanism with outcomes similar to our bargaining protocol, much of the analysis would remain unchanged.³¹ Parameters would exist that would force monitoring on blacks but not whites, the black labor market would churn, and it would produce higher unemployment duration and lower lifetime wages for blacks. In this formulation, black workers succeeding at monitoring would only be as good as whites who had never been monitored; therefore a churned market does not necessarily produce better long-run matches or higher wages for the successfully monitored.

However, this alternate model would imply that some workers are purely parasitic and cannot be matched well, but rather aspire simply to find a job where their lack of productivity is undiscovered. An equivalent of (C6) cannot hold here and as a result we cannot rule out equilibria where negotiations sometimes break down and separation occurs without monitoring producing information.

Stigma and degeneration into lower-skilled jobs

Our model unrealistically assumes that employers have no information regarding the time that workers have been in the labor market or the number of jobs they have held. If the other aspects of our model were a rough representation of reality, it is implausible that firms would not recognize that some workers were unlikely to be new entrants and therefore very likely to β types. Suppose also that if a worker is

³¹Unfortunately, this alternate model would add a lot of complexity and require additional assumptions for uniqueness, due to the lack of a single posterior following successful monitoring. Bargaining strategies would have a much more tangled relation to beliefs and wages, and S3 would not be an apt tool to facilitate the task.

sufficiently likely to be a β , it is not efficient to employ or monitor him. Then workers who do not find a good match sufficiently quickly will be permanently barred from the monitoring sector.

Somewhat more formally, as an extension to the model, we can relax the assumption that past history is entirely unobservable. Assume instead that each separation has a probability ζ of becoming public common knowledge. Any worker who has a revealed separation is known to be of type β in any new match. Thus, a newly hired worker who does not have such a stigma will be of average quality $\theta'_B = [\beta + g\zeta(1 - \beta)] / [g\beta + (1 - g) + g\zeta(1 - \beta)]$. If we assume θ'_B satisfies (C5), churning can persist but will be primarily a phenomenon for relatively young workers.

But what will happen to workers revealed to be β s? It is straightforward to extend the model to allow for a second occupation type (q', c') lacking monitoring technology³² that is less skill intensive than the task described so far, i.e. $q > q'$ and $q - \lambda c < q' - \lambda c'$. As unrevealed β types are strictly better off than revealed ones in a new match of the first task, there must be q', c' such that the revealed β types prefer to enter the job market for the second occupation but the unrevealed ones do not.

In this scenario, a fraction of black workers are relegated to low-wage jobs while white workers with similar skills can always get better jobs. Furthermore, since the low-wage jobs are not monitored, they are a terminal state, with no possibility of promotion or escape.

Changing screening and monitoring technology

Autor and Scarborough (2008) examine the effect of bringing in a new screening process. They find that the screening process raised the employment duration of both

³²Or, more palatably, the same technology but without the incentives to use it, as in the case of a small enough c' .

black and white workers with no noticeable effect on minority hiring. In our model, we can think of this technology as allowing the firm to screen for job match quality prior to employment. This increases the proportion of hired blacks who become permanent workers since some bad matches are not hired. If the screening mechanism is good enough, the firm will choose not to monitor the black workers it hires, and all black workers will be permanent. Formally, since all white workers are permanent in the absence of the screen, the screen does not affect this proportion. Informally, if poor matches are more likely to depart even without monitoring, then there will also be positive effects on white employment duration.³³ Similarly, Wozniak (2015) shows that drug testing increases black employment and reduces the wage gap; we interpret this as confirming evidence for the notion that employers are more uncertain about the quality of black workers, and therefore that black workers benefit more from early resolution of such uncertainty.³⁴

We note that improved technology appears to have reduced monitoring costs. This is unambiguously good for blacks who share the cost of being monitored. Unless the reduction shifts whites into the monitoring equilibrium, they are unaffected by the cost reduction. However, if firms begin monitoring, α workers and the firm will initially be better off. On the other hand, β workers will generally be worse off as they will not be able to oppose the use of monitoring without revealing themselves. In a collective bargaining setting, the union might resist monitoring. The more interesting point is that since monitoring creates an externality, it is easy to develop an example in which monitoring makes both types of workers and capital worse off in the long run.³⁵

³³Formally, the model would have to be modified to ensure that some β workers are never perfectly matched and/or that some β workers are still in bad matches when they exit the labor force.

³⁴Wozniak (2015) is not to be interpreted as evidence that monitoring is good for black workers on the aggregate. As in the present paper, it can be beneficial on an individual level; our model, however, shows it can also create a worse externality.

³⁵Suppose that g_0 is just sufficient to sustain a no-monitoring equilibrium. A small reduction in b puts the labor market into a monitoring equilibrium. If there were no subsequent churning, α

2.5.6 Additional Empirical Content

This paper's aim is to explain the employment duration differential; its chief extra predictions are longer unemployment duration and lower lifetime income for black workers. In the interest of falsifiability, we posit here additional empirical implications that we view as following from our explanatory hypothesis and are relatively model-free.

We expect that when jobs vary with respect to the cost or accuracy of monitoring technology, black workers would skew more heavily towards those that favor monitoring. If the firm's cost of monitoring is lower, the initial match surplus with black workers is greater with benefits split between higher initial wages and lower labor costs. If the monitoring is faster at a firm,³⁶ α workers can reveal their ability and reap better wages sooner, while the firm will keep bad matches for less time. Either case produces a comparative advantage for this firm in hiring black workers. Jobs with high monitoring potential could for example be recording employee-customer interactions, as incoming call centers do.

Our explanation for lower black employment duration involves learning about match quality. In our model, the separation hazard into unemployment at time t is $h_t = (1 - \theta_B)\lambda e^{-\lambda t} / (\theta_B + (1 - \theta_B)e^{-\lambda t})$ for black workers and, rather starkly, 0 for white workers. More realistically, we expect the gap in the hazard rate to be declining in seniority.

2.6 Example

Here we provide a simple numerical example satisfying our conditions.

workers and firms would experience a slight gain, but the churning will wipe this out and more. Firms always make zero profit on vacancies, but if we allow for a distribution of vacancy costs, then the rents earned by firms with low costs of creating vacancies will also fall.

³⁶while keeping type productivity constant; that is, the firm has $\lambda' > \lambda$ and $\lambda'c' = \lambda c$.

Take $r = .05$, suggesting a unit of time of about a year. $\lambda = 2$ so that the average unsuccessful job lasts six months. $\beta = .2$ and $g = .95$, implying that most workers are good matches. $c = .5$ and $q = 1$ so that bad matches produce expected flow output of 0, making separation efficient regardless of unemployment duration. Finally, $b = 1$ so that $b/\lambda = 1/2$, making the expected cost of monitoring roughly equal to the value of six months of output from a well-matched worker.

It is readily verified numerically that the numerical conditions hold.

These parameters result in a θ_W of .96 and churning produces a θ_B of .833. The value of a filled vacancy in the white market is 9.6 and in the black labor market 8.9. Postulating a Cobb-Douglas matching function with elasticity $\eta = .75$, the model predicts a black-white unemployment duration ratio of 1.25. We can now compute the ratio of white to black income PDV at birth to be 1.11.

This example illustrates that a churning equilibrium is possible even if the proportion of type β workers is quite low in the population, and can generate reasonable income and unemployment disparities while doing so; one would do well to bear in mind, however, that our model applies *conditional* on observables and therefore cannot be calibrated to make economy-wide predictions.

2.7 Conclusions

We have developed a model that explains the black-white employment duration differential, and in the process have uncovered a mechanism that both reproduces standard empirical findings and makes novel predictions.

Our model in some ways resembles models of adverse selection in the labor market. Displaced workers are worse, on average, than a randomly selected worker. However, in contrast with standard adverse selection models, firms cannot distinguish between displaced workers and other unemployed workers. Therefore displaced workers depress

the wage of all unemployed workers. At the same time, our approach does not generate the asymmetry of information among firms that drives adverse selection models. If the worker is known to be good at a particular job, he will not leave for another job even if he knows that, on average, he will be good at other jobs. If the worker is bad at this particular job, he separates immediately.

To keep the analysis simple, our model assumes that workers who turn out to be well-matched remain employed forever. At first blush, this suggests that it applies only to new entrant unemployment because the market will surely recognize that a fifty-year old worker is not a new entrant. We believe it is more realistic to assume that the market cannot tell whether a fifty-year old worker who was laid-off six months ago has just been unlucky and not had any matches or has had a match that turned out to be bad. The market often cannot tell how long the worker has been unemployed. Thus we think the model is more general. In addition, it provides some insight into the scarring effect of unemployment.

Unlike most, perhaps all, existing models, ours can explain a number of empirical regularities regarding discrimination simultaneously:

1. Black workers have shorter employment durations.
2. Black workers have longer unemployment durations.
3. Black workers have lower lifetime earnings.

As written, the model has infinitely lived matches and agents so there is no unemployment rate. Allowing for deaths, we would have well-defined unemployment rates and would predict the rate for black workers is higher.

More generally, we view the main message of this paper as robust to many of the modeling decisions. The key element is that blacks are subject to more scrutiny or to a higher standard than white workers. This leads to more blacks being returned

to the unemployment pool, lowering the quality of that pool and, completing the equilibrium, making tougher scrutiny optimal.

Our model also strongly suggests that history matters and that equality of opportunity is not enough to eliminate racial disparities in the labor market even if this concept is used very expansively. The fact that blacks historically had low skills leads to an equilibrium in which the pool of black job seekers has lower skills than the pool of white job seekers even when the distribution of skills among all workers is identical for blacks and whites. While, over time, a human capital-based policy could mitigate labor market discrimination, achieving equality in human capital may be insufficient to eliminate racial disparities in the labor market.

Chapter 3

An Impossibility Theorem in Repeated Games (with Sambuddha Ghosh)

3.1 Introduction

Repeated games provide the standard model to show when and how concerns about the future of a repeated interaction can overcome myopic incentives to cheat. The canonical model with perfect monitoring is well studied. It involves a finite number of players choosing simultaneously from finite actions spaces; actions are revealed at the end of the period; total payoffs are defined as the discounted sum of the period payoffs. Perhaps best known among the repeated games results are the so-called folk theorems, which pin down the set of equilibrium payoff vectors as the common discount factor converges to unity (equivalently, the discount rate goes to zero). Friedman (1971) uses reversion to Nash equilibrium to produce ‘trigger strategies’ that form equilibria at any discount factor great enough. Such strategies can deliver any feasible average payoff better than some stage-game Nash equilibrium for each player. The first comprehensive folk theorem is that of Fudenberg and Maskin (1986), henceforth FM: Any payoff profile that is feasible and strictly individual rational (henceforth FSIR) in the

stage-game is the average discounted payoff of a subgame-perfect Nash equilibrium (SPNE) of the infinitely repeated game if each player is sufficiently patient; from the best response property inherent in any Nash equilibrium it is clear that one cannot get strictly lower payoffs, which justifies the label ‘comprehensive’.

It is usual to assume in repeated games that the exact discount factors are known to all players and to the game theorist. In reality the game theorist who recommends a course of action, for example by designing mechanisms for repeated auctions, might not know the discount factor of the players exactly. If strategies are very sensitive to discount factors, small amounts of ignorance can have dramatic effects. Indeed, one can also imagine that in real life players themselves cannot be certain that the others discount the future in exactly the same way as they do.

Kalai and Stanford (1988) provides a notion of discount robustness that is of use to us: strategy profiles form a Discount Robust Subgame Perfect (DRSP) equilibrium of a game if they are also an equilibrium of the same game with ever so slightly different discount factors. This is a *local* notion of robustness, which they show is useful in proving several results. We will additionally be interested in a stronger notion, which we dub Blackwell-Nash (BN) equilibrium, in which strategies are additionally required to form an equilibrium if players had weakly greater discount factors than initially supposed. This second notion allows a game theorist to design equilibrium strategies if all she has available is a lower bound on each player’s discount factor.

Nash reversion, introduced in Friedman (1971), can give a folk theorem for BN equilibria covering part of the feasible set. However, folk theorems such as the one in FM that cover the entire FSIR set preclude deviations from equilibrium by postulating severe punishments off-path. In particular, deviating players are subjected to ‘minmaxing’, multiple periods where they are given their worst possible individually rational payoff. The players executing the punishment are then rewarded, in part for

their mixing over actions that may offer varied stage-game payoffs in the punishment phase; they are rewarded more for using actions that give them less immediate payoff. Such rewards would be impossible to calibrate if the discount factor of each player was not precisely known. One may think, then, that as the strategies sketched by these folk theorems were not designed with discount robustness in mind, an alternative technique could support payoffs in the whole of the FSIR robustly.

This paper establishes that this is not, in general, the case. We prove a negative result, showing the existence of a critical point in the FSIR, so that points not in the positive orthant defined by it cannot be robustly supported under either notion of robustness. That is to say, our impossibility theorem shows that not only do current techniques fail to construct such strategies, but that such strategies cannot ever be constructed.

Having precluded robust equilibria in a region of the FSIR, we then set out to salvage the rest of the FSIR. We show that payoffs in (the interior of) the complement of the excluded region can always be supported robustly, with respect to both notions. In particular, we prove a discount-robustness folk theorem: for each such payoff, if players are sufficiently patient, there exist equilibrium strategies that provide that payoff that are robust to discount vectors, both locally and to all weakly greater ones.

It is perhaps surprising that although we are using two nested robustness concepts, the regions of the FSIR over which they can hold are the same. However, the two concepts share an important commonality: both require robustness to an uncountable set of discount vectors. The intuition for our results lies with the ability to calibrate punishments for deviations. If players' discount factors are fixed, players participating in the punishment of another can later be rewarded for mixing between actions that do not give the same stage-game payoff. However, if the discount factors may vary

over an uncountable set, it is not possible to keep indifference under all possibilities. Therefore, certain kinds of punishments are impossible; as players will not accept payoffs that are not supported by worse punishments, they cannot get such payoffs in equilibrium. We thus show that the ability to punish, developed for (constructive) folk theorems, can therefore be thought of as a necessary condition, not merely a sufficient one, in thinking about payoffs feasible in equilibrium.

3.2 Preliminaries

We consider a standard infinitely repeated game with perfect monitoring and possibly unequal discounting. At each $t \in \mathbb{Z}_+$ the (finite) stage-game $G = \langle I; (A_i)_i; (g_i)_i \rangle$ is played, where I is the set of players $\{1, \dots, n\}$, A_i is player i 's finite set of actions, $A := \times_{i \in I} A_i$ is the set of all pure action profiles, and $g_i : A \rightarrow \mathbb{R}$ is player i 's payoff function. A mixed action of i is $\alpha_i \in \Delta A_i$, where ΔE is the set of all probability distributions on a set E . Let $\mathbf{a}^{(t)} \in A$ be the (realized) action profile played at time t .¹ When player i discounts future payoffs using the discount factor δ_i , player i 's average discounted utility defined over infinite sequences of pure actions in A is

$$u_i(\{\mathbf{a}^{(t)}\}_{t=0}^{\infty}) := (1 - \delta_i) \sum_{t=0}^{\infty} \delta_i^t g_i(\mathbf{a}^{(t)}).$$

Under perfect monitoring the public history at the end of period t is $h^t = (\mathbf{a}^{(0)}, \dots, \mathbf{a}^{(t)}) \in A^{t+1}$ (starting with the empty history h^{-1}). A pure strategy of i is $s_i(t+1) : H^t \rightarrow A_i$ (for $t = -1, 0, 1, \dots$) where H^t denotes the set of histories at the end of period t ; mixed strategies are analogous, except that they map to the corresponding mixed actions ΔA_i . We also allow strategies to be conditioned on

¹In what follows vectors are boldfaced while scalars and sets are not. Sequence indices are denoted by superscripts and sometimes they are enclosed in parentheses to distinguish them from exponents or from another sequence denoted by the same letter; for example, \mathbf{c}^l denotes the l -th vertex of a polytope C , while $\{\mathbf{c}^{(t)}\}$ denotes an infinite sequence of vertices each element of which is a \mathbf{c}^l for some l . Coordinates of vectors are denoted by subscripts.

the realization of a publicly observable random variable (henceforth PRD), but do not explicitly model it. This describes the repeated game $G^\infty(\boldsymbol{\delta})$, where the vector $\boldsymbol{\delta} = (\delta_1, \dots, \delta_n)$ is referred to as the discount factor vector. In the special case where each player discounts the future at the rate δ , we obtain the game $G^\infty(\delta)$.

Player i 's minmax value is

$$w_i := \min_{\boldsymbol{\alpha}_{-i} \in \times_j (\Delta A_j)} \max_{a_i} g_i(a_i, \boldsymbol{\alpha}_{-i}),$$

while her pure-action minmax value is

$$w_i^p := \min_{\boldsymbol{\alpha}_{-i} \in A_{-i}} \max_{a_i} g_i(a_i, \boldsymbol{\alpha}_{-i}).$$

Let $\mathbf{m}^i \in \times_j (\Delta A_j)$ be a strategy profile that minmaxes i , with player i playing a best response; for the pure minmax case we define the corresponding action profile as $\mathbf{m}^{p,i} \in A$. The feasible set is the convex hull of the set of pure-action stage-game payoffs, $F := \text{co}(g(A))$, and the feasible strictly individually rational (FSIR) set is $F^* := \{\mathbf{x} \in F \mid x_i > w_i\}$; by analogy with it we can define $F^p := \{\mathbf{x} \in F \mid x_i > w_i^p\}$. The lower boundary of F is $\underline{\partial}F := \{\mathbf{x} \in F : \nexists \mathbf{y} \in F \text{ such that } \mathbf{y} \ll \mathbf{x}\}$.

FM shows that for any $\mathbf{v} \in F^* \setminus \underline{\partial}F$ there exists a $\underline{\delta} \in (0, 1)$ such that if the common discount factor δ satisfies $\delta > \underline{\delta}$, there is a subgame perfect Nash equilibrium $(\sigma_i)_{i \in I}$ of $G^\infty(\delta)$ such that for each $i \in I$, player i 's payoff is v_i .²

3.3 Discount Robustness

Folk theorems provide a surprisingly large set of possible equilibrium payoffs. Motivated by this, a variety of later work first questions and then proves the robustness of the folk theorem to a variety of alternative assumptions.

In this work, we ask to what extent payoffs in F^* can be supported by strategies

²This result is extended to $\underline{\partial}F$ in Abreu, Dutta and Smith (1991).

that form equilibria not at just a single discount vector, but rather a suitably larger set. However, it is not obvious what definition of robustness is appropriate. We will use two such definitions — one weak, and one strong. We borrow the first definition from Kalai and Stanford (1988).

Definition 1. [Kalai and Stanford (1988)] A **Discount Robust Subgame Perfect (DRSP)** equilibrium of $G^\infty(\boldsymbol{\delta})$ is a strategy profile σ such that there exists a neighborhood B of $\boldsymbol{\delta}$ for which if $\boldsymbol{\delta}' \in B$ then σ is an SPNE strategy profile of the game $G^\infty(\boldsymbol{\delta}')$.

Intuitively, this is a weak ‘local’ notion as it requires equilibrium strategies of the game $G^\infty(\boldsymbol{\delta})$ to remain equilibria if the discount factor is close enough to $\boldsymbol{\delta}$. In Kalai and Stanford (1988) it is shown that if a unique stage-game best response to opponents’ pure actions always exists, then a pure-strategy DRSP equilibrium involves each player playing a unique best-response strategy to opponent strategies.³ However, one may desire, in the spirit of the folk theorems, robustness to not just nearby discounting factors but all higher discount factors. This leads to the following concept, first formally introduced in Dasgupta and Ghosh (2015).⁴

Definition 2. A **Blackwell-Nash (BN)** equilibrium of $G^\infty(\boldsymbol{\delta})$ is a strategy profile σ such that there exists a $\underline{\boldsymbol{\delta}} \ll \boldsymbol{\delta}$ with the property that if $\boldsymbol{\delta}' \gg \underline{\boldsymbol{\delta}}$ then σ is a SPNE of $G^\infty(\boldsymbol{\delta}')$.

This is a natural multi-player version of the single-player notion from dynamic programming, where a player’s strategy in a dynamic choice problem is said to be Blackwell optimal if there exists a cutoff such that it is optimal for the single player to follow the strategy when the discount factor of the player exceeds the cutoff. If

³In fact, this proposition is proven using a corollary of the same power-series result we will use to prove one of our theorems.

⁴The word ‘formally’ is far from redundant. During a conversation at Yale, Ghosh had asserted that the equilibria constructed in Dasgupta and Ghosh (2015) had the merit of not needing to be fine-tuned, being robust to small variations in discount factors. This led Johannes Hörner to verbally propose to Ghosh the notion of Blackwell-Nash equilibrium.

the game theorist knows that each player's discount factor exceeds a critical value, she can design strategy profiles from which no player would want to deviate. This is the sort of robustness of Friedman (1971)'s folk theorem. A BN equilibrium not only provides local robustness near $\boldsymbol{\delta}$ but also to higher discount factors for each player, that is, it remains an equilibrium for all $\delta'_i \in [\delta_i, 1)$. Clearly, a BN equilibrium strategy profile is also a DRSP equilibrium strategy profile, but the reverse need not be true.

Our main question is *to what extent is each payoff \boldsymbol{v} in F^* obtainable in DRSP/BN equilibria of $G^\infty(\boldsymbol{\delta})$ for sufficiently large $\boldsymbol{\delta}$?* Surprisingly, despite the use of nested notions of robustness, the answers to these two questions coincide.

The key to our results is that these questions turn on whether or not players need future incentives to randomize. Recall from Fudenberg and Maskin (1986) that the issue of punishments becomes relevant only when it comes to punishing a player for a deviation. If players are asked to mix over actions that will give them the same stage-game payoff, given opponents' current actions, and future play is independent of which pure action is played, players are willing to mix. On the other hand, if players are to mix among actions that give them different current payoffs, they must be compensated later for taking a worse action today. This is easy enough if the discount factor is known,⁵ but it is harder to do so for multiple discount factors. In fact, we show that when the set of potential discount factors is uncountable, which is required for both DRSP and BN, it is impossible to calibrate such rewards to keep the player indifferent no matter what her discount factor is.

Motivated by the above, it is clear that payoff vectors fall into one of two types, those that can be given by equilibrium strategies without mixing over actions giving different myopic payoffs, and those that can not be. If player i is to get payoff v_i in a SPNE, it must be the case that she is unable to myopically best-respond to

⁵See, for example, Fudenberg and Maskin (1986).

opponents' stage-game actions and do better. But if randomization can only be used when myopically indifferent, opponent actions are restricted. For instance, it may be impossible to incentivize i 's opponents to minmax player i and give her the standard minmax w_i . With a limited potential to punish in either kind of discount-robust equilibrium, it should come as no surprise that not all of F^* remains incentive compatible.

3.3.1 The restricted minmax

To proceed along these lines, we must find the worst punishments that can be delivered with robustness to discount factors. To that end, we first need to define the set of action profiles where mixing does not require future rewards to generate indifference. If the action profile α is to make players myopically indifferent, each α_i should involve mixing only over elements in A_i that give the same payoff to i given α_{-i} ; that is,

Definition 3. An action profile α for stage-game G has the **myopic indifference property** if for each player i we have $a_i, a'_i \in \text{supp}(\alpha_i) \Rightarrow g_i(a_i, \alpha_{-i}) = g_i(a'_i, \alpha_{-i})$.

Let $Q \subset \times_i \Delta A_i$ denote the set of all action profiles with the myopic indifference property. All pure action profiles trivially satisfy the definition. Nash equilibria of G also have the myopic indifference property as they specify actions that give the same (maximum) payoff given other players' actions. In fact, one way to conceive of Q is as the collected Nash equilibria of the family of stage games with pure action sets $\{A' | A' = \times_i A'_i, A'_i \subset A_i\}$. That is, if $\alpha \in Q$, we can modify G into some G' by removing pure actions in such a way that α is a Nash equilibrium of G' .

We define player i 's restricted minmax (payoff) as

$$r_i = \min_{\alpha \in Q} \max_{a_i \in A_i} g_i(a_i, \alpha_{-i}) \quad (3.1)$$

That is, the restricted minmax for player i is the highest payoff she can get if other

players are using the worst possible actions for her, with the added constraint that those actions must form a myopically indifferent action profile when coupled with some action α_i for i . Notice that the action profile used to compute r_i is not necessarily myopically indifferent. Therefore it is both the case that r_i is (weakly) higher than the lowest payoff i attains in Q and (weakly) lower than the lowest payoff i attains in Q when he is best responding.

Now, define

$$F^{res} := \{\mathbf{x} \in F \mid \mathbf{x} \gg \mathbf{r}\}. \quad (3.2)$$

As stated, pure actions profiles all trivially satisfy myopic indifference. Therefore (with some abuse of notation) we have that

$$\times_i A_i \subseteq Q \subseteq \times_i \Delta A_i \quad (3.3)$$

and therefore that

$$\min_{\mathbf{a}_{-i} \in A_{-i}} \max_{a_i \in A_i} g_i(a_i, \boldsymbol{\alpha}_{-i}) \geq \min_{\boldsymbol{\alpha} \in Q} \max_{a_i \in A_i} g_i(a_i, \boldsymbol{\alpha}_{-i}) \quad (3.4)$$

$$\geq \min_{\boldsymbol{\alpha}_{-i} \in \times_j (\Delta A_j)} \max_{a_i \in A_i} g_i(a_i, \boldsymbol{\alpha}_{-i}) \quad (3.5)$$

$$\Rightarrow w_i^p \geq r_i \geq w_i. \quad (3.6)$$

The reader may be interested in the question of whether both inequalities in (3.6) can be strict at the same time; we provide an example where this is the case in Section 3.6. We are now prepared to present the paper's main results.

3.4 An Impossibility Theorem

Our first main result is that if $\mathbf{v} \notin \overline{F^{res}}$ then \mathbf{v} cannot be supported by DRSP equilibrium strategies. Quite simply, if some player i were to get below her restricted minmax payoff r_i , she could do better by best-responding on a period-by-period

basis; her opponents would not have a punishment apparatus available in a DRSP equilibrium that is powerful enough to dissuade her.

Theorem 7. *If $\mathbf{v} \notin \overline{F^{res}}$, then $\nexists \sigma, \delta$ such that σ is a DRSP equilibrium strategy profile of $G^\infty(\delta)$ with equilibrium payoff vector \mathbf{v} .*

The impossibility theorem precludes the existence of a DRSP equilibrium with payoff vector $\mathbf{v} \notin \overline{F^{res}}$ and therefore also disallows the existence of equilibria satisfying the stronger BN condition. Most surprisingly, the proof shows that if \mathbf{v} is an equilibrium payoff outside $\overline{F^{res}}$ supported by some strategy profile σ , for any compact set of discount factors $D \subset (0, 1)^n$, only up to a finite subset of D can support σ as an equilibrium strategy profile. As $(0, 1)^n$ can be written as the countable union of a collection of its compact subsets, that means that only up to a countable set of discount vectors can support σ , and those discount vectors form a totally disconnected set. That is to say, if $\mathbf{v} \notin \overline{F^{res}}$ is delivered by an SPNE σ of $G(\delta)$, not only is there no neighborhood of δ where σ remains an equilibrium, but there exist neighborhoods of δ where σ is an equilibrium *only* at δ .

3.5 Blackwell-Nash Equilibrium Payoffs

We now ask what payoffs can be supported in Blackwell-Nash equilibria. Let $F^{*p} := \{\mathbf{v} \in F \mid v_i > w_i^p \ \forall i\}$, where w_i^p is the pure strategy minmax, where players are restricted to pure strategies. It should be clear that all such payoffs can be supported in Blackwell-Nash equilibria exactly using the standard FM strategies; during the punishment phase of i all other players pick $\mathbf{m}^{p,i} \in A_{-i}$, the action profile against which i 's best response gives her her pure minmax payoff w_i^p . Player j 's failure to punish a deviator is immediately detected and met with punishment (pure minmax) and forgone rewards for enforcing punishment.

Indeed one can do better. We show that all payoffs in F^{res} may be obtained in

BN equilibria and therefore in DRSP equilibria. In light of the impossibility result presented above, this really is the best one could hope for. Perhaps counter-intuitively, there is no region of the feasible set for which DRSP strategies can be constructed, but BN strategies cannot; this is despite the fact that in general, equilibrium strategies with the former property but not the latter exist. Note that as $\times_i A_i \subseteq Q$, the theorem presented below can nest as a special case pure min-max punishment for payoff vectors $\mathbf{v} \gg \mathbf{w}^p$.

The particular theorem presented makes use of a Public Randomization Device. This allows on-path payoffs to be produced by an i.i.d. randomization over pure actions for each player every period; hence, the ex-ante expected payoff for each player in equilibrium will not depend on her discount factor. That is, not only does the theorem deliver strategies that form a Blackwell-Nash equilibrium delivering \mathbf{v} in $G^\infty(\boldsymbol{\delta})$; they deliver \mathbf{v} in $G^\infty(\boldsymbol{\delta}')$ for all $\boldsymbol{\delta}' \in [\delta, 1]^n$ as well. However, proving a BN-robustness folk theorem for F^{res} does not require a PRD. A version of this theorem which constructs BN strategies delivering \mathbf{v} in $G^\infty(\boldsymbol{\delta})$ and arbitrarily close to \mathbf{v} in $G^\infty(\boldsymbol{\delta}')$ for all $\boldsymbol{\delta}' \in [\delta, 1]^n$ and does not use a PRD is also possible.

The proof of the folk theorem presented here uses a construction based on the simple strategy profiles of Abreu (1988). Following any deviation by a player i , continuation play is identical. Naively, one may think in analogy with minmax punishments that a punishment is only feasible if the punishers are myopically indifferent on the support of their actions when the punished player is myopically best-responding. However, this is not the case. The key innovation in the proof is that if players are convinced the target of their punishment will cooperate in the punishment rather than play a myopic best-response, they can be induced to deliver harsher punishments than otherwise. By then ensuring that cooperating in one's punishment is preferable to myopically best-responding, as the latter prolongs the punishment phase, we deter

such deviation from punishment. The punished player provides the punishing players the incentives to punish, under threat that she will be required to do so for longer if she deviates. In this way, the present theorem differs substantially from previous folk theorems: the ability to incentivize the punishers *and* the punished is what restricts the equilibrium payoff set.

Theorem 8. *Let F be full-dimensional. Fix $\mathbf{v} \in F^{res}$ and not on the lower boundary of the set. There exists a cutoff discount factor $\underline{\delta}$ and a strategy profile $\sigma[\mathbf{v}]$ such that $\sigma[\mathbf{v}]$ is a Blackwell-Nash equilibrium of $G^\infty(\underline{\delta})$, with discounted average payoff \mathbf{v} when evaluated at $\delta \geq \underline{\delta}$.*

Thus, with Theorem 7 precluding robust equilibria with payoffs outside $\overline{F^{res}}$ and Theorem 8 producing BN equilibria for all points of F^{res} not on the lower boundary, we have achieved a characterization of discount robust equilibrium payoffs (modulo a set of measure zero).

Notice the strategies produced in Theorem 7 are invariant to the discount factors. Unlike FM, we do not adjust payoffs in Phase III based on the actions taken in Phase II. However, that adjustment was required in FM as they use ‘exact’ minimaxing. Alternatively, Gossner (1995) shows that deviations may be deterred by schemes with ‘approximate’ minmaxing. Such ‘codes of conduct’ specify on-path play exactly, but allow strategies to vary in Phase II with the parameters of the model. In such a way, one can give sharp descriptions of play in Phases I and III; play during Phase II is left unspecified, but players compute a statistical test to determine whether they have approximately minmaxed the deviator, and if not, whose punishment needs to be triggered.

It should not come as a surprise that although robust equilibria for $\mathbf{v} \in F \setminus \overline{F^{res}}$ do not exist, one can construct equilibria with payoffs \mathbf{v} such that only off-path play need vary with the discount factor. This further reinforces the basic intuition that discount robustness has bite because it affects the ability to punish deviations.

3.6 Example

This section provides an example game, for which we compare the feasible region F^* and the robustly feasible region F^{res} . Consider the stage-game given by the following payoff matrix.

		Player 2		
		L	M	R
Player 1	T	(5, 0)	(0, 2)	(5, 1)
	D	(0, 4)	(0, 5)	(6, 0)
	B	(8, 0)	(6, 3)	(0, 2)

Table 3.1: Example stage game G .

As player 2 has a strictly dominant strategy of the stage-game in M , it is easy to see that her pure, restricted and mixed minmax values coincide at $r_2 = w_2 = w_2^p = 2$ via (T, M) . For player 1, the pure minmax is $w_1^p = 6$ and is achieved with either (B, M) or (D, R) . Her mixed minmax is $w_1 = 3$ and given by $(qD + (1 - q)B, \frac{1}{2}M + \frac{1}{2}R)$ for any $q \in [0, 1]$. However, as M is a strictly dominant strategy in the stage-game, player 2 will never be myopically indifferent between M and R . Therefore, player 1 can never be given such a low payoff in a BN equilibrium. Instead, the restricted minmax involves player 2 mixing between L and R . A quick computation shows that $\alpha^i = (\frac{1}{3}D + \frac{2}{3}B, qL + (1 - q)R)$, where $q \in [\frac{1}{6}, \frac{5}{8}]$. Player 1's best response to $(qL + (1 - q)R)$ is T , which would yield a payoff of 5, so $r_1 = 5$.

Therefore, $F^* = \{x \in F | x \gg (3, 2)\}$ and $F^{res} = \{x \in F | x \gg (5, 2)\}$. The folk theorem in FM implies that, for instance, $\mathbf{v} = (3, 4)$ can be delivered in an SPNE of $G^\infty(\delta)$ for high enough δ , but Theorem 7 in this paper shows that no such SPNE is discount-robust.

Figure 3.1 shows F , F^* and F^{res} for game G . Although $r_1 = w_2$ means that the minimum payoff for player 2 is the same in F^{res} and F^* , the fact that $r_1 > w_1$ limits player 2's maximum payoff in robust equilibria.




Figure 3-1: The payoff space for game G . Starred region: payoffs that can be delivered by a BN equilibrium. Hatched region: SPNEs exist, but they are not robust.

3.7 Conclusion

We have shown that a sizable part of the FSIR, the set of payoffs proven to be equilibrium payoffs in folk theorems, cannot be supported by discount-robust strategies. Even a local notion of robustness is unattainable for equilibria with payoffs in that region. However, we have shown in a robust folk theorem that the entire complementary payoff region can be supported by both locally and strongly discount-robust equilibria. We do this by construction, designing simple strategy profiles that deliver any point in the robust region.

Furthermore, in the course of addressing the robustness question, we discover a new point of interest, the restricted minmax. We have shown that it describes the

ability to punish when punishers' incentives cannot be finely tuned via future rewards for mixing. This intuition explains why certain regions of the payoff space cannot be robustly attained, and additionally, how it is possible to construct equilibria in the robustly feasible region.

Therefore, the present research establishes that care must be taken when discount factors are not precisely known. While we present a tool to get around this problem, it - and all other tools - may not always be applicable.

Appendix A

Proofs of results in Chapter 1

A.1 Proof of Theorem 2

Recall the *belief function* \hat{p}_n that takes opponent type-wise acceptance probabilities $(A_m)_{m \neq n}$ into the Bayesian posterior for a visiting applicant at firm n and the function \hat{A}_n that takes beliefs for firm n into optimal actions. The function $\hat{A}_n \circ \hat{p}_n$ thus takes opponent type-wise acceptance probabilities into (unique) optimal acceptance probabilities for firm n and is therefore a best-response function. As the composition of continuous functions, it is continuous. Forming the N -firm best response function $\hat{A} \circ \hat{p} : [0, 1]^{|Θ|N} \rightarrow [0, 1]^{|Θ|N}$ using each firm's best response functions, it inherits continuity and is defined over a compact¹, convex set. Therefore the Brouwer Fixed Point Theorem gives us the existence of a fixed point in $\hat{A} \circ \hat{p}$ and hence a Nash Equilibrium. ■

A.2 Proof of Theorem 3

Once a firm acquires information, it will learn the applicant's type and accept iff $\theta > 0$. Therefore, there are only three pure strategies and their mixtures to consider:

¹To be precise, the initial action space is not necessarily compact; however, best responses always fit Woodford (2008)'s description.

staying uninformed (choosing signal g_{no} , the completely uninformative signal structure) and rejecting all applicants, staying uninformed and accepting all applicants, and becoming informed to accept iff $\theta > 0$. Using formula (1.2) and noticing g_{no} offers no information and therefore $H(p_n) = H(p_n(\cdot|s))$ when s is drawn according to g_{no} , it is the case that $M(p_n, g_{no}) = 0$. Therefore, remaining uninformed and rejecting the applicant gives a payoff of 0. Accepting the applicant without information results in a payoff of

$$\sum_{\theta \in \Theta} \theta p_n(\theta) - kM(p_n, g_{no}) = \sum_{\theta \in \Theta} \theta p_n(\theta). \quad (\text{A.1})$$

Becoming informed by choosing g_{all} (the completely informative signal structure) eliminates all uncertainty in the posterior, and therefore $H(p_n(\cdot|s)) = 0$; therefore $M(p_n, g) = H(p_n) = -\sum_{\theta \in \Theta} p_n(\theta) \ln p_n(\theta)$. Thus, becoming informed and accepting iff $\theta > 0$ results in a payoff of

$$\sum_{\theta > 0} \theta p_n(\theta) - kM(p_n, g_{all}) = \sum_{\theta > 0} \theta p_n(\theta) + k \sum_{\theta \in \Theta} p_n(\theta) \ln p_n(\theta). \quad (\text{A.2})$$

First, suppose that the equilibrium belief for each firm n , p_n , is equal to the prior p_0 . Then, if (A.1) is greater than (A.2), as (1.9) implies that if $\forall m \neq n, \forall \theta A_m(\theta) = 1$, and that in turn delivers that for each n , $\hat{p}_n((A_m)_{m \neq n}) = p_0$. Therefore, we have an equilibrium in pure strategies where $\forall n, g_n = g_{no}$ and $\forall s, a_n(s) = 1$. That is, it is an equilibrium that no firm becomes informed and all applicants are always accepted.

Otherwise, uninformed firms will reject the applicant. Each firm chooses a probability z_n of becoming informed. The set of optimal information acquisition strategies, as a function of p_n , is therefore given by

$$Z_n(p_n) = \begin{cases} \{0\} & \sum_{\theta>0} \theta p_n(\theta) + k \sum_{\theta \in \Theta} p_n(\theta) \ln p_n(\theta) < 0 \\ [0, 1] & \sum_{\theta>0} \theta p_n(\theta) + k \sum_{\theta \in \Theta} p_n(\theta) \ln p_n(\theta) = 0 \\ \{1\} & \sum_{\theta>0} \theta p_n(\theta) + k \sum_{\theta \in \Theta} p_n(\theta) \ln p_n(\theta) > 0 \end{cases}$$

which is convex-valued and closed-graph.

Optimal actions following information acquisition for firm n satisfy $\forall \theta < 0, A_n(\theta) = 0$ and $\forall \theta > 0, A_n(\theta) = z_n$. Thus, using (A.2) we can write the set of optimal strategies in terms of \hat{A}_n , a convex-valued, closed-graph correspondence in p_n . Recall that \hat{p}_n given by (1.9) is continuous. Therefore, $\hat{A}_n \circ \hat{p}_n : [0, 1]^{N-1} \rightarrow 2^{[0,1]}$ mapping opponent acceptance rates into optimal own acceptance rates is a best-response correspondence for firm n and as a composition of a closed-graph, convex-valued correspondence on a continuous function, in turn convex-valued and closed-graph. This property is inherited by the N -firm best-response correspondence $\hat{A} \circ \hat{p} : [0, 1]^N \rightarrow 2^{[0,1]^N}$ and hence as $[0, 1]^N$ is a compact, convex subset of \mathbb{R}^N , the Kakutani Fixed Point Theorem applies, providing the existence of a fixed point and thereby an equilibrium. ■

A.3 Proof of Theorem 4

The proof proceeds inductively. Theorem (1) establishes that that profitability for firm n implies that $\sum_{\theta} p_n(\theta) e^{\frac{\theta}{k}} > 1$; a monopolist has $p_n = p_0$.

Let \hat{p}_m^N be the belief function for firm m of N firms as defined in (1.9), and let $\hat{A}_m^N : \Delta\Theta \rightarrow [0, 1]^{|\Theta|}$ be firm m of N 's optimal acceptance probability function, given by Theorem 1, so that $\hat{A}_m^N \circ \hat{p}_m^N$ is firm m 's best response function in the N -firm game. Let $\hat{A}^N \circ \hat{p}^N$ be the N -firm best response function, so that an equilibrium corresponds to a fixed point of the function. For each $N \in \mathbb{N}$ such a fixed point is guaranteed to exist by Theorem 2.

Suppose, for induction, that every fixed point of $\hat{A}^N \circ \hat{p}^N$ has $\forall \theta, n, A_n(\theta) > 0$ and

therefore positive expected profits by Theorem 1. Recall that if in equilibrium $A_n(\theta) > 0$, firm n gets a strictly positive payoff. Assume, for contradiction, an equilibrium given by $(A_n, p_n)_{n \leq N+1}$ in the market with $N + 1$ firms where firm $N + 1$ gets a payoff of 0. Therefore $(A_n)_{n \leq N+1}$ is a fixed point of $\hat{A}^{N+1} \circ \hat{p}^{N+1}$ and $\forall \theta, A_{N+1}(\theta) = 0$. Thus, using (1.9), for each firm $m \leq N + 1$, $\hat{p}_m^{N+1}((A_n)_{n \leq N+1}) = \hat{p}_m^N((A_n)_{n \leq N})$. Therefore, as $\hat{A}_n^N = \hat{A}_n^{N+1}$, $(A_n)_{n \leq N}$ is a fixed point of $\hat{A}^N \circ \hat{p}^N$. As by the inductive hypothesis every such fixed point has $\forall \theta \in \Theta, \forall n \leq N, A_n(\theta) > 0$, each firm $n \leq N$ is getting a strictly positive payoff in the $N + 1$ firm equilibrium as well.

What remains to be shown is that firm $N + 1$ has a profitable deviation. As each firm $n \leq N$ gets a positive payoff, from Theorem (1),

$$\sum_{\theta \in \Theta} p_n(\theta) e^{\frac{\theta}{k}} > 1 \quad (\text{A.3})$$

and A_n , the acceptance probability vector for n , if not identically 1, is given by

$$\frac{A_n(\theta)}{1 - A_n(\theta)} = \frac{\bar{A}_n}{1 - \bar{A}_n} e^{\frac{\theta}{k}} \quad (\text{A.4})$$

$$\bar{A}_n = \sum_{\theta \in \Theta} p_n(\theta) A_n(\theta). \quad (\text{A.5})$$

Combining the above, and manipulating, we get

$$\bar{A}_n = \sum_{\theta \in \Theta} \frac{p_n(\theta) \frac{\bar{A}_n}{1 - \bar{A}_n} e^{\frac{\theta}{k}}}{1 + \frac{\bar{A}_n}{1 - \bar{A}_n} e^{\frac{\theta}{k}}} \quad (\text{A.6})$$

$$1 = \sum_{\theta \in \Theta} \frac{p_n(\theta) e^{\frac{\theta}{k}}}{1 - \bar{A}_n + \bar{A}_n e^{\frac{\theta}{k}}} \quad (\text{A.7})$$

The distribution of types rejected by firm $n \leq N$, q_n , is given by:

$$q_n(\theta) = \frac{p_n(\theta)(1 - A_n(\theta))}{\sum_{\theta' \in \Theta} p_n(\theta')(1 - A_n(\theta'))} = \frac{p_n(\theta)(1 - A_n(\theta))}{1 - \bar{A}_n}. \quad (\text{A.8})$$

Substituting for $A_n(\theta)$ from (A.4) we get

$$q_n(\theta) = \frac{p_n(\theta)}{\left(1 + \frac{\bar{A}_n}{1 - \bar{A}_n} e^{\frac{\theta}{k}}\right) (1 - \bar{A}_n)} = \frac{p_n(\theta)}{1 - \bar{A}_n + \bar{A}_n e^{\frac{\theta}{k}}} \quad (\text{A.9})$$

and using (A.7) we have

$$\sum_{\theta \in \Theta} q_n(\theta) e^{\frac{\theta}{k}} = \sum_{\theta \in \Theta} \frac{p_n(\theta) e^{\frac{\theta}{k}}}{1 - \bar{A}_n + \bar{A}_n e^{\frac{\theta}{k}}} = 1 \quad (\text{A.10})$$

As A_{N+1} is identically 0 by hypothesis, the distribution of rejected types for $n \leq N + 1$ is independent of whether the applicant has visited n . Therefore, we can write the type distribution of applicants visiting firm $N + 1$, p_{N+1} as a weighted sum of the distribution of rejected applicants at firms $n \leq N$ and the new applicant distribution p_0 :

$$p_{N+1}(\theta) = \sum_{n \leq N} \phi_n q_n(\theta) + \left(1 - \sum_{n \leq N} \phi_n\right) p_0(\theta) \quad (\text{A.11})$$

where ϕ_n is the probability an applicant visited firm n last before visiting firm $N + 1$, conditional on firm $N + 1$ being visited. If a firm sets $\bar{A}_n = 1$, it has no rejected applicants and thus for that n , $\phi_n = 0$. The remaining probability $1 - \sum_{n \leq N} \phi_n$, is the probability (conditional on $N + 1$ being visited) that firm $N + 1$ is the first firm visited; note $1 - \sum_{n \leq N} \phi_n \geq \frac{1}{N+1}$. Then,

$$\sum_{\theta \in \Theta} p_{N+1}(\theta) e^{\frac{\theta}{k}} = \sum_{n \leq N} \phi_n \sum_{\theta \in \Theta} q_n(\theta) e^{\frac{\theta}{k}} + \left(1 - \sum_{n \leq N} \phi_n\right) \sum_{\theta \in \Theta} p_0(\theta) e^{\frac{\theta}{k}} \quad (\text{A.12})$$

so that substituting (A.8) we get

$$\sum_{\theta \in \Theta} p_{N+1}(\theta) e^{\frac{\theta}{k}} = \sum_{n \leq N} \phi_n \sum_{\theta \in \Theta} q_n(\theta) e^{\frac{\theta}{k}} + \left(1 - \sum_{n \leq N} \phi_n\right) \sum_{\theta \in \Theta} p_0(\theta) e^{\frac{\theta}{k}} \quad (\text{A.13})$$

and using (A.10) we attain

$$\sum_{\theta \in \Theta} p_{N+1}(\theta) e^{\frac{\theta}{k}} = \sum_{n \leq N} \phi_n + (1 - \sum_{n \leq N} \phi_n) \sum_{\theta \in \Theta} p_0(\theta) e^{\frac{\theta}{k}}. \quad (\text{A.14})$$

Now, since by hypothesis $\sum_{\theta \in \Theta} p_0(\theta) e^{\frac{\theta}{k}} > 1$, we have that

$$\sum_{\theta \in \Theta} p_{N+1}(\theta) e^{\frac{\theta}{k}} > 1 \quad (\text{A.15})$$

but this is the condition under which Theorem 1 guarantees firm $N + 1$ a positive payoff, a contradiction. Therefore, all firms in the $N + 1$ firm market get positive payoff. By PMI, this holds for all $N \in \mathbb{N}$, and the theorem is proven. ■

A.4 Proof of Theorem 5

First, suppose the payoff to becoming informed at the prior $\Pi = \sum_{\theta > 0} p_0(\theta) \theta + k \sum_{\theta \in \Theta} p_0(\theta) \ln p_0(\theta)$ is weakly negative. Then no firm (even a monopolist) can ever profit and $\bar{N} = 0$ trivially satisfies both sides. The rest of the proof assumes that the monopolist would make strictly positive profits..

Firms always reject known $\theta < 0$, and by hypothesis $k \sum_{\theta \in \Theta} p_0(\theta) \ln p_0(\theta) > \sum_{\theta < 0} \theta p_0(\theta)$ so that acceptance without information is not optimal; therefore, no $\theta < 0$ is ever accepted. Thus, we can write the equilibrium posterior for firm n as a function of x , the total probability that a $\theta > 0$ applicant was not accepted at a previous firm:

$$p_n(\theta) = \begin{cases} \frac{x p_0(\theta)}{x \sum_{\theta > 0} p_0(\theta) + \sum_{\theta < 0} p_0(\theta)} & \theta > 0 \\ \frac{p_0(\theta)}{x \sum_{\theta > 0} p_0(\theta) + \sum_{\theta < 0} p_0(\theta)} & \theta < 0. \end{cases} \quad (\text{A.16})$$

The payoff to becoming informed by choosing g_{all} (and accepting $\theta > 0$) at p_n is

$$x \sum_{\theta>0} \theta p_0(\theta) - kH(p_n(\theta)) \quad (\text{A.17})$$

so that substituting in (A.16) and manipulating, we get

$$\begin{aligned} \Pi(x) = & \left[x \sum_{\theta>0} p_0(\theta) + \sum_{\theta<0} p_0(\theta) \right]^{-1} \cdot \\ & \left[x \sum_{\theta>0} p_0(\theta)\theta + xk \sum_{\theta>0} p_0(\theta) (\ln x + \ln p_0(\theta)) + k \sum_{\theta<0} p_0(\theta) \ln p_0(\theta) \right] \\ & - k \ln \left(x \sum_{\theta>0} p_0(\theta) + \sum_{\theta<0} p_0(\theta) \right). \end{aligned}$$

Notice $\Pi(0) \leq 0$ and that, as $x = 1$ corresponds to the prior and it is assumed that $\sum_{\theta>0} p_0(\theta)\theta + \sum_{\theta \in \Theta} p_0(\theta) \ln p_0(\theta) > 0$, we have $\Pi(1) > 0$. Now, consider the derivative of the payoff to becoming informed wrt x :

$$\begin{aligned} \Pi'(x) = & \left[\left(\sum_{\theta<0} p_0(\theta) \right) \left(\sum_{\theta>0} p_0(\theta)(\theta + k \ln p_0(\theta)) \right) \right. \\ & \left. - \left(\sum_{\theta>0} p_0(\theta) \right) \left(\sum_{\theta<0} p_0(\theta)k \ln p_0(\theta) \right) \right. \\ & \left. + k \ln x \cdot \left(\sum_{\theta>0} p_0(\theta) \right) \left(\sum_{\theta<0} p_0(\theta) \right) \right] \left[x \sum_{\theta>0} p_0(\theta) + \sum_{\theta<0} p_0(\theta) \right]^{-2} \end{aligned}$$

As $\Pi'(x)$ limits to $-\infty$ as $x \rightarrow 0$ from the right, as it has at most a single root in x , given that $\Pi(0) \leq 0$ and $\Pi(1) > 0$, the payoff to becoming informed has a single root in $(0, 1)$, at a point I denote \bar{x} .

Set $\bar{N} = \frac{1}{\bar{x}}$. For (a), suppose there are $N < \bar{N}$ firms in the market. Then, from the proof of Theorem 3, for $\theta < 0$, $A_m(\theta) = 0$ and for $\theta > 0$, $A_m(\theta) = z_m$. Beliefs for firm n are given therefore given by (A.16) where $x \geq \frac{1}{\bar{N}}$ (the value taken when all other firms set $z_m = 1$). Therefore from $\frac{1}{N} > \frac{1}{\bar{N}} = \bar{x}$, we have $\Pi(\frac{1}{N}) > 0$ and thus firm n must necessarily profit.

For (b), suppose there are $N > \bar{N}$ firms getting positive payoff in an equilibrium. A firm making a profit uses a strategy with $z_m = 1$ (due to non-indifference) so that the equilibrium beliefs of firm n are given by (A.16) where $x \leq \frac{1}{N} < \frac{1}{\bar{N}} = \bar{x}$. Therefore, $\Pi(\frac{1}{N}) < 0$, so that firm n cannot get a positive profit, a contradiction. ■

A.5 Proof of Lemma 1

Without loss of generality, take a realization x of X_t , written as $x : \mathbb{R}_+ \rightarrow \Delta\Theta$ and denote its stopping time by τ . Supposing x is continuous, and since $\tau = \inf\{t | t \in \mathcal{S}(t)\}$, $x(0) = p_n \in \Delta\Theta \setminus \mathcal{S}(0)$ and \mathcal{S} is closed-graph, we have $\tau \neq 0$. As $\forall t, \mathcal{S}(0) \subseteq \mathcal{S}(t)$, the image under x of $[0, \tau)$ satisfies $x([0, \tau)) \subseteq \Delta\Theta \setminus \mathcal{S}(0)$. If x is continuous, the image under x of a set's closure is contained in the closure of its image; $\overline{\Delta\Theta \setminus \mathcal{S}(0)} \supseteq x(\overline{[0, \tau)}) = x([0, \tau])$. Thus, as the realized path x is almost surely continuous, $x_\tau \in \overline{\Delta\Theta \setminus \mathcal{S}(0)}$ almost surely.

A.6 Proof of Theorem 6

Let $a_n : \Delta\Theta \rightarrow [0, 1]$ denote acceptance probability for each posterior once the process X has stopped. Let $\hat{A}_n^N : \Delta\Theta \rightarrow [0, 1]^{|\Theta|}$ be the function mapping initial beliefs p_n for firm n of N into an ex-ante vector of type-specific acceptance probabilities. Additionally use \hat{p}_n^N as given by (1.9), the belief function for player n of N , to map opponents' acceptance probabilities into beliefs for firm n . Write the N -firm acceptance probability function as $\hat{A}^N : (\Delta\Theta)^N \rightarrow [0, 1]^{|\Theta|^N}$ and the N -firm belief function as $\hat{p}^N : [0, 1]^{|\Theta|^N} \rightarrow (\Delta\Theta)^N$; an equilibrium for N firms then corresponds to each fixed point of $\hat{A}^N \circ \hat{p}^N$.

For each firm n , if $p_n \in \Delta\Theta \setminus \mathcal{S}(0)$, write as q_n the average over X_t 's paths posterior distribution for rejected applicants. As from Lemma (1) each posterior is almost surely in $\overline{\Delta\Theta \setminus \mathcal{S}(0)}$ and that set is convex as the closure of a convex set,

$$q_n \in \overline{\Delta\Theta \setminus \mathcal{S}(0)}.$$

The proof will proceed with induction on the number of firms. With 1 firm in the market, its beliefs are $p_1 = p_0 \in \Delta\Theta \setminus \mathcal{S}(0)$ so the firm will get a positive payoff. Assume for induction that every equilibrium with N firms gives positive payoff to every firm; that is, $\forall n \leq N, p_n \in \Delta\Theta \setminus \mathcal{S}(0)$ or A_n identically 1.

Assume for contradiction that in some equilibrium described by $(A_n)_{n \leq N+1}$ of the market with $N + 1$ firms, firm $N + 1$ gets weakly negative expected payoff. By equilibrium, $\hat{A}^{N+1} \circ \hat{p}^{N+1}((A_n)_{n \leq N+1}) = (A_n)_{n \leq N+1}$. As A_{N+1} is identically 0, $\hat{p}^{N+1}((A_n)_{n \leq N+1}) = \hat{p}^{N+1}((A_n)_{n \leq N}, \vec{0})$. Using (1.9), we have $\hat{p}^{N+1}((A_n)_{n \leq N}, \vec{0}) = (\hat{p}^N(A_n^N)_{n \leq N}, p_{N+1})$ for some p_{N+1} ; so $(A_n)_{n \leq N}$ is a fixed point of $\hat{A}^N \circ \hat{p}^N$ and therefore defines beliefs in an N firm equilibrium.

From the inductive hypothesis, in each such N -firm equilibrium, each firm $n \leq N$ is getting a strictly positive expected payoff. Thus for each firm n either A_n is identically 1 or $p_n \in \Delta\Theta \setminus \mathcal{S}(0)$ and therefore Lemma (1) applies to each $n \leq N$. Then, firm $N + 1$ is either receiving an applicant last rejected by firm n where A_n is not identically 1, or an applicant that has never visited another firm. Therefore

$$p_{N+1} = \sum_{n \leq N: \bar{A}_n \neq 1} \phi_n q_n(\theta) + \left(1 - \sum_{n \leq N: \bar{A}_n \neq 1} \phi_n\right) p_0(\theta) \quad (\text{A.18})$$

for some probability weights ϕ , where $\sum_{n \leq N: \bar{A}_n \neq 1} \phi_n \leq N/(N + 1)$. As $p_0 \notin \mathcal{S}(0)$ and $\mathcal{S}(0)$ closed, $p_0 \in \text{int}(\Delta\Theta \setminus \mathcal{S}(0))$.

But as Lemma (1) gives that $q_n \in \overline{\Delta\Theta \setminus \mathcal{S}(0)}$, as $p_0 \in \text{int}(\Delta\Theta \setminus \mathcal{S}(0))$ and $\Delta\Theta \setminus \mathcal{S}(0)$ is convex,

$$p_{N+1} \in \Delta\Theta \setminus \mathcal{S}(0) \quad (\text{A.19})$$

and therefore firm $N + 1$ gets a strictly positive payoff, a contradiction. Thus, by PMI, for any $N \in \mathbb{N}$, in every equilibrium of the market with N firms, each firm gets

strictly positive profit, proving the theorem. ■

Appendix B

Proofs of results in Chapter 2

B.1 Proof of Lemma 2

Consider an equilibrium match that has just been revealed to be good at t . For revelation to have just occurred, the currently active offer involves monitoring.

If renegotiation occurs as per case 2 in Section 2.3.4 the proposer will receive $q/(r(1 + e^{-r\Delta}))$. The payoff to triggering renegotiations is obtained by discounting this by $e^{-r\Delta}$.

Assume that renegotiation never occurs in equilibrium; then the current monitoring offer persists forever, yielding a total surplus of $\frac{q-b}{r}$. Assuming that neither player wants to reopen negotiations, if the current wage in place is w , we must have that

$$\min\left\{\frac{w}{r}, \frac{q-b-w}{r}\right\} \geq \frac{1}{1+e^{-r\Delta}} \frac{q}{r}$$

For any current wage w , the greatest $\min\{w/r, (q-b-w)/r\}$ can be is $(q-b)/(2r)$; thus for renegotiation to never occur we require that $(q-b)/(2r) \geq q/(r(1 + e^{-r\Delta})) \Leftrightarrow (1 + e^{-r\Delta})b < (1 - e^{-r\Delta})q$, which is ruled out by (C2).

But as one's opponent reopening negotiations gives the receiver's share of the new bargain,

$e^{-r\Delta}qe^{-r\Delta}/(1 + e^{-r\Delta})$, it becomes even harder to satisfy the requirement to not renegotiate instantly with probability 1 if one's opponent may trigger renegotiation; therefore both instantly triggering renegotiation is the only equilibrium. As this means that each player has a probability 1/2 of being first proposer following revelation, each player at the instant of revelation has an expected payoff of

$$\frac{1}{2} \cdot \frac{qe^{-r\Delta}}{1 + e^{-r\Delta}} + \frac{1}{2} \cdot e^{-r\Delta} \frac{qe^{-r\Delta}}{1 + e^{-r\Delta}} = \frac{qe^{-r\Delta}}{2}. \quad \square$$

B.2 Proof of Lemma 3

Define the quantities

ξ	Flow mass of workers born per unit time
A	Mass of unemployed black type α workers
B	Mass of unemployed black type β workers
Λ	Mass of currently monitored black type β workers

As g is the fraction of new workers that is type α and unemployed α workers are becoming employed each at a Poisson rate μ and never separate, A obeys

$$\frac{dA}{dt} = \xi g - \mu A$$

Similarly, a proportion $(1 - g)$ of new workers is type β and such unemployed workers are also being hired at a Poisson rate μ each. However, as Λ workers who are of type β are being monitored, a flow mass $\Lambda\lambda(1 - \beta)$ of black β workers are separating after monitoring reveals a bad match are also coming in to the black unemployed pool. Hence, B obeys

$$\frac{dB}{dt} = \xi(1 - g) - \mu B + \Lambda\lambda(1 - \beta)$$

Finally, unemployed β workers are becoming employed with monitoring at a Poisson rate μ and once they are employed they cease being monitored when match quality is revealed, which occurs at a rate λ . Thus the mass of monitored black β workers Λ must satisfy

$$\frac{d\Lambda}{dt} = \mu B - \Lambda \lambda$$

Steady state implies that

$$\frac{dA}{dt} = \frac{dB}{dt} = \frac{d\Lambda}{dt} = 0$$

Solving, we obtain

$$A = \frac{\xi g}{\mu}$$

$$B = \frac{\xi(1-g)}{\mu\beta}$$

and therefore the proportion of α workers in the unemployed pool is

$$\frac{A}{A+B} = \frac{\frac{\xi g}{\mu}}{\frac{\xi g}{\mu} + \frac{\xi(1-g)}{\mu\beta}} = \frac{g}{g + \frac{1}{\beta}(1-g)}.$$

Thus, a new match from the black job-seeker pool is of average quality

$$\frac{g}{g + \frac{1}{\beta}(1-g)} \cdot 1 + \left(1 - \frac{g}{g + \frac{1}{\beta}(1-g)}\right) \cdot \beta = \frac{\beta}{\beta g + (1-g)} \equiv \theta_B$$

As $\beta < 1$ this is less than θ_W . \square

B.3 Proof of Proposition 1

The equilibrium wage proposed is

$$w_{N\theta_W}^{work} = \frac{1}{1 + e^{-r\Delta}}(q - \lambda(1 - \theta_W)c)$$

if the worker proposes first and

$$w_{N\theta_W}^{firm} = \frac{e^{-r\Delta}}{1 + e^{-r\Delta}}(q - \lambda(1 - \theta_W)c)$$

if the firm proposes first. As there will be no revelation, these shares split the expected output (using firm beliefs) equally. This equilibrium is supported by firm beliefs that are invariant to all contingencies before revelation.

At a pre-revelation history where an off-path offer is on the table, or where one is already in place, it is accepted/not renegotiated if it does not involve monitoring and the wage w satisfies

$$w_{N\theta_W}^{work} \geq w \geq w_{N\theta_W}^{firm}$$

If this condition does not hold at the off-path history in question, then, if in negotiations, the current proposer plays the equilibrium offer; otherwise, both players' strategy is to instantly reopen negotiations; and when they do, the equilibrium offer will be proposed.

At off-path histories where the match is revealed to be good, if the wage w satisfies $\frac{1}{1+e^{-r\Delta}}q \geq w \geq \frac{e^{-r\Delta}}{1+e^{-r\Delta}}q$ it stays in place; otherwise, play proceeds as in Lemma 2, granting an expected $qe^{-r\Delta}/(2r)$ to each party.

As discussed in Section 2.3.4, off-path histories that led to the revelation of a bad match lead to termination of the match.

A party who deviates before revelation can at most, therefore, transition from the receiver's share to the proposer's share of the match surplus, as one's opponent's strategy will not accept worse offers. Doing so, however, occasions a single delay, which discounts the payoff from such a deviation to exactly the receiver's payoff, which is the least the deviator could have started with. Therefore, there is no deviation that will strictly increase the agents' payoff and the strategies described are mutual best responses.

To show that there is no S3-type deviation that proposes monitoring, it remains to show that an α worker or a firm cannot propose a mutually beneficial monitoring regime.

Lemma 2 pins down continuation payoffs from being in a match revealed to be good. Thus, an α worker making an offer of w_u with monitoring in place yields to this worker, in the absence of renegotiation until match quality revelation,

$$V_{M\alpha}(w_u) = \frac{w_u + \lambda \frac{qe^{-r\Delta}}{2r}}{\lambda + r}. \quad (\text{B.1})$$

The requirement for S3 is stated in terms of beliefs remaining constant, in this case at θ_W . Thus, an employer accepting this offer expects a payoff of

$$F_{M\theta_W}(w_u) = \frac{q - b - (1 - \theta_W)\lambda c - w_u + \lambda\theta_W \frac{qe^{-r\Delta}}{2r}}{\lambda + r}. \quad (\text{B.2})$$

Summing (B.1) and (B.2), we get

$$\frac{q - b - (1 - \theta_W)\lambda c + \lambda(1 + \theta_W) \frac{qe^{-r\Delta}}{2r}}{\lambda + r} \quad (\text{B.3})$$

For such a deviation to violate S3 necessarily (B.3) has to be greater than equilibrium payoffs; this can only be the case if

$$\frac{b}{\lambda} < (1 - \theta_W) \frac{\lambda c}{r} - \frac{q(2 - e^{-r\Delta}(1 + \theta_W))}{2r} \quad (\text{B.4})$$

which is precluded by assumption (C4).

In the limit as $\Delta \downarrow 0$, the equilibrium shares of the first proposer and receiver equalize; the limiting wage is $w_{N\theta_W} = .5q - .5(1 - \theta_W)\lambda c$.

B.4 Proof of Proposition 2

Consider strategies for the firm and each type of worker that may in principle involve renegotiation and different monitoring for each type of worker. Call \tilde{m}_α and \tilde{m}_β the expected discount at the point of match quality revelation for α and β workers given these strategies and let \tilde{W}_α and \tilde{W}_β be the total expected discounted wages before such revelation occurs.

Equilibrium requires incentive compatibility: each worker weakly prefers the strategy they actually adopt to the other type's; hence, α 's IC requires

$$\tilde{V}_\alpha = \tilde{W}_\alpha + \tilde{m}_\alpha \frac{qe^{-r\Delta}}{2r} \geq \tilde{W}_\beta + \tilde{m}_\beta \frac{qe^{-r\Delta}}{2r} \quad (\text{B.5})$$

and the β 's IC imposes

$$\tilde{V}_\beta = \tilde{W}_\beta + \tilde{m}_\beta \left(\beta \frac{qe^{-r\Delta}}{2r} + (1 - \beta)U_\beta \right) \geq \tilde{W}_\alpha + \tilde{m}_\alpha \left(\beta \frac{qe^{-r\Delta}}{2r} + (1 - \beta)U_\beta \right) \quad (\text{B.6})$$

Combining the two and rearranging terms

$$\begin{aligned} (\tilde{m}_\beta - \tilde{m}_\alpha)(1 - \beta) \left(U_\beta - \frac{qe^{-r\Delta}}{2r} \right) &\geq \\ \left(\tilde{W}_\alpha + \tilde{m}_\alpha \frac{qe^{-r\Delta}}{2r} \right) - \left(\tilde{W}_\beta + \tilde{m}_\beta \frac{qe^{-r\Delta}}{2r} \right) &\geq 0 \end{aligned} \quad (\text{B.7})$$

Since from (C3) $U_\beta < qe^{-r\Delta}/2r$, we have that $\tilde{m}_\alpha \geq \tilde{m}_\beta$.

The firm's payoff in any such equilibrium given $(\tilde{W}_\alpha, \tilde{m}_\alpha, \tilde{W}_\beta, \tilde{m}_\beta)$ is ¹

¹Here it is relevant to point out that if the expected discount on revelation is m , then the expected cost of monitoring until revelation is mb/λ . This is for the following reason: to say the expected discount on revelation is m is to say that for probability $M(t)$ of monitoring at time t conditional on no revelation by t , $\int_0^\infty e^{-rt} e^{-\int_0^t \lambda M(t') dt'} \lambda M(t) dt = m$; but the amount spent on monitoring at each time is the probability no revelation occurs until that time multiplied by the probability monitoring occurs at that time, the discount and the flow cost of monitoring; hence the total cost is $\int_0^\infty e^{-rt} e^{-\int_0^t \lambda M(t') dt'} bM(t) dt = mb/\lambda$.

$$\begin{aligned} \tilde{F} = & g\left(\frac{q}{r}(1 - \tilde{m}_\alpha) - \tilde{W}_\alpha + \tilde{m}_\alpha \frac{qe^{-r\Delta}}{2r} - \tilde{m}_\alpha \frac{b}{\lambda}\right) + \\ & (1 - g)\left(\frac{q - (1 - \beta)\lambda c}{r}(1 - \tilde{m}_\beta) - \tilde{W}_\beta + \tilde{m}_\beta \beta \frac{qe^{-r\Delta}}{2r} - \tilde{m}_\beta \frac{b}{\lambda}\right) \end{aligned} \quad (\text{B.8})$$

S3 requires that deviating first offers cannot be made that improve the payoff of an α worker and the firm. As the wage can transfer surplus freely, this implies that the sum of candidate equilibrium payoffs are weakly greater than the sum of those feasible by a no-monitoring offer²:

$$\tilde{V}_\alpha + \tilde{F} \geq \frac{q - (1 - \theta_W)\lambda c}{r} \quad (\text{B.9})$$

Expanding,

$$\begin{aligned} \tilde{V}_\alpha + \tilde{F} = & (1 - g)(\tilde{W}_\alpha - \tilde{W}_\beta) + \tilde{m}_\alpha((1 + g)qe^{-r\Delta} - \frac{gb}{\lambda} - \frac{gq}{r}) \\ & + \tilde{m}_\beta(1 - g)\left(-\frac{q - (1 - \beta)\lambda c}{r} + \frac{\beta qe^{-r\Delta}}{2r} - \frac{b}{\lambda}\right) \end{aligned} \quad (\text{B.10})$$

Retrieving

$$\tilde{W}_\alpha - \tilde{W}_\beta \leq (\tilde{m}_\beta - \tilde{m}_\alpha)\left(\beta \frac{qe^{-r\Delta}}{2r} + (1 - \beta)U_\beta\right)$$

by rearranging (B.6), we substitute it into (B.10) we arrive at an expression weakly greater than the LHS of (B.9) where the coefficient of \tilde{m}_β is

$$(1 - g)\frac{1}{r}\left[\lambda c(1 - \beta) - q(1 - e^{-r\Delta}\beta) + r(1 - \beta)U_\beta - \frac{rb}{\lambda}\right] \quad (\text{B.11})$$

which we know from (C6) is positive. Therefore, up to the constraint imposed by (B.5) we have an upper bound of the LHS of (B.9) increasing in \tilde{m}_β . So if (B.9) holds

²Notice that while the candidate equilibrium is allowed to generate value from screening worker types by strategies and therefore apply monitoring more efficiently, the deviations S3 checks against are not. That it turns out such deviations are enough to destroy all equilibria but one is a product of the β workers' incentives to not reveal themselves.

for some $(\tilde{m}_\alpha, \tilde{m}_\beta)$, it must hold for $(\tilde{m}_\alpha, \tilde{m}_\alpha)$. Making this substitution, we arrive at

$$\frac{q - (1 - \theta_W)\lambda c}{r}(1 - \tilde{m}_\alpha) + \tilde{m}_\alpha \frac{(1 + \theta_W)qe^{-r\Delta}}{2r} - \tilde{m}_\alpha \frac{b}{\lambda} \geq \frac{q - (1 - \theta_W)\lambda c}{r} \quad (\text{B.12})$$

which due to (C4) can only occur if $\tilde{m}_\alpha = 0$.

Therefore, regardless of the first proposer, all equilibria in the white labor market lack monitoring for both types of workers. We can further exclude equilibria with delay, as an S3-type deviation giving the receiver his equilibrium utility and the proposer taking the excess would be payoff-increasing in those cases.

Finally, by S3, no deviation by a first receiver that gives the first proposer his payoff when proposing, discounted, is gainful. Therefore, the first proposer's share cannot be greater than $\frac{1}{1+e^{-r\Delta}}$. Similarly, the initial proposer i cannot be getting $x < \frac{1}{1+e^{-r\Delta}}$, lest j have a deviating offer in his own role as first proposer giving i his discounted value, $e^{-r\Delta}x$ and j a share of $1 - xe^{-r\Delta} > \frac{1}{1+e^{-r\Delta}}$.

Thus, all equilibria of the white labor market reach immediate agreement with a no-monitoring offer; the wage splits the surplus along the Rubinstein shares and therefore the equilibrium of Proposition 1 is essentially (up to off-path behavior and beliefs) unique.

B.5 Proof of Proposition 3

The initial equilibrium wage proposed is

$$w_{M\theta_B}^{work} = [q - b - \lambda(1 - \theta_B) - (e^{-r\Delta} - \theta_B)qe^{-r\Delta}/(2r)] / (1 + e^{-r\Delta})$$

if the worker proposes first and

$$w_{M\theta_B}^{firm} = [e^{-r\Delta}(q - b - \lambda(1 - \theta_B)) - (1 - e^{-r\Delta}\theta_B)qe^{-r\Delta}/(2r)] / (1 + e^{-r\Delta})$$

if the firm proposes first. Monitoring is in use until revelation, and no renegotiation takes place until then.

If the worker rejects a firm monitoring with a wage in $[w_{M\theta_B}^{firm}, w_{M\theta_B}^{work}]$, opens renegotiation when a monitoring regime with a wage in that interval is in place, or makes or accepts a non-monitoring offer before revelation, the firm immediately believes the worker to be type β . This change is irreversible. Otherwise, the firm has beliefs constant at θ_B .

When the firm starts believing the worker to be type β , both parties immediately renegotiate to the equilibrium in Section 2.3.4.4 with monitoring and a lower wage.

As long as beliefs are θ_B , agents in the role of proposer offer their $w_{M\theta_B}^i$. Monitoring offers with wages in $[w_{M\theta_B}^{firm}, w_{M\theta_B}^{work}]$ wages are accepted by either party without renegotiation until revelation. Other offers are rejected or renegotiated, and the next offer is the proposer's $w_{M\theta_B}^i$.

If revelation occurs, bad matches separate; good matches renegotiate as per Lemma 1.

Clearly, workers don't want to deviate to propose in $(w_{M\theta_B}^{firm}, w_{M\theta_B}^{work})$ as they will lead to acceptance but a lower payoff; also, they don't propose wages outside $[w_{M\theta_B}^{firm}, w_{M\theta_B}^{work}]$ as the firm will reject and propose $w_{M\theta_B}^{firm}$ in addition to suffering the delay. Thus, always proposing $w_{M\theta_B}^{work}$ is optimal. Workers won't reject offers in $[w_{M\theta_B}^{firm}, w_{M\theta_B}^{work}]$ or accept or propose non-monitoring offers as they don't want to be treated as β s as per (C6) and (C7).

Firms know that by the workers' strategy, the highest offer they can get accepted is $w_{M\theta_B}^{firm}$ and that higher ones, or ones below $w_{M\theta_B}^{work}$, will be rejected and that the worker will counter-offer $w_{M\theta_B}^{work}$, in addition to the firm suffering a delay. Firm offers in $(w_{M\theta_B}^{firm}, w_{M\theta_B}^{work}]$ will be accepted but yield a lower payoff than $w_{M\theta_B}^{firm}$; thus always proposing $w_{M\theta_B}^{firm}$ is optimal for the firm. Given this, the firm accepts offers in

$[w_{M\theta_B}^{firm}, w_{M\theta_B}^{work}]$; but will reject higher ones because it can do better as proposer, and lower ones because it knows renegotiation will be imminent once they are in place. If production is occurring, the firm can gain by renegotiating if either (a) the worker will instantly renegotiate, and the firm's first offer here is a lower wage than the worker's (so as above, if $w < w_{M\theta_B}^{firm}$), or if the firm's payoff from making its offer, $w_{M\theta_B}^{firm}$, with delay, is preferable to the current payoff; but that is precisely when the current wage $w > w_{M\theta_B}^{firm}$.

That there is no S3-type deviation that proposes no monitoring follows from (C4).

In the limit as $\Delta \downarrow 0$, the equilibrium shares of the first proposer and receiver equalize; the limiting wage is $w_{M\theta_B} = \frac{1}{2} [q - b - \lambda c(1 - \theta_B)] - \frac{(1 - \theta_B)}{2} \lambda \frac{q}{2r}$.

B.6 Proof of Proposition 4

The proof here proceeds in the same fashion as that in B.4. Instead of comparing to a no-monitoring deviating offer we compare to a monitoring offer; therefore instead of B.12 we have

$$\frac{q - (1 - \theta_B)\lambda c}{r}(1 - \tilde{m}) + \tilde{m} \frac{(1 + \theta_B)qe^{-r\Delta}}{2r} - \tilde{m} \frac{b}{\lambda} \geq \frac{q - (1 - \theta_B)\lambda c - b + \frac{(1 + \theta_B)qe^{-r\Delta}}{2r}}{\lambda + r} \quad (\text{B.13})$$

which due to (C5) can only be true if $\tilde{m} \geq \frac{\lambda}{\lambda + r}$; but as \tilde{m} is an expected discount of a variable that at most arrives as a Poisson with rate λ , this constitutes an upper bound to \tilde{m} and corresponds to full monitoring and no delay.

Therefore, only fully monitoring equilibria exist in the black labor market. Within such candidate equilibria, S3 would allow for deviation from any initial offer not corresponding to that in Proposition 3, therefore that equilibrium is unique.

Appendix C

Proofs of results in Chapter 3

C.1 Proof of Theorem 7

Proof. Let σ be an SPNE of $G^\infty(\delta)$ with equilibrium payoff vector $u(\sigma) = \mathbf{v} \notin \overline{F^{res}}$. From (3.2) we have a player i such that $v_i < r_i$. Let $\bar{\sigma}_i$ be a strategy for player i that myopically best responds to σ_{-i} in every period, that is,

$$\bar{\sigma}_i(h^{t-1}) \in BR_i(\sigma_{-i}(h^{t-1})) = \arg \max_{a_i \in A_i} g_i(a_i, \sigma_{-i}(h^{t-1})).$$

Since σ was assumed to be an SPNE, the optimality of σ_i implies that the myopic best response strategy must also give player i less than r_i against the opponents' strategies: $u_i(\bar{\sigma}_i, \sigma_{-i}) \leq u_i(\sigma_i, \sigma_{-i}) < r_i$. Therefore, by definition of r_i in (3.1) it follows that there is a period t such that the myopic best response gives player i less than r_i :

$$\exists h^{t-1} \in H^{t-1} : \max_{a_i \in A_i} g_i(a_i, \sigma_{-i}(h^{t-1})) < r_i = \min_{\alpha \in Q} \max_{a_i^* \in A_i} g_i(\alpha_i^*, \alpha_{-i}) \quad (\text{C.1})$$

This immediately means that there is no α_i that gives $(\alpha_i, \sigma_{-i}(h^{t-1}))$ the myopic indifference property. In other words, there is some player j who is not myopically indifferent at h^{t-1} , i.e. for some pair of actions

$$\exists a_j, a'_j \in \text{supp}(\sigma_j(h^{t-1})) \text{ such that } g_j(a_j, \sigma(h^{t-1})) \neq g_j(a'_j, \sigma(h^{t-1})). \quad (\text{C.2})$$

There are two strategies for i starting at h^{t-1} , denoted $\sigma_j|_{h^{t-1}}$ and $\sigma'_j|_{h^{t-1}}$ such that one plays a_j and the other a'_j at the starting history h^{t-1} ; since j plays both a_j and

a'_j with positive probability, we must have

$$u_j(\sigma_j |_{h^{t-1}}, \sigma_{-j} |_{h^{t-1}}) = u_j(\sigma'_j |_{h^{t-1}}, \sigma_{-j} |_{h^{t-1}}). \quad (\text{C.3})$$

Thus their induced player- j expected stage-game payoff sequences $g_j^{(\tau)}$ and $g'_j^{(\tau)}$ satisfy

$$(1 - \delta_j) \sum_{\tau \geq t} \delta_j^{\tau-t} g_j^{(\tau)} = (1 - \delta_j) \sum_{\tau \geq t} \delta_j^{\tau-t} g'_j^{(\tau)} \quad (\text{C.4})$$

and therefore

$$\sum_{\tau \geq t} \delta_j^{\tau-t} (g_j^{(\tau)} - g'_j^{(\tau)}) = 0. \quad (\text{C.5})$$

If σ is an equilibrium strategy profile for a discount factor δ' it must be the case that

$$\sum_{\tau \geq t} \delta_j'^{\tau-t} (g_j^{(\tau)} - g'_j^{(\tau)}) = 0. \quad (\text{C.6})$$

Regarded as a function of δ'_j ,

$$f(\delta'_j) \equiv \sum_{\tau \geq t} \delta_j'^{\tau-t} (g_j^{(\tau)} - g'_j^{(\tau)}) \quad (\text{C.7})$$

is a power series with bounded coefficients (as F is compact). As such, it converges absolutely for $\delta'_j \in (0, 1)$. A well-known result from complex analysis asserts that a power series either

1. has finitely many roots in any compact subset of the interior of its disk of convergence; or
2. has all coefficients equal to 0.¹

As (C.2) implies that $g_j^t \neq g'_j^t$, the coefficient on $\delta_j'^0$ is not 0, ruling out (2). But then (1) is true, implying that there are only finitely many roots of f in $[\frac{1}{2}\delta_j, \frac{1}{2}(1 + \delta_j)]$ and so (C.6) will not hold for all points of any neighborhood of δ_j . Therefore there can be no neighborhood of δ on which σ is an SPNE, so that σ cannot be a DRSP equilibrium strategy profile. \square

¹This is a corollary of the fact that a non-constant analytic function's zeros are isolated (the set of roots is totally disconnected). A function defined by a power series is analytic on the interior of its radius of convergence.

C.2 Proof of Theorem 8

Proof. Fix $\mathbf{v} \in F^{res} \setminus \underline{\partial}F^{res}$, where $\underline{\partial}F^{res}$ is the lower boundary of the set.² Pick $\mathbf{v}' \in F^*$ and $\epsilon > 0$ s.t.

$$\forall i, \quad r_i < v'_i < v'_i + \epsilon < v_i. \quad (\text{C.8})$$

For each i define a vector \mathbf{v}^i giving each $j \neq i$ a ‘reward’ of ϵ :

$$v_j^i = \begin{cases} v'_i & j = i \\ v'_j + \epsilon & \forall j \neq i. \end{cases}$$

Let $\boldsymbol{\alpha}^i \in Q$ be such that player i could earn her restricted minmax if she were to play a best response to the action of the other players; that is,

$$\boldsymbol{\alpha}^i \in \arg \min_{\boldsymbol{\alpha} \in Q} \max_{a_i} g_i(a_i, \boldsymbol{\alpha}_{-i}) \quad (\text{C.9})$$

Choose $N \in \mathbb{N}$ s.t.

$$\forall i, \quad \max_{\mathbf{a} \in A} g_i(\mathbf{a}) + N g_i(\boldsymbol{\alpha}^i) < \min_{\mathbf{a} \in A} g_i(\mathbf{a}) + N v'_i \quad (\text{C.10})$$

which is possible by inequality (C.8).

Strategies.

We construct a Simple Strategy Profile *a la* Abreu (1988): If j deviates at any point unilaterally then impose *Phase II(j)* followed by *Phase III(j)*.

Phase I: In each period play the correlated action $p \in \Delta A$ such that $\mathbf{v} = \sum_{\mathbf{a} \in A} p(\mathbf{a})g(\mathbf{a})$.

Phase II(i): Play $\boldsymbol{\alpha}^i$ for N periods.

Phase III(i): Play $p^i \in \Delta A$ such that $\mathbf{v}^i = \sum_{\mathbf{a} \in A} p^i(\mathbf{a})g(\mathbf{a})$ at each period.

Note that the strategy, including parameters in it, is set independently of the discount factor(s).³ Also note that we do not ask a player to best respond during her own punishment phase, as that would not leave the others willing to mix.

Checking subgame perfection.

Step 1. Player i does not deviate from *Phase I* if

$$(1 - \delta) \max_{\mathbf{a} \in A} g_i(\mathbf{a}) + \delta[(1 - \delta^N)r_i + \delta^N v'_i] \leq (1 - \delta) \min_{\mathbf{a} \in A} g_i(\mathbf{a}) + \delta v_i. \quad (\text{C.11})$$

²The boundary condition is common to FM. Abreu, Dutta and Smith (1991) showed that exclusion of the lower boundary may be dispensed with.

³In contrast, the third phase in FM depends on the exact discount factor when mixing is involved.

As $\delta \rightarrow 1$, $LHS \rightarrow v'_i$ and $RHS \rightarrow v_i$. In other words inequality (C.11) holds for high δ because $v'_i < v_i$ by (C.8).

Step 2. Player i does not deviate from *Phase II(i)* if

$$\forall \tau \in \{1, \dots, N\}, \quad (1 - \delta^\tau)g_i(\boldsymbol{\alpha}^i) + \delta^\tau v'_i \geq (1 - \delta)r_i + \delta[(1 - \delta^N)g_i(\boldsymbol{\alpha}^i) + \delta^N v'_i],$$

which as $\delta \rightarrow 1$ reduces to $r_i \leq v_i$, which clearly holds with strict inequality; thus, following Abreu, we see that restarting the minmax phase suffices to deter a deviation by i .

Step 3. Player i does not deviate from *Phase III(i)* if

$$(1 - \delta) \max_{\mathbf{a} \in A} g_i(\mathbf{a}) + \delta[(1 - \delta^N)g_i(\boldsymbol{\alpha}^i) + \delta^N v'_i] \leq (1 - \delta) \min_{\mathbf{a} \in A} g_i(\mathbf{a}) + \delta v'_i.$$

As $\delta \rightarrow 1$, both RHS & LHS $\rightarrow v'_i$. Hence we rearrange so that $LHS - RHS =$

$$\begin{aligned} &= (1 - \delta) \left\{ \max_{\mathbf{a} \in A} g_i(\mathbf{a}) - \min_{\mathbf{a} \in A} g_i(\mathbf{a}) \right\} + \delta(1 - \delta^N)(g_i(\boldsymbol{\alpha}^i) - v'_i) \\ &= (1 - \delta) \left\{ \max_{\mathbf{a} \in A} g_i(\mathbf{a}) - \min_{\mathbf{a} \in A} g_i(\mathbf{a}) \right\} + \delta(1 + \delta + \dots + \delta^{N-1})(1 - \delta)(g_i(\boldsymbol{\alpha}^i) - v'_i) \\ &= (1 - \delta) \left[\max_{\mathbf{a} \in A} g_i(\mathbf{a}) - \min_{\mathbf{a} \in A} g_i(\mathbf{a}) + (\delta + \dots + \delta^N)(g_i(\boldsymbol{\alpha}^i) - v'_i) \right]. \end{aligned}$$

As $\delta \rightarrow 1$, the term in square brackets tends to

$$\max_{\mathbf{a} \in A} g_i(\mathbf{a}) - \min_{\mathbf{a} \in A} g_i(\mathbf{a}) + N(g_i(\boldsymbol{\alpha}^i) - v'_i) = (\max_{\mathbf{a} \in A} g_i(\mathbf{a}) + N g_i(\boldsymbol{\alpha}^i)) - (\min_{\mathbf{a} \in A} g_i(\mathbf{a}) + N v'_i),$$

which is negative by inequality (C.10), implying that *Phase III(i)* is an equilibrium strategy phase for large enough δ .

Step 4. Player i does not deviate (observably) from *Phase II(j)* or *Phase III(j)* because as $\delta \rightarrow 1$ deviation payoff $\rightarrow v'_i$, whereas the equilibrium payoff $\rightarrow v'_i + \varepsilon$.

Step 5. Although player i does not deviate observably we need to show that she mixes as required in *Phase II(j)*. Recall our definition of Q ; since $\boldsymbol{\alpha}^j \in Q$, mixing only occurs between myopically indifferent actions according to $\boldsymbol{\alpha}^j$; as future play does not vary over i 's actions on $\text{supp}(\boldsymbol{\alpha}^j)$, she is indifferent. \square

Appendix D

References

- Abreu, Dilip. “On the theory of infinitely repeated games with discounting.” *Econometrica: Journal of the Econometric Society* (1988): 383-396.
- Abreu, Dilip, Prajit K. Dutta, and Lones Smith. “The folk theorem for repeated games: a NEU condition.” *Econometrica* 62.4 (1994): 939-948.
- Autor, David H., and David Scarborough. “Does Job Testing Harm Minority Workers? Evidence from Retail Establishments.” *Quarterly Journal of Economics* (2008): 219-277.
- Altonji, Joseph G., and Rebecca M. Blank. “Race and gender in the labor market.” *Handbook of labor economics* 3 (1999): 3143-3259.
- Becker, Gary S. *The economics of discrimination*. University of Chicago press, 2010.
- Bergemann, Dirk, Xianwen Shi, and Juuso Välimäki. “Information acquisition in interdependent value auctions.” *Journal of the European Economic Association* 7.1 (2009): 61-89.
- Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch. “A theory of fads, fashion, custom, and cultural change as informational cascades.” *Journal of Political Economy* (1992): 992-1026.

- Bjerk, David. "Glass ceilings or sticky floors? Statistical discrimination in a dynamic model of hiring and promotion." *The Economic Journal* 118.530 (2008): 961-982.
- Black, Dan A. "Discrimination in an equilibrium search model." *Journal of Labor Economics* 13.2 (1995): 309-334.
- Bowlus, Audra J., and Zvi Eckstein. 2002. "Discrimination and Skill Differences in an Equilibrium Search Model." *International Economic Review*, 43(2): 1309-1345.
- Bowlus, A. J., Kiefer, N. M. and Neumann, G. R. 2001. "Equilibrium Search Models and the Transition from School to Work." *International Economic Review*, 42(2), 317-343.
- Bowlus, Audra J., and Shannon N. Seitz. "Search friction in the US labor market: equilibrium estimates from the PSID." *Panel Data and Structural Labour Market Models*. Emerald Group Publishing Limited, 2000. 145-169.
- Bronars, Stephen G., and Melissa Famulari. "Wage, tenure, and wage growth variation within and across establishments." *Journal of Labor Economics* 15.2 (1997): 285-317.
- Charles, Kerwin Kofi, and Jonathan Guryan. "Prejudice and wages: an empirical assessment of Becker's The Economics of Discrimination." *Journal of Political Economy* 116.5 (2008): 773-809.
- Coate, Stephen, and Glenn C. Loury. "Will affirmative-action policies eliminate negative stereotypes?." *American Economic Review* (1993): 1220-1240.
- Coates, Ta-Nehisi. "Fear of a Black President." *The Atlantic*. 2012.

www.theatlantic.com/magazine/archive/2012/09

[/fear-of-a-black-president/309064](http://www.theatlantic.com/magazine/archive/2012/09/fear-of-a-black-president/309064) . Accessed January 12th 2016.

D' Amico, Ronald, and Nan L. Maxwell. "The Impact of Post-School Joblessness on Male Black-White Wage Differentials." *Industrial Relations: A Journal of Economy and Society* 33.2 (1994): 184-205.

Dasgupta, Aniruddha, and Sambuddha Ghosh. *Folk Theorems without Public Randomization: A New Approach*. Working Paper.

Denti, Tommaso. *Unrestricted Information Acquisition*. Working Paper.

Fairlie, Robert W., and William A. Sundstrom. "The emergence, persistence, and recent widening of the racial unemployment gap." *International Labor Relations Review* 52.2 (1999): 252-270.

Friedman, James W. "A non-cooperative equilibrium for supergames." *Review of Economic Studies* 38.1 (1971): 1-12.

Fryer, Roland G. "Belief flipping in a dynamic model of statistical discrimination." *Journal of Public Economics* 91.5 (2007): 1151-1166.

Fudenberg, Drew, and Eric Maskin. "The folk theorem in repeated games with discounting or with incomplete information." *Econometrica: Journal of the Econometric Society* (1986): 533-554.

Fudenberg, Drew, Philipp Strack, and Tomasz Strzalecki. *Stochastic choice and optimal sequential sampling*. Working Paper.

Geanakoplos, John, David Pearce, and Ennio Stacchetti. "Psychological games and sequential rationality." *Games and Economic Behavior* 1.1 (1989): 60-79.

- Gentzkow, Matthew, and Emir Kamenica. "Costly persuasion." *American Economic Review* 104.5 (2014): 457-462.
- Gossner, Olivier. "The folk theorem for finitely repeated games with mixed strategies." *International Journal of Game Theory* 24, no. 1 (1995): 95-107.
- Guardian. "Prince: Black People Don't Get Second Chances," Feb. 25 2014, <http://www.theguardian.com/music/2014/feb/25/prince-black-people-dont-get-second-chances> .Accessed January 24, 2016.
- Hall, Robert E., and Paul R. Milgrom. "The limited influence of unemployment on the wage bargain." *American Economic Review* 98.4 (2008): 1653-1674.
- Hébert, Benjamin, and Michael Woodford. *Rational Inattention with Sequential Information Sampling*. Working Paper.
- Inderst, Roman. "Matching markets with adverse selection." *Journal of Economic Theory* 121.2 (2005): 145-166.
- Kalai, Ehud, and William Stanford. "Finite rationality and interpersonal complexity in repeated games." *Econometrica: Journal of the Econometric Society* (1988): 397-410.
- Lang, Kevin. "A language theory of discrimination." *The Quarterly Journal of Economics* 101.2 (1986): 363-382.
- Lang, Kevin, Jee Lehmann, and K. Yeon. "Racial discrimination in the labor market: Theory and empirics." *Journal of Economic Literature* 50.4 (2012): 959-1006.
- Lang, Kevin, and Michael Manove. "Wage announcements with a continuum of worker types." *Annales d'Economie et de Statistique* (2003): 223-244.

- Lang, Kevin, and Michael Manove. "Education and labor market discrimination." *American Economic Review* 101.4 (2011): 1467-1496.
- Lang, Kevin, Michael Manove, and William T. Dickens. "Racial discrimination in labor markets with posted wage offers." *American Economic Review* 95.4 (2005): 1327-1340.
- Lauermann, Stephan, and Asher Wolinsky. "Search with adverse selection." *Econometrica* 84.1 (2016): 243-315.
- Lundberg, Shelly J., and Richard Startz. "Private discrimination and social intervention in competitive labor market." *American Economic Review* 73.3 (1983): 340-347.
- Mabry, Marcus. "Twice as good: Condoleezza Rice and her path to power." Rodale, 2008.
- MacLeod, William Bentley. "Optimal contracting with subjective evaluation." *American Economic Review* 93.1 (2003): 216-240.
- Matejka, Filip, and Alisdair McKay. "Rational inattention to discrete choices: A new foundation for the multinomial logit model." *American Economic Review* 105.1 (2014): 272-298.
- Moro, Andrea, and Peter Norman. "A general equilibrium model of statistical discrimination." *Journal of Economic Theory* 114.1 (2004): 1-30.
- Mortensen, Dale T., and Christopher A. Pissarides. "Job creation and job destruction in the theory of unemployment." *Review of Economic Studies* 61.3 (1994): 397-415.

- Neal, Derek A., and William R. Johnson. "The role of premarket factors in black-white wage differences." *Journal of Political Economy* 104.5 (1996): 869-895.
APA
- Persico, Nicola. "Information acquisition in auctions." *Econometrica* 68.1 (2000): 135-148.
- Peski, Marcin, and Balázs Szentes. "Spontaneous discrimination." *American Economic Review* 103.6 (2013): 2412-2436.
- Ravid, Doron. *Bargaining with rational inattention*. Working paper.
- Reid, Alvin. "Black Head Coaches Don't Get Second Chances," *St. Louis American*, 31 Aug. 2015, <http://newpittsburghcourieronline.com/2015/08/31/black-head-coaches-dont-get-second-chances> . Accessed January 24, 2016.
- Rosén, Åsa. "An equilibrium search-matching model of discrimination." *European Economic Review* 41.8 (1997): 1589-1613.
- Rubinstein, Ariel. "Perfect equilibrium in a bargaining model." *Econometrica: Journal of the Econometric Society* (1982): 97-109.
- Shannon, C. E. "A Mathematical Theory of Communication." *Bell System Technical Journal* 27 (1948): 379-423.
- Sims, Christopher A. "Implications of rational inattention." *Journal of Monetary Economics* 50.3 (2003): 665-690.
- Spencer, DeShuna. 2014. "For Black Youth in America, There Are No Second Chances," *emPower*, 5 Feb. 2014, <http://www.empowermagazine.com/black-youth-america-second-chances> . Accessed January 24, 2016.

Woodford, Michael. *Inattention as a source of randomized discrete adjustment*.
Working Paper.

Wozniak, Abigail. “Discrimination and the effects of drug testing on black employment.” *Review of Economics and Statistics* 97.3 (2015): 548-566.

Yang, Ming. “Coordination with flexible information acquisition.” *Journal of Economic Theory* 158 (2015): 721-738.

Zhu, Haoxiang. “Finding a good price in opaque over-the-counter markets.” *Review of Financial Studies* 25.4 (2012): 1255-1285.

CURRICULUM VITAE

