

UCLA

UCLA Electronic Theses and Dissertations

Title

Essays in Information Economics

Permalink

<https://escholarship.org/uc/item/48d293p0>

Author

Aleksenko, Stepan

Publication Date

2024

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA
Los Angeles

Essays in Information Economics

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Economics

by

Stepan Aleksenko

2024

© Copyright by
Stepan Aleksenko
2024

ABSTRACT OF THE DISSERTATION

Essays in Information Economics

by

Stepan Aleksenko

Doctor of Philosophy in Economics

University of California, Los Angeles, 2024

Professor Moritz Meyer-ter-Vehn, Chair

This dissertation comprises three chapters that study information and learning aspects in various game-theoretical models.

The first chapter studies how firms manage their reputation for quality via price-dependent consumer reviews. Pricing decisions are crucial for managing a firm's reputation and maximizing profits. Consumer reviews reflect both the product quality and its price, with more favorable reviews being left when a product is priced lower. We study whether such review behavior can induce a firm to manipulate the review process by underpricing its product, or pricing it below current consumers' willingness to pay. We introduce an equilibrium model with a privately informed firm repeatedly selling its product to uninformed but rational consumers who learn about the quality of the product from past reviews and current prices. We show that underpricing can arise only when the firm reputation is low and then only under a specific condition on consumers' taste shock distribution, which we fully characterize. Rating manipulation unambiguously benefits consumers, because it operates via underpricing.

The second chapter studies how delegated recruitment shapes talent selection. Firms typically pay recruiters via refund contracts, which specify a payment upon the hire of a suggested candidate and a refund if a candidate is hired but terminated during an initial period of employment. We develop a model where refund contracts naturally arise and

show that delegation leads to statistical discrimination, where the recruiter favors candidates with more precise productivity information. This is misaligned with direct hiring, where the firm has option value from uncertain candidates. Under tractable parametric assumptions, we characterize the unique equilibrium in which candidates with lower expected productivity but more informative signals (“safe bets”) are hired over candidates with higher expected productivity but less informative signals (“diamonds in the rough”).

The third chapter studies the efficiency of information aggregation in the DeGroot learning model. We introduce a social planner in the DeGroot model who aims to improve the time asymptotic information aggregation in finite observational networks. We show that in any connected network, it is possible to achieve the best information aggregation by reassigning the attention individuals pay to each others’ opinions. We provide an algorithm that constructs a solution to this problem. We also identify the necessary and sufficient condition on the network for achieving the best information aggregation in the average-based updating learning model for homogeneous private signals. Finally, we demonstrate an approach to increasing the speed of learning.

The dissertation of Stepan Aleksenko is approved.

Sushil Bikhchandani

Alexander W. Bloedel

Simon Adrian Board

William R. Zame

Moritz Meyer-ter-Vehn, Committee Chair

University of California, Los Angeles

2024

TABLE OF CONTENTS

Acknowledgments	viii
Vita	ix
1 Reputational Underpricing	1
1.1 Introduction	1
1.2 Literature Review	4
1.3 Model	7
1.3.1 Model Discussion	9
1.3.2 Equilibrium Concept	10
1.4 Analysis	14
1.4.1 Uniform Case	16
1.4.2 General Case	19
1.5 Quality and Taste Differentiation	25
1.6 Discussion and Future Work	26
1.6.1 Welfare Analysis	26
1.6.2 Price Dynamics	27
1.6.3 Price Signaling	28
1.6.4 Perfect Bad News	29
1.7 Conclusion	30
1.8 Appendix [Latest Version of Online Appendix]	31
1.8.1 Section 4.1 Proofs	31
1.8.2 Section 4.2 Proofs	31
1.8.3 Section 4.3 Proofs	35

1.8.4	Section 5 Proofs	35
1.8.5	Section 3 Proofs	37
	References	39
2	Delegated Recruitment and Statistical Discrimination	41
2.1	Introduction	41
2.1.1	Literature	43
2.2	Model	46
2.2.1	Model Comments	49
2.3	Nonparametric Analysis	51
2.3.1	Candidate Posteriors	51
2.3.2	The First-Best, Equilibrium, and Suggestion Decisions	52
2.3.3	Individual and Group Statistical Discrimination	54
2.3.4	Contract Design in Equilibrium	58
2.4	Parametric Analysis without Commitment	61
2.4.1	First-Best Suggestion Region	62
2.4.2	The Equilibrium Contract	62
2.4.3	Hiring Outcomes by Group	66
2.5	Comparative Statics	69
2.6	Robustness and Extensions	72
2.6.1	Three-Part Contracts without Commitment	72
2.6.2	Multiple Candidates and Delegated Search	74
2.6.3	Nonparametric Full Surplus Extraction with Commitment	74
2.7	Conclusion	76
2.8	Appendix	77

2.8.1	Bayesian Updating Under Pareto-Uniform Information Structure . .	77
2.8.2	Derivations for the Proof of Theorem 2	78
2.8.3	Optimal Contract with Prior Heterogeneity	82
2.8.4	Proof of Theorem 3	83
2.8.5	Proofs for Comparative Statics	85
2.8.6	Qualitative Interviews	86
2.8.7	Proof of General Risk Attitudes	89
	References	90
3	Efficient Information Aggregation in DeGroot Model	92
3.1	Introduction	92
3.2	Literature	96
3.3	Model	97
3.4	First best learning	102
3.5	Improving DeGroot	104
3.5.1	Correcting the Interaction Matrix	107
3.5.2	Deleting Links	115
3.6	Speed of Convergence	118
3.7	Conclusion	121
	References	123

ACKNOWLEDGMENTS

I am thankful to my advisors, UCLA departments of Economics, my co-author, friends, and family, for helping and supporting me throughout my Ph.D. path.

VITA

2013–2017	B.A. (Economics), New Economic School and Higher School of Economics
2017–2018	Invited Lecturer, Economics Department, Higher School of Economics
2018–2020	M.A. (Economics), UCLA
2020	Graduate Summer Research Mentorship, UCLA
2019–2022	Teaching Assistant, Economics Department, UCLA
2022–2023	Dissertation Year Fellowship, UCLA
2024	(expected) Ph.D. in Economics, UCLA

CHAPTER 1

Reputational Underpricing

with Jakob Kohlhepp (UNC Chapel Hill)¹

1.1 Introduction

Consumer reviews often reflect not just product quality but also value for the price. Current consumers leave better reviews for higher-quality goods than for low-quality ones, but the reviews are worse if the same good is sold at a higher current price ([Abrate et al. \(2021\)](#), [Luca and Reshef \(2021\)](#)). Future consumers then make purchase decisions based on a firm’s reputation for quality, which is shaped by these reviews ([Chevalier and Mayzlin \(2006\)](#)). In settings where consumers observe past reviews but not past prices of a product (such as Amazon, Google Maps, Airbnb, etc.), the firm selling the product may strategically manipulate its own reputation through prices, a phenomenon that is supported by empirical evidence ([Sorokin \(2021\)](#)).

In this paper, we propose a novel model to analyze how firms make pricing decisions in the presence of reputational incentives driven by price-dependent consumer reviews. We derive a necessary and sufficient condition for these reputational incentives to cause

¹First version: February 2023. Previous Version: November 2023. This paper supersedes an earlier version circulated as “Consumer Reviews and Dynamic Price Signaling” ([Aleksenko and Kohlhepp \(2023\)](#)). We are grateful for insightful comments from Sushil Bikhchandani, Alex Bloedel, Simon Board, Daniel Clark, Felipe Goncalves, Jay Lu, Moritz Meyer-ter-Vehn, Ichiro Obara, Luciano Pomatto, Tomasz Sadzik, Eran Shmaya, and William Zame. We thank participants in the ASSA 2024 Annual Meeting; SEA 2023 Annual Meeting; All-UC Theory Conference; 35th Stony Brook International Conference on Game Theory; South-Western Economic Theory Conference; UCLA, UNC Chapel Hill, and SBU theory proseminars; and CalTech reading group where this work was presented.

the firm to price its product below consumers' willingness to pay, a phenomenon we call underpricing.

In our model, an infinite stream of short-lived consumers decide whether to purchase the product based on their expected utility from consuming it. After purchasing the product, a consumer can leave a review for it depending on the quality of the product and her realized utility, which additionally depends on the current price of the product and the consumer's idiosyncratic taste shock. Future consumers rationally learn about the quality of the product from past consumer reviews and the current price of the product, but they do not observe past prices. Our technical contribution is solving this reputation model with observable current actions (prices) and dynamic price-signaling.

To solve the model, we analyze the firm's trade-off between the reputational and the myopic pricing incentives. Consumers' ability to observe past reviews but not past prices of a product creates a reputation-management channel for the firm selling that product. Specifically, if the firm lowers the price of its product today, the firm will receive better consumer reviews today and build a better future reputation. Future consumers will have higher beliefs about the quality of the firm's product because they cannot distinguish whether the firm's better reviews are due to its higher product quality or lower past prices. The downside of lowering the price today is a lower current profit, either directly, via price, or indirectly, via signaling a lower quality today and lowering the current demand for that product.

We show that the consumer taste shock distribution, specifically, its *adjusted hazard rate*, determines whether underpricing occurs in equilibrium. For a wide range of primitives, including the case when consumer taste shocks are distributed uniformly, there is a unique equilibrium where the firm prices its product at consumers' willingness to pay at all reputation levels and never tries to build its reputation via underpricing. For another range of primitives, underpricing occurs only at lower reputation levels. Underpricing is unlikely to occur when the taste shock distribution is sufficiently dispersed. Additionally, we show that the reputation-management channel, and therefore underpricing, is more prominent when consumers arrive more frequently.

Although many insights can be extended to other review processes, we focus on the case when the reviews are modeled as perfect good news. Consumers leave only good reviews, and they can leave a review only for a high-quality product.² Each consumer's probability of leaving a good review is increasing in the utility delivered net of the price and is therefore decreasing in the price. We model this feature by introducing an explicit review rule: the consumer leaves a review if the overall ex-post utility depending on the price of the product and the consumer's idiosyncratic (ex-post) taste shock exceeds the given threshold.³ The taste shocks are i.i.d. mean zero distributed and realized only after the decision to purchase the product. These taste shocks can be interpreted as after-purchase idiosyncratic experiences from consuming the product. Examples include faster or slower delivery of the product or service in the restaurant on a given date, or horizontal matching shocks unobserved prior to purchase.

To build intuition and show that underpricing need not occur, we first consider the case when ex-post taste shocks are distributed uniformly. We show that reputational incentives are always dominated by the static profit motive. At lower reputation levels, the firm is forced to charge low prices, whether via underpricing or not. Under the uniform shock distribution, the reputational incentives become insensitive to price reductions relative to the overall speed of reputation building. Therefore, future reputation-building crowds out any incentives to underprice a product today. This means that both the high-quality and low-quality types of firm have the same pricing incentives, so price signaling never occurs. Thus, both types of firm pool at consumers' maximum willingness to pay at any rating and underpricing does not occur.

More generally, for an arbitrary taste shock distribution, there do not exist equilibria where firms pool and price their products below consumers' willingness to pay. Either there is a unique equilibrium where both firm types pool at consumers' willingness to

²There is empirical evidence for mostly positive reviews on online platforms including AirBnB (Carnehl et al. (2021a)) and Amazon (Hu et al. (2009)).

³This implies only people with extreme experiences select into leaving a review, an idea with empirical support (Schoenmüller et al. (2019), Lafky (2014), Marinescu et al. (2021)).

pay or the two types separate at low reputation levels and pool at high reputation levels. In the second case, for a given reputation level below a threshold, the high-quality type of firm (“high type”) prices its product below that of the low-quality type of firm (“low type”). Then, prices are also fully informative about the quality of the product at low reputation levels. We derive necessary and sufficient conditions on the ex-post taste shock distribution and the consumer arrival rate that determine which type of equilibrium arises. Underpricing does not occur in equilibrium if the density of marginal consumers that the firm can win by using underpricing is low relative to the mass of consumers who will leave reviews regardless of whether underpricing is used.

With equilibrium pricing behavior in hand, we analyze the welfare implications. Contrary to the standard intuition that review manipulation harms consumers, underpricing, when it occurs, makes consumers better off relative to a myopic benchmark for two main reasons. First, underpricing directly increases the consumer surplus. Second, underpricing by the high-quality firm increases the arrival rate of good news, which speeds up consumer learning and effectively transfers information rent from the low-quality firm to consumers. The welfare consequences for the high-quality firm are ambiguous because it cuts prices at low reputation levels but is sooner differentiated from the low-quality firm.

The rest of the paper is organized as follows. We discuss the related literature and our contribution to it in Section 1.2. We introduce the model in Section 1.3 and show the equilibrium characterization and main results in Section 1.4. The robustness of our results to a number of extensions and the welfare implications are discussed in Section 1.6. Section 1.7 concludes the paper.

1.2 Literature Review

Our paper is related to a set of papers that study how a firm sets prices over time in the presence of a consumer review system. In this literature, papers typically consider two types of reviews: those that do not depend on product prices (He and Chen (2018)) and those that do. Our model belongs to the second category. Many papers in this literature

consider consumers which are behavioral rather than Bayesian. [Shin et al. \(2021\)](#) consider consumers who choose a single quality and a single price faced by past consumers that rationalize the observed current average rating. [Carnehl et al. \(2021b\)](#) consider consumers who form beliefs about product quality that rationalize the current rating at the current price. [Crapis et al. \(2017\)](#) consider a firm that sets a price once and for all and consumers who assume all past consumers had the same information.

Behavioral assumptions on consumers are typically used because “a fully rational consumer would have to solve a dynamic signaling game with rating systems, which is a highly complicated problem” ([Carnehl et al. \(2021b\)](#)). Our contribution to this literature is an analysis of fully rational consumers that make Bayesian inferences from the current price and full history of reviews of a product. One important difference between our model and those in the literature is that we allow for static price signaling, and it occurs in equilibrium. Of the four papers mentioned in the previous paragraph, three explicitly assume price signaling does not occur.

A small literature ([Huang et al. \(2022\)](#), [Martin and Shelegia \(2021\)](#)) considers pricing incentives and learning in the presence of both consumer reviews and Bayesian consumers, similar to our paper. One key difference between the models in these papers and ours is that they consider a single period of building reputation via reviews, while we consider an infinite number of periods. This is one reason underpricing is more common in the models in this literature: there is only one attempt at building reputation, while underpricing generically need not occur in our model, because the firms might simply wait for their future selves to build their reputation.

Our paper is also related to work on sequential learning through review systems. In line with this literature, one goal of our paper is to understand how and what consumers learn from rating systems. However, we focus on how dynamic pricing undertaken by a forward-looking strategic firm interacts with consumer learning. In contrast, other work in this literature, including [Acemoglu et al. \(2017\)](#) and [Koh and Li \(2023\)](#), focuses on the case when prices are given but other questions are of interest, including how the selection of consumers into purchase impacts learning.

Our paper contributes to the literature on reputation management and is most related to [Holmström \(1999\)](#) and [Board and Meyer-ter Vehn \(2013\)](#). We discuss these papers using the language of our model, for clarity. In [Holmström \(1999\)](#), the firm has a quality that is fixed but initially unknown by everyone. The firm makes costly effort choices that impact the utility delivered to consumers. Importantly, while the full history of utility delivered is observed by consumers, effort choices both today and in the past are unobserved. In [Board and Meyer-ter Vehn \(2013\)](#), the firm can make costly effort choices that determine quality when it is redrawn. Importantly, while the firm observes its quality and its effort choices, consumers do not; instead, they observe signals about quality via a process that the firm cannot influence.

Our paper is also related to the strand of reputation literature with observable actions (such as the seminal paper by [Fudenberg and Levine \(1989\)](#), and most closely to [Pei \(2020\)](#)). The main differences between our paper and this literature are that in our model (1) we do not have the commitment type, (2) current rather than past actions are observable, (3) the long-lived player is not necessarily patient, (4) we analyze a different class of stage games.

In our model, the firm cannot influence the quality of its product; instead, the firm makes costly effort choices (prices) that impact the utility delivered to consumers. Reviews are left when the firm is high quality, but the firm can influence the arrival rate by changing the prices of its product. Importantly, consumers observe the full history of reviews and the current price, but not the full history of prices. Our paper is like [Holmström \(1999\)](#) (and unlike [Board and Meyer-ter Vehn \(2013\)](#)) in that the firm can influence learning about quality but not quality itself. Our paper is like [Board and Meyer-ter Vehn \(2013\)](#) (and unlike [Holmström \(1999\)](#)) in that the firm knows its quality, but consumers do not. Our paper is different from both in that the firm takes an observable action (price). Similar to [Board and Meyer-ter Vehn \(2013\)](#), we find that when incentives for underpricing exist, equilibria take a partition form with investment (underpricing) at low reputation levels and shirking (pricing at consumers' willingness to pay) at high reputation levels. Differ-

ent from Board and Meyer-ter Vehn (2013), we find that for a wide class of primitives, there is a unique equilibrium where investment never occurs.

Finally, our paper is related to a literature where individuals signal by choosing an information structure. A few examples are Rodríguez Barraquer and Tan (2022) (tasks on the job), Degan and Li (2021) (precision of information), and Daley and Green (2014) (grades across levels of education). In our case, the firm's choice of price influences the arrival rate of future information. In some cases, this gives rise to a form of endogenous single crossing where high-quality firms separate themselves from low-quality firms by choosing to price their products low, to invest in reputation. The biggest difference from our model is that we assume a repeated dynamic structure where the choice of signal structure is observed today, but its realization is observed in the future. For this reason, (semi-)separating equilibria are less common in our model.

1.3 Model

Firm. A single long-lived firm repeatedly sells a single product. Time is continuous, and the firm posts a price $p_t \in [0, 1]$ at each moment of time $t \in \mathbb{R}_+$. Production is costless and the future is discounted at rate r . The quality of the firm's product θ_t is low or high $\theta_t \in \{L, H\}$, with a prior probability $q_0 \in (0, 1)$ of being high at $t = 0$. Quality is exogenously redrawn from the same prior distribution at a Poisson rate χ .⁴ High quality is normalized to $H = 1$, and low quality is assumed to be strictly positive ($L > 0$).

Consumers. The market is composed of a stream of short-lived consumers that arrive at Poisson rate λ . When a consumer arrives, she decides whether to buy a single unit of the product. Consumer utility from purchasing a product of quality θ_t at price p_t is equal to $u_t = \theta_t - p_t + \varepsilon_t$, where the consumer's idiosyncratic ex-post taste shock ε_t is drawn i.i.d. from a symmetric, unimodal, and mean-zero distribution with CDF F_ε and

⁴Almost all results in the paper will be for a positive but "small" χ .

PDF f_ε .⁵ This shock is realized only after the good is purchased, i.e., the consumer makes the purchase decision based only on the expected quality net of the price. We normalize the consumer's outside option to 0, which implies each consumer purchases the good if her expected utility from consumption of that product is greater than 0, with indifference resolved in favor of purchasing.

Consumer Reviews. Consumer reviews are modeled by perfect good news. Specifically, a consumer only leaves a good review at t if (1) she purchases a high-quality product, $\theta_t = H$, and (2) her realized ex-post utility exceeds threshold \bar{u} : $u_t = H - p_t + \varepsilon_t \geq \bar{u}$. This implies that conditional on the product of high-quality being purchased, the probability of a good review being generated is decreasing in the price of the product:

$$\Pr(\varepsilon_t > \bar{u} - (1 - p_t)) = 1 - F_\varepsilon(\bar{u} - 1 + p_t)$$

Consumers never leave good reviews for a low-quality product, but as long as the product is truly high quality, consumers are more likely to leave good reviews if their expected utility from consumption is higher. The firm's review history $h^{t-} = \{t, \tau_1, \dots, \tau_n\}$ is a public history of good review arrival times before time t ($\tau_i < t$) that also tracks the current calendar time t .

Information. A consumer at t observes the review history h^{t-} and the currently posted price p_t of a product, but not past prices, and forms an expectation about the firm's current quality $\tilde{\theta}(p_t, h^{t-}) \in [L, H]$. Then, she buys the product if her expected utility from the consumption of that product is weakly positive: $\tilde{\theta}(p_t, h^{t-}) - p_t \geq 0$. The firm is privately informed about the quality of its product, but consumers are initially uncertain of it. The firm also observes the review history prior to setting a price $p_t = p(\theta^t, h^{t-})$.

Firm's Problem. Even though consumers arrive at discrete times, the Poisson structure implies that expected discounted profit can be expressed as if consumers were arriving as a flow. From the firm's perspective, at any small interval of time (dt), a consumer arrives with probability λdt . Thus the firm's expected profit during dt is equal to

⁵Normal, logistic, uniform, and type-1 extreme value random variables are examples.

$\mathbf{1}_{\{\tilde{\theta}(p_t, h^{t-}) \geq p_t\}} \lambda p_t dt$, and we can write the firm's expected discounted value at the $t = 0$ as an integral:

$$\max_{\{p(\theta^t, h^{t-})\}} \mathbb{E} \left[\int_0^{+\infty} e^{-rt} \mathbf{1}_{\{\tilde{\theta}(p(\theta^t, h^{t-}), h^{t-}) \geq p(\theta^t, h^{t-})\}} p(\theta^t, h^{t-}) \lambda dt \right]$$

1.3.1 Model Discussion

Several ingredients in our model allow us to highlight the key economic forces present in the environment without sacrificing tractability. We discuss these in turn.

The product quality in our model is exogenous and changing but highly persistent over time: $\chi > 0$ but small. This is mainly a technical assumption that guarantees the continuity of value functions at all points. It also guarantees that consumers do not learn the product quality perfectly and allows us to produce more realistic price dynamics. Intuitively, there may be some idiosyncratic changes in the firm's supply chain that impact its product quality that are observable by the firm but out of its control.

A binary type with perfect good news is a tractable benchmark model in the reputation literature. To model the idea that consumers are reciprocal and more likely to leave good reviews when the product is priced lower, we introduce an additional condition that a consumer leaves a good review only if the overall utility, including an idiosyncratic taste shock, is above some threshold. This model provides a micro-foundation and is also isomorphic to modeling the arrival rate function of perfect good news explicitly as a primitive. The ex-post taste shock ε_t captures unexpected differences in the consumer's individual experience of the product, such as faster or slower delivery of the product or service in the restaurant.

Consumers in our model are fully Bayesian in the way they update their beliefs based on the review history and the current price. However, they do not observe the full history of past prices. These ingredients imply that the firm can invest in its reputation by choosing off-path prices, and because the different types of firm may have different incentives,

price signaling becomes a possibility. In this way, we allow for price signaling to occur in any given period instead of assuming it away ex-ante.

1.3.2 Equilibrium Concept

We define a pure Markov Perfect Bayesian Equilibrium (MPBE) with the current firm's quality and the public belief about this quality, which we call the firm's reputation, as the Markovian states.

Definition 1 *The firm's reputation q is the public belief that the firm's quality is high:*

$$q(h^{t^-}) := \frac{\tilde{\theta}(h^{t^-}) - L}{H - L} \in [0, 1]$$

The firm's reputation is a single sufficient statistic that summarizes the whole review history and describes the quality distribution in the market for this review history. Under the Markov assumption, the firm's prices and consumers' beliefs depend on the review history only via the firm's reputation. Definition 2 below formalizes that the firm's reputation is Markovian and the full review history is not necessary for updating the firm's reputation in equilibrium.

To proceed with defining the equilibrium concept, we first introduce an auxiliary function, the **good news arrival rate**, which is the arrival rate of good news for the high-quality product sold at price p , conditional on that the high-quality firm sells its product at price p if a consumer arrives:

$$\lambda_g(p) := \lambda \cdot (1 - F_\varepsilon(\bar{u} - 1 + p))$$

Now we can formalize the equilibrium concept.

Definition 2 *A pure Markov Perfect Bayesian Equilibrium (MPBE) consists of*

1. Piecewise continuous⁶ in q firm's pricing strategies: $p(\theta, q) : \{L, H\} \times [0, 1] \rightarrow [0, 1]$;
2. Value functions: $V(\theta, q) : \{L, H\} \times [0, 1] \rightarrow \mathbb{R}_+$;
3. Consumers' belief about prices: $\tilde{p}(\theta, q) : \{L, H\} \times [0, 1] \rightarrow [0, 1]$;
4. Consumers' expectations about the firm's quality: $\tilde{\theta}(p, q) : [0, 1]^2 \rightarrow [L, H]$

such that:

(a) The value functions $V(\theta, q)$ solve the Hamilton-Jacobi-Bellman (HJB) equations (1.1) and (1.2).⁷

(b) Prices $p(\theta, q)$ maximize the right-hand sides of HJB equations (1.1) and (1.2):

$$rV(H, q) = \max_p \left\{ \mathbf{1}_{\{\tilde{\theta}(p, q) \geq p\}} \cdot [\lambda p + \lambda_g(p) \cdot (V(H, 1) - V(H, q))] + V_q(H, q) \cdot \frac{dq}{dt} + \chi(1 - q_0)(V(L, q) - V(H, q)) \right\} \quad (1.1)$$

$$rV(L, q) = \max_p \left\{ \mathbf{1}_{\{\tilde{\theta}(p, q) \geq p\}} \cdot \lambda p + V_q(L, q) \cdot \frac{dq}{dt} + \chi q_0(V(H, q) - V(L, q)) \right\} \quad (1.2)$$

where q jumps to 1 at rate $\mathbf{1}_{\{\tilde{\theta}(\tilde{p}(q), q) \geq \tilde{p}(q)\}} \cdot \lambda_g(\tilde{p}(H, q))$, and otherwise drifts as

$$\frac{dq}{dt} = -\mathbf{1}_{\{\tilde{\theta}(\tilde{p}(q), q) \geq \tilde{p}(q)\}} \cdot \lambda_g(\tilde{p}(H, q)) \cdot q(1 - q) + \chi \cdot (q_0 - q) \quad (1.3)$$

(c) Consumers' expectations about the firm's quality $\tilde{\theta}(p, q)$ is Bayesian for the on-equilibrium-path prices $\{p(L, q), p(H, q)\}$.

(d) Consumers' belief about price is correct $\tilde{p}(\theta, q) = p(\theta, q)$.

We now explain each component of this definition:

- **Strategies and beliefs.** Definition 2 formalizes that the firm's price and continuation value depend only on its current quality and reputation: $p(\theta_t, q(h^{t-})) = p(\theta^t, h^{t-})$.

⁶A condition that guarantees its integrability and differentiability of the value function.

⁷Where $V_q(\theta, q)$ is a left or right derivative $\frac{\partial V(\theta, q)}{\partial q}$ depending on whether $\frac{dq}{dt}$ is negative or positive.

Consumers' belief about prices for each type of firm depends only on the current reputation of the firm, and their expectations about the firm's quality depend only on the current reputation of the firm and the price of the product: $\tilde{\theta}(p_t, q(h^{t-})) = \tilde{\theta}(p_t, h^{t-})$.

- **HJB.** Equations (1.1) and (1.2) are recursive formulations of the low- and high-quality firms' problems. The first line of (1.1) includes the revenue stream as well as the possibility of getting a value jump as the reputation jumps from q to 1 after receiving a good review. It is multiplied by an indicator function because the firm gets the revenue and reviews only if a consumer buys the firm's product at the chosen price. The second line reflects how the future continuation value drifts down without good reviews and might also jump to a different type's value if the quality is redrawn.
- **Law of motion of reputation.** To derive the law of motion for the firm's reputation (1.3), we need to understand how the consumers form their beliefs about the firm's quality based on the firm's review history. The review history process is governed by the prices chosen by the firm. Consumers do not observe past prices, so they use their beliefs about those prices in order to update their belief about the firm's quality in the absence of good news. Intuitively, consumers "fill in" unobserved past prices using their understanding of equilibrium and the review history.

Consumers believe that the high-quality firm charges price $\tilde{p}(H, q)$ and receives a good review at arrival rate $\lambda_g(\tilde{p}(H, q)) \cdot \mathbf{1}\{\tilde{\theta}(\tilde{p}(H, q), q) \geq \tilde{p}(H, q)\}$. We derive the law of motion for the firm's reputation q using the fact that without redrawing the state, it is a martingale and it jumps to 1 immediately after a good review.

Otherwise, since the time of the last review, the reputation drifts down. An additional term $\chi \cdot (q_0 - q_t)$ represents the mean reversion of the firm's quality because the quality is stochastically redrawn at rate χ . The HJB equations include all these events to calculate the expected continuation value of each type of firm.

A set of *acceptable prices* is a set of prices at which consumers purchase the good:

$$\mathcal{P}_q := \{p \in [0, 1] | \tilde{\theta}(p, q) \geq p\} \quad (1.4)$$

Any type of firm at any reputation level would only choose a price for its product among the acceptable prices $p \in \mathcal{P}_q$ (1.4), because selling today adds weakly positive revenue to a stream of payoffs and allows a possibility of getting good news, which increases revenue in the future.⁸ Thus the law of motion of the firm's reputation in equilibrium can be expressed as

$$\frac{dq}{dt} = -\lambda_g(p(H, q)) \cdot q(1 - q) + \chi \cdot (q_0 - q) \quad (1.5)$$

There is a signaling game at every reputation level that endogenously defines both the firm's prices and consumers' expectations, which prohibits us from solving for full HJB equations paths for a given consumers' expectation function. Thus, to derive value functions and MPBE, we first need to define and solve an **auxiliary signaling game** as a static version of our model played at a single moment for a given reputation level, with the firm's payoffs derived from the HJB equations (1.1) and (1.2).

To be a signaling equilibrium at q for a given value function $V(H, q)$, the firm's prices must be optimal

$$\begin{aligned} p(H, q) &\in \arg \max_{p \in \mathcal{P}_q} \{\lambda p + \lambda_g(p)(V(H, 1) - V(H, q))\} \\ p(L, q) &= \max \mathcal{P}_q \end{aligned} \quad (1.6)$$

and the consumers' expectations function is correct for the equilibrium prices, i.e.,

1. $\tilde{\theta}(p(L, q), q) = L$ and $\tilde{\theta}(p(H, q), q) = H$, if $p(L, q) \neq p(H, q)$ OR
2. $\tilde{\theta}(p(L, q), q) = \tilde{\theta}(p(H, q), q) = qH + (1 - q)L$, if $p(L, q) = p(H, q)$.

⁸We show in online appendix Section 1.8.5 that $\forall q : V(H, 1) \geq V(H, q)$

Equation (1.6) includes only the parts of the firm's HJB (1.1) that depend on the chosen price and therefore reflect the same pricing incentives for the firm. Equation (1.6) illustrates the main trade-off in the model. Conditional on choosing a price that signals a quality high enough to sell today, a higher price increases the revenue today but decreases the arrival rate of good news (for the high type), which decreases future payoffs.

The low type's arrival rate of good news is always zero, and therefore the payoff is increasing, and $p(L, q) = \max \mathcal{P}_q$. We continue the analysis of this trade-off for the high type and how it affects the signaling equilibrium and MPBE structure in Section 1.4.

Therefore MPBE conditions (a)-(d) are satisfied if and only if, (A) the firm's prices and consumers' expectations satisfy the equilibrium conditions of the signaling games at (almost) every q for the given value functions, and (B) the value functions are derived from HJB equations (1.1), (1.2) for given price functions with the law of motion for q given by equation (1.5).

1.4 Analysis

In this section, we derive necessary and sufficient conditions for *underpricing* in equilibria: pricing below the consumer's expected value of the product at the given reputation of the firm $\tilde{\theta}(q) := qH + (1 - q)L$.

First, we analyze the case with a uniform taste shock distribution, which allows us to illustrate the main forces in the model and introduce a continuity equilibrium refinement. In this case, we show that underpricing will never arise in equilibrium and different firm types pool at consumers' willingness to pay $\tilde{\theta}(q)$.

Second, we solve the general case of the model and characterize the necessary and sufficient condition under which underpricing occurs in equilibrium. A key element of the model is the taste shock CDF F_ε and a related primitive object that we call the *adjusted hazard rate*:

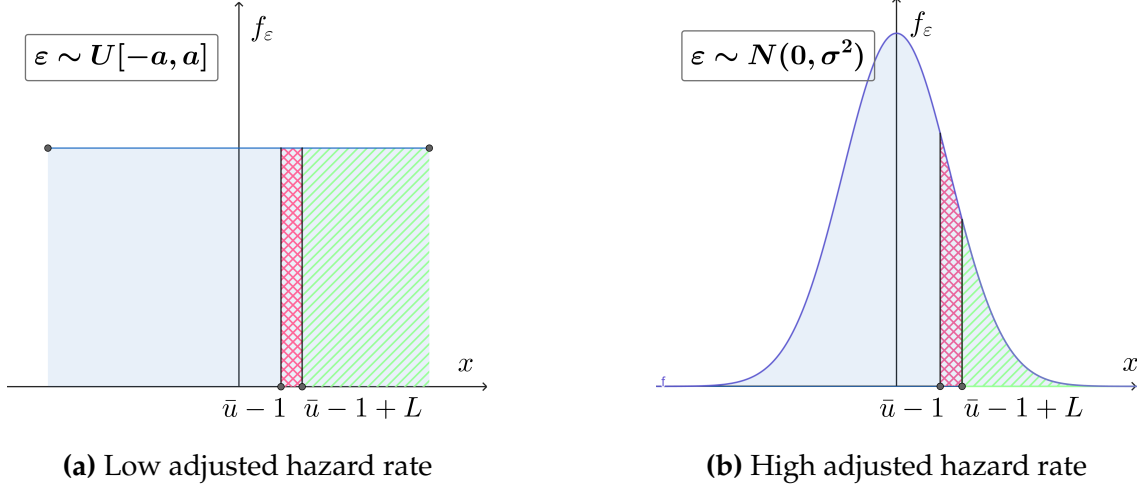


Figure 1.1

$$h_\varepsilon := \frac{\lambda(F_\varepsilon(\bar{u} - 1 + L) - F_\varepsilon(\bar{u} - 1))/L}{\lambda(1 - F_\varepsilon(\bar{u} - 1 + L)) + r} \quad (1.7)$$

The adjusted hazard rate represents the value of underpricing to the firm. The numerator of this expression is equal to the density of reviewers at $p = 0$ if $L \rightarrow 0$, and the average density of reviewers between $p = 0$ and $p = L$ if $L > 0$. It reflects the density of reviewers the firm can attract through underpricing when its reputation is lowest. The denominator includes the mass of consumers who leave reviews regardless of the firm's underpricing at its lowest reputation level. Thus, this expression is equal to the ratio of the marginal to the inframarginal reviewers (see Figure 1.1).

Additionally, the adjusted hazard rate is affected by the discount rate and the consumer arrival rate. A higher discount rate makes the current marginal reviewers less valuable. A higher consumer arrival rate makes the review channel, and therefore the marginal reviewers, more important.

In Theorem 1, we show that there is underpricing in (every) MPBE if and only if the adjusted hazard rate is high ($h_\varepsilon > h^*$). Otherwise, there is no underpricing in the unique MPBE.

1.4.1 Uniform Case

In this subsection, we analyze a version of our model where the taste shock ε is distributed uniformly. In this natural case, we show that the possibility of good news in the future crowds out any incentive of the firm to underprice its product today, and the reputational incentives do not lead to underpricing.

Assumption 1 *The taste shock has a uniform distribution with sufficiently large support ($a \geq \bar{u} \geq 1 - a$):*

$$\varepsilon \sim U[-a, a]$$

Under Assumption 1, we can show that λ_g is linear in price:

$$\lambda_g(p) = \lambda(1 - F_\varepsilon(\bar{u} - 1 + p)) = \lambda \cdot \frac{1 + a - \bar{u} - p}{2a} \quad (1.8)$$

We will show that underpricing does not occur, in two steps: First, we will solve the firm's problem and derive the pricing incentives for different types of firm in order to characterize all possible equilibria of the auxiliary game. Second, we will characterize a unique (in terms of prices) equilibrium in this model under a continuous belief refinement.

We can see three different components of the pricing incentives in the high-quality firm's objective function in equation (1.6) and (1.4). The reputational incentive is the combination of the probability of the good news with the value jump $\lambda_g(p) \cdot (V(H, 1) - V(H, q))$. The myopic profit maximization incentive is reflected by λp . The signaling incentive and the demand confound the other two incentives, because firms must choose a price at which the consumer's belief is sufficiently high to purchase the good ($p \in \mathcal{P}_q$).

We will analyze the first-order condition (FOC) of the firm's problem formulated in (1.6) subject to the selling-price constraint \mathcal{P}_q to determine the optimal price $p(H, q)$. We determine the pricing incentives by differentiating the firm's objective. Because the firm's

problem is linear in p , the sign of the FOC will determine if H chooses the maximum or the minimum price in \mathcal{P}_q :

$$\frac{\partial}{\partial p}[\lambda p + \lambda_g(p) \cdot (V(H, 1) - V(H, q))] = \lambda - \frac{\lambda}{2a} \cdot (V(H, 1) - V(H, q)) \quad (1.9)$$

Lemma 1 ⁹ For small χ ¹⁰, the high-type firm always prefers the highest possible price, i.e., (1.9) is positive and $p(H, q) = \max \mathcal{P}_q$ (1.4).

This lemma implies that for a low redrawing rate χ , the high-quality firm has no incentives to underprice its product. A high-quality firm gains the most from receiving a good review, $(V(H, 1) - V(H, q))$, when its reputation q and the current value $V(H, q)$ are low. However, when the firm's reputation is low, the product sells only at low prices. Because the firm always sells in any profit-maximizing strategy, the only deviations we need to consider are those from already low prices to even lower prices.

The marginal benefit of reducing the price today is decreasing the expected time until getting a good review and jumping to a high reputation level. Without good reviews at low reputation levels, the firm also has to charge low prices in the future. For a uniform distribution of taste shocks, the density of marginal consumers that the firm can win by underpricing, $\frac{\lambda}{2a}$, is low relative to the mass of consumers who will leave reviews regardless of whether the firm uses underpricing today or will use it in the near future (because the density in the tails of the uniform distribution is high). Thus, underpricing does not significantly reduce the expected time until getting a good review, and the benefit of it is lower than the cost of sacrificing the profit today.

From an algebraic perspective, one could try to break this result by making the $\lambda_g(p)$ function steeper, i.e., increasing the density of marginal consumers $\lambda'_g(p) = \frac{\lambda}{2a}$ in order to make the continuation value more sensitive to price. However, this effect is completely counteracted by also increasing the reviews arrival rate at $q \approx 0$, when the firm is forced

⁹The proof is provided in the online appendix, Section 1.8.1.

¹⁰There exists $\chi^* > 0$ such that for any $\chi < \chi^*$

to charge $p \approx 0$ no matter what: $\lambda_g(0) \geq \frac{\lambda}{2a}$. Increasing this rate decreases the value jump $V(H, 1) - V(H, 0)$, because a review will arrive very soon when $q = 0$ and q will jump to 1 almost immediately. For this reason, the firm has no incentives to decrease the price when the firm's reputation level is low because the good news will arrive very soon regardless of whether the firm uses underpricing.¹¹

Thus, the high and low types have the same pricing incentives and, as a result, all equilibria under uniform taste shocks will be pooling equilibria. The problem is that both types of firm are choosing the same price for each equilibrium belief, so on-path behavior pins down only one point of the equilibrium belief function. Thus there is a continuum of equilibria with pooling at any given price point and fully pessimistic off-path beliefs at all other prices ($\tilde{\theta}(p, q) = \mathbf{1}_{\{p=p(H, q)\}}\tilde{\theta}(q)$). Further, many common refinements do not help select an equilibrium because both types of firm have exactly the same preferences over actions for any consumers' belief functions. To remedy this problem, we introduce a continuity refinement.¹²

Assumption 2 *For all q , the expectation function $\tilde{\theta}(p, q)$ is continuous in p .*

This continuity refinement requires that small differences in price do not cause large jumps in perceived quality. This requirement is reasonable in the context of online marketplaces such as Amazon: We do not expect consumers to believe a product priced at \$99.99 has a significantly different quality than a product priced at \$100. Under this refinement, there is a unique equilibrium.

Proposition 1 *Under Assumptions 1 and 2, and for small χ , no underpricing is the unique MPBE: the high and low types pool at the willingness to pay, $p(L, q) = p(H, q) = \tilde{\theta}(q)$, at any reputation level q .*

¹¹This and the previous arguments can be thought of as a form of the one-shot deviation principle, where we consider a price path that is equal to consumers' willingness to pay at every reputation level of the firm, and then show one-shot price cuts are not profitable.

¹²A similar continuity refinement is used in [Gertz \(2014\)](#).

Proof. We show that any pooling equilibrium with the pooling price strictly below consumers' belief ($p^* < \tilde{\theta}(q)$) cannot be an equilibrium of the signaling game if the belief function is continuous by contradiction. If it is, then by definition of the equilibrium, the expectations $\tilde{\theta}(p, q)$ should be correct on-path (for p^*) and strictly above the price $\tilde{\theta}(p^*, q) > p^*$. By continuity of $\tilde{\theta}(p, q)$, for a small range of prices around p^* , this consumers' expectations about the quality of the product is still above the price. Thus, consumers would also buy the good for a price a little higher than the equilibrium price and both types of firm would prefer to deviate to that higher price (see Figure 1.2). This contradicts optimality. Therefore, in the unique equilibrium of the auxiliary game at any q , both types of firm always charge the consumers' willingness to pay: $\tilde{\theta}(q)$. Thus, the unique MPBE features no underpricing. ■




Figure 1.2

1.4.2 General Case

In this section, we derive the main results of the paper for a general class of unimodal taste shock distributions and show that underpricing occurs if and only if the adjusted hazard rate, h_{ε} , is sufficiently high. We show that any equilibrium is defined as a partition of the firm reputation interval, with no underpricing at higher reputation levels and possibly full underpricing at lower reputation levels.

If underpricing occurs, the high-quality firm charges lower prices than the low-quality firm, and prices are fully informative about quality. This happens because the high-

quality firm values establishing its reputation, while the low-quality firm does not. The high-quality firm knows that its quality is high and it only needs to “convince” consumers that this is so, to start collecting high profits. Thus, the high-quality firm prefers to underprice its product heavily when its current reputation is low.

In contrast, the low-quality firm at the lower reputation level cannot “convince” consumers that the quality of its product is high, not even by charging low prices. Thus, the best the firm can do is collect as much profit as possible given its reputation, which entails selling the product at face value to consumers, whereas the high-quality firm prefers to price its product even lower.

Throughout this section, we maintain two important assumptions. First, we retain our continuity belief refinement introduced in the linear case (Assumption 2). Second, we assume $\bar{u} \geq 1$, which implies that reviews are sufficiently selected, and therefore the review arrival rate $\lambda_g(u)$ is convex (Lemma 2). This assumption is equivalent to requiring that the fraction of consumers leaving reviews be never above $1/2$. This requirement is consistent with empirical evidence that a very small fraction (1 out of 1000) of consumers leave a review (Hu et al. (2017)).

Lemma 2 *If $\bar{u} \geq 1$, $\lambda_g(p)$ is decreasing and convex.*

Proof. Ex-ante utility, u_t , is bounded by 1 from above in any equilibrium. Thus $u_t - \bar{u}$ is below zero if $\bar{u} \geq 1$. A random variable that is unimodal has a CDF that is concave above its mode (zero for ε); thus F_ε is concave over the relevant domain of $\bar{u} - 1 + p$. Finally, $1 - F_\varepsilon(\bar{u} - 1 + p)$ is a linear increasing function composed within a convex decreasing function and is thus decreasing and convex. ■

We show that convex good news arrival rate functions can induce underpricing and separating equilibria at lower reputation levels. Intuitively, a convex arrival rate means that the good reviews arrival rate function is more sensitive to price when the price is low.

In Theorem 1, we will show precisely which distributions F_ε lead to underpricing for small χ . The necessary and sufficient condition for underpricing is that h_ε (1.7) be sufficiently high or F_ε be convex “enough” around the marginal consumer point, or the taste

shock be sufficiently concentrated around its mean. If F_ε is linear or nearly linear, then there is no underpricing. On the other hand, making F_ε steeper between $\bar{u}-1$ and $\bar{u}-1+L$ generates underpricing.

Theorem 1 *An equilibrium exists.*

1. If $h_\varepsilon < \frac{1}{1-L}$, then for small χ , no-underpricing is the unique MPBE:

$$\forall q : p(L, q) = p(H, q) = \tilde{\theta}(q)$$

2. If $h_\varepsilon > \frac{1}{1-L}$, then for small χ , there must be underpricing at low reputation levels and no underpricing at high reputation levels in any MPBE, i.e., $\exists 0 < q^* < q^{**} < 1$:

$$\forall q < q^* : p(L, q) = L, p(H, q) = 0$$

$$\forall q > q^{**} : p(L, q) = p(H, q) = \tilde{\theta}(q)$$

High h_ε corresponds to a case where there is a large density of consumers who can be convinced to leave a review after the firm cuts the price of its product from L to 0 and a small mass of consumers who leave reviews even if the firm does not cut the price. Whenever this is the case, the high-quality firm with a low reputation exploits the opportunity because it knows that if a customer leaves a review, the product will be revealed for what it truly is: high quality. The high-quality firm with a good reputation does not cut the price of its product, because the gains it obtains from a good review are small when the firm's reputation and the profit stream are already high.

In contrast, the low-quality firm, at any reputation level, knows that regardless of how low it cuts the price of its product, no customer will leave a good review. Therefore it sets the price of its product as high as it can to exploit its current reputation as much as possible. In this way, the two types of firm engage in separate pricing strategies at low reputation levels, but pool and engage in the same pricing strategy at high reputation levels.

Proof. We prove Theorem 1 in three steps: First, we discuss how the high-quality firm's pricing incentives determine the signaling-game equilibria at different reputation levels and characterize all possible MPBEs as a partition of the reputation interval. Second, we prove equilibrium existence and show that more generally there is a dichotomy: either there is a unique equilibrium with no underpricing or all equilibria have underpricing. Third, we derive the condition on λ_g along which the dichotomy occurs and translate it into a condition on the primitives.¹³

First, similarly to 1.4.1, we first need to analyze the high-quality firm's pricing incentives for a given reputation q . Let us fix all the future strategies of the firm and consumers, which determines $V(H, q)$ and $V(H, 1)$, and analyze an auxiliary signaling game at q . The high-type (H) firm's problem is characterized by (1.6):

$$p(H, q) \in \arg \max_{p \in \mathcal{P}_q} \{ \lambda p + \lambda_g(p)(V(H, 1) - V(H, q)) \}$$

Because $\lambda_g(p)$ is convex, the whole objective function in (1.6) is also convex in p , for any given $(V(H, 1) - V(H, q))$. This implies that the optimal solution to (1.6) is bang-bang: $p(H, q) \in \{0, \max \mathcal{P}_q\}$.

Given that the low type always plays $p(L, q) = \max \mathcal{P}_q$ and the high type plays 0 or $\max \mathcal{P}_q$, there can be two possible equilibria in the auxiliary game: separating or pooling. Knowing the consumer's preference, we can characterize the two equilibria. In any separating equilibrium, $p(H, q) = 0$ and $p(L, q) = L$ because the consumers are not ready to buy an obviously low-quality good for any price above L , and they are ready to buy any kind of good for any price weakly below L . In the pooling equilibrium $p(L, q) = p(H, q) = \tilde{\theta}(q)$ because of the continuity refinement.

Because all these equilibrium prices are within $[0, \tilde{\theta}(q)]$ interval, we need to determine the high type's preferences over prices on this interval, or, more specifically, its preferences over 0, L , and $\tilde{\theta}(q)$. Both the continuation value and the public belief about the quality q (1.3) are continuous functions of time, so the value function $V(\theta, q)$ is also a continu-

¹³The full proof is provided in the online appendix, Section 1.8.2.




Figure 1.3: Pricing Incentives

ous function in any MPBE. To simplify the next argument, we assume just for now that $V(H, q)$ is monotone increasing in q .¹⁴

$$\frac{\partial}{\partial p} \left(\lambda p + \lambda_g(p)(V(H, 1) - V(H, q)) \right) = \lambda + \lambda'_g(p)(V(H, 1) - V(H, q)) \quad (1.10)$$

We start by showing that pooling is a unique equilibrium of the auxiliary game at high reputation levels. When q is high and $V(H, 1) - V(H, q)$ is small (since $V(H, q)$ is continuous), the static profit motive dominates the reputational incentive and (1.10) is positive for any p . Therefore, the objective in (1.6) is monotone increasing in p and H always prefers higher prices to lower prices (see Figure 1.3a), thus $p(H, q) = \max \mathcal{P}_q$ and pooling at $\tilde{\theta}(q)$ is a unique equilibrium of the auxiliary game. Because of the continuity of $V(H, q)$, there must be a non-empty interval of high reputations at which this is a unique equilibrium of the auxiliary game.

Next, we show that separating is a unique equilibrium of the auxiliary game for low reputation levels. When reputation q is low and the value gap $V(H, 1) - V(H, q)$ is large, the high type prefers $p = 0$ to $p = \tilde{\theta}(q)$ and therefore to all prices in $[0, \tilde{\theta}(q)]$ (see Figure 1.3b). This happens because the reputational incentive becomes more significant than the static motive in (1.10). In this case, the high type unambiguously chooses 0 in any

¹⁴We relax this assumption in the full proof in the online appendix, Section 1.8.2.

equilibrium \mathcal{P}_q , and separating $p(H, q) = 0$, $p(L, q) = L$ is a unique equilibrium of the auxiliary game.

Finally, if q is intermediate, such that both $\tilde{\theta}(q)$ and $V(H, 1) - V(H, q)$ are sufficiently large for the high type to prefer $p = \tilde{\theta}(q)$ to $p = 0$ to $p = L$ (see Figure 1.3c). In this case, both pooling and separating equilibria are possible in the auxiliary game. In the pooling equilibrium, both types charge $p = \tilde{\theta}(q)$ and have no profitable deviations. In the separating equilibrium, the high type would like to deviate from 0 to $\tilde{\theta}(q)$, but it is not available because $\max \mathcal{P}_q = L$. Thus, neither type has a profitable deviation from $p(H, q) = 0$ and $p(L, q) = L$.

This auxiliary-game equilibria characterization suggests that equilibria can be described as partitions with separating, or underpricing, at low reputation levels, pooling at high reputation levels, and multiple equilibria at intermediate reputation levels (we formalize how this partition and the thresholds q^* and q^{**} are defined in the appendix).

Second, we prove equilibrium existence by constructing a partition equilibrium with underpricing below some reputation threshold and no underpricing above it. We start with a no-underpricing strategy profile and keep increasing the underpricing interval of the partition until it becomes an equilibrium (a formal proof is provided in the appendix and relies on the intermediate value theorem).

Then we show generally that if there exists an equilibrium with complete pooling and no underpricing, then it is unique. Otherwise, there must be some non-empty interval of low reputations where the auxiliary game exhibits separating equilibria. We show it by proving that a complete pooling equilibrium generates the highest values for the high type at any q , including $V(H, 1)$ (the full proof relies on comparing equilibria values by imitating the MPBE with underpricing as a long deviation from the MPBE without underpricing). It also generates the largest gain from a positive review, which is the value gap $V(H, 1) - V(H, 0)$ (this follows from the 1.1). Underpricing cannot occur in any equilibrium, because the high type has no incentive to use underpricing even when the value gap is largest.

Finally, the condition separating between cases (1) and (2) of the theorem follows from the HJB equation (1.1). If χ is small, then approximately $V(H, 1) = \lambda/r$. Then $V(H, 0) = (\lambda_g(L) \cdot V(H, 1) + \lambda L)/(\lambda_g(L) + r)$. Then for the high type to prefer $p = L$ to $p = 0$, the absolute value of the average slope of λ_g between them $(\lambda_g(1) - \lambda_g(L))/L$ should be smaller than $\lambda/(V(H, 1) - V(H, 0))$. We can rewrite this condition as

$$\frac{\lambda_g(1) - \lambda_g(L)}{L} < \frac{\lambda_g(L) + r}{1 - L}$$

or as equation 1.11 in terms of the model primitives:

$$\frac{F_\varepsilon(1 - \bar{u}) - F_\varepsilon(1 - L - \bar{u})}{L} < \frac{F_\varepsilon(1 - L - \bar{u}) + r/\lambda}{1 - L} \quad (1.11)$$

which is equivalent to $h_\varepsilon < \frac{1}{1-L}$. ■

1.5 Quality and Taste Differentiation

In this section, we discuss how the spread of the taste shock distribution, consumer arrival rate, and discount factor affect the occurrence of underpricing in equilibrium. A larger horizontal differentiation relative to vertical differentiation leads to a less likely occurrence of underpricing in equilibrium due to the highly heterogeneous consumer review behavior, which cannot be easily affected by the price. The following corollary gives a partial characterization of the set of situations when underpricing does not occur.

Corollary 1.1 ¹⁵ *Take a set of primitives $L, q_0, \lambda, r, F_\varepsilon$. Then*

1. *There exists $\alpha^* < +\infty$, such that for any $\alpha > \alpha^*$ and $\varepsilon' = \alpha\varepsilon$ no underpricing is the unique MPBE (for small χ).*
2. *There exists $\lambda^* > 0$, such that for any $\lambda < \lambda^*$ no underpricing is the unique MPBE (for small χ).*

¹⁵The full proof is in the online appendix.

3. *There exists $r^* < \infty$, such that for any $r > r^*$ no underpricing is the unique MPBE (for small χ).*
4. *There exists $L^* < 1$, such that for any $L > L^*$ no underpricing is the unique MPBE (for small χ).*

The first statement of Proposition 1.1 states that increasing the consumer taste shock spread above some threshold (by multiplying it by a constant) guarantees that there is no underpricing in equilibrium. The same is true if the discount rate is sufficiently large, the consumer arrival rate is sufficiently small, or the low quality is sufficiently high.

$$\frac{(F_\varepsilon(\bar{u} - 1 + L) - F_\varepsilon(\bar{u} - 1))/L}{1 - F_\varepsilon(\bar{u} - 1 + L) + r/\lambda} < \frac{1}{1 - L} \quad (1.12)$$

To prove show these comparative statics, we rewrite the no-underpricing condition from Theorem 1 as (1.12). Increasing (resp., decreasing) any of the parameters from 1–4 above (resp., below) some threshold decreases the left-hand side of (1.12) below its right-hand side $\frac{1}{1-L}$ and thus guarantees the equilibrium from the first case of Theorem 1¹⁶.

1.6 Discussion and Future Work

1.6.1 Welfare Analysis

Underpricing benefits consumers at the expense of the low-quality firm. To illustrate this, compare our model to a myopic benchmark. Consider a model where the firm ignores its ability to influence future reviews through prices. In this myopic model, the first-order conditions of the low- and high-quality firm are the same, with both desiring to price their products at consumers' willingness to pay, which implies zero consumer surplus.

Whenever the adjusted hazard rate is high, consumer surplus is unambiguously higher under our baseline model than in the myopic benchmark via the direct benefit of paying

¹⁶See online appendix Section 1.8.3 for a detailed proof.

a strictly lower price for a range of reputation levels. Moreover, the high-quality firm sets a lower price to speed up the arrival rate of good reviews. When this happens, the low-quality firm is revealed to be low quality and forced to charge exactly its market value.

Although the welfare benefits of underpricing seem straightforward ex-post, they are not obvious ex-ante. This is because once a good review arrives, the price increases to 1, and consumer surplus is 0. Thus revealing the quality of the product does not directly lead to high consumer surplus.

1.6.2 Price Dynamics

In our model, a combination of quality changes, review arrivals, and (in some cases) strategic underpricing generates nontrivial price dynamics over time. We highlight several examples in turn.

Consider the case with uniform consumer taste shocks so that there is no underpricing. Then the price moves downward over time until a good review arrives, at which point the price jumps to 1. The true quality of the firm's product matters only indirectly via the reputation dynamics.

Next, consider the case when the adjusted hazard rate is high enough so that there is underpricing. If a good review arrives, the price jumps immediately to 1 and then moves downward until it reaches a critical threshold. At this point, what occurs depends on the quality of the firm's product. If it is high quality, the firm will engage in underpricing and the price will suddenly drop to 0 and stay there until a good review arrives. If the product is low quality, the price will still drop, but only to $L > 0$. The price will remain at L until the quality becomes high, at which point the firm will drop the price further to 0 and engage in underpricing.

The taste shock distribution is the primitive of the model which determines whether underpricing occurs. Because preference distributions are typically unobserved, we may want an alternative empirical test for underpricing. Price dynamics provide such a test. When underpricing is occurring, we should observe both upward jumps in prices when

good reviews arrive and downward jumps in prices when the firm's reputation is low. When underpricing is not occurring, we should observe only upward jumps in prices at all levels of a firm's reputation.

In many models with uncertain quality, firms price their products low early on in order to build their reputation later, a strategy called introductory pricing.¹⁷ Although there are conceptual differences between our model and many of those in the literature, underpricing in our model can be viewed as a form of introductory pricing. To see this, consider when high-quality products are uncommon but nevertheless a firm begins its life with a high-quality product. Then the firm's reputation will start out low, and in some cases, the firm will underprice its product initially in order to build reputation, which it will exploit via higher prices when a good review arrives. Crucially, introductory pricing of this type occurs only for certain consumer taste shock distributions.

1.6.3 Price Signaling

In our model, when price signaling occurs, the firm with a high-quality product signals with a low price. This may seem to clash with past work, where high-type firms typically signal with higher prices (Milgrom and Roberts (1986)). However, to better understand our result, recall that the high-quality firm prices its product lower than the low-quality firm only conditional on reputation. However, from an unconditional perspective, high-quality firms are more likely to have high reputations at any given moment of time because they have some probability of receiving good reviews, while low-quality firms do not. As a result, there is still a sense in which firms with high-quality products set generally higher prices: they enjoy a high reputation for longer.

Even with this qualification, it is still true that when price signaling occurs, a price of 0 signals high quality, while a price of $L > 0$ signals low quality. To understand why this occurs, we can ask what single crossing is supporting separation. Recall that there is no difference in the costs of producing high- and low-quality products, and conditional on a

¹⁷See for example, Shapiro (1983).

firm's reputation, there is no difference in demand. The only source that could generate single-crossing is the review process. This is why the high-quality firm has an incentive to underprice its product, while the low-quality firm does not: using underpricing today can increase the arrival rate of good reviews in the future.

1.6.4 Perfect Bad News

In our main model, we consider a review process with perfect good news. Under this model, the interesting strategic choices and forces operate when the firm has a high-quality product. When underpricing occurs, it is used by a firm with a high-quality product attempting to improve its low reputation. One could also consider a perfect-bad-news review process, where consumers leave reviews only for low-quality products. Our framework can be easily adjusted to accommodate such an extension.

In this alternative model, the interesting strategic choices and forces operate when the firm has a low-quality product. If underpricing occurs, it will be used by the low-quality firm with a good reputation attempting to preserve its reputation. In this case underpricing can harm consumers, because it allows the low-quality firm to slow down consumer learning.

In general, the forces in this model will differ from those in the perfect-good-news model. However, when idiosyncratic taste shocks follow a uniform distribution, it continues to be true that no underpricing occurs.¹⁸ In this sense, one of the main results of the paper is robust to some alternative review processes.

In terms of price dynamics, our benchmark model produces a smooth downward trend in price punctuated by sudden upward jumps when good reviews are left. An alternative model with perfect bad news produces the opposite: a smooth upward trend in price punctuated by sudden downward jumps when bad reviews are left.

¹⁸See online appendix Section 1.8.4 for formal results.

1.7 Conclusion

This paper proposes a model of dynamic pricing, where a firm privately informed about the quality of its product faces a market of rational consumers and a rating system that depends on both product quality and net utility delivered. The model allows us to fully account for three important economic forces at play in many online markets: static profit maximization, price signaling, and ratings-based reputation building. We characterize a simple necessary and sufficient condition for when such a review system leads to underpricing in equilibrium.

1.8 Appendix [Latest Version of Online Appendix]

1.8.1 Section 4.1 Proofs

Proof of Lemma 1. We prove this lemma by contradiction. First, we notice that the high-quality firm's value $V(H, q)$ at any q cannot be lower than that of charging $p = 0$ until receiving a good review. The high type can always charge zero price at any reputation level, because the consumers are ready to buy a product of any quality at any price weakly below L .

Moreover, $V(H, 1) < \lambda/r$; that is, selling the product to every consumer arriving at the maximum price consumers are possibly ready to pay.

Then for small χ , rearranging (1.1) together with (1.3) under Assumption 1 implies

$$\frac{\lambda}{2a} \cdot (V(H, 1) - V(H, q)) \leq \lambda_q(0) \cdot (V(H, 1) - V(H, q)) \leq rV(H, 0) \leq rV(H, 1) \leq r \cdot (\lambda/r) = \lambda.$$

Therefore,

$$\lambda - \frac{\lambda}{2a} \cdot (V(H, 1) - V(H, q)) \geq 0.$$

This inequality immediately implies that 1.9 is positive for any q . ■

1.8.2 Section 4.2 Proofs

To proceed with the characterization of the set of all possible equilibria, we require additional notation. First, we need to determine what happens to the firm's reputation on path in any given equilibrium.

Definition 3 *The lowest rating is $\underline{q} := \sup\{q \in (0, 1) | dq/dt \geq 0\}$ for a given equilibrium price $\tilde{p}(H, q)$.*

Lemma 3 *The lowest rating is well defined, and $\underline{q} \in (0, q_0)$. Without good reviews, the firm's rating drifts down until it reaches \underline{q} , where it stays forever.*

Proof. By definition of \underline{q} , $dq/dt < 0$ for $q > \underline{q}$ and q drifts down without good reviews. When q drifts down to \underline{q} , it does not drift up or down anymore and stays at \underline{q} without good news. Thus, all reputation levels $q < \underline{q}$ are off-path. ■

Proof of Theorem 1. In equilibrium, H 's price $\tilde{p}(H, q)$ determines \underline{q} , and given \underline{q} , equilibrium prices should be defined on $[\underline{q}, 1]$. We are now ready to define equilibrium partition thresholds for a given equilibrium and unify those partitions across all equilibria as q^* and q^{**} :

$$q^* = \inf_{\text{all equilibria}} \inf\{q \in [\underline{q}, 1] | p(L, q) > L\}$$

$$q^{**} = \sup_{\text{all equilibria}} \sup\{\{q \in [\underline{q}, 1] | p(L, q) < \tilde{\theta}(q)\} \cup \{\underline{q}\}\}.$$

The auxiliary-game equilibria characterization suggests that equilibria can be described as a partition with separating, or underpricing, at low reputation levels, pooling at high reputation levels, and multiple equilibria at intermediate reputation levels. By the definition of thresholds q^* and q^{**} , $[0, q^*)$ and $(q^{**}, 1]$ are separating and pooling regions, respectively; that is, in any equilibrium, high and low types always pool at $\tilde{\theta}(q)$ for any $q \in (q^{**}, 1]$ and always separate at $p(L, q) = L$, $p(H, q) = 0$, respectively, for any on-path $q \in [0, q^*)$. For a range of reputations $[q^*, q^{**}]$, the prices can vary across multiple equilibria. Now we need to show the existence of MPBE and characterize the equilibrium dichotomy in terms of these thresholds.

Lemma 4 *MPBE exists.*

Proof. We prove the existence of MPBE by constructing a bi-partition equilibrium with underpricing below some reputation threshold and no underpricing above it: $q^* = q^{**}$. We start with a no-underpricing strategy profile $\underline{q} = q^* = q^{**}$ and check if it is an equilibrium from the firm's optimality perspective. If it is not, then the high type wants to underprice at low ratings, and we start increasing the underpricing region by increasing $q^* = q^{**}$. Our goal is to find a threshold with which the H 's pricing incentives are

consistent (there are multiple such thresholds), for instance, at which H is indifferent between $p = 0$ and $p = \tilde{\theta}(q)$ (this would be just one possible equilibrium). We know that at $q^* = q^{**} = \underline{q}$ H strictly prefers 0. Also, if $q^* = q^{**} = 1$, then $V(H, q) = 0$ for any q and H prefers any positive price to 0. By the continuity of the values in $q^* = q^{**}$, there exists a threshold at which H is indifferent between $p = 0$ and $p = \tilde{\theta}(q)$, and this will be a threshold for which the strategy profile is an equilibrium. Thus, an equilibrium exists. ■

Then we show generally that if there exists an equilibrium with complete pooling and no underpricing, then it is unique: $q^* = q^{**} = \underline{q}$. Otherwise, there must be some non-empty interval of low reputations where the auxiliary game exhibits separating equilibria: $q^* > \sup \underline{q} > 0$. We show that this interval exists by proving that a no-underpricing equilibrium generates the highest values for the high type at any q , including $V(H, 1)$ and therefore the largest gap $V(H, 1) - V(H, \underline{q})$ (this follows from the 1.1). Then, underpricing cannot occur in any equilibrium, because the high type has no incentive to underprice in the NUP MPBE, where the value gap is largest.

Lemma 5 *If there are multiple MPBE and no-underpricing is one of them, then it generates the largest $V(H, 1)$ among all MPBE.*

Proof. Consider any other possible equilibrium with underpricing at some reputation levels. We can recreate it as a long off-path deviation of H from the NUP equilibrium, because pricing at 0 is always allowed. This deviation is not profitable for H , because we assumed that complete pooling is an equilibrium and no single-shot or longer deviations are profitable. Thus the pooling equilibrium value is higher than the off-path deviation one at any q . The off-path deviation value, on the other hand, is higher than the value in the underpricing equilibrium we picked because the prices are the same for every q , but the public belief q drifts down slower in the off-path deviation, because the consumers do not expect lower prices and a higher good news arrival rate in the past, which benefits H . This implies that the value in the pooling equilibrium is higher than in the underpricing one. The proof is concluded. ■

Thus if no-underpricing is an equilibrium, it generates the largest $V(H, 1)$ and thus the largest gap $V(H, 1) - V(H, \underline{q})$ because

$$V(H, 1) - V(H, \underline{q}) = \frac{(r^2 + r\chi)V(H, 1) - (r + \chi)\lambda L}{r^2 + r\chi + (r + \chi q_0)\lambda_g(L)},$$

which is increasing in $V(H, 1)$. The expression follows from equations (1.1 and (1.2).

The next step is to show that if no-underpricing is an MPBE, it is unique. Here we explicitly rely on small χ . Specifically, when χ is small, \underline{q} is small and $\tilde{\theta}(\underline{q}) \approx L$. Thus the choice between $p = 0$ and $p = L$ is nearly the same as between $p = 0$ and $p = \tilde{\theta}(\underline{q})$. Then by the continuity of the problem in χ , if χ is small enough, H prefers $p = L$ to $p = 0$ because it prefers $p = \tilde{\theta}(\underline{q})$ to $p = 0$. Therefore, there can be no underpricing signaling equilibria with the largest gap $V(H, 1) - V(H, \underline{q})$, which implies that there cannot be any underpricing signaling equilibria (because all other value gaps $V(H, 1) - V(H, q)$ are smaller).

The previous point implies that the condition separating between cases (1) and (2) of the theorem can be characterized as an explicit condition for when no-underpricing is an MPBE. This condition follows directly from the *HJB*.

If χ is small, then $V(H, 1) \approx \lambda/r$ and $V(H, 0) \approx (\lambda_g(L) \cdot V(H, 1) + \lambda L)/(\lambda_g(L) + r)$. Then for the high type to prefer $p = \tilde{\theta}(\underline{q}) \approx L$ to $p = 0$, the absolute value of the average slope of λ_g between them $(\lambda_g(1) - \lambda_g(L))/L$ should be smaller than $\lambda/(V(H, 1) - V(H, 0))$. We can rewrite this condition as

$$\frac{\lambda_g(1) - \lambda_g(L)}{L} < \frac{\lambda_g(L) + r}{1 - L}$$

or as equation 1.11 in terms of the model primitives:

$$\frac{F_\varepsilon(1 - \bar{u}) - F_\varepsilon(1 - L - \bar{u})}{L} < \frac{F_\varepsilon(1 - L - \bar{u}) + r/\lambda}{1 - L},$$

which is equivalent to $h_\varepsilon < \frac{1}{1-L}$.

Finally, when relaxing the monotonicity of $V(H, q)$ assumption, we need to show that the underpricing and no-underpricing regions in case (2) of Theorem 1 are non-empty. This follows from the continuity of $V(H, q)$ in q . Specifically, if $h_\varepsilon > \frac{1}{1-L}$, then H prefers to underprice at \underline{q} and at least some small interval of q 's around it. ■

1.8.3 Section 4.3 Proofs

Proof of Comparative Statics. Let us rewrite h_ε (1.7) in an alternative form:

$$h_\varepsilon := \frac{(F_\varepsilon(\bar{u} - 1 + L) - F_\varepsilon(\bar{u} - 1))/L}{1 - F_\varepsilon(\bar{u} - 1 + L) + r/\lambda}.$$

We want to show that we can make this expression below $1 < \frac{1}{1-L}$ in either part of Corollary 1.1.

1. Increasing α above a large threshold increases $1 - F_{\alpha\varepsilon}(\bar{u} - 1 + L)$ and decreases density $f_{\alpha\varepsilon}(x)$ for any $x \in [\bar{u} - 1, \bar{u} - 1 + L]$. By decreasing the density at any point of this interval below $\frac{r}{\lambda(1-L)}$, we make $h_{\alpha\varepsilon}$ below $\frac{1}{1-L}$.
2. The adjusted hazard rate h_ε is increasing in λ and $\lim_{\lambda \rightarrow 0} h_\varepsilon = 0$. Thus there is a threshold λ^* above which there is no underpricing.
3. The adjusted hazard rate h_ε is decreasing in r and $\lim_{r \rightarrow \infty} h_\varepsilon = 0$. Thus there is a threshold r^* above which there is no underpricing.
4. $h_\varepsilon \leq \lim_{L \rightarrow 0} h_\varepsilon < \infty$. By making $L < L^*$ and $\frac{1}{1-L}$ above this limit, we guarantee no underpricing.

1.8.4 Section 5 Proofs

Perfect Bad News with Uniform Taste Shock.

Assume a model modification where consumers leave only bad reviews if the quality is low and the overall utility is below \underline{u} , and the taste shocks are distributed i.i.d. uniformly

between $-a$ and q with $a > \max\{L - \underline{u}, 1 - (L - \underline{u})\}$. Then the bad news arrival rate is linear in the price:

$$\lambda_b(p) = \lambda \cdot \frac{\underline{u} - L + a + p}{2a}$$

Proposition 2 For small χ , the no-underpricing MPBE is the unique MPBE.

Proof. L 's pricing incentives are given by

$$\frac{\partial V}{\partial p} = \lambda - \frac{\lambda}{2a} \cdot (V(L, q) - V(L, 0)).$$

Given that $V(L, 0) \approx 0$ and $rV(L, 1) \approx \max_p\{\lambda p + \lambda_b(p)(V(L, 0) - V(L, 1))\}$, we can show that

$$V(L, 1) - V(L, 0) \leq \frac{\lambda}{r + \lambda_b(1)}$$

and

$$\frac{\partial V}{\partial p} \geq \lambda - \frac{\lambda}{2a} \cdot \frac{\lambda}{r + \frac{\lambda}{2a}} > 0$$

Therefore, L always prefers $\max \mathcal{P}_q$, and pooling at the consumers' willingness to pay is a unique signaling equilibrium at any q . Thus, the no-underpricing MPBE is the unique MPBE.

Consumer Arrival Rate Depending on q .

Assume a model modification where consumers are more likely to arrive for a firm with a higher reputation (i.e., $\lambda(q)$ is increasing). We want to show that the analysis remains the same as in the main model, and we can easily characterize the condition for the case when there is underpricing in equilibrium.

Notice that for any given value function $V(H, q)$ and q , $\lambda(q)$ multiplies both the static and the dynamic parts of H 's objective. Therefore, signaling equilibrium characterization remains the same.

However, $\lambda(q)$ affects the overall solution of the HJB equation (1.1), $V(H, q)$. Therefore, we need to verify the condition on the taste shock distribution for the lowest q and $\lambda(0)$,

$$h_\varepsilon = \frac{\lambda(0) \cdot (F_\varepsilon(\bar{u} - 1 + L) - F_\varepsilon(\bar{u} - 1))/L}{\lambda(0) \cdot (1 - F_\varepsilon(\bar{u} - 1 + L)) + r}.$$

However, given the threshold which follows from $V(H, 1) \approx \frac{\lambda(1)}{r}$, which is equal to $\frac{1}{\frac{\lambda(1)}{\lambda(0)} - L}$ instead of $\frac{1}{1-L}$ in the benchmark model. Thus, if $\lambda(0) = \lambda$ from the benchmark model, then the reputation-based consumer arrival rate (popularity-based demand) increases the possibility of underpricing, which is quite intuitive given that the higher reputation increases the profit stream via both price and demand. ■

1.8.5 Section 3 Proofs

Law of motion of q without redrawing the state:

$$q_t = q_t \cdot \lambda_g(\tilde{p}(q_t))dt \cdot 1 + (1 - q_t \cdot \lambda_g(\tilde{p}(q_t))dt)q_{t+dt}$$

$$q_t = q_t \cdot \lambda_g(\tilde{p}(q_t))dt \cdot 1 + (1 - q_t \cdot \lambda_g(\tilde{p}(q_t))dt)(q_t + dq_t)$$

$$dq_t \cdot (1 - q_t \cdot \lambda_g(\tilde{p}(q_t))dt) = q_t(1 - q_t)\lambda_g(\tilde{p}(q_t))dt$$

$$\frac{dq_t}{dt} = q_t(1 - q_t)\lambda_g(\tilde{p}(q_t)) \cdot \frac{1}{(1 - q_t \cdot \lambda_g(\tilde{p}(q_t))dt)}$$

$$\frac{dq_t}{dt} = \lim_{dt \rightarrow 0} q_t(1 - q_t)\lambda_g(\tilde{p}(q_t)) \cdot \frac{1}{(1 - q_t \cdot \lambda_g(\tilde{p}(q_t))dt)} = q_t(1 - q_t)\lambda_g(\tilde{p}(q_t))$$

Adding mean reversion and a possibility of selling or not selling at $\tilde{p}(q)$ implies the law of motion in equation (1.3).

HJB for high type (1.1). We show an intuitive way of deriving this HJB, simplifying some aspects, such as demand, for ease of notation. This HJB can also be derived more formally from the continuation value of the firm introduced in Section 3.1:

$$V(H, q) = \lambda p dt + (1 - r dt) [(1 - \chi(1 - q_0) dt) \lambda_g(p) dt V(H, 1) \\ + (1 - \chi(1 - q_0) dt) (1 - \lambda_g(p) dt) V(H, q + dq) + \chi(1 - q_0) dt V(L, q + dq)]$$

$$rV(H, q) dt = \lambda p dt + (1 - r dt) [(1 - \chi(1 - q_0) \lambda_g(p) dt (V(H, 1) - V(H, q)) \\ + (1 - \chi(1 - q_0) dt) (1 - \lambda_g(p) dt) dV(H, q) + \chi(1 - q_0) dt (V(L, q) + dV(L, q) - V(H, q))] \quad (1.13)$$

$$rV(H, q) = \lim_{dt \rightarrow 0} \text{RHS}(1.13) \\ = \lambda p + \lambda_g(p) (V(H, 1) - V(H, q)) + \frac{dV(H, q)}{dt} + \chi(1 - q_0) (V(L, q) - V(H, q)).$$

Incorporating the demand $\mathbf{1}_{\{\tilde{\theta}(p, q) \geq p\}}$ into this equation gives us equation (1.1).

Finally, we want to show that for small χ , $V(H, 1) \geq V(H, q) \forall q$. That follows from the fact that $\lim_{\chi \rightarrow 0} V(H, 1) = \lambda/r$ and $\forall q: V(H, q) < \lambda/r$.

References

- Abrate, Graziano, Sarah Quinton, and Rebecca Pera (2021) "The relationship between price paid and hotel review ratings: Expectancy-disconfirmation or placebo effect?," *Tourism Management*, 85, 104314, [10.1016/J.TOURMAN.2021.104314](https://doi.org/10.1016/J.TOURMAN.2021.104314).
- Acemoglu, Daron, Ali Makhdoumi, Azarakhsh Malekian, and Asuman E. Ozdaglar (2017) "Fast and Slow Learning From Reviews," *SSRN*, [10.3386/W24046](https://doi.org/10.3386/W24046).
- Aleksenko, Stepan and Jacob Kohlhepp (2023) *Consumer Reviews and Dynamic Price Signaling*: SSRN.
- Board, Simon and Moritz Meyer-ter Vehn (2013) "Reputation for Quality," *Econometrica*, 81 (6), 2381–2462, [10.3982/ECTA9039](https://doi.org/10.3982/ECTA9039).
- Carnehl, Christoph, Maximilian Schaefer, André Stenzel, and Kevin Ducbao Tran (2021a) "Value for Money and Selection: How Pricing Affects Airbnb Ratings."
- Carnehl, Christoph, André Stenzel, and Peter Schmidt (2021b) "Pricing for the Stars," *SSRN Electronic Journal*, [10.2139/SSRN.3305217](https://doi.org/10.2139/SSRN.3305217).
- Chevalier, Judith A and Dina Mayzlin (2006) "The effect of word of mouth on sales: On-line book reviews," *Journal of marketing research*, 43 (3), 345–354.
- Crapis, Davide, Bar Ifrach, Costis Maglaras, and Marco Scarsini (2017) "Monopoly pricing in the presence of social learning," *Management Science*, 63 (11), 3586–3608.
- Daley, Brendan and Brett Green (2014) "Market signaling with grades," *Journal of Economic Theory*, 151 (1), 114–145, [10.1016/J.JET.2013.10.009](https://doi.org/10.1016/J.JET.2013.10.009).
- Degan, Arianna and Ming Li (2021) "Persuasion with costly precision," *Economic Theory*, 72 (3), 869–908.
- Fudenberg, Drew and David K Levine (1989) "Reputation and equilibrium selection in games with a patient player," *Econometrica*, 57 (4), 759–778.
- Gertz, Christopher (2014) "Quality uncertainty with imperfect information acquisition," *Institute of Mathematical Economics Working Paper* (487).
- He, Qiao Chu and Ying Ju Chen (2018) "Dynamic pricing of electronic products with consumer reviews," *Omega (United Kingdom)*, 80, 123–134, [10.1016/J.OMEGA.2017.08.014](https://doi.org/10.1016/J.OMEGA.2017.08.014).
- Holmström, Bengt (1999) "Managerial Incentive Problems: A Dynamic Perspective," *The Review of Economic Studies*, 66 (1), 169–182, [10.1111/1467-937X.00083](https://doi.org/10.1111/1467-937X.00083).
- Hu, Nan, Paul A Pavlou, and Jie Zhang (2017) "On self-selection biases in online product reviews," *MIS quarterly*, 41 (2), 449–475.

- Hu, Nan, Jie Zhang, and Paul A Pavlou (2009) "Overcoming the J-shaped distribution of product reviews," *Communications of the ACM*, 52 (10), 144–147.
- Huang, Yangguang, Chenyang Li, and Si Zuo (2022) "Price Signaling and Reputation Building: Evidence from a Consulting Platform," *Available at SSRN 4080448*.
- Koh, Andrew and Ricky Li (2023) "Competition and Social Learning," *Available at SSRN*.
- Lafky, Jonathan (2014) "Why do people rate? Theory and evidence on online ratings," *Games and Economic Behavior*, 87, 554–570.
- Luca, Michael and Oren Reshef (2021) "The Effect of Price on Firm Reputation," *SSRN Electronic Journal*, [10.2139/SSRN.3625175](https://ssrn.com/abstract=3625175).
- Marinescu, Ioana, Andrew Chamberlain, Morgan Smart, and Nadav Klein (2021) "Incentives can reduce bias in online employer reviews.," *Journal of Experimental Psychology: Applied*.
- Martin, Simon and Sandro Shelegia (2021) "Underpromise and overdeliver?-Online product reviews and firm pricing," *International Journal of Industrial Organization*, 79, 102775.
- Milgrom, Paul and John Roberts (1986) "Price and advertising signals of product quality," *Journal of political economy*, 94 (4), 796–821.
- Pei, Harry (2020) "Reputation effects under interdependent values," *Econometrica*, 88 (5), 2175–2202.
- Rodríguez Barraquer, Tomás and Xu Tan (2022) "A model of competitive signaling with rich message spaces," *Review of Economic Design*, 1–43, [10.1007/S10058-021-00279-Y/FIGURES/2](https://doi.org/10.1007/S10058-021-00279-Y/FIGURES/2).
- Schoenmüller, Verena, Oded Netzer, and Florian Stahl (2019) "The extreme distribution of online reviews: Prevalence, drivers and implications," *Columbia Business School Research Paper* (18-10).
- Shapiro, Carl (1983) "Optimal pricing of experience goods," *The Bell Journal of Economics*, 497–507.
- Shin, Dongwook, Stefano Vaccari, and Assaf Zeevi (2021) "Dynamic Pricing with Online Reviews," *SSRN Electronic Journal*, [10.2139/SSRN.3439836](https://ssrn.com/abstract=3439836).
- Sorokin, D (2021) "Reaching for The Stars: Discounts and Review Tier Transitions in the Video Games Market," Technical report, working paper.

CHAPTER 2

Delegated Recruitment and Statistical Discrimination

with Jacob Kohlhepp (UNC Chapel Hill)¹

2.1 Introduction

The use of delegated recruitment has increased in the past few decades. Between 1991 and 2022, the fraction of American workers who found their job through a recruiter or headhunter rose from 4.9% to 14.3%, at the expense of direct applications (Black et al. (2022)). How does delegation impact the choice of candidates hired?

To answer this question, we develop a model where a firm (the principal) wishes to fill a position. A recruiter (the agent) can choose to suggest a candidate based on imperfect private information about the candidate's productivity. Candidates differ in their underlying productivity and the amount of information available about their productivity. If a candidate is suggested by a recruiter, the firm pays a hiring cost, observes the candidate's true productivity, and then chooses whether to terminate or retain the candidate. Given that employment is observable but the candidate's productivity and the recruiter's signals are not, all feasible contracts can be expressed as refund contracts, typical in the recruiting industry. Such contracts consist of a payment when a candidate suggested by the recruiter is hired by the firm, and a refund when a candidate is hired by the firm but terminated during an initial period of employment.

¹We are grateful for comments and suggestions from Alex Bloedel, Simon Board, Moritz Meyer-ter-Vehn, Tara Sinclair, and participants in the UCLA Theory Proseminar, UNC Theory Seminar, the SEA Annual Meeting, and SOLE Annual Conference. This project benefited from the support of the Institute for Humane Studies. We thank Hanqing Ye for excellent research assistance.

We demonstrate that refund contracts generate misalignment over whom to hire. Because refund contracts penalize recruiters for terminated candidates but do not reward recruiters for suggesting outstanding performers (“stars”), the recruiter prefers candidates whose productivity is likely to exceed the threshold. This leads to statistical discrimination in favor of more information when the contract features strong screening incentives: the recruiter requires a lower expected productivity threshold to suggest candidates whose resumes the recruiter understands better.

In contrast, in a first-best benchmark with direct recruitment, the firm has the option value of trying candidates with large residual uncertainty because of the firm’s ability to terminate poor performers and retain stars. This leads to group statistical discrimination in favor of less information: the firm requires a lower expected productivity threshold to suggest candidates whose resumes the firm understands less.

We further show that misalignment occurs when the firm designs an optimal contract under a tractable Pareto-uniform information structure. The unique equilibrium contract features strong screening incentives: the recruiter suggests only candidates who will certainly not be terminated. As a result, delegation reverses the direction of statistical discrimination compared to the first-best benchmark. Furthermore, low-expected-productivity candidates about whom the recruiter is better informed (“safe bets”) are inefficiently hired at the expense of high-expected-productivity candidates about whom the recruiter is less informed (“diamonds in the rough”). Thus, delegation via refund contracts systematically disadvantages groups of candidates. All groups with less informative signals than a cutoff are hired with a lower probability than in the first-best benchmark, while all groups with more informative signals than a cutoff are hired with a higher probability than in the first-best benchmark.

Information heterogeneity across candidates is a key driver of inefficiency in equilibrium. When the recruiter observes the same amount of information about all candidates, the first-best benchmark is achieved. Information heterogeneity can reflect both differences in a candidate’s profile (e.g., one candidate may have worked at a well-known company or may have obtained a college degree from a well-known university, while another may

have not) and differences in a recruiter's ability to interpret a candidate's profile (e.g., a recruiter may be better at interpreting the resumes of candidates that share the recruiter's race, gender, or socioeconomic backgrounds).² Under this interpretation, our results imply recruitment delegation favors the candidates who can afford to acquire precise but costly productivity signals and those who share demographic characteristics with the recruiter.

By considering extensions of the baseline model, we show that the qualitative results derived under the parametric information structure without commitment are robust. When the firm has a third payment it can use to extract all surplus from the agent, profit weakly increases but the first-best benchmark is not achieved even when information heterogeneity is small. When the recruiter can search for other candidates, the contract additionally serves to encourage unobserved effort, but different types of statistical discrimination between the first-best benchmark and equilibrium remain. When the firm can commit to a termination threshold, full extraction occurs and misalignment persists due to the fundamental tension between first- and second-stage screening.

Using a series of comparative statics in productivity and signal distributions, we show how delegation induces various spillovers across groups of candidates in equilibrium. Better information about some candidates makes screening unambiguously more valuable and therefore increases the hiring bar and decreases the probability of being hired for the rest of the candidates. However, better productivity for some candidates can either increase or decrease the hiring bar, depending on the quality of the information available for those candidates.

2.1.1 Literature

Our work relates to several strands of literature. Similar to many papers in the literature on delegation to an expert, our model features an agent with an information advantage making decisions for a principal. However, unlike in many papers in this literature

²For an example of racial homophily in hiring, see [Giuliano et al. \(2009\)](#).

(e.g., [Frankel \(2014\)](#), [Kundu and Nilssen \(2020\)](#), [Szalay \(2005\)](#)), in our paper the principal cannot constrain the action set of the agent directly and must exert control indirectly through monetary payments. Our paper shows refund contracts are a way to screen out recruiters with low-productivity candidates, similar to how debt contracts can be used to screen borrowers ([Gale and Hellwig \(1985\)](#)). Additionally, the agent's bias stems from not internalizing the principal's hiring cost, similar in spirit to [Che et al. \(2013\)](#) where the agent does not internalize the firm's outside option. Notably, [Che et al. \(2013\)](#) find a bias towards "conditionally better looking" projects, similar to how we find a bias towards "safe-bet" applicants.

Our paper is also related to the literature on delegated information acquisition. As in [Chade and Kovrijnykh \(2016\)](#), [Inderst and Ottaviani \(2012\)](#), and [Szalay \(2009\)](#), the principal must consider how the agent will act when faced with varying amounts of information. However, while in these papers the amount of information acquired is endogenous and homogeneous, in ours it is exogenous but heterogeneous. Moreover, the focus of the principal in this literature is to encourage information acquisition, while in our paper it is to achieve the best screening behavior on average across differently informed agents.

Our model features an agent who is privately informed about the candidates' expected productivity and the amount of information available about it. As a result, our paper is related to the literature on multidimensional screening ([Carroll \(2017\)](#), [Yang \(2021\)](#)). An essential aspect of our setting is that the agent is privately informed about the quality of information available. Even though productivity is the only object the firm cares about, the distribution of productivity across interim agent types differs in two dimensions. Additionally, the agent does not have intrinsic preferences over the dimensions being screened, as in problems of screening buyers. Instead, the way the agent engages with the contract offered by the firm is impacted by the agent's private type, and the firm exploits this in equilibrium to screen out recruiters with low-expected-productivity candidates. Another difference from this literature is that we do not allow the firm to offer a menu of contracts. This is because our model is application driven, and to our knowledge contract menus are not commonly observed in the recruiting industry.

Our paper differs from most past work in that the principal commits to a monetary contract, but this contract explicitly depends on an action taken by the principal after all information is revealed. There is a form of limited commitment built into the model that is inspired by the specific application to recruiters. One exception is [Levitt and Snyder \(1997\)](#) where, as in our model, an agent observes a private signal about a project's success and the principal can cancel the project based on the agent's interim advice. A key difference is that in [Levitt and Snyder \(1997\)](#) the contract is accepted before information is observed. The authors find that the principal's ability to influence the final contractible outcome undermines incentives, and commitment can help rectify the situation.

Our paper also contributes to the literature that studies what occurs when the hiring decision is delegated. [Frankel \(2021\)](#) considers delegation to a hiring manager, where the trade-off is between using the manager's soft information and indulging the manager's bias towards soft information relative to hard information. Two key differences are that in [Frankel \(2021\)](#) the firm has access to hard information and uses it to limit the manager's actions directly. [Cowgill and Perkowski \(2020\)](#) propose a theoretical framework and study empirically how recruiters select applicants. They show in a two-sided audit that recruiters over-interview candidates from elite schools and big companies, a fact the authors interpret as evidence of a reputational effect. Our paper proposes another interpretation: even conditional on the same expected productivity of candidates, recruiters are biased towards candidates from elite schools or big companies because of the way refund contracts are structured.

After presenting the model, the paper proceeds as follows. In [Section 2.3.1](#), we derive hiring thresholds in a first-best benchmark where the firm does not delegate but rather observes the candidate directly. In [Section 2.3.2](#), we derive the equilibrium contract. In [Section 2.3.3](#), we compare the set of hired candidates in the first-best benchmark and equilibrium and derive the main result that refund contracts induce misalignment and different types of statistical discrimination. In [Section 2.4](#), we discuss the forces that lead to the main result. Finally, we analyze several comparative statics in [Section 2.5](#), discuss the robustness of our results in [Section 2.6](#), and conclude the paper in [Section 2.7](#).

2.2 Model

Players and Actions. A risk-neutral firm wishes to hire one candidate, and a risk-neutral recruiter has one candidate with uncertain productivity (a). The firm proposes a contract to the recruiter. After receiving private information about the candidate's productivity, the recruiter either accepts the contract by suggesting the candidate or rejects the contract by not suggesting the candidate. If the candidate is suggested, the firm incurs a hiring cost (c), fully and privately observes the candidate's productivity, and then decides whether to retain or terminate the candidate.³ If the candidate is retained, the firm receives the candidate's productivity. Finally, all contract transfers are realized.

If the candidate is not suggested, both the firm and the recruiter receive their outside option. The outside option of the firm is 0, and that of the recruiter is $\bar{u} \geq 0$. The candidate is not a strategic player.

Candidates and Information. A candidate is characterized by a productivity a distributed according to a common prior with a CDF F_a and an information type (or candidate's group) $i \in \{1, \dots, N\}$ with probabilities $\{p_1, \dots, p_n\}$, which are drawn independently. Prior to deciding whether to suggest the candidate to the firm, the recruiter privately observes the candidate's i and a signal about productivity, $x \in \mathbb{R}^{\tau_i}$, distributed according to a CDF $G_i(\cdot|a)$ for the candidate with productivity a and information type i .⁴ With a slight abuse of notation, we call G_i an information structure for the productivity signal that the recruiter observes for a candidate from group i . Throughout, we index candidate information types in descending order of their informativeness in the Blackwell sense, i.e., G_i is Blackwell more informative than G_j for any $i < j$. To make these concepts concrete, we provide two parametric examples below.

Normal example. The prior distribution of the candidate's productivity is normal: $a \sim N(\mu_a, \sigma_a^2)$; formally, F_a is the CDF of a normal random variable with mean μ_a and variance

³After suggestion but before paying c , the firm does not observe any additional information about the candidate.

⁴The candidate's information type i (or group) is not payoff relevant and serves only as a part of information structure description and, hence, the productivity signal.

σ_a^2 . For each information type, the recruiter observes a single productivity signal ($\tau_i = 1$), which is the sum of true productivity and independent normal noise: $x = a + \varepsilon$, where $\varepsilon|a, i \sim N(0, \sigma_i^2)$. Using our notation, $G_i(\cdot|a)$ is a CDF of a normal random variable with mean a and variance σ_i^2 . Groups with lower i have higher signal precision ($\sigma_1^2 < \sigma_2^2 < \dots < \sigma_N^2$). Conditional on observing signal x for a candidate with information type i , the recruiter's posterior belief about productivity is $N\left(\frac{\sigma_i^2}{\sigma_a^2 \sigma_i^2} \cdot \mu_a + \frac{\sigma_a^2}{\sigma_a^2 \sigma_i^2} \cdot x, \frac{\sigma_a^2 \sigma_i^2}{\sigma_a^2 + \sigma_i^2}\right)$.

Pareto example. The prior distribution of the candidate's productivity is Pareto: $a \sim \text{Pareto}(\bar{a}, k)$. The recruiter observes a productivity signal consisting of τ_i unidimensional signals: $x = (x_1, \dots, x_{\tau_i})$. Groups with lower i have more signals: $\tau_1 > \tau_2 > \dots > \tau_N$. Conditional on productivity a , signals are drawn i.i.d. from a uniform distribution with minimum 0 and maximum a . Using our notation, $G_i(\cdot|a)$ is a CDF of a multivariate uniform distribution on $[0, a]^{\tau_i}$. Conditional on observing signal x for a candidate with information type i , the recruiter's posterior belief about productivity is $\text{Pareto}(\max\{\bar{a}, \{x_t\}_{t=1}^{\tau_i}\}, \tau_i + k)$.

Contracts and Payoffs. The productivity signal x is the recruiter's private information and is not contractible. The productivity a is privately observed only by the firm and is also not contractible. Even though the information type i is sometimes observable by both parties, we will frequently interpret it as representing demographic characteristics (race, gender, age, etc.). Because in many countries these characteristics are illegal to contract on, we assume they are also not contractible. This leaves only two possible contractible outcomes after the contract is accepted: whether the candidate is eventually retained or terminated.

Given the space of contractible outcomes, we can express all contracts as *refund contracts*, which consist of a transfer from the firm to the recruiter if the candidate is suggested and retained ($\alpha \in \mathbb{R}$) and a refund from the recruiter to the firm if the candidate is terminated ($\beta \in \mathbb{R}$). Under these contracts, the ex post profit of the firm is

$$\pi = \mathbf{1}_{\{\text{suggested}\}} \cdot \left(-c - \alpha + \mathbf{1}_{\{\text{retained}\}} \cdot a + \mathbf{1}_{\{\text{terminated}\}} \cdot \beta \right),$$

and the ex post utility of the recruiter is

$$u = \mathbf{1}_{\{suggested\}} \cdot \left(\alpha - \beta \cdot \mathbf{1}_{\{terminated\}} \right).$$

In Section 2.2.1, we argue that these contracts are commonly observed in the recruiting industry. We specify that the recruiter chooses to accept the contract if indifferent.

We consider two versions of the model in this paper. **A.** In the commitment case, the firm can commit to a termination threshold γ (for all candidates regardless of their information type i), include it as a part of the contract, and then terminate the candidate if and only if the candidate's realized productivity a is below γ . The contracts cannot depend on the realized productivity in any other way apart from the termination rule.⁵ **B.** In the no-commitment case, the firm cannot commit to the termination rule and decides whether to terminate the candidate in the last stage of the game. Thus, the firm prefers to terminate the candidate if and only if the candidate's productivity a is below the refund β regardless of the candidate's information type i (thus, the ex post compatible termination threshold in the no-commitment case is $\gamma = \beta$).

Equilibrium. The equilibrium concept we use is weak Perfect Bayesian Equilibrium with the assumption that the recruiter has passive beliefs about the productivity of the candidate in the contract offered by the firm. We require passive beliefs in order to have well-defined beliefs under off-path contracts.⁶ We assume that when indifferent, the recruiter suggests the candidate. Finally, we assume that there exists a feasible interior suggestion

⁵If the contract could freely depend on the productivity, the firm would prefer to lie about its realization conditional on retaining or terminating the candidate, in order to increase its payoff.

⁶Without passive beliefs, the recruiter does not have to be Bayesian when interpreting the signals under off-path contracts because those information sets are of zero probability. Thus, for any realization of signals for any off-path contract, the recruiter can be extremely pessimistic about the candidate's productivity and suggest no one. Therefore, any contract can be an equilibrium one under wPBE. To fix that, we use passive beliefs and force the recruiter to make a Bayesian inference from the productivity signals in every information set (including those off-path).

strategy⁷ which delivers strictly greater total surplus than always or never suggesting a candidate.

2.2.1 Model Comments

Refund Contracts. Refund contracts are a natural formalization of the contract space in our model. They closely resemble what is called guarantee contracts, which are the main form of contract used by external recruiters, according to evidence from qualitative interviews, industry materials, and surveys. We interviewed one mid-career headhunter and one early-career recruiter. Both stated they were paid if a candidate they suggested was placed, but they had to provide a refund or a free replacement if the candidate left the firm in the first 90 days of employment. This refund is called in the industry a “guarantee.” The early-career recruiter confirmed that the refund was given for any reason, including termination of the candidate by the company. The relevant sections of the interview transcript are provided in Appendix Section 2.8.6.

Recruiters that use this structure of compensation are called contingent recruiters, and they represent the majority of recruiters, according to estimates (Finlay and Coverdill (2007)). The contingent compensation structure with a guarantee discussed by the recruiters we interviewed is also mentioned in a variety of sources, including a report on recruiting practices in the hospitality industry (Dingman (1993)), how-to books about starting an executive recruiting firm (Press (2007), Perry and Haluska (2017)), a guide for lawyers working with headhunters (Steinberg and Machlowitz (1989)), a guide for managing financial service companies (Arslanian (2016)), and an academic article (Florea (2014)). Further, the American Staffing Association provides a “Model Recruiting Agreement” which includes sample language for refund and replacement guarantees (Association (2014)).

⁷Feasible by the firm via refund contracts, and “interior” meaning the probability of suggestion is not 0 or 1.

A survey of recruiters by Top Echelon found that 96% of recruiters offered some form of guarantee. The most popular guarantee time frame was 90 days, consistent with our interviews. When asked about the form of the refund, 61% responded that they offered a replacement, and 26% responded that they offered money back (Deutsch (2019)).⁸

Information Types. The information type of the candidate may in some cases be observed by the firm, but it is unnatural and often illegal to condition on it in a contract. For example, the information type of a candidate may be determined by the candidate's demographics (age, gender, race, etc.). In the U.S. it is illegal to recruit based on such characteristics, much less write them explicitly in a contract.⁹ In our model, the information type by itself is uninformative about the candidate's productivity. Alternatively, the information heterogeneity can also come from differences in the candidates' experiences or certifications, which might be unobserved by the firm and therefore cannot be included in the contract.

Firm and Recruiter. The hiring cost can be interpreted as the cost to the firm of interviewing the candidate¹⁰ or the cost of employing the candidate for a probationary period. Within the model, the hiring cost operates as the cost to the firm of fully learning the candidate's productivity. This is an important ingredient in the model because it makes the recruiter's private information valuable to the firm. The firm's main goal is to use the recruiter's private information to screen candidates. Because the recruiter does not bear any intrinsic cost of suggesting a candidate, the recruiter has a natural tendency not to screen candidates and instead suggest everyone. For this reason, the contract must be designed to reduce this tendency.

It may appear that the only reason the firm uses a recruiter is that the recruiter has a candidate and the firm does not. However, by setting the payment from suggestion to be \bar{u} and

⁸Of the remaining 13%, 11% responded that they offered a guarantee that did not fit into the two aforementioned categories, and 2% gave no response.

⁹Per EEOC (2023): "It is also illegal for an employer to recruit new employees in a way that discriminates against them because of their race, color, religion, sex (including gender identity, sexual orientation, and pregnancy), national origin, age (40 or older), disability or genetic information."

¹⁰For example, the cost of having current employees conduct on-site interviews.

setting $\beta = 0$, the firm can always design a contract where the recruiter suggests everyone (i.e., no screening). Thus, if the firm had an outside option of paying the hiring cost c to learn the productivity of a random candidate or using a recruiter, it will always weakly prefer to use the recruiter. Recruiters are useful to the firm because of their information advantage prior to hiring, not because they have a monopoly on candidates.

2.3 Nonparametric Analysis

In this section, we define notions of statistical discrimination and show that discrimination goes in opposite directions in the first-best benchmark and equilibrium under a wide set of contracts. We begin by analyzing the space of posteriors about the candidate. Second, we define the first-best benchmark and equilibrium and derive the conditions under which the candidate is suggested. Third, we introduce directed notions of individual and group statistical discrimination and provide conditions under which the direction of statistical discrimination is different in the first-best benchmark and equilibrium. We finally show what forces distort the equilibrium outcome relative to the first-best benchmark.

2.3.1 Candidate Posteriors

Characterizing the recruiter's (or firm's) optimal decisions for each possible recruiter's information set is the same as characterizing the optimal suggestion decisions over the space of the candidate's group and productivity signal realizations. We start by parameterizing this space.

For any information type i and any realized signal x , we can derive the posterior productivity distribution $F_{i,x}(a)$. We assume that each group's information structure is such that no two signal realizations have the same posterior μ . This assumption allows us to uniquely map every posterior distribution of productivity $F_{i,x}$ to its mean μ and the information type of the candidate i , and therefore to parameterize the set of all candidates'

productivity posterior distributions by i and μ ($F_{i,x} \leftrightarrow F_{i,\mu}$). Further, we describe the candidates upon observing their productivity signal and group by (i, μ) .

To define intuitive risk attitudes for the firm and the recruiter, we require that the space of candidates satisfy the following assumption.

Assumption

- (a) $F_{i,\mu_1}(a)$ first-order stochastically dominates $F_{i,\mu_2}(a)$ for any $(\mu_1 > \mu_2, i)$,
- (b) $F_{i,\mu}(a)$ second-order stochastically dominates $F_{j,\mu}(a)$ for any $(\mu, i < j)$, and
- (c) $\forall \mu, i < j \exists a^* \geq F_{j,\mu}^{-1}(1/2)$, s.t. $\forall a < (>)a^* : F_{i,\mu}(a) \leq (\geq) F_{j,\mu}(a)$.

Part (a) of this assumption means that there is an unambiguous order of signal realizations within an information type, where a “better” signal means that the candidate’s posterior productivity distribution first-order dominates all posteriors with “worse” signals (conditional on the information type of the candidate). Parts (b) and (c) of this assumption require a stronger order of information structures than the Blackwell order. Besides requiring that the signal for a group j be a garbling of the signal for a group $i < j$, it requires that each posterior $F_{j,\mu}$ be a single-mean-preserving spread of $F_{i,\mu}$ and have a fatter lower tail than $F_{i,\mu}$.

Many common prior distributions and information structures satisfy this assumption, including (1) normal prior distribution with normal signals of various precision (*parametric example A*); (2) the Pareto prior distribution with various numbers of uniform signals on $[0, a]$ (*parametric example B*); and (3) any information structure with posteriors that can be expressed as $\mu + \sigma_i \varepsilon$, where ε is a symmetric mean-zero random variable.

2.3.2 The First-Best, Equilibrium, and Suggestion Decisions

First-Best. We define the first-best benchmark as a hypothetical situation where the firm possesses the recruiter’s private information about the candidate and makes the hiring decision directly based on this information.

In the first-best benchmark (shortened to “first-best”), there is no need for the firm to worry about the payments to the recruiter when making the termination decision. The firm terminates the candidate only if $a < 0$ because the hiring cost is already paid. The firm’s ex post value from a candidate with known productivity a is $\max\{a, 0\}$. The firm’s expected value from hiring a candidate with information type i and expected productivity μ is equal to $\mathbb{E}[\max\{a, 0\}|i, \mu]$. Thus the firm suggests (and hires) a candidate if and only if: $\mathbb{E}[\max\{a, 0\}|i, \mu] \geq c$.

The first-order stochastic order in μ implies monotonicity of posterior truncated expectations in μ . We denote the lowest-posterior-expectation candidate from group i whom the firm suggests in the first-best as $\mu^*(i) = \inf_m \{m | \mathbb{E}[\max\{a, 0\}|i, \mu = m] \geq c\}$. Then the firm suggests a candidate (i, μ) if and only if $\mu \geq \mu^*(i)$. We focus our attention on non-trivial cases where the firm prefers to screen out a positive share of candidates in the first-best.¹¹

Equilibrium. We begin our analysis of equilibrium by considering the firm’s decision to terminate or retain the candidate. In the commitment case, the firm terminates the candidate if and only if the realized productivity a is below the termination threshold γ . In the no-commitment case, because the firm cannot commit to a termination rule, the decision is made after the contract is designed, hiring costs are sunk, and the firm has fully learned the productivity of the candidate. Thus the firm retains the candidate if the realized productivity exceeds the refund (β) that could be obtained from termination: $a \geq \beta$. This is the firm’s incentive-compatible retention rule (IC). Thus, the termination threshold is given by $\gamma = \beta$ in the no-commitment case. For now, we proceed with a general γ (not necessarily IC).

Anticipating the firm’s termination or retention decision, the recruiter has an ex post value from suggesting a candidate with known productivity a that is $\alpha - \beta \mathbb{I}_{a < \gamma}$. The recruiter’s expected value from suggesting a candidate with information type i and expected productivity μ is equal to $\alpha - \beta \cdot \Pr(a < \gamma | i, \mu)$. Thus the recruiter suggests (and the firm hires) a candidate if and only if the interim value exceeds the recruiter’s out-

¹¹If the firm prefers not to screen anyone out in the first-best, then the solution is trivial and is implementable in equilibrium.

side option: $\alpha - \beta \cdot \Pr(a < \gamma|i, \mu) \geq \bar{u}$. The recruiter's decision to suggest a candidate in equilibrium is therefore fully determined by whether a given candidate's termination probability is above or below a threshold: $\Pr(a < \gamma|i, \mu) < p^*$, where $p^* = \frac{\alpha - \bar{u}}{\beta}$.

With these results established, and recalling that our assumptions imply monotonicity of $\Pr(a < \gamma|i, \mu)$ in μ , we denote the lowest-posterior-expectation candidate from group i whom the recruiter suggests under a refund contract with parameters (α, β, γ) as $\mu_{\alpha, \beta, \gamma}(i) = \inf_m \{m | \Pr(a < \gamma|i, \mu = m) < \frac{\alpha - \bar{u}}{\beta}\}$. Then the recruiter suggests a candidate (i, μ) if and only if $\mu \geq \mu_{\alpha, \beta, \gamma}(i)$. We focus on non-degenerate cases where the firm's optimal contract induces some screening in equilibrium.¹²

2.3.3 Individual and Group Statistical Discrimination

In this section, we define individual and group statistical discrimination and show how delegation to the recruiter can generate both. We demonstrate that while the firm is endogenously risk-loving in the first-best, refund contracts often make the recruiter endogenously risk-averse.

We define individual discrimination based on how two candidates with the same expected productivity (μ) but different information types are treated.

Definition 4 *A decision maker engages in individual statistical discrimination in favor of more (less) information for candidates with expected productivity μ and information types $i < j$ if the candidate with the more (less) informative type i (j) is always suggested whenever the candidate with the less (more) informative type j (i) is suggested.*

To make this definition concrete, suppose there are two candidates (label them 1 and 2) from different information types (say, 1 and 2) but identical posterior productivity means: $(1, \mu)$ and $(2, \mu)$. Because candidate 1 comes from a lower information type, we are more informed about candidate 1. Specifically, $F_{2, \mu}(a)$ is a single-mean-preserving spread of

¹²If the firm prefers not to screen anyone out in the equilibrium, then the solution is trivial and leads to the recruiter suggesting all candidates and the firm retaining those with productivity above zero.

$F_{1,\mu}(a)$. If candidate 1 being suggested implies that candidate 2 is suggested but not the other way around, then there is individual statistical discrimination in favor of less information. Intuitively, candidate 1 gives a higher value to the decision maker than does candidate 2. The next result establishes that individual statistical discrimination is quite different in the first-best compared to delegation via refund contracts.

Proposition 3

1. *In the first-best, the firm always engages in individual statistical discrimination in favor of less information (for all μ, i, j).*
2. *In equilibrium, the recruiter engages in individual statistical discrimination in favor of more (less) information if $\gamma \leq (\geq) a^*$, where $a^* : F_{i,\mu}(a^*) = F_{j,\mu}(a^*)$.*

Proof. The first-best result follows directly from the fact that, due to the firm's ability to terminate candidates, the firm's value from a candidate in the first-best is a convex function of realized productivity ($\max\{0, a\}$). The refund contract result follows from the fact that information types are single-mean-preserving spreads and under refund contracts the recruiter value from a candidate depends only on the probability that a candidate is below a threshold. This value is exactly the CDF, so which candidate the recruiter prefers depends only on the termination threshold relative to the single-crossing point a^* . ■

Proposition 3 captures the key tension which drives misalignment. Absent delegation, the firm is endogenously risk-loving; it appreciates the option value of uncertain candidates because of its ability to terminate candidates. However, refund contracts frequently cause the recruiter to be endogenously risk-averse. Specifically, the recruiter receives no additional reward for suggesting candidates with outstanding productivity but is penalized for suggesting candidates whose productivity falls below a threshold.

Whether the recruiter statistically discriminates in one direction or the other about a particular pair of candidates depends on the termination threshold. When the termination threshold γ is low (relative to the single crossing-point for that particular pair), the recruiter statistically discriminates in favor of candidates that the recruiter understands better. When the firing threshold γ is high (relative to the single crossing-point for that

particular pair) the recruiter statistically discriminates in favor of candidates the recruiter understands less. This situation happens precisely because the recruiter's goal is to minimize the probability of termination for the suggested candidate.

The relation between the termination threshold (γ) and the single-crossing point a^* determines the amount of screening in the second stage of the hiring process (after the candidate was suggested). If γ is high (low) relative to a^* , a candidate is more (less) likely to be terminated (screened out) after being suggested by the recruiter. With little second-stage screening, the recruiter prefers less noise and more certainty in the candidate's productivity to make sure that it likely exceeds the threshold γ . With a lot of second-stage screening and many candidates that are unlikely to make the cut, the recruiter prefers more noise in the candidate's productivity, to have at least some chance of a high realization exceeding the threshold γ .

This result about individual discrimination is incomplete. The misalignment between the firm and the recruiter over any particular candidate depends on the termination threshold γ relative to the single-crossing point a^* . When this threshold is low, there is misalignment; when it is high, there is not. Because the single-crossing point between any two candidates (a^*) depends on the expectation and information type of those candidates, it is unclear how this individual alignment or misalignment of discrimination, which differs across specific pairs of candidates, translates into systemic differential treatment of groups of candidates.

To address these concerns, we introduce the notion of group discrimination. Recall that $\mu^*(i)$ is the lowest-posterior-expectation candidate from group i that is suggested in the first-best, while $\mu_{\alpha,\beta,\gamma}(i)$ is the lowest-posterior-expectation candidate from group i that is suggested by the recruiter under a refund contract with parameters (α, β, γ) .

Definition 5 *A decision maker engages in group statistical discrimination in favor of more (less) information in the first-best if $\mu^*(i)$ is increasing (decreasing) in i , and in equilibrium under refund contract (α, β, γ) if $\mu_{\alpha,\beta,\gamma}(i)$ is increasing (decreasing) in i .*

Group statistical discrimination captures whether the “hiring bar” differs across groups of candidates. The next result characterizes when differential group statistical discrimination occurs across the first-best and equilibrium.

Proposition 4

1. In the first-best, the firm always engages in group statistical discrimination in favor of less information.
2. In equilibrium, the recruiter engages in group statistical discrimination in favor of more information if $\frac{\alpha - \bar{u}}{\beta} < \frac{1}{2}$.

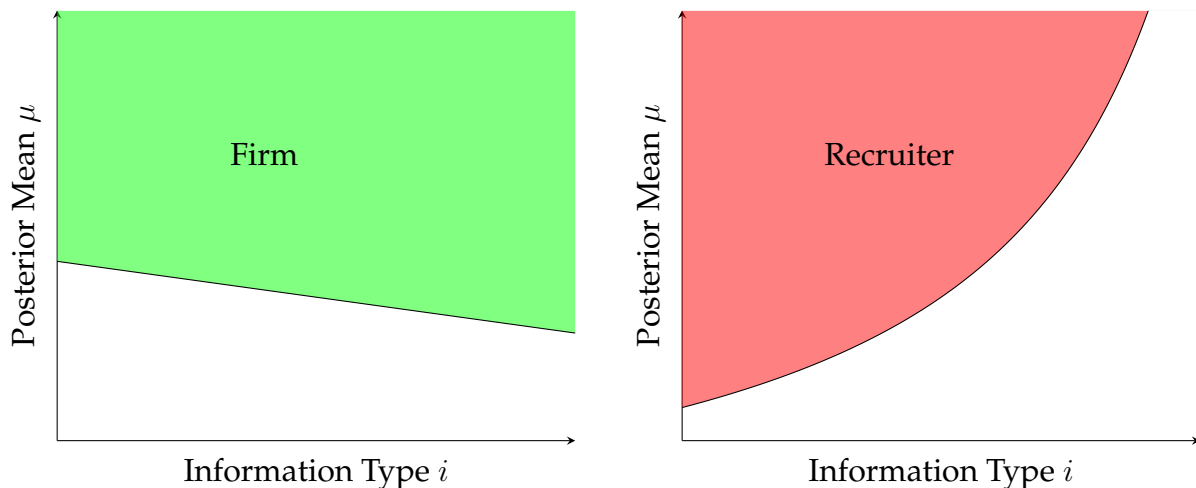
Proof. The fact that the first-best favors groups of candidates with less information follows directly from the previous proposition. For equilibrium, consider that the recruiter suggests any candidate when the recruiter’s payoff from suggestion exceeds the outside option:

$$\alpha - \beta Pr(a < \gamma | i, \mu) \geq \bar{u} \leftrightarrow \frac{\alpha - \bar{u}}{\beta} \geq Pr(a < \gamma | i, \mu).$$

The threshold $\mu(i)$ is given by the points where this inequality holds with equality. When $\frac{\alpha - \bar{u}}{\beta} < \frac{1}{2}$, all candidates with less informative types will have a higher threshold $\mu(i) \leq \mu(j)$ due to single-crossing. ■

Figure 2.1 illustrates Proposition 4 for $\frac{\alpha - \bar{u}}{\beta} < \frac{1}{2}$.

Figure 2.1: Candidate Types Hired in Equilibrium vs. First-Best



Proposition 4 shows that under a range of contracts, there are opposite types of statistical discrimination not only across individual pairs of candidates but also across groups of candidates. The firm sets a lower hiring bar for groups the firm understands less, while the recruiter sets a lower hiring bar for groups the recruiter understands more. Whenever this result holds, some groups will be disadvantaged in terms of hiring by delegated recruitment relative to direct recruitment.

The second part of Proposition 4 specifies that misalignment occurs under contracts where $\frac{\alpha - \bar{u}}{\beta} < \frac{1}{2}$. One way to interpret this condition is that the recruiter offers an economically significant refund. Notice that β is exactly the part of the contract that induces first-stage screening, i.e., screening in the suggestion decision. When $\beta = 0$, the recruiter suggests all candidates. When $\beta \rightarrow \infty$, the recruiter suggests only candidates that are retained with certainty. If β is high and therefore $\frac{\alpha - \bar{u}}{\beta}$ is low, many candidates are screened out in the first stage by the recruiter, and few of them are terminated after the probation period. Similar to individual discrimination, the key condition for misalignment with respect to group discrimination is little second-stage screening implied by the strong first-stage screening incentives.

An important observation from this condition is that the termination threshold does not directly determine whether group statistical discrimination in favor of more information occurs in equilibrium. However, the termination threshold will impact the way the firm designs the other contract parameters, so it will still play an indirect role. For this reason, in the next few sections, we analyze how the equilibrium contract is designed.

2.3.4 Contract Design in Equilibrium

The analysis up to this point treats the features of the contract (α, β, γ) as exogenous parameters. In this section, we discuss the key forces which determine how the contract is designed in equilibrium. We will also discuss the tensions between these forces that arise due to information heterogeneity within the candidate pool.

Corollary 1.

1. $\mu_{\alpha,\beta,\gamma}(i)$ is decreasing in α (with or without commitment),
2. $\mu_{\alpha,\beta,\gamma}(i)$ is increasing in β (with or without commitment),
3. $\mu_{\alpha,\beta,\gamma}(i)$ is increasing in γ (with commitment).

From the contract designer's (the firm's) perspective, the contract serves three purposes. We first consider the case where the firm can commit to a termination threshold (γ is a parameter of the contract). We conclude the section discussing the differences that arise without commitment ($\gamma = \beta$).

First-Stage Screening. The contract determines the set of candidates the recruiter suggests to the firm. The firm wishes to design the contract so that the equilibrium set of suggested candidates is close to the first-best set of suggested candidates. By increasing the suggestion payment α , the firm expands the set of candidates suggested across all information types. By increasing the refund β or the termination threshold γ , the firm shrinks the set of candidates suggested across all information types.

Replicating the first-best is often not possible because of the coarseness of refund contracts. They give the firm only three choice variables with which to control the suggestion or hiring region. This leads to a mechanical impossibility of achieving the first-best for more than three information types with commitment or more than two information types without commitment. Although we numerically verify that the first-best suggestion region is feasible with two information types, the firm does not choose to implement the first-best region in this case.

Second-Stage Screening. After hiring a candidate and paying the cost c , the firm fully learns the candidate's productivity and makes a termination decision (either mechanically with commitment or endogenously without). This can be viewed as a perfect but costly second stage of screening. The firm wishes to design the contract to minimize inefficient terminations: candidates with positive productivity ($a \geq 0$) that are hired at cost c and then terminated. Holding fixed the set of suggested candidates, reducing the positive termination threshold γ reduces these inefficient terminations.

Since the termination threshold γ affects both the hiring region (first-stage screening) and the termination threshold (second-stage screening), there is tension when choosing the optimal two-stage screening procedure. Raising the termination threshold improves first-stage screening but requires the firm to inefficiently terminate candidates after they are hired.

Surplus Extraction. Because the recruiter selects into the contract after observing signals about the candidate, the recruiter has information rents. One goal of the contract is to extract surplus from the recruiter and transfer it to the firm. This goal is generally in tension with the first-stage screening goal. Intuitively, trying to approximate the first-best suggestion region requires setting $\beta > 0$ and $\alpha > \bar{u}$, which means that the recruiter gets the outside option for the candidates on the boundary of the suggestion region but gets a strictly greater expected utility for all other suggested candidates. As we show in Section 2.6, although many specifications of the model allow the firm to extract all surplus, the first-best is still not achieved.

Commitment. When the firm can commit to a termination rule, γ is an additional degree of freedom that can be used to navigate the tensions outlined previously. However, when the firm cannot commit, $\gamma = \beta$ because the firm terminates low-productivity candidates in order to obtain the refund. In this case several equilibrium forces compete to determine β .

Recall that β relative to $\alpha - \bar{u}$ is a measure of the recruiter's incentive to screen candidates. With commitment, as β rises, the recruiter engages in more intense screening. Without commitment, increasing β has two additional impacts. First, it amplifies screening incentives because it increases not just the refund the recruiter has to pay but also the probability the recruiter has to pay the refund. Second, increasing β shifts the firm's payoff from individual candidates because the firm's ex post value function is $\max\{a, \beta\}$.

2.4 Parametric Analysis without Commitment

In this section, we make a parametric assumption under which there is a unique equilibrium contract without commitment, which we obtain in closed-form. We illustrate that the misalignment discussed in Section 2.3 occurs even though the firm designs the equilibrium contract to maximize total surplus.

Assumption 3 *The prior distribution of the candidate's productivity is $a \sim \text{Pareto}(\bar{a}, k)$. Conditional on productivity a , signals are drawn i.i.d. from a uniform distribution with minimum 0 and maximum a . Information type i has τ_i signals: $x = (x_1, \dots, x_{\tau_i})$, where $\tau_1 > \tau_2 > \dots > \tau_N$.*

Conditional on observing signal x for a candidate with information type i , the recruiter's posterior belief about productivity is $\text{Pareto}(\max\{\bar{a}, \{x_t\}_{t=1}^{\tau_i}\}, \tau_i + k)$.¹³ Under this information structure, the recruiter's posterior belief about a candidate depends only on the maximum of the observed signals and the number of signals. This information structure satisfies our nonparametric ordering assumption in Section 2.3.

The candidate's information type matters only in determining the number of signals observed about the candidate (τ_i). For this reason, we occasionally suppress the information type subscript i and consider an arbitrary information type with τ signals. The posterior distribution depends only on the maximum of observed signals and \bar{a} , an object we denote by $x_{max}^\tau := \max\{\bar{a}, \{x_t\}_{t=1}^\tau\}$. Whenever we observe a maximum signal of x_{max}^τ , we know for sure that the candidate's productivity is at least x_{max}^τ .

Given x_{max}^τ , the posterior mean productivity is equal to $\mathbb{E}[a|x_{max}^\tau] = \frac{\tau+k}{\tau+k-1} \cdot x_{max}^\tau$. For the same maximum signal x_{max}^τ , the posterior distribution keeps the same minimum (x_{max}^τ) regardless of τ , but the expected productivity decreases in τ via the shape parameter of the posterior distribution ($\tau + k$). Intuitively, if we observe more productivity signals with the same overall maximum, we become more pessimistic about the candidate's productivity.

¹³It is well-known in the statistics literature that a Pareto prior and uniform signals are a conjugate family, meaning beliefs remain Pareto after updating (Fink (1997)). For completeness, we provide a proof in Appendix Section 2.8.1.

2.4.1 First-Best Suggestion Region

Recall that $\mu^*(i)$ is the lowest-posterior-expectation candidate from group i that the firm suggests in the first-best. Under this information structure we have that

$$\mu^*(i) = \inf_m \{m | \mathbb{E}[\max\{a, 0\} | i, \mu = m] \geq c\} = c.$$

The firm suggests and hires only candidates it expects to be worth the hiring cost c . Thus the hiring bar in terms of expected productivity is the same across all information types, and the firm is endogenously risk-neutral. In terms of group-based statistical discrimination, the firm is also neutral, technically engaging in both types of statistical discrimination.

2.4.2 The Equilibrium Contract

Recall that $\mu_{\alpha, \beta, \gamma}(i)$ is the lowest-posterior-expectation candidate from group i whom the recruiter suggests under a refund contract with parameters (α, β, γ) . Under this information structure we have that

$$\mu_{\alpha, \beta, \gamma}(i) = \inf_m \left\{ m \left| \Pr(a < \gamma | i, \mu = m) < \frac{\alpha - \bar{u}}{\beta} \right. \right\} = \gamma \frac{\tau_i + k}{\tau_i + k - 1} \left(1 - \frac{\alpha - \bar{u}}{\beta} \right)^{\frac{1}{\tau_i + k}}.$$

Without commitment, the firm terminates all candidates with productivity below $(\gamma = \beta)$, and we have that

$$\mu_{\alpha, \beta}(i) = \beta \frac{\tau_i + k}{\tau_i + k - 1} \left(1 - \frac{\alpha - \bar{u}}{\beta} \right)^{\frac{1}{\tau_i + k}}.$$

This expression shows that as the termination threshold (γ) rises, the suggestion threshold of the recruiter is scaled up for all information types. In general, the upfront fee must exceed to outside option $(\alpha > \bar{u})$, so as the refund β rises, the suggestion threshold rises.

Notice that $\frac{\tau_i + k}{\tau_i + k - 1}$ is increasing in i , while $\left(1 - \frac{\alpha - \bar{u}}{\beta} \right)^{\frac{1}{\tau_i + k}}$ is weakly decreasing in i . As a result, the direction of group-based statistical discrimination is ambiguous: $\mu_{\alpha, \beta}(i)$ can be either upward sloping or downward sloping in information type i . In the first-best, $\mu(i)$

is constant in i . This is the tension between first-stage screening, inefficient terminations, and surplus extraction outlined in Section 2.3.4. If the firm wants to approximate the shape of the first-best suggestion region, the firm must set $\alpha > \bar{u}$. However, this requires both leaving surplus for the recruiter and hiring candidates who are terminated with no realized benefit to the firm. There is a uniquely optimal way to navigate this trade-off.

Theorem 2 *The unique equilibrium contract is a refund contract with*

$$\beta^* = \frac{\mathbb{E}_\tau \left[\frac{\tau}{\tau+k} \right]}{\mathbb{E}_\tau \left[\frac{\tau}{\tau+k-1} \right]} c, \quad \alpha^* = \bar{u}.$$

A candidate is suggested if and only if the candidate's maximum signal exceeds β^ , that is, when the candidate's posterior expected productivity exceeds $\mu_{\alpha^*, \beta^*}(i) := \frac{\tau_i+k}{\tau_i+k-1} \beta^*$.*

Proof of Theorem 2.

- **Transfers are positive.** When specifying the contract space, we did not restrict the sign of the refund or the suggestion payment. We now show that in equilibrium, the transfer from the firm to the recruiter for suggestion is at least the outside option ($\alpha \geq \bar{u}$) and the transfer from the recruiter back to the firm is weakly positive ($\beta \geq 0$). Suppose for the sake of contradiction that in an optimal contract $\beta < 0$. Because the firm cannot commit to a termination rule, the firm terminates candidates only when $a \geq \beta$. This means all candidates are not terminated if they were suggested by a recruiter. Anticipating this situation, the recruiter either suggests all candidates if $\alpha \geq \bar{u}$, yielding negative profit, or suggests no candidates if $\alpha \leq \bar{u}$, yielding at most 0 profit. The firm can obtain strictly positive profit by offering an alternative contract $\alpha' = \bar{u}, \beta' = c$, which leaves the recruiter with 0 surplus and induces the recruiter to suggest all candidates who have productivity above the hiring cost. This contradicts optimality.

Suppose for the sake of contradiction that in an optimal contract $\alpha < \bar{u}$. Because $\beta \geq 0$, even if the recruiter knows for sure the candidate will not be terminated,

the recruiter will not suggest the candidate, because $\alpha - \beta \cdot 0 < \bar{u}$. Therefore, no candidates are suggested, yielding 0 profit for the firm. The firm can obtain strictly positive profit by offering the alternative contract $\alpha' = \bar{u}, \beta' = c$. This contradicts optimality.

- **Full surplus extraction.** We now argue that any equilibrium contract extracts all surplus from the recruiter, by setting $\alpha = \bar{u}$. In Appendix Section 2.8.2 we show that profit is strictly decreasing in α for all β . Therefore, $\alpha = \bar{u}$ in any profit-maximizing contract. When $\alpha = \bar{u}$,

$$\mu_{\alpha,\beta}(i) = \beta \frac{\tau_i + k}{\tau_i + k - 1}.$$

Note that the posterior mean maps to the maximum of τ signals in the following way: $\mu = \frac{\tau+k}{\tau+k-1}x_{max}^\tau$. Thus we have that the minimum suggested posterior expectation implies a minimum suggested maximum signal that is the same across all information types (x_{EQ}):

$$x_{EQ} = \frac{\tau_i + k - 1}{\tau_i + k} \mu_{\alpha,\beta}(i) = \beta \frac{\tau_i + k}{\tau_i + k - 1} \frac{\tau_i + k - 1}{\tau_i + k} = \beta.$$

Since only candidates with a maximum signal above x_{EQ} are suggested, this has the additional implication that only candidates who are retained for sure are suggested. Thus the recruiter's payoff is \bar{u} , all surplus is extracted, and the firm maximizes total surplus:

$$\mathbb{E}[\mathbb{I}\{x_{max}^\tau \geq \beta\}(\mathbb{I}\{a \geq \beta\}a + \mathbb{I}\{a \leq \beta\}\beta - c - \alpha)] = \mathbb{E}[\mathbb{I}\{x_{max}^\tau \geq x_{EQ}\}(a - c - \bar{u})]. \quad (2.1)$$

- **Optimal payments.** To complete the proof, we derive the unique refund payment (β^*) which maximizes (2.1). We do this by solving the first-order condition for β when $\alpha = \bar{u}$. We then show the second-order condition is satisfied at β^* . Appendix Section 2.8.2 contains these calculations. ■

All three forces from Section 2.3.4 (first stage screening, second stage screening and surplus extraction) play a role in determining the unique equilibrium contract. By setting β^* to be the ratio of two expectations, the firm approximates the first-best suggestion region as best it can given the other two goals and the coarseness of the contract. By setting $\alpha = \bar{u}, \beta > 0$ the firm simultaneously extracts all surplus from the recruiter and eliminates inefficient terminations.

Strong screening incentives are key to full surplus extraction and elimination of inefficient terminations. The suggestion payment is exactly equal to the recruiter's outside option, and the refund is economically meaningful (close to the hiring cost). The recruiter suggests only candidates the recruiter is certain will not be terminated. Because screening incentives are strong, the condition for group statistical discrimination in favor of more information is satisfied:

$$\frac{\alpha^* - \bar{u}}{\beta^*} = 0 < \frac{1}{2}.$$

Corollary 2.1 *The firm engages in group statistical discrimination in favor of less information in the first-best, and the recruiter engages in group statistical discrimination in favor of more information under the unique equilibrium contract.*

Under the unique equilibrium contract, the directions of statistical discrimination are misaligned in the first-best and equilibrium. In the first-best, the firm suggests and hires all candidates with posterior expectations above the hiring cost. In equilibrium, the recruiter suggests and the firm hires candidates with posterior expectations that exceed group-specific hiring bars. The bar is higher for groups that are less well-understood by the recruiter.

Beyond refund contracts themselves, a pivotal ingredient in preventing achievement of the first-best is information heterogeneity across candidates. To see this, note that when there is no information heterogeneity (i.e., there is only one group with τ signals), we can drop the expectation operators on the numerator and denominator of β^* , yielding

$$\beta^* = \frac{\tau+k-1}{\tau+k}c:$$

$$\mu_{\alpha^*, \beta^*}(i) = \beta^* \frac{\tau+k}{\tau+k-1} \frac{\tau+k-1}{\tau+k} c = c.$$

Thus when the pool of candidates is homogeneous in terms of information quality, the recruiter suggests only candidates with posterior expectations above the hiring cost, and the first-best is achieved.¹⁴

2.4.3 Hiring Outcomes by Group

Because the Pareto-uniform information structure without commitment yields a unique closed-form equilibrium contract, we can make stronger statements about how delegation impacts candidate groups differently. We present and discuss these results in this section, after defining the following cutoff.

Definition 6 *The cutoff number of signals τ^* is such that*

$$\frac{\tau^* + k - 1}{\tau^* + k}c = \beta^*.$$

A few observations follow immediately: (1) any information type with τ_i equal to τ^* will be hired at the same rates in the first-best and equilibrium; (2) τ^* is always unique and well-defined; (3) τ^* can be non-integer, but the analysis can be generalized to $\tau_i \in \mathbb{R}_+$, via a completion of the analysis for integer $\tau_i \in \mathbb{Z}_+$; (4) τ^* is strictly between τ_N and τ_1 .

The equilibrium contract from Theorem 1 and the first-best hiring threshold generate misalignment between the two suggestion regions. Further, delegation systematically improves the hiring outcomes of groups of candidates who are better understood by recruiters.

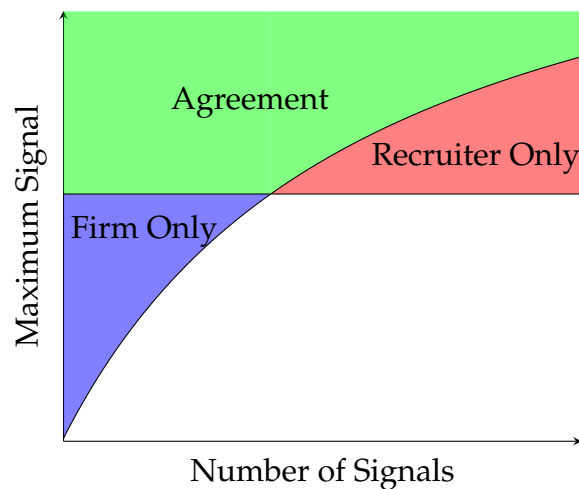
Theorem 3 *All information types with more signals than τ^* , i.e., $i < \min\{i' | \tau_{i'} \leq \tau^*\}$, are hired with higher probability in equilibrium than in the first-best, while all information types with fewer signals than τ^* , i.e., $i > \max\{i' | \tau_{i'} \geq \tau^*\}$, are hired with lower probability than in the first-best.*

¹⁴There is full surplus extraction, and the first-best outcome is generally achieved in equilibrium when there is only one information type.

The complete proof is provided in Appendix Section 2.8.4. Under the equilibrium contract, there is a common threshold for hired candidates in terms of the maximum observed signal: the refund payment (β^*). This corresponds to the horizontal line in Figure 2.2. We show via induction that this threshold is both strictly below the first-best threshold for the information type with the most signals and strictly above the first-best threshold for the information type with the least signals. This implies that the blue and red regions in Figure 2.2 are not empty.

There is a positive share of candidates who are not hired in equilibrium but that the firm would like to hire because their maximum signal is high enough given the small number of signals observed. These candidates are depicted as the “Firm Only” (blue) region in Figure 2.2. There is also a positive share of candidates who are hired in equilibrium that the firm would prefer not to hire because their maximum signal is not high enough given the large number of signals observed. These candidates are depicted as the “Recruiter Only” (red) region in Figure 2.2.

Figure 2.2: Types of Candidates Hired

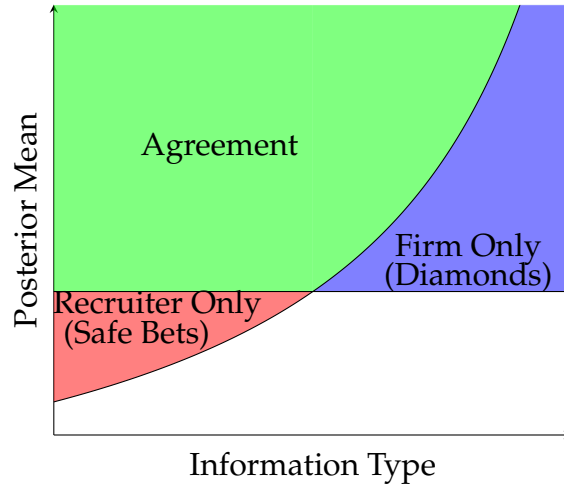


Note: The space of candidates is depicted in terms of their information type (number of signals) and expected productivity. The green region includes all candidates hired in both the first-best and equilibrium. The red region includes candidates hired only in equilibrium. The blue region includes candidates hired only in the first-best.

In the first-best benchmark only candidates with posterior expected productivity greater than the hiring cost (c) are hired. In equilibrium, the trade-offs induced by delegation are such that the firm optimally sets the suggestion bar to a common threshold in terms of signals. Because different signal realizations imply different expected productivity across information types, the recruiter suggests some candidates with posterior expected productivity less than c and does not suggest some candidates with expected productivity above c . First-stage screening is distorted in equilibrium relative to the first-best.

To better understand this result, in Figure 2.3 we depict the set of suggested and hired candidates in terms of their information types and posterior means, (i, μ) . The key misalignment is that the equilibrium suggestion region excludes candidates who have high posterior expected productivity but a high residual uncertainty in less informative groups (so-called diamonds in the rough, depicted in blue in Figure 2.3), and includes candidates who have low posterior expected productivity and a low residual uncertainty from more informative groups (so-called safe bets, depicted in red in Figure 2.3). The optimal contract always generates a positive share of the two types of candidates.

Figure 2.3: Candidate Types Hired in Equilibrium vs. First-Best



Note: Candidates are depicted in terms of their posterior expectation and information type. Diamonds in the rough are candidates that are poorly understood (high information type) but are expected to have high productivity. Safe bets are candidates that are well-understood (many signals) but are expected to have low productivity. The green region includes all candidates hired in both the first-best and equilibrium. The red region includes candidates hired only in equilibrium. The blue region includes candidates hired only in the first-best.

2.5 Comparative Statics

In the baseline model, underlying productivity (\bar{a}) is distributed similarly across information types. In this section we relax this assumption and allow candidate types to vary in the scale parameter of their prior productivity (\bar{a}_i). We also ask how improvements in information, group size, and productivity of one information type impact the other types. We show that in the first-best, hiring outcomes of information types are independent, while in equilibrium delegation via refund contracts causes spillovers.

We begin by considering the first-best. Following the same logic as in the baseline model, in the first-best, the firm hires a candidate if $\mu^*(i) = c$.

Theorem 1' *The unique equilibrium contract is a refund contract with*

$$\beta^* = \frac{\mathbb{E}\left[\frac{\bar{a}_i^k \tau_i}{\tau_i + k}\right]}{\mathbb{E}\left[\frac{\bar{a}_i^k \tau_i}{\tau_i + k - 1}\right]} c, \quad \alpha^* = \bar{u}.$$

The proof of Theorem 1' requires only minor additions to the proof of Theorem 2 and is given in Appendix Section 2.8.3. A key observation is that heterogeneity in prior productivity essentially re-weights information types within the expectations that determine the optimal refund. Notice that $\frac{\alpha^* - \bar{u}}{\beta^*} = 0 < \frac{1}{2}$ and thus we have group statistical discrimination in favor of more information as in the baseline model. We now analyze how an increase in information for one information type affects the other types in the population.

Proposition 5 *Suppose more signals become available for information type i with $\tau_i > \tau^*$, i.e., τ_i increases. Then, β^* increases and candidates from any type that is not i are hired with lower unconditional probability (for all other information types together, each type separately, or any given candidate of any ability and of any type that is not i).*

The proof is provided in Appendix Section 2.8.5.1. In equilibrium, additional information for one information type spills over onto other information types because it raises the refund payment. In the first-best, each type has its own threshold signal and information improvements for one type do not impact the other types. Consider the case when there are only two information types, one with many signals and one with few. If we add a signal to the information type that already has many signals, we increase the threshold we require of all types, further widening the difference in hiring probabilities between the first-best and equilibrium for the low type. This happens because of two forces acting in the same direction: (1) more signals increase the individual first-best threshold for this type, increasing the “average” equilibrium threshold; (2) by improving information for the type, we make correctly screening the type more important, which means that we should put “more weight” on this individual type first-best threshold, which is above the

“average” equilibrium threshold. In this sense, we increase delegation-based statistical discrimination.

Second, we analyze how first-order stochastic improvements of ability in one information type affect the hiring probabilities for all other information types. The proofs for the remaining two results are provided in Appendix Section [2.8.5.2](#).

Proposition 6 *Suppose the prior productivity of information type i improves, i.e., \bar{a}_i increases. If τ_i is greater (less) than τ^* , then β^* increases (decreases) and candidates from any other type (not i) are hired with lower (higher) unconditional probabilities (for all other information types together, each type separately, or any given candidate of any productivity and of any type that is not i).*

This comparative static highlights that the equilibrium contract is always attempting to balance the inclusion of safe bets and the exclusion of diamonds in the rough. In any equilibrium, diamonds in the rough come exclusively from information types below the cutoff τ^* , while safe bets come exclusively from information types above the cutoff. A productivity improvement for an information type below the cutoff makes excluding diamonds in the rough more expensive, and the firm responds by reducing β in order to reduce screening. A productivity improvement for an information type above the cutoff makes including safe bets more expensive, and the firm responds by increasing β in order to improve screening.

Finally, we analyze how a change in the size of an information type impacts the hiring probabilities of the other information types.

Proposition 7 *Assume that information type i becomes larger relative to the other information types, i.e., p_i increases and all other p_j proportionally decrease. If τ_i is greater (less) than τ^* , then β^* increases (decreases) and candidates from any other type (not i) are hired with lower (higher) unconditional probabilities (for all other types together, each type separately, or any given candidate of any productivity and of any type that is not i).*

The similarity between Propositions [6](#) and [7](#) illustrates that productivity improvements have similar effects to changing the relative sizes of information types. The intuition for

this result is similar in spirit to Proposition 6. Increasing the size of information types below the cutoff means there are more diamonds in the rough, so the firm reduces screening. Increasing the size of information types above the cutoff means that there are more safe bets, so the firm improves screening.

2.6 Robustness and Extensions

In this section, we show that the qualitative results derived under the parametric information structure without commitment are robust. We do this by considering several variations of the model. Without commitment, if the contract space is enriched to include a retainer payment which can be used to extract all surplus, the first-best is still not achieved. When the recruiter has a choice over an intensive margin of screening (modeled using costly sequential search), misalignment and different types of statistical discrimination still arise. With commitment, full extraction generally occurs, but misalignment persists due to the fundamental tension between first- and second-stage screening.

2.6.1 Three-Part Contracts without Commitment

In this subsection we consider the case without commitment ($\gamma = \beta$), and we show the need to use (α, β) to extract surplus is not pivotal in preventing the achievement of the first-best. To do this, we enrich the space of contracts to include an additional transfer which is paid out prior to the recruiter observing the information type and productivity signals of candidates. This additional transfer functions as a form of retainer, and it can be used directly by the firm to extract all surplus from the recruiter, freeing up other parts of the contract (α, β) to play other roles. We continue to assume the Pareto-uniform information structure for this result.

Proposition 8 *Suppose the firm designs three-part contracts, with an additional transfer before the recruiter sees the productivity signals. Then the following are true:*

1. *Profit is weakly higher than in the baseline equilibrium.*

2. *The first-best profit and set of hired and suggested candidates are not achieved.*

Proof. Recall that the recruiter cannot commit to a suggestion decision rule, because the signals are not contractible. In the model we consider, \bar{u} is the outside option for rejecting the contract and also the outside option for not suggesting any candidate after the contract is accepted. So, if the recruiter accepts the contract but does not suggest a candidate, the recruiter's payoff is the upfront transfer plus \bar{u} . The firm can replicate the profit from the baseline model by setting the additional transfer to be 0 and setting the other contract payments to be the same. Therefore, the firm's profit is weakly higher.

The additional transfer occurs before the signals are realized, so it impacts only whether the recruiter accepts the contract. In any equilibrium, the firm sets this transfer such that the recruiter's expected utility is equal to the recruiter's outside option \bar{u} . This implies the firm maximizes total surplus in order to maximize profit. Since total surplus depends only on the set of candidates suggested and terminated, achieving the first-best requires achieving the same set of suggested and terminated candidates.

Suppose for the sake of contradiction that the first-best hiring and suggestion regions are achieved. In the first-best, all candidates who are suggested are not terminated. To achieve this scenario in equilibrium, the recruiter must not suggest anyone who is later terminated. This requires $\beta > 0$ and the upfront payment equal to the outside option $\alpha = \bar{u}$. However, this implies that the suggestion threshold in terms of signals $x_{EQ}^\tau(\alpha, \beta) = \beta$, which is the same for all information types. This is a contradiction, because in the first-best the suggestion thresholds are different for different types: $x_{FB}^\tau = \frac{\tau+k-1}{\tau+k}c$. ■

The proposition demonstrates that the need to extract rent is not pivotal in terms of preventing achievement of the first-best. The proof of the proposition further confirms that the deeper tension is between achieving the first-best termination level and achieving the first-best suggestion region.

2.6.2 Multiple Candidates and Delegated Search

In our baseline model, the recruiter either suggests the candidate or not. This situation, which can be viewed as screening along an extensive margin, captures many realistic scenarios where the recruiter has a candidate in hand who they can choose to suggest for an opportunity. In other situations, the recruiter may have the ability to search more or less hard for a candidate that will satisfy the firm's hiring requirements. This situation can be viewed as screening along an intensive margin.

Consider a model where the recruiter can pay a search cost to sample additional candidates. After sampling a candidate, the recruiter views the candidate's information type and signals. In such a situation, the refund contract determines not just the set of suggested candidates but also the amount of costly effort the recruiter spends searching for candidates. How does this impact misalignment and statistical discrimination?

In such a sequential search model with normal signals and two information types, the spirit of our results continues to hold. The candidate from the more uncertain group is hired more often in the first-best and equilibrium. Intuitively, adding an intensive margin introduces moral hazard issues. The contract must have strong incentives not just so that the right candidates are chosen but also so that enough search effort is exerted by the recruiter. In general, the firm will choose to encourage less search in this setting because the selection of candidates is distorted.

2.6.3 Nonparametric Full Surplus Extraction with Commitment

In this subsection, we show the equilibrium contract with commitment extracts all surplus from the recruiter for any information structure which satisfies our nonparametric ordering assumption. We also discuss how despite full surplus extraction, the tension between first- and second-stage screening generally persists and prevents achievement of the first-best. This illustrates that the full extraction and failure to achieve the first-best that we observe in our parametric analysis without commitment is a robust feature of equilibrium contracts.

Proposition 9 *The optimal contract with commitment features full extraction.*¹⁵

Proof

- **Transfers are positive.** We prove that in equilibrium, the transfer from the firm to the recruiter for suggestion of a candidate is at least the outside option ($\alpha \geq \bar{u}$), and the transfer from the recruiter back to the firm is weakly positive ($\beta \geq 0$) by contradiction. Suppose that in an optimal contract $\beta < 0$. Then the recruiter will suggest a candidate with termination probability above threshold $p^* = \frac{\alpha - \bar{u}}{\beta}$, or $\mu < \mu_i$, which yields a negative selection among suggested candidates and contradicts optimality.

If $\beta=0$, then the recruiter either suggests any candidate if $\alpha \geq \bar{u}$ or suggests no candidate otherwise. Either situation is suboptimal by assumption. If $\alpha < \bar{u}$ and $\beta \geq 0$, the recruiter will never suggest the candidate, because $\alpha - \beta \Pr(a < \gamma|i, \mu) < \bar{u}$. Thus, the contract yields 0 profit for the firm, which again contradicts optimality.

- **Refund is equal to zero and the firm extracts the total surplus.** Assume any contract with $\beta > 0$, $\alpha > \bar{u}$, and some γ . The recruiter suggests all candidates with the equilibrium termination probability $\Pr(a < \gamma|i, \mu) \leq p^* = \frac{\alpha - \bar{u}}{\beta}$ and gets strictly positive expected utility for those candidates with termination probability strictly below p^* . We propose increasing the firm's profit by a sequence of contracts with the same suggestion and termination rules (and thus the surplus) that converges to full extraction of the surplus from the recruiter:

$$(\beta_m, \alpha_m = \bar{u} + \beta_m \cdot \frac{\alpha - \bar{u}}{\beta}, \gamma_m = \gamma)_{m=1}^{\infty}$$

and $\lim_{m \rightarrow \infty} \beta_m = 0$ ($\beta_m > 0 \forall m$). Each of these contracts has $p_m^* = p^*$ and thus leads to the same suggesting decisions. Moreover, γ is fixed across all contracts, which means that the sets of hired and the set of retained candidates do not change. Finally,

¹⁵There is a sequence of contracts that converges to the firm's supremum profit and features full extraction in the limit.

the recruiter's overall expected utility decreases and converges to \bar{u} (as it is bounded between \bar{u} and α), which, given the fixed surplus, implies that the firm's expected payoff increases. ■

The proposition confirms that under commitment, all equilibria feature full surplus extraction. The proof holds economic insight: The firm can vary β, α in such a way that it fully extracts all surplus while maintaining the same set of suggested candidates.

We can interpret this result in light of the three goals of the firm outlined in Section 2.3.4: first-stage screening, second-stage screening, and surplus extraction. Proposition 9 implies that with commitment the firm can perfectly extract all surplus and still be left with two degrees of freedom to accomplish the other two goals. However, the remaining two goals are still in conflict. Replicating the first-best suggestion region generally requires inefficient terminations compared to the first-best. Thus, the first-best will not generally be achieved even with commitment.

2.7 Conclusion

This paper asks how refund contracts, which are used to delegate recruitment to recruiters, shape the types of candidates ultimately hired by the firm. We show that relative to a first-best benchmark where a firm recruits directly, delegation induces endogenous risk aversion. Refund contracts make the recruiter less likely to suggest and the firm less likely to hire candidates that are less well-understood by the recruiter.

We call this statistical discrimination in favor of more information. We show that it occurs under many refund contracts and many variants of the model. In the first-best, the firm generally statistically discriminates in favor of less information, because it can terminate poor performers. Despite this result, we show that in one tractable case, the unique equilibrium contract features misalignment. Thus, the firm is typically unable to design a contract which avoids misalignment of statistical discrimination.

We show that the presence of information differences across candidates is a main driver of misalignment. When candidates are all understood equally well, the first-best can be achieved using refund contracts. This could give rise to a vicious cycle, the nature of which we sketch here but leave the study of for future work. Suppose candidates that are not hired get discouraged and exit an industry. Suppose also that firms can recruit directly at an opportunity cost or hire a recruiter. Because candidates about whom recruiters are less informed are less likely to be hired, they will exit at higher rates. In subsequent periods, the pool of candidates will be more homogeneous, and delegation will become attractive for a greater share of firms. This would then encourage more firms to delegate, thus continuing the cycle.

We see two veins for future work. First, the theoretical impact of refund contracts should be tested empirically using a two-sided audit study, where recruiters are hired to fill a position, and the experiment varies both the types of candidates the recruiters have access to and the type of compensation scheme. Second, how candidates behave in the presence of widespread delegation should be explored. Job candidates often make costly choices about which signals to acquire, and these choices are shaped by the labor market returns to these signals. We have shown in this paper that delegated recruitment changes the return to additional signals. It remains unclear if job candidates strategically change how they build their resumes and careers as recruiters and headhunters become more common.

2.8 Appendix

2.8.1 Bayesian Updating Under Pareto-Uniform Information Structure

To derive the posterior distribution, note that $x|a \sim U[0, a]$. Therefore the joint probability density function of τ signals given a is:

$$\prod_{t=1}^{\tau} f_{x|a}(x_t) = \frac{1}{a^{\tau}} \prod_{t=1}^{\tau} \mathbb{I}\{x_t \leq a\}$$

Notice that $\prod_{t=1}^{\tau} \mathbb{I}\{x_t \leq a\} = 1$ if and only if $x_{max}^{\tau} \leq a$. Therefore we can re-write as:

$$\prod_{t=1}^{\tau} f_{x|a}(x_t) = \frac{1}{a^{\tau}} \mathbb{I}\{x_{max}^{\tau} \leq a\}$$

The joint distribution of $a, \{x_t\}_{t=1}^{\tau}$ is then the product of this distribution and the prior distribution of a :

$$\begin{aligned} f(a, \{x_t\}_{t=1}^{\tau}) &= \frac{1}{a^{\tau}} \mathbb{I}\{x_{max}^{\tau} \leq a\} \mathbb{I}\{a \geq \bar{a}\} \frac{k\bar{a}^k}{a^{k+1}} \\ &= \frac{k\bar{a}^k}{a^{\tau+k+1}} \mathbb{I}\{\max\{x_{max}^{\tau}, \bar{a}\} \leq a\} \end{aligned}$$

Notice that this depends only on the maximum signal. Therefore x_{max}^{τ} is a sufficient statistic for $\{x_t\}_{t=1}^{\tau}$ for a . Note that the conditional PDF will be proportional to the joint PDF for any fixed x_{max}^{τ} . We can then solve for the multiplicative constant A that makes the conditional PDF integrate to 1 over the support:

$$f(a|\{x_t\}_{t=1}^{\tau}) = \int_{\max\{x_{max}^{\tau}, \bar{a}\}}^{\infty} \frac{A}{a^{\tau+k+1}} da \leftrightarrow \frac{A}{(\tau+k) \max\{x_{max}^{\tau}, \bar{a}\}^{\tau+k}} = 1$$

$$\begin{aligned} A &= (\tau+k) \max\{x_{max}^{\tau}, \bar{a}\}^{\tau+k} \leftrightarrow f(a|\{x_t\}_{t=1}^{\tau}) \\ &= \mathbb{I}\{\max\{x_{max}^{\tau}, \bar{a}\} \leq a\} \frac{(\tau+k) \max\{x_{max}^{\tau}, \bar{a}\}^{\tau+k}}{a^{\tau+k+1}} \end{aligned}$$

Thus, $a|\{x_t\}_{t=1}^{\tau} \sim Pareto(\max\{x_{max}^{\tau}, \bar{a}\}, \tau+k)$.

2.8.2 Derivations for the Proof of Theorem 2

This section proves that under refund contracts, firm profit is strictly decreasing in α . It then shows that after $\alpha = \bar{u}$, β^* given in the theorem solves the first-order condition. It concludes by showing that β^* satisfies the second-order condition. For historical reasons and to reduce algebra, derivations are performed by making the change of variables: $\tilde{\alpha} = \beta - \alpha + \bar{u}$, $\tilde{\beta} = \beta$. $\tilde{\beta}$ is a bonus, and $\tilde{\alpha}$ is an upfront payment incorporating the outside

option. The recruiter suggests a candidate if $\tilde{\beta}Pr(a \geq \beta | x_{max}^\tau) - \tilde{\alpha} \geq 0$, where note $\tilde{\alpha}$ is taken from the recruiter. The restriction that α be weakly greater than the outside option translates to $\beta \geq \tilde{\alpha}$. We now suppress the tilde notation. The firm chooses a contract to maximize profit:

$$\max_{\beta \geq 0, \alpha \leq \beta} \mathbb{E} \left[\mathbb{I}\{x_{max}^\tau \geq x_{EQ}^\tau(\alpha, \beta)\} \left(\mathbb{I}\{a \geq \beta\}(a - \beta) - c + \alpha \right) \right]$$

This can be broken into three parts. First, the expected productivity benefit of hiring a candidate:

$$\mathbb{E}[\mathbb{I}\{x_{max}^\tau \geq x_{EQ}^\tau(\alpha, \beta)\} \mathbb{I}\{a \geq \beta\} a]$$

Second, the expected suggestion costs:

$$\mathbb{E}[\mathbb{I}\{x_{max}^\tau \geq x_{EQ}^\tau(\alpha, \beta)\}(\alpha - c)]$$

Third, the expected bonus cost:

$$-\mathbb{E}[\mathbb{I}\{x_{max}^\tau \geq x_{EQ}^\tau(\alpha, \beta)\} \mathbb{I}\{a \geq \beta\} \beta]$$

The productivity benefit (2.8.2) can be written as:

$$\begin{aligned} & E[\mathbb{I}\{a > \beta\} \mathbb{I}\{x_{max}^\tau \geq \beta \left(\frac{\alpha}{\beta}\right)^{\frac{1}{\tau+k}}\} a] \\ &= \frac{\tau}{\tau+k-1} \frac{k\bar{a}^k}{k-1} \beta^{1-k} + \frac{k\bar{a}^k}{\tau+k-1} \beta^{1-\tau-k} \left[\beta^\tau - \left(\beta \left(\frac{\alpha}{\beta}\right)^{\frac{1}{\tau+k}} \right)^\tau \right] \\ &= \bar{a}^k \frac{k}{\tau+k-1} \beta^{1-k} \left[\frac{\tau}{k-1} + 1 - \left(\frac{\alpha}{\beta}\right)^{\frac{\tau}{\tau+k}} \right] \\ &= \beta \left(\frac{\bar{a}}{\beta}\right)^k \frac{k}{\tau+k-1} \left[\frac{\tau}{k-1} + 1 - \left(\frac{\alpha}{\beta}\right)^{\frac{\tau}{\tau+k}} \right] \end{aligned}$$

The expected suggestions costs (2.8.2) can be written as:

$$(-c + \alpha)Pr(x_{\tau}^{max} \geq \beta \left(\frac{\alpha}{\beta}\right)^{\frac{1}{\tau+k}}) = \frac{\tau}{\tau+k} \left(\frac{\bar{a}}{\alpha^{\frac{1}{\tau+k}}}\right)^k \beta^{\frac{k}{\tau+k}-k} (-c + \alpha) \quad (2.2)$$

$$= \frac{\tau}{\tau+k} \left(\frac{\bar{a}}{\beta}\right)^k \left(\frac{\beta}{\alpha}\right)^{\frac{k}{\tau+k}} (-c + \alpha) \quad (2.3)$$

$$= \left(\frac{\alpha}{\beta}\right)^{\frac{\tau}{\tau+k}} \beta \left(\frac{\bar{a}}{\beta}\right)^k \frac{\tau}{\tau+k} - c \frac{\tau}{\tau+k} \left(\frac{\bar{a}}{\beta}\right)^k \left(\frac{\beta}{\alpha}\right)^{\frac{k}{\tau+k}} \quad (2.4)$$

The expected bonus cost (2.8.2) can be written as:

$$- \beta Pr(a \geq \beta, x_{\tau}^{max} \geq x_{EQ}^{\tau}(\alpha, \beta)) \quad (2.5)$$

$$= -\beta Pr(x_{max}^{\tau} \geq \beta) - \beta Pr(x_{EQ}^{\tau}(\alpha, \beta) \leq x_{\tau}^{max} \leq \beta, a \geq \beta) \quad (2.6)$$

$$= -\beta \frac{\tau}{\tau+k} \left(\frac{\bar{a}}{\beta}\right)^k - \beta \int_{\beta(\frac{\alpha}{\beta})^{\frac{1}{\tau+k}}}^{\beta} \left(\frac{x}{\beta}\right)^{\tau+k} \frac{k\tau \bar{a}^k}{(\tau+k)x^{k+1}} dx \quad (2.7)$$

$$= -\beta \frac{\tau}{\tau+k} \left(\frac{\bar{a}}{\beta}\right)^k - \frac{\tau\beta k \bar{a}^k}{(\tau+k)\beta^{\tau+k}} \int_{\beta(\frac{\alpha}{\beta})^{\frac{1}{\tau+k}}}^{\beta} x^{\tau-1} dx \quad (2.8)$$

$$= -\beta \frac{\tau}{\tau+k} \left(\frac{\bar{a}}{\beta}\right)^k - \frac{\beta k}{\tau+k} \left(\frac{\bar{a}}{\beta}\right)^k \left(1 - \left(\frac{\alpha}{\beta}\right)^{\frac{\tau}{\tau+k}}\right) \quad (2.9)$$

$$= \beta \left(\frac{\bar{a}}{\beta}\right)^k \left[-1 + \frac{k}{\tau+k} \left(\frac{\alpha}{\beta}\right)^{\frac{\tau}{\tau+k}}\right] \quad (2.10)$$

Profit is the sum of (2.2), (2.4), and (2.10):

$$\beta \left(\frac{\bar{a}}{\beta}\right)^k \left[\left(\frac{\alpha}{\beta}\right)^{\frac{\tau}{\tau+k}} \frac{\tau-1}{\tau+k-1} + \frac{1}{k-1} - \frac{c}{\beta} \frac{\tau}{\tau+k} \left(\frac{\beta}{\alpha}\right)^{\frac{k}{\tau+k}} \right]$$

Note that this expression is strictly increasing in α , and it is strictly increasing in α for all valid values of τ , therefore total profit is increasing in α . Thus in the optimal contract the upfront payment is set to its maximal possible value: $\alpha = \beta$. The expression now becomes only a function of β :

$$\beta \left(\frac{\bar{a}}{\beta}\right)^k \left[\frac{\tau-1}{\tau+k-1} + \frac{1}{k-1} - \frac{c}{\beta} \frac{\tau}{\tau+k} \right]$$

Simplifying:

$$\pi_i = \left(\frac{\bar{a}}{\beta}\right)^k \left[\frac{\tau k}{(k-1)(\tau+k-1)}\beta - \frac{c\tau}{\tau+k} \right]$$

The first order condition for profit from information type i is:

$$\begin{aligned} \frac{\partial \pi_i}{\partial \beta} &= \frac{-k}{\beta} \left(\frac{\bar{a}}{\beta}\right)^k \left[\frac{\tau k}{(k-1)(\tau+k-1)}\beta - \frac{c\tau}{\tau+k} \right] + \left(\frac{\bar{a}}{\beta}\right)^k \frac{\tau k}{(k-1)(\tau+k-1)} \\ &= k \left(\frac{\bar{a}}{\beta}\right)^k \left\{ \frac{c}{\beta} \frac{\tau}{\tau+k} + \frac{\tau}{(k-1)(\tau+k-1)} - \frac{\tau k}{(k-1)(\tau+k-1)} \right\} \\ &= k \left(\frac{\bar{a}}{\beta}\right)^k \left\{ \frac{c}{\beta} \frac{\tau}{\tau+k} - \frac{\tau(k-1)}{(k-1)(\tau+k-1)} \right\} \\ &= k \left(\frac{\bar{a}}{\beta}\right)^k \left\{ \frac{c}{\beta} \frac{\tau}{\tau+k} - \frac{\tau}{\tau+k-1} \right\} \end{aligned}$$

$$\frac{\partial \pi_i}{\partial \beta} = k \left(\frac{\bar{a}}{\beta}\right)^k \left\{ \frac{c}{\beta} \frac{\tau}{\tau+k} - \frac{\tau}{\tau+k-1} \right\}$$

Taking the weighted average of the first order condition for each information type gives the total profit first-order condition:

$$\begin{aligned} \sum_i p_i k \left(\frac{\bar{a}}{\beta}\right)^k \left[-\frac{\tau_i}{\tau_i+k-1} + \frac{c}{\beta} \frac{\tau_i}{\tau_i+k} \right] &= k \left(\frac{\bar{a}}{\beta}\right)^k \left[-\sum_i p_i \frac{\tau_i}{\tau_i+k-1} + \frac{c}{\beta} \sum_i p_i \frac{\tau_i}{\tau_i+k} \right] \\ &= k \left(\frac{\bar{a}}{\beta}\right)^k \left(-\mathbb{E} \left[\frac{\tau_i}{\tau_i+k-1} \right] + \frac{c}{\beta} \mathbb{E} \left[\frac{\tau_i}{\tau_i+k} \right] \right) = 0 \end{aligned}$$

Which yields the optimal value of β :

$$\beta = \frac{\mathbb{E} \left[\frac{\tau_i}{\tau_i+k} \right]}{\mathbb{E} \left[\frac{\tau_i}{\tau_i+k-1} \right]} c$$

Returning to our original change of variables, $\tilde{\beta} = \beta$ and $\tilde{\alpha} = \tilde{\beta}$ thus $\alpha = \bar{u}$. We show the second-order condition is satisfied at β^* for the case with prior productivity heterogeneity (which nests the baseline model) in the next section.

2.8.3 Optimal Contract with Prior Heterogeneity

Suppose additionally that information types also differ in prior productivity. That is they differ in their Pareto shift parameter \bar{a}_i in addition to the number of signals. We can still use the same argument as above to say that $\alpha = \beta$ because the argument is true for each information type individually regardless of the values of \bar{a}_i, k . We can adapt the first-order condition (2.8.2) to incorporate heterogeneity:

$$\begin{aligned} \sum_i p_i k \left(\frac{\bar{a}_i}{\beta} \right)^k \left\{ \frac{c}{\beta} \frac{\tau_i}{\tau_i + k} - \frac{\tau_i}{\tau_i + k - 1} \right\} &= 0 \\ \sum_i p_i \left[\frac{\bar{a}_i^k \tau_i}{\tau_i + k - 1} \beta - \frac{\bar{a}_i^k c \tau_i}{\tau_i + k} \right] &= 0 \\ \beta \sum_i p_i \frac{\bar{a}_i^k \tau_i}{\tau_i + k - 1} &= \sum_i p_i \frac{\bar{a}_i^k c \tau_i}{\tau_i + k} \\ \beta^* &= \frac{\sum_i p_i \frac{\bar{a}_i^k c \tau_i}{\tau_i + k}}{\sum_i p_i \frac{\bar{a}_i^k \tau_i}{\tau_i + k - 1}} \\ \beta^* &= \frac{\mathbb{E} \left[\frac{\bar{a}_i^k \tau_i}{\tau_i + k} \right]}{\mathbb{E} \left[\frac{\bar{a}_i^k \tau_i}{\tau_i + k - 1} \right]} c \end{aligned}$$

We can also check the second-order condition, to make sure that the solution is a local maximum. It is then also the global maximum, since the function is everywhere differentiable, and the solution to the first-order condition is unique.

$$\frac{\partial^2 \pi}{\partial \beta^2} = \sum_i \frac{p_i k \tau_i}{\beta} \cdot \left(\frac{\bar{a}_i}{\beta} \right)^k \left\{ -\frac{c}{\beta} \frac{k+1}{\tau_i + k} + \frac{k}{\tau_i + k - 1} \right\}$$

Now we can show that for any $\beta < c$ and for any i , the expression inside the brackets is negative, which will imply that the second-order derivative is negative, the function is concave, and the second-order condition is satisfied.

$$\frac{c}{\beta} \cdot \frac{k+1}{\tau_i+k} > \frac{k+1}{\tau_i+k} = \frac{k}{(\tau_i+k-1)+1} > \frac{k}{\tau_i+k-1}$$

■

2.8.4 Proof of Theorem 3

Recall that information types are indexed by their number of signals (τ_i) from most to least signals. Then we have that:

$$\frac{\tau_1+k-1}{\tau_1+k} > \dots > \frac{\tau_i+k-1}{\tau_i+k} > \dots > \frac{\tau_N+k-1}{\tau_N+k}$$

Therefore:

$$\frac{p_1 \frac{\tau_1}{\tau_1+k}}{p_1 \frac{\tau_1}{\tau_1+k-1}} > \dots > \frac{p_i \frac{\tau_i}{\tau_i+k}}{p_i \frac{\tau_i}{\tau_i+k-1}} > \dots > \frac{p_N \frac{\tau_N}{\tau_N+k}}{p_N \frac{\tau_N}{\tau_N+k-1}}$$

Note that:

$$\frac{a}{b} > \frac{c}{d} \implies \frac{a}{b} > \frac{a+c}{b+d} > \frac{b}{d}$$

Therefore:

$$\begin{aligned} \frac{p_1 \frac{\tau_1}{\tau_1+k}}{p_1 \frac{\tau_1}{\tau_1+k-1}} &> \frac{p_2 \frac{\tau_2}{\tau_2+k}}{p_2 \frac{\tau_2}{\tau_2+k-1}} \implies \\ \frac{p_1 \frac{\tau_1}{\tau_1+k}}{p_1 \frac{\tau_1}{\tau_1+k-1}} &> \frac{p_1 \frac{\tau_1}{\tau_1+k} + p_2 \frac{\tau_2}{\tau_2+k}}{p_1 \frac{\tau_1}{\tau_1+k-1} + p_2 \frac{\tau_2}{\tau_2+k-1}} > \frac{p_2 \frac{\tau_2}{\tau_2+k}}{p_2 \frac{\tau_2}{\tau_2+k-1}} \end{aligned}$$

We can then repeat this process with information type 3:

$$\begin{aligned} \frac{p_1 \frac{\tau_1}{\tau_1+k} + p_2 \frac{\tau_2}{\tau_2+k}}{p_1 \frac{\tau_1}{\tau_1+k-1} + p_2 \frac{\tau_2}{\tau_2+k-1}} &> \frac{p_3 \frac{\tau_3}{\tau_3+k}}{p_3 \frac{\tau_3}{\tau_3+k-1}} \implies \\ \frac{p_1 \frac{\tau_1}{\tau_1+k} + p_2 \frac{\tau_2}{\tau_2+k}}{p_1 \frac{\tau_1}{\tau_1+k-1} + p_2 \frac{\tau_2}{\tau_2+k-1}} &> \frac{p_1 \frac{\tau_1}{\tau_1+k} + p_2 \frac{\tau_2}{\tau_2+k} + p_3 \frac{\tau_3}{\tau_3+k}}{p_1 \frac{\tau_1}{\tau_1+k-1} + p_2 \frac{\tau_2}{\tau_2+k-1} + p_3 \frac{\tau_3}{\tau_3+k-1}} > \frac{p_3 \frac{\tau_3}{\tau_3+k}}{p_3 \frac{\tau_3}{\tau_3+k-1}} \end{aligned}$$

Continuing until the final information type N we have that:

$$\frac{p_1 \frac{\tau_1}{\tau_1+k}}{p_1 \frac{\tau_1}{\tau_1+k-1}} > \frac{\sum_i p_i \frac{\tau_i}{\tau_i+k}}{\sum_i p_i \frac{\tau_i}{\tau_i+k-1}} > \frac{p_N \frac{\tau_N}{\tau_N+k}}{p_N \frac{\tau_N}{\tau_N+k-1}}$$

The middle expression is the ratio of expectations observed in the equilibrium suggestion threshold. With this realization, we can multiply all expressions by the hiring cost c to obtain:

$$x_{FB}^1 = \frac{\tau_1 + k - 1}{\tau_1 + k} c > \frac{\mathbb{E} \left[\frac{\tau_i}{\tau_i+k} \right]}{\mathbb{E} \left[\frac{\tau_i}{\tau_i+k-1} \right]} c > \frac{\tau_N + k - 1}{\tau_N + k} c = x_{FB}^N \quad (2.11)$$

Thus we have that the first-best threshold for the information type with the most signals is greater than the equilibrium threshold while the first-best threshold for the information type with the least signals is less than the equilibrium threshold. Now consider the ratio of the probability of hire in the second best compared to the first-best for a given information type i :

$$R_i := \frac{Pr_i^{EQ}(\text{hire})}{Pr_{FB}^i(\text{hire})} = \left(\frac{x_{FB}^i}{x_{EQ}} \right)^k$$

The denominator is the same for all information types and k is a positive power, therefore the ratio is monotone decreasing in i (information type 1 has the highest ratio). From (2.11) we know that $x_{FB}^1/x_{EQ} > 1$ and $x_{FB}^N/x_{EQ} < 1$. Thus the information type with the most signals are hired at higher rates in the equilibrium than the first-best, while the information type with the least is hired less in the equilibrium than the first-best. Further, because the ratio is monotone decreasing in the number of signals, all types with a number of signals above some cut-off will be hired at higher rates, while all those below some cut-off will be hired at lower rates. ■

2.8.5 Proofs for Comparative Statics

2.8.5.1 Proof of Proposition 5

Consider all information types but i , and denote τ_{-i}^* the threshold τ for these information types.

$$\beta^* = \frac{(1-p_i) * \frac{\tau_{-i}^*}{\tau_{-i}^* + k} + p_i * \frac{\tau_i}{\tau_i + k}}{(1-p_i) * \frac{\tau_{-i}^*}{\tau_{-i}^* + k - 1} + p_i * \frac{\tau_i}{\tau_i + k - 1}}$$

$$\begin{aligned} \frac{d\beta^*}{d\tau_i} &= \frac{p_i}{(1-p_i) * \frac{\tau_{-i}^*}{\tau_{-i}^* + k - 1} + p_i * \frac{\tau_i}{\tau_i + k - 1}} * \frac{\partial \frac{\tau_i}{\tau_i + k}}{\partial \tau_i} \\ &\quad - \frac{\left((1-p_i) * \frac{\tau_{-i}^*}{\tau_{-i}^* + k} + p_i * \frac{\tau_i}{\tau_i + k} \right) * p_i}{\left((1-p_i) * \frac{\tau_{-i}^*}{\tau_{-i}^* + k - 1} + p_i * \frac{\tau_i}{\tau_i + k - 1} \right)^2} * \frac{\partial \frac{\tau_i}{\tau_i + k - 1}}{\partial \tau_i} \\ &\propto \frac{\frac{\partial \frac{\tau_i}{\tau_i + k}}{\partial \tau_i}}{\frac{\partial \frac{\tau_i}{\tau_i + k - 1}}{\partial \tau_i}} - x_{EQ}^* \\ &= \frac{k * (\tau_i + k - 1)}{(k - 1) * (\tau_i + k)} \cdot x_{FB}^i - x_{EQ}^* \\ &> x_{FB}^i - x_{EQ}^* > 0 \end{aligned}$$

since $\tau_i > \tau^*$ and $k * (\tau_i + k - 1) > (k - 1) * (\tau_i + k)$. This implies that $\frac{d\beta^*}{d\tau_i} > 0$. ■

2.8.5.2 Proof of Propositions 6 and 7

$$\beta^* = \frac{\mathbb{E} \left[\frac{\bar{a}_i^k \tau_i}{\tau_i + k} \right]}{\mathbb{E} \left[\frac{\bar{a}_i^k \tau_i}{\tau_i + k - 1} \right]} c = \frac{\sum_i p_i \bar{a}_i^k \cdot \left[\frac{\tau_i}{\tau_i + k} \right]}{\sum_i p_i \bar{a}_i^k \cdot \left[\frac{\tau_i}{\tau_i + k - 1} \right]} c$$

Let us define the following measure function over i : $\tilde{p}_i = p_i \bar{a}_i^k$. Then β^* is given by the formula for the homogeneous \bar{a} case

$$\beta^*/c = \frac{\mathbb{E}_{\tilde{p}} \left[\frac{\tau_i}{\tau_i+k} \right]}{\mathbb{E}_{\tilde{p}} \left[\frac{\tau_i}{\tau_i+k-1} \right]}$$

Increase in p_i or \bar{a}_i is equivalent to increase in \tilde{p}_i . Assume \tilde{p}_i increased to \tilde{p}'_i . The following inequalities are equivalent.

$$\begin{aligned} \frac{\tilde{p}'_{-i} \frac{\tau_{-i}^*}{\tau_{-i}^*+k} + \tilde{p}'_i \frac{\tau_i}{\tau_i+k}}{\tilde{p}'_{-i} \frac{\tau_{-i}^*}{\tau_{-i}^*+k-1} + \tilde{p}'_i \frac{\tau_i}{\tau_i+k-1}} &> \frac{\tilde{p}_{-i} \frac{\tau_{-i}^*}{\tau_{-i}^*+k} + \tilde{p}_i \frac{\tau_i}{\tau_i+k}}{\tilde{p}_{-i} \frac{\tau_{-i}^*}{\tau_{-i}^*+k-1} + \tilde{p}_i \frac{\tau_i}{\tau_i+k-1}} \\ (\tilde{p}'_i - \tilde{p}_i) * \frac{\tau_i}{\tau_i+k} * \frac{\tau_{-i}^*}{\tau_{-i}^*+k-1} &> (\tilde{p}'_i - \tilde{p}_i) * \frac{\tau_{-i}^*}{\tau_{-i}^*+k} * \frac{\tau_i}{\tau_i+k-1} \\ (\tilde{p}'_i - \tilde{p}_i) * \frac{\frac{\tau_i}{\tau_i+k}}{\frac{\tau_i}{\tau_i+k-1}} &> (\tilde{p}'_i - \tilde{p}_i) * \frac{\frac{\tau_{-i}^*}{\tau_{-i}^*+k}}{\frac{\tau_{-i}^*}{\tau_{-i}^*+k-1}} \\ (\tilde{p}'_i - \tilde{p}_i) * \left(\frac{\frac{\tau_i}{\tau_i+k}}{\frac{\tau_i}{\tau_i+k-1}} - \frac{\frac{\tau_{-i}^*}{\tau_{-i}^*+k}}{\frac{\tau_{-i}^*}{\tau_{-i}^*+k-1}} \right) &> 0 \\ (\tilde{p}'_i - \tilde{p}_i) * \left(\frac{\tau_i+k-1}{\tau_i+k} - \frac{\tau_{-i}^*+k-1}{\tau_{-i}^*+k} \right) &> 0 \\ (\tilde{p}'_i - \tilde{p}_i) * (\tau_i - \tau_{-i}^*) &> 0 \end{aligned}$$

Since x_{FB}^τ is increasing in τ . Therefore an increase in p_i or \bar{a}_i increases x_{EQ}^* iff $\tau_i > \tau_{-i}^*$, which is equivalent to $\tau_i > \tau^*$. ■

2.8.6 Qualitative Interviews

We conducted three recorded interviews with people in the recruiting industry. The interviews occurred in 2019 and lasted around one hour each. The interviews asked a series of open-ended questions about recruiting in general and did not focus on the refund contract

in particular. Audio recordings are available but require approval from the interviewees prior to release.

One interview was with an early career recruiter who works for an external recruiting firm. One interview was with a mid-career headhunter who at the time managed his own external recruiting firm. One interview was with an early career internal recruiter and human resource professional. The interview with the internal recruiter did not yield much information about external recruiters so it is not discussed much in this paper.

Both the mid-career external headhunter and the early-career external recruiter stated that they utilized a contract structure where they got a fee that was contingent on the candidate staying for a certain period of time. Here is the excerpt from the mid-career headhunter:

Jacob (Author): “Do you ever consider the probability of termination or probability of leaving or retention when you are choosing who to suggest? Because you are paid based on placement but is it contingent on them staying for a certain period of time?”

Headhunter: “Yes most of the time it is anywhere from a 30 day to 90 day replacement. Sometimes it is full replacement no cost, sometimes it is a refund. We as recruiters may get paid 16,000 for a fee. But we can’t..you know. You just don’t want to send that money. Or you just won’t get paid until those 90 days are up.”

Here is the excerpt from the early-career recruiter:

Jacob (Author): “Do you consider probability of termination, probability of separation or retention when you are considering someone?”

Recruiter: “Yes.”

Jacob (Author): “What is the window you get paid for generally?”

Recruiter: “Generally it is 90 days. We get paid upfront. One of two things happen. We either have the next placement for free or we return the money.”

Jacob (Author): “And that’s if they leave for what reasons?”

Recruiter: “Any reason.”

Jacob (Author): "Even if the company fires them?"

Recruiter: "Yes."

Jacob (Author): "Do you ever contest why the company fired them?"

Recruiter: "No because the amount of times; I do not know if in my two years if we have ever had a candidate fired in the first 3 months."

Jacob (Author): "It is usually because they left or something else like that?"

Recruiter: "But even then it has happened only a few times. If a candidate leaves after the 90 day mark there are times when we may provide a discount on the next placement. But we are not tied to it by any means"

Jacob (Author): "Do you have any idea why it is 90 days?"

Recruiter: "It is just our good faith policy to our clients who we work with to say that hey we trust the candidates we are putting in front of you."

Jacob (Author): "It seems like everyone is 90 days, so. Like of all the people I have talked to, it seems like everyone is 90 days. Is there a reason for that number?"

Recruiter: "I have never had a conversation about it, my understanding would likely just be that turnover in tech is just so high we would not want it to be past 3 months, but 3 months is enough time for someone to get caught up to speed and get to work and get connected."

Jacob (Author): "One month is too short?"

Recruiter: "One month is too short. Six months is... there is too many reasons things could fall apart."

Additionally, when asked why companies use recruiter, one reason the early-career recruiter gave is: "Another reason is essentially we are free to use unless success happens. It doesn't cost any money unless they want to hire someone."

2.8.7 Proof of General Risk Attitudes

Proof. The firm's payoff for a suggested candidate with productivity a is equal to $\max\{a, 0\} - c$, which is a convex function. Let us take a candidate (μ, σ) , which the firm decides to suggest. Then any candidate to the upper right from this one, with $\mu' \geq \mu$ and $\sigma' \geq \sigma$ is also suggested since the firm's payoff function is increasing and convex, and therefore a FOSD shift and an MPS increase its expectation. Thus, the firm suggests everyone above some threshold $\tilde{\mu}_{FB}(\sigma)$ that is decreasing in σ . This means that the firm is risk-loving since it prefers candidates with a mean-preserving spread of productivity. ■

Proof. The recruiter suggests anyone with (μ, σ) such that

$$\alpha - \beta \Pr(a < \beta | \mu, \sigma) \geq \bar{u}$$

$$\Pr(a < \beta | \mu, \sigma) \leq \frac{\alpha - \bar{u}}{\beta} \leq q$$

$\Pr(a < \beta | \mu, \sigma)$ is the conditional CDF of a given μ, σ , which we denote $F_{(\mu, \sigma)}(a)$. Then we have that: $\beta \leq F_{(\mu, \sigma)}^{-1}(q)$. Then for a fixed σ , consider all candidates with $\mu = \tilde{\mu}_{EQ}(\sigma)$ and $\sigma' < \sigma$. Since the candidates are q -lower-tail-risk ordered and $\beta \leq F_{(\mu, \sigma)}^{-1}(q)$:

$$\Pr(a < \beta | \mu, \sigma') < \Pr(a < \beta | \mu, \sigma) < q$$

And the (μ, σ') candidate should also be suggested. All candidates with σ and $\mu' > \tilde{\mu}_{EQ}(\sigma)$ are also suggested since the FOSD implies

$$\Pr(a < \beta | \mu', \sigma) < \Pr(a < \beta | \mu, \sigma) < q$$

This together implies that if a candidate (μ, σ) is suggested then all candidates to the upper left, $(\mu' \geq \mu, \sigma' \leq \sigma)$ should also be suggested, which implies that $\tilde{\mu}_{EQ}(\sigma)$ is increasing in σ . This result, on the other hand, can be thought of as risk aversion because the recruiter who suggests a candidate (μ, σ) would also suggest any other candidate (μ', σ') that second-order stochastically dominates (μ, σ) candidate. ■

References

- Arslanian, Henri (2016) "Talent and Hiring," *Entrepreneurship in Finance: Successfully Launching and Managing a Hedge Fund in Asia*, 183–193.
- Association, American Staffing (2014) , 9, https://americanstaffing.net/wp-content/uploads/2014/09/NAPS_ASA_recruiting_agrmnt.doc.
- Black, Ines, Sharique Hasan, and Rembrand Koning (2022) "Hunting for talent: Firm-driven labor market search in the United States," *Available at SSRN*.
- Carroll, Gabriel (2017) "Robustness and separation in multidimensional screening," *Econometrica*, 85 (2), 453–488.
- Chade, Hector and Natalia Kovrijnykh (2016) "Delegated information acquisition with moral hazard," *Journal of Economic Theory*, 162, 55–92.
- Che, Yeon-Koo, Wouter Dessein, and Navin Kartik (2013) "Pandering to persuade," *American Economic Review*, 103 (1), 47–79.
- Cowgill, Bo and Patryk Perkowski (2020) "Delegation in Hiring: Evidence from a Two-Sided Audit," *Columbia Business School Research Paper* (898).
- Deutsch, Matthew (2019) "Most popular type of Recruitment Guarantee: Agency Guarantee period," 3, <https://www.topechelon.com/blog/recruiter-guarantees-once-i-am-paid-i-do-not-refund-any-money/>.
- Dingman, H Bruce (1993) "Contingency recruiting in the hospitality industry," *Cornell Hospitality Quarterly*, 34 (6), 30.
- EEOC (2023) "Prohibited Employment Policies/Practices," 8, <https://www.eeoc.gov/prohibited-employment-policiespractices>.
- Fink, Daniel (1997) "A compendium of conjugate priors," *See http://www.people.cornell.edu/pages/df36/CONJINTRnew%20TEX.pdf*, 46.
- Finlay, William and James E Coverdill (2007) *Headhunters: Matchmaking in the labor market*: Cornell University Press.
- Florea, Nicoleta Valentina (2014) "Using recruitment agencies to obtain the best candidates," *Land Forces Academy Review*, 19 (1), 80.
- Frankel, Alexander (2014) "Aligned delegation," *American Economic Review*, 104 (1), 66–83.
- (2021) "Selecting Applicants," *Econometrica*, 89 (2), 615–645.
- Gale, Douglas and Martin Hellwig (1985) "Incentive-compatible debt contracts: The one-period problem," *The Review of Economic Studies*, 52 (4), 647–663.

- Giuliano, Laura, David I Levine, and Jonathan Leonard (2009) "Manager race and the race of new hires," *Journal of Labor Economics*, 27 (4), 589–631.
- Inderst, Roman and Marco Ottaviani (2012) "Competition through commissions and kick-backs," *American Economic Review*, 102 (2), 780–809.
- Kundu, Tapas and Tore Nilssen (2020) "Delegation of regulation," *The Journal of Industrial Economics*, 68 (3), 445–482.
- Levitt, Steven D and Christopher M Snyder (1997) "Is no news bad news? Information transmission and the role of "early warning" in the principal-agent model," *The RAND Journal of Economics*, 641–661.
- Perry, David E and Mark J Haluska (2017) *Executive Recruiting For Dummies*: John Wiley & Sons.
- Press, Entrepreneur (2007) *Start Your Own Executive Recruiting Service: Your Step-By-Step Guide to Success*: Entrepreneur Press.
- Steinberg, Sheryl L and David S Machlowitz (1989) "How to Work with a Headhunter," *Legal Econ.*, 15, 43.
- Szalay, Dezsö (2005) "The economics of clear advice and extreme options," *The Review of Economic Studies*, 72 (4), 1173–1198.
- (2009) "Contracts with endogenous information," *Games and Economic Behavior*, 65 (2), 586–625.
- Yang, Frank (2021) "Costly multidimensional screening," *arXiv preprint arXiv:2109.00487*.

CHAPTER 3

Efficient Information Aggregation in DeGroot Model

Solo-authored¹

3.1 Introduction

People have different influence on the society. Social influence of an individual is largely defined by one's position in the social network. One can influence people directly or through other individuals in the network. Social influence of individual may be mathematically described by the centrality of the corresponding node in the graph characterizing the network.

¹We are grateful for insightful comments from Moritz Meyer-ter-Vehn, Simon Board, Ichiro Obara, Tomasz Sadzik, Sushil Bikhchandani, and William Zame. We thank participants in the UCLA theory pros-
eminar where this work was presented.

Figure 3.1: Social network and node centrality ²



In social learning models, the node/individual centrality defines how much the individual's opinion matters for the others. Fully rational agents could aggregate the information efficiently regardless of the structure of the network, however this requires the agents to be highly sophisticated. Trying to avoid this difficult computational problem, people tend to act according to some simpler heuristic learning rules. In this case the social importance of one's private information is no more defined by its quality but by the centrality of this person in the network.

One of conventionally accepted heuristic learning models in Economics is DeGroot model. For DeGroot model it has been shown that eigenvector centrality reflects the individuals' opinions importance in the society (Golub and Jackson, 2010). The information aggregation may be suboptimal if the centrality of the individual does not correspond to the quality of this individual's information. This would mean that people listen and pay a lot of attention to somebody who is not very knowledgeable in the question, but they

²https://en.wikipedia.org/wiki/Network_theory

ignore somebody who is rather advanced in the subject. Recent experience of learning details about Covid'19 illustrates this idea rather well. It happened that among central (influential) people in the society – like politicians – there are a few who did not possess deep knowledge about the medical aspects of the pandemic, however still disseminated some false or inaccurate facts about it, and those inaccurate opinions spread widely in the society due to the high centrality of their sources. On the other hand, the medical professionals were not heard by the public since they are not the central individuals in the society.

Individuals may have difficulties with considering the network structure which affects the correlation of the opinions or news that they observe. They also may not know with certainty who provides more and who provides less accurate facts about the matter. A Facebook (or any other online social network) user may not know if her friends are connected, and therefore their opinions are correlated. In this case, considering their opinions independent would be wrong as well as updating her own opinion based on this assumption. The user may also not know how trustworthy her friends are when sharing their opinions, and therefore make a mistake on how much weight to assign to their opinions too. Hence, the question I address in this paper is whether a benevolent Social Planner knowing the network structure and the knowledgeability of individuals can improve the results of the heuristic learning in this network.

More precisely, I consider that given an exogenous observational network Social Planner tries to improve the time asymptotic results of the DeGroot model. SP cannot create new links in the network but only can affect the weights individuals assign to each others' opinions for a given set of people they can observe. Facebook can choose which posts of my friends to show at the top and which of them at the bottom; what users to ban at all for spreading disinformation. Amazon (any internet retail platform) may genuinely care whether the customers efficiently learn about the quality of different items since it helps to increase sales volume. In order to do it, Amazon can show reviews from seemingly trustworthy users and ban those who apparently posts fake reviews. Some ways of censorship may also speak to the situation, if SP tries to prevent untrustworthy individuals

from sharing their opinions. We will focus on a situation where social planner knows how precise the private individual signals are and can reassign the individual's attention across her neighbors in the network.

The crucial assumption in this paper is that SP cannot create new links in the network. If it was not a case, apparently connecting everyone to everyone and assigning everyone a weight corresponding to the individual signal quality/precision would allow to aggregate information perfectly and immediately. (If the graph is not complete, it is not generally optimal to assign weights proportional to the information quality.) This however means that SP can make an individual to believe some strangers, and this is a too strong requirement. It is much easier to believe that in some cases SP can reshape the individual's attention across the other individuals she knows (her neighbors in the network).

To approach solving SP's problem, I start with characterization of the first best learning (the best possible learning given the private signals). The main contribution of this paper is demonstration that in any connected finite network it is always possible to achieve the first best learning. I provide an algorithm solving this problem and discuss the degrees of freedom in the construction of the solution. I also discuss a benchmark with simple average-based updating learning, where SP is only free to delete links between individuals in the network. In this model for homogeneous initial signals, the first best learning is attainable if and only if the initial graph contains a connected regular subgraph in it. Finally, I introduce a method of increasing the speed of convergence for a fixed asymptotic result.

Another view on the contribution of the paper may be considered in a pure statistical area. Since DeGroot interaction matrices are closely related to Markov processes the results may be interpreted as follows. For any connected aperiodic Markov process, changing the probabilities for existing transitions between states may help to establish any stationary probability distribution.

The rest of the paper is organized as follows. Section 2 discusses the related literature and the contribution to it. Section 3 introduces the model and relevant definitions distinguishing social network from interaction matrix. In section 4, I define the first best learning and discuss its relation to SP's problem solution. Section 5 contains the main results regarding achieving the first best learning. Section 5.2 addresses a problem under average-based updating learning and characterizes the necessary and sufficient conditions for achieving the first best learning. Section 6, I introduce a method of increasing the speed of convergence in DeGroot model. Section 7 concludes and discusses the limitations of the paper and the possibilities for the future research.

3.2 Literature

Social learning has been an important topic in Economic Theory for the last 30 years. A significant strand of the literature was focused on large population asymptotic analysis in rational sequential learning models. Starting with classic papers [Banerjee \(1992\)](#) and [Bikhchandani et al. \(1992\)](#), researchers were analyzing the properties of information aggregation and asymptotic behavior. In these models agents observe their own private signals and actions of some of the previous agents before choosing their own action. The structure of observational network, boundedness of signals, and coarseness of actions play crucial role for such asymptotic properties in sequential learning as herding, wisdom of crowds, information diffusion, and others. [Smith and Sørensen \(2000\)](#) show the asymptotic learning in the complete network happens if and only if the beliefs based on private signals are unbounded. [Acemoglu et al. \(2011\)](#) show that information diffuses if learning subnetworks length diverge and they are independent (even if the private beliefs are bounded).

Another strand of literature focuses on repeated learning in social networks, which features better tractability in many situations. Due to the complexity of Bayesian updating and strategic interaction, heuristic learning is often a more appealing choice for the analysis in repeated updating models. DeGroot model introduced by [DeGroot \(1974\)](#) is one

of the most common heuristic learning models. After receiving private signals at the beginning of the time, agents use weighted average of their neighbors beliefs to update their own belief every day. [DeMarzo et al. \(2003\)](#) provide a micro foundation for DeGroot model, which follows as a result of naive Bayesian repeated updating in their model. Convergence to consensus in DeGroot model follows closely from the Markov chain properties and was stated in many papers including [DeGroot \(1974\)](#), [DeMarzo et al. \(2003\)](#), and [Golub and Jackson \(2010\)](#). This results show that social influence of an individual is defined by her eigenvector centrality in the network. For a sequence of expanding networks, the necessary and sufficient condition for the wisdom of crowd is that everyone's centrality converges to zero, i.e. there remain now overinfluential people in the network [Golub and Jackson \(2010\)](#). Many papers including [Golub and Jackson \(2010\)](#) also show that the speed of convergence to the consensus is bounded by a number proportional to the second largest eigenvalue and the size of the network. [DeMarzo et al. \(2003\)](#) demonstrate that in the fully rational model with jointly Gaussian distribution of the true state and the private signals the agents would aggregate information in at most n^2 periods.

This paper differs from the previously mentioned in two directions. Firstly, I will focus on the finite population learning quality. Considering DeGroot model I will analyze the quality of the time asymptotic information aggregation comparatively to the best possible aggregation given the private signals. Secondly, since the learning is generally sub-optimal, I will introduce a Social Planner, who aims to improve the asymptotic learning subject to a give observational network constraint. The problem of increasing the speed of learning by SP will be also addressed in this paper.

3.3 Model

A social network consists of n individuals, observational links between them and is given by an undirected graph $G = (N, E)$, where $N = \{1, \dots, n\}$. The links define a possibility of observational learning between individuals, i.e. $(i, j) \in E$ means that individual i can pay attention to j 's opinion and vice versa. Also, the graph contains loops on every node

– $(i, i) \in E$ – which simply indicates that every individual has a memory about her own past opinion on a subject. A set of the nodes connected to i will be called i 's neighborhood in the rest of the paper and denoted by N_i .

As in [Golub and Jackson \(2010\)](#), DeGroot learning is defined as a mechanical updating of unidimensional opinions. An individual i updates her opinion $x_i \in \mathbb{R}$ over time according to a linear formula with time-independent weights of attentions she pays to her neighbors' yesterday opinions. These weights are given by an interaction matrix W , where $W_{i,j}$ means how much weight individual i puts on individual j 's yesterday opinion.

$$x_i(t) = \sum_{j=1}^n W_{i,j} * x_j(t-1) \quad (3.1)$$

Time is discrete $t \in \{0, 1, 2, \dots\}$, and therefore the vector of all opinions evolves according to the following formula

$$x(t) = W * x(t-1) = W^t * x(0) \quad (3.2)$$

The initial opinions $x(0)$ are noisy signals about the true state of the world θ . θ has an improper uniform distribution, and individual signals $x_i(0) = \theta + \epsilon_i$ include individual conditionally independent normal errors $\epsilon_i \sim N(0, \sigma_i^2) \perp (\theta, \{\epsilon_j\}_{j \in N \setminus \{i\}})$.

As described above, social interactions are given by two different objects – network G and interaction matrix W . The network G is a primitive fundamental element of the setup, which is defined exogenously and cannot be changed. However, learning improvement will involve changing the interaction matrix W subject to the given network. This means that the social planner can reshape attention people assign across their neighbors in order to improve learning, but cannot make individuals pay attention to somebody who they are not connected to in the network.

Definition 7 *An interaction matrix W compatible with network G is a positive valued matrix, such that*

1. $W_{i,j} = 0$ if $(i, j) \notin E$,
2. $\forall i, \sum_{j=1}^n W_{i,j} = 1$

Even though, G is undirected, a compatible with G matrix W is not required to be symmetric, e.g. i can put a different weight on j than j puts on i .

Definition 8 *The average-based updating interaction matrix W of the network G is defined as*

$$W_{i,j} = \frac{\mathbf{1}\{(i, j) \in E\}}{\sum_{k=1}^n \mathbf{1}\{(i, k) \in E\}} = \frac{\mathbf{1}\{j \in N_i\}}{\#N_i}$$

which in fact means that today individual's opinion is equal to a simple average of her neighbors' (including herself) yesterday opinions.


These two definitions formally describe the ways social planner can affect learning. Subject to the given social network and DeGroot learning model, SP can correct the weights one assigns to the opinions of people she knows. This models the assumption that learning remains heuristic and the individuals are not sophisticated for Bayesian learning; however, SP can consider the information and the network structure and change the interaction matrix in order to improve learning results. The main feature of both Definition 1 and 2 is that SP cannot make individuals pay attention to somebody who they are not connected in the original network. SP can reshape one's attention across her neighbors – people she knows – but cannot make her listen to somebody she does not know/trust.

This way of correcting the interaction matrix speaks to numerous examples. Social networks on the internet may affect the way they show me my friends' posts and choose which of them to show at the top and which of them to show at the bottom, or even not to show at all. They can completely ban users who spread some unreliable or poor information. Internet retail platforms may choose which reviewers are the most trustworthy, in order to show their opinions first and help other consumers' to learn better.

I will consider two benchmarks with respect to the freedom of choosing the interaction matrix. The first benchmark allows complete freedom of choosing the weights people

assign to their neighbors' opinions (there can not be a positive weight put on an absent link in the graph). This benchmark corresponds to Definition 1 of W compatible with G . The second benchmark is more restrictive and allows less freedom in choosing the weights. Particularly, I will consider that given a set of neighbors, an agent will always use a simple average-based updating rule. Therefore, the only freedom Social Planner has in correcting W is deleting links between agents. This corresponds to Definition 2 of the average-based updating interaction matrix W of G . The difference may be illustrated by the following example.

Figure 3.2: Social network G



$$W' = \begin{pmatrix} 0.5 & 0.3 & 0.2 & 0 & 0 & 0 \\ 0.4 & 0.6 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0.5 & 0.5 & 0 \\ 0 & 0 & 0 & 0.3 & 0.4 & 0.3 \\ 1 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

W' is an example of an interaction matrix compatible with G . It is also worth to notice that the presence of the link does not require to put a positive weight on it. W'' is another compatible with G interaction matrix, which also can be obtained as an average-based updating interaction matrix of G after deleting links $(3, 4)$ and $(1, 6)$ from it.

$$W'' = \begin{pmatrix} 1/3 & 1/3 & 1/3 & 0 & 0 & 0 \\ 1/2 & 1/2 & 0 & 0 & 0 & 0 \\ 1/3 & 0 & 1/3 & 0 & 1/3 & 0 \\ 0 & 0 & 0 & 1/2 & 1/2 & 0 \\ 0 & 0 & 1/4 & 1/4 & 1/4 & 1/4 \\ 0 & 0 & 0 & 0 & 1/2 & 1/2 \end{pmatrix}$$

Finally, it is worth mentioning why the sum of weights an agent puts on yesterday neighbors' opinions is equal to 1 (every row in the interaction matrix adds up to 1). This is a natural assumption corresponding to the fact that the opinion/estimate of the true state of the world should not inflate or deflate over time. This requirement also allows for a clear analysis of the time asymptotic results of the DeGroot learning, which I am going to focus on.

The problem that SP solves is to minimize the following loss function

$$L = \sum_i \left[\limsup_{t \rightarrow \infty} \mathbb{E}[(x_i(t) - \theta)^2] \right] \quad (3.3)$$

In other words, SP tries to minimize the sum of individual asymptotic mean square errors about the true state of the world. The fact that SP weights all individuals equally does not matter in terms of the results; however, considering time asymptotic results is crucial. It gives a lot of tractability since it is important to aggregate as much information from the individuals' signals as possible eventually. This also allows using of the previous results and characterization of time asymptotic opinions in DeGroot model ([Golub and Jackson, 2010](#)).

[Golub and Jackson \(2010\)](#) also show that in expanding networks individuals will learn the true state if everyone's influence vanishes as the network gets bigger. In this paper, on the other hand, I analyze finite networks and the question of how to make time asymptotic learning as precise as it can be. I also focus on the construction of interaction matrices solving this problem.

Lastly, I am considering that the initial signals have possibly heterogeneous precision, but contain uncorrelated errors. Heterogeneity of the private signals precision is interesting since it allows me to model the situations when there are more and less knowledgeable individuals. In the average-based updating benchmark, I consider only homogeneous signals due to the limitations on choosing W .

3.4 First best learning

For solving SP's problem in the next section, it is useful to derive the upper bound on the learning quality and under which conditions this upper bound can be attained. In order to do this, I will ignore the learning process and characterize the best information aggregation from the private signals, that minimizes SP's loss function. As it was said before, the true state $\theta \sim U[-\infty; +\infty]$ and the initial signals are independent normal unbiased estimates of θ

$$(x_1(0), \dots, x_n(0))^T \sim N((\theta, \dots, \theta)^T, \text{diag}(\sigma_1^2, \dots, \sigma_n^2))$$

$$\min_f \mathbb{E}[(f(x(0)) - \theta)^2]$$

The solution of the reduced problem is the true state conditional expectation on all initial signals, i.e. the best possible learning is if every individual opinion will converge to this conditional expectation.

$$\lim_{t \rightarrow \infty} x_i(t) = \mathbb{E}[\theta | x(0)] = \sum_{i=1}^n \left(\frac{\sigma_i^{-2}}{\sum_{j=1}^n \sigma_j^{-2}} * x_i(0) \right) \quad (3.4)$$

The solution follows from the joint distribution of the signals and the state of the world; it can be interpreted as follows. Since every signals is an unbiased estimate of the true state, then any weighted average of them will be too. Thus, the only question is how to weight

them in order to make the result the least noisy, and the optimal way to do it is to assign weights proportional to the precision of the signals, which is captured by (4). This defines the best possible inference from the signals (in terms of mean square error minimization), and sets the lower bound on the value of the SP's loss function.

Definition 9 *The first best learning is $\pi^* \cdot x(0)$, where*

$$\pi^* = \left(\frac{\sigma_1^{-2}}{\sum_j \sigma_j^{-2}}, \dots, \frac{\sigma_n^{-2}}{\sum_j \sigma_j^{-2}} \right) \quad (3.5)$$

Proposition 10 *SP's loss function L takes the lowest possible value if and only if all individual opinions converge to the first best learning.*

This is a simple result giving a base for the next analysis of improving DeGroot learning results. It also gives a clear comparison of the DeGroot and Bayesian time asymptotic results for a fixed network and information structure. [DeMarzo et al. \(2003\)](#) and [Mossel et al. \(2016\)](#) show that the myopic Bayesian learners will converge to the first best learning in at most n^2 periods. Therefore, the first best learning becomes a major baseline for comparison. In the next section I show that given any connected social network, SP can correct the interaction matrix to attain the first best learning. I also provide an algorithm constructing the solution to this problem.

Finally, it is useful to illustrate what is the first best learning on two examples. First, let us suppose that there are only two agents 1 and 2, with independent noise in their signals $x_1(0) = \theta + N(0, 1)$ and $x_2(0) = \theta + N(0, 2)$. Given that, precisions of their signals are 1 and 1/2 respectively, and therefore the first best learning is $(2/3) * x_1(0) + (1/3) * x_2(0)$ and $\pi^* = (2/3, 1/3)$. In the other important example we will often consider later, we will suppose that all agents receive signals with IID errors (homogeneous signals). Since the precision of their signals is identical, the first best learning is a simple average of their signals and $\pi^* = (1/n, \dots, 1/n)$.

3.5 Improving DeGroot

To proceed with solving SP's problem, it is useful to recall the results regarding convergence in DeGroot model (these or similar results are demonstrated in many papers including [Golub and Jackson \(2010\)](#), [DeMarzo et al. \(2003\)](#), and [DeGroot \(1974\)](#)). Since the first best learning is the the best potential solution to SP's problem, we will try to identify under which conditions opinions in DeGroot model converge to the first best learning.

The first best learning assigns a positive weight on every private signal. Therefore, for every opinion to converge to the first best learning, it is necessary that there is a path between every two nodes (every private signal reaches every node through the network). Thus, we require the graph to be connected (or strongly connected in case of directed graph). If it is not, there exist two nodes with no path from the first to the second one, and therefore the signal from the first one will never reach the second one through the network. The social network was introduced as an undirected graph earlier; however, since the interaction matrix compatible with the network is asymmetric and can also have zero weights on some links, the eventual graph may become directed. Therefore, strong connectedness of W is required for achieving the first best learning.

After we limit the analysis to strongly connected interaction matrices, [Golub and Jackson \(2010\)](#) state that if the interaction matrix is aperiodic³ then all opinions always converge to the consensus defined below

$$x_i(\infty) = \left(\lim_{t \rightarrow \infty} W^t \cdot x(0) \right)_i = \pi \cdot x(0) \quad (3.6)$$

Where π is a unique left eigenvector of W corresponding to eigenvalue 1 with sum of elements equal to 1.

³An interaction matrix is periodic if all cycles in it have a common divisor greater than one. In this case some information will just travel around with common periodicity but the opinions will not settle down to a consensus. The existence of at least one self-loop resolves this problem since it creates a cycle of length 1 and makes the matrix aperiodic. Since the introduced network contains all self-loops, to keep at least one of them will not be a problem.


$$\pi \cdot W = \pi$$

Vector π can be a vector of eigenvector centralities of nodes in a graph defined by interaction matrix W , where π_i is the centrality of node i . (This remark can be made with the only adjustment that in this case I take a left rather than a right eigenvector. Given that interaction matrix is very similar to Markov process matrix, it becomes apparent that the spectral radius of W is equal to 1 and π is an eigenvector corresponding to the maximum eigenvalue.) With sum of elements of π normalized to one, we can think of it as a vector of social influences of individuals in the network, where π_i is the weight of i 's initial opinion in the asymptotic social consensus.

The individual eigenvector centrality/social influence π_i depends on the network structure, however it may be excessively large for some individuals whose initial opinions are not very accurate but small for those whose information is actually rather precise. Another possible scenario is that everybody's private signals are identically precise, but the network structure is such that there are individuals much more central and influential than others. In this case, the information aggregation (results if learning) will be also suboptimal. Example 1 will illustrate such situation.

Example 1 *Let us consider a quite extreme example of social structure capturing natural variation in individual centralities – a star network. For $n = 7$ the network looks like in Figure 2.*

Figure 3.3: Star network ($n = 7$)



Let us consider that the private signals have identical precision $\sigma_i^{-1} = 1$ and that the individuals use a simple average-based updating rule, so that the interaction matrix is

$$W = \begin{pmatrix} 1/7 & 1/7 & 1/7 & 1/7 & 1/7 & 1/7 & 1/7 \\ 1/2 & 1/2 & 0 & 0 & 0 & 0 & 0 \\ 1/2 & 0 & 1/2 & 0 & 0 & 0 & 0 \\ 1/2 & 0 & 0 & 1/2 & 0 & 0 & 0 \\ 1/2 & 0 & 0 & 0 & 1/2 & 0 & 0 \\ 1/2 & 0 & 0 & 0 & 0 & 1/2 & 0 \\ 1/2 & 0 & 0 & 0 & 0 & 0 & 1/2 \end{pmatrix}$$

Given this interaction matrix we can derive the eigenvector centrality vector.

$$\pi = \left(\frac{7}{19}, \frac{2}{19}, \frac{2}{19}, \frac{2}{19}, \frac{2}{19}, \frac{2}{19}, \frac{2}{19} \right)$$

Now let us see how this network aggregates information comparatively to the first best learning. All opinions will converge over time to the consensus $\pi \cdot x(0)$ and its mean square error about the true state is equal to $\mathbb{E}[(\pi \cdot (\varepsilon_1, \dots, \varepsilon_7))^2] = 73/361$, which is equivalent that the social learning has precision equal to $4\frac{69}{73}$. This is close to the precision of the efficient aggregation of 5 independent individual signals. The first best learning here is a simple average aggregation of all private signals $\pi^* = (1/7, \dots, 1/7)$ (since they are equally precise), and therefore the precision of the first best learning is equal to the sum of the individual precisions, which is equal to 7. Hence, the actual learning loses 2 signals out of 7 in terms of precision comparatively to the first best learning. Unfortunately, it is quite frequent that individual's centrality does not reflect her knowledgeability. This brings us to the question of how to correct the interaction matrix in order to fix this learning inefficiency.

In this section, I will consider two ways of achieving the first best learning in the DeGroot model. Firstly, I will analyze the general problem with full freedom in changing the interaction matrix compatible with the network, and after that, I will analyze the situation

when the agents only use simplified average-based learning and SP can only delete the links in order to affect W .

3.5.1 Correcting the Interaction Matrix

To approach the problem, I first analyze the situation with identically precise individual signals (initial opinions). It is a simple version of the problem considered later however capturing the main features of both of them. The preview of results is that in both problems the first best learning is attainable, therefore the question can be reduced to how to match the social influence vector π to the first best learning vector π^* , or, in other words, how to construct the interaction matrix W compatible with the network G , such that

$$\pi^* \cdot W = \pi^* \tag{3.7}$$

This means that it is desired to construct a non-negative valued matrix W with every row's elements adding up to 1, and positive elements only corresponding to the links in the network G , but 0 on the positions corresponding to the pairs of nodes that are not connected directly, such that π^* is an eigenvector of W corresponding to the unit eigenvalue. We will solve equation (7) for W by matching the LHS entries (depending on W) to the RHS one by one. It is also crucial to preserve strong connectedness and aperiodicity of W , so that all opinions converge to consensus.

Remark 1 *If all elements of RHS equation (7) are equal to the corresponding RHS elements except for maybe one, then this last element is also equalized. This can be proved by the following logic. A proper interaction matrix is also a proper Markov probability transition matrix. Thus the operation on the LHS of (7) is a probability transition in Markov chain and a proper probability distribution today transits to a proper probability distribution tomorrow. This implies that the sum of LHS entries must be equal one. Therefore, if all except one of them are the same as the RHS entries, then this last element is also the same and π^* is a proper stationary distribution for W .*

In terms of solving equation (7), Remark 1 means that going over the interaction matrix and correcting its columns and therefore LHS vector elements is a valid strategy – it will work until the last column, which there is no freedom to change anymore due to the restrictions on W , but there will be no need to do this since the last element will be equalized automatically.

Proposition 11 *For any connected graph G (with all selfloops) there exists a strongly connected aperiodic interaction matrix W , s.t. the DeGroot time asymptotic results are the same as the first best learning, if the initial signals are identically precise.*

Proof. In case of identical precise signals, the first best learning is given by $\pi^* = (1/n, \dots, 1/n)$. For the simplicity of future notations, lets us denote by W_i , and $W_{,j}$ row i and column j of interaction matrix W respectively. Then we can rewrite equation (7) as

$$(1/n, \dots, 1/n) \cdot W_{,j} = 1/n$$

Or

$$\sum_{i=1}^n W_{ij} = 1$$

For every i . This condition does not seem to be very restrictive, which suggests that solution will not be unique. However, we do not aim to characterize all of them, but rather to find at least one of them to show that the first best learning is attainable. Thus, a simple algorithm described below is sufficient for the proof.

1. Fill in all possibly positive elements in the first row with equal numbers adding up to 1.
2. Transpose it to the first column.
3. Fill in the rest of the possibly positive elements in the second row by the same principle. Transpose.

4. Iterate until the last row/column.

This is a well defined algorithm, which makes every column of W add up to 1 by the construction, and therefore, produces W solving (7). Also, it does not remove any links from the original network, so it remains connected and aperiodic. ■

To see how it works we run it for the Example 1.

$$W = \begin{pmatrix} 1/7 & 1/7 & 1/7 & 1/7 & 1/7 & 1/7 & 1/7 \\ 1/7 & 6/7 & 0 & 0 & 0 & 0 & 0 \\ 1/7 & 0 & 6/7 & 0 & 0 & 0 & 0 \\ 1/7 & 0 & 0 & 6/7 & 0 & 0 & 0 \\ 1/7 & 0 & 0 & 0 & 6/7 & 0 & 0 \\ 1/7 & 0 & 0 & 0 & 0 & 6/7 & 0 \\ 1/7 & 0 & 0 & 0 & 0 & 0 & 6/7 \end{pmatrix}$$

For W constructed by the algorithm, DeGroot learning leads to the first best learning $\pi = \pi^* = (1/7, \dots, 1/7)$.

3.5.1.1 Heterogeneous signals

Now, as we captured the idea of how to solve equation (7) and the flavor of how the algorithm may work, we can move to a general case with heterogeneously precise signals. We first introduce Lemma 1 allowing us to reduce the problem to subgraphs and to make the induction step later.

Lemma 6 *Let graph G can be parted into two intersecting subgraphs G^1 and G^2 on the sets of nodes $\{1, \dots, i\}$ and $\{i, \dots, n\}$ such that there are no links between $\{1, \dots, i-1\}$ and $\{i+1, \dots, n\}$, i.e. the subgraphs are connected only through node i . Let $\pi^1 = (\pi_1, \dots, \pi_i)$ and $\pi^2 = (\pi_i, \dots, \pi_n)$ be parts of the first best learning vector π^* ; centrality vector conditional on subgraph G^k is $\tilde{\pi}^k = \frac{\pi^k}{\sum_{j \in V^k} \pi_j^k}$, and W^1 and W^2 are proper (strongly connected and aperiodic) interaction matrices compatible with G^1 and G^2 solving the conditional parts of equation (7) on the subgraphs*

$$\tilde{\pi}^k \cdot W^k = \tilde{\pi}^k \quad (3.8)$$

for $k = 1, 2$. Then there exists strongly connected aperiodic W compatible with G solving equation (7).

Proof. Firstly, let us notice that equation (8) is equivalent to the following

$$\pi^k \cdot W^k = \pi^k$$

Then, W solving equation (7) can be constructed from W^1 and W^2 as follows. Merge W^1 and W^2 in a diagonal block matrix intersecting on the cell $W_{i,i}$, with zero elements off the blocks. Adjust $W_{i,i}^* = 1 - \sum_{j \neq i} W_{i,j}$, so that row i elements add up to 1 as well as for every other row

$$W = \begin{pmatrix} W_{1,1}^1 & \dots & \dots & W_{1,1}^1 & 0 & \dots & 0 \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ W_{i-1,1}^1 & \dots & \dots & W_{i-1,i}^1 & 0 & \dots & 0 \\ W_{i,1}^1 & \dots & W_{i,i-1}^1 & W_{i,i}^* & W_{i,i+1}^2 & \dots & W_{i,n}^2 \\ 0 & \dots & 0 & W_{i+1,i}^2 & \dots & \dots & W_{i+1,n}^2 \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & \dots & 0 & W_{n,i}^2 & \dots & \dots & W_{n,n}^2 \end{pmatrix}$$

For all elements $j = 1, \dots, i-1$, since W^1 solved the equation (8) on subgraph G^1

$$\pi \cdot W_{,j} = \pi^1 \cdot W_{,j}^1 = \pi_j^1 = \pi_j$$

The same can be told for $j = i+1, \dots, n$, so that in equation (7) the constructed W equalizes all elements of the LHS and RHS except for maybe i . Moreover, Remark 1 then implies that entry i is equalized too since every row's elements of W sum up to 1. So W solves equation (7) and is strongly connected and aperiodic as soon as W^1 and W^2 are.


The last issue with the construction is that $W_{i,i}^*$ may be negative. This does not violate the previous arguments, but questions whether W is positive valued and therefore a proper interaction matrix. If this element is negative, then by construction $W_{i,i}^* \geq -1$. Therefore, there exists $\alpha \in [1/2, 1]$ such that $W' = \alpha \cdot W + (1 - \alpha) \cdot I$ is positive valued (where I is identity matrix) and still solves equation (7)

$$\pi \cdot W' = \pi \cdot (\alpha \cdot W + (1 - \alpha) \cdot I) = \alpha \cdot \pi \cdot W + (1 - \alpha) \cdot \pi \cdot I = \alpha \cdot \pi + (1 - \alpha) \cdot \pi = \pi$$

This last corrections guarantees that W' is a strongly connected aperiodic interaction matrix compatible with G maintaining the first best learning. (All requirements to W remain satisfied after weighting with identity matrix since it simply inflates diagonal elements and deflates the rest. The positive elements remain being positive, and all self-loops are preserved in graph G .) ■

Example 2 To illustrate how Lemma 1 and the algorithm in it works let us consider the following example.

Figure 3.4: Graph G



Let the vector of the signals' variances is $(\sigma_1^2, \dots, \sigma_6^2) = (3, 3, 6, 2, 6, 6)$ and therefore the first best learning is defined by $\pi^* = (0.2, 0.2, 0.1, 0.3, 0.1, 0.1)$.

Let $i = 1$ and subgraphs G^1 and G^2 contain nodes $\{1, 2, 3\}$ and $\{3, 4, 5, 6\}$ respectively. Then the conditional centrality vectors are $\pi^1 = (0.4, 0.4, 0.2)$ and $\pi^2 = (1/6, 1/2, 1/6, 1/6)$. Let also W^1 and W^2 given below solve the problems on the subgraphs

$$W^1 = \begin{pmatrix} 0.4 & 0.4 & 0.2 \\ 0.4 & 0.4 & 0.2 \\ 0.4 & 0.4 & 0.2 \end{pmatrix}$$

$$W^2 = \begin{pmatrix} 1/2 & 1/2 & 0 & 0 \\ 1/6 & 1/2 & 1/6 & 1/6 \\ 0 & 1/2 & 1/2 & 0 \\ 0 & 1/2 & 0 & 1/2 \end{pmatrix}$$

Then according to Lemma 1, firstly we can construct


$$W = \begin{pmatrix} 0.4 & 0.4 & 0.2 & 0 & 0 & 0 \\ 0.4 & 0.4 & 0.2 & 0 & 0 & 0 \\ 0.4 & 0.4 & -0.3 & 1/2 & 0 & 0 \\ 0 & 0 & 1/6 & 1/2 & 1/6 & 1/6 \\ 0 & 0 & 0 & 1/2 & 1/2 & 0 \\ 0 & 0 & 0 & 1/2 & 0 & 1/2 \end{pmatrix}$$

This W solves equation (7) but it is not actually a proper interaction matrix since $W_{3,3} = -0.3 < 0$. Then we will make a last step and take a convex combination of this matrix with the identity to eliminate the issue

$$W' = \frac{1}{2} \cdot W + \frac{1}{2} \cdot I = \begin{pmatrix} 0.7 & 0.2 & 0.1 & 0 & 0 & 0 \\ 0.2 & 0.7 & 0.1 & 0 & 0 & 0 \\ 0.2 & 0.2 & 0.35 & 1/4 & 0 & 0 \\ 0 & 0 & 1/12 & 3/4 & 1/12 & 1/12 \\ 0 & 0 & 0 & 1/4 & 3/4 & 0 \\ 0 & 0 & 0 & 1/4 & 0 & 3/4 \end{pmatrix}$$

This is a well-defined strongly connected aperiodic interaction matrix which leads to the first best learning in the example.

Figure 3.5: Graph G''



For the next proofs it is also important to notice that for graph G'' (that contains graph G and some additional edges) Lemma 1 can be also applied in a slightly modified way. G^1 and G^2 are still connected aperiodic subgraphs in G' connected between each other through vertex 3. Therefore, we can construct the same interaction matrix W' for G' as before, even though there are some additional edges between G^1 and G^2 . In this case these edges – $(2, 5)$ and $(1, 4)$ – will be just automatically deleted from the graph (the resulting interaction matrix).

We discussed the induction step for the next proof, so we can state Theorem 1 and prove it.

Theorem 4 For any connected graph G (with all selfloops) and for any vector of initial signal variances $(\sigma_1^2, \dots, \sigma_n^2)$, there exists a strongly connected aperiodic interaction matrix W compatible with G , s.t. the DeGroot time asymptotic results are the same as the first best learning.

Proof. To prove Theorem 1 it is just needed to understand how to apply Lemma 1 to reduce solving equation (7) on some degenerate subgraphs. One of the ways to do it is to pick any tree containing all the vertices of G (together with self-loops), to solve (7) for every link connecting two vertices in this tree, and to solve equation (7) for the whole graph by induction, adding nodes one by one and applying Lemma 1 for every iteration. Firstly, since G is connected it contains at least one tree on all its nodes as a subgraph. Choosing this tree is the first element of freedom in the construction of W resulting in the first best learning.

Second, the reduced problem (on a subgraph) can be solved easily for any clique (subgraph which is complete) including a clique of only two adjacent nodes – an edge. Let $1'$ and $2'$ be two adjacent nodes in the original graph and G' be a graph containing only this two nodes (together with the link between them and the self-loops). Let $\pi' = (\pi'_{1'}, \pi'_{2'}) \equiv \left(\frac{\pi_{1'}}{\pi_{1'} + \pi_{2'}}, \frac{\pi_{2'}}{\pi_{1'} + \pi_{2'}} \right)$ be the conditional on G' desired centrality vector. The the reduced problem is

$$\pi' * \begin{pmatrix} W'_{1',1'} & W'_{1',2'} \\ W'_{2',1'} & W'_{2',2'} \end{pmatrix} = \pi'$$

or equivalently

$$\frac{W'_{1',2'}}{W'_{2',1'}} = \frac{\pi'_{2'}}{\pi'_{1'}}$$

Solving this equation is the second element of freedom in the construction of W . However, for the proof it is just enough to notice that the following interaction matrix will solve it for the subgraph

$$W' = \begin{pmatrix} \pi'_{1'} & \pi'_{2'} \\ \pi'_{1'} & \pi'_{2'} \end{pmatrix}$$

The final step of the proof is to formalize an induction step. Given any connected subtree in the tree (including one link) for which the problem is already solved, we can choose another link in the tree adjacent to this subtree and connect them, applying Lemma 1 in order to solve conditional (7) in this bigger subtree. Lemma 1 here can be applied since we always can reorder the nodes in the subgraphs to satisfy the conditions for the lemma, and it is true that every additional adjacent link to a subtree is connected to it only through one node. Also, in step 2 of the proof we have already shown how to solve a problem for one link, and since we have a tree this algorithm will eventually include all nodes in it and therefore in the initial graph.

anymore since given the corrected network individuals use the average-based updating rule

$$x_i(t) = \frac{\sum_{j \in N_i} x_j(t-1)}{\#N_i} \quad (3.9)$$

Thus, social planner who cannot create new social connections (edges) can only delete them in order to improve the learning results. This stylized assumption captures that the individuals perform some stubbornness in their learning rules, but SP still can restrict information spreading through some channels.

Given this learning model, DeGroot time asymptotic results can be characterized more clear in terms of the graph structure, rather than the interaction matrix which is less interesting in this setup since it is obtained mechanically from the graph. Let us denote by the node degree $d_i = \#N_i$ the number of the nodes it connected too, possibly including itself. This is slightly different from the conventional definition of the vertex degree in undirected graph since according to the definition in this paper, a self-loop on the node adds only 1 to its degree rather than 2.

Lemma 7 *Let G be an undirected connected graph. Then the eigenvector centralities vector for the average-based updating interaction matrix W of graph G is*

$$\pi = \left(\frac{d_1}{\sum_i d_i}, \dots, \frac{d_n}{\sum_i d_i} \right) \quad (3.10)$$

Proof. The interaction matrix is given by $W_{i,j} = \frac{\mathbf{1}\{j \in N_i\}}{d_i}$. Then we can verify that π defined in (10) is actually an eigenvector centrality of average-based updating interaction matrix W of G :

$$\pi_i = \pi \cdot W_{:,j} = \sum_{j \in N_i} \frac{1}{d_j} * \pi_j = \sum_{j \in N_i} \frac{1}{d_j} * \frac{d_j}{\sum_k d_k} = \sum_{j \in N_i} \frac{1}{\sum_k d_k} = \frac{d_i}{\sum_k d_k} \quad \blacksquare$$

Lemma 2 states that in average-based updating learning processes the social influence of the node is simply proportional to its degree centrality – how many individuals one


connected to. This is a natural concept of social influence in graphs however implies some restrictions on the asymptotic learning. Particularly it is easy to notice that vector π changes very discretely and clearly cannot take all possible real values. In a directed multigraph any interaction matrix from Section 5.1 could be attained by duplicating or deleting directed links; however, we will exclude such possibility from the analysis and focus on undirected simple graphs. Given that, obviously the first best learning is not always attainable for heterogeneous signals and thus we will simplify the analysis restricting it to homogeneous signals case. To characterize the results it will be useful to remember the definition of regular graph.

Definition 10 *The graph is regular if every vertex has the same degree.*

Proposition 12 *Let G be an undirected connected graph and the initial signals are identically precise. Then under average-based updating learning, the first best learning is attainable by deleting some links from G if and only if there exists a regular connected subgraph in G including all vertices of G .*

Proof. The proposition is quite straightforward from Lemma 2. As soon as there is a regular connected subgraph, the rest of the links can be deleted and average-based updating learning will lead to the first best learning. Moreover, if there is no such subgraph, the now correction will lead to even degree distribution, and therefore, the first best learning is unattainable under homogeneous signals. ■

Finally we can illustrate the condition from Proposition 3 on the difference between the following two graphs:



The star network, graph (a), does not contain a connected regular subgraph on all of its vertices – deleting any link causes disconnectedness, and the original graph is clearly not regular. Therefore, SP cannot attain the first best learning by deleting links under average-based updating. Graph (b) however can be corrected to result in the first best learning. For instance, deleting links (2, 3) and (3, 4) leaves only Hamiltonian cycle, which is clearly a regular graph. Every vertex degree is equal to 2, and the centrality vector is $\pi = (0.2, \dots, 0.2) = \pi^*$.

The latest example and Proposition 3 show that in case of average-based updating the first best learning (in case of homogeneous signals) is attainable only if the structure of network G is not too asymmetric. If some nodes have too high influence relatively to others in the original network, it may be hard to reduce it without disconnecting the network or reducing some other nodes' centralities.

3.6 Speed of Convergence

It is also interesting to make a few remarks regarding possibility of increasing speed of learning in DeGroot model. For this discussion it is important to recall the results on the speed of convergence in the DeGroot model from [Golub and Sadler \(2017\)](#). In Proposition 8 of this paper, they state that the second largest eigenvalue of the interaction matrix W is closely related to the speed of convergence (the largest eigenvalue is equal 1).

$$\sup_{x(0) \in [0,1]^n} \|x(t) - x(\infty)\|_{\infty} \leq (n-1)|\lambda_2|^t$$

The smaller (in absolute value) the second largest eigenvalue is, the faster individual opinions will converge to the consensus. Some theoretical results explaining it are provided below.

Given the properties of the interaction matrix it can be decomposed as

$$W = \Pi^{-1} \cdot \Lambda \cdot \Pi \tag{3.11}$$

Where Π is the row eigenvector matrix of W and $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$ a diagonal matrix of eigenvalues. For simplicity let reorder eigenvalues from the largest $\lambda_1 = 1$ in absolute value, to the smallest one. Then the convergence of the opinions in the DeGroot model relies on the fact that

$$W^t = \Pi^{-1} \cdot \Lambda^t \cdot \Pi = \Pi^{-1} \cdot \text{diag}(\lambda_1^t, \dots, \lambda_n^t) \cdot \Pi$$

And the speed of convergence depends on how soon $\Lambda^t \rightarrow \text{diag}(1, 0, \dots, 0)$, which depends on the second largest eigenvalue (in absolute values), since the rest of them converge to zero even faster. This also explains why the opinions converge to the consensus $\pi \cdot x(0)$

$$\lim_{t \rightarrow \infty} W^t \cdot x(0) = \Pi^{-1} \cdot \text{diag}(1, 0, 0, \dots, 0) \cdot \Pi \cdot x(0) = \Pi_{,1}^{-1} \cdot \pi \cdot x(0) = (\pi \cdot x(0), \dots, \pi \cdot x(0)) \quad (3.12)$$

Which follows from $\Pi_{,1}^{-1} = (1, \dots, 1)$, since every other eigenvector except for π sums up to 0 and $\Pi \cdot \Pi_{,1}^{-1} = (1, 0, \dots, 0)^T$.

Now let us say the λ_2 is the second largest eigenvalue for W and π' is the corresponding eigenvector

$$\pi' \cdot W = \lambda_2 \pi'$$

Let β be the minimal of the diagonal elements of W and $\alpha \in \left(-\frac{\beta}{1-\beta}, 1\right)$, then the following equation is true for the interaction matrix $W' = (1-\alpha) * W + \alpha * I$

$$\pi' \cdot W' = [(1-\alpha)\lambda_2 + \alpha] \cdot \pi'$$

So for the new interaction matrix W' the speed of convergence will be associated with $\lambda'_2 = (1-\alpha)\lambda_2 + \alpha$. This eigenvalue is smaller than λ_2 in absolute value if $\alpha * \lambda < 0$ is

negative and therefore the speed of convergence is higher. Let us illustrate it on the next example.

Example 3

$$W = \begin{pmatrix} 5/6 & 1/6 \\ 1/3 & 2/3 \end{pmatrix} = \begin{pmatrix} 1 & 1/3 \\ 1 & -2/3 \end{pmatrix} \cdot \begin{pmatrix} 1 & 0 \\ 0 & 1/2 \end{pmatrix} \cdot \begin{pmatrix} 2/3 & 1/3 \\ 1 & -1 \end{pmatrix}$$

The second largest eigenvalue here $\lambda_2 = 1/2$ and the smallest diagonal element is $\beta = 2/3$. So let us pick $\alpha = -1 \in (-2, 1)$, for example. Then the new interaction matrix leading to the same asymptotic learning will be

$$W' = 2 \cdot W - I = \begin{pmatrix} 2/3 & 1/3 \\ 2/3 & 1/3 \end{pmatrix} = \begin{pmatrix} 1 & 1/3 \\ 1 & -2/3 \end{pmatrix} \cdot \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \cdot \begin{pmatrix} 2/3 & 1/3 \\ 1 & -1 \end{pmatrix}$$

Now the speed of convergence is higher, more precisely the convergence to the same asymptotic consensus is immediate since $\lambda'_2 = 0$ and for any $t \geq 1$ $x_i(t) = 2/3 * x_1(0) + 1/3 * x_2(0)$.

The intuition of the method accelerating learning is as follows. If we take an interaction matrix (with a positive second largest eigenvalue) leading to the first best learning and take a convex combination of it with the identity matrix, it will slow down the learning because all eigenvalues will take convexified with 1. So, if on the other hand we can decompose a proper interaction matrix W as a convex combination of the identity matrix and another proper interaction matrix W' , then this W' will also lead to the first best learning but faster. This is exactly the case shown by the Example 3.

Proposition 4. Let $W_{i,i} \in (0, 1) \forall i$ and $\lambda_2 \neq -\lambda_3$ (if $n > 2$). Then there exists W' such that (1) opinions converge to the same consensus as for W , (2) asymptotic speed of convergence to the consensus is higher than it is for W .

More generally, such transformation may help speed up learning in two following way: (1) decreasing self-weights can make “stubborn” individual learn faster from their neighbors; (2) increasing self-weights can prevent from opinions circulating around the network if individuals are too susceptible to other opinions.

3.7 Conclusion

Efficient learning in the networks is crucial for many economic outcomes in the modern world. The paper shows the possibility of efficient asymptotic aggregation of information in connected finite networks in DeGroot model. The algorithm from Section 5.1 constructs an interaction matrix leading to the first best learning. I also address the problem of increasing the speed of convergence. However, the paper has a number of limitations opening the possibilities for the future research.

Firstly, the freedom to correct the weights in interaction matrix in any possible way may be arguable. This may be a too strong assumption in real world, where SP for some reason can be limited in the ways of controlling learning. The model considered in the paper does not provide particular foundations for the naive learning with SP, therefore does not explain how exactly SP can affect the attention individuals pay to each other’s opinion. Introducing such a naive learning foundation model with SP and some particular methods of “censoring” information can be a good start for improving theory in this direction.

Another limitation of this paper is that the analysis is mostly focused on time asymptotic learning and has a little to say about making the convergence to the first best learning fast. Section 6 provides one possible way of doing it, however does not tell much about the fastest ways of learning. One problem that could be considered is what interaction matrix leads to the first best learning with the highest possible asymptotic speed of convergence. Another possible approach is to introduce a discounted loss function, which brings to the speed of convergence issue automatically since such loss function includes earlier periods in it.

Finally, the paper considers that the agents are unsophisticated in the way they update their opinions, and do not know much about the network structure. However, the paper presumes that SP knows the network and the information structure perfectly. This is a rather strong assumption, which can be a subject to questions. Hence, it could be interesting to consider a naive learning model in the network where both individuals and SP have limited information about the network and/or information structure. Addressing this and latter extensions would provide important insights to the theory of centralized improving the aggregation of information in naive learning models.

References

- Acemoglu, Daron, Munther Dahleh, Ilan Lobel, and Asuman Ozdaglar (2011) "Bayesian learning in social networks," *The Review of Economic Studies*, 78, 1201–1236.
- Bala, Venkatesh and Sanjeev Goyal (1998) "Learning from neighbours," *The Review of Economic Studies*, 65, 595–621.
- Banerjee, Abhijit (1992) "A simple model of herd behavior," *The Quarterly Journal of Economics*, 107, 797–817.
- Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992) "A theory of fads, fashion, custom, and cultural change as information cascades," *The Journal of Political Economy*, 100, 992–1026.
- DeGroot, Morris H. (1974) "Reaching a consensus," *Journal of the American Statistical Association*, 69, 118–121.
- (2005) *Optimal Statistical Decisions*, Hoboken, New Jersey: Wiley Classics Library, Wiley.
- DeMarzo, Peter M., Dimitri Vayanos, and Jeffrey Zwiebel (2003) "Persuasion bias, social influence, and unidimensional opinions," *The Quarterly Journal of Economics*, 118, 909–968.
- Gale, Douglas and Shachar Kariv (2003) "Bayesian learning in social networks," *Games and Economic Behavior*, 45, 329–346.
- Golub, Benjamin and Matthew O. Jackson (2010) "Naïve learning in social networks and the wisdom of crowds," *American Economic Journal: Microeconomics*, 2, 112–149.
- (2012) "How homophily affects the speed of learning and best-response dynamics," *The Quarterly Journal of Economics*, 127, 1287–1338.
- Golub, Benjamin and Evan Sadler (2017) "Learning in Social Networks."
- Hann-Caruthers, Willem, Valentin V. Martynov, and Omer Tamuz (2018) "The speed of sequential asymptotic learning," *Journal of Economic Theory*, 173, 383–409.
- Jackson, Matthew O. (2010) *Social and Economic Networks*, Princeton, NJ: Princeton University Press.
- Lobel, Ilan, Daron Acemoglu, Munther Dahleh, and Asuman Ozdaglar (2009) "Rate of convergence of learning in social networks," in *Proceedings of the American Control Conference*.
- Lobel, Ilan and Evan Sadler (2014) "Preferences, homophily, and social learning."

- (2015) “Information diffusion in networks through social learning,” *Theoretical Economics*, 10, 807–851.
- Meyer, Carl D. (2000) *Matrix Analysis and Applied Linear Algebra*, Philadelphia: SIAM.
- Mossel, Elchanan, M. Mueller-Frank, Allan Sly, and Omer Tamuz (2020) “Social learning equilibria,” *Econometrica*, 88 (3), 1235–1267.
- Mossel, Elchanan, Nadav Olsman, and Omer Tamuz (2016) “Efficient Bayesian learning in social networks with Gaussian estimators,” Available at <http://arxiv.org/abs/1002.0747>.
- Mossel, Elchanan, Allan Sly, and Omer Tamuz (2015) “Strategic learning and the topology of social networks,” *Econometrica*, 83, 1755–1794.
- Smith, Lones and Peter Sørensen (2000) “Pathological outcomes of observational learning,” *Econometrica*, 68, 371–398.