Information and Communication - Essays on Review Systems, Front-of-Package Labelling and Unstructured Bargaining

Inaugural-Dissertation
zur Erlangung des Grades
Doctor oeconomiae publicae (Dr. oec. publ.)
an der Ludwig-Maximilians-Universität München

2020
vorgelegt von

Hung-Ni Chen

Referent:	Prof. Dr. Klaus M. Schmidt
Korreferent:	Prof. Dr. Florian Englmaier
Promotionsabschlussberatung:	03.02.2021

Datum der mündlichen Prüfung: 18. Januar 2021
Namen der Berichterstatter: Klaus Schmidt, Florian Englmaier, Aniko Öry

Acknowledgements

First, I would like to express my gratitude towards my supervisors. I would like to thank my first supervisor Klaus Schmidt for always being accessible and involved, and for his guidance, discipline and invaluable advice. I am truly grateful to have him as my supervisor. I would also like to thank my second supervisor Florian Englmaier, who has been kind and encouraging since my first year and takes time discovering the potential in my research when I fail to. My third supervisor Aniko Öry broadened my horizon with her inspiring courses and wonderful research. I am thankful for her support in various ways, and I feel very lucky to have a great role model like her during my journey in academia.

Second, I would like to thank my co-authors Colin Camerer, Po-Hsuan Lin, Gideon Nave, Alec Smith and Joseph Tao-Yi Wang, for teaching, guiding and working together with me. I would also like to thank Ishita Chakraborty for constantly brain-storming with me, and for reminding me how much fun research can be.

Furthermore, I would like to thank my colleagues, Benjamin Häusinger, Christoph Schwaiger, Lion Henrich and Hoa Ho, who provided rigorous discussions, helpful advice and recharging lunch and coffee breaks. I also thank Peter Schwardmann for all the great discussions and fruitful comments, and our secretaries Silke Englmaier, Manuela Beckstein and Sabine Wilhelm-Kauf for their solid support on administrative work. I acknowledge the financial support of the German Research Foundation (DFG) and mentoring from our GRK director Carsten Eckel.

This dissertation would not have been possible without my friends and family. I would like to thank Felix, Ling, Else and other friends I made in Germany for enriching my life and for making me feel welcome in a foreign country in the last four years. Special thanks go to my dearest friends in Taiwan: Tingwun, Hsiao-Chuan, Amelia and Xin-Ying as well as my brother Hung-Chi. Thank you for always making time for me, and for nourishing and believing in me.

Last but not least, I would like to thank my partner Andreas, for his company, care, love and for bringing laughter in my life with his excellent sense of humor.

Contents

List of Figures

List of Tables

Preface

In the last half-century, information asymmetry has been one of the most important topics in economics. While it may be beneficial for individuals to withhold information sometimes, very often, communication helps prevent undesired outcomes such as market breakdowns or bargaining impasses.

This dissertation consists of three self-contained chapters that investigate communication in different aspects. The first chapter investigates the impacts of review systems, which are crucial in online marketplaces today. The second chapter discusses Front-of-Package labelling, which is a highly debated current event. The third chapter analyzes bargaining processes with machine learning methods. In the following, I provide a brief overview on each chapter.

Chapter 1: In the first chapter, I analyze how review systems influence firms' incentives to join a collaborative brand and in turn affect market outcomes. I investigate a case where firms invest in product quality at heterogeneous costs and the investment levels are private. Product information can be communicated in three ways. Firms can signal through prices and brand names, and consumers can inspect on quality. I model review systems as an instrument which lowers consumers' inspection costs.

I show that review systems encourage cost-efficient firms to invest in product quality, which results in higher market efficiency. Furthermore, brand names become less valuable when review systems are present. Consumers do not rely on the information brand names provide, and firms have incentives to detach from a brand and build their own reputation. In addition, while lack of review systems undermines firms' incentives to invest, review systems and market powers of brands lead to inefficiency due to over-investment.

Chapter 2: In the second chapter, I discuss regulations on Front-of-Package labels. I construct a model where firms can voluntarily adopt a label which gives incomplete information about their products. I first provide an explanation why the classic unravelling results where all firms adopt the label ((Grossman [1981], Milgrom and Roberts [1986])) may not take place. In my model, products have multiple attributes and consumers have heterogeneous preferences. A firm whose products are only valued by certain consumers lack the incentive to adopt the label. Instead, they prefer that their targeted consumers conduct search and find them.

I continue with the discussion whether mandating Front-of-Package labels increases welfare. I show that if firms lack the incentive to adopt labels for the reason described above, mandating the label decreases welfare because consumers have to search more. This result comes from pooling on prices, which weakens the price signal.

Chapter 3: The third chapter explores how machine learning methods help to extract information in communication process and to predict bargaining outcomes. Subjects in a lab experiment are paired to engage in a 10-second bargaining game, and one person in each pair learns about the bargaining surplus. Two theoretical predictions are tested. Our results suggest that subjects with private information have the tendency to use equal splits as focal points and to signal on the bargaining surplus. Moreover, we find that as process data become richer at a later stage of bargaining, the machine learning model outperforms static models which only make predictions based on the size of the surplus. We identify important features and document their effects using LASSO regression. Among other things, we confirm classic results such as deadline effects.

While the three chapters have very different applications, they all center around the topic of communication and complement one another if put in a broader context. The first chapter investigates how cost reduction in one communication channel affects communication incentives in other channels. The second chapter explain how a new information source may counteract with existing sources and lead to little information transmission. Moreover, in a market setting where firms' pricing strategies react to regulations, mandating firms to disclose information can do harm to social welfare. The third chapter sets focus on the content of communication and pioneers in extracting information using novel methods.

Chapter 1

Tow Big Brands or Walk Alone - The Impacts of Review Systems

1.1 Introduction

Review systems are ubiquitous in online markets today. They are extensively adopted on online platforms and widely accepted by consumers. Rating platforms profit from their reviews and intermediaries place an emphasis on product reviews as well. For example, the annual revenue of TripAdvisor, the largest online review platform, has increased by more than 5 times in the last decade.[1] Also, large platforms such as Amazon, Booking.com and Expedia place star ratings as an important attribute in their ranking algorithms.[2] Review systems are an inherent element in modern market places, and they impact market dynamics in numerous industries.

Existing literature has classified reviews as online word-of-mouth, which is

[1]According to Statistica, TripAdvisor's annual revenue in 2008 was 298.25 million, whereas that in 2018 was 1.6 billion.

[2]See the A9 algorithm of Amazon and https://partner.booking.com/en-us for details on Booking.com

often considered as a substitute for certification, brand-signalling and advertising. (Hollenbeck et al. [2019], Dranove and Jin [2010], Jansen et al. [2009]) In this paper, we take a new approach in the interpretation of review systems, and model them as an instrument to lower consumer inspection costs. We ask how review systems influence firms' incentives to join a collaborative brand and their impacts on market outcomes. We construct a game-theoretic model with a brand-holder, many firms and many consumers. The brand-holder owns an established trademark and contracts with the firms on its usage. Firms with heterogeneous costs consider whether to join the brand or not, and choose investment levels and prices. Upon arrival, consumers observe prices and trademark if available. Before making purchasing decisions, consumers can incur an inspection cost and learn about the product qualities.

We set the benchmark as the case without review systems, and compare it with the case where reviews are available. First, there is a non-monotonic relationship between cost-efficiency and branding decisions. Firms with moderate costs adopt the trademark while firms with high or low costs stand alone. Complementing empirical findings in Hollenbeck [2018], Fang [2019], Newberry and Zhou [2019], we show that the value of a brand name decreases because product information is more accessible.

As consumers rely less on brand-signalling, firms also find it less attractive to contract with the brand-holder. Review systems encourage cost-efficient firms to detach from the brand. Without the brand name, consumers cannot tell apart high and low quality products immediately. However, review systems allow firms to build their own reputation and counter-signal their quality by not joining the brand. Second, with more information disclosed, firms have stronger incentives to invest in product quality. Without review systems, firms generally under-invest in quality. Interestingly, an inefficiency arises from over-investment when review systems are present. A fraction of firms over-invest in quality to build their reputation. The profits they earn in this case are still higher than if they join the brand and pay a significant contract fee.

We find that review systems unambiguously improve welfare, but when compared to the first-best case, an efficiency gap persists even when the cost of reading reviews becomes arbitrarily small. Cost-efficient firms tend to over-invest and cost-inefficient firms tend to under-invest. This efficiency gap results from the market power of the brand-holder and can only be closed if there is sufficient competition among brand-holders.

The remainder of this paper is structured as follows. Section 1.2 discusses related literature. Section 1.3 introduces the set-up of our model. Section 1.4 presents our results. Section 1.5 concludes.

1.2 Related Literature

Our paper adds to the vast signalling literature. Signalling games are applied in various contexts to explain education choices (Spence [1978]), advertising (Nelson [1970]), gift-exchange (Camerer [1988]), and more recently on voting behavior (DellaVigna et al. [2016]). While most work suggests a monotonic relationship between signalling and senders' types, in our paper, only firms with moderate costs opt for brand-signalling. The first paper to document a similar non-monotonic relationship is Feltovich et al. [2002]. In a setting with one sender and one receiver, Feltovich et al. [2002] establish that if the receiver has additional information source which is noisy, high and low types will pool on the same message. In our setting, however, noisy information is not the key. Rather, it is firms' endogenous choices of quality investment and the market power of the brand-holder which drive the result.

In terms of modelling choices, our paper is most related to Stahl and Strausz [2017], which extends from Bester and Ritzberger [2001]. Stahl and Strausz [2017] compare buyer and seller certification and show that seller certificate improves efficiency because no certificate is also a kind of revelation. In a similar model setting, our model incorporates both buyer and seller certification. The brand-holder serves as a third-party certifier who maximizes profit. Instead of comparing buyer and seller certification, we discuss market outcomes and focus on change in sellers' incentive to acquire certificates when buyers' certification costs decrease.

1.3 Set-up

There is a brand-holder and a continuum of firms and consumers. The brand-holder writes a contract for the usage of its brand name. Firms decide whether they join the brand or not, and how much to invest in the quality of their products. Firms then set prices for their products. Consumers arrive and observe prices. They can pay an inspection cost and acquire information about product quality before they make purchase decisions.

1.3.1 Firms

There is a continuum of firms with mass 1, and each firm produces one unit of good. Each firm first takes a branding choice, which is denoted by $b_i \in \{0, 1\}$, where $b_i = 1$ means firm i joins the brand. After firms took their branding decisions, they decide on how much to invest in quality and set prices for their products. Each firm has a firm-specific cost component denoted by c_i, which is uniformly distributed on $[0, 1]$. A product can be of three quality levels: high, medium and low, which are denoted by q_h, q_m, and q_l respectively, with $q_h > q_m > q_l$. Denote e_i firm i's investment choice. Then $e_i \in \{e_h, e_m, e_l\}$, where $e_h > e_m > e_l$. Firms' investment choices translate into quality with certainty. For example, if firm i chooses $e_i = e_h$, then the product is of high quality q_h. The total cost of firm i is $c_i e_i$, where both components are private information to the firm. We assume that $q_l \geq e_l$ so that it is socially optimal for all firms to stay in the market. We allow for mixed pricing strategies, but firms are restricted to randomizing over countably many prices. [3] We denote σ_i the mixed strategy of firm i and p_i the price firm i sets if it is a pure strategy.

1.3.2 Brand-holder

There is one brand-holder in this world. The brand-holder does not participate in manufacturing. It owns a trademark which firms can use upon its agreement. For simplicity, we abstract away the monitoring behavior of the brand-holder. Also, we assume that the brand-holder cannot impose investment choice e_h. Therefore, the brand-holder offers a contract which entails fixed fee F and minimal investment level e_m. The brand-holder can only offer one contract to all firms. Namely, screening through contracts is not possible. For the ease in analysis, we assume that a firm accepts the contract if it is indifferent.

We focus on the case where the brand-holder can only secure the investment level e_m. This setting corresponds to the case where monitoring costs are high, or when the brand-holder cannot contract on every aspect that contributes to consumers' perception of quality.

1.3.3 Consumers

Without loss of generality, we assume that exactly one consumer arrives at each firm. Consumers observe listed prices and firms' branding decisions upon arrival. Denote consumers' beliefs $\mu(p, b) = (\mu_h, \mu_m, \mu_l)$, where p and b are price and branding choices and $\mu_h + \mu_m + \mu_l = 1$. Before making purchase decisions, consumers can choose whether to inspect the quality of the good. If a consumer inspects, she incurs inspection cost k and fully learns about

[3] This restriction is often adopted to avoid measure problems. See Bester and Ritzberger [2001] and Stahl and Strausz [2017]

the quality. We allow consumers to randomize their inspection strategy and denote $\gamma(p, b)$ the probability that a consumer inspects upon receiving price signal p and branding signal b. If a consumer purchases without inspection, her expected utility is $E[q] - p$, where $E[q] = \mu_h q_h + \mu_m q_m + \mu_l q_l$. If she purchase after learning about the quality being q, her utility is $q - p - s$. If consumers do not purchase from the firm they arrive at, they leave the market with reservation utility normalized to 0. We assume that consumers purchase if they are indifferent.

Our results do not hinge on specific relationships between exogenous parameters. However, to prevent from discussing a degenerate case, we assume that

$$1 > \frac{q_m - q_l}{e_m - e_l} > \frac{q_h - q_m}{e_h - e_m} > 0. \tag{1.1}$$

The inequalites in 1.1 guarantees that in social optimum, all investment levels are chosen by some firms, and that all firms should stay in the market. To see this, observe that if every firm can set the price exactly to the quality it chooses, a firm i with $c_i = \frac{q_m - q_l}{e_m - e_l}$ is indifferent between investing in low quality and medium quality since $\pi_i(e_m) = q_m - c_i e_m = q_l - c_i e_l = \pi_i(e_l)$. Similarly, a firm with $c_i = \frac{q_h - q_m}{e_h - e_m}$ is indifferent between choosing high and medium quality. Therefore, it is optimal for a firm to choose medium quality if its cost-component is between $\frac{q_m - q_l}{e_m - e_l}$ and $\frac{q_h - q_m}{e_h - e_m}$. Likewise, $\frac{q_m - q_l}{e_m - e_l} < 1$ implies that some firms should produce low-quality products in this case, and $\frac{q_h - q_m}{e_h - e_m} > 0$ guarantees that choosing high quality is optimal for some firms as well.

1.3.4 Time Structure

The game has 4 stages.[4] At the first stage, the brand-holder designs the contract and sets the contract fee F. In the second stage, firms first decide whether to accept the contract or not. If a firm accepts the contract, then the investment level has to be at least e_m. If a firm does not accept the contract, it can freely decide on the investment level. Firms then set prices at the second stage. At the third stage, consumers arrive and observe prices and branding choices, then they decide whether to inspect the quality at cost k or not. At the last stage, consumers make purchase decisions and payoffs realize accordingly.

We focus on Perfect Bayesian Equilibria. A Perfect Bayesian Equilibrium consists of $(F^*, \{b_i^*\}, \{\sigma_i^*\}, \{e_i^*\}, \mu^*, \gamma^*, Y^*)$, where the brand-holder's contract fee, every firm's branding, pricing, investment choices and every consumer's inspection and purchasing decisions are optimal, and consumers' beliefs are consistent. All proofs are included in the appendix.

1.4 Results

In this chapter, we first explore the social optimal allocation. We then establish common properties of the equilibria in this game. We set the case where inspection is impossible as our benchmark, and compare it to the case where consumer inspection is available with a cost. Finally, we show comparative statics and asymptotic properties of the equilibrium we focus on.

[4]Mathematically, an equilibrium in our game is equivalent to a stationary equilibrium in the repeated game where consumers can incur costs to read reviews generated in the previous period. In this dynamic setting, even though firms' incentives are not affect by the reviews, if discount rates are not too low, they still maintain the same quality for future profit concerns.

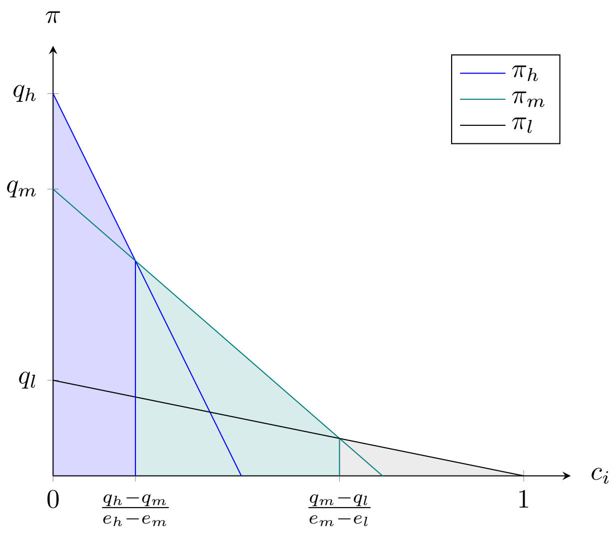

Figure 1.1: Social Optimum

1.4.1 Social Optimum

Before we begin with our analysis, it is helpful to explore the socially optimal case. In this case, prices and contract fee are irrelevant as they are transfers among players. If a benevolent social planner has the power to impose investment decisions, firms with lower investment costs should invest more in quality. If a firm's cost component c_i is $\frac{q_h-q_m}{e_h-e_m}$, whether it invest in high or medium quality yields the same level of welfare gain because $q_h - \frac{q_h-q_m}{e_h-e_m} \cdot e_h = q_m - \frac{q_h-q_m}{e_h-e_m} \cdot e_m$. Similarly, whether a firm with the cost component level $\frac{q_m-q_l}{e_m-e_l}$ invests either in medium or low quality does not affect the welfare level. Therefore, firms whose cost components c_i are less than $\frac{q_h-q_m}{e_h-e_m}$ should invest in high quality, those with c_i in between $\frac{q_h-q_m}{e_h-e_m}$ and $\frac{q_m-q_l}{e_m-e_l}$ should invest in medium quality, and those with c_i greater than $\frac{q_m-q_l}{e_m-e_l}$ should invest in low quality. Figure 1.1 summarizes the welfare and investment levels of each firm in the social optimal allocation. The sum of the shaded area is the welfare level in social optimum.

1.4.2 Preliminary Results

In this section, we establish properties of equilibria in this game. For simplicity, if a firm chooses to join the brand, we say it is in the brand market. Otherwise, it is in the stand-alone market.

Lemma 1.1. *In equilibrium, given any price p and branding decision b, consumers assign non-zero weight on the lowest possible quality in their beliefs, i.e. $\mu_m(p, 1) > 0$ and $\mu_l(p, 0) > 0$ for all p.*

Lemma 1.1 states that rational consumers should believe there can be low-quality goods in the stand-alone market, and medium-quality goods in the brand market. If a firm chooses $b_i = 1$, then the lowest possible investment level is medium. Otherwise, the lowest possible investment level is low. Since production costs increase with quality, it is straight-forward that if consumers assign 0 probability to the lowest quality, firms have an incentive to deviate to a lower investment level.

Lemma 1.2. *If consumers inspect with zero probability at stage 3, then any belief which induces such action must assign probability 1 to the lowest possible quality in equilibrium.*

Lemma 1.2 states that for any price the consumer observes, if she decides not to search, she believes that the quality is medium if she is in the brand market, and low if she is in the stand-alone market. If consumers do not inspect, the unique best response of firms is to invest as little as possible. Therefore, if consumers do not inspect in equilibrium, they must expect the lowest possible quality.

1.4.3 Benchmark Case: No Inspection

As benchmark, we discuss the case where inspection is not possible. Proposition 1.1 summarizes the unique PBE in this case.

Proposition 1.1. *If inspection is not possible, there exists a unique PBE where*

- *only firms with* $c_i \leq \hat{c} = \frac{1}{2}\frac{q_m - q_l}{e_m - e_l}$ *contract with the brand-holder,*
- *firms who contract with the brand-holder choose* e_m *and set prices at* q_m*,*
- *stand-alone firms choose* e_l *and set prices at* q_l*,*
- *the brand-holder sets fixed fee* $F = \frac{q_m - q_l}{2}$*.*

Proposition 1.1 follows immediately from Lemma 1.1 and 1.2. When inspection is impossible, Lemma 1.2 states that the only equilibrium is where all branded firms produce medium-quality goods and all stand-alone firms produce low-quality goods. Since there is no competition among firms, branded firms set prices at q_m and stand-alone firms q_l. Expecting these outcomes at later stages, the brand-holder's problem is equivalent to a cut-off problem. If a firm with cost-component c can profit more in the brand market, then all firms with $c_i < c$ should join the brand market as well. Therefore, the brand-holder's problem is

$$\begin{aligned} \max_{\{F\}} \quad & \hat{c}F \\ \text{s.t.} \quad & q_m - c_i e_m - F \geq q_l - c_i e_l \qquad \forall c_i \leq \hat{c} \\ & q_l - c_i e_l \geq q_m - c_i e_m - F \qquad \forall c_i > \hat{c} \end{aligned}$$

The incentive constraints pin down the contract fee F to $q_m - q_l - \hat{c}(e_m - e_l)$. Solving for the problem, we obtain $\hat{c} = \frac{1}{2}\frac{q_m - q_l}{e_m - e_l}$.

Our benchmark analysis shows that without inspection, no firm invests in high quality even though it can be welfare improving. Furthermore, the

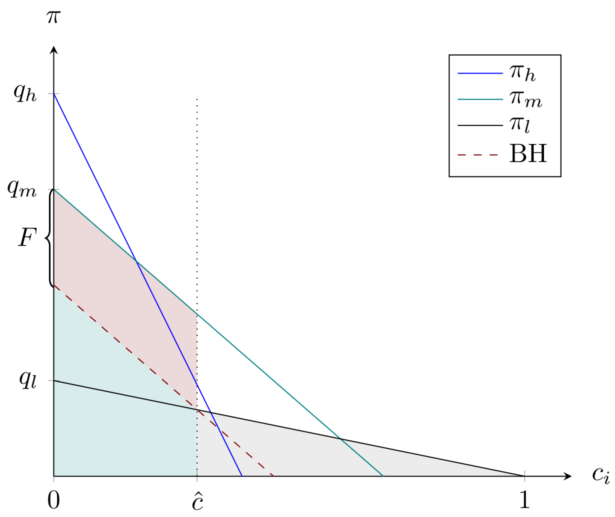

Figure 1.2: Benchmark Equilibrium

mass of firms who offer at least medium quality is lower than that in social optimum as well. Firms with c_i between $\hat{c} = \frac{1}{2}\frac{q_m - q_l}{e_m - e_l}$ and $\frac{q_m - q_l}{e_m - e_l}$ are left out in the stand-alone market. However, in the equilibrium price scheme, it is welfare improving for them to offer medium-quality goods. Welfare losses arise for two reasons. First, when inspection is not available, firms lose the incentive to invest in high quality. Second, since the brand-holder is the monopolist of the brand signal, optimal contract fee is high, which drives some firms away from the brand market.

1.4.4 Inspection

We now turn to the case where inspection is possible. The benchmark equilibrium in the previous section is still an equilibrium when inspection is possible. However, it is no longer unique. When inspection is possible and inspection costs are sufficiently low, there exists an equilibrium where high quality goods are produced in the stand-alone market. To solve for the game systematically, we first characterize a crucial property of equilibrium beliefs in Lemma 1.3:

Lemma 1.3. *In any equilibrium belief, if a certain quality is assigned non-*

zero probability in the brand market, it must be assigned zero probability in the stand-alone market, and vice versa. In other words, a certain quality level can only exist in at most one market.

Note that we do not allow for randomization in branding decisions. Hence, fixing the effort level, costs of a firm are identical in both markets. Fixing the investment level, when firms make branding decisions, they only take into account the revenue gain in the brand market and the associated contract fee. Therefore, if one firm finds it worthwhile to join the brand, all other firms producing on the same quality level should also join the brand.

Proposition 1.2 establishes the existence of another equilibrium, which is the main focus of this paper.[5]

Proposition 1.2. *When inspection cost is sufficiently low, there exists a PBE with cut-offs* $\underline{c}, \bar{c}$*, where*

- *a firm joins the brand if and only if its cost component is in the range* $[\underline{c}, \bar{c}]$*,*
- *if a firm joins the brand, it invests in medium quality,*
- *if a firm is in the stand-alone market, it invests in high quality if its cost component is less than* $\underline{c}$ *and low quality otherwise*
- *High-quality goods are of high price, medium-quality goods are of medium price, and the price of low-quality goods is randomized between the high price and a low price.*
- *Consumers only inspect upon high price in the stand-alone market.*

We now sketch the proof for Proposition 1.2. Denote the equilibrium high, medium and low prices p_h, p_m, and p_l, respectively. Denote σ the probability that a low-quality good is listed with a high price in equilibrium, then

[5] Other equilibria are discussed in the appendix.

consumers' beliefs upon observing price p and brand b is

$$\mu(p,b) = \begin{cases} (0,1,0) & \text{if} \quad (p,b) = (p_m, 1) \\ (\rho, 0, 1-\rho) & \text{if} \quad (p,b) = (p_h, 0) \\ (0,0,1) & \text{otherwise} \end{cases}$$

where $\rho = \frac{\underline{c}}{\underline{c}+\sigma(1-\bar{c})}$ is the probability that a high-price good is of high quality.

We show the existence of this equilibrium ("inspection equilibrium" hereafter) by backward induction. First, at the consumption stage, consumers purchase if the expected quality is greater than the listed price. On the equilibrium path, branded firms set prices at p_m, and stand-alone firms set prices at either p_h or p_l. If consumers observe p_l from an stand-alone firm, they believe the good is of low quality. If the price is p_m and the firm joins the brand, $\mu(p_m, 1) = (0, 1, 0)$. Therefore, consumers' purchasing decision can be simplified to purchase if $p_l \leq q_l$, $p_m \leq q_m$ and if $p_h \leq \rho q_h + (1-\rho)q_l$ without inspection, and $p_h \leq q_h$ after inspection and $q = q_h$.

In the inspection stage, given the beliefs stated above, consumers do not inspect if a stand-alone firm charges a low price or if a branded firm charges a medium price. In these cases, consumers assign probability 1 to a certain quality level, and therefore have no incentive to inspect. When a stand-alone firm charges a high price p_h, consumers inspect if

$$\rho(q_h - p_h) - k \geq \rho q_h + (1-\rho)q_l - p_h. \tag{1.2}$$

Denote $\gamma^*(p_h, 0)$ the optimal probability that consumers inspect in this case. When equation A.3 holds with equality (consumers' indifference condition for inspection hereafter), consumers are indifferent between inspection and no inspection. Therefore, any randomization between inspection and no inspection is a best response. Therefore, it is also optimal for consumers to inspect with probability $\gamma^*(p_h, 0) = \frac{p_h - p_l}{p_h}$ to support the equilibrium.

Expecting optimal inspection and purchase strategies, firms choose optimal prices and quality levels. Since there is no competition among firms, it is straight-forward that branded firms price at $p^*_m = q_m$ and choose e_m in this equilibrium. Consumers believe that a branded firm produces at medium quality with certainty, so branded firms can charge at consumers' willingness to pay for the medium-quality good. Moreover, because consumers do not inspect a branded firm, it is optimal for them to choose the lowest possible quality e_m. On the other hand, equilibrium strategies in the stand-alone market construct a semi-separating equilibrium. Regarding the low price, given that consumers purchase if $p_l \leq q_l$, the optimal low price $p^*_l = q_l$.

Hence, for any stand-alone firm i, the expected profit

$$E[\pi_i(p_i, e_i)] \begin{cases} p - c_i e_l & \text{if} \quad (p_i, e_i) = (q_l, e_l) \\ p - c_i e_h & \text{if} \quad (p_i, e_i) = (p_h, e_h) \\ (1-\gamma^*)p - c_i e_l & \text{if} \quad (p_i, e_i) = (p_h, e_l) \end{cases}$$

where p_h satisfies equation A.3.

Note that if γ^* is equal to $\frac{p_h - q_l}{p_h}$, then $(1-\gamma^*)p_h = q_l$, so a firm who chooses $e_i = e_l$ is indifferent between setting prices at p_h and q_l. In the optimization problem of the firms, consumers are indifferent between inspection, purchasing without inspection and leaving the market when the price is p_h in the stand-alone market. This is because firms who invest in high quality cannot differentiate themselves through price signalling. Since firms who invest in low quality will always have the incentive to mimic them, high quality products have to be set at a price which is equal to consumer's expected utility. Solving for this problem, we obtain the equilibrium high price $p^*_h = \frac{q_h + q_l + \sqrt{(q_h - q_l)(q_h - q_l - 4s)}}{2}$ and low-quality products are set at a igh price

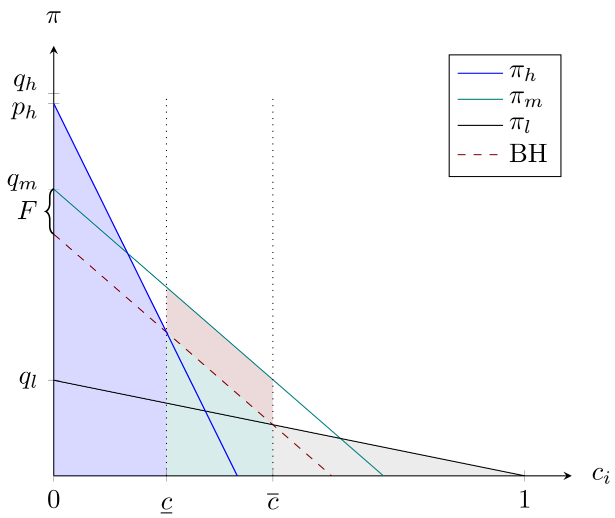

Figure 1.3: Inspection Equilibrium

with probability $\sigma^* = \frac{\underline{c}\left[q_h - q_l - 2s - \sqrt{(q_h-q_l)(q_h-q_l-4s)}\right]}{2s(1-\bar{c})}$. We defer the proof that these strategies are incentive compatible in the appendix.

Finally, in the contract stage, the brand-holder faces the following maximization problem:

$$\begin{aligned} &\max_{\{F\}} \quad (\bar{c} - \underline{c})F \\ &\text{s.t.} \quad q_m - c_i e_m - F \geq max.\{p_h - c_i e_h, q_l - c_i e_l\} \quad \forall \quad c_i \in [\underline{c}, \bar{c}], \end{aligned} \tag{1.3}$$

which result in $\underline{c} = \frac{1}{2}[\frac{p_h - q_m}{e_h - e_m} + \frac{p_h - q_l}{e_h - e_l}]$ and $\bar{c} = \frac{1}{2}[\frac{q_m - q_l}{e_m - e_l} + \frac{p_h - q_l}{e_h - e_l}]$.

Figure 1.3 illustrates the equilibrium strategies of firms and the brand-holder graphically. In the inspection equilibrium, firms with very low and high cost components ($c_i < \underline{c}$ or $c_i > \bar{c}$) are in the stand-alone market. Firms with low c_i invest in high quality and price at $p_h = \frac{q_h + q_l + \sqrt{(q_h-q_l)(q_h-q_l-4s)}}{2}$. Firms with high c_i invest in low quality and mimic the high-quality firms with probability $\sigma = \frac{\underline{c}[q_h - q_l - 2s - \sqrt{(q_h-q_l)(q_h-q_l-4s)}]}{2s(1-\bar{c})}$ in pricing. When consumers learn that they are in the stand-alone market, they inspect with probability

$\frac{p_h - q_l}{p_h}$ if the listed price is p_h. Consumers do not inspect in the stand-alone market if they face price q_l since it is obvious that the good is of low quality. If consumers do not inspect, they always purchase the good. If consumers inspect, then they purchase if the good is of high quality. In the brand market, on the other hand, the brand name already signals medium quality to consumers. Therefore, consumers do not inspect and always purchase the good because a branded firm always prices at q_m. Expecting the equilibrium strategies in the later stages, the brand-holder sets the contract fee F to miximize its profit.

Similar to the benchmark case, firms with high costs are in the stand-alone market because of the high contract fee. Unlike in the benchmark case, firms with low costs self-select into the stand-alone market because consumers can inspect on quality. Since consumers can verify product quality, these firms have incentives to produce high-quality goods in order to extract higher mark-ups. Firms who join the brand are those who have moderate firm-specific cost components c_i. They do not find it worthwhile to build their own reputation in the stand-alone market, but their costs are not that high that they are driven out of the brand market.

1.4.5 Comparison between Benchmark and Inspection Equilibrium

In this section, we compare the results in the benchmark equilibrium and the inspection equilibrium more in depth.

Proposition 1.3. *Welfare level in the inspection equilibrium is unambiguously higher than in the benchmark equilibrium.*

There are two reasons for the result in Proposition 1.3. First, firms which are more cost-efficient offer high-quality goods in the inspection equilibrium. Recall in the socially optimal case, the social planner will impose high-quality investment for firms whose cost components are below the threshold $\frac{q_h - q_m}{e_h - e_m}$. While these firms lack the incentives to invest in high quality in the benchmark equilibrium, the reduction of consumers' inspection costs in the inspection equilibrium enables them to profit through investments. Second, less firms which should invest in medium quality are left out in the stand-alone market and produce low-quality goods. In the benchmark case, the brand-holder contracts with firms who have lower costs, and leave out a sufficient mass of firms with its high contract fee. In the inspection equilibrium, the low cost firms will not join the brand. Therefore, the brand-holder has to adjust its contract fee to attract firms. This results in a reduction of contract fee, and more of the firms who invest in medium quality in social optimum joining the brand.

Proposition 1.4. *Compared to the benchmark case, the profit of every firm is weakly higher. Moreover, the difference decreases with the firm-specific cost component.*

Proposition 1.4 states that firms' profits are weakly higher in the inspection equilibrium. This result suggests that firms only profit from low inspection costs. The intuition is as follows. In the inspection equilibrium, the contract fee is lower. The marginal firm $\underline{c}$ which is indifferent between joining the brand and producing a high quality good in the stand-alone market is strictly better-off. If the marginal firm would join the brand in the benchmark equilibrium, the firm benefits from the lowered contract fee. If it does not join the brand in the benchmark case, its profit is equal to selling a low quality good at price q_l, which is strictly lower than producing a high-quality good in the inspection equilibrium. Therefore, all firms who have a lower cost component than $\underline{c}$ have a higher profit, and since their revenues are the same, a lower cost component implies a higher profit. Moreover, $\bar{c}$, the cost component of the marginal firm which is indifferent between joining the brand and producing a low-quality good in the stand-alone market, is strictly larger than $\hat{c}$, the marginal firm-specific cost component level in the benchmark case. Therefore, for firms with cost components greater than $\underline{c}$, they either profit from a lower contract fee if their cost components are lower than $\bar{c}$, or they stay indifferent in the stand-alone market.

1.4.6 Comparative Statics and Asymptotic Properties

In this section we discuss the comparative statics and asymptotic properties with respect to inspection cost k.

Proposition 1.5. *In the inspection equilibrium, the high price p_h in the stand-alone market increases as inspection costs decrease. As the inspection cost approaches 0, this price approaches q_h.*

Recall that $p_h = \frac{q_h+q_l+\sqrt{(q_h-q_l)(q_h-q_l-4k)}}{2}$. It is straight-forward to see that p_h decreases in k, and that as k approaches 0, p_h approaches q_h. It is intuitive that as the inspection cost decreases, firms who produce high-quality goods are more likely to build reputation in the stand-alone market. When inspection is almost free, these firms can almost charge at consumers' willingness to pay for the high-quality goods q_h.

Corollary 1.1. *As the inspection cost decreases, the fraction of branded firms, given by $\bar{c} - \underline{c}$, decreases. As the inspection cost approaches 0, this fraction approaches to $\frac{1}{2}[\frac{q_m-q_l}{e_m-e_l} - \frac{q_h-q_m}{e_h-e_m}]$, which is one half of the socially optimal level.*

Both $\underline{c}$ and $\bar{c}$ are increasing functions of p_h, which is a decreasing function of the inspection cost k. Therefore, $\underline{c}$ and $\bar{c}$ also decrease with k. This means that as the inspection cost decreases, there will be more low cost firms going in the stand-alone market, and more high cost firms accepting the contract. However, the marginal impact of inspection costs k on $\underline{c}$ is greater than that on $\bar{c}$. Therefore, as inspection costs k decreases, the fraction of firms joining the brand decreases.

Corollary 1.1 shows that for any arbitrarily small inspection cost k, an inefficiency arises and some firms do not choose their social optimal investment levels. There are more firms in the stand-alone market than is socially optimal. Also, there are more high-equality goods as well as low-quality goods.

Figure 1.3 illustrates the inefficiency graphically. We know by Lemma 1.3 that on equilibrium path, consumers believe that medium-quality goods only exist in the brand market. Therefore, the brand-holder acts as a monopolist of medium-quality signalling technology. The brand-holder faces a linear demand by firms who want to access such technology. Consequently, the brand-holder sets the contract fee F to extract monopolist profit. If there is competition among brand-holders, such dead-weight loss can be eliminated.

Corollary 1.2. *As the inspection cost decreases, the brand-holder's profit decreases.*

Both the fraction of branded firms and the contract fee decrease as the inspection cost decreases. Therefore, the brand-holder's profit unambiguously decreases.

1.5 Conclusion

We set up a game-theoretic model to analyze the impact of reviews systems on firms' branding choices and market outcomes. In our model, emergence of review systems is equivalent to a decrease in consumers' costs of quality inspection. We set our benchmark as the case where such costs are so high that inspection becomes impossible. When inspection is not available, firms have stronger incentives to join a brand in order to signal quality. A brand-holder, who has established its reputation in the market, is able to extract a significant share of profits from firms who use its trademark. Inefficiencies arise for two reasons in this case: First, firms who join the brand lack incentives to provide products with higher quality than demanded by the brand-holder. Second, due to the high contract fees, there are too few branded firms.

Our first finding is that review systems mitigates the aforementioned efficiency problems. We show that for sufficiently low inspection costs, there exists an equilibrium where high-quality goods are produced by stand-alone firms who do not join the brand, and that more firms are included in the brand. High quality goods are pooled in the stand-alone market with low quality goods. However, low inspection costs enable consumers to search and tell these products apart. Moreover, firms profit from review systems, and the increase in profits increases with their cost efficiency. The value of a brand name is lower, for the brand-holder's profit decreases.

We conduct comparative statics analysis and find that as the inspect cost reduces, the number of firms joining the brand decreases. Moreover, the price of high-quality goods increases, and the brand-holder's revenue decreases. Interestingly, we find that inefficiency arises from over-investment in quality. When the brand-holder has market power, a fraction of cost-efficient firms over-invest in quality because profits are higher to stand alone than to join the brand.

Our results complement existing empirical findings of the decreased values of brand names. We emphasize that in addition to information disclosure, review systems also affect firms' branding choices and their incentives to invest. However, future research is required to verify other mechanisms of review systems as well as measuring the magnitude of effects we document.

Chapter 2

Front-of-Package Labelling: Adoptions and Regulations

2.1 Introduction

Front-of-Package (FoP) labels are designed to help consumers understand nutrition contents better and faster. According to WHO, 17 million people die from cardiovascular diseases each year. Many of those premature deaths could have been circumvented by improving diets (Anand et al. [2015]). Diet improvement requires better nutrition choices. Yet, a survey conducted by Christoph et al. [2018] shows that only one third of people read nutrition contents when buying groceries.

The health benefits of FoP labels can be considerable. A field experiment study suggests that the mere effect of offering FoP labels results in reduction of calories intake at lunch by 10%. Similar supporting evidence is also found in Ducrot et al. [2016], Khandpur et al. [2018], and Talati et al. [2019]. However, take-up rates of FoP labels are low among Eurpean countries.[1]

[1] In 2019, the take-up rate of Nutri-score in France is only 25%, and in countries like Germany and Italy, strong resistance against Nutri-score persists. For details, see the report "Food Labelling System, Nutri-Score, Gains Momentum in Europe" on Quality

In this paper, we provide an explanation for the low take-up rate of FoP labels, which contradicts the classic unravelling result in the information disclosure literature (Grossman [1981], Milgrom and Roberts [1986]). We show that low take-up rates can be a result of the interplay between FoP labels and existing nutrition facts labels.

FoP labels face a trade-off between informativeness and interpretability. In order to provide consumers nutrient information at a glance, the labels often do not contain exact nutrition amounts but rather letter grades, warning signs or traffic lights. While nutrition facts tables on the back of packages seem to complement FoP labels well, we show that they can in fact result in low FoP adoption rates.

We construct a game-theoretic model and establish that (i) heterogeneous consumer preferences and (ii) potential other sources of information result in low FoP adoption rates. These assumptions are satisfied in the grocery industry as well as many general settings.

In our model, a product has two attributes. Consumers assign different weightings among attributes, but within each attribute, the product is vertically differentiated. Specifically, every attribute can be of high or low quality. One attribute is valued by all types of consumers while the other attribute is only valued by a certain type of consumer. For example, a body-builder may value protein content more than an average consumer, and a person on kidney dialysis may be more cautious on sodium content.

If firms can only signal through the FoP label, all firms adopt the label voluntarily. This result holds even though the label cannot fully disclose the quality of each attribute in our model. Likewise, it continues to hold if consumers can search for nutrition information at a low cost, FoP label adoption is costless and consumers' preferences are homogeneous. However, when both

Assurance and Food Safety magazine and reports on criticism in Germany and Italy on https://www.foodnavigator.com/.

of our key assumptions hold, only the products with a high quality attribute that is valued by all consumers value will adopt FoP labels. Producers with products which have high quality only on the other attribute find it more profitable to avoid the FoP label. Instead of showing the FoP label, these producers induce consumers who value their quality to search. They charge a higher price for their targeted consumers and earn a higher profit.

We continue with the welfare analysis of mandatory FoP labels. While madatory FoP labels are in demand in Europe[2], we show that mandating FoP labels may lead to a decrease in social welfare. When FoP labels are mandatory, price signalling yields more diluted information. Products grouped together by FoP labels will have similar prices, and consumers who value a certain attribute has to search more frequently.

The contributions of our paper are two-fold. First, we offer a theoretical explanation of the observed low adoption rates of FoP labels. Second, we compare two different regulation schemes: mandatory and voluntary FoP label adoption. Our results not only establish a fundamental cause of low adoption rates but also highlights firms' adaptation to regulations.

Despite its focus on the application of Front-of-Package labels, the insight of our paper can be extended to a broader setting. In fact, it can be applied to any market with an aggregated signal of multiple attribute. E.g., on an online platform with a 5-star review system, those reviews serve as a simplified signal similar to FoP labels. Another example are GPAs that serve as a vague signal of the students' personal attributes. Moreover, for products whose attribute qualities cannot be easily transmitted to consumers, such as laptops and smartphones, a potential demand for such simplified signals can be expected.

The remainder of this paper is structured as follows. Section 2.2 discusses

[2]In April 2020, a joint letter requesting mandatory Nutri-score was sent to the European commission by European parliament members, consumer organizations, public health scholars as well as a list of retailers. For the letter content, see the European Consumer website for reference number: Ref.: BEUC-X-2020-029/MGO/cm.

related literature. Section 2.3 introduces the set-up of our model. Section 2.4 presents our results. Section 2.5 concludes.

2.2 Related Literature

This paper extends the vast literature on information disclosure. Under reasonable assumptions, Grossman and Hart [1980], Grossman [1981] and Milgrom and Roberts [1986] establish that a firm will voluntarily disclose its private information to consumers. In a vertically differentiated market, a firm with high quality products can charge a higher mark-up through quality disclosure. If consumers are rational, other firms will also disclose information in order to separate their products from those with lower qualities. This result is known as the "unravelling" result in information disclosure literature.

While existing literature focuses on how certain market structures or irrationality of consumers may jeopardize disclosure[3], our paper suggests that an incomplete information structure also undermines disclosure incentives. The most related work to our paper is Ghosh and Galbreth [2013]. Similar to our setting, in Ghosh and Galbreth [2013], communication is two-sided. Firms can disclose product quality and consumers can conduct search as well. While Ghosh and Galbreth [2013] find that firms disclose less information when search costs are high, we document how low search costs lead firms to counter-signal by avoiding disclosure.

Our paper also adds to the vast discussion of mandatory disclosure. Even though mandatory disclosure ensures more information transmission, the lit-

[3]For example, Jovanovic [1982] shows that if disclosure is not costless, the classic unravelling result is not obtained. In Fishman and Hagerty [2003], consumers may be uninformed even after firms disclose. Board [2009] shows competition among firms may hinder disclosure. Levin et al. [2009] discuss disclosure incentives in a horizontally differentiated market. Guo and Zhao [2009] compare cases where firms disclose sequentially or simultaneously.

erature generally agrees that whether mandating disclosure increases welfare heavily depends on the context.[4]

In terms of modelling choices, our model is closest to Mayzlin and Shin [2011], who investigate why firms engage in uninformative advertisement. As an alternative to the "money burning" story in Nelson [1974], Mayzlin and Shin [2011] show that when both the firm and the consumer can invest in communication, a producer of high-quality goods has an incentive to engage in uninformative advertisement, inducing consumers to conduct search. While the research questions are different, the intuition behind our main results is similar to that in Mayzlin and Shin [2011]. If a firm's product is highly valued by consumers, the firm may have the incentive to avoid disclosure, inviting its targeted consumers to study the product.

2.3 Set-up

There is one firm and one consumer. The firm produces a product with uncertain quality which is private to the firm. The firm can voluntarily opt for a label that indicates, but does not fully disclose product quality. The consumer learns about the price of the product and, if available, the information about product quality through the label. Before making purchasing decisions, the consumer considers whether or not to engage in a costly search after which product quality is fully disclosed.

[4]Positive effects of mandatory disclosure are documented theoretically in Fishman and Hagerty [2003], and empirically in Mathios [2000] in the salad dressing industry and Frondel et al. [2020] in the German housing market. Negative effects are documented theoretically in Jovanovic [1982]. Ispano and Schwardmann [2018] also discovers that if consumers are cursed (Eyster and Rabin [2005]), partial mandatory disclosure decreases welfare. In a financial setting, Jayaraman and Wu [2019] document lower efficiency resulting from mandatory disclosure empirically.

2.3.1 Firm

The firm produces one unit of a product with two attributes, denoted by α and β. Both attributes can be of either high or low quality with equal probabilities. Assume that there is no correlation between quality levels of the two attribute. We write $\alpha \in \{A, a\}$ and $\beta \in \{B, b\}$, where capital letters represent high qualities. Therefore, the firm's type θ, characterized by its product, can be one of the following: $(A, B), (A, b), (a, B)$, or (a, b), and the *ex ante* probabilities of each type is $\frac{1}{4}$. For simplicity, we say the firm is of high type ($\theta = h$) if both attributes are of high quality, and low type ($\theta = l$) if neither is of high quality. If the product is of high quality only on the α (β) attribute, we say the firm is of m_α (m_β) type ($\theta = m_\alpha$ (m_β)).

After the firm learns about its type, it can choose whether or not to adopt a label. The information structure of the label will be discussed in detail. The adoption of the label is costless. Denote s_θ the label adoption strategy of the firm with type θ with $s_\theta \in \{0, 1\}$ for $\theta \in \{h, m_\alpha, m_\beta, l\}$ where 1 denotes adoption and 0 no adoption. After the firm decides on the labelling strategy, it sets a price for its product. For simplicity, there are no production costs. Thus, the firm is a profit maximizer and its profit depends solely on the price the consumer pays if transaction take place.

2.3.2 Consumer

The consumer can be of two types: the niche type (N) or the mass type (M). Denote ω the probability that the consumer is the niche type. The type-dependent utility functions of the consumer are as follows:

$$\begin{aligned} U_N &= \lambda \cdot \mathbb{1}(\alpha = A) && + \mathbb{1}(\beta = B) + u_0 \\ U_M &= && \mathbb{1}(\beta = B) + u_0 \end{aligned}$$

First, every product generates base utility $u_0 > 0$ for the consumer.[5] Second, regardless of the consumer's type, a product with high quality on the β attribute generates unit utility. Finally, the niche type derives utility from a product with high quality on the α dimension, but the mass type does not. $\lambda > 0$ represents the weighting of the α attribute relative to the β attribute for the niche type consumer.[6]

For simplicity, we call the α attribute the "niche attribute", and the β attribute the "mass attribute" hereafter. Also, we call the consumer of type N the "niche consumer", and type M the "mass consumer".

Upon arrival, the consumer observes the price of the product and the label if available. Regardless of the label, the consumer can learn the quality of attributes at a search cost $c > 0$. The consumer then decides on whether to purchase the product or not. If transaction does not take place, the consumer is left with reservation utility normalized to 0. We assume that the consumer buys if she is indifferent between buying and leaving the market.

2.3.3 Information Structure

The mapping from product attributes to label is summarized in Table 2.1:

Product Quality	Label($\mathcal{L}$)
(A, B)	2
(A, b)	1
(a, B)	1
(a, b)	0

Table 2.1: Information Structure

We model a simple information structure for product quality. Essentially,

[5]In general, u_0 can be normalized to zero. This assumption simplifies the discussion of the equilibrium we focus on. See appendix for details.

[6]Our results are still valid if the α attribute also generates utility for the mass type consumer. The key assumption for our results is the heterogeneity of relative attribute weightings among consumers.

the label shows the number of high-quality attributes of the product. If both α and β attributes are of high quality, the label is 2 ($\mathcal{L} = 2$). If only one attribute is of high quality, regardless of which attribute it is, the label marks 1 ($\mathcal{L} = 1$). Finally, if neither attribute is of high quality, the label shows 0 ($\mathcal{L} = 0$).

Our choice of modelling not only simplifies the analysis but also preserves the essence of Front-of-Package labels as well as other uni-dimensional signals such as reviews. A product with low fat and high sugar content many be grouped together with another product with high fat and low sugar content. A restaurant which has a mediocre star rating may have not-so-outstanding food or service. Thus, the signal structure is a simple representation which preserves the trade-off between informativeness and interpretability.

2.3.4 Timeline

The timeline of the game is summarized in Table 2.2:

Period 0	Nature draws the type of the firm and the consumer.
Period 1	The firm learns about its type θ and chooses labelling strategy s_θ.
Period 2	The firm sets prices p_θ.
Period 3	The consumer observes product price and label if available.
Period 4	The consumer decides whether to investigate (at cost c) quality or not.
Period 5	The consumer makes purchase decisions. Payoffs realize accordingly.

Table 2.2: Timeline

2.3.5 Equilibrium Concept

Throughout this paper, we focus on Perfect Bayesian Nash Equilibria (PBE) which survive D1-criterion refinement (Banks and Sobel [1987]). We allow the firm to set mixed pricing strategies among finite prices and the consumer to play mixed investigation strategies.

2.4 Results

In this chapter, we present our findings. We first show that unravelling is an equilibrium if consumers have homogeneous preferences, or if they cannot conduct search. We then characterize the equilibrium we focus on, which only exists if label adoption is voluntary. In this equilibrium, the firm with type m_α does not adopt the label in order to induce the niche consumer to conduct costly search. The type m_α firm can therefore target the niche consumer and earn a higher profit. We then discuss the case where label adoption is mandatory. We show that making the adoption mandatory may decrease welfare. Proofs to all results in this paper can be found in the Appendix.

2.4.1 Preliminary Results

Before presenting our main results, we show that either of our two main assumptions, heterogeneous preferences or alternative source of information, does not undermine disclosure. Rather, it is the interplay of the two that gives rise to low adoption rates.

Proposition 2.1. *If the consumer cannot investigate product quality, then there exists an equilibrium where all types except the l type firm adopt the label.*

Proposition 2.1 shows that the classic unravelling result can be obtained if the consumer cannot search for quality. In this case, the label is the only credible source to signal quality. Even though the label does not offer full revelation, firms of high and medium quality products still have the incentive to separate themselves from the worst.

Proposition 2.2. *If the consumer is for sure a niche type, namely, if ω is equal to 1, then there exists an equilibrium where all types except the l type firm adopt the label.*

If consumers do not value product attributes differently, the unravelling results are still obtained. The intuition is similar to that in Proposition 2.1. In this case, products are differentiated vertically. Since it is a dominant strategy for type h firm to disclose, type m_α's best response is also to disclose, which in turn induces m_β to disclose as well.

2.4.2 Voluntary Label Adoption

We now establish the equilibrium which explains low label adoption rates. As this equilibrium only exists when label adoption is voluntary, we call this equilibrium the "voluntary equilibrium" hereafter. Proposition 2.3 summarizes the voluntary equilibrium:

Proposition 2.3. *If $\lambda \geq \frac{1}{1-c}$, $c < \frac{\lambda}{4}$ and $\omega \geq \bar{\omega} = \frac{2+2u_0}{\lambda+2u_0+\sqrt{\lambda(\lambda-4c)}}$, then there exists a PBE where*

- *the firm adopts the label if and only if it is of type h or m_β,*
- *the mass consumer never searches,*
- *the niche consumer searches only when the product does not have a label and the price is $p_\alpha = u_0 + \frac{\lambda+\sqrt{\lambda(\lambda-4c)}}{2}$,*
- *the type-dependent prices p_θ are deterministic for $\theta \in \{h, \alpha, \beta\}$. p_l is randomized between p_α and u_0 with probability $\sigma = \frac{\lambda-2c-\sqrt{\lambda(\lambda-4c)}}{2c}$.*

We now illustrate the intuition of Proposition 2.3 and describe all equilibrium strategies in detail. First, we solve for the equilibrium strategies by backward induction. At the consumption stage, the consumer's strategy is simple. Since the consumer is rational, she purchases if the expected utility she can derive from this product is at least as high as the product price. Next, in the search stage, since the consumer's utility gain is type-dependent, benefits from search differs for the two consumer types. For both types of consumer, if the label is present, beliefs about product attributes are deterministic in this equilibrium. If the signal is 2, then the product is (A, B). If the signal is 1, the product must be (a, B). Moreover, since only the l type firm sets price at u_0 with positive probability, if the price is equal to u_0, then the consumer knows that the product must be (a, b). Therefore, in all of the above cases, regardless of her type, the consumer does not have an incentive to search.

When the price is p_α and there is no label, the consumer forms the belief that the product is (A, b) with probability $\frac{1}{1+\sigma}$, and (a, b) with probability $\frac{\sigma}{1+\sigma}$, where σ is the probability that the l type firm prices at p_α. For the mass consumer, both products yield the utility level of u_0, which leaves her no incentive to search. For the niche consumer, (A,b) is worth λ more than (a,b), which only provides the base utility u_0. Hence, the niche consumer searches if

$$\frac{1}{1+\sigma}(\lambda + u_0 - p_\alpha) - c \geq \frac{1}{1+\sigma}(\lambda + u_0 + \sigma u_0) - p_\alpha. \tag{2.1}$$

As we will see in the following analysis, in equilibrium, the niche consumer is indifferent between search and no search. It is therefore a best response if the niche consumer searches with probability $\gamma = \frac{\omega p_\alpha - u_0}{\omega p_\alpha}$. To summarize the search strategies of the consumer, the mass consumer never searches, and the niche consumer only searches with positive probability if the price is p_α and no FoP label is available.

Next, we solve for the optimal pricing and signalling strategies of the firm. In period 2, the optimal pricing strategy in this equilibrium is as follows: First, if the firm is a high type, its attribute qualities are fully communicated through the label. In the range of parameter specified, the high type firm will set the price equal to the willingness to pay of the niche consumer. This is due to the fact that the niche attribute generates sufficiently high utility, or equivalently, the probability that the consumer is of niche type is sufficiently high. Therefore, the high type firm sets the price at $p_h = \lambda + 1 + u_0$. The pricing strategy for the m_β type firm is straight-forward. In this equilibrium, m_β also fully reveals its type. Since the consumer values its product at $1 + u_0$ regardless of her type, m_β type firm sets the price at $p_\beta = 1 + u_0$.

As for the pricing strategies of types who do not adopt the label, the m_α type firm is pooled with the l type firm. It is a best response for the l type to randomize between u_0 and p_α, the price that the m_α type firm sets. In this equilibrium, both pricesu_0 and p_α yield the same expected profit for the l type firm because the niche consumer randomizes her search strategy.

The profit maximizing pricing strategy for the m_α type firm can be obtained by solving the following equations:

$$\frac{1}{1+\sigma}(\lambda + u_0 - p_\alpha) - c = \frac{1}{1+\sigma}(\lambda + u_0 + \sigma u_0) - p_\alpha = 0$$

Since the m_α type firm cannot prevent the l type firm from mimicking, it has to set a price which makes the niche consumer indifferent between buying without search and searching, while leaving the niche consumer as well-off as exiting the market. Solving for the equations, we have

$$\begin{aligned} \sigma &= \frac{\lambda - 2c - \sqrt{\lambda(\lambda - 4c)}}{2c} \\ p_\alpha &= u_0 + \frac{\lambda + \sqrt{\lambda(\lambda - 4c)}}{2} \end{aligned}$$

We provide only the intuition of incentive compatibility here and defer the full examination to the appendix. An off-equilibrium belief which survives D1-criterion is where the consumer assigns probability 1 to any off-equilibrium signal bundle of price and institutional signal. With such off-equilibrium beliefs, it is straightforward to see that the h and l types do not deviate. Both type h and l earn the monopolist profit for the value of their products.

Type l cannot mimic any type other than m_α in equilibrium, for all other types have verifiable signals through the label.

Considering the incentives of the m_α type firm, note that in the range of parameters we specify in Proposition 2.3, the probability that the consumer is of niche type is sufficiently high, or equivalently, the niche attribute generates sufficiently high utility. Therefore, firm m_α can set a sufficiently high price such that expected profit ωp_α, the probability of the consumer being niche type times the price it sets, is higher than the profit of firm m_β. In such a case, firm m_α has no incentive to deviate and adopt the label.

The m_β type firm has no incentive to deviate and mimic the m_α type firm either. Note that in equilibrium, the consumer randomizes search with a probability such that the l type firm is indifferent between setting price at u_0 and p_α. Therefore, if firm m_β does not adopt the label and sets the price at p_α, the expected profit will be exactly u_0, since the niche consumer will not purchase (a, B) at price p_α. Clearly, this profit is lower than the m_β type firm's equilibrium profit level $1 + u_0$.

In the voluntary equilibrium, the classic unravelling result is not obtained. The key to this result is that the label is not the only source of information. If the consumer is willing to incur the search cost, she can learn more about the attribute qualities. In such a case, if consumers have heterogeneous preferences, firms have an incentive to target the consumers who value their products more. One way to achieve targeting is to avoid the label and induce the targeted consumers to conduct search so that they find these products. By doing so, these firms earn a higher profit than they would if they adopt the label, but the consumer has to bear the search cost.

2.4.3 Mandatory Label Adoption

One intuitive solution for the low label adoption rate is to make the label mandatory. In this section, we first characterize the unique equilibrium in the same parameter range as in the voluntary equilibrium. We then show in the next section that if the label adoption policy is the only thing that changes and consumers still have other information sources, then making the label mandatory may decrease welfare. Proposition 2.4 presents the only equilibrium ("mandatory equilibrium" hereafter) which survives D1-criterion refinement when label adoption is mandatory in the parameter range of the voluntary equilibrium.

Proposition 2.4. *Assume that label adoption is mandatory. If* $\lambda \geq \bar{\lambda}^M = \frac{u_0^2+6u_0-4c+8}{4-4c+2u_0}$, $c < \frac{\lambda-1}{4}$ *and* $\omega \geq \bar{\omega}^M = \frac{1+u_0}{1+u_0+\frac{\lambda-1+\sqrt{(\lambda-1)(\lambda-1-4c)}}{2}}$, *there exists a PBE where*

- *the type-dependent prices* p_θ *are deterministic for* $\theta \in \{h, m_\alpha, l\}$,
- *the* m_β *type firm randomizes the price between* p_α^M, *the price that type* m_α *sets, and* $1 + u_0$, *with probability* $\sigma^M = \frac{\lambda-1-2c-\sqrt{(\lambda-1)(\lambda-1-4c)}}{2c}$,
- *the mass consumer never searches*
- *the niche consumer searches only when the price is* $p_\alpha^M = 1 + u_0 + \frac{\lambda-1+\sqrt{(\lambda-1)(\lambda-1-4c)}}{2c}$.

We show in the appendix that an equilibrium where the m_α and m_β type firms pool on pricing strategy does not survive D1-criterion refinement. Moreover, for the parameter range where the voluntary equilibrium exists, only a semi-separating equilibrium characterized in Proposition 2.4 exists.

We now discuss the incentives of the firm. When label adoption is mandatory, the high and low types are fully revealed to the consumer. The pricing strategies for these two types are therefore straight-forward. As in the voluntary equilibrium, the price that the high type firm sets p_h^M is equal to the niche consumer's willingness to pay, $\lambda + 1 + u_0$. The low type firm sets the price p_l^M at the base utility u_0.

When the probability that the consumer is a niche type ω is sufficiently high, or equivalently, when λ is sufficiently high, the m_β type has an incentive to mimic the m_α type. The construction of the mandatory equilibrium is similar to that of the voluntary equilibrium. It is incentive compatible for the m_α type firm to set a price where the niche consumer is indifferent between purchasing without search, searching and leaving the market. The m_β type firm randomizes between this price p_α^M and $1 + u_0$, the willingness to pay of the consumer for its product.

In terms of the consumer's optimal strategies, it is easy to see that the consumer does not search if the label signal s is 2 or 0. In these cases, information on product attributes is fully conveyed. On the other hand, uncertainty may still persist if the label signal s yields 1. If the listed price is $1 + u_0$, both the niche and mass type consumers purchase without search, since the product attributes are clearly (a, B). The optimal search behaviors of the niche and mass consumers differ, however, when the label signal $s = 1$ and the price is p_α^M. Since p_α^M is greater than $1+u_0$, the mass consumer leaves the market. The niche consumer, as described in the firms' optimal pricing strategy above, is indifferent between searching and buying without search. Therefore, it is a best response for the niche consumer to randomize search at a rate such that the m_β type firm earns equal profits between prices p_α^M and $1 + u_0$.

2.4.4 Mandatory vs. Voluntary Adoption

In this section, we compare the mandatory and voluntary equilibria. For the sake of clarity, all strategies played in the mandatory equilibrium are denoted with an upper script M. For example, the price that the m_α firm sets in the mandatory equilibrium is denoted as $p_{m_\alpha}^M$.

Lemma 2.1. *When label adoption is mandatory, the m_α type firm is mim-*

icked more frequently. In other words, $\sigma_\alpha^M > \sigma_\alpha$.
Proof: see appendix.

Corollary 2.1. *The* m_α *type firm sets a lower price in the mandatory equilibrium than in the voluntary equilibrium. In other words,* $p_\alpha^M < p_\alpha$.
Proof: see appendix.

From Lemma 2.1 and Lemma 2.1, we can derive Proposition 2.5.

Proposition 2.5. *Given signals that indicate the firm may be of* m_α *type, the consumer searches less in the mandatory equilibrium. Namely,* $\gamma^M < \gamma$. *However, the ex ante probability of consumer search is higher in the mandatory equilibrium.*
Proof: see appendix.

Proposition 2.5 establishes our first inefficiency result in the mandatory equilibrium. It may be intuitive to postulate that consumers incur less search costs when label adoption is mandatory. However, in the case where consumers have heterogeneous preferences and the simplified label cannot fully separate all products, enforcing mandatory label adoption results in consumers searching more frequently. The drive of this result, as we show in Lemma 2.1, is that prices provide a weaker signal. In our model, when label adoption is mandatory, the m_β type firm mimicks the m_α type firm more frequently than the l type firm does so in the voluntary equilibrium. Indeed, given that the label already signals one high quality attribute, the consumer does not have to conduct search as frequently. However, since label adoption is mandatory, the niche consumer encounters the signals which induces search more frequently. In reality, this maps to the case where many products have a mediocre label and similar prices. Consumers who value certain qualities will spend a higher cost searching.

Proposition 2.6. *Welfare is lower in the mandatory equilibrium than in the voluntary equilibrium.*
Proof: see appendix.

Proposition 2.6 presents our main result in this section. As we show in the voluntary equilibrium, it is true that if consumers have heterogeneous preferences and other sources of information, label adoption may not be appealing to some firms. Such reaction sabotages the well-meaning introduction of the label. If very few firms adopt the label, then consumers still don't have additional information about most products. However, forcing firms to adopt the label does not necessarily improve efficiency. As we show in the appendix, when label adoption becomes mandatory, the probability of trade taking place decreases. Note that the m_α firm sets a lower price to compensate for the niche consumer's search cost. Therefore, in equilibrium, the lowered price p_α^M does not increase efficiency. Since the m_β type firm mimicks the m_α type firm at a higher frequency, the likelihood that the niche consumer purchases after search is lower. In other words, the niche consumer is more likely to leave the market after search in the mandatory equilibrium than in the voluntary equilibrium.

2.5 Conclusion

In marketplaces today, consumers often face the problem of having too much information to process. It is therefore important to present information in a comprehensive way. Front-of-Package labels is one example to help consumers better understand nutrition contents of grocery products. In contrast to the nutrition facts table on the back of the package, front-of-package labels are shown at the front so that consumers can easily see them and make comparisons between products. Moreover, for readability, Front-of-Package labels are often a simple sign or a one-dimensional traffic light. A trade-off between informativeness and interpretability emerges. The labels are immediately readable, but some information is inevitably lost in the process of simplification.

In this paper, we construct a signalling game to discuss simplified signals such as Front-of-Package labels. First, we provide an explanation to an empirical puzzle: why label adoption rates are low. In our setting, all players in the game are rational. We show that low adoption rates result from consumers' heterogeneous preferences and existing information sources available to consumers. If consumers have homogeneous preferences, or if the label is the only information available to consumers, firms will have an incentive to adopt the label. However, since these conditions often do not hold, low adoption rates of FoP labels may occur. Some firms may avoid labels in order to induce their targeted consumers to search for them.

Second, we discuss adoption regulations for Front-of-Package labels. We show that making adoption mandatory may be welfare-decreasing. Mandatory adoption not only changes the label of the products, but also their prices. Specifically, mandatory adoption results in firms setting a similar price to other products in their tier. Consumers who value certain product attributes will have to search among a potentially larger pool. If search is costly for the consumer, which is the case if Front-of-Package labels can be helpful, then making it mandatory for firms to adopt these labels may in fact create more search cost on the consumer side.

Our paper is the first attempt to our knowledge offering a theoretical explanation for why label adoption rates are low, and how mandating the label may not be the solution, but much more can be explored in future research. There are many more potential contributing factors to low adoption rate. Consumers may exhibit behavioral biases which allow firms to profit without adopting the label. Moreover, in our model, a firm produces only one product. In reality, a firm may product multiple products with different nutrition contents, which can change their incentives to unravel. Empirical research can help to identify the main causes of low FoP adoption rates. Finally, when it comes to Front-of-Package labels, although effects of nutrition improvement of products are often expected, few consider other equilibrium effects such as pricing, signalling and label adoption. Future research applying knowledge in industrial organization, mechanism design and behavioral economics is in need to bridge these gaps.

Chapter 3

Using Machine Learning to understand Bargaining Experiments

[1]

3.1 Introduction

This chapter is about a general class of bargaining games in which there is private information about the amount that is being bargained over (often called the "pie size"). This class is most common in everyday bargaining. It is also interesting in both *theory* and *practice*.

Theory is interesting because when there is private information and people are self-interested, theories based on individual rationality typically predict an inevitable loss of efficiency. That is, even when a bargain is mutually beneficial for both sides, they will not always come to agreement.

Private information bargaining is interesting in *practice* because, while

[1] This chapter is based on the joint work with Colin Camerer, Po-Hsuan Lin, Gideon Nave, Alec Smith and Joseph Tao-Yi Wang.

inefficiencies are predicted by theory, it is also known that if there are observable statistical proxies for the hidden private information, then sets of rules (mechanisms) which use this information can improve efficiency [McAfee and Reny, 1992, Crémer and McLean, 1985, 1988]. Therefore, it is possible that methods for measuring private information can improve efficiency, even when bargainers voluntarily participate in systems using those measures.

There is a long history of using highly controlled laboratory experiments to study bargaining. A brief description of this history helps explain why we are enthusiastic about modern applications of machine learning.

3.1.1 A Brief History of Bargaining Experiments

The experimental literature on bargaining is vast, so we focus only on those studies closely related to ours.[2] Before theoretical breakthroughs in understanding structured bargaining, most experiments used unstructured communication. The main focus of interest was process-free solution concepts such as the Nash bargaining solution [Nash Jr, 1950], and important extensions to those concepts [e.g. Kalai et al., 1975].

We will refer to the amount of surplus available to share as the "pie". Many bargains [Nydegger and Owen, 1974, Roth and Malouf, 1979] led to an equal split of the pie. Roth suggested that "bargainers sought to identify initial bargaining positions that had some special reasons for being credible... that served as *focal points* that then influenced the subsequent conduct of negotiation" [Roth, 1985]. Under informational asymmetries, disagreements may arise due to coordination difficulties. Several papers by Roth and colleagues then explored what happens when players bargain over points which have different financial value to players [Roth and Malouf, 1979, Roth et al., 1981, Roth and Murnighan, 1982, Roth, 1985]. In theory, there should be no

[2]For longer reviews, see Kennan and Wilson [1993], Ausubel et al. [2002], Thompson et al. [2010]

disagreements in these games but a modest percentage of trials (10-20%) did result in disagreement, which seems to involve differences about which "focal points" are acceptable.[3] Roth et al. [1988] also drew attention to the fact that the large majority of agreements are made just before a (known) deadline, an observation called the "deadline effect".

Two pioneering papers, Ståhl [1972] and Rubinstein [1982] showed how noncooperative game theory might be used to improve the apparent precision of bargaining theories. Since then, almost all experimental studies have tested what happens in highly structured settings using variants of those early game structures; for a review see Ausubel et al. [2002]. In these theories and experiments, "structure" means that the rules of how bargaining proceeds are clearly specified in the theory; put differently, bargainers have no freedom to time their offers or use natural language. The rules typically define when bargaining must be completed (either a deadline or an infinite horizon), who can offer or counteroffer and at what time, when offers are accepted, whether communication is allowed (and in what form), and so on.

Theoretical predictions of outcomes and payoffs depend sensitively on these structural features.[4] That structural-sensitivity proved to be enticing, because it created a cornucopia of interesting experiments testing whether bargaining was sensitive to structured features as theory predicted. This led to a burst of progress in experimental literature testing these theories.[5]

Many other experiments have observed what happens in semi-structured bargaining in which there is *two-sided* private information [Valley et al., 2002].

---

[3]See Schelling [1960], Roth [1985], Kristensen and Gärling [1997], Janssen [2001], Binmore and Samuelson [2006], Janssen [2006], Bardsley et al. [2009], Isoni et al. [2014, 2013], Hargreaves Heap et al. [2014].

[4]See Cramton [1984], Chaussees [1985], Rubinstein [1985], Grossman and Perry [1986], Gul and Sonnenschein [1988], Ausubel and Deneckere [1993].

[5]See Ochs and Roth [1989], Camerer et al. [1993], Mitzkewitz and Nagel [1993], Güth et al. [1996], Kagel et al. [1996], Güth and Van Damme [1998], Rapoport et al. [1998], Kagel and Wolfe [2001], Srivastava [2001], Croson et al. [2003], Johnson et al. [2002], Kriss et al. [2013].

The term "semi-structured" means that there is structure about bargainers' valuations and beliefs, but players may make offers at any time, and offers can be accompanied by natural language. The typical finding is that in face-to-face and unstructured communication via message-passing, there are *fewer* disagreements than predicted by theory. (A comparable finding in sender-receiver games is that senders willingly share more private information than is selfishly rational; see Crawford [2003], Cai and Wang [2006], Wang et al. [2010a].) However, in the highly structured case in which the bargainers can only make a single offer and no natural language is allowed, disagreements are more common, and the key predictions of theory hold surprisingly well [Radner and Schotter, 1989, Rapoport et al., 1995, Rapoport and Fuller, 1995, Daniel et al., 1998].

Since the rise of structured bargaining theories, experimentation in economics on unstructured bargaining has all but disappeared. Our paper returns to this less popular route, exploring unstructured bargaining with one-sided private information in an experiment.

There are three good reasons to revive the study of unstructured bargaining now.

First, most natural two-player bargaining is *not* highly structured. Conventional methods for conducting bargaining do emerge in natural settings, but these methods are rarely constrained because there are no penalties for deviating from conventions. Studying unstructured bargaining is of particular importance, as strategic behavior may substantially differ between static and dynamic environments that allow continuous-time interaction [Friedman and Oprea, 2012]. There may also be clear empirical regularities in unstructured bargaining—such as deadline effects [Roth et al., 1988, Gächter and Riedl, 2005]—that are evident in the data but not predicted by most theories (though see [Fuchs and Skrzypacz, 2013]). Establishing these regularities can *lead* theorizing, rather than test theory.

Second, even when bargaining is unstructured, theory can still be applied to make clear, interesting predictions. A natural intuition is that when bargaining methods are unstructured, no clear predictions can be made, as if the lack of structure in the bargaining protocol must imply a lack of structure (or precision) in predictions. This intuition is just not right. In the case we study, clear predictions about unstructured bargaining do emerge, thanks to the wonderful "revelation principle" [Myerson, 1979, 1984]. This principle has the useful property of implying empirical predictions for non-cooperative equilibria, independently of the bargaining protocol, based purely on the information structure. For example, the application of the revelation principle in our setting leads to the prediction that strikes will become less common as the amount of surplus the players are bargaining over grows. This type of prediction is non-obvious and can be easily tested. Furthermore, if additional assumptions are made about equilibrium offers, and combined with the revelation principle, then exact numerical predictions about offers and strike rates can be made. That is, even if the bargaining protocol lacks structure, predictions can have plenty of restricted "structure" thanks to the beautiful game theory.

Third, unstructured bargaining creates a large amount of interesting data during the bargaining process. Players can make offers at any time, retract offers, and so on. Natural language can be analyzed, perhaps including vocal properties in verbal communication. Self-reported and biological measures of emotion, cognitive effort, visual attention to display elements, and even neural activity can also be gathered.

Our view is that theoretical and experimental economists regarded these types of data as a nuisance—a "bug" in an experimental design rather than a "feature." Such data seem like a nuisance if one does not have a theory to say anything about them. However, if process variables are systematically associated with outcomes, these empirical regularities both challenge simple equilibrium theories and invite new theory development.

To this end, we use a very limited type of process data in a new way: To predict which bargaining trials will results in deals and strikes. We use a penalized regression approach from machine learning, to select predictive features from a large number of process features. Over-fitting is controlled by making out of sample, cross-validated predictions. We find that a machine-learned predictive model based only on process features can predict strikes about as accurately as the pie sizes can. Adding both process and pie size together makes even better predictions.

Process data are also useful because practical negotiation advice often consists of simple heuristics about how to bargain well [Pruitt, 2013]. For example, negotiation researchers have long ago postulated that initial offers might serve as bargaining anchors and that various psychological manipulations, such as perspective taking, could potentially bias bargaining outcomes.[6] Advice like this can be easily tested by carefully controlled experimental designs that allow unstructured bargaining while keeping the process fully tractable, such as our paradigm.

The closest precursor to our design is Forsythe, Kennan, and Sopher (henceforth FKS), who studied unstructured bargaining with one-sided private information about the sizes of two possible pies [Forsythe et al., 1991b]. They used mechanism design to identify properties shared by all Bayesian equilibria of any bargaining game, using the revelation principle [Myerson, 1979, 1984]. This approach gives a "strike condition" predicting when disagreements would be ex-ante efficient. They tested their theory by conducting several experimental treatments, with free-form communication. The results qualitatively match the theory. We generalize their earlier model to capture

[6] See Kristensen and Gärling [1997], Galinsky and Mussweiler [2001], Van Poucke and Buelens [2002], Mason et al. [2013], Ames and Mason [2015].

any finite number of pie sizes. Because there are several different pie sizes, equilibria which maximize efficiency or equality create different predictions, which we test. Our experimental design uses 6 pie sizes with rapid bargaining (10 seconds per trial), where bargaining occurs only through visible offers and counter-offers, with no other restrictions. They also did not analyze their process data at all, whereas we use machine learning analysis of the process features to predict strikes on a trial-by-trial basis.

Another branch of literature that is related to our study is the experimental work investigating how humans resolve trade-offs between equality and efficiency. While this question is still under a (heated) debate[7], it is largely accepted that people are heterogeneous with respect to how they prioritize these factors.[8]

A few recent papers have investigated highly structured strategic interactions [De Bruyn and Bolton, 2008, Blanco et al., 2011, López-Pérez et al., 2015, Jacquemet and Zylbersztejn, 2014], and some have examined free form bargaining with full information [Herreiner and Puppe, 2004, Galeotti et al., 2018]. We extend this literature by deriving theoretical predictions and test empirically how humans resolve the equality-efficiency trade-off in a dynamic strategic environment with informational asymmetry.

Finally, our study closely relates to negotiation research [Pruitt, 2013], a branch of social psychology and organizational behavior research. In contrast to economic theories that typically describe behavior in equilibrium (i.e., when players best respond to each other's actions), negotiation theories assume that bargainers are not in equilibrium and focus on prescriptive models, in which adopting certain strategies improves negotiation outcomes. Nego-

[7] See Kritikos and Bolle [2001], Charness and Rabin [2002], Engelmann and Strobel [2004, 2006], Fehr et al. [2006], Bolton and Ockenfels [2006], Durante et al. [2014], El Harbi et al. [2015].

[8] For example, economics students are inclined to favor efficiency over equality, females are more egalitarian than males, and political preferences do not seem to have an effect [Engelmann and Strobel, 2004, Fehr et al., 2006].

tiation researchers take into account the process of bargaining by studying psychological constructs such as aspirations, defined as "the highest valued outcome at which the negotiator places a non-negligible likelihood that that value would be accepted by the other party" [White and Neale, 1994]. Aspirations play an important role in determining the bargainers' initial offers, and were shown to influence bargaining outcome variables such as disagreement rates and surplus division.[9]

The reminder of this paper is organized as follows. In section 3.2.1, we use mechanism design theory to derive general qualitative properties of bargaining in equilibrium. We show that in our setting, no equilibrium satisfies both equality and efficiency in all states of the world, and propose two equilibria that solve this trade-off by either favoring the former or the latter. We present a novel experimental design in section 3.2.2, and summarize its general results in section 3.3.1. We use machine learning to examine how bargaining process data can be associated with bargaining outcome variables in section 3.3.2, and conclude in section 3.4.

3.2 Theory and Experiments

In this paper, we adopt the theoretical framework from Camerer et al. [2019] who extend the two state model developed in Kennan and Wilson [1990] and Forsythe et al. [1991b] to an arbitrary finite number of states. This extension yields comparative statics predictions regarding the frequency of disagreements in each state with only the game structure, incentive compatibility (IC) and individual rationality (IR) constraints. However, the mechanism design approach only characterizes the class of possible equilibria rather than predicts specific outcomes. Thus, Camerer et al. [2019] further take advan-

[9]See Yukl [1974], White and Neale [1994], White et al. [1994], Kristensen and Gärling [1997], Galinsky and Mussweiler [2001], Van Poucke and Buelens [2002], Buelens and Van Poucke [2004], Mason et al. [2013], Ames and Mason [2015].

tage of the focal points in this game to obtain testable predictions about both deal rates and payoffs in each state. In this section, we first introduce the theoretical framework and then the experiments.

3.2.1 Theoretical Framework

In this unstructured bargaining game, two players bargain over an economic surplus or "pie," which is a random variable denoted by π. The finite set of true states indexed by $k \in \{1, 2, \ldots, K\}$, and the pie amount in state k is π_k. Without loss of generality, we assume $\pi_k > \pi_j$ when $k > j$. The informed player knows the true pie amount. The uninformed player does not know the pie amount, but knows the informed player knows it. The probability distribution over pie sizes $\Pr(\pi_k) = p_k$ is common knowledge. The players bargain over the payoff of the uninformed player, denoted by w, by continuously communicating their bids within a certain amount of time T—which is also a common knowledge. If the players agree on a payoff for the uninformed player w, then the informed player gets the rest of the pie $\pi - w$. If they do not agree on an allocation before the deadline, both players get nothing and we refer to this outcome as a disagreement, or in keeping with the motivation of Forsythe et al. [1991b], as a *strike*, while successful bargaining outcomes are *deals*.

From a mechanism design perspective, we can view this bargaining game as a process of transmitting the private information about the pie size from the informed player to the uninformed player. By the revelation principle (Myerson [1979, 1984]), we know that every Nash equilibrium in this bargaining game can be implemented in an incentive compatible direct mechanism where the informed player truthfully reports the actual state to a neutral mediator and the player's payoffs are equal to their payoffs in the original bargaining game.

Following Forsythe et al. [1991b] and Camerer et al. [2019], in the direct mechanism the informed player announces that the state is $j \in \{1, \dots, K\}$. Given the announcement, the neutral mediator determines the deal probability (γ_j) and the payoff to the the uninformed player (x_j). The informed player gets the rest of the pie ($\gamma_j \pi_k - x_j$). Thus a mechanism involves $2K$ parameters, $\{\gamma_k, x_k\}_{k=1}^K$.

A mechanism is incentive compatible (IC) if it is optimal for players to reveal their private information. In our setting, this means that the informed player's expected payoff must be (weakly) maximized in the direct mechanism when she announces the true size of the pie. This requires

$$\gamma_k \pi_k - x_k \geq \gamma_j \pi_k - x_j, \quad \forall k \text{ and } \forall j \neq k. \tag{IC}$$

An IC-mechanism is individually rational (IR) when both players prefer to participate in it. Assuming the players' payoffs from not participating are zero, this means that for every state k the expected payoff to each player is positive, so that

$$\gamma_k \pi_k - x_k \geq 0, \tag{IR-1}$$

$$x_k \geq 0. \tag{IR-2}$$

Based on the IC, IR-1, IR-2 and conditions, Camerer et al. [2019] prove the following two lemmas.

Lemma 3.1. *If the bargaining mechanism satisfies the IC, IR-1, and IR-2 conditions, then:*

1. *Deal rates are monotonically increasing in the pie size x_k.*

2. *The uninformed player's payoffs are monotonically increasing in the pie size.*

3. *The uninformed player's payoff is identical for all states in which the deal probability is 1.*

Note that the payoff of the uniformed player x_k in the direct mechanism is equivalent to the expected payoff of the uninformed player in state k of the bargaining game: $x_k = \gamma_k w_k$, where w_k is the uninformed player's payoff conditional on a deal being made in state k.

The direct mechanism is interim-efficient [Holmström and Myerson, 1983] if the payoff profile is Pareto-optimal for the set of $K+1$ agents: the informed player in each of the K possible states of the world and the uninformed player (in expectation). Interim efficiency implies the following *strike condition*:

Lemma 3.2. *For any mechanism that satisfies the IR-1, IR-2 and IC conditions, strikes in state k are interim-efficient if*

$$\frac{\pi_k}{\pi_{k+1}} < \frac{\left(1 - \sum_{j=1}^{k} p_j\right)}{\left(1 - \sum_{j=1}^{k-1} p_j\right)} = \frac{\Pr(\pi \geq \pi_{k+1})}{\Pr(\pi \geq \pi_k)}.$$

For the proof, see Camerer et al. [2019].

The IC, IR-1, IR-2, and strike conditions limit the scope of possible bargaining outcomes and predict when strikes are likely to occur. However, they are not sufficient to pin down the strike rates $1 - \gamma_k$ and the equilibrium payoffs w_k in each state. To make a more precise prediction, Camerer et al. [2019] use an equilibrium selection approach which assumes that equal payoff splits are natural focal points. In the experiments, the possible states, π, takes on values that are the integer dollar amounts between \$1 and \$6 with equal probability. Therefore, we can restrict the state space to $\{\$1, \ldots, 6\}$.

The importance of focal points has been well-studied in the literature (Schelling [1960], Roth [1985], Kristensen and Gärling [1997], Janssen [2001], Binmore and Samuelson [2006], Janssen [2006], Bardsley et al. [2009], Isoni et al. [2013, 2014]). Absent other salient features of bargaining, the natural focal point is an equal split (i.e., $w_k = \frac{\pi_k}{2}$). Indeed, equal splits often emerge in bargaining experiments (e.g. Lin et al. [2018]). Based on players' tendency to coordinate on the equal-split allocation, Camerer et al. [2019] propose that the equilibrium payoff of the uninformed player, conditional on a deal, will equal half of the pie size ($w_k = \frac{\pi_k}{2}$) as long as an equal split satisfies the IR and IC conditions (Lemma 3.1), and subjects to efficiency conditions. By either prioritizing the former or the latter, Camerer et al. [2019] derive two competing equilibrium predictions, which are **the efficient equilibrium** and **the equal split equilibrium**.

The Efficient Equilibrium

The IC conditions and Lemma 3.1 show that if there exists a cutoff state π_c where the deal rate γ_c is equal to 1, then strikes are inefficient in all states π_k such that $k \geq c$. Lemma 3.1 also predicts that the uninformed player's payoff must be the same in all states where no disagreements occur.

The efficient equilibrium prioritizes efficiency over equality. In this equilibrium, the deal rate is assumed to be 1 whenever the strike condition (Lemma 3.2) does not hold. To obtain a precise prediction about the equilibrium uninformed payoffs w_k^*, we assume that players divide the pie equally in lower-value pie states given this constraint:

$$w_k^* = \begin{cases} \frac{\pi_k}{2} \ , \forall \pi_k \leq \pi_c, \\ \frac{\pi_c}{2} \ , \forall \pi_k > \pi_c. \end{cases}$$

In our experiment, π is an integer dollar between 1 and 6 with equal likelihood. It follows numerically that the strike condition (Lemma 3.2) holds for pies of size 1 and 2. When $\pi = 3$, the two sides of the inequality are equal, so the strike rate is indeterminate. When $\pi \geq 4$, there should be no strikes. Therefore, we can first pin down that the deal rates in large pies would satisfy $\gamma_4 = \gamma_5 = \gamma_6 = 1$. Furthermore, as we plug w_k^* into the IC constraint, we can obtain the following predictions regarding payoffs and deal rates in the efficient equilibrium:

$$(w_1, w_2, w_3, w_4, w_5, w_6) = \left(\tfrac{1}{2}, 1, \tfrac{3}{2}, 2, 2, 2\right),$$
$$(\gamma_1, \gamma_2, \gamma_3, \gamma_4, \gamma_5, \gamma_6) = \left(\tfrac{2}{5}, \tfrac{3}{5}, \tfrac{4}{5}, 1, 1, 1\right).$$

The Equal Split Equilibrium

On the other hand, some efficiency must be sacrificed in order to achieve equality for each pie size. In the equal split equilibrium, we first assume that players split the pie equally in all states, and only then maximize efficiency, so that

$$w_k^{**} = \frac{\pi_k}{2}.$$

in all states.

Because deal rates are increasing with the pie size and the uninformed player's payoff must be identical in all states where there are no strikes (Lemma 3.1), full equality implies that efficiency (i.e., no strikes) can only be achieved in the largest pie. Thus, to pin down exact numerical predictions of deal rates in the equal split equilibrium, we set $\gamma_6 = 1$. Inserting $w_k^{**} = \frac{\pi_k}{2}$ into the IC constraint, we derive the predicted payoffs and deal rates in the equal split equilibrium:

$$(w_1, w_2, w_3, w_4, w_5, w_6) = \left(\tfrac{1}{2}, 1, \tfrac{3}{2}, 2, \tfrac{5}{2}, 3\right),$$
$$(\gamma_1, \gamma_2, \gamma_3, \gamma_4, \gamma_5, \gamma_6) = \left(\tfrac{2}{7}, \tfrac{3}{7}, \tfrac{4}{7}, \tfrac{5}{7}, \tfrac{6}{7}, 1\right).$$

Thus the two types of equilibria imply different predictions regarding payoffs and deal rates.

3.2.2 Experiments

Camerer et al. [2019] developed a novel experimental paradigm of dynamic bargaining that allows both parties to communicate offers whenever they please, while keeping their behavior tractable. This experiment was first conducted by Camerer et al. [2019] (Experiment 1), which is the baseline treatment. We also report results from a follow-up experiment with same design but with different treatments (Experiment 2). In this section, we first introduce the experimental design and then the treatments.

Design

The experimental design is a continuous-time bargaining game with one-sided private information. At the beginning of the experiment, subjects are assigned to one of the two roles: the informed player or the uninformed player. Players' roles are fixed for the session's 120 bargaining rounds.

In each round, each informed player is randomly matched with an uninformed player to bargain over a pie with a size unknown to the uninformed player. The pie size is an integer from 1 to 6, i.e. $\pi \in \{\$1, 2, 3, 4, 5, 6\}$ and drawn from a commonly known discrete uniform distribution. The informed player would know the pie size for that round after the draw is made.

Each pair bargained over the uninformed player's payoff, denoted by w. Both players communicate their offers, in multiples of $0.1[10], using a mouse click on a graphic interface which was programmed with z-Tree software (Fischbacher [2007]). During the first two seconds, both players can decide their initial bargaining position without seeing the opponent's position (Figure 3.1A). The initial cursor location is randomized.

[10]In Camerer et al. [2019] (Experiment 1), they set the resolution to be in multiples of $0.2, which is a compromise between $0.1 (too fine a resolution for coordinating in a short game) and $0.5 (a resolution that would not allow for testing the use of focal points, as every possible offer would be a half of an integer pie). However, the result in Experiment 1 shows that players are able to coordinate in such a short period, so we increase the resolution to be in multiples of $0.1 in Experiment 2.

After initial locations are set, the players start a 10-second bargaining round. They communicate the offers with mouse clicks (Figure 3.1B). As both players' positions match, a green vertical stripe would appear on the screen (Figure 3.1C), and this position would become the final deal if there is no change on the position in the following 1.5 seconds (or if the period ends, which ever came first)[11]. If no deal is reached within 10 seconds, both players earn nothing. After each round, the players would be notified their payoffs and the actual pie size (Figure 3.1D).

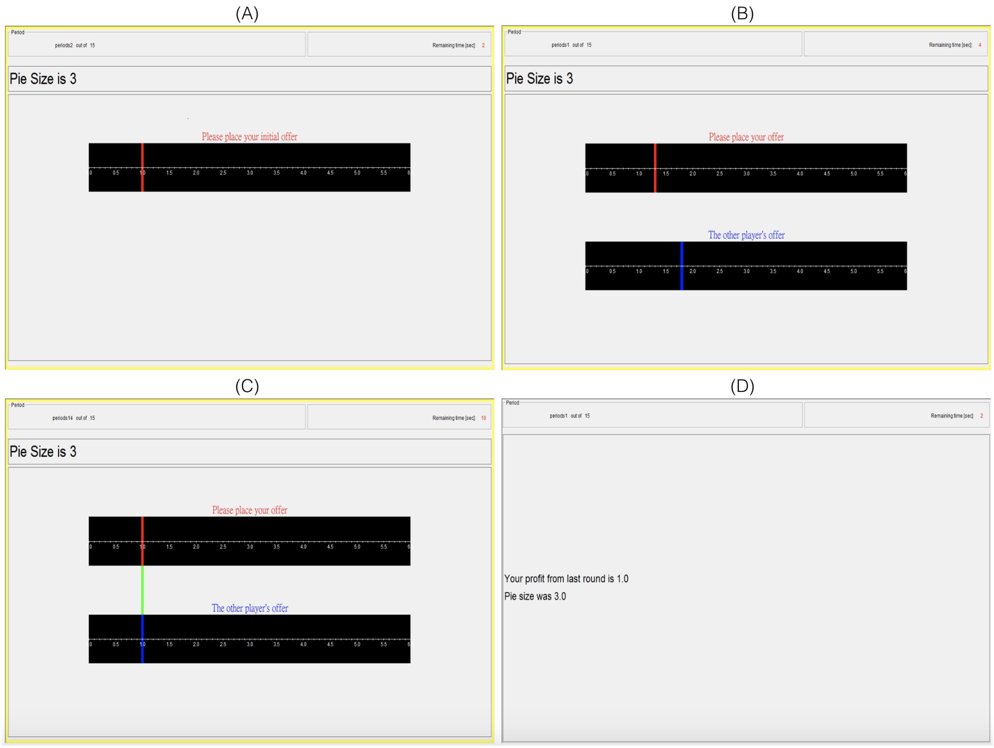

Figure 3.1: (A) Initial offer screen: in the first two seconds of bargaining, both players can set their initial positions without revealing to the opponent. The pie size is on the top left corner and it only appears on informed player's screen. (B) Players communicate their offers using mouse click on the interface. (C) When two players' positions match, the green vertical stripe appears and this would be the deal if there is no change in the following 1.5 seconds. (D) After the bargaining round, both players would be notified about their payoffs and the pie size.

[11] In Experiment 1, the offers have to match for 1.5 seconds in order to make a deal. In other words, the latest time where the players' bids can match is $t = 8.5$ seconds.

Experiment 1

Camerer et al. [2019] conduct eight experimental sessions, five in the Social Science Experimental Laboratory (SSEL) at Caltech and three in the California Social Science Experimental Laboratory (CASSEL) at UCLA. In the beginning of each session, subjects are randomly assigned to isolated computer workstations and are handed printed versions of the instructions (see Appendix D in Camerer et al. [2019]). The instructions are also read aloud by the experimenter (who is the same person in all sessions). All of the participants complete a short quiz to check their understanding of the task. Subjects play 15 practice rounds to become familiar with the game and the interactive interface before the actual play of 120 rounds. Participants' payoffs are based on their profits in a randomly chosen 15% of the rounds, plus a show-up fee of $5. Each session lasts between 70 and 90 minutes (including check-in, reading of instructions, experimental task, and payment).

Experiment 2

The follow-up experiment is conducted in the Taiwan Social Science Experimental Laboratory (TASSEL) at National Taiwan University. We conduct eight experimental sessions. Three sessions are female-informed sessions where female subjects take the role of informed players and played against uninformed male subjects. Another three sessions are male-informed sessions which have the opposite design to the female-informed sessions. In the female-informed and male-informed sessions, we require an equal number of male and female subjects. Subjects are only notified of this requirement when entering the experiment. In addition, we conduct one experienced session and one high-stake session in order test whether our results are robust to experience and stakes. In the experienced session, we recruit subjects who have participated one of the six previous sessions. In the high-stakes session, we multiply the stakes by 5. Notice that there is no gender constraint in the experienced and high-stake session.

The experimental procedures are the same in Experiment 1 and Experiment 2. In Experiment 2, participants' payoffs are based on their profits in a randomly chosen 10% of the rounds, plus a show-up fee of NT$ 100. Payoffs in the experiments are converted into NT$ according to a pre-set exchange rate (1 ECU = NT$15) specified in the instructions. In the high-stake session, the exchange rate is 1 ECU = NT$75 while the exchange rate is 1 ECU = NT$30 in the experienced session.

After 120 rounds of the bargaining game, we measure subjects' risk preferences and loss aversion by Dynamically Optimized Sequential Experimentation (DOSE) developed by Wang et al. [2010b]. In each round, subjects are asked to choose from 2 lotteries. Lottery 1 is a risky asset, while lottery 2 yields a fixed amount. There are 3 practice rounds and 40 paid rounds. At the end of the experiment, 12 rounds from the bargaining game and 1 round from DOSE would be drawn and realized. Before undergoing DOSE, all subjects evaluated their subjective willingness to take risk on a scale from 0 (not willing to take any risk at all) to 10 (willing to take any risk). The evaluation would not affect the payoff. Each session lasts around 2.5 hours.

3.3 Experimental Results

3.3.1 Basics

In this section, we focus on analyzing the deal rates across different treatments. See Camerer et al. [2019] and Online Appendix for further analysis on the payoffs and the bargaining dynamics.

Table 3.1 provides the summary statistics of average bargaining outcomes in different treatments. The average bargaining outcomes are similar across treatments. Differences in the average payoffs across treatments are less than \$0.1 and differences of average deal rates are within 5%. We high-light some of our findings in the following: The average surplus loss is the lowest in the experienced treatment and the highest in the male-informed treatment. Turning to the information value, which can be interpreted as the advantage of knowing the pie size, we observe that it is the largest in the experienced treatment and lowest in the baseline treatment. Bargaining outcomes are generally robust across different treatments and stakes on the aggregated level.

Table 3.1: Summary Statistics for Different Treatments

Treatment	Baseline	Female	Male	Experienced	High-Stake
Informed Payoff[a]	2.01	2.08	2.09	2.10	2.04
	(0.03)	(0.02)	(0.06)	–	–
Uninformed Payoff[a]	1.49	1.42	1.41	1.40	1.46
	(0.03)	(0.02)	(0.06)	–	–
Deal Rate	0.61	0.66	0.62	0.66	0.65
	(0.03)	(0.02)	(0.02)	–	–
Surplus Loss[b]	1.13	1.02	1.18	0.96	1.11
	(0.08)	(0.09)	(0.09)	–	–
Information Value[c]	0.40	0.51	0.49	0.54	0.42
	(0.03)	(0.05)	(0.06)	–	–

Means and standard errors (which are shown in parentheses) are calculated by treating each session's mean as a single observation. Since there is only one session for experienced and high-stake treatment, the standard errors for these two treatments are not computable.
[a] Averages are calculated for deal games only.
[b] Surplus loss = the mean expected loss of pie due to strikes.
[c] Information value = the mean difference between the informed and uninformed payoffs.

Next, we break down deal rates according to different pie sizes for different treatments. Figure 3.2 and Figure 3.3 show that in all treatments, deal rates increase with the pie size. This confirms our theoretical prediction in Lemma 3.1. Moreover, deal rates in female-informed sessions and the experienced session are higher than the baseline sessions in all pies (except the largest pie). On the other hand, deal rates in male-informed sessions and the high-stake session are higher than the baseline in small pies ($\pi \leq 3$), but lower in large pies ($\pi \geq 4$).

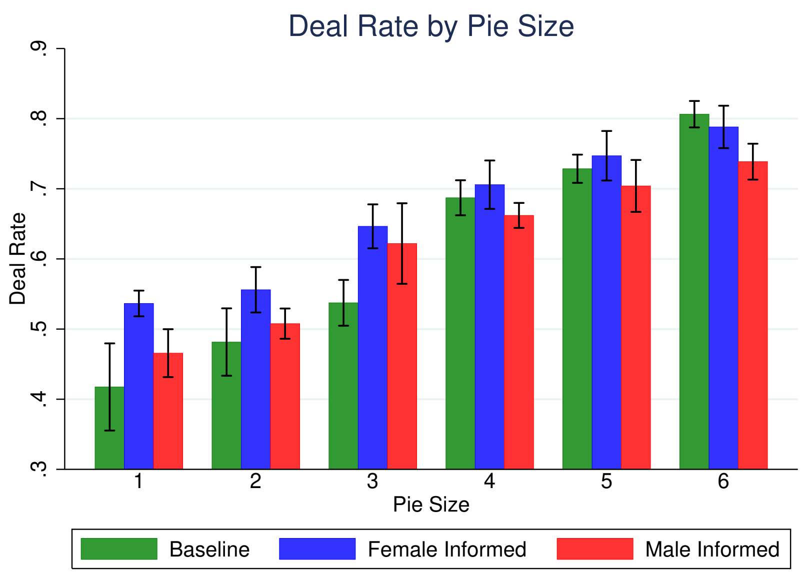

Figure 3.2: The deal rates under different pie sizes and treatments. The green bars stand for the average deal rates of baseline sessions at different pie size. The blue and red bars are for female-informed sessions and male-informed sessions, respectively. The standard errors (overlaid on the bars) are calculated at the session level.

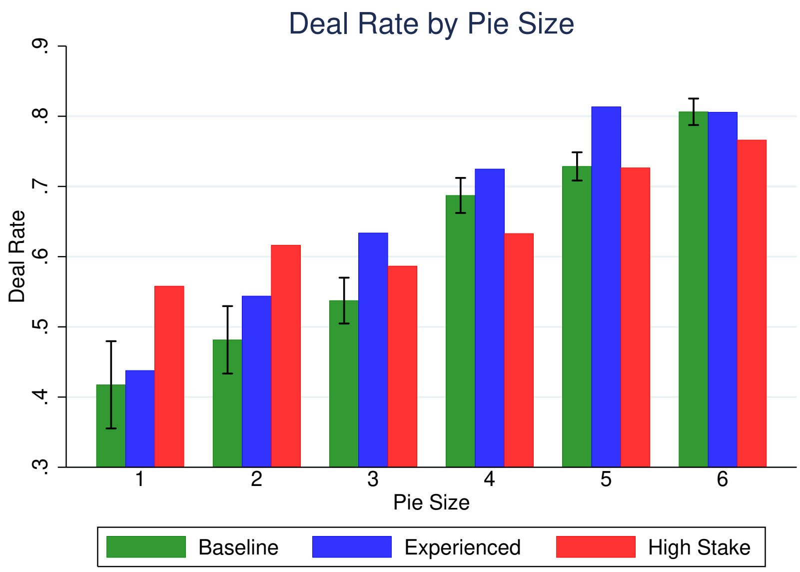

Figure 3.3: The deal rates under different pie sizes and treatments. The green bars stand for the average deal rates of baseline sessions at different pie size. The blue and red bars are for the experienced sessions and high-stake sessions, respectively. The standard errors (overlaid on the bars) are calculated at the session level. Since there is only one session for experienced and high-stake treatment, the standard errors are not computable for these two treatments.

We defer further results from Experiment 2 to the Appendix. These results include analyses of the bargaining dynamics (see Online Appendix C.1) and testing predictions in Lemma 3.1 (see Online Appendix C.2). In general, the results in Camerer et al. [2019] are replicated by Experiment 2. Besides the monotone increase of deal rates and payoffs, we also observe that the equal-split allocation is the most salient focal point. Regarding the dynamics, we observe that the informed players' offers increase, and the uninformed players' demands decrease with time (within a trial). There is also a strong deadline effect—most of the deals are reached close to the deadline. Lastly, we analyze the differences in equilibrium selections using regression.

3.3.2 Outcome Prediction via Machine Learning

The unstructured paradigm established by Camerer et al. [2019] records, in addition to initial demands and offers, a large amount of bargaining process data that may be used to predict disagreements before the deadline. For instance, suppose that at the five-second mark, neither player has changed her offer for more than three seconds. This mutual stubbornness might be associated with an eventual strike. We consider a large number of such candidate observable features in search of a small set that is predictive, using cross-validation (Stone [1974]) to control for over-fitting. This machine learning approach has been used in many applications in computer science and neuroscience, and starts to be more widely adopted in economics (e.g. Krajbich et al. [2009]) and other social sciences (e.g. Dzyabura and Hauser [2011]).

In this paper, we treat Experiment 2 as the lockbox test for the predictive model built in Camerer et al. [2019]. Therefore, in this section we report the results from directly feeding the data from Experiment 2 into the model. First of all, we briefly introduce the algorithm here. We follow Camerer et al. [2019] to choose 35 behavioral features recorded during bargaining. Examples of features are the current difference between the offer and demand, the time since the last position change, and an indicator denoting which player had changed his or her position in the game first. The full list of features is presented in Camerer et al. [2019]. For each of the eight sessions in Experiment 2, we trained a model to classify trials into disagreements or deals using the data of the remaining seven sessions. The classification is done by estimating a logistic regression with a least absolute shrinkage and selection operator (LASSO) penalty (Tibshirani [1996]). By applying these trained models, we then made out-of-sample predictions of the binary bargaining outcomes for each of the experimental session.

To assess the predictive power of process data, we estimate three strike prediction models at eight different points in the bargaining process, separated by one-second intervals (i.e., $1, 2, \ldots, 8$ seconds after bargaining starts). The first model relies only on the pie size, the second uses only process features and the third uses both pie size and process features. For each time stamp, predictions are carried out using the following nested cross-validation procedure: For each of the eight sessions, we train a linear model with the seven other sessions to predict the outcome (a deal or a strike) by fitting a logistic LASSO regression. The tuning parameter, λ, is optimized via ten-fold cross-validation, performed within each training set. Finally, using that trained model, we conduct out-of-sample predictions for the holdout sessions.

To compare the three models, we use the "receiver operating characteristic" (ROC) curves (Hanley and McNeil [1982], Bradley [1997]), a standard tool in signal detection theory which quantifies the performance of a binary classifier under different trade-offs between type I and type II errors. For a random classifier, the true positive and false positive rates are identical (the 45-degree line in Figure 3.4). A good classifier increases the true positive rate (moving up on the y axis) and decreases the false positive rate (moving left on the x axis). The difference between the ROC and the 45-degree line, in the upper-left direction, also known as the "area under the curve" (AUC) is an index of how well the classifier does.

Figure 3.4 shows the ROC curve at $t = 2, 5, 7$ seconds for both Experiment 1 and 2. The ROC analysis indicates that process data do better than random for every time stamp in both experiments. Moreover, the fitness of models with process data increase with time, but the same is not true for the model with pie size only.

While patterns of AUC are similar in Experiment 1 and 2, there are still some subtle differences. In Experiment 1, the model with pie and process features always has the best predictive power and the other two models are not so distinguishable in later seconds. On the other hand, even though the model with pie and process feature is the best model among the three, its

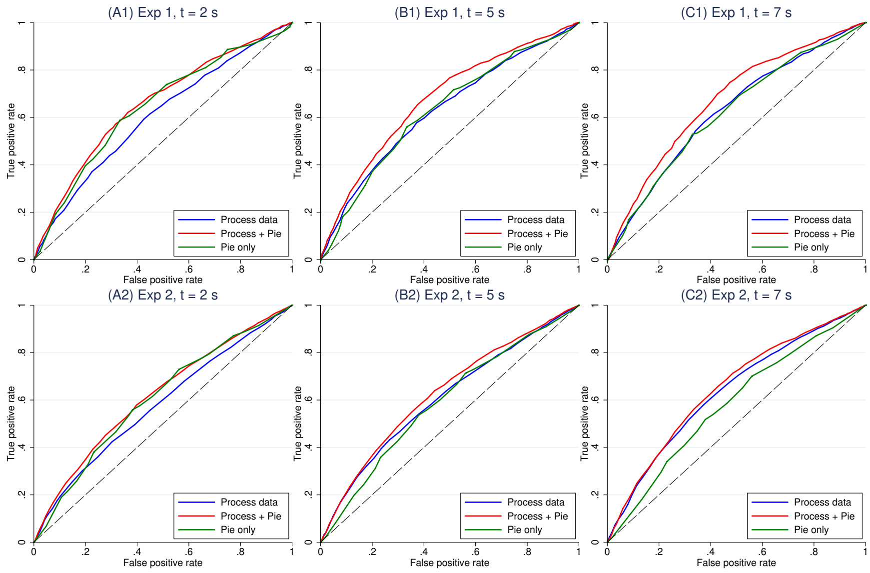

Figure 3.4: Receiver operating characteristic (ROC) for predicting disagreements, two, five and seven seconds into the bargaining game. The dashed lines represent the false and true positive rates of a random classifier. (A1–C1) show the data from Camerer et al. [2019] (Experiment 1) and (A2–C2) plot the result from Experiment 2.

predictive power is not significantly stronger than the model with process features only.

To further investigate which behavioral process features predict strikes, we follow Camerer et al. [2019] and use a "post-LASSO" procedure proposed by Belloni et al. [2013, 2012]. Figure 3.6 summarizes the marginal effects of all process features (z-scored for every time point) in both experiments. The general feature patterns in Experiment 2 are consistent with those in Experiment 1.

The current informed player's offer (positively correlated with a deal) and the current difference between the players' bargaining positions (positively correlated with a strike) are the most predictive process features. One surprising finding in Camerer et al. [2019] is that initial bargaining positions contain predictive information regarding the chance of reaching a deal, even as the deadline approaches, and even after controlling for current offers. In

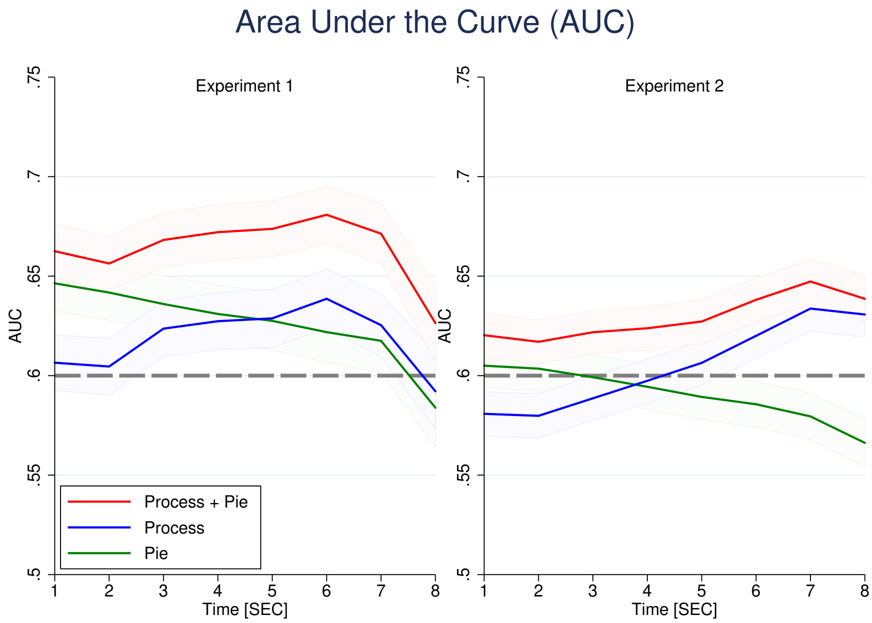

Figure 3.5: Area under the curve (AUC) of disagreements classifiers using process data, pie size, and the two combined. Note that the classifier's input included only trials that were still in progress (when a deal has not yet been achieved), and excluded trials in which the offers and demand were equal at the relevant time stamp. The left figure is the original result from Camerer et al. [2019] (Experiment 1) and the right one is the result from Experiment 2.

Experiment 2, such effect of initial positions is even stronger. We also find a negative interaction between initial offer and initial demand and a negative interaction between initial and current offer, which again confirms the arguments in Camerer et al. [2019].

3.4 Conclusion

Experimentation in economics on bargaining mostly abandoned unstructured bargaining because there is too little experimental control over all the things that bargainers can do. Unstructured bargaining seems to hand over the reins of endogenous "treatments" to the experimental subjects.

If the goal is prediction rather than theory-testing, however, having a

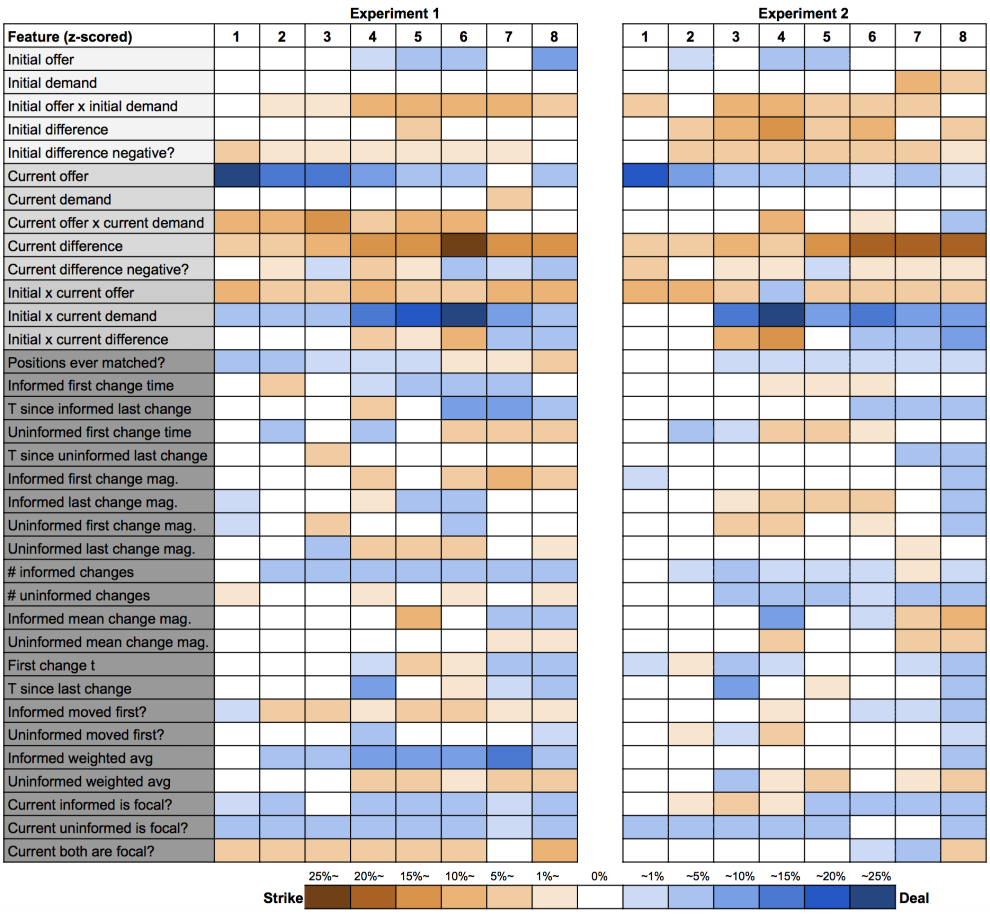

Figure 3.6: Bargaining Process Features Selected by the Classifier for Outcome Prediction (Deal= 1) and Their Estimated Marginal Effects. The left panel is the result from Camerer et al. [2019] (Experiment 1) and the right panel is the result from Experiment 2. The pie sizes are excluded.

large amount of data is terrific. For machine learning applications, there is (almost) no such thing as too much data.

Furthermore, a kind of theory-testing can still be done in a machine learning framework. In our example, the revelation principle, along with other restrictions, still delivers predictions about what will happen in equilibrium which are highly independent of the unstructured behavior. Everything depends on pie sizes. A lean predictive "machine" using only pie sizes should therefore predict as accurately as one with many more features.

Our main finding, using National Taiwan University subjects and some small design changes, is a close replication of earlier results using Caltech subjects. Agreements are often equal splits, even though the exact pie size is only known to one side. Deal rates do increase with pie size, but there is a lot of inefficiency—deal rates are too low—compared to revelation principle predictions. However, theory also predicts a break for uninformed offers for pies of $4-6 compared to lower pie amounts, and this break is evident in the data. There are some experience effects (deal rates go up across trials in an experimental session). One session with twice-experienced subjects—repeating the entire experimental session—did not produce results much closer to equilibrium (to our surprise).

There are only weak effects of gender. When females are informed, the deal rate is a bit higher and uninformed (males) get a little less, but the evidence is not statistically strong. While gender effects in bargaining are interesting, a lot more statistical power is probably needed to establish whether there are differences or not [12].

Finally, we hope these data and method inspire other experimenters from a range of social sciences to measure a lot more about what goes in the bargainers' bodies and brains, and results from their typing or talking, on during bargaining. For example, Forsythe et al. [1991a] allowed subjects to transmit verbal messages during bargaining. At the time, methods of analyzing natural language processing (NLP) were so primitive that they did not do any sophisticated analysis of those rich data. What they wrote at the time was:

> Because of the unconstrained messages which pass between the players, our bargaining game is too complicated to allow a detailed strategic analysis. However, by the revelation principle, any Nash equilibrium of this game is equivalent to some direct mechanism,

[12] And of course, gender differences are likely to vary wildly across the globe, so a serious attempt to understand such differences must look at the influences of developmental life cycle, biological factors such as hormones, and cultural variation.

> which specifies whether a strike should occur and how much each bargainer should get as a function of the informed bargainer's announcement of the size of the pie.

That is, While they allowed messages and recorded them, they did not analyze them at all because the resulting game—treating messages as strategy choices—is too complicated to solve. Using the messages as data in machine learning does not test a theory either, but it provides preliminary evidence of how features of messages influence agreement rates. Such evidence could provide inspiration for theory.

It is also notable that recording messages is very easy technically. NLP is one area of machine learning which is now hugely successful and improving by leaps and bounds every year. In general, machine learning methods are hungry for any such choice process data. And now we know what to do with them.

Appendix A

Appendix to Chapter 1

A.1 Appendix: Omitted Proof

A.1.1 Proof of Lemma 1.1

First, we show that Lemma 1 holds for any p on the equilibrium path. In the brand market, if $\mu_m^*(p,1) = 0$, then $\mu_h^*(p,1) = 1$, $\gamma^* = 0$ and $\pi_i^* = \pi_i^*(p,1,e_h)$. However, for firm i, $\pi_i'(p,1,e_m) > \pi_i^*$.

In the stand-alone market, if $\mu_l^*(p,0) = 0$, then $\mu^* = (1-\mu_m, \mu_m, 0)$ for some $\mu_m \in [0,1]$. If μ_m induces $\gamma^* = 0$, the argument for the brand market applies. If μ_m induces some $\gamma^* > 0$, $(1-\mu_m)q_h + \mu_m q_m - p = (1-\mu_m)\mathbb{1}(q_h \geq p)(q_h - p) + \mu_m \mathbb{1}(q_m \geq p)(q_m - p) - s \geq 0$. For any $s > 0$, it follows that $p \in (q_m, q_h)$. For firm j where $\pi_j^* = \pi_j^*(p,0,e_m) = (1-\gamma^*)p - c_i e_m$, $\pi_j'(p,0,e_l) = (1-\gamma^*)p - c_i e_l > \pi_j^*$. μ_m can never induce $\gamma^* = 1$ because firms can only choose e_h and there will be no incentive for consumers to inspect.

A.1.2 Proof of Lemma 1.2

Consider the brand market first. Suppose there exists p^* such that $\gamma^*(p^*,1) = 0$ and $\mu_m^*(p^*,1) < 1$. For p^* to construct an equilibrium, it must be that $\Sigma_{k\in\{h,m\}}\mu_k^* u(k) - p^* \geq 0$. In addition, there exists i such that $\pi_i^* = \pi_i^*(p^*,1,e_h) =$

$p^* - c_i e_h$. However, $\pi_i' = \pi_i'(p^*, 1, e_m) = p^* - c_i e_m > \pi_i^*$. e_h cannot be the best response of any firm. Similarly, in the individual market, the only best response of effort is e_l if $\gamma^* = 0$.

A.1.3 Proof of Lemma 1.3

Suppose some quality-level good x is produced in both markets. Denote the expected revenue in brand market p^b and that in individual market p^i. Then $\exists c_1, c_2$ such that

$$p^b - c_1 e_x - F \geq p^i - c_1 e_x \text{ and } p^b - c_2 e_x - F < p^i - c_2 e_x.$$

The two inequalities contradict each other.

A.1.4 Proof of Proposition 1.1

Lemma 1.1 and 1.2 show that when inspection is impossible, the only consistent belief $\mu^*(p, 1) = (0, 1, 0)$ and $\mu^*(p, 0) = (0, 0, 1)$. Given this belief, in the consumption stage, consumers purchase if $p \leq q_l$ in the stand-alone market, and if $p \leq q_m$ in the brand market. Since there is no competition, if firm i joins the brand, $(p_i^*, e_i^*) = (q_m, e_m)$. Otherwise, $(p_i^*, e_i^*) = (q_l, e_l)$.

In the contract stage, the brand-holder faces the following maximization problem:

$$\begin{aligned} \max_{\{F\}} \quad & \hat{c}F \\ \text{s.t.} \quad & q_m - c_i e_m - F \geq q_l - c_i e_l \quad \forall c_i \leq \hat{c} \\ & q_l - c_i e_l \geq q_m - c_i e_m - F \quad \forall c_i > \hat{c} \end{aligned} \tag{A.1}$$

The constraints yield $F = q_m - q_l - (e_m - e_l)\hat{c}$. The maximization problem can be re-written as the following:

$$\max_{\{\hat{c}\}} \quad \hat{c}[q_m - q_l - (e_m - e_l)\hat{c}] \tag{A.2}$$

The objective function is a concave quadratic function of c, so global maximum exists. The first-order condition yields $q_m - q_l - 2(e_m - e_l)c = 0$. Therefore, $\hat{c} = \frac{1}{2}\frac{q_m - q_l}{e_m - e_l}$ and $F^* = q_m - q_l - (e_m - e_l)\hat{c} = \frac{q_m - q_l}{2}$.

A.1.5 Proof of Proposition 1.2

We first show the existence of the inspection equilibrium. Denote the equilibrium high, medium and low prices p_h, p_m, and p_l, respectively. In the consumption stage, consumers purchase if $\Sigma_{k\in\{h,m,l\}}\mu_k^I u(k) - p \geq 0$. If $p = p_l$, $\mu = (0,0,1)$. If $p = p_m$, $\mu = (0,1,0)$. Therefore, consumers' purchasing decision can be simplified to purchase if $p_l \leq q_l$, $p_m \leq q_m$ and if $p_h \leq \mu_h q_h + \mu_m q_m$ with no inspection and $p_h \leq q_h$ with inspection and the quality is q_h.

$$\mu(p,b) = \begin{cases} (0,1,0) & \text{if} \quad (p,b) = (p_m, 1) \\ (\rho, 0, 1-\rho) & \text{if} \quad (p,b) = (p_h, 0) \\ (0,0,1) & \text{otherwise} \end{cases}$$

where $\rho = \frac{c}{\underline{c} + \sigma(1-\bar{c})}$ is the probability that a high-price good is of high quality.

Therefore, the optimal inspection decision $\gamma(p,b) = 0$ if $(p,b) = (p_l, 0)$ or $(p_m, 1)$.

If $(p,b) = (p_h, 0)$, consumers inspect if

$$\rho(q_h - p_h) - k \geq \rho q_h + (1-\rho)q_l - p_h. \tag{A.3}$$

If the inequality holds with equality, consumers are indifferent between inspection and no inspection, any $\gamma \in [0,1]$ is optimal. Hence, in equilibrium, $(1-\rho)(p_h - q_l) = s$. Supporting the equilibrium, consumers may as well choose $\gamma = \frac{p_h - p_l}{p_h}$.

Next, we solve for firms' optimal decisions. It is clear to see that branded firms price at $p_m = q_m$ and choose effort level e_m. For individual firms, expected profit

$$E[\pi_i(p_i, e_i)] \begin{cases} p - c_i e_l & \text{if} \quad (p_i, e_i) = (q_l, e_l) \\ p - c_i e_h & \text{if} \quad (p_i, e_i) = (p_h, e_h) \\ (1-\gamma^*)p - c_i e_l & \text{if} \quad (p_i, e_i) = (p_h, e_l) \end{cases}$$

Given consumer's indifference constraint in inspection, firms with high effort can at most extract $p_h = \rho q_h + (1-\rho) q_l$. An off-equilibrium belief which supports this equilibrium is that for any price $p' > p_h$, consumers assign probability 1 to quality q_l. The upper bound of the high price and consumers' indifference condition in inspection pin down high price p_h and the probability that low-quality firms charge a high price σ:

$$(p_h, \sigma) = \left(\frac{q_h + q_l + \sqrt{(q_h - q_l)(q_h - q_l - 4k)}}{2}, \frac{\underline{c}\left[q_h - q_l - 2k - \sqrt{(q_h - q_l)(q_h - q_l - 4k)}\right]}{2k1 - \bar{c})} \right)$$

or

$$\left(\frac{q_h + q_l - \sqrt{(q_h - q_l)(q_h - q_l - 4k)}}{2}, \frac{\underline{c}\left[q_h - q_l - 2k + \sqrt{(q_h - q_l)(q_h - q_l - 4s)}\right]}{2k1 - \bar{c})} \right)$$

if $k \leq \frac{q_h - q_l}{4}$.

We will show in the end that these choices of firms are incentive compatible and that without further assumptions on the external parameters, only the first bundle $(\frac{q_h + q_l + \sqrt{(q_h - q_l)(q_h - q_l - 4k)}}{2}, \frac{\underline{c}[q_h - q_l - 2k - \sqrt{(q_h - q_l)(q_h - q_l - 4k)}]}{2k1 - \bar{c})})$ is guaranteed to construct an equilibrium for arbitrarily small k.

In the contract stage, the brand-holder's problem is

$$\begin{aligned} \max_{\{F\}} \quad & (\bar{c} - \underline{c})F \\ \text{s.t.} \quad & q_m - c_i e_m - F \geq max.\{p_h - c_i e_h, q_l - c_i e_l\} \quad \forall \quad c_i \in [\underline{c}, \bar{c}] \end{aligned} \tag{A.4}$$

The constraint can be simplified to $F = -(p_h - q_m) + (e_h - e_m)\underline{c}$ and

$\bar{c} = \frac{p_h - q_l}{e_m - e_l} - \frac{e_h - e_m}{e_m - e_l}\underline{c}$. The objective function can be re-written as a uni-variable function of $\underline{c}$: $(\frac{p_h-q_l}{e_m-e_l} - \frac{e_h-e_l}{e_m-e_l}\underline{c})(-(p_h-q_m)+(e_h-e_m)\underline{c})$. $\underline{c} = \frac{1}{2}[\frac{p_h-q_m}{e_h-e_m} + \frac{p_h-q_l}{e_h-e_l}]$.

Next we check for incentive compatibility for individual firms. To support the equilibrium, the following inequalities must hold:

$$(e_h - e_l)\underline{c} \leq \gamma p_h \leq (e_h - e_l)\bar{c}$$

Since $\gamma = \frac{p_h-q_l}{p_h}$, the inequalities can be simplified to $\frac{p_h-q_l}{e_h-e_l} \in [\underline{c}, \bar{c}]$. Moreover, $\underline{c} = \frac{1}{2}[\frac{p_h-q_m}{e_h-e_m} + \frac{p_h-q_l}{e_h-e_l}]$ and $\bar{c} = \frac{1}{2}[\frac{q_m-q_l}{e_m-e_l} + \frac{p_h-q_l}{e_h-e_l}]$. Therefore, it is sufficient to show that $\frac{p_h-q_l}{e_h-e_l} \in [\frac{p_h-q_m}{e_h-e_m}, \frac{q_m-q_l}{e_m-e_l}]$. It is easy to see that $\frac{p_h-q_l}{e_h-e_l} < \frac{q_h-q_l}{e_h-e_l} < \frac{q_m-q_l}{e_m-e_l}$. Also, $\frac{p_h-q_l}{e_h-e_l} - \frac{p_h-q_m}{e_h-e_m} = \frac{q_h-q_l}{e_h-e_l} - \frac{q_h-p_h}{e_h-e_l} - \frac{q_h-q_m}{e_h-e_m} + \frac{q_h-p_h}{e_h-e_m} > 0$. The incentive constraints are satisfied. Moreover, for p_h to be higher than q_m, it must be that $k < \frac{1}{4[q_h-q_l]}\{[q_h-q_l]^2 - [(q_m-q_l)-(q_h-q_m)]^2\}$ if $q_h - q_m < q_m - q_l$.

Finally, we establish the condition for which the optimal probability of price randomization σ is bounded in $[0,1]$. Recall that $\sigma = \frac{\underline{c}[q_h-q_l-2s-\sqrt{(q_h-q_l)(q_h-q_l-4k)}]}{2k(1-\bar{c})}$ and $\rho = \frac{\underline{c}}{\underline{c}+\sigma(1-\bar{c})}$ is the probability that an individual good with price p_h is of high quality. $\rho q_h + (1-\rho)q_l - p_h = \rho(q_h - p_h) - k = 0$. Then $\rho = \frac{k}{q_h-p_h} = \frac{1+\sqrt{1-\frac{4k}{q_h-q_l}}}{2}$. We can rewrite $\sigma = \frac{\rho}{1-\rho}\frac{\underline{c}}{1-\bar{c}} = \frac{\rho}{1-\rho}\frac{\lambda}{1-\lambda}$, where $\lambda = \frac{\underline{c}}{1-\bar{c}+\underline{c}}$ is the fraction of high quality firms in the individual market. Then $q \in [0,1]$ if $\rho \geq \lambda$.

$\rho \geq \lambda$

$\Leftrightarrow \frac{1+\sqrt{1-\frac{4k}{q_h-q_l}}}{2} \geq \lambda$

$\Leftrightarrow \sqrt{1-\frac{4k}{q_h-q_l}} \geq 2\lambda - 1$.

If $\lambda \leq \frac{1}{2}$, the inequality always holds. If $\lambda > \frac{1}{2}$, the inequality holds if and only if $k < [q_h - q_l]\lambda(1-\lambda)$. Observe that $\frac{d\lambda(1-\lambda)}{d\lambda} = 1 - 2\lambda$, which is negative if $\lambda > \frac{1}{2}$. If λ is bounded above by some $\bar{\lambda} < 1$, $\lambda(1-\lambda)$ is bounded below by $\bar{\lambda}(1-\bar{\lambda})$. Then as long as $k \leq \bar{k} = [q_h - q_l]\bar{\lambda}(1-\bar{\lambda})$, inspection equilibrium exists. Since $\frac{\underline{c}}{1-\bar{c}}$ decreases with k and

$$\lim_{k\to 0}\underline{c}=\frac{1}{2}[\frac{q_h-q_l}{e_h-e_l}+\frac{q_h-q_m}{e_h-e_m}]$$

,

$$\lim_{k\to 0}\bar{c}=\frac{1}{2}[\frac{q_h-q_l}{e_h-e_l}+\frac{q_m-q_l}{e_m-e_l}]$$

,

λ is decreasing in k and bounded above by

$$\bar{\lambda}=\frac{\frac{1}{2}[\frac{q_h-q_l}{e_h-e_l}+\frac{q_h-q_m}{e_h-e_m}]}{1-\frac{1}{2}[\frac{q_h-q_l}{e_h-e_l}+\frac{q_m-q_l}{e_m-e_l}]+\frac{1}{2}[\frac{q_h-q_l}{e_h-e_l}+\frac{q_h-q_m}{e_h-e_m}]}<1$$

Denote $\bar{k}=[q_h-q_l]\bar{\lambda}(1-\bar{\lambda})$ and $\tilde{k}=\frac{1}{4[q_h-q_l]}\{[q_h-q_l]^2-[(q_m-q_l)-(q_h-q_m)]^2\}$, then inspection equilibrium exists if $k\leq k^*$, where

$$k^*=\begin{cases}\bar{k} \quad \text{if} \quad q_h-q_m\geq q_m-q_l\\ min.\{\bar{k},\tilde{k}\} \quad \text{otherwise}\end{cases}$$

A.1.6 Proof of Proposition 1.3

Denote the inspection equilibrium welfare $V_I=\int_0^{\underline{c}}p_h-e_h\cdot cdc+\int_{\underline{c}}^{\bar{c}}q_m-e_m\cdot cdc+\int_{\bar{c}}^1 q_l-e_l\cdot cdc$ and benchmark equilibrium welfare $V_B=\int_0^{\hat{c}}q_m-e_m\cdot cdc+\int_{\hat{c}}^1 q_l-e_l\cdot cdc$. If $\hat{c}>\underline{c}$,

$$\begin{aligned}V_I-V_B&=(p_h-q_m)\underline{c}-\frac{1}{2}(e_h-e_m)\underline{c}^2+(q_m-q_l)(\bar{c}-\hat{c})-\frac{1}{2}(e_m-e_l)(\bar{c}^2-\hat{c}^2)\\&=\frac{\underline{c}}{4}\left[p_h-\frac{1}{2}(e_h-e_m)\frac{1}{2}[\frac{p_h-q_m}{e_h-e_m}+\frac{p_h-q_l}{e_h-e_l}\right]+(q_m-q_l)(\bar{c}-\hat{c})-\frac{1}{2}(e_m-e_l)(\bar{c}^2-\hat{c}^2)\\&=\frac{\underline{c}}{4}\left[2p_h+\frac{e_m-e_l}{e_h-e_l}p_h+q_m+\frac{e_h-e_m}{e_h-e_l}q_l\right]+\frac{\bar{c}-\hat{c}}{4}\left[2(q_m-q_l)-\frac{e_m0e_l}{e_h-e_l}(p_h-q_l)\right]\\&=\Delta V\\&>0\end{aligned}$$

because $\frac{p_h-q_l}{e_h-e_l} < \frac{q_m-q_l}{e_m-e_l}$

If $\hat{c} < \underline{c}$,

$$\begin{aligned} V_I - V_B &= \Delta V + (p_h - q_l)\hat{c} - \frac{1}{2}(e_h - e_l)\hat{c}^2 \\ &= \Delta V + \frac{\hat{c}(e_h - e_l)}{4}\left[4\frac{p_h - q_l}{e_h - e_l} - \frac{q_m - q_l}{e_m - e_l}\right] \end{aligned}$$

Since $\hat{c} < \underline{c}$, $\frac{q_m-q_l}{e_m-e_l} < \frac{p_h-q_l}{e_h-e_l} + \frac{p_h-q_m}{e_h-e_m} < 2\frac{p_h-q_l}{e_h-e_l}$. Therefore, $V_I - V_B > 0$.

A.1.7 Proof of Proposition 1.4

To prove Proposition 1.4, it is helpful to establish that the contract fee is lower in the inspection equilibrium. Denote the contract fee in the inspection equilibrium F_I and that in benchmark equilibrium F_B. $F_I = -(p_h-q_m)+(e_h-e_m)\underline{c} = -\frac{1}{2}(p_h-q_m)+\frac{1}{2}\frac{e_h-e_m}{e_h-e_l}(p_h-q_l)$. It is straight-forward that F_I decreases in p_h. Since p_h is in between q_m and q_h, $F_I < -\frac{1}{2}(q_m-q_m)+\frac{1}{2}\frac{e_h-e_m}{e_h-e_l}(q_m-q_l) < \frac{q_m-q_l}{2} = F_B$.

Since the contract fee is lower in the inspection equilibrium, the firm with component $\underline{c}$ is strictly better-off. If this marginal firm would join the brand in the benchmark equilibrium, the firm benefits from the lowered contract fee. If it does not join the brand in the benchmark case, its profit is $q_l - \underline{c}e_l$, which is strictly lower than $q_h - \underline{c}e_h$. In fact, all firms who have c_i lower than $\underline{c}$ have a higher profit in the benchmark case, the difference $p_h - q_m - (e_h - e_m)c_i$ decreases with c_i. Therefore, the lower the cost component, the more profit gain in the inspection equilibrium.

Moreover, for firms with cost components greater than $\underline{c}$, they either profit from a lower contract fee if their cost components are lower than $\bar{c}$, or they stay indifferent in the stand-alone market with profit $q_l - c_ie_l$.

Therefore, we can claim that firms earn a weakly higher profit in the inspection equilibrium, and the increase in profit is a decreasing function of c_i.

A.1.8 Proof of Proposition 1.5

Recall that $p_h = \frac{q_h+q_l+\sqrt{(q_h-q_l)(q_h-q_l-4k)}}{2}$. $\frac{dp_h}{dk} = -\left(\frac{q_h-q_l}{q_h-q_l-4k}\right)^{\frac{1}{2}} < 0$. Moreover,

$$\lim_{k\to 0} p_h = \lim_{k\to 0} \frac{q_h + q_l + \sqrt{(q_h - q_l)(q_h - q_l - 4k)}}{2} = \frac{q_h + q_l + \sqrt{q_h - q_l}^2}{2} = q_h.$$

A.1.9 Proof of Corollary 1.1 and 1.2

$\underline{c}$ and $\bar{c}$ can be expressed as $\frac{1}{2}[\frac{p_h-q_m}{e_h-e_m} + \frac{p_h-q_l}{e_h-e_l}]$ and $\frac{1}{2}[\frac{q_m-q_l}{e_m-e_l} + \frac{p_h-q_l}{e_h-e_l}]$. Both $\underline{c}$ and $\bar{c}$ are increasing functions of p_h, which is a decreasing function of k. Therefore, $\underline{c}$ and $\bar{c}$ decrease with k. However, the marginal impact of inspection cost k on $\underline{c}$ is greater than that on $\bar{c}$. Therefore, as k decreases, the fraction of firms joining the brand decreases.

$\bar{c} - \underline{c} = \frac{1}{2}[\frac{q_m-q_l}{e_m-e_l} - \frac{p_h-q_m}{e_h-e_m}]$. As k approaches 0, p_h approaches q_h, and $\bar{c} - \underline{c}$ approaches $\frac{1}{2}[\frac{q_m-q_l}{e_m-e_l} - \frac{q_h-q_m}{e_h-e_m}]$. Recall that in the first best case, the fraction of firms investing in medium quality is $\frac{q_m-q_l}{e_m-e_l} - \frac{q_h-q_m}{e_h-e_m}$. Therefore, the fraction of firms investing in medium quality approaches one half the socially optimal mass as k approaches 0.

In terms of the brand-holder's profit, recall that $F = -(p_h - q_m) + (e_h - e_m)\underline{c}$. $\frac{dF}{dp_h} = -1 + (e_h - e_m) \cdot \frac{1}{2}\left[\frac{1}{e_m-e_l} + \frac{1}{e_h-e_l}\right] < 0$. The contract fee F decreases with p_h, which decreases with k. Therefore, as the inspection cost decreases, the contract fee decreases. As we established in Corollary 1.1, $\bar{c}-\underline{c}$ also decreases as k decreases. Therefore, as k decreases, the profit of the brand-holder $(\bar{c} - \underline{c}) \cdot F$ decreases.

A.1.10 Other Equilibria

Lemma 1.3 implies that there may exist one other equilibrium other than the benchmark equilibrium and the inspection equilibrium. Indeed, there is another equilibrium when k is sufficiently low and $\frac{1}{2}\frac{q_m-q_l}{e_m-e_l} > \frac{q_h-q_m}{e_h-e_m}$. In this equilibrium, the branded goods are of either high or medium quality, and the individual goods are always with low quality. Firms who invest in high quality charge a high price. Firms who invest in medium quality randomize between high and medium price. Firms who invest in low quality always charge a low price. Consumers do not inspect if the price is medium or low. If the price is high, consumers inspect with probability γ_1.

Denote $\underline{c}_1$ the fraction of firms choosing high quality, $\hat{c}_1$ the marginal type of firm who is indifferent between accepting and not accepting the contract, σ_1 the probability that a medium-quality good is listed a high price and $\rho_1 = \frac{\underline{c}_1}{\underline{c}_1+\sigma_1(\hat{c}_1-\underline{c}_1)}$ the probability that a high-price good is of high quality.

Likewise, we solve for the equilibrium strategies with backward induction. In the consumption stage, consumers purchase if $p_l \leq q_l$, $p_m \leq q_m$ and $p_h \leq \rho_1 q_h + (1-\rho_1)q_m$ without inspection and $p_h \leq q_h$ with inspection and $q = q_h$.

In the inspection stage, consumers' belief on the equilibrium path can be characterized by

$$\mu(p,b) = \begin{cases} (0,0,1) & \text{if} \quad (p,b) = (p_l,0) \\ (0,1,0) & \text{if} \quad (p,b) = (p_m,1) \\ (\rho_1, 1-\rho_1, 0) & \text{if} \quad (p,b) = (p_h,1) \end{cases}$$

Consumers inspect upon high price if

$$\rho_1(q_h - p) - k \geq \rho_1 q_h + (1-\rho_1)q_m - p$$

Therefore, $k = (1-\rho_1)(p_h - p_m)$. When the inequality holds with equality, consumers inspect with probability $\gamma_1 = \frac{p_h - q_m}{p_h}$ upon observing p_h of a branded firm.

In firms' decision stage, since consumers believe that all individual firms are of low quality, stand-alone firms chooses investment level e_l and price at

$p_l = q_l$. Also, $p_m = q_m$ and $p_h = \frac{\underline{c}_1 q_h + \sigma_1(\hat{c}_1 - \underline{c}_1) q_m}{\underline{c}_1 + \sigma_1(\hat{c}_1 - \underline{c}_1)}$. If $\underline{c}_1$ satisfies

$$-F + p_h - \underline{c}_1 e_h = -F + q_m - \underline{c}_1 e_m,$$

which yields $\underline{c}_1 = \frac{p_h - q_m}{e_h - e_m}$, then investment and pricing decisions are incentive compatible. $p_h = \frac{\underline{c}_1 q_h + \sigma_1(\hat{c}_1 - \underline{c}_1) q_m}{\underline{c}_1 + \sigma_1(\hat{c}_1 - \underline{c}_1)}$ and $k = \frac{\sigma_1(\hat{c}_1 - \underline{c}_1)}{\underline{c}_1 + \sigma_1(\hat{c}_1 - \underline{c}_1)}$ pin down p_h and σ_1: $\sigma_1 = \frac{\underline{c}_1[q_h - q_m - 2s - \sqrt{(q_h - q_m)(q_h - q_m - 4s)}]}{2s(\hat{c}_1 - \underline{c}_1)}$ and $p_h = \frac{q_h + q_m + \sqrt{(q_h - q_m)(q_h - q_m - 4s)}}{2}$ if $k \leq \frac{q_h - q_m}{4}$. Substituting into $\underline{c}_1$, we have $\underline{c}_1 = \frac{q_h + q_m - \sqrt{(q_h - q_m)(q_h - q_m - 4s)}}{2}$.

Finally, it is straight-forward that $\hat{c}_1$ is equal to $\hat{c}$ in the benchmark equilibrium. In the contract stage, for the medium quality firms to randomize between high and medium price, the profits that these two prices yield must be the same. Moreover, consumers believe that individual firms must be of low quality. Therefore, the brand-holder faces the same problem, and therefore chooses the same cut-off $\hat{c}_1 = \hat{c} = \frac{1}{2}\frac{q_m - q_l}{e_m - e_l}$. $\sigma_1 = \frac{\underline{c}_1}{\hat{c}_1 - \underline{c}_1} = \frac{\lambda_1}{1 - \lambda_1}\frac{1 - \rho_1}{\rho_1}$, where $\lambda_1 = \frac{\underline{c}_1}{\hat{c}_1}$ is the fraction of high quality branded firms in the brand. Similar to the condition in the inspection equilibrium, σ_1 is bounded between $[0, 1]$ if $\rho_1 > \lambda_1$. $\rho_1 > \lambda_1$ if $k \leq [q_h - q_m]\lambda_1(1 - \lambda_1)$. Note that $\hat{c}_1$ is a function of exogenous parameters and $\underline{c}_1$ is increasing in p_h, which is a decreasing function of k. Therefore, λ_1 is bounded above by $\bar{\lambda}_1 = \frac{\frac{q_h - q_m}{e_h - e_m}}{\frac{1}{2}\frac{q_m - q_l}{e_m - e_l}}$. The equilibrium exists if $k \leq [q_h - q_m]\bar{\lambda}_1(1 - \bar{\lambda}_1)$.

Appendix B

Appendix to Chapter 2

B.1 Appendix: Figures

Figure B.1: Nutri-score

B.2 Proof of Proposition 2.1

When consumer search is not possible, an equilibrium where all types of firms except the l type adopt the label exists. If all types of firms adopt label, then the consumer forms the belief that the firm is of type h if the label says $\mathcal{L} = 2$, assigns equal probability weight to m_α and m_β when the label says $\mathcal{L} = 1$. The optimal pricing strategies for the firm of different types would be:

- $p_h = \begin{cases} \lambda + 1 + u_0 & \text{if} \quad \lambda \geq \frac{1-\omega}{\omega}(1 + u_0) \\ 1 + u_0 & \text{otherwise.} \end{cases}$

- $p_\alpha = p_\beta = \frac{\lambda+1}{2} + u_0$

- $p_l = u_0$

It is straight-forward to show that in this equilibrium, no type has an incentive to deviate to other labelling strategies. If the h, m_α or m_β doesn't adopt the label, then their profit after deviation will be $\pi' = u_0$, which is less than their equilibrium profits. The l type has no incentive to deviate, since adopting the label sends a signal to the consumer that the product has low qualities on both dimensions.

B.3 Proof of Proposition 2.2

When the consumer is of niche type with probability 1, the h type firm adopts the label and sets the price at $p_h = \lambda + 1 + u_0$, which is the willingness to pay of the niche consumer. If both m_α and m_β types adopt the label, then l type prices at $p_l = u_0$, and doesn't benefit from deviation since having a label which says both attributes are of low quality does not increase profits. Neither m_α nor m_β has the incentive to give up on the label. If they do so, the consumer infers that they are of type l, and the profit following such deviation will be lower than the equilibrium profits.

In fact, this equilibrium is a special case of the mandatory equilibrium with ω being equal to 1. We therefore defer the rest of the strategies to the proof for Proposition 2.4.

B.4 Proof of Proposition 2.3

Given the equilibrium beliefs, we solve for the equilibrium strategies and show that these strategies are incentive compatible. First we set up some notations. Denote the FoP Label $\mathcal{L} \in \{0, 1, 2, \phi\}$ where the numbers represent the number of high-quality attributes and ϕ denotes no label and let the consumer's posterior after observing price signal p and label $\mathcal{L}$ be $\mu((\alpha, \beta)|p, \mathcal{L})$. For simplicity, we sometimes use the following notation instead:

Notation	Conditional Beliefs
μ_h	$\mu((A, B)\|p, \mathcal{L})$
μ_α	$\mu((A, b)\|p, \mathcal{L})$
μ_β	$\mu((a, B)\|p, \mathcal{L})$
μ_l	$\mu((a, b)\|p, \mathcal{L})$

B.4.1 Consumption Stage

In period 5, the consumer makes purchase decisions. If the consumer conducted search in the previous period, she learns the quality of the product. In this case, the consumer purchases if and only if

$$U_i(\alpha, \beta) \geq p$$

where $i \in \{N, M\}$. If the consumer did not search, then the she purchases if and only if the expected utility derived from the product is at least as high as the price:

$$\Sigma_{\alpha \in \{A,a\}} \Sigma_{\beta \in \{B,b\}} \ \mu((\alpha, \beta)|p, s) U_i(\alpha, \beta) \geq p$$

B.4.2 Search Stage

In the search stage in period 4, the consumer's beliefs on the equilibrium path is summarized in Table B.1 The consumer only conducts search when

$(p, \mathcal{L})$	μ_h	μ_α	μ_β	μ_l
$(p_h, 2)$	1	0	0	0
$(p_\beta, 1)$	0	0	1	0
(p_α, ϕ)	0	$\frac{1}{1+\sigma}$	0	$\frac{\sigma}{1+\sigma}$
(p_l, ϕ)	0	0	0	1

Table B.1: Posterior Beliefs on the Equilibrium Path

the signals are $(p, \mathcal{L}) = (p_\alpha, \phi)$. It is clear to see that in all other cases, the posterior beliefs assign probability 1 to a certain type. Therefore, the consumer has no incentive to search.

Next we show that the mass consumer never searches in equilibrium. Recall that the consumer's type-dependent utility functions are

$$\begin{aligned} U_N &= \lambda \mathbb{1}(\alpha = A) + \qquad && \mathbb{1}(\beta = B) + u_0 \\ U_M &= && \mathbb{1}(\beta = B) + u_0 \end{aligned}$$

When the signals $(p, \mathcal{L}) = (p_\alpha, \phi)$, the product can be either (A, b) or (a, b). The mass consumer derives the same utility level from these two products. Therefore, she has no incentive to search. The niche consumer, on the other hand, searches if

$$\frac{1}{1+\sigma}(\lambda + u_0 - p_\alpha) - c \geq \frac{1}{1+\sigma}(\lambda + u_0) + \frac{\sigma}{1+\sigma}u_0 - p_\alpha.$$

The left-hand-side is the expected utility gain of the niche consumer if she searches, which is the probability that the product is (A, b) and she purchases at the listed price p_α minus the search cost c. The right-hand-side is the expected utility gain if she buys without searching. In this case, the expected utility gain from the product is $\frac{1}{1+\sigma}(\lambda + u_0) + \frac{\sigma}{1+\sigma}u_0$, and she pays p_α with certainty.

We will see the next stage that the niche consumer is indifferent between search and no search, so one best response of hers is to search with probability $\gamma = \frac{\omega p_\alpha - u_0}{\omega p_\alpha}$.

B.4.3 Pricing Stage

We now solve the optimal pricing strategies of the firm.

If the firm is of high type ($\theta = h$), it expects that with the institutional signal, it type is revealed. If it sets the price at $1 + u_0$, the consumer will purchase the product regardless of her type. If it sets a price higher than this level, only the niche consumer will purchase the product. Therefore, it sets the price at the niche consumer's willingness to pay if this level is sufficiently high, or equivalently, if the fraction of the mass consumer is sufficiently low. Otherwise,it sets the price at the mass consumer's willingness to pay so that the consumer purchases with certainty. To sum up the high type firm's pricing strategy, we have:

$$p_h = \begin{cases} \lambda + 1 + u_0 & \text{if } \lambda \geq \frac{1-\omega}{\omega}(1+u_0) \\ 1 + u_0 & \text{otherwise.} \end{cases}$$

Similarly, the m_α type firm expects its type to be fully communicated. Regardless of the type, the consumer derives utility level of $1 + u_0$ from the product (a, B). Therefore, the m_α type firm's optimal pricing strategy is

$$p_\beta = 1 + u_0.$$

The m_α type firm is pooled with the low type firm in equilibrium. If the consumer searches with certainty, the low type firm does not have the incentive to mimic the m_α type in the price signal. If the consumer does not search in this case, then the low type firm mimics with certainty. Therefore, the only equilibrium in this subgame when the signals are (p_α, ϕ) is for the consumer to randomize search and the low type firm to randomize the price it sets.

The low-type firm only randomizes between two prices: p_α, the price the m_α type firm sets, and $p_l = u_0$, the consumer's willingness to pay for its product. If the low-type firm sets any other price $p_l > u_0$, no transaction will take place due to the D1 criterion refinement. For any price deviation $p_l' \in (u_0, p_\alpha)$, the mass consumer certainly won't buy. Therefore, regardless of the response of the niche consumer, the m_α type firm cannot benefit from this deviation. D1 criterion requires that off-equilibrium beliefs for these deviations should assign probability 1 to type l. Moreover, the off-equilibrium beliefs assigning probability 1 on type l for price deviations $p' > p_\alpha$ also survives D1 criterion refinement. If $p_l < u_0$, the low type firm can increase its profit by setting a higher price, and the consumer will still buy the product with certainty.

On the other hand, expecting the niche consumer to randomize search, the m_α type firm sets a price as high as possible, which leaves the niche consumer with zero surplus. Combining the niche consumer's indifference condition and the m_α type firm's incentive to leave her zero surplus, we have

$$\frac{1}{1+\sigma}(\lambda + u_0 - p_\alpha) - c = \frac{1}{1+\sigma}(\lambda + u_0) + \frac{\sigma}{1+\sigma}u_0 - p_\alpha = 0.$$

Solving for p_α and σ, we have [1]

$$\sigma = \frac{\lambda - 2c - \sqrt{\lambda(\lambda - 4c)}}{2c}$$
$$p_\alpha = u_0 + \frac{\lambda + \sqrt{\lambda(\lambda - 4c)}}{2}$$

[1] There is another set of solution with $\sigma = \frac{\lambda - 2c + \sqrt{\lambda(\lambda-4c)}}{2c} > 1$, but in our setting, σ is the probability that the low type firm sets its price at p_α, which should be in between 0 and 1.

B.4.4 Signaling Stage

As we explained in the previous stage, one off-equilibrium belief which satisfies the D1 criterion refinement is to assign probability 1 to all deviations $(p, \mathcal{L}) = (p', \phi)$ for any $p' \notin \{p_\alpha, u_0\}$. With such off-equilibrium belief, regardless of the type, the firm finds no incentive to deviate to any of these strategies. The only potential concern is for a type to deviate to the equilibrium strategies of other types. Therefore, we check for incentive compatibility for all types of firm and show that they cannot benefit from any deviation. First, the m_α type does not want to mimic the m_β type because its equilibrium profit ωp_α is greater than that of the m_β type, $1 + u_0$. To see this,

$$\begin{aligned}\omega p_\alpha &\geq \bar{\omega} p_\alpha \\ &= \frac{2 + 2u_0}{\lambda + 2u_0 + \sqrt{\lambda(\lambda - 4c)}}(u_0 + \frac{\lambda + \sqrt{\lambda(\lambda - 4c)}}{2}) \\ &= 1 + u_0\end{aligned}$$

The m_β type firm does not deviate and mimic the m_α type firm because $p_\alpha > 1 + u_0$, so the consumer will not purchase a m_α product after search. The m_α firm's expected profit if it mimics the m_α firm is therefore $\omega(1-\gamma)p_\alpha$, which is equal to the low type firm's equilibrium profit u_0.

Finally, it is straight-forward to see that $\gamma = \frac{\omega p_\alpha - u_0}{\omega p_\alpha}$ is in between 0 and 1 since $\omega p_\alpha > 1 + u_0 > u_0$.

In our setting, u_0 can also be normalized to 0. Then there exist equilibria where consumers search with certainty, and the low type firm randomizes prices between p_α and 0, both giving it an expected profit of 0.

B.5 Proof of Proposition 2.4

When label adoption is mandated, the strategies of the h and l types are simple since information about their product qualities are fully transmitted to the consumer. In the parameter range specified in Proposition 2.4, since ω is sufficiently high, the h type firm prices at $p_h = \lambda + 1 + u_0$. To see this, denote $\pi_h(p)$ the profit of h type if it sets the price at p. Then

$$\begin{aligned}
\pi_h(\lambda + 1 + u_0) &= \omega(\lambda + 1 + u_0) \\
&\geq \bar{\omega}^M(\lambda + 1 + u_0) \\
&= \frac{1 + u_0}{1 + u_0 + \frac{\lambda-1+\sqrt{(\lambda-1)(\lambda-1-4c)}}{2}}(\lambda + 1 + u_0) \\
&= (1 + u_0)\frac{\lambda + 1 + u_0}{1 + u_0 + \frac{\lambda-1+\sqrt{(\lambda-1)(\lambda-1-4c)}}{2}} \\
&\geq 1 + u_0 \\
&= \pi_h(1 + u_0)
\end{aligned}$$

since $\lambda + 1 + u_0 - (1 + u_0 + \frac{\lambda-1+\sqrt{(\lambda-1)(\lambda-1-4c)}}{2}) = \frac{\lambda+1+\sqrt{(\lambda-1)(\lambda-1-4c)}}{2} > 0$. The mass consumer leaves the market upon observing $(p, \mathcal{L}) = (p_h, 2)$. It is a best response for the niche consumer to buy without search. In fact, since the label says $\mathcal{L} = 2$, neither type of the consumer has the incentive to search. The search incentive is the same for l type firm. The l type firm can only price at $p_l = u_0$, which is the willingness to pay the consumer has for its product.

We now show that the strategies for the m_α and m_β types are also incentive compatible. In the search stage, when the consumer observes $(p, \mathcal{L}) = (1 + u_0, 1)$, she assigns probability 1 to type m_β. Therefore, she has no incentive to search. When the label reads $\mathcal{L} = 1$ and the price is p_α^M, the consumer forms the belief that the the firm is of type m_α with probability $\frac{1}{1+\sigma^M}$ and type m_β with probability $\frac{\sigma^M}{1+\sigma^M}$. As we will see later, p_α^M is greater than

$1 + u_0$. Therefore, the mass consumer leaves the market in this case. Just as in Proposition 2.3, the niche consumer will be indifferent between buying directly and searching, so it is also a best response for her to randomize search with probability $\gamma^M = 1 - \frac{1+u_0}{\omega p_\alpha^M}$. This probability of searching makes the m_β type firm indifferent between setting the price at $1 + u_0$ and p_α^M.

Similar to the pricing strategies in the voluntary equilibrium, here m_α and m_β types' pricing strategies have to satisfy

$$\frac{1}{1+\sigma^M}(\lambda + u_0 - p_\alpha^M) - c = \frac{1}{1+\sigma^M}(\lambda + u_0) + \frac{\sigma^M}{1+\sigma^M}(1 + u_0) - p_\alpha^M = 0.$$

where p_α^M is the price that the m_α type firm sets and σ^M the probability that the m_β type firm sets the price at p_α^M in the mandatory equilibrium.

Solving for these equations, we obtain

$$\sigma^M = \frac{\lambda - 1 - 2c - \sqrt{(\lambda-1)(\lambda-1-4c)}}{2c}$$
$$p_\alpha^M = 1 + u_0 + \frac{\lambda - 1 + \sqrt{(\lambda-1)(\lambda-1-4c)}}{2}$$

To check for incentive compatibility, first, it is straight-forward that the h and l types have optimal strategies. Since the label is mandatory, they cannot mimic other types, and they receive all the gains from trade. The m_α type firm does not want to mimic the m_β type firm because its equilibrium profit $\pi_\alpha^* = \omega p_\alpha^M \geq \bar{\omega}^M p_\alpha^M = 1+u_0$. The m_β type firm is indifferent between the two prices p_α^M and $1 + u_0$ because the niche consumer searches with probability γ^M.

To summarize, in the mandatory equilibrium, the h type firm sets the price at $p_h = \lambda + 1 + u_0$, and the l type firm sets the price at u_0. The m_α type sets the price at $p_\alpha^M = 1+u_0+\frac{\lambda-1+\sqrt{(\lambda-1)(\lambda-1-4c)}}{2}$, and the m_β type firm randomizes between p_α^M and $1 + u_0$. In equilibrium, the niche consumer will buy from the h and l type firm or when the label says $\mathcal{L} = 1$ and the price is $1 + u_0$. When the label says $\mathcal{L} = 1$ and the price is p_α^M, the niche consumer will search with probability $\gamma^M = 1 - \frac{1+u_0}{\omega p_\alpha^M}$ and buys if the product is (A, b). The mass consumer buys if the label says $\mathcal{L} = 1$ and the price is $1 + u_0$, or if the price is u_0. In all other cases, the mass consumer exits the market.

B.6 Other Equilibria

In this section we discuss all other possible equilibria in this game. First of all, adopting the label is a dominant strategy for the type h firm. Therefore, in terms of label adoption strategy, we only need to discuss $2^3 = 8$ different cases. Table B.2 summarizes these cases:

Adopt	Don't Adopt	Note
m_α, m_β, l		Unravelling
m_α, m_β	l	Unravelling
m_α, l	m_β	l deviates
m_β, l	m_α	l deviates
m_α	m_β, l	m_β deviates
m_β	m_α, l	Voluntary Equilibrium
l	m_α, m_β	l deviates
	m_α, m_β, l	

Table B.2: Possible Equilibria

In fact, we can rule out 5 out of 8 cases. Going through these cases, we show that in the parameter range where the voluntary equilibrium exists, there is no other equilibrium that survives D1 criterion refinement.

First, in the unravelling cases, the pricing strategies of type h and l are straight-forward. Type h sets the price $p_h = \lambda + 1 + u_0$ and type l $p_l = u_0$. An equilibrium where type m_α and m_β separate with prices does not exist, since in general the profits of the two types do not coincide. The type who earns a lower profit will have an incentive to mimic the pricing strategy of the other type. Since the consumer does not search on the equilibrium path, this is a valid deviation.

In the parameter range of our discussion, a semi-separating equilibrium corresponds to the mandatory equilibrium.

We show in the following that an equilibrium where m_α and m_β pool on the same price does not survive D1-criterion refinement. Upon observing a label that shows signal $s = 1$, the niche consumer searches if

$$\begin{aligned}
&\mu_\alpha(\lambda + u_0 - p) - c \geq \mu_\alpha + (1 - \mu_\alpha)(1 + u_0) - p \qquad \text{and} \\
&\mu_\alpha(\lambda + u_0 - p) - c \geq 0
\end{aligned}$$

We can further simplify the conditions to the following:

$$\begin{aligned}
p &\in [1 + u_0 + \frac{c}{1 - \mu_\alpha}, \lambda + u_0 - \frac{c}{\mu_\alpha}] \\
\mu_\alpha &\in [\frac{1 - \sqrt{1 - \frac{4c}{\lambda - 1}}}{2}, \frac{1 + \sqrt{1 - \frac{4c}{\lambda - 1}}}{2}] \\
c &< \frac{\lambda - 1}{4}
\end{aligned}$$

Likewise, we can simplify the mass consumer's search conditions to

$$\begin{aligned}
p &\in [u_0 + \frac{c}{\mu_\alpha}, 1 + u_0 - \frac{c}{1 - \mu_\alpha}] \\
\mu_\alpha &\in [\frac{1 - \sqrt{1 - 4c}}{2}, \frac{1 + \sqrt{1 - 4c}}{2}] \\
c &< \frac{1}{4}
\end{aligned}$$

If the m_α and m_β firm pool on the same price p^*, when the consumer sees that price, she should expect that the probabilities of the firm being type m_α and m_β are both $\frac{1}{2}$. Fixing $\mu_\alpha = \frac{1}{2}$, we can derive that the niche consumer searches if and only if $p^* \in [1 + u_0 + 2c, \lambda + u_0 + 2c]$ from the search conditions above.

Therefore, $p^* \leq 1+u_0+2c$, for any price above this level will yield equilibrium profit of 0 for the m_β type. Similarly, the mass consumer searches if $p^* \in [u_0+2c, 1+u_0-2c]$. Therefore, p^* cannot be in the range $(u_0+2c, 1+u_0-2c)$. If p^* is indeed in that range, $\pi^*_\alpha = \omega p^*$ because the mass consumer will search and will not buy from the m_α type firm. The m_α type can deviate to a price $p' = 1+u_0-2c$ and increase its profit. p^* cannot be in between $(1+u_0-2c$ and $1+u_0$ either. If this is the case, the mass consumer exits the market and the niche consumer buys with certainty. Then both types of firm can deviate to the price $1+u_0$, which the niche consumer still accepts.

We now discuss all other possible pooling equilibria. If $p^* \in [u_0, u_0+2c]$, the conditions for p^* to be an equilibrium are the following:

- ω is sufficiently low such that $\omega(1+u_0) < p^*$
- For any deviation $p' > p^*$,
  - i if $p' \in (1+u_0, u_0 + \frac{\lambda-1+\overline{\lambda-1-4c}}{2}) \cup (\lambda+u_0 - \frac{\lambda-1-\overline{\lambda-1-4c}}{2}, \lambda+u_0)$,
    $\mu_\alpha(p') < \frac{p'-1+u_0}{\lambda-1}$
  - ii if $p' \in [u_0 + \frac{\lambda-1+\overline{\lambda-1-4c}}{2}, \lambda+u_0 - \frac{\lambda-1-\overline{\lambda-1-4c}}{2}]$,
    $\mu_\alpha(p') < \frac{c}{\lambda+u_0-p'}$.
- For any deviation $p' < p^*$ and $p' > \omega p^*$,
  - i if $p' \in [u_0 + \frac{1+\sqrt{1-4c}}{2}, 1+u_0) \cup (u_0, u_0 + \frac{1-\sqrt{1-4c}}{2})$,
    $\mu_\alpha(p') > 1+u_0-p'$
  - ii if $p' \in [u_0 + \frac{1-\sqrt{1-4c}}{2}, u_0 + \frac{1+\sqrt{1-4c}}{2}]$,
    $\mu_\alpha(p') > 1 - \frac{c}{1+u_0-p'}$.

In the range of deviation $[u_0 + \frac{\lambda-1+\overline{\lambda-1-4c}}{2}, \lambda+u_0 - \frac{\lambda-1-\overline{\lambda-1-4c}}{2}]$, compared to the m_β type firm, the m_α type firm benefits for a larger set of best responses of the niche consumer. Therefore, D1-criterion requires that the off-equilibrium beliefs in this range should assign probability 1 to the m_α type. There is a continuum of pooling equilibria $p^* \in [u_0, u_0 + 2c]$ if ω is sufficiently small, but none survives D1-criterion refinement. Similar contradiction can also be found in the range where the mass consumer searches, which we omit here. The other possible range of pooling prices is $p^* \in [1 + u_0, 1 + u_0 + 2c]$. The violation for D1 criterion refinement is similar. In order to support the equilibrium, the off-equilibrium beliefs are required to have sufficiently low weight on m_α type when the niche consumer may want to conduct search. However, D1-criterion requires that in this range, the off-equilibrium belief should have all the probability weight on the m_α type.

Finally, we discuss the last case where only type h adopts the label. An equilibrium where m_α and m_β pool on the same price does not exist for similar reasons as previously shown. For sufficiently small μ_l, there are similar best response conditions with different thresholds. The range of deviation which may induce the consumer to search is therefore non-empty. As long as this is the case, D1-criterion requires to put probability 1 on the type which can still profit after the consumer searches. Therefore, no equilibrium where m_α and m_β type pool on the same price survives D1-criterion refinement. There is no semi-separating equilibrium either, since type l has an incentive to mimic every price, and in general the conditions for such an equilibrium to hold are not satisfied. Likewise, an equilibrium where m_α and m_β type set different prices does not exist. Since type l will have an incentive to mimic both types, m_β will deviate to a slightly higher price than its equilibrium price.

We have examined all possible equilibria and shown that in the parameter range where the voluntary equilibrium exists, the only other equilibrium which survives D1-criterion is the mandatory equilibrium.

B.7 Comparision between the Mandatory and Voluntary Equilibrium

B.7.1 Proof of Lemma 2.1

In the mandatory equilibrium, the m_β type firm sets the price at p_α^M with probability $\sigma^M = \frac{\lambda-1-2c-\sqrt{(\lambda-1)(\lambda-1-4c)}}{2c}$, whereas the l type firm mimicks the m_α type firm with probability $\sigma = \frac{\lambda-2c-\sqrt{\lambda(\lambda-4c)}}{2c}$. Since the *exante* probability of firms types are equal, we only need to show that $\sigma^M > \sigma$.

$$\begin{aligned}
\sigma^M > \sigma &\Leftrightarrow \sqrt{\lambda(\lambda-4c)} > \sqrt{(\lambda-1)(\lambda-1-4c)} + 1 \\
&\Leftrightarrow \lambda^2 - 4c\lambda > 1 + 2\sqrt{(\lambda-1)(\lambda-1-4c)} + \lambda^2 - \lambda - 4c\lambda - \lambda + 1 + 4c \\
&\Leftrightarrow \sqrt{(\lambda-1)(\lambda-1-4c)} < \lambda - 1 - 2c \\
&\Leftrightarrow 4c^2 > 0
\end{aligned}$$

Since consumer's search cost c is positive, $\sigma^M > \sigma$ always holds. Lemma 2.1 is proven.

We now show Corollary 2.1.

$$\begin{aligned}
p_\alpha^M < p_\alpha &\Leftrightarrow 1 + u_0 + \frac{\lambda - 1 + \sqrt{(\lambda-1)(\lambda-1-4c)}}{2} < u_0 + \frac{\lambda + \sqrt{\lambda(\lambda-4c)}}{2} \\
&\Leftrightarrow \sqrt{(\lambda-1)(\lambda-1-4c)} + 1 < \sqrt{\lambda(\lambda-4c)}
\end{aligned}$$

The inequality in the last line is shown to hold in Lemma 2.1. Therefore, we can conclude that the m_α type firm sets a lower price in the mandatory equilibrium.

B.7.2 Proof of Proposition 2.5

The first part of Proposition 2.5 states that $\gamma^M < \gamma$, but the expected search cost of the consumer is higher in the mandatory equilibrium.

First, we show that $\gamma^M < \gamma$. Note that γ^M is such that the m_β type firm is indifferent between setting the price at p_α^M and $1+u_0$. Therefore, γ^M satisfies $(1-\gamma^M)p_\alpha^M = 1+u_0$. Similarly, γ satisfies $(1-\gamma)p_\alpha = u_0$.

$$\begin{aligned}\gamma^M < \gamma &\Leftrightarrow 1 - \frac{1+u_0}{\omega p_\alpha^M} < 1 - \frac{u_0}{\omega p_\alpha} \\ &\Leftrightarrow \frac{1+u_0}{\omega p_\alpha^M} > \frac{u_0}{\omega p_\alpha}\end{aligned}$$

Since we know from Corollary 2.1 that $p_\alpha^M < p_\alpha$ and $1+u_0 > u_0$, it is straight-forward to see that $\frac{1+u_0}{\omega p_\alpha^M} > \frac{u_0}{\omega p_\alpha}$. Therefore, $\gamma^M < \gamma$. Conditioning on the niche consumer being indifferent between search and buying directly, the probability that the niche consumer searches is lower in the mandatory equilibrium than in the voluntary equilibrium.

We now show the second part of Proposition 2.5. In expectation, the consumer searches more in the mandatory equilibrium than in the voluntary equilibrium. In the voluntary equilibrium, the expected probability that the consumer will encounter the signal bundle which induces search is $\frac{1}{4}(1+\sigma)$. The m_α type firm always produces such a signal bundle, and the l type firm does so with probability σ. Therefore, the probability that the niche consumer searches is $\frac{1}{4}(1+\sigma)\gamma$. Likewise, for the niche consumer, the expected probability of search in the mandatory equilibrium is $\frac{1}{4}(1+\sigma^M)\gamma^M$. In the following, we will show that $(1+\sigma^M)\gamma^M > (1+\sigma)\gamma$, which completes our proposition. First, note that σ satisfies $\frac{1}{1+\sigma}(\lambda+u_0) + \frac{\sigma}{1+\sigma}u_0 = p_\alpha$. After rearrangement, we have $(1+\sigma) =$. Similarly, $(1+\sigma^M) = \frac{\lambda-1}{p_\alpha^M-(1+u_0)}$. Therefore,

$$
\begin{aligned}
(1+\sigma^M)\gamma^M > (1+\sigma)\gamma \Leftrightarrow & \frac{\lambda-1}{p_\alpha^M-(1+u_0)}\frac{\omega p_\alpha^M-(1+u_0)}{\omega p_\alpha^M} > \frac{\lambda}{p_\alpha-u_0}\frac{\omega p_\alpha-u_0}{\omega p_\alpha} \\
\Leftrightarrow & \frac{\lambda-1}{\frac{\lambda-1+\sqrt{(\lambda-1)(\lambda-1-4c)}}{2}}\frac{\omega[1+u_0+\frac{\lambda-1+\sqrt{(\lambda-1)(\lambda-1-4c)}}{2}]-(1+u_0)}{\omega[1+u_0+\frac{\lambda-1+\sqrt{(\lambda-1)(\lambda-1-4c)}}{2}]} \\
& > \frac{\lambda}{\frac{\lambda+\sqrt{\lambda(\lambda-4c)}}{2}}\frac{\omega[u_0+\frac{\lambda+\sqrt{\lambda(\lambda-4c)}}{2}]-u_0}{\omega[u_0+\frac{\lambda+\sqrt{\lambda(\lambda-4c)}}{2}]}
\end{aligned}
$$

To simplify the expression, denote $x = \lambda - 1 + \sqrt{(\lambda-1)(\lambda-1-4c)}$ and $y = \lambda + \sqrt{\lambda(\lambda-4c)}$. Then the inequality can be written as

$$
\begin{aligned}
(1+\sigma^M)\gamma^M > (1+\sigma)\gamma \Leftrightarrow & \frac{\lambda-1}{x}\frac{\omega x-2(1-\omega)(1+u_0)}{2(1+u_0)+x} > \frac{\lambda}{y}\frac{\omega y-2(1-\omega)u_0}{y+2u_0} \\
\Leftrightarrow & \frac{\lambda-1}{x+2(1+u_0)} - \frac{(1-\omega)(\lambda-1)}{x} > \frac{\lambda}{y+2u_0} - \frac{(1-\omega)\lambda}{y} \\
\Leftrightarrow & \frac{\omega(\lambda-1)}{x(x+2(1+u_0))} > \frac{\omega\lambda}{y(y+2u_0)} \\
\Leftrightarrow & y(y+2u_0)(\lambda-1) > x(x+2(1+u_0))\lambda \\
\Leftrightarrow & (\lambda-1)[\lambda^2-2c\lambda+u_0\lambda+(\lambda+u_0)\sqrt{\lambda(\lambda-4c)}] \\
& > \lambda[\lambda^2-\lambda-2c\lambda+u_0\lambda+(\lambda+u_0)\sqrt{(\lambda-1)(\lambda-1-4c)}+2c-u_0] \\
\Leftrightarrow & (\lambda-1)\sqrt{\lambda(\lambda-4c)} > \lambda\sqrt{(\lambda-1)(\lambda-1-4c)} \\
\Leftrightarrow & (\lambda-1)^2(\lambda^2-4c\lambda) > \lambda^2[(\lambda-1^2)-4c\lambda+4c] \\
\Leftrightarrow & \lambda^2-(\lambda-1)^2-\lambda > 0 \\
\Leftrightarrow & \lambda-1 > 0
\end{aligned}
$$

Since it is assumed that $\lambda > 1$, we can conclude that the expected search cost in the mandatory equilibrium is higher than that in the voluntary equilibrium.

B.7.3 Proof of Proposition 2.6

Since there is no competition, total welfare is simply the producer's surplus. Denote the welfare in the voluntary equilibrium

$$V = \frac{1}{4}\omega(\lambda + 1 + u_0) + \frac{1}{4}\omega p_\alpha + \frac{1}{4}(1 + u_0) + \frac{1}{4}[\sigma\omega(1-\gamma)p_\alpha + (1-\sigma)u_0]$$

and that in the mandatory equilibrium

$$V^M = \frac{1}{4}\omega(\lambda+1+u_0) + \frac{1}{4}\omega p_\alpha^M + \frac{1}{4}[\sigma^M\omega(1-\gamma^M)p_\alpha^M + (1-\sigma^M)(1+u_0)] + \frac{1}{4}u_0.$$

It can be shown that the difference of total welfare between the mandatory and the voluntary equilibrium, $V^M - V = \frac{1}{4}(p_\alpha^M - p_\alpha)$ is negative, since we have established in Corollary 2.1 that $p_\alpha^M < p_\alpha$.

Appendix C

Appendix to Chapter 3

C.1 Basics for Experiment 2

C.1.1 Summary Statistics for Different Pies and Treatments

In this section, we report the summary statistics for different pies and treatments in Supplementary Table C.1. From this table, we can see that no matter in which treatment the probability of disagreements decreases with pie size. Also, conditional on reaching a deal, players' payoffs increase with pie size. This pattern can be observed in Supplementary Figure C.1 and C.2 as well.

In addition, surplus loss due to disagreements and the value of private information generally increase with pie size. Similar empirical patterns occur in all treatments. Therefore, we confirm that the first result in Camerer et al. [2019] is robust from different treatments.

Result C.1. *Deal rates and payoffs are increasing with the pie size.*

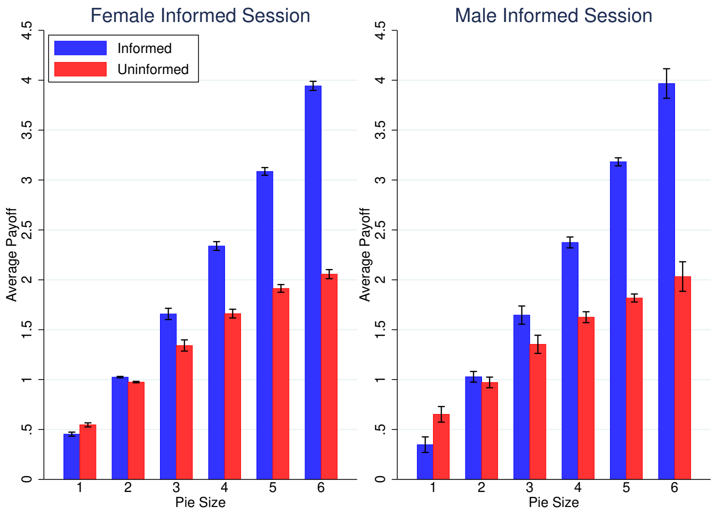

Figure C.1: Mean payoffs by pie size and subject type, rounds ending in a deal. The blue bars represent the average payoff of informed players and the red bars are for uninformed players. The standard errors (overlaid on the bars) are calculated at the session level.

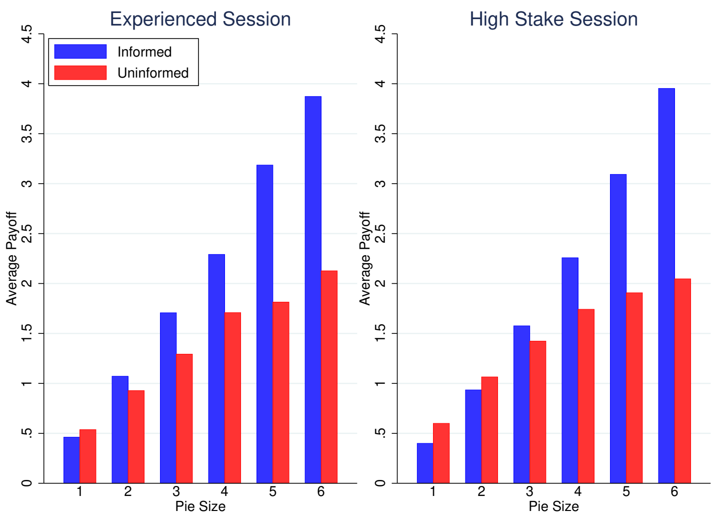

Figure C.2: Mean payoffs by pie size and subject type, rounds ending in a deal. The blue bars represent the average payoff of informed players and the red bars are for uninformed players. Since there is only one session for each treatment, the standard errors can not be computed.

Table C.1: Summary Statistics for Different Pies and Treatments

Treatment	Pie Size	1	2	3	4	5	6	Mean
Baseline Treatment	Informed Payoff[a]	0.37	0.95	1.56	2.23	3.07	3.87	2.01
		(0.03)	(0.04)	(0.04)	(0.03)	(0.05)	(0.06)	
	Uninformed Payoff[a]	0.63	1.05	1.44	1.77	1.93	2.13	1.49
		(0.03)	(0.04)	(0.04)	(0.03)	(0.05)	(0.06)	
	Deal Rate	0.42	0.48	0.54	0.69	0.73	0.81	0.61
		(0.06)	(0.05)	(0.03)	(0.02)	(0.02)	(0.02)	
	Surplus Loss[b]	0.58	1.04	1.39	1.25	1.36	1.16	1.13
		(0.06)	(0.10)	(0.10)	(0.10)	(0.10)	(0.11)	
	Information Value[c]	-0.11	-0.05	0.05	0.31	0.83	1.39	0.40
		(0.03)	(0.03)	(0.04)	(0.04)	(0.07)	(0.10)	
Female Informed	Informed Payoff[a]	0.45	1.02	1.66	2.34	3.09	3.94	2.08
		(0.02)	(0.01)	(0.06)	(0.04)	(0.04)	(0.05)	
	Uninformed Payoff[a]	0.55	0.98	1.34	1.66	1.91	2.06	1.42
		(0.02)	(0.01)	(0.06)	(0.04)	(0.04)	(0.05)	
	Deal Rate	0.54	0.56	0.65	0.71	0.75	0.79	0.66
		(0.02)	(0.03)	(0.03)	(0.03)	(0.04)	(0.03)	
	Surplus Loss[b]	0.46	0.89	1.06	1.18	1.26	1.27	1.02
		(0.02)	(0.06)	(0.09)	(0.14)	(0.18)	(0.18)	
	Information Value[c]	-0.05	0.03	0.21	0.48	0.87	1.49	0.51
		(0.02)	(0.01)	(0.08)	(0.07)	(0.06)	(0.13)	
Male Informed	Informed Payoff[a]	0.35	1.03	1.65	2.37	3.18	3.97	2.09
		(0.08)	(0.05)	(0.09)	(0.05)	(0.04)	(0.15)	
	Uninformed Payoff[a]	0.65	0.97	1.35	1.63	1.82	2.03	1.41
		(0.08)	(0.05)	(0.09)	(0.05)	(0.04)	(0.15)	
	Deal Rate	0.47	0.51	0.62	0.66	0.70	0.74	0.62
		(0.03)	(0.02)	(0.06)	(0.02)	(0.04)	(0.03)	
	Surplus Loss[b]	0.53	0.98	1.13	1.35	1.48	1.57	1.18
		(0.03)	(0.04)	(0.17)	(0.07)	(0.19)	(0.15)	
	Information Value[c]	-0.13	0.03	0.17	0.49	0.96	1.41	0.49
		(0.06)	(0.05)	(0.09)	(0.06)	(0.08)	(0.18)	
Experienced	Informed Payoff[a]	0.46	1.07	1.71	2.29	3.19	3.87	2.10
	Uninformed Payoff[a]	0.54	0.93	1.29	1.71	1.81	2.13	1.40
	Deal Rate	0.44	0.54	0.63	0.72	0.81	0.81	0.66
	Surplus Loss[b]	0.56	0.91	1.10	1.10	0.93	1.17	0.96
	Information Value[c]	-0.03	0.08	0.26	0.42	1.12	1.41	0.54
High Stake	Informed Payoff[a]	0.40	0.93	1.58	2.26	3.09	3.95	2.04
	Uninformed Payoff[a]	0.60	1.07	1.42	1.74	1.91	2.05	1.46
	Deal Rate	0.56	0.62	0.59	0.63	0.73	0.77	0.65
	Surplus Loss[b]	0.44	0.77	1.24	1.47	1.37	1.40	1.11
	Information Value[c]	-0.11	-0.08	0.09	0.33	0.86	1.46	0.42

Means and standard errors (shown in parentheses) are calculated by treating each session's mean as a single observation. Since there is one session for experienced and high-stake treatment, the standard errors for these two treatments are not computable. [a] Averages are calculated for deal games only. [b] Surplus loss = the mean expected loss of pie due to strikes. [c] Information value = the mean difference between the informed and uninformed payoffs.

C.1.2 Focal Points

Supplementary Figure C.3 to Supplementary Figure C.6 show distributions of uninformed player's payoff conditional on reaching a deal in different treatments. See Camerer et al. [2019] for the distribution of the baseline. Focusing on the female-informed and male-informed sessions, as we pool the data from all pairs that reached a deal, we can observe 84.9% of the payoffs are 0.5, 1, 1.5, 2, 2.5 or 3, matching values that are exactly halves of possible pie sizes.

In addition, equal-split is the most common outcome. Among female-informed and male-informed treatments, 49.45% of the outcomes are equal-splitting when pie size is small or medium ($\pi \leq \$4$). This is even more common in female-informed sessions (55.98%) than male-informed sessions (42.94%). Notice that when pie size is small, the predicted payoffs of the efficient and equal-split equilibrium are the same (the top two rows in Supplementary Figure C.3 to Supplementary Figure C.6).

When pie size is large ($\pi \geq \$5$), the efficient equilibrium and equal-split equilibrium predict differently. The bottom row in Figure C.3 to Figure C.6 show that the mode is at \$2 and there are second modes at half the pie. Pooling the data from female-informed and male-informed treatments, 25.59% of payoffs are \$2. When pie size is \$5, 18.9% of the payoffs are at the equal-split equilibrium, while 13.0% of the payoffs are equal-splits when pie size is \$6. This empirical pattern is observed in all treatments and coincides with the finding in Camerer et al. [2019]. Thus, the second result in Camerer et al. [2019] is also replicated in Experiment 2.

Result C.2. *In all treatments, when pie size is small or medium ($\pi \leq \$4$), the modes of the uninformed player's payoff distributions equal to half the pie; when pie size is large ($\pi \geq \$5$), the modes are at \$2, though there are second modes at half the pie.*

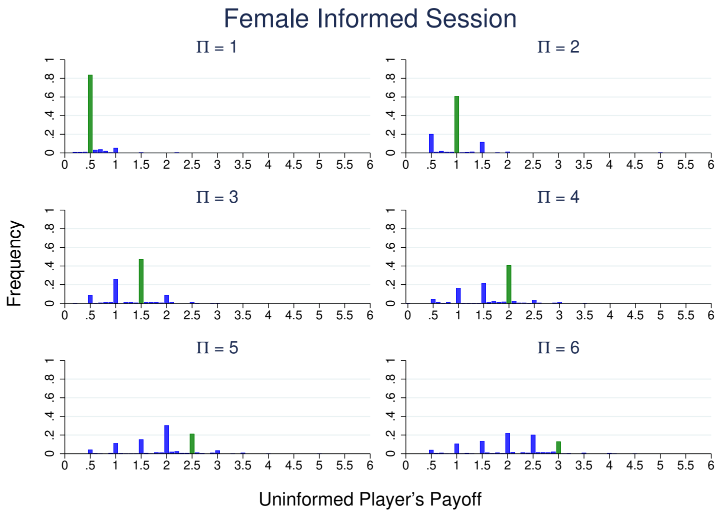

Figure C.3: Uninformed player's payoff relative frequencies in female-informed sessions (conditional on reaching a deal). The bin size is 0.1 and the green bar locates half the pie in each distribution.

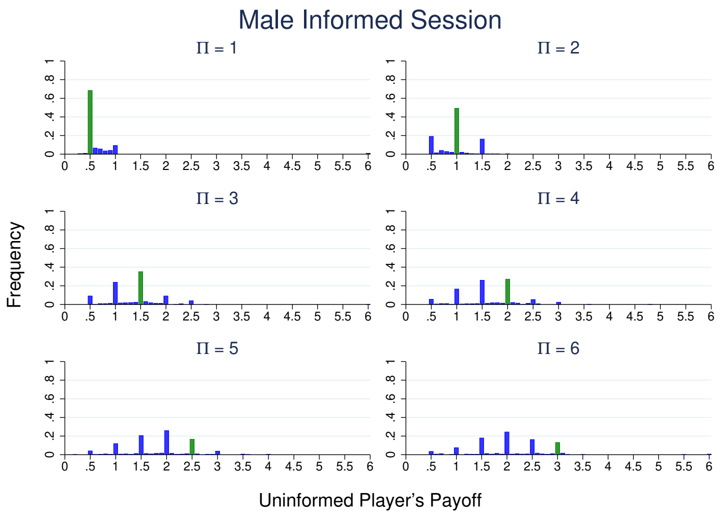

Figure C.4: Uninformed player's payoff relative frequencies in male-informed sessions (conditional on reaching a deal). The bin size is 0.1 and the green bar locates half the pie in each distribution.

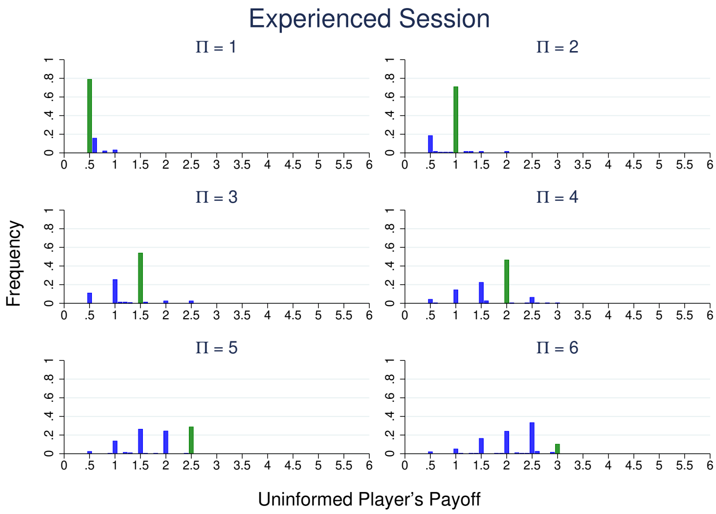

Figure C.5: Uninformed player's payoff relative frequencies in the experienced session (conditional on reaching a deal). The bin size is 0.1 and the green bar locates the half of the pie in each distribution.

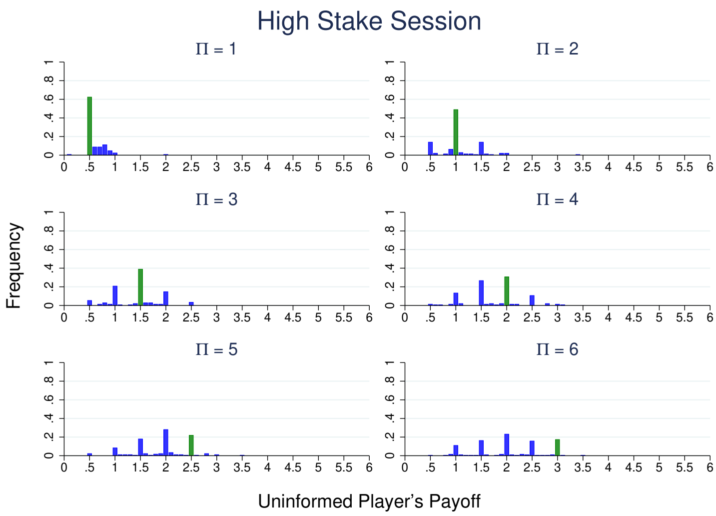

Figure C.6: Uninformed player's payoff relative frequencies in the high stake session (conditional on reaching a deal). The bin size is 0.1 and the green bar locates the half of the pie in each distribution.

C.1.3 Bargaining Dynamics

Supplementary Figure C.7 to Supplementary Figure C.10 show the dynamics of mean bargaining positions across different points of time. From these figures, we can see that no matter in which treatment, the informed player's offer increases with time, while the uninformed players' demand decreases. Therefore, we can first conclude that the empirical pattern found in Camerer et al. [2019] is robust to the gender difference, experience and stakes.

Result C.3. *In all treatments, the informed player's offer increases, and the uninformed player's demand decreases with time.*

Although the general dynamic changes are similar, there are still some differences in information transmission among treatments. Comparing the female-informed and male-informed treatments, information transmission in male-informed sessions is better than that in female-informed sessions since in the latter the final positions are not clearly separated for pie size greater than \$4.

On the other hand, from Supplementary Figure C.9, we can see that the experienced treatment has the worst information transmission because the range of the average final position is only about \$0.6. By contrast, the high-stake session has the best information transmission when pie size is small, in which there is clear separation in the uninformed player's final position.

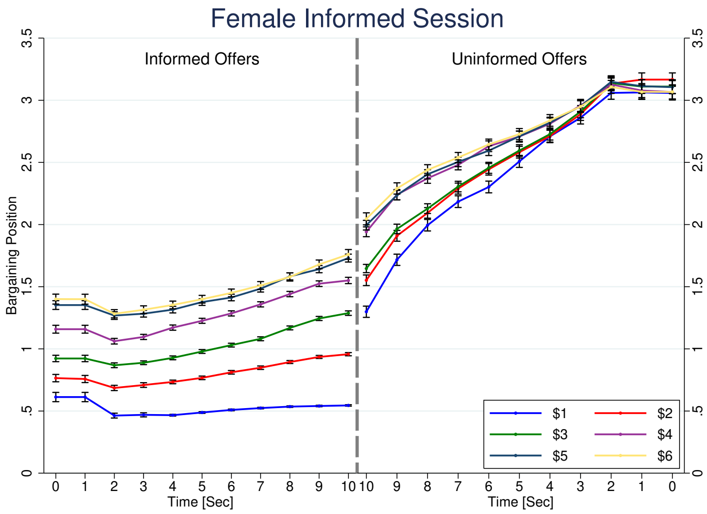

Figure C.7: Mean Bargaining Position for All Pie Sizes in Female Informed Session (All Rounds Pooled). The mean position is sampled at every second and standard errors are overlaid.

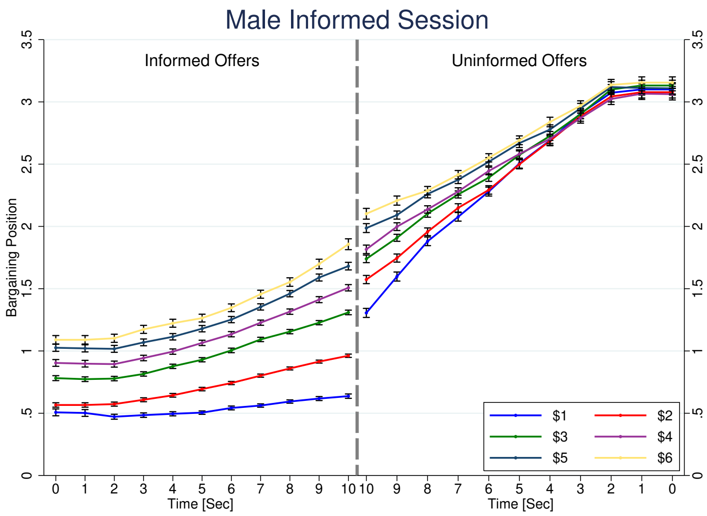

Figure C.8: Mean Bargaining Position for All Pie Sizes in Male Informed Session (All Rounds Pooled). The mean position is sampled at every second and standard errors are overlaid.

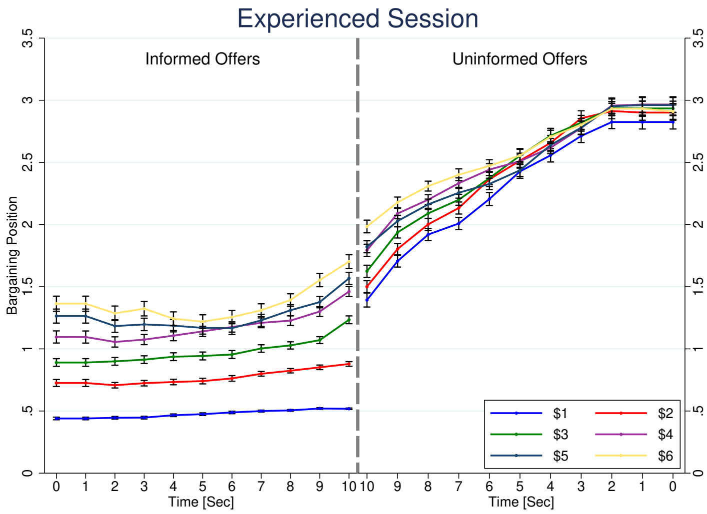

Figure C.9: Mean Bargaining Position for All Pie Sizes in the experienced session (All Rounds Pooled). The mean position is sampled at every second and standard errors are overlaid.

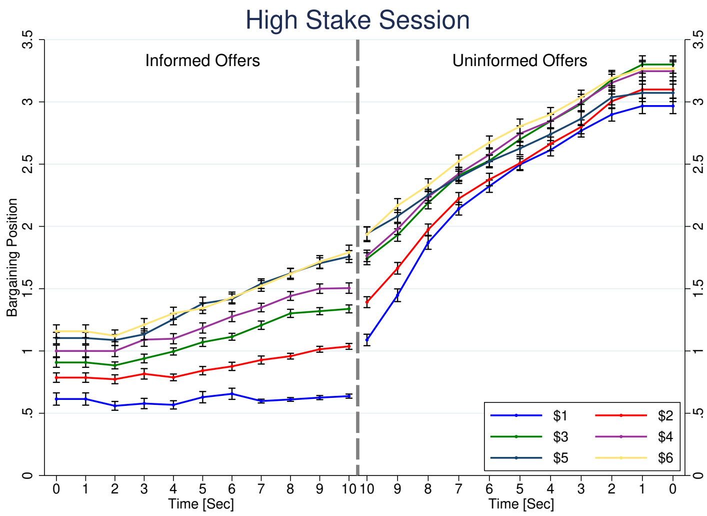

Figure C.10: Mean Bargaining Position for All Pie Sizes in high stake session (All Rounds Pooled). The mean position is sampled at every second and standard errors are overlaid.

C.1.4 Deadline Effect

Supplementary Figure C.11 and Supplementary Figure C.12 show the CDF of deals over time , which sharply increases as the deadline approached. This "endgame effect" is common to all pie sizes in both treatments. Basically, more than half of the deals are made in the last two seconds no matter in which session. Furthermore, deals are reached sooner when the pie is larger. This confirms the last empirical trend identified in Camerer et al. [2019]. One thing worth noticing here is that in the experienced session, players are more likely to reach a deal in the early seconds when pie size is large.

Result C.4. *Most deals are made close to the deadline.*

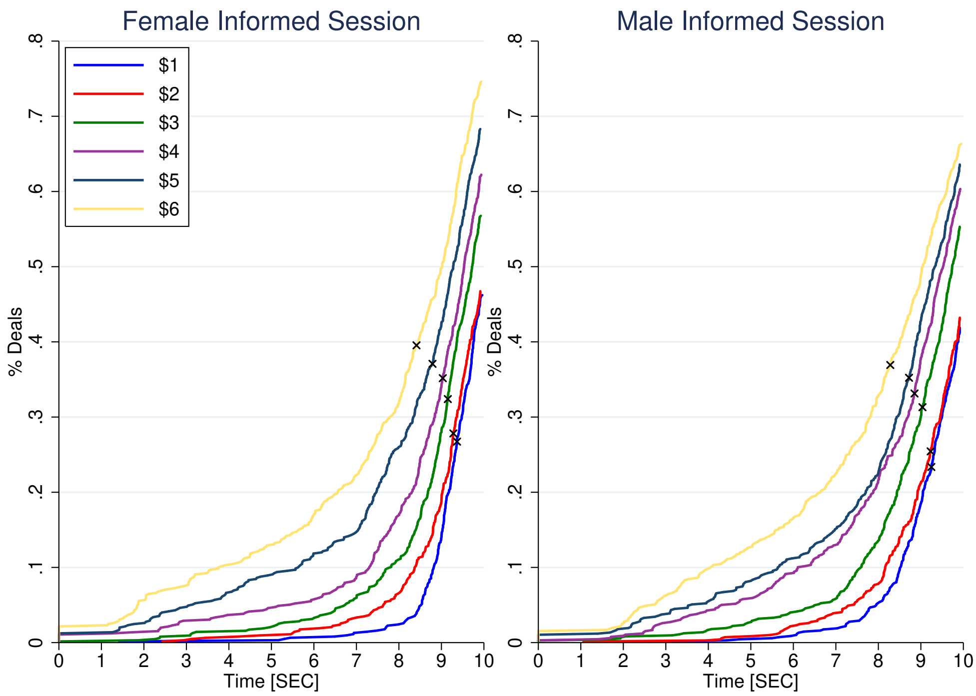

Figure C.11: Cumulative Distribution of Deal Times by Pie Size in Female and Male Informed Sessions. Median deal times in different pie sizes are marked by a cross.

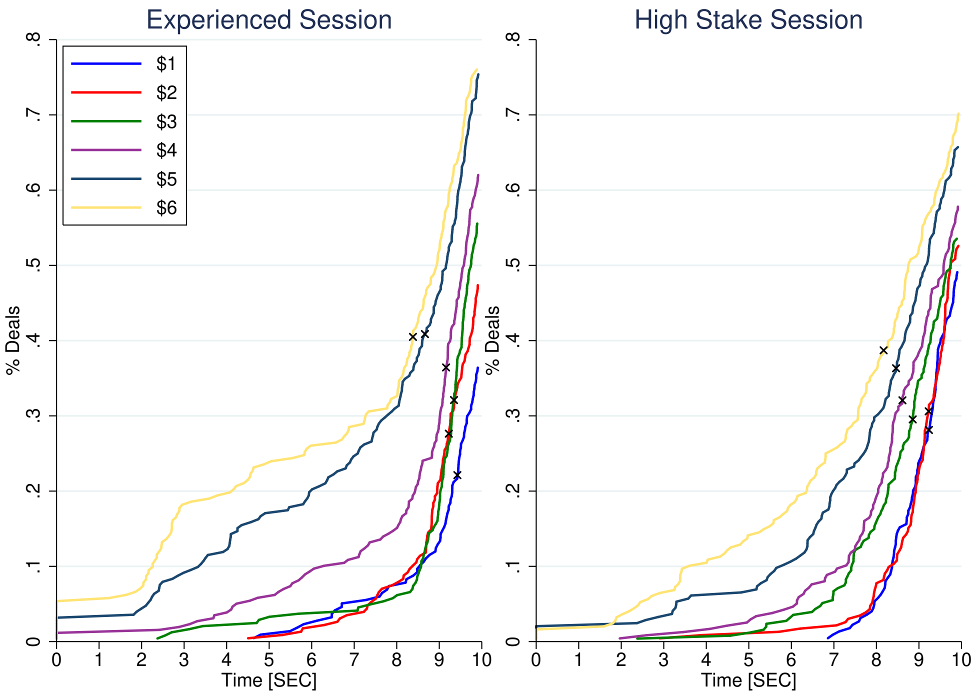

Figure C.12: Cumulative Distribution of Deal Times by Pie Size in the experienced and high stake. Median deal times in different pie sizes are marked by a cross.

C.2 Comparison with Theoretical Predictions

C.2.1 Non-parametric Hypothesis Tests

In this section, we conduct non-parametric hypothesis tests on the predictions of Lemma 1. First of all, Lemma 1 predicts that deal rates increase with pie sizes. Table 2 and Figure 2 show this monotonic pattern in both female-informed and male-informed sessions. A non-parametric Wilcoxon-type test for trend in deal rates over pie sizes reject the null hypothesis of no trend (female-informed sessions: $z = 11.98, p < 0.001$, two-sided; male-informed sessions: $z = 12.63, p < 0.001$, two-sided).

The second prediction of Lemma 1 is that the uninformed player's payoff monotonically increases with pie size, conditional on reaching a deal. A non-parametric Kruskal-Wallis test (with adjustments for ties) rejects the null hypothesis that distributions of uninformed player's payoff are the same for each pie size in both treatments (female-informed sessions: $\chi^2(5) = 1284.828, p < 0.001$, two-sided; male-informed sessions: $\chi^2(5) = 1140.427, p < 0.001$, two-sided). In addition, a non-parametric Wilcoxon-type test rejects the null hypothesis of no trend in the payoffs across pie size conditional on reaching a deal (female-informed sessions: $z = 34.92, p < 0.001$, two-sided; male-informed sessions: $z = 32.93, p < 0.001$, two-sided).

Lastly, Lemma 1 also predicts that the uninformed player's payoff is identical for all pie sizes where the deal rate is 1. In particular, the efficient equilibrium predicts that the deal rates would be 1 for all $\pi \geq \$4$. However, the equal-split equilibrium predicts the deal rate is 1 only when pie size is 6. Table 1 and Figure 2 show that strikes are common even when pie size is 6 in both female-informed and male-informed sessions. A non-parametric Kruskal–Wallis test for equality of payoff distributions (with corrections for ties) rejects the null hypothesis that mean payoffs of the uninformed player, conditional on reaching a deal, are the same for pie sizes 4, 5, and 6 (female-informed Session: $\chi^2(2) = 95.623, p < 0.001$, two-sided; male-informed Sessions: $\chi^2(2) = 79.255, p < 0.001$, two-sided).

C.2.2 Regression Analyses

In this section, we perform two sets of linear regressions to test the theoretical predictions on deal rates and the uninformed player's payoffs conditional on reaching a deal.

Deal Rate Regression

First of all, we predict whether a deal is reached or not using the following specification:

$$y_{iust} = \alpha_0 + \alpha_1 \pi_{iust} + \alpha_2 d_{iust} \left(\pi_{iust} - 4\right) + \mathbb{X}_{iust}\beta + \epsilon_{iust}.$$

Here, y_{iust} is the dummy variable for reaching a deal between informed player i and uninformed player u in period t of session s. The spline term $d_{iust} \left(\pi_{iust} - 4\right)$ consists of two parts: the dummy variable d_{iust} for pie sizes greater or equal to 4 and $(\pi_{iust} - 4)$ as the increment of pie size beyond 4. The efficient equilibrium predicts that the deal rate is $\frac{2}{5}$ when pie size is 1, increasing by $\frac{1}{5}$ per unit in pie size. Hence, the deal rate would be 1 when pie size is greater or equal to 4. In contrast, the equal-split equilibrium predicts the deal rate to be $\frac{2}{7}$ when pie size is 1, increasing by $\frac{1}{7}$ per unit in pie size. Therefore, the efficient equilibrium predicts $(\alpha_0, \alpha_1, \alpha_2) = (\frac{1}{5}, \frac{1}{5}, -\frac{1}{5})$, while the equal-split equilibrium predicts that $(\alpha_0, \alpha_1, \alpha_2) = (\frac{1}{7}, \frac{1}{7}, 0)$.

Table C.2 reports the regression results. All models include standard errors clustered at the session level, to account for dependence in residuals within a particular session. Model A provides the baseline results predicting deal rate with pie size and the spline term, using pooled data from both female-informed and male-informed sessions. Models B includes session controls, and Model C employs controls at the level of individual subject pairs (the smallest grouping available). Model D drops these controls and adds an indicator term controlling for female-informed sessions (Female = 1), as well as an indicator term for the last 60 rounds of the experiment (rounds 61–120) to capture the effect of experience. Model E adds interactions between female-informed/experience and pie size/spline. In model F, we further control for informed and uninformed player's risk preference ($ln(\rho_i), ln(\rho_u)$) and degree of loss aversion (λ_i, λ_u), as well as the initial offer and initial demand. Model G drops the female-informed session indicator and other controls but adds session-level controls. Model H adds back the controls for risk preference, loss aversion and the initial bargaining positions.

In all models, the coefficient on pie size is always significantly positive, ranging from 6.3% to 8.1%. This estimate is robust to different specification, which supports the prediction of the IC condition in Lemma 1. Yet, the slope coefficient on pie size is smaller than predicted from either equilibrium. The constant term from Model A is about 0.42 which is larger than predicted from either equilibrium. The spline term is negatively significant only when we control for interactions with experience, indicating that the players tend to coordinate on the efficient equilibrium at the beginning, but later move toward the equal-split equilibrium. In fact, experience plays an important role in reaching a deal. Players are roughly 8.5% more likely to reach a deal in later rounds. Also, the slope coefficient on the pie size is also significantly smaller in later rounds. On the other hand, it seems like whether the informed player is female is not a significant factor to predict deal rate. The dummy variable for female-informed sessions is only barely significant in Model E. Lastly, the initial position of the uninformed player would also affect deal rate.

Wage Regression

Similarly, we perform linear regressions to test the focal-split predictions regarding payoffs conditional on players reaching a deal. The efficient equilibrium predicts equal-splits when the pie is small, and the uninformed player's conditional payoff would be 2 when the pie size is 4 or greater. The equal-split equilibrium predicts a 50/50 split for all pie sizes. We test these predictions with the following specification:

$$w_{iust} = \alpha_0 + \alpha_1 \pi_{iust} + \alpha_2 d_{iust} \left(\pi_{iust} - 4\right) + \mathbb{X}_{iust}\beta + \epsilon_{iust},$$

where w_{iust} is the uninformed payoff (conditional on reaching a deal) agreed upon by informed player i and uninformed player u in period t of session s. The efficient equilibrium predicts that $(\alpha_0, \alpha_1, \alpha_2) = (\frac{1}{2}, \frac{1}{2}, -\frac{1}{2})$, yet the equal-split equilibrium predicts $(\alpha_0, \alpha_1, \alpha_2) = (\frac{1}{2}, \frac{1}{2}, 0)$.

In all specifications, the coefficient of pie size is positively significant, ranging from 0.28 to 0.35, but these point estimates are lower than the theoretical prediction. The spline term is roughly -0.16 (Model A to D), which is negatively significant. However, as we control for interactions between whether the informed player is female and experience, the spline term is no longer significant (Model E and F). This indicates that players in female-informed sessions are slightly more likely to coordinate on the efficient equilibrium than those in male-informed sessions. On the other hand, the constant term ranges from 0.24 to 0.34 (Model A to D), which is also lower than the theoretical prediction. Model F and H includes risk preferences, degrees of loss aversion and initial positions and find that initial offers play a significant role in determining final payoffs of the uninformed.

Table C.2: Linear Regressions—Predictors of Deals

	Model A	Model B	Model C	Model D	Model E	Model F	Model G	Model H
	Coef./SE	Coef./SE	Coef./SE	Coef./SE	Coef./SE	Coef./SE	Coef./SE	Coef./SE
Pie	0.0649***	0.0649***	0.0634***	0.0649***	0.0809***	0.0786***	0.0756***	0.0729***
	(0.0106)	(0.0106)	(0.0084)	(0.0106)	(0.0133)	(0.0131)	(0.0088)	(0.0075)
Spline at $\pi = \$4$	-0.0253	-0.0260	-0.0183	-0.0258	-0.0509**	-0.0491**	-0.0412**	-0.0394**
	(0.0169)	(0.0169)	(0.0139)	(0.0168)	(0.0148)	(0.0154)	(0.0137)	(0.0119)
Female				0.0451	0.0722*	0.0638		
				(0.0275)	(0.0343)	(0.0369)		
Rounds $61-120$				0.0270	0.0872***	0.0893***	0.0846***	0.0864***
				(0.0185)	(0.0180)	(0.0186)	(0.0184)	(0.0193)
Female × Pie					-0.0107	-0.0118		
					(0.0206)	(0.0201)		
Female × Spline					0.0211	0.0206		
					(0.0322)	(0.0324)		
Rd.61-120 × Pie					-0.0215**	-0.0213**	-0.0207**	-0.0204**
					(0.0065)	(0.0067)	(0.0062)	(0.0066)
Rd.61-120 × Spline					0.0298	0.0296	0.0297	0.0295
					(0.0208)	(0.0204)	(0.0208)	(0.0206)
$ln(\rho_i)$						-0.0585		-0.0579
						(0.0458)		(0.0422)
λ_i						0.0039		0.0017
						(0.0047)		(0.0061)
$ln(\rho_u)$						0.0231		-0.0000
						(0.0254)		(0.0150)
λ_u						0.0081		0.0077
						(0.0062)		(0.0066)
Initial demand						-0.0302*		-0.0327*
						(0.0135)		(0.0130)
Initial offer						0.0160		0.0144
						(0.0107)		(0.0112)
Constant	0.4245***	0.4246***	0.4261***	0.3893***	0.3457***	0.4122***	0.3812***	0.4549***
	(0.0238)	(0.0290)	(0.0233)	(0.0210)	(0.0260)	(0.0573)	(0.0259)	(0.0511)
Observations	7,920	7,920	7,920	7,920	7,920	7,920	7,920	7,920
AIC	10,553.52	10,499.13	8,623.804	10,532.82	10,527.62	10,455.81	10,494.74	10,416.78
BIC	10,574.45	10,513.09	8,637.758	10,567.71	10,562.51	10,490.70	10,529.63	10,451.67
Session Controls	No	Yes	No	No	No	No	Yes	Yes
Pair Controls	No	No	Yes	No	No	No	No	No

Notes. Coef., coefficient; SE, standard errors. Standard errors (in parentheses) are clustered at the session level.

$*p < 0.1$; $**p < 0.05$; $***p < 0.01$.

Table C.3: Linear Regressions—Predictors of Uninformed Payoffs Conditional on Deal

	Model A	Model B	Model C	Model D	Model E	Model F	Model G	Model H
	Coef./SE	Coef./SE	Coef./SE	Coef./SE	Coef./SE	Coef./SE	Coef./SE	Coef./SE
Pie	0.3527***	0.3518***	0.3299***	0.3534***	0.3208***	0.2831***	0.3402***	0.2938***
	(0.0180)	(0.0186)	(0.0194)	(0.0181)	(0.0218)	(0.0293)	(0.0173)	(0.0233)
Spline at $\pi = \$4$	-0.1641**	-0.1619**	-0.1534**	-0.1653**	-0.1192	-0.1246	-0.1390**	-0.1390**
	(0.0447)	(0.0455)	(0.0452)	(0.0455)	(0.0725)	(0.0785)	(0.0358)	(0.0445)
Female				0.0060	-0.1249	-0.1591		
				(0.0599)	(0.0868)	(0.1006)		
Rounds $61 - 120$				0.0338	-0.0217	0.0116	-0.0292	0.0031
				(0.0217)	(0.0289)	(0.0273)	(0.0335)	(0.0333)
Female $\times$ Pie					0.0424	0.0260		
					(0.0322)	(0.0379)		
Female $\times$ Spline					-0.0475	-0.0390		
					(0.0918)	(0.0931)		
Rd.61-120 $\times$ Pie					0.0215	0.0158	0.0240	0.0191
					(0.0178)	(0.0143)	(0.0196)	(0.0160)
Rd.61-120 $\times$ Spline					-0.0433	-0.0127	-0.0474	-0.0202
					(0.0483)	(0.0374)	(0.0496)	(0.0378)
$ln(\rho_i)$						-0.2145		-0.2226
						(0.1611)		(0.1592)
λ_i						-0.0134		-0.0307
						(0.0703)		(0.0771)
$ln(\rho_u)$						-0.0156		0.0150
						(0.0483)		(0.0337)
λ_u						0.0178**		0.0149
						(0.0067)		(0.0099)
Initial demand						0.0365*		0.0344*
						(0.0154)		(0.0139)
Initial offer						0.2857***		0.2887***
						(0.0482)		(0.0455)
Constant	0.2617***	0.2636***	0.3408***	0.2394**	0.3350***	0.0890	0.2773***	0.0539
	(0.0503)	(0.0489)	(0.0566)	(0.0832)	(0.0645)	(0.1541)	(0.0445)	(0.1280)
Observations	5,058	5,058	5,058	5,058	5,058	5,058	5,058	5,058
AIC	8,647.621	8,572.667	5,678.255	8,647.008	8,637.756	7,743.683	8,572.392	7,644.838
BIC	8,667.207	8,585.725	5,691.312	8,678.651	8,670.400	7,776.327	8,605.036	7,677.481
Session Controls	No	Yes	No	No	No	No	Yes	Yes
Pair Controls	No	No	Yes	No	No	No	No	No

Notes. Coef., coefficient; SE, standard errors. Standard errors (in parentheses) are clustered at the session level.

$*p < 0.1$; $**p < 0.05$; $***p < 0.01$.

Bibliography

D. R. Ames and M. F. Mason. Tandem anchoring: Informational and politeness effects of range offers in social exchange. *Journal of personality and social psychology*, 108(2):254, 2015.

S. S. Anand, C. Hawkes, R. J. De Souza, A. Mente, M. Dehghan, R. Nugent, M. A. Zulyniak, T. Weis, A. M. Bernstein, R. M. Krauss, et al. Food consumption and its impact on cardiovascular disease: importance of solutions focused on the globalized food system: a report from the workshop convened by the world heart federation. *Journal of the American College of Cardiology*, 66(14):1590–1614, 2015.

L. M. Ausubel and R. J. Deneckere. Efficient sequential bargaining. *The Review of Economic Studies*, 60(2):435–461, 1993.

L. M. Ausubel, P. Cramton, and R. J. Deneckere. Bargaining with incomplete information. *Handbook of game theory with economic applications*, 3:1897–1945, 2002.

J. S. Banks and J. Sobel. Equilibrium selection in signaling games. *Econometrica: Journal of the Econometric Society*, pages 647–661, 1987.

N. Bardsley, J. Mehta, C. Starmer, and R. Sugden. Explaining focal points: cognitive hierarchy theory versus team reasoning. *The Economic Journal*, 120(543):40–79, 2009.

A. Belloni, D. Chen, V. Chernozhukov, and C. Hansen. Sparse models and methods for optimal instruments with an application to eminent domain. *Econometrica*, 80(6):2369–2429, 2012.

A. Belloni, V. Chernozhukov, et al. Least squares after model selection in high-dimensional sparse models. *Bernoulli*, 19(2):521–547, 2013.

H. Bester and K. Ritzberger. Strategic pricing, signalling, and costly information acquisition. *International Journal of Industrial Organization*, 19 (9):1347–1361, 2001.

K. Binmore and L. Samuelson. The evolution of focal points. *Games and Economic Behavior*, 55(1):21–42, 2006.

M. Blanco, D. Engelmann, and H. T. Normann. A within-subject analysis of other-regarding preferences. *Games and Economic Behavior*, 72(2):321–338, 2011.

O. Board. Competition and disclosure. *The Journal of Industrial Economics*, 57(1):197–213, 2009.

G. E. Bolton and A. Ockenfels. Inequality aversion, efficiency, and maximin preferences in simple distribution experiments: Comment. *American Economic Review*, 96(5):1906–1911, December 2006.

A. P. Bradley. The use of the area under the roc curve in the evaluation of machine learning algorithms. *Pattern recognition*, 30(7):1145–1159, 1997.

M. Buelens and D. Van Poucke. Determinants of a negotiator's initial opening offer. *Journal of Business and Psychology*, 19(1):23–35, 2004.

H. Cai and J. T.-Y. Wang. Overcommunication in strategic information transmission games. *Games and Economic Behavior*, 56(1):7–36, 2006.

C. Camerer. Gifts as economic signals and social symbols. *American journal of Sociology*, 94:S180–S214, 1988.

C. F. Camerer, E. Johnson, T. Rymon, and S. Sen. Cognition and framing in sequential bargaining for gains and losses. *Frontiers of game theory*, 104: 27–47, 1993.

C. F. Camerer, G. Nave, and A. Smith. Dynamic unstructured bargaining with private information: theory, experiment, and outcome prediction via machine learning. *Management Science*, 65(4):1867–1890, 2019.

G. Charness and M. Rabin. Understanding social preferences with simple tests. *The Quarterly Journal of Economics*, 117(3):817–869, 2002.

E. Chaussees. Infinite-horizon models of bargaining with one-sided incomplete information. *Game-theoretic models of bargaining, Cambridge University Press, West Nyack, NY*, pages 73–98, 1985.

M. J. Christoph, N. Larson, M. N. Laska, and D. Neumark-Sztainer. Nutrition facts panels: who uses them, what do they use, and how does use relate to dietary intake? *Journal of the Academy of Nutrition and Dietetics*, 118 (2):217–228, 2018.

P. C. Cramton. Bargaining with incomplete information: An infinite-horizon model with two-sided uncertainty. *The Review of Economic Studies*, 51(4): 579–593, 1984.

V. P. Crawford. Lying for strategic advantage: Rational and boundedly rational misrepresentation of intentions. *American Economic Review*, 93 (1):133–149, 2003.

J. Crémer and R. P. McLean. Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. *Econometrica*, 53:345–361, 1985.

J. Crémer and R. P. McLean. Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica: Journal of the Econometric Society*, pages 1247–1257, 1988.

R. Croson, T. Boles, and J. Murnighan. Cheap talk in bargaining experiments: Lying and threats in ultimatum games. *Journal of Economic Behavior & Organization*, 51(2):142–159, 2003.

T. E. Daniel, D. A. Seale, and A. Rapoport. Strategic play and adaptive learning in the sealed-bid bargaining mechanism. *Journal of Mathematical Psychology*, 42(2-3):133–166, 1998.

A. De Bruyn and G. E. Bolton. Estimating the influence of fairness on bargaining behavior. *Management Science*, 54(10):1774–1791, 2008.

S. DellaVigna, J. A. List, U. Malmendier, and G. Rao. Voting to tell others. *The Review of Economic Studies*, 84(1):143–181, 2016.

D. Dranove and G. Z. Jin. Quality disclosure and certification: Theory and practice. *Journal of Economic Literature*, 48(4):935–63, 2010.

P. Ducrot, C. Julia, C. Méjean, E. Kesse-Guyot, M. Touvier, L. K. Fezeu, S. Hercberg, and S. Péneau. Impact of different front-of-pack nutrition labels on consumer purchasing intentions: a randomized controlled trial. *American journal of preventive medicine*, 50(5):627–636, 2016.

R. Durante, L. Putterman, and J. Weele. Preferences for redistribution and perception of fairness: An experimental study. *Journal of the European Economic Association*, 12(4):1059–1086, 2014.

D. Dzyabura and J. R. Hauser. Active machine learning for consideration heuristics. *Marketing Science*, 30(5):801–819, 2011.

S. El Harbi, I. Bekir, G. Grolleau, and A. Sutan. Efficiency, equality, positionality: What do people maximize? experimental vs. hypothetical evidence from tunisia. *Journal of Economic Psychology*, 47:77–84, 2015.

D. Engelmann and M. Strobel. Inequality aversion, efficiency, and maximin preferences in simple distribution experiments. *American Economic Review*, 94(4):857–869, September 2004.

D. Engelmann and M. Strobel. Inequality aversion, efficiency, and maximin preferences in simple distribution experiments: Reply. *American Economic Review*, 96(5):1918–1923, December 2006.

E. Eyster and M. Rabin. Cursed equilibrium. *Econometrica*, 73(5):1623–1672, 2005.

L. Fang. The effects of online review platforms on restaurant revenue, survival rate, consumer learning and welfare. 2019.

E. Fehr, M. Naef, and K. M. Schmidt. Inequality aversion, efficiency, and maximin preferences in simple distribution experiments: Comment. *American Economic Review*, 96(5):1912–1917, December 2006.

N. Feltovich, R. Harbaugh, and T. To. Too cool for school? signalling and countersignalling. *RAND Journal of Economics*, pages 630–649, 2002.

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

M. J. Fishman and K. M. Hagerty. Mandatory versus voluntary disclosure in markets with informed and uninformed customers. *Journal of Law, Economics, and organization*, 19(1):45–63, 2003.

R. Forsythe, J. Kennan, and B. Sopher. Dividing a shrinking pie: an experimental study of strikes in bargaining games with complete information. *Research in experimental economics*, 4:223–68, 1991a.

R. Forsythe, J. Kennan, and B. Sopher. An experimental analysis of strikes in bargaining games with one-sided private information. *The American Economic Review*, pages 253–278, 1991b.

D. Friedman and R. Oprea. A continuous dilemma. *American Economic Review*, 102(1):337–63, 2012.

M. Frondel, A. Gerster, and C. Vance. The power of mandatory quality disclosure: Evidence from the german housing market. *Journal of the Association of Environmental and Resource Economists*, 7(1):181–208, 2020.

W. Fuchs and A. Skrzypacz. Bargaining with deadlines and private information. *American Economic Journal: Microeconomics*, 5(4):219–43, 2013.

S. Gächter and A. Riedl. Moral property rights in bargaining with infeasible claims. *Management Science*, 51(2):249–263, 2005.

F. Galeotti, M. Montero, and A. Poulsen. Efficiency versus equality in bargaining. *Journal of the European Economic Association*, 17(6):1941–1970, 2018.

A. D. Galinsky and T. Mussweiler. First offers as anchors: the role of perspective-taking and negotiator focus. *Journal of personality and social psychology*, 81(4):657, 2001.

B. Ghosh and M. R. Galbreth. The impact of consumer attentiveness and search costs on firm quality disclosure: A competitive analysis. *Management Science*, 59(11):2604–2621, 2013.

S. J. Grossman. The informational role of warranties and private disclosure about product quality. *The Journal of Law and Economics*, 24(3):461–483, 1981.

S. J. Grossman and O. D. Hart. Disclosure laws and takeover bids. *The Journal of Finance*, 35(2):323–334, 1980.

S. J. Grossman and M. Perry. Sequential bargaining under asymmetric information. *Journal of Economic Theory*, 39(1):120–154, 1986.

F. Gul and H. Sonnenschein. On delay in bargaining with one-sided uncertainty. *Econometrica: Journal of the Econometric Society*, pages 601–611, 1988.

L. Guo and Y. Zhao. Voluntary quality disclosure and market interaction. *Marketing Science*, 28(3):488–501, 2009.

W. Güth and E. Van Damme. Information, strategic behavior, and fairness in ultimatum bargaining: An experimental study. *Journal of mathematical Psychology*, 42(2-3):227–247, 1998.

W. Güth, S. Huck, and P. Ockenfels. Two-level ultimatum bargaining with incomplete information: An experimental study. *The Economic Journal*, 106(436):593–604, 1996.

J. A. Hanley and B. J. McNeil. The meaning and use of the area under a receiver operating characteristic (roc) curve. *Radiology*, 143(1):29–36, 1982.

S. Hargreaves Heap, D. Rojo Arjona, and R. Sugden. How portable is level-0 behavior? a test of level-k theory in games with non-neutral frames. *Econometrica*, 82(3):1133–1151, 2014.

D. K. Herreiner and C. Puppe. Equitable allocations in experimental bargaining games: Inequality a version versus efficiency. Technical report, Bonn econ discussion papers, 2004.

B. Hollenbeck. Online reputation mechanisms and the decreasing value of chain affiliation. *Journal of Marketing Research*, 55(5):636–654, 2018.

B. Hollenbeck, S. Moorthy, and D. Proserpio. Advertising strategy in the presence of reviews: An empirical analysis. *Marketing Science*, 38(5):793–811, 2019.

B. Holmström and R. B. Myerson. Efficient and durable decision rules with incomplete information. *Econometrica*, pages 1799–1819, 1983.

A. Isoni, A. Poulsen, R. Sugden, and K. Tsutsui. Focal points in tacit bargaining problems: Experimental evidence. *European Economic Review*, 59: 167–188, 2013.

A. Isoni, A. Poulsen, R. Sugden, and K. Tsutsui. Efficiency, equality, and labeling: An experimental investigation of focal points in explicit bargaining. *American Economic Review*, 104(10):3256–87, 2014.

A. Ispano and P. Schwardmann. Competition over cursed consumers. 2018.

N. Jacquemet and A. Zylbersztejn. What drives failure to maximize payoffs in the lab? a test of the inequality aversion hypothesis. *Review of Economic Design*, 18(4):243–264, 2014.

B. J. Jansen, M. Zhang, K. Sobel, and A. Chowdury. Twitter power: Tweets as electronic word of mouth. *Journal of the American society for information science and technology*, 60(11):2169–2188, 2009.

M. C. Janssen. Rationalizing focal points. *Theory and Decision*, 50(2):119–148, 2001.

M. C. Janssen. On the strategic use of focal points in bargaining situations. *Journal of Economic Psychology*, 27(5):622–634, 2006.

S. Jayaraman and J. S. Wu. Is silence golden? real effects of mandatory disclosure. *The Review of Financial Studies*, 32(6):2225–2259, 2019.

E. J. Johnson, C. Camerer, S. Sen, and T. Rymon. Detecting failures of backward induction: Monitoring information search in sequential bargaining. *Journal of Economic Theory*, 104(1):16–47, 2002.

B. Jovanovic. Selection and the evolution of industry. *Econometrica: Journal of the Econometric Society*, pages 649–670, 1982.

J. H. Kagel and K. W. Wolfe. Tests of fairness models based on equity considerations in a three-person ultimatum game. *Experimental Economics*, 4(3):203–219, 2001.

J. H. Kagel, C. Kim, and D. Moser. Fairness in ultimatum games with asymmetric information and asymmetric payoffs. *Games and Economic Behavior*, 13(1):100–110, 1996.

E. Kalai, M. Smorodinsky, et al. Other solutions to nash's bargaining problem. *Econometrica*, 43(3):513–518, 1975.

J. Kennan and R. Wilson. Can strategic bargaining models explain collective bargaini. *The American Economic Review*, 80(2):405, 1990.

J. Kennan and R. Wilson. Bargaining with private information. *Journal of Economic Literature*, 31(1):45–104, 1993.

N. Khandpur, P. d. M. Sato, L. A. Mais, A. P. B. Martins, C. G. Spinillo, M. T. Garcia, C. F. U. Rojas, and P. C. Jaime. Are front-of-package warning labels more effective at communicating nutrition information than traffic-light labels? a randomized controlled experiment in a brazilian sample. *Nutrients*, 10(6):688, 2018.

I. Krajbich, C. Camerer, J. Ledyard, and A. Rangel. Using neural measures of economic value to solve the public goods free-rider problem. *Science*, 326(5952):596–599, 2009.

P. H. Kriss, R. Nagel, and R. A. Weber. Implicit vs. explicit deception in ultimatum games with incomplete information. *Journal of Economic Behavior & Organization*, 93:337–346, 2013.

H. Kristensen and T. Gärling. The effects of anchor points and reference points on negotiation process and outcome. *Organizational Behavior and Human Decision Processes*, 71(1):85–94, 1997.

A. Kritikos and F. Bolle. Distributional concerns: equity-or efficiency-oriented? *Economics Letters*, 73(3):333–338, 2001.

D. Levin, J. Peck, and L. Ye. Quality disclosure and competition. *The Journal of Industrial Economics*, 57(1):167–196, 2009.

P.-H. Lin, A. L. Brown, T. Imai, J. T.-y. Wang, S. Wang, and C. Camerer. General economic principles of bargaining and trade: Evidence from 2,000 classroom experiments. *Available at SSRN 3250495*, 2018.

R. López-Pérez, Á. Pintér, and H. J. Kiss. Does payoff equity facilitate coordination? a test of schelling's conjecture. *Journal of Economic Behavior & Organization*, 117:209–222, 2015.

M. F. Mason, A. J. Lee, E. A. Wiley, and D. R. Ames. Precise offers are potent anchors: Conciliatory counteroffers and attributions of knowledge in negotiations. *Journal of Experimental Social Psychology*, 49(4):759–763, 2013.

A. D. Mathios. The impact of mandatory disclosure laws on product choices: An analysis of the salad dressing market. *The Journal of Law and Economics*, 43(2):651–678, 2000.

D. Mayzlin and J. Shin. Uninformative advertising as an invitation to search. *Marketing Science*, 30(4):666–685, 2011.

R. P. McAfee and P. J. Reny. Correlated information and mecanism design. *Econometrica: Journal of the Econometric Society*, pages 395–421, 1992.

P. Milgrom and J. Roberts. Price and advertising signals of product quality. *Journal of political economy*, 94(4):796–821, 1986.

M. Mitzkewitz and R. Nagel. Experimental results on ultimatum games with incomplete information. *International Journal of Game Theory*, 22(2):171–198, 1993.

R. B. Myerson. Incentive compatibility and the bargaining problem. *Econometrica*, pages 61–73, 1979.

R. B. Myerson. Two-person bargaining problems with incomplete information. *Econometrica*, pages 461–487, 1984.

J. F. Nash Jr. The bargaining problem. *Econometrica: Journal of the Econometric Society*, pages 155–162, 1950.

P. Nelson. Information and consumer behavior. *Journal of political economy*, 78(2):311–329, 1970.

P. Nelson. Advertising as information. *Journal of political economy*, 82(4): 729–754, 1974.

P. Newberry and X. Zhou. Heterogeneous effects of online reputation for local and national retailers. *International Economic Review*, 60(4):1565–1587, 2019.

R. V. Nydegger and G. Owen. Two-person bargaining: An experimental test of the nash axioms. *International Journal of game theory*, 3(4):239–249, 1974.

J. Ochs and A. E. Roth. An experimental study of sequential bargaining. *The American Economic Review*, pages 355–384, 1989.

D. G. Pruitt. *Negotiation behavior*. Academic Press, 2013.

R. Radner and A. Schotter. The sealed-bid mechanism: An experimental study. *Journal of Economic Theory*, 48(1):179–220, 1989.

A. Rapoport and M. A. Fuller. Bidding strategies in a bilateral monopoly with two-sided incomplete information. *Journal of Mathematical Psychology*, 39 (2):179–196, 1995.

A. Rapoport, I. Erev, and R. Zwick. An experimental study of buyer-seller negotiation with one-sided incomplete information and time discounting. *Management Science*, 41(3):377–394, 1995.

A. Rapoport, T. E. Daniel, and D. A. Seale. Reinforcement-based adaptive learning in asymmetric two-person bargaining with incomplete information. *Experimental Economics*, 1(3):221–253, 1998.

A. E. Roth. Toward a focal point theory of bargaining. *Game-Theoretic Models of Bargaining*, 259:265–67, 1985.

A. E. Roth and M. W. Malouf. Game-theoretic models and the role of information in bargaining. *Psychological review*, 86(6):574, 1979.

A. E. Roth and J. K. Murnighan. The role of information in bargaining: An experimental study. *Econometrica: Journal of the Econometric Society*, pages 1123–1142, 1982.

A. E. Roth, M. W. Malouf, and J. K. Murnighan. Sociological versus strategic factors in bargaining. *Journal of Economic Behavior & Organization*, 2(2): 153–177, 1981.

A. E. Roth, J. K. Murnighan, and F. Schoumaker. The deadline effect in bargaining: Some experimental evidence. *The American Economic Review*, 78(4):806–823, 1988.

A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica: Journal of the Econometric Society*, pages 97–109, 1982.

A. Rubinstein. A bargaining model with incomplete information about time preferences. *Econometrica: Journal of the Econometric Society*, pages 1151–1172, 1985.

T. Schelling. The strategy of conflict harvard university press cambridge, 1960.

M. Spence. Job market signaling. In *Uncertainty in economics*, pages 281–306. Elsevier, 1978.

J. Srivastava. The role of inferences in sequential bargaining with one-sided incomplete information: Some experimental evidence. *Organizational behavior and human decision processes*, 85(1):166–187, 2001.

I. Ståhl. Bargaining theory. *PhD diss. Economic Research Institute at the Stockholm School of Economics (EFI)*, 1972.

K. Stahl and R. Strausz. Certification and market transparency. *The Review of Economic Studies*, 84(4):1842–1868, 2017.

M. Stone. Cross-validatory choice and assessment of statistical predictions. *Journal of the Royal Statistical Society: Series B (Methodological)*, 36(2): 111–133, 1974.

Z. Talati, M. Egnell, S. Hercberg, C. Julia, and S. Pettigrew. Food choice under five front-of-package nutrition label conditions: an experimental study across 12 countries. *American Journal of Public Health*, 109(12):1770–1775, 2019.

L. L. Thompson, J. Wang, and B. C. Gunia. Negotiation. *Annual review of psychology*, 61:491–515, 2010.

R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.

K. Valley, L. Thompson, R. Gibbons, and M. H. Bazerman. How communication improves efficiency in bargaining games. *Games and Economic Behavior*, 38(1):127–155, 2002.

D. Van Poucke and M. Buelens. Predicting the outcome of a two-party price negotiation: Contribution of reservation price, aspiration price and opening offer. *Journal of Economic Psychology*, 23(1):67–76, 2002.

J. T.-y. Wang, M. Spezio, and C. F. Camerer. Pinocchio's pupil: using eyetracking and pupil dilation to understand truth telling and deception in sender-receiver games. *American Economic Review*, 100(3):984–1007, 2010a.

S. W. Wang, M. Filiba, and C. F. Camerer. Dynamically optimized sequential experimentation (dose) for estimating economic preference parameters. *Manuscript submitted for publication*, 2010b.

S. B. White and M. A. Neale. The role of negotiator aspirations and settlement expectancies in bargaining outcomes. *Organizational Behavior and Human Decision Processes*, 57(2):303–317, 1994.

S. B. White, K. L. Valley, M. H. Bazerman, M. A. Neale, and S. R. Peck. Alternative models of price behavior in dyadic negotiations: Market prices, reservation prices, and negotiator aspirations. *Organizational Behavior and Human Decision Processes*, 57(3):430–447, 1994.

G. Yukl. Effects of the opponent's initial offer, concession magnitude and concession frequency on bargaining behavior. *Journal of Personality and Social Psychology*, 30(3):323, 1974.

Eidesstattliche Versicherung

Ich versichere hiermit eidesstattlich, dass ich die vorliegende Arbeit selbstständig und ohne fremde Hilfe verfasst habe. Die aus fremden Quellen direkt oder indirekt übernommenen Gedanken sowie mir gegebene Anregungen sind als solche kenntlich gemacht. Die Arbeit wurde bisher keiner anderen Prüfungsbehörde vorgelegt und auch noch nicht veröffentlicht. Sofern ein Teil der Arbeit aus bereits veröffentlichten Papers besteht, habe ich dies ausdrücklich angegeben.

17.09.2020
Hung-Ni Chen