The Market for Conflicted Advice*

Briana Chang[†] Martin Szydlowski[‡]

November 30, 2017

Abstract

We present a model of the market for advice, where advisers have conflicts of interest and compete for heterogeneous customers through information provision. The competitive equilibrium features information dispersion: advisers with expertise in more information-sensitive assets attract less informed customers, provide worse information, and earn higher rents. Even though distorted information leads to lower returns, investors choose to trade through advisers, which rationalizes empirical findings. Banning conflicted payments only improves the information quality but not customers' welfare. It is the underlying distribution of financial literacy that determines welfare, and the fee structure is irrelevant.

Keywords: Information Provision, Conflicted Advice, Financial Literacy, Sorting
JEL Codes: G2, D1,D8

*We would like to thank Dean Corbae, Mark Egan, Simon Gervais, Pablo Kurlat, Ricardo Lagos, Rasmus Lentz, Gregor Matvos, Espen Moen, Dimitry Orlov, Christine Parlour, Erwan Quintin, Adriano A. Rampini, Mark Ready, Shouyong Shi, Gunther Strobl, Vish Viswanathan, and Randy Wright for their useful discussions, as well as participants at 9th NYU search theory workshop, Finance Theory Group Workshop Spring 2017, UNC/DUKE Corporate Finance Conference, Fall 2016 Midwest Economic Theory Conference, the Wisconsin School of Business, 8th Annual Conference on Money, Banking and Asset Markets, the Frankfurt School of Finance and Management, 2016 Summer Workshop on Money, Banking, Payments and Finance, University of Wisconsin, and the University of Minnesota for helpful comments.

[†]University of Wisconsin-Madison. Email: briana.chang@wisc.edu

[‡]University of Minnesota. Email: szydl002@umn.edu.

1 Introduction

Providing advice is an important function of intermediaries. The compensation structure of intermediaries, however, often leads to conflicts of interest. For example, broker-dealer firms are compensated by commissions and fund distribution fees, and realtors receive fees only if they close a deal. Financial advice, in particular, has received much attention in the empirical literature and from regulators. Specifically, it is well documented that, while many households rely on advice, conflicted advice leads to lower investment returns, imposing a substantial loss to households.[1] However, fundamentally, how markets for advice function and their consequences for investors remain unknown. Specifically, can market discipline correct these misaligned incentives? Given that household may differ in their sophistication, how do advisers' recommendations actually influence their decisions and affect their welfare?

To answer these questions, we develop a tractable framework where the quality of advice along with customers' investment decisions as well as choice of advisers are determined jointly in a competitive market. To capture conflicts of interest, we assume that advisers (brokers or realtors) are compensated only if they successfully convince their customers to take an action (buy the asset). This aims to capture the reality that, in practice, brokers are compensated by distribution fees set by fund issuers and do not charge a separate fee for their information service.[23] To attract customers, advisers then compete through information

[1]Financial advice especially has received much interest in the empirical literature, which shows that financial advisers drive customers to chase returns (Linnainmaa et al. (2015)), steer them towards high-fee, actively managed investments (Mullainathan et al. (2012)), and recommend unsuitable products when they earn high commissions (Anagol et al. (2013)). As a result, portfolios of advised customers underperform (e.g. Bergstresser et al. (2009), Chalmers and Reuter (2010) and Hoechle et al. (2013)). This has drawn the attention of regulators (CEA (2015)), who have found a substantial welfare loss from conflicted advice in the United States. As a result, the department of labor has instituted a new rule holding financial advisers to the fiduciary duty standard. Other countries have enacted a variety of measures, such as banning payments from product providers to advisers (Australia, Netherlands, UK) or mandating disclosures about conflicts of interest (Canada, Germany).

[2]See, for example, *Understanding Your brokerage and Investment advisery Relationships* by Morgan Stanley (Dec 2014): "In addition to taking your orders, executing your trades and providing custody services, we also provide investor education, investment research, financial tools and professional, personalized information about financial products and services, including recommendations to our brokerage clients about whether to buy, sell or hold securities. We do not charge a separate fee for these services because these services are part of, or "incidental to," our brokerage services.

[3]On the other hand, there is another type of finical advisers who charge flat fees for their service. We compare our results to these fee-only advisers In Section 5.1.

provision. Taking these misaligned incentives as given,[4] the goal of the paper is to understand to what extent competition can discipline advisers and generate valuable information for customers.

To further sheds light on how such conflicts of interest affect different agents in the economy, this paper advances a modeling framework that takes into account the heterogeneity of both customers and advisers. Customers differ in the quality of their ex-ante information and advisers differ in their value of expertise. Precisely, each adviser is perfectly informed about one particular type of asset, and assets differ in their information sensitivity. Hence, an adviser with expertise in a more information-sensitive asset effectively has more valuable knowledge.

In our setting, advisers compete for customers by posting information policies. Customers understand the misaligned incentives of advisers and choose their advisers optimally. After choosing an adviser, the customer receives the information, rationally updates his beliefs, and makes his investment decision. To attract a particular customer, an adviser must provide sufficiently valuable information so that the customer does not prefer to match with someone else. Each adviser is subject to a capacity constraint in the sense that they can at most match only with one customer and vice versa. Thus, the environment can be understood as a matching model with two-sided heterogeneity, where information provision determines the gains for both sides. The key equilibrium objects are the distribution of information quality, the value of information to customers, the profits of advisers, and the matching patterns.

Intuitively, since all advisers have incentives to deceive customers, they prefer to attract the least-informed ones. This drives up the value of information received by these customers. We show that, in equilibrium, the adviser with the most valuable knowledge attracts the least-informed customer. As a result, less-informed customers receive less valuable information than better-informed customers, while still gaining more information than better-informed customers relative to deciding on their own. This also implies that these customers' matching advisers—those with more valuable knowledge—obtain a higher profit. Thus, competition under heterogeneity leads to different rents for different agents. Information quality is dispersed, and advisers generally provide partial disclosure.

Our framework provides a direct link between information quality and the distribution

[4] Inderst and Ottaviani (2009) and Inderst and Ottaviani (2012b) provide micro-foundations for the existence for this type of compensation structure.

of customers' informedness and advisers' expertise. We show that when other customers in the economy become better informed, less-informed customers benefit and receive better information, although their own level of informedness does not change.

Conceptually, the economics behind this result are analogous to the standard assignment model with price competition. The exact quality of information a customer receives in equilibrium is pinned down such that his matching adviser would not profit from attracting his next-best competitor, a customer who is slightly better informed than him. Hence, when other customers become better-informed, meaning that his competitors become less attractive from the viewpoint of advisers, the customer becomes relatively scarce and thus receives a higher rent in the form of better information.

A fundamental difference from the standard model with price competition is that the value of information is not perfectly transferable. Hence, technically, our environment falls into the class of matching problems with imperfectly transferable utilities.[5] This suggests that the decentralized outcome may not maximize aggregate surplus. Indeed, we show that information provision is always distorted.

The existence of distortion further highlights the cost of advisers' compensation structures that we often observe, where brokers are compensated by distribution fees set by fund issuers. As shown in Section 5.1, if fees were set competitively (e.g., if advisers could commit to refunding 12b-1 fees to investors) or for the case for fee-only advisers, the environment could then be solved as the standard matching model with transferal utilities, which always guarantees the efficient outcome.

We further establish two new insights on how conflicted payments affect customers. First, counterintuitively, we show that higher advisers' fees lead to better information received by customers. This is because information and fees are substitutes from the viewpoints of customers. With competition, an adviser must provide better information to the customer in order to compensate for the higher fee.

Second, we establish an equivalence of consumers' utilities under information vs. price competition. That is, even though information is distorted in our environment, consumers' utilities generally are not. Specifically, as long as advisers are relatively scarce (i.e. they are on the short side of market), we show that all customers receive the same utilities across

[5] Legros and Newman (2007) provides a general sorting condition in an environment with imperfect transferable utilities.

the two settings. That is, customers will be indifferent between obtaining better information with a higher fee and obtaining distorted information.

Our results thus have important policy implications in the market of financial advice. When regulators limit the advisers' ability to receive kickbacks, in equilibrium advisers respond by charging a higher fee. Our irrelevance result suggests that this effect balances out perfectly. Thus, while regulation aimed at conflicted payments can improve information, contrary to the conventional wisdom, it does not actually increase consumer welfare.

Last but not least, since whether trading through intermediaries is the endogenous choice of customers, our model gives predictions on which type of customers will actually trade through advisers, and on such choice affect their return (i.e., self-directed vs. broker-client trades). We show that, in equilibrium, investment returns could actually be lower for broker-client trades, rationalizing the puzzling empirical findings. Precisely, in the parameter regime where customers only invest if they receive a positive signal which may have false negatives, customers may forego some investment opportunities (i.e., less willing to take risks) without further information. Seeking advice thus has two effects: first, it helps customers to better identify the investment opportunities, so that customers are more willing to take risks, which increases their expected utility. However, since the adviser may oversell the product due to the conflict of interest, customers may invest even when the payoff is negative.

Consistent with empirical findings, the model thus predicts that broker clients are more willing to take risk but have a lower return than self-directed investors. This result is related to the existing models based on trust. For example, Gennaioli et al. (2015) considers the environment where money managers can decrease the customers' risk aversion (i.e., anxiety). While we do not model trust, one can interpret how much a customer trusts an adviser's recommendation is endogenously determined by the quality of information. Since this value is endogenous, our framework sheds light on how returns and welfare would change in response to policy or different fee structures.

Related Literature Methodologically, our work is built on the Bayesian Persuasion approach (put forth by Kamenica and Gentzkow (2011))[6] and the literature on matching

[6]Kamenica and Gentzkow (2011) analyzes the environment with one single sender with monopoly power and one receiver.

markets.[7] Compared to models that study competition within the Bayesian Persuasion frameworks, our model has two important distinctions: first, our paper is the first to study information provision in a model of decentralized competition with heterogeneity.[8] In particular, the question in Gentzkow and Kamenica (2011) and Au and Kawai (2015) is about how the degree of competition affects information provision in a setting with multiple senders and one receiver.[9] In our framework, the market is perfectly competitive, but advisers are competing for heterogeneous customers. The question we focus on is how the underlying distribution of financial literacy affects the equilibrium value of information and how it affects different consumers with different sophistication. This is of first-order importance for policy, since generally, less sophisticated customers are perceived as being more at risk.

Second, building on matching models, competition in our framework is captured by agents' matching decisions, and within the match, the information provision is exclusive.[10] This modeling makes our framework very tractable. Despite allowing for a general message space, we provide a closed form characterization of the optimal policy. Specifically, the value of information to a customer is uniquely pinned down so that it is indeed optimal for him to stay within the match, taking into account the value provided by other advisers. Our paper is also related to Board and Lu (2015) who study a random search model in which sellers compete through persuasion.[11] The continuation value in their model shares a similar spirit

[7] Building on Becker (1973), Shapley and Shubik (1971), most works in this literature after analyzes matching patterns with transferable utilities, with few exceptions: Legros and Newman (2007) provides a general sorting condition in an environment with imperfect transferable utilities. Chiappori and Reny (2006) studies a risk-sharing problem in a matching model.

[8] Outside of Bayesian Persuasion, Ostrovsky and Schwarz (2010) studies the school's optimal disclosure of students' abilities in order to maximize their average job placement, taking into account that the job market features assortative matching. We share similar conflicts of interest in the sense that advisers want to maximize customers' purchasing rates. The focus in Ostrovsky's paper, however, is how one information provider affects the matching patterns between students (assets) and jobs (customers). The matching in our model, on the other hand, is between different information providers and customers.

[9] Gentzkow and Kamenica (2011) study a model where senders whose preferences differ from those a single receiver can simultaneously disclose information. They show that adding more senders or making their preferences less aligned improves the quality of information in equilibrium. Au and Kawai (2015) study a game with one receiver and multiple senders with possibly different priors and characterize how equilibrium information changes with each sender's prior.

[10] This also distinguishes our paper from the literature on selling information, where information is generally not exclusive. See Admati and Pfleiderer (1988) for a seminal contribution and García and Sangiorgi (2011) and Malenko and Malenko (2016) for recent work.

[11] The main tradeoff in their paper is whether a seller should provide information to discourage the buyer from continuing to search, which will depend on the value of information offered by other sellers in equilib-

as our equilibrium utilities, although matching is frictionless in our model. [12]

Regarding the literature on financial advice more specifically, a series of papers by Inderst and Ottaviani (Inderst and Ottaviani (2009) and Inderst and Ottaviani (2012b)) study the optimal compensation for a direct marketing agent (i.e., advisers) and show that a conflict of interest arises endogenously.[13] They consider a setting where producers compete through commissions paid to an adviser and advisers are assumed to have concerns for suitability for the customers' needs. We, however, take these misaligned interests as given, and focus on how advisers compete for heterogeneous consumers. Advisers do not care about the customer's well-being, but they generate valuable information purely because of competition.Stoughton et al. (2011) study the use of kickbacks to advisers when customers differ in wealth levels. They find that kickbacks always decrease customer welfare, whereas we find that customer welfare remains the same with and without kickbacks.[14]

One key difference from these previous works is that we allow customers for investing all assets without contacting an adviser (i.e., intermediary). This distinction allows us to provide empirical predictions on customers' returns between different channels: broker-client vs. self-directed channel. The economics here is also very different, since the value provided by an adviser is purely informational (not about choosing the asset itself).[15]

2 Model

Customers There is a mass of heterogeneous customers who are different in terms of their ex-ante information quality. They are indexed by type $b \in B \equiv \left[\underline{b}, \overline{b}\right]$, which is observable with $0 \leq \underline{b} < \overline{b} < 1$. Let $Q(b)$ denote the measure of customers with types weakly below b

rium.

[12]The special case of ours where advisers are homogeneous and they are on the long side, our model would yield full disclosure, which is similar to the case of vanishing search costs in their framework.

[13]Inderst and Ottaviani (2009) characterizes a monopoly problem and Inderst and Ottaviani (2012a) extend the setting to two competing manufacturers. Inderst and Ottaviani (2012b) further extends the monopoly case to two consumer types: naive ones who always believe advice and rational ones who understand adviser incentives.

[14]Additionally, in their model, advisers directly choose how to allocate each customer's money. In equilibrium advisers choose the same allocation for each customer, which we can interpret as all customers receiving the same advice. In our paper, advisers can only provide information and customers ultimately decide whether to invest. We also characterize how different customers receive different advice.

[15]This feature also distinguishes us from the credence goods literature, where sellers simultaneously offer a good and advice about its suitability. See Dulleck and Kerschbamer (2006) for a recent survey.

and let $Q(\bar{b})$ denote the overall measure of customers. $Q(b)$ admits a differentiable density $q(b)$. All customers have access to a set of assets $L \equiv [\underline{l}, \bar{l}]$, which is indexed by l. The payoff of each asset l is determined by the asset specific random variable $s \in [0,1]$ and is given by

$$y(s,l) = \begin{cases} r & \text{if} \quad s \geq \lambda \\ -l & \text{if} \quad s < \lambda, \end{cases} \tag{1}$$

where s is distributed with strictly positive and continuous pdf $f(s)$. Thus, the asset gives a positive payoff $r > 0$ only if $s \geq \lambda$. Otherwise, it results in a loss $l > 0$. This assumption implies that if a customer perfectly observes the state s, he invests if only if $s \geq \lambda$.

Each customer b can potentially invest in all assets and receives a private signal about each asset's quality, $x_l(s,b) = \mathbb{1}\{s \geq b\lambda\}$. Hence, by construction, if an asset is good ($s \geq \lambda$), customers always receive a positive signal ($x = 1$). However, there are false positives if $b\lambda \leq s < \lambda$. A customer with higher b is thus more informed, since his signal has a lower probability of false positives.

Likewise, an asset with higher l is more information sensitive: information reduces the likelihood of false positives, which is more valuable whenever l is higher.

Advisers There is a unit mass of heterogeneous advisers on the other side of the market. Each adviser has expertise in one particular type of asset l and is perfectly informed about the realized state s of this asset. Since this is the only dimension in which advisers differ, we use l to denote the type of an adviser. The distribution of adviser types is given by measure $G(l)$, which has differentiable density $g(l)$ and domain L.

The two-sided heterogeneity in our model captures two important dimensions of the market for information: customers (the demand side) differ in their ex-ante information and hence their needs for information; advisers (the supply side) differ in the value of the information they can provide, since an adviser who has expertise in a more information sensitive asset (i.e. a higher l) effectively has more valuable information.

To capture the conflict of interest between advisers and customers, we assume that an adviser receives a positive payoff $\alpha > 0$ whenever a customer invests after matching with her, independently of the realized state. This payoff captures a conflict of interest that makes advisers have an incentive to persuade customers to invest, despite them not benefiting from fix

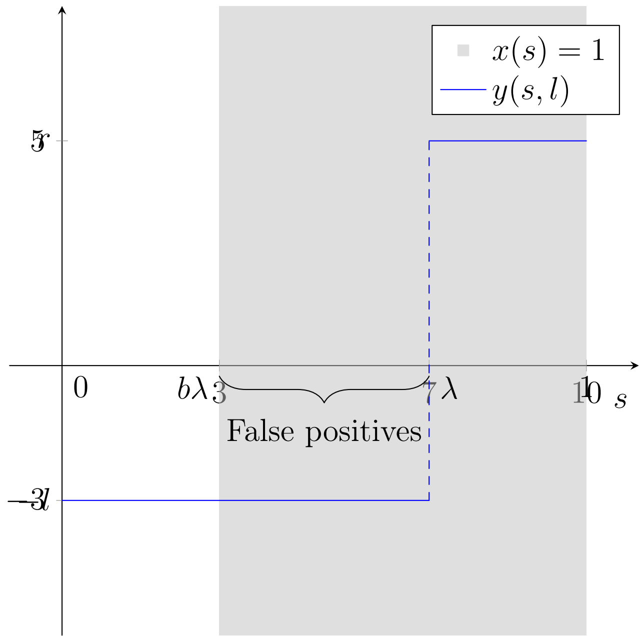

Figure 1: Payoff and Information

it.

?

In the context of financial advisers, this payoff represents a commission paid by customers. Specifically, in the case of mutual funds, customers pay additional distribution fees (e.g., loads and 12b-1 fees) when they buy funds through brokers or advisers, instead of buying fund directly.[16] More generally, for doctors, it may represent a kickback received by a pharmaceutical company, the cost of which is not directly passed on to the customer,[17] while for lawyers it may represent the reputational gain of fighting a large trial. We consequently allow the payoff to have two components: $\alpha = \alpha_c + \alpha_0$, where $\alpha_c \geq 0$ represents the portion paid by customers, and $\alpha_0 \geq 0$ represents the portion that is not.

[16] For example, Bergstresser et al. (2009) estimates that broker-channel mutual fund consumers may have paid as much as \$3.6 billion in front end loads in 2002, \$2.8 billion in back-end loads and another \$8.8 billion in 12b-1 fees, in additional to the investment management fees.

[17] For example, if insurance covers the cost of brand medication and a doctor receives a kickback from the manufacturer which is priced into the cost of the medicine, the patient does not bear this cost when making his decision.

Skip α_c ?

Market for advice We consider a decentralized market for advice, where different advisers compete for different customers through information provision. Specifically, we assume that each adviser can only serve one customer, and each customer can only obtain advice from one adviser (i.e., information about one type of asset l). To model the information provision of advisers, we assume that each adviser can post (and commit to) an information policy before the state is realized. The information policy consists of a signal $\sigma_b : [0,1] \to \Delta(\mathbb{M})$ that maps the state into a distribution over messages $m \in \mathbb{M}$.[18]

Given the posted information policies, customers make two choices. The first decision is to choose from which adviser to seek advice. This is the matching decision. The second one is whether to invest. After the match is formed, the customer observes first his private signal x_l for asset l and then the message m sent by the adviser. He then uses Bayes rule to form a posterior belief $\mu \in \Delta([0,1])$ about the state of asset l, conditional on both x and m, and decides whether or not to invest. If he decides to invest, he buys the asset l through the adviser l, and the adviser receives payoff α.[19]

For assets about which he does not receive advice, each customer makes his investment decision based on his own private signal. In other words, we assume that all investors can buy assets directly; hence, whether a customer purchases through advisers or not is an endogenous outcome.[20] We use $u^0(b,l)$ to denote the expected value of customer b for trading asset l by himself, which is given by

$$\begin{aligned} u^0(b,l) &\equiv (1-F(b\lambda))\left(\frac{1-F(\lambda)}{1-F(b\lambda)}r - \left(1-\frac{1-F(\lambda)}{1-F(b\lambda)}\right)l\right) \\ &= (1-F(\lambda))r - (F(\lambda)-F(b\lambda))l. \end{aligned} \tag{2}$$

In Equation (2), with probability $1-F(b\lambda)$, the customer's private signal is high and he chooses to invest. His expected payoff is the term in brackets. The first expression is the probability that $s \geq \lambda$ conditional on $s \geq b\lambda$ times payoff r. The second term is the probability that $s < \lambda$ conditional on $s \geq b\lambda$ times the loss l.

Recall that, when matching with an adviser, a customer needs to pay α_c whenever he chooses to invest. We can thus interpret parameter α_c as the additional cost if a customer

[18]The particular message space $\mathbb{M}$ is irrelevant, as long as it has more than two elements. This is because the customer's action (invest or don't invest) is binary. See Kamenica and Gentzkow (2011), Prop. 1.

[19]We provide a more explicit construction of the customer's problem in Appendix C.

[20]In Appendix **??**, we consider the case where advisers act as gatekeepers for comparison.

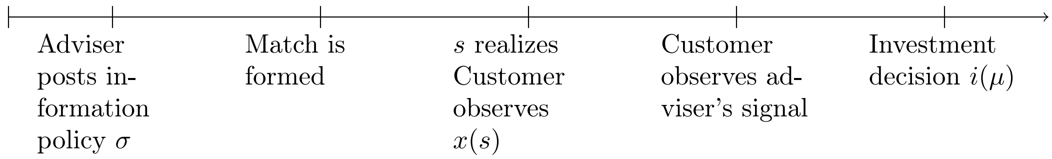

Figure 2: Timeline

chooses to invest through an adviser instead doing so on his own. To make our information market more interesting, we focus on the following parameter restrictions throughout the paper.

Assumption. *(A1)* $\left(\frac{1-F(\lambda)}{1-F(\underline{b}\lambda)}r - \left(1 - \frac{1-F(\lambda)}{1-F(\underline{b}\lambda)}\right)\bar{l}\right) > \alpha_c$. *(A2)* $\frac{F(\lambda)-F(\bar{b}\lambda)}{1-F(\lambda)}\underline{l} > \alpha_c$

Assumption (A1) guarantees that the asset is sufficiently valuable so that, after the cost α_c, all customers will invest even without receiving further information from advisers. This allows us to simplify our exposition. Assumption (A2) guarantees that all advisers can provide positive value to all customers, so that all matches are potentially profitable. Under these two assumptions, if we had a monopolistic adviser, she would choose an information policy that extracts all value from customers and leaves them with their autarky value $u^0(b, l)$. This highlights how our results are driven by competition.

Figure 2 summarizes the timeline of our setup.

Equilibrium Our equilibrium concept is the core of the assignment game: it requires that advisers are matched to customers in a stable way. That is, any competitive equilibrium must satisfy no surplus condition: if any two agents agree to match, they cannot be better off by matching with others or using different information policies.

Denote $v(b, l, u)$ as the utility of adviser l, conditional on providing a customer b with additional information value of u, which is given by:

$$v(b, l, \bar{u}) = \max_{\sigma} \tilde{V}(b, l, \sigma)$$
$$\tilde{U}(b, l, \sigma) - u^0(b, l) \geq \bar{u} \tag{3}$$

That is, by choosing an optimal policy σ, $v(b, l, \bar{u})$ is the payoff to an adviser conditional on leaving customer b with gain $\bar{u}$. This function represents the utility possibility frontier, and in equilibrium, agents will always achieve an allocation on this frontier. Each customer can invest on his own, regardless of which adviser he chooses. Hence, he will choose the adviser who provides him the most valuable information, which is the difference between the expected value under policy σ, $\tilde{U}(b, l, \sigma)$, and the expected value of self-directed trade, $u^0(b, l)$. We let $u(b, l, \bar{v})$ denote the maximum gain to customer b if he matches with an adviser l, who receives an utility of $\bar{v}$ within the match.

The equilibrium features of a payoff function for each agent, $U(b)$ and $V(l)$. Each agent takes the equilibrium utility of others as given and decides whom to match with. We let $H(b, l)$ denote the measure of advisers below l that match to customers with type below b on the product of $\{B \cup \emptyset\} \times \{L \cup \emptyset\}$. Let $H^A(l)$ and $H^C(b)$ denote the marginals of this distribution, respectively. That is, $H^A(l) \equiv H(\bar{b}, l)$ denotes the measure of advisers with type below l, and $H^C(b) \equiv H(b, \bar{l})$ denotes the measure of customers with type below b. The matching is feasible if and only if the marginals coincide with the distribution of advisers and customers:

$$H^C(b) = Q(b) \; H^A(l) = G(l). \tag{4}$$

This marginal condition is often called the market clearing condition.

Definition 1. An equilibrium consists of matching decisions $H(b, l)$ and payoff functions $U : B \to \mathbb{R}_+$ and $V : L \to \mathbb{R}_+$ that satisfy the following conditions.

(1) Advisers' optimality:

$$V(l) = \max_{\tilde{b} \in B \cup \{\emptyset\}} v(\tilde{b}, l, U(\tilde{b})) \tag{5}$$

and $(b, l) \in \operatorname{supp} H$ if and only if b maximizes (5) given l.

(2) Customers' optimality:

$$U(b) = \max_{\tilde{l} \in L \cup \{\emptyset\}} u(b, \tilde{l}, V(\tilde{l})) \tag{6}$$

and $(b, l) \in \operatorname{supp} H$ if and only if l maximizes (6) given b.

(3) Feasibility Condition in Equation (4).

The definition does not explicitly define the information provided in equilibrium. However, given equilibrium payoffs $V(l)$ and $U(b)$, it is implicitly defined by Problem (3) and it is characterized in detail below.

3 Characterization

We solve the model in two steps. First, we show that within a pair, without loss of generality, the optimal policy in Problem 3 can be characterized by a cutoff $\hat{s} \in [b\lambda, \lambda]$ such that the adviser recommends investing if and only if $s \geq \hat{s}$ (Proposition 1). As a result, despite allowing for a general message space, we obtain a simple functional form for $v(b, l, u)$. Based on this, we then analyze the sorting decision in this two-sided market and pin down the equilibrium payoff and information policies for all agents.

3.1 Information Policies

Our information structure implies that whenever a customer receives a bad signal ($x = 0$), he knows for sure that the asset is bad (i.e., $x = 0$ whenever $s < b\lambda$ which implies that $y(s, l) < 0$ with certainty). Then, investing is never optimal for any message received from an adviser. We can therefore limit attention to $x = 1$ when describing the optimal information policy. Conditional on $x = 1$, the customer learns that $s \geq b\lambda$, and her belief about s is given by the pdf $\frac{f(s)}{1-F(b\lambda)}$ with domain $[b\lambda, 1]$. We write μ_1 for the corresponding measure.

Reformulation Given any policy σ and realized message m a customer forms a posterior belief $\mu \in \Delta([0, 1])$. Since the message is random, σ induces a distribution P_σ on the space of posteriors beliefs. Based on Kamenica and Gentzkow (2011), Prop. 1, every distribution on the space of posterior beliefs $P \in \Delta(\Delta([b\lambda, 1]))$ can be induced by a signal, provided that P satisfies the Bayes-plausibility condition

$$E_P \mu = \mu_1. \tag{7}$$

We can therefore represent the adviser's problem 3 as maximizing over Bayes plausible distributions of posteriors after $x = 1$ is realized and identify any information policy with P.

We let $i(\mu) \in \{0, 1\}$ denote the investment decision conditional on $x = 1$ and posterior μ, so that $i = 1$ whenever the customer's expected value from investing is positive and $i = 0$ otherwise, i.e.,

$$i(\mu) = \mathbb{1}\{E_\mu [y(s, l)] - \alpha_c \geq 0\}. \tag{8}$$

Given information policy P, the probability of a customer investing is then given by $E_P [i(\mu)]$.

Hence, the payoff of an adviser under such policy is

$$\tilde{V}(b, l, P) = \alpha (1 - F(b\lambda)) E_P [i(\mu)].$$

That is, conditional on the customer receiving a positive private signal, which occurs with probability $1 - F(b\lambda)$, the adviser receives α with probability $E_P[i(\mu)]$. The expected utility of a customer under information policy P is then his expected payoff conditional on investing, which is

$$\tilde{U}(b, l, P) = (1 - F(b\lambda)) E_P [i(\mu)(E_\mu [y(s, l)] - \alpha_c)].$$

To summarize, the optimization problem of an adviser in Equation 3 can be written as

$$\begin{aligned} v(b, l, \bar{u}) \quad = \quad & \max_{P \in \Delta(\Delta([b\lambda, 1]))} \alpha (1 - F(b\lambda)) E_P [i(\mu)] \\ s.t. \quad & (1 - F(b\lambda)) E_P [i(\mu)(E_\mu [y(s, l)] - \alpha_c)] - u^0(b, l) \geq \bar{u} \\ & E_P \mu = \mu_1. \end{aligned} \tag{9}$$

Threshold Policy We now solve for the optimal policy and show that it is a threshold policy. The following Lemma shows that the optimal policy induces at most two posterior beliefs.

Lemma 1. *For any information policy P such that $\tilde{U}(b, l, P) > u^0(b, l)$, there exists an equivalent policy $\hat{P}$, which only puts weight*

$$p := P(\mu : i(\mu) = 1)$$

on the posterior $\mu_I = E_P[\mu | i(\mu) = 1]$ and $(1 - p)$ on the posterior $\mu_N = E_P[\mu | i(\mu) = 0]$.[21]

Based on Lemma 1, it is without loss of generality to put mass on two posteriors μ_I and μ_N, which are the posteriors conditional on investing and not investing. The Bayes

[21] In equilibrium, all customers who participate in the market will receive a value strictly higher than $u^0(b, l)$. Throughout this Section, we therefore focus attention on the case $\tilde{U}(b, l, P) > u^0(b, l)$. If $\tilde{U}(b, l, P) = u^0(b, l)$, the customer receives no information. P then simply puts probability one on the customer's prior.

Skip, too technical?

plausibility condition then becomes

$$p\mu_I + (1-p)\,\mu_N = \mu_1.$$

This immediately implies that μ_I and μ_N are absolutely continuous with respect to μ_1. Therefore, they both must admit densities, which, with slight abuse of notation, we denote with $\mu_I(s)$ and $\mu_N(s)$. For any given state s, the density $\mu_N(s)$ can be expressed as a function of p and $\mu_I(s)$ using the Bayes plausibility condition. This implies

$$\mu_N(s) = \frac{\frac{f(s)}{1-F(b\lambda)} - p\mu_I(s)}{1-p} \geq 0,$$

which is strictly decreasing in $\mu_I(s)$. Since the density function must be positive, this equation defines an upper bound for $\mu_I(s)$.

Hence, the adviser's problem can be further reduced to maximize over $\mu_I(s)$ and p, subject to

$$\mu_I(s) \in \left[0, \frac{1}{p}\frac{f(s)}{1-F(b\lambda)}\right].$$

Recall that an adviser only cares about the probability of investing (which is represented by p). However, for a given p and conditional on the investment taking place, a customer is strictly better off if an adviser puts a higher weight on the states that give a positive payoff, i.e. $s \geq \lambda$. This immediately suggests that for any optimal policy that solves Problem 3, $\mu_I(s)$ must hit its upper bound for any $s \geq \lambda$. Otherwise, there exists a Pareto improvement. However, whenever $s < \lambda$, customers receive a constant payoff of $-l$. Hence, on $[0, \lambda]$ it is irrelevant for customers whether $\mu_I(s)$ hits is upper bound. Only the total mass on $s < \lambda$ matters.

As a result, without loss of generality, the optimal policy can be characterized by a cutoff $\hat{s} \in [b\lambda, \lambda]$ such that, for any $s \geq \hat{s}$, the density $\mu_I(s)$ hits its upper bound and is zero otherwise:

$$\mu_I(s) = \begin{cases} \frac{1}{p}\frac{f(s)}{1-F(b\lambda)} & s \geq \hat{s} \\ 0 & s < \hat{s}. \end{cases}$$

The probability of investing is then given by $p = 1 - F(\hat{s})$. This leads to our Proposition 1.

Proposition 1. *Without loss of generality, the optimal policy can be characterized by a*

Skip more, need to get to results faster...?

cutoff $\hat{s} \in [b\lambda, \lambda]$ *such that the adviser recommends investing if and only if* $s \geq \hat{s}$.

The implementation of this policy is a message $m(s) = \mathbb{1}\{s \geq \hat{s}\}$. That is, when the adviser recommends (not) investing, the state s must be (lower) higher than a cutoff $\hat{s} \in [b\lambda, \lambda]$. In the following, we identify each information policy with the cutoff $\hat{s}$.

The higher the threshold $\hat{s}$, the better the information, as the customer is less likely to have a false positive. The gain of a customer who seeks advice about asset l with information policy $\hat{s}$, relative to self-directed trade, thus yields

$$\tilde{U}(b, l, \hat{s}) - u^0(b, l) \quad = \quad (F(\hat{s}) - F(b\lambda))\, l - \alpha_c(1 - F(\hat{s})). \tag{10}$$

The first term captures the additional gain from information $\hat{s}$. Specifically, under the cutoff rule $\hat{s} \in [b\lambda, \lambda]$, a customer b makes a bad investment and receives a negative payoff $-l$ with probability $F(\lambda) - F(\hat{s})$. If a customer would have traded based on his own signal, he would have made a mistake with probability $F(\lambda) - F(b\lambda)$. Hence, the added value of cutoff rule $\hat{s}$ to a customer b is given by $(F(\hat{s}) - F(b\lambda))l$, which represents the value of reducing the customer's mistakes. The second term represents the expected cost by investing through an adviser. That is, the customer needs to pay additional cost $\alpha_c > 0$ whenever he invests, which occurs with probability $1 - F(\hat{s})$.

The adviser's optimization problem can now be conveniently rewritten as

$$v(b, l, \bar{u}) \equiv \max_{\hat{s} \in [b\lambda, \lambda]} \alpha(1 - F(\hat{s})) \tag{11}$$
$$s.t. (F(\hat{s}) - F(b\lambda))l - \alpha_c(1 - F(\hat{s})) \geq \bar{u}.$$

An adviser is better off with a lower $\hat{s}$, as it implies a higher probability of investment. A customer however is worse off with a lower $\hat{s}$, as it implies a higher probability of making mistakes. The solution is then simply the threshold that provides customer b with utility $\bar{u}$.

3.2 Matching Patterns

We now analyze the matching pattern. First, as we can observe from Problem 3, conditional on providing $\bar{u}$ to her customer, an adviser would prefer to match with a less informed customer. Intuitively, fixing any policy $\hat{s}$, a less informed customer gains more as he has

State results upfront here

too detailed?

worse information when he trades the asset himself. Hence, conditional on promising gain $\bar{u}$ to a customer, an adviser is able to give worse information (i.e. a lower threshold $\hat{s}$) if the customer is less informed. Thus, she obtains a higher payoff. Second, clearly, the payoff to an adviser decreases with the promised utility of the customer (i.e., $v_u(b, l, \bar{u}) < 0$).

In equilibrium, an adviser takes customers' utilities $U(b)$ as given and chooses a customer optimally:

$$V(l) = \max_{\tilde{b}} v(\tilde{b}, l, U(\tilde{b})). \tag{12}$$

Given that all advisers prefer to match with less informed customers, competition implies that less informed customers must be promised a higher utility $U(b)$ in equilibrium. Otherwise, everyone wants to choose a lower b, which cannot occur in equilibrium. This thus suggests that the equilibrium gain $U(b)$ must decrease in b. Similarly, an adviser with more valuable expertise is more attractive from the viewpoint of customers and thus the equilibrium profit $V(l)$ must increase in l.

In equilibrium, it is the adviser that can offer the most valuable information that attracts the least informed customer. To understand the sorting outcome, observe that, fixing a policy $\hat{s}$ for any pair (l, b), the information gain of customer is given by $(F(\hat{s}) - F(b\lambda))\, l$. That is, there is complementarity between a higher l and a less informed customer for any given information. This suggests that it is less costly for an adviser with higher l to attract the least informed customer, since he can lie more (i.e., providing a lower $\hat{s}$) conditional on providing the same level of promised utility.

One can see this from the customers' viewpoint as well. Notice that, conditional on the threshold, an adviser does not care about the type of customers. Hence, in equilibrium, one can simply think about each adviser l posting a threshold $s^*(l)$. Since an adviser with higher l provides less valuable information, a customer is effectively facing a trade-off between choosing more informative information (a lower $\hat{s}$) vs. more valuable knowledge (a higher l). Formally, from equation (10), we have

$$U(b) = \max_{l} \left(F(s^*(l)) - F(b\lambda)\right) l - \alpha_c(1 - F(s^*(l))). \tag{13}$$

Observe that, conditional on $\hat{s}$, there is complementarity between an adviser with more

Up-front?

valuable expertise and a customer's demand for information. Specifically, less informed customers are more willing to accept a less precise policy in exchange for more valuable expertise. That is, if a more informed customer b' weakly prefers an adviser with higher l', then a less informed customer must strictly prefer such an adviser, which explains the sorting in equilibrium.[22]

Proposition 2 (Assortative Matching)**.** *A less informed customer matches to an adviser with more valuable expertise.*

3.3 Equilibrium Disclosure and Payoff

We now characterize the assignment function, equilibrium payoffs $U(b)$ and $V(l)$, and the information policy. Let $\ell : B \to L \cup \{\emptyset\}$ denote the type of adviser l from whom customer b receives information. Lemma 2 suggests that the assignment function $\ell(b)$ must be weakly decreasing.

Moreover, we know that (1) less informed customers must gain more by participating in the market (i.e., $U(b)$ decreases in b) and (2) advisers with more valuable expertise must obtain a higher equilibrium payoff (i.e., $V(l)$ increases in l). Hence, in equilibrium, there exists a highest active customer (denoted by b^*) and a lowest active adviser (denoted by l^*) such that $Q(b^*) = 1 - G(l^*)$.

For agents that are actively matched, the assignment function must then solve the following market-clearing condition: $\int_{\underline{b}}^{b} dQ(\tilde{b}) = \int_{\ell(b)}^{\bar{l}} dG(\tilde{l})$. That is, given any b, the measure of customers below b (LHS) equals the measure of advisers above $\ell(b)$ (RHS). This yields the differential equation

$$\frac{d\ell(b)}{db} = -\frac{dQ(b)}{dG(\ell(b))} \leq 0, \tag{14}$$

with initial condition $\ell(\underline{b}) = \bar{l}$. That is, the least informed customer must match with the adviser with the highest l.

Furthermore, from Equation (13), by envelope, we have

$$U'(b) = -F'(b\lambda)\lambda\ell(b) < 0. \tag{15}$$

[22] One can establish the single crossing property for $v(b, l, U(b))$ as well. That is, if an adviser with a lower l prefers a less informed customer, then an adviser with a higher l must strictly prefer such a customer. See Appendix.

too detailed still??

Equation (14) and (15) thus give a system of ODE. Given that the least informed customers must match with the adviser with highest $\bar{l}$, the boundary condition for the assignment function is then $\ell(\underline{b}) = \bar{l}$. The gain for the marginal buyer $U(b^*)$ then depends on whether customers or advisers are on the short side of the market. As is standard, if advisers are on the short side of the market, this customer will gain zero by participating in the market, i.e. $U(b^*) = 0$. However, if customers are on the short side of the market, the marginal adviser l^* is given by $Q(\bar{b}) = 1 - G(l^*)$. Since customers are scarce in this case, the marginal adviser must provide her customer the highest utility gain (i.e., full disclosure). Otherwise, other inactive advisers can outbid her by providing better information to attract her customer.

The equilibrium information policy can be solved accordingly. Specifically, let $s^*(b)$ denote the equilibrium information policy that an active customer b receives, which must provide $U(b)$ to customer b when he matches with adviser $\ell(b)$. This yields a differential equation for $s^*(b)$:

$$\frac{ds^*(b)}{db} = -\frac{\left(F(s^*(b)) - F(b\lambda)\right)\ell'(b)}{F'(s^*(b))(\ell(b) + \alpha_c)} > 0. \tag{16}$$

That is, a more informed customer must receive more precise information (i.e., a higher cutoff $\hat{s}$) in equilibrium, despite having lower gain $U(b)$. However, an adviser with more information sensitive asset matches with a less informed customer. Thus, she provides worse information (i.e. a lower cutoff) and earns a higher profit. Formally, each adviser's payoff is given by

$$V(l) = \alpha\left(1 - F(s^*(\ell^{-1}(l))\right),$$

where $\ell^{-1}(l)$ denotes the inverse of ℓ.

Proposition 3 then characterizes the assignment function and the information received by each customer directly. These two functions can then easily be mapped to the and equilibrium payoffs $U(b)$ and $V(l)$. To facilitate our characterization, we further define $s^0(b, l)$ as the policy under which the participation constraint of a customer with type b binds.[23]

[23] Our assumption (A2) guarantees the solution $s^0(b, l) \in [b\lambda, \lambda]$ exists. To see this, let $z(\hat{s}, l, b) \equiv (F(\hat{s}) - F(b\lambda))\, l - \alpha_c(1 - F(\hat{s})) = 0\ \forall b, \forall l$. Given that $z_s = (l + \alpha_c) > 0$, Assumption (A2) implies that $z(\lambda, \underline{l}, \bar{b}) > 0$ and that there exists $\hat{s} \in [b\lambda, \lambda]$ such that $z(\hat{s}, \underline{l}, \bar{b}) = 0$. Furthermore, given that $z_b < 0$, and $z_l > 0$, $s^0(b, l) < s^0(\bar{b}, \underline{l}) < \lambda\ \forall b, l$.

That is,

$$s^0(b,l) \equiv \min\{\hat{s} \in [b\lambda, \lambda] : \tilde{U}(b,l,\hat{s}) - u^0(b,l) \geq 0\}. \tag{17}$$

The IR constraint then simply requires that $s^*(b) \geq s^0(b, \ell(b))\ \forall b$. The equilibrium characterization is summarized in Proposition 3.

Proposition 3. *The equilibrium is characterized by a marginal customer b^*, an assignment function $\ell(b):\ B \to L \cup \{\emptyset\}$, and an information policy $s^*(b):\ [0, b^*] \to [0, \lambda]$ such that*

(1) The marginal customer is given by $b^* = \begin{cases} Q^{-1}(1) & if\ Q(\bar{b}) > 1 \\ \bar{b} & if\ Q(\bar{b}) \leq 1. \end{cases}$

(2) For any $b \leq b^*$, $\ell(b)$, and $s^*(b)$ solve the system of differential equations (14) and (16) with initial conditions $\ell(\underline{b}) = \bar{l}$ and

$$s^*(b^*) = \begin{cases} s^0(b^*, \ell(b^*)) & if\ Q(\bar{b}) > 1 \\ \lambda & if\ Q(\bar{b}) < 1 \\ \zeta & if\ Q(\bar{b}) = 1, \end{cases}$$

where $\zeta \in [s^0(b^*, \ell(b^*)), \lambda]$. For $b > b^*$, customers remain inactive, i.e., $\ell(b) = \emptyset$.

Note that the initial condition pinning down $s^*(b^*)$, which represents the information received by the marginal customer, depends on $U(b^*)$ and thus whether customers or advisers are on the short side of the market. Hence, the marginal buyer receives his autarky value (full disclosure) when advisers are on the short (long) side of the market. Lastly, when $Q(\bar{b}) = 1$, the customers' share is not uniquely pinned down. Thus, any information between $s^0(b^*, \ell(b^*))$ and λ can be supported as an equilibrium. This explains condition (2) in the Proposition.

In summary, as information is the only currency, each adviser attracts customers by promising a certain information value. In equilibrium, each adviser must prefer attracting her customer at his equilibrium pay level (in terms of information) to attracting any other customers at their pay level.

The equilibrium information policy is illustrated in Figure 3.

fig

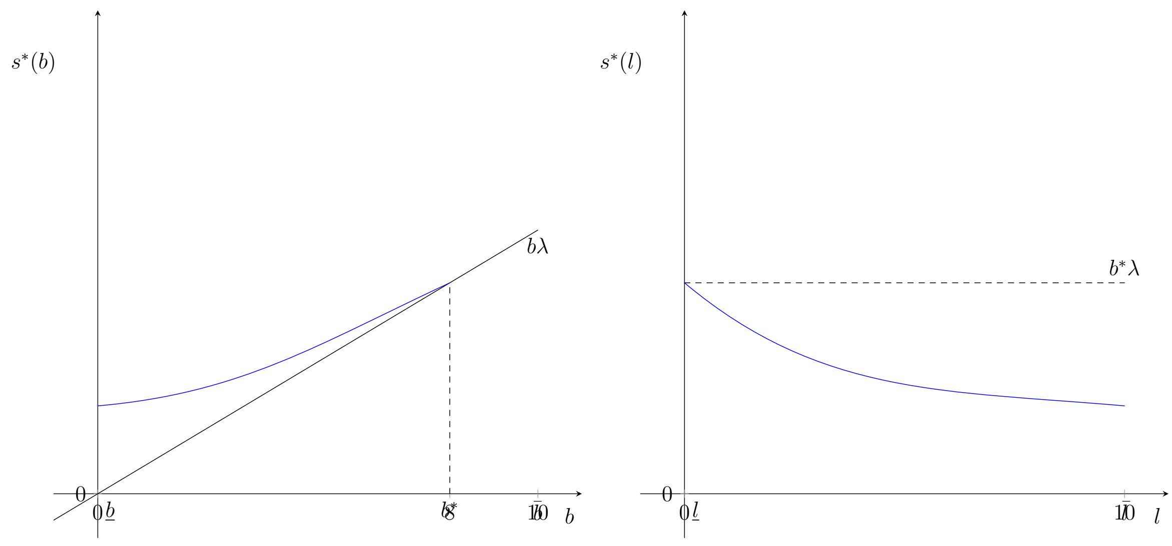

Figure 3: Equilibrium Information Policy with Scarce Advisers ($\alpha_0 > 0$,$\alpha_c = 0$)

3.4 Information Distortion and Welfare

Our equilibrium definition guarantees that the equilibrium allocation is always on the Pareto frontier. However, it is well known that when utilities are not fully transferable, the equilibrium outcome of matching markets might not maximize total surplus. To see this explicitly in our model, consider a social planner who designs the allocation function and the policy function $s^e(b,l)$ to maximize the pair-wise surplus, taking into account the gains for agents on both sides of the market.

Within each pair, the information policy $s^e(b,l)$ thus solves

$$\Omega^e(b,l) = \max_{s\in[b\lambda,\lambda]} \alpha_0(1-F(s)) + (F(s)-F(b\lambda))l. \tag{18}$$

Since α_c only affects the transfer within the pair, it does not affect the pair-wise surplus. The solution is then simply given by

$$s^e(b,l) = \begin{cases} \lambda & \text{if } l \geq \alpha_0 \\ b\lambda & \text{if } l < \alpha_0. \end{cases} \tag{19}$$

That is, when the benefit of the adviser α_0 is strictly smaller than the loss l, the information

Not that interesting?

policy that maximizes the joint surplus is full disclosure. If α_0 is larger than l, no disclosure is optimal. Hence, the surplus function yields:

$$\Omega^e(b,l) = \begin{cases} \alpha_0(1-F(\lambda)) + (F(\lambda)-F(b\lambda))l & \text{if } l \geq \alpha_0 \\ \alpha_0\,(1-F(b\lambda)) & \text{if } l < \alpha_0. \end{cases}$$

For either case, the surplus function is sub-modular, i.e., $\Omega^e_{bl} \leq 0$, so the maximal surplus can be achieved by negative sorting. Thus, the assignment function is the same under the decentralized equilibrium and the planner problem.

However, the information policy generally is distorted in the decentralized equilibrium, as the information policy generally lies in between full disclosure and no disclosure. In other words, the decentralized equilibrium can feature either too little information (when $\alpha_0 < l$) or too much (when $\alpha_0 > l$). Furthermore, given that the sorting pattern in the decentralized market will be the same as the solution to the social planner problem (i.e., the same the assignment function $\ell(b)$ and same marginal customer b^*), the welfare loss is then summarized by the distortion on the information policy relative to the planner's solution, which yields $L = \int_{\underline{b}}^{b^*} |s^e(b,\ell(b)) - s^*(b)|\ell(b)dQ(b)$.

4 Comparative Statics

Composition Effects One unique feature of our market with two-sided heterogeneity is that competition is about *which* customer to serve and *which* adviser to ask for advice. In equilibrium, all agents are compensated so that it is indeed optimal for them to match to their counterparty instead of others. Hence, the distribution of other customers and advisers, which represents the potential competitors, is the key determinant of information provision. In fact, if all advisers were the same, the equilibrium will always feature one information policy and the composition of distribution does not matter.[24]

To further explore the effects of heterogeneity, we now consider comparative statics on the distribution of customer types and analyze how the composition affects the quality of information in equilibrium. We consider a change in the shape of distribution $Q(b)$ within the

[24] Similar to the logic that there exists one price that clears the market, there will be one information policy that makes sure the marginal customer is willing to participate. That is, if advisers (customers) are at the short side, $s^*(b) = b^*\lambda$ ($s^*(b) = \lambda$).

Too many explanations and mechanics

active matches, fixing the marginal type b^*. Formally, consider two distributions $Q_1(b)$ and $Q_2(b)$: both of them have the same range B, but $Q_1(b)$ first-order stochastically dominates $Q_2(b)$, i.e., $Q_1(b) \leq Q_2(b)\ \forall b \in B$. Intuitively, $Q_2(b)$ has more customers who are relatively uninformed.

To see the effect on information provision more clearly, it is convenient to look at the information policy as a function of adviser l. The differential equation for information policies in (16) can be rewritten as

$$\frac{ds^*(l)}{dl} = -\frac{F(s^*(l)) - F(\ell^{-1}(l)\lambda)}{F'(s^*(l))(l+\alpha_c)}, \tag{20}$$

where the initial condition for the marginal adviser l^* is given by $s^0(b^*, \ell(b^*))$, which is the same for given fixed b^*.

When there are more customers who are less informed, any given adviser l must match with a customer with a (weakly) lower b. Thus, $\ell_2^{-1}(l) \leq \ell_1^{-1}(l)$, where $\ell_j^{-1}(l)$ represents the assignment under distribution $Q_j(b)$. Hence, one can show via a comparison theorem, that the solution to Equation (20) under $Q_2(b)$ must be weakly lower. That is, when the economy has a larger portion of less informed customers, information quality is worse *for all customers*, not just the ones who are less informed. This is in contrast to settings with either a monopoly or with homogeneous advisers. In those environments, the aggregate distribution will not affect the information provided by an adviser l, so there is no composition effect.

This result highlights how competitive forces shape information provision. Advisers compete for less informed customers, since they are the ones who value information most and are also easier to deceive. As a result, these customers earn rents. Now, fixing a customer $\hat{b}$, when all other customers above him become more informed, meaning that his other competitors becomes less attractive, such a customer then must earn a higher surplus, thus receiving better information.

One can also consider a similar exercise by changing the shape of the distribution of adviser types $G(l)$. For example, suppose $G_1(l)$ first-order stochastically dominates $G_2(l)$. That is, there are fewer advisers with valuable information under $G_2(l)$. The same intuition as before applies: an adviser with expertise in higher l becomes more scarce and as a result, she must earn a higher profit, meaning worse information for customers. The comparative statics are summarized in Lemma 2.

stress more? gets drowned ...

Lemma 2. *(1) Consider two distributions with the same range and* $Q_1(b) \leq Q_2(b)$*, then* $s^*(l)$ *is weakly lower under* $Q_2(b)$*. (2) Consider two distributions with the same range and* $G_1(l) \leq G_2(l)$*, then* $s^*(l)$ *is weakly lower under* $G_2(l)$*.*

Advisers Fees Intuitively, one might expect that a higher fee worsens the information provided in equilibrium, because it implies a more severe conflict of interest. This argument is inaccurate under a competitive market. From equation (13), we see that commission fee α_c and information are substitutes from the viewpoint of customers. Hence, a customer must be compensated with better information whenever there is a higher commission fee α_c. In particular, given any commission fees $\alpha_c > 0$, the schedule of equilibrium information $s^*(l; \alpha_c)$ is pinned down in a way so that all customers must choose their advisers $\ell(b)$. This suggests that the higher the commission fees, the better the information being provided in equilibrium.

Lemma 3. *Let* $s^*(l; \alpha_c)$ *denote the equilibrium information under commission* α_c*. The higher the fee, the better the information: for any* $\alpha_c' > \alpha_c$*,* $s^*(l; \alpha_c') > s^*(l; \alpha_c)$*.*

Scaling Effects We now turn to another source of the conflict of interest: α_0 which captures other potential gains of advisers that are not costly to customers. From Equation (20), we can immediately show that the information policy is independent of α_0. In other words, a higher α_0 only scales up advisers' profits. Similarly, when $\alpha_c = 0$, if all assets l are scaled by a factor $z > 0$ (i.e., $l_z = zl$), the customer of each adviser remains the same. One can then show that the equilibrium information policy remains the same as well, i.e., $s_z^*(l_z) = s^*(\frac{l_z}{z})$. In summary, when α_0 increases (or when the loss of each asset increases), the payoff to all advisers (or to customers) simply increases by the same amount.

Lemma 4. *For any* $z > 0$*, (1) If* $\alpha_c = 0$ *and* $l_z \equiv zl$*, the information policy stays the same. We have* $s_z^*(b) = s^*(b)$*,* $U_z(b) = zU(b)$*, and* $V_z(l_z) = V(l)$*. (2) If* $\alpha_z \equiv z\alpha$*, the information policy stays the same and we have* $s_z^*(b) = s^*(b)$*,* $U_z(b) = U(b)$*, and* $V_z(l) = zV(l)$*.*

5 Information vs. Price Competition

Advisers in our main model are compensated by a commission set by the mutual funds (e.g., loads and 12b-1 fees), but they are, however, competing on information provision in order to

Interesting?

Very late?

attract customers. As discussed in a White House report (CEA (2015)), other common types of advisers in the industry are called "fee-only" advisers, who are compensated by flat fees (such as hourly rates). That is, their compensation does not directly depend on customers' investment decisions. The competition among this type of advisers is then the standard price competition.

To understand how these different compensations affect customers, we now compare our result to an environment with fee-only advisers. The difference between fee-only and commission-based advisers thus represents the cost of a compensation structure that involves conflicts of interests.

The conventional wisdom is that customers should look for someone who is paid as a "Fee-Only" adviser to eliminate as many conflicts of interest as possible. Perhaps surprisingly, we show that despite information being distorted in our main model, customer welfare is not harmed as long as advisers are on the short side of the market. Our result thus invalidates the conventional argument that banning conflicted payment improves customer welfare.

5.1 Price Competition for Fee-Only advisers

In an environment where fees can be set by advisers, the fees will be determined in equilibrium together with the information policy. Suppose an adviser can post fees γ in addition to the information policy σ. As before, customers make their matching decisions to maximize their expected utility of the match. This setting can then be solved as the classical assignment model with transferable utilities, where the optimal information policy simply maximizes the pairwise surplus.[25] Since we are in the case that $\alpha_0 = 0$, the optimal policy $s^e(b, l)$ is full disclosure as shown in equation (19).

Hence, the pairwise surplus is given by $\Omega^e(b, l) = (F(\lambda) - F(b\lambda))\, l$. Note that, since $\Omega_{bl} < 0$, the equilibrium still features negative sorting. The utility of a customer is

$$U^{FB}(b) \equiv \max_{l} \Omega^e(b, l) - V(l),$$

[25]For example, Terviö (2008).

where, by the envelope theorem,

$$\frac{dU^{FB}(b)}{db} = -F'(b\lambda)\lambda\ell(b). \tag{21}$$

Equation (21) shows that for any given assignment function $\ell(b)$, the utility of a customer $U(b)$ is pinned down up to a constant. Specifically, given the utility level of the marginal customer $U^{FB}(b^*)$, customers' utilities and thus total payments are pinned down uniquely.

The level of the marginal customer is pinned down in the standard way: when advisers are scarce, the marginal customer must gain zero after the fee. That is, he pays $(F(\lambda) - F(b\lambda))\ell(b^*)$ for full disclosure. However, when customers are scarce, the marginal customer receives the whole surplus. That is, he receives full disclosure for free.

Lemma 5. *With competitive fees, all advisers provide full disclosure, and the gain of customers in equilibrium (after fees) is characterized by equation (21) with initial condition:*

$$U^{FB}(b^*) = \begin{cases} 0 & if\ Q(\bar{b}) > 1 \\ \left(F(\lambda) - F(b\lambda)\right)\ell(b^*) & if\ Q(\bar{b}) < 1 \\ U^0(\bar{b}) & if\ Q(\bar{b}) = 1 \end{cases}$$

where $b^* = \min\{Q^{-1}(1), \bar{b}\}$, $1 - G(\ell(b^*)) = Q(b^*)$, *and* $U^0(\bar{b}) \in [0, \left(F(\lambda) - F(\bar{b}\lambda)\right)\underline{l}\,]$.

5.2 Irrelevance of Fee Structures

When fees cannot be freely adjusted, as shown in our baseline model, advisers compete on another dimension: information. We now use our results with conflicted fees $\alpha_c > 0$ to answer two questions: (1) What is the cost to customers under conflicted payments? (2) How does the fee affect information quality?

Equivalence of Customers' Utilities Recall that the utility of a customer in equilibrium is characterized by Equation (15), $U'(b) = F'(b\lambda)\lambda\ell(b)$, which is in fact identical to our competitive benchmark in equation (21). As a result, the difference (if any) across these two environments is captured by the utility level of the marginal customer, which is summarized by Proposition 4:

earlier?

Proposition 4. *For any conflicted payment α_c that satisfies A1 and A2, the utility of each customer is given by*

$$U(b) = \begin{cases} U^{FB}(b) & if\ Q(\bar{b}) \geq 1 \\ U^{FB}(b) - \alpha_c(1 - F(\lambda)) & if\ Q(\bar{b}) > 1, \end{cases}$$

where $U^{FB}(b)$ denotes customer utilities under the competitive benchmark.

The proposition establishes a connection between competition through fees and competition through information provision. To see this more clearly, consider any given commission fee α_c and information policy $s^*(l)$. The equilibrium gain of customer b can rewritten as:

$$U(b) = \max_{l} -F(b\lambda)l + \tau(l),$$

where $\tau(l) \equiv F(s^*(l))l - \alpha_c(1 - F(s^*(l)))$. That is, $\tau(l)$ can be interpreted as the net benefit provided by adviser l, which includes the value of information minus the fees.

In other words, the net benefit can take the form of either an informational or monetary reward. Nevertheless, any schedule $\tau(l)$ that guarantees that all customers optimally match to their own adviser $\ell(b)$ must satisfy the following FOC of customers' optimization:

$$-F'(b\lambda)\lambda\ell(b) = \tau'(\ell(b)).$$

Hence, the schedule $\tau(l)$ is uniquely pinned down up to a constant.[26] In fact, this argument holds more generally for any exogenous fee structure that induces the same assignment function $\ell(b)$.[27] In this case, consumer utilities can only be different up to a constant compared to the competitive benchmark, which is captured by the difference of the utility of the marginal customer $U(b^*)$.

The utility of the marginal customer is pinned down as usual: when the customers are in the short side of the market, the adviser of the marginal customer $\ell(b^*)$ must provide him the highest payoff possible. In this case, it means full disclosure minus the exogenous fees

[26] This argument is analogous to revenue equivalence result in the standard pricing problem where customers have quasi-linear utility functions.

[27] In Appendix, we further allow for fees depending on types and provide condition under which $(\alpha_c(l), \gamma(l))$ leads to negative sorting in equilibrium.

$\alpha_c(1-F(\lambda))$. Since a marginal customer would have received full disclosure for free under the benchmark, this implies that the net benefit of all customers will be lower by exactly the same amount, $\alpha_c(1-F(\lambda))$.

However, if advisers are on the short side, the marginal customer always receives zero gain. This is true as long as the fees are small enough so that the marginal customer can still gain by entering the market (i.e.,under A1 and A2). Hence, *all* customers must obtain the same net benefit $\tau(\ell(b))$ across the two settings: full disclosure but a higher fee or distorted information with a lower payment. In other words, how fees are set is in fact irrelevant for customers' utilities.

Corollary 1 (Irrelevance)**.** *If advisers are on the short side, the net benefit received by all customers – a combination of fees and information – is the same under any fee structure that induces the same assignment function.*

However, since the total surplus is strictly higher under the competitive benchmark (i.e., full disclosure), the profit of advisers must be higher, because customers' utilities are weakly lower. In other words, advisers are the ones who bear the cost of information distortions under fixed fees. Combining with Lemma 3, the higher fee leads to a lower information distortion, thus a higher joint surplus and a higher profit to advisers.

Corollary 2. *If advisers are on the short side, the information distortion decreases the profit of advisers, but not of consumers. A higher fee α_c leads to higher profits for all advisers.*

6 Extensions

6.1 False Negatives

In the environment where customers may have false positives, we establish that less informed customers must match with advisers with more valuable expertise (Proposition 2).We now show that our main result holds for the case with false negatives as well.

Formally, the signal structure for customers remains the same: $x = \mathbf{1}\left\{s \geq b\lambda\right\}$. However, instead of having $b \leq 1$ as in the false positive case, we now assume that $b \in [\underline{b}, \bar{b}]$, where $\underline{b} > 1$ and $\bar{b} \leq \frac{1}{\lambda}$. That is, when $x = 1$, the customer knows that the asset pays r for sure. However, when customers observe a negative signal $x = 0$, the assets can give a positive return (i.e., false negatives). Note that with this assumption, higher b means the customer is more likely to have a false negative and thus is "less informed".

An asset is characterized by the return and the downside risk, denoted by (r, l). Since the information value is now about the upside, we analyze the case where all assets have the same downside risk l but differ in their upside return r with distribution $G(r)$. In other words, the adviser who knows an asset that pays a higher return r is the type who has more valuable expertise.

Analogous to our Assumption (A1), we assume that customers never invest upon receiving a negative signal ($x = 0$). That is, $r\left(F\left(b\lambda\right) - F\left(\lambda\right)\right) - lF\left(\lambda\right) < 0 \forall b$. Under this assumption, the customer's outside option is $u^0\left(b, r\right) = \left(1 - F\left(b\lambda\right)\right) r$. That is, he only invests if he receives the signal $x = 1$ and then receives r for sure. Intuitively, this captures the idea that, without any further information, customers will forego some investment opportunities.

As before, we can show that the information policy of the adviser can again be described with a threshold $\hat{s} \leq \lambda$. Equation (10), which represents the gain of a customer who seeks advice about asset r with information policy $\hat{s}$, relative to self-directed trade, thus yields

$$U\left(b, r, \hat{s}\right) - u^0(b, r) = r\left(F\left(b\lambda\right) - F\left(\lambda\right)\right) - l\left(F\left(\lambda\right) - F\left(\hat{s}\right)\right) - (1 - F(\hat{s}))\alpha_c. \tag{22}$$

Advice now has two roles: on the one hand, the adviser helps a customer identify positive investment opportunities. That is, he provides value by ensuring that the agent invests whenever the asset pays r. Specifically, without advice, customers would have foregone investment opportunities, which is captured by the first term $r\left(F\left(b\lambda\right) - F\left(\lambda\right)\right)$. On the other hand, to extract value from the customer, the adviser also sometimes oversells, which is represented by the second term, $l\left(F\left(\lambda\right) - F\left(\hat{s}\right)\right)$. The last term, as before simply captures the additional cost of the commission compared to the direct-trading channel.

By the same logic as before, advisers with more valuable expertise (i.e., a higher r in this case) must earn a higher payoff and thus must give worse information in equilibrium. That is, $s^*(r)$ must decrease in r. From the viewpoint of customers, choosing an adviser is then effectively trading off between the upside return r and information quality $s^*(r)$, where

$$U^*(b) = \max_r \left\{U(b, r, s^*(r)) - u^0(b, r)\right\}. \tag{23}$$

Observe from Equation (22) that there is complementarity between the investors' information needs (i.e., higher b in the case of false negative) and the value of expertise r. That is, for any $r' > r$, if customer b prefers an adviser with a more valuable expertise r', then a less informed customer $b' > b$ must prefer adviser r' as well.[28] This thus shows that Proposition 2 remains intact. Same as before, less informed customers (a higher b in this case) must match with advisers with more valuable expertise. The assignment function is then given by the following market-clearing condition: $\int_b^{\bar{b}} dQ(\tilde{b}) = \int_{\ell(b)}^{\bar{r}} dG(\tilde{r})$. which yields

$$\frac{d\ell(b)}{db} = \frac{dQ(b)}{dG(\ell(b))}. \tag{24}$$

Given the sorting, $U^*(b)$ can then be pined down as before. Specifically, a less informed customer must gain more. Furthermore, as we can observe from from Equation (23), the marginal utility gain only depends on the allocation. By the envelope theorem,

$$\frac{dU^*(b)}{db} = F'(b\lambda)\lambda\ell(b), \tag{25}$$

where $\ell(b)$ here represents the type of adviser r that a customer b matches with in equilibrium, analogous to Equation (15).

Thus, the characterization is the same as before, except that $U^*(b)$ and $\ell(b)$ are now characterized by Equation (25) and (24) instead. The key result in our basic model remains the same: less informed customers (the ones who are more likely to forego investment opportunities) match with advisers with more valuable expertise and thus receive worse information policy.

6.2 Assets with Multidimensional Heterogeneity

For both information environments, we so far assume that assets only differ in one dimension. That is, downside risk (l) for the false positive case and upside return (r) for the false negative case. We now show that our result can be easily extended for more general asset payoffs.

For both cases, the dimension of heterogeneity is designed to capture the value of information. Indeed, in the case for false positives, one can see that, according to Equation (10),

[28] Formally, for any $b' > b$ and $r' > r$, if $U(b, r, s^*(r')) - u^0(b, r') \geq \left[U(b, r, s^*(r)) - u^0(b, r)\right]$, then $U(b', r, s^*(r')) - u^0(b', r') > U(b', r, s^*(r)) - u^0(b', r)$.

only the downside risk l is relevant for customers' value, but not the upside return r. This is because choosing an adviser is not about choosing which asset to invest in[29] but purely about maximizing the information gain, which is only about reducing the downside risk in this case. Thus, as long as the underlying returns satisfy Assumption A1, the upside return is payoff irrelevant, and thus all the characterization remains the same.

In the case for false negatives, as shown in Equation (22), one can see that both return r and l matter for customers' utilities. However, what matters for the complementarity is again also the upside return r. In other words, the sorting result will not be affected even when we allow for heterogeneous downside risk. Thus, Equation (25) remains the same, and thus so do customers' utilities $U^*(b)$.

However, since a higher l suggests that making mistakes is more costly, the thresholds strategy will then be affected by the value of l. Specifically, let (r, l_r) denote the payoff of the asset r.[30] Given $U^*(b)$, the information policy for each customer then solves:

$$U^*(b) = r\left(F\left(b\lambda\right) - F\left(\lambda\right)\right) - l_r\left(F\left(\lambda\right) - F\left(s^*(b)\right)\right) - (1 - F(s^*(b)))\alpha_c.$$

That is, conditional on giving $U^*(b)$ to customers b, $s^*(b)$ is higher (lower) when the asset r has a higher (lower) downside risk l_r. Intuitively, if the asset has a higher (lower) downside risk l_r, overselling is more (less) costly, and thus the adviser r must give better (worse) information to compensate the customer.

6.3 Customers with Heterogeneous Wealth

So far, we focus on the case where customers differ in their informedness. Our framework nevertheless can be applied to different notions of heterogeneity. For example, another important dimension of heterogeneity in the financial market is the wealth of customers. To capture this, assume that, for any given asset l, customers differ in the amount of capital that they can invest, which we denote with w, but they have the same level of information.

[29] Recall that customers can always invest the asset themselves.

[30] Again, we allow for any arbitrary pair (r, l_r) such that customers never invest upon receiving a negative signal: $r\left(F\left(b\lambda\right) - F\left(\lambda\right)\right) - l_r F\left(\lambda\right) < 0\ \forall b$.

The Pareto frontier within each match is then given by

$$v(w,l,\bar{u}) = \max_{\hat{s}\in[b\lambda,\lambda]} \alpha_c (1-F(\hat{s}))\, w \tag{26}$$
$$s.t. \quad w\{(F(\hat{s})-F(b\lambda))\, l - \alpha_c(1-F(\hat{s}))\} = \bar{u}$$

Clearly, all advisers would like to attract customers with higher wealth (i.e., $v_w > 0$). Intuitively, a customer with a higher wealth is the one who has a higher demand for information. Thus, fixing any information quality, wealthy customers benefit most as they have more skin in the game. This thus suggests that they must gain more in equilibrium: $U^*(w)$ must increase in w.

Moreover, observe that there is complementary between information demand (w) and the value of expertise (l). Hence, the adviser with the most valuable expertise attracts the wealthiest customer. Given the sorting, the allocation $\ell(w)$ and customer's utilities $U^*(w)$ can be pinned down as before. The information received by customer w in equilibrium is then such that customers receive value $U^*(w)$ when he matches with $\ell(w)$.[31]

Proposition 5. *When customers differ in wealth, a wealthier customer gains more by participating in the market for advice, and he is matched with an adviser with more valuable expertise.*

7 Empirical Implications

Empirically, it has been well documented that conflicted advice leads to lower investment returns (e.g., Bergstresser et al. (2009), Chalmers and Reuter (2010) and Hoechle et al. (2013)).[32] If customers anticipate that advice destroys value, since it leads to lower returns, they should never seek advice in the first place. The difficulty of reconciling the evidence with customers who behave rationally has lead researchers to attribute the outcomes in the market for advice to customer naiveté. As we show now, our model can generate lower investment returns conditional on receiving advice, even though customers are perfectly rational and

[31] See detailed derivation in Appendix B. Note that one key difference here is that, conditional on information policy $\hat{s}$, advisers do care the type of customer. Thus, each adviser will post a menu of information policy that is conditional on the type of customers.

[32] See the detailed summary in White House report (CEA (2015)).

choose advisers and investments optimally. That is, naiveté is not necessary to rationalize the empirical findings.

The key insight here is that investment returns do not fully reflect the customer's value from participating in the market for advice. What matters is instead is how much advice can improve on the investment decisions they would have been making on their own. In our model, less informed customers rationally accept advice that yields lower returns, because it improves on what they could achieve by investing on their own.

To see this clearly, consider the expected return from investing in asset l after receiving advice. This return is

$$R\left(l\right)=\frac{(1-F(\lambda))r-\left(F(\lambda)-F(s^{*}\left(l\right))\right)l}{(1-F(s^{*}\left(l\right)))}-\alpha_{c}. \tag{27}$$

Compare this to the return of a customer who does not receive any advice at all (i.e., $b\geq b^{*}$in Proposition 3) and invests in the same asset. For this customer, the return is

$$\hat{R}\left(l\right)=\frac{1-F\left(\lambda\right)r-\left(F\left(\lambda\right)-F\left(b\lambda\right)\right)l}{1-F\left(b\lambda\right)}.$$

This return is necessarily higher.[33] However, this does not mean advice destroys value. The gain from the customer who invests in asset l after receiving advice is strictly positive while the gain from the one investing without advice is exactly zero.[34] Thus, while advice seems to lower returns, this does not mean it actually destroys value.[35]

In the case of false negatives in Section 6.1, this result is even more stark. There, if a customer invests without receiving any advice, his return is always r. When he receives advice however, the adviser sometimes recommends him to invest in an asset with a negative return, so conditional on receiving advice, the return is always lower. However, advice still improves value for customers because without advice, they would reject too many investments that have positive returns.

A common way to measure the cost of conflicted advice is to compare the return across fee-only and commission-based advisers. The difference in these returns is often interpreted

[33]From Proposition 3, we can see that $s^{*}\left(l\right)\leq b^{*}\lambda$, which implies $R\left(l\right)\leq\hat{R}\left(l\right)$.

[34]Again this follows from Proposition 3.

[35]In Section 6.2 we discussed how our model can be generalized when asset quality is multidimensional and each asset can be characterized by a pair (r,l). All results discussed in this section extend to this case.

as cost of conflicted payment, as discussed in CEA (2015). Our theory establishes that this interpretation is misleading. As we have shown in Proposition 4, if we compare a setting where advisers charge only upfront fees to one where they receive commissions, customer utility will be the same. Intuitively, customers can pay for advice either by higher fees or by receiving worse information,[36] and the expected payoff of seeking advice is the same, i.e.,

$$(1 - F(s^*(l)))\, R(l) = (1 - F(\lambda))\, r - \gamma(l).$$

The difference across these two only reflects the information quality, but customers can simply pay a higher fee for better information. Moreover, our theory also shows that the realized return "before" fees, $R(s, l) + \alpha_c$, is the more precise measure of information, and not "after" fees.

Finally, our model delivers new testable predictions: First, according to Lemma 3, the threshold must be higher (better information) when there is a higher fee so that customers' utilities remain the same: $U^*(b) = (1 - F(s^*(r, l)))\, R(r, l)$. Hence, our model predicts that an increase in conflicted fees leads to lower probability of investment but a higher realized return $R(r, l)$. Second, less informed customers or more wealthy customers are the ones who have higher demand for advice and are thus more likely to trade through advisers, consistent with the findings in Chalmers and Reuter (2010). More importantly, while less informed customers do receive worse information (i.e., it may appear that they are being exploited), they are actually the ones who gain most by trading through advisers since their outside option is much worse.[37]

8 Conclusion

The paper studies competitive markets for information in which advisers are subject to conflicts of interests. In contrast to existing models that assume naïve customers, our model analyzes the quality of information in a competitive market with rational customers with different levels of sophistication. We are thus able to establish new insights and predictions

[36] The exercise here assumes that everything else is equal across these two types of structure. In particular, both $G(l)$ and $Q(b)$ are the same.

[37] As emphasized in Chalmers and Reuter (2010), in order to measure the value of advice, one would need to take into account investors' utilities when they invest on their own.

regarding customers' welfare and asset returns. Specifically, we show that it is the underlying distribution of financial literacy that determines the consumers' welfare. When advisers are scarce, the fee structure of advisers is irrelevant for the welfare of consumers. We further show that the existing empirical measures of conflicted advice can be misleading if they fail to take into account the selection effect of the characteristics of funds and customers.

References

Admati, A. R. and P. Pfleiderer (1988). Selling and trading on information in financial markets. *The American Economic Review 78*(2), 96–103.

Aliprantis, C. D. and K. Border (2006). *Infinite dimensional analysis: a hitchhiker's guide*. Springer Science & Business Media.

Anagol, S., S. A. Cole, and S. Sarkar (2013). Understanding the advice of commissions-motivated agents: Evidence from the Indian life insurance market. *Harvard Business School Finance Working Paper* (12-055).

Au, P. H. and K. Kawai (2015). Competition in Information Disclosure. *Available at SSRN 2705326*.

Becker, G. S. (1973). A theory of marriage: Part I. *The Journal of Political Economy*, 813–846.

Bergstresser, D., J. M. Chalmers, and P. Tufano (2009). Assessing the costs and benefits of brokers in the mutual fund industry. *Review of financial studies 22*(10), 4129–4156.

Board, S. and J. Lu (2015). *Competitive information disclosure in search markets*. Unpublished.

CEA (2015). The Effects of Conflicted Investment Advice on Retirement Savings. Technical report, Executive Office of the President.

Chalmers, J. and J. Reuter (2010). What is the Impact of Financial Advisors on Retirement Portfolio Choices & Outcomes? Technical report, National Bureau of Economic Research.

Chiappori, P.-A. and P. Reny (2006). Matching to share risk. *manuscript http://home. uchicago. edu/˜ preny/papers/matching-05-05-06. pdf*.

Dantzig, G. B. and A. Wald (1951). On the fundamental lemma of Neyman and Pearson. *The Annals of Mathematical Statistics 22*(1), 87–93.

Dulleck, U. and R. Kerschbamer (2006). On doctors, mechanics, and computer specialists: The economics of credence goods. *Journal of Economic literature 44*(1), 5–42.

García, D. and F. Sangiorgi (2011). Information sales and strategic trading. *Review of Financial Studies 24*(9), 3069–3104.

Gennaioli, N., A. Shleifer, and R. Vishny (2015). Money doctors. *The Journal of Finance 70*(1), 91–114.

Gentzkow, M. and E. Kamenica (2011). Competition in persuasion. Technical report, National Bureau of Economic Research.

Hoechle, D., S. Ruenzi, N. Schaub, and M. Schmid (2013). Don't answer the phone: financial advice and individual investors' performance. Technical report, Working paper.

Inderst, R. and M. Ottaviani (2009). Misselling through agents. *The American Economic Review 99*(3), 883–908.

Inderst, R. and M. Ottaviani (2012a). Competition through commissions and kickbacks. *The American Economic Review 102*(2), 780–809.

Inderst, R. and M. Ottaviani (2012b). How (not) to pay for advice: A framework for consumer financial protection. *Journal of Financial Economics 105*(2), 393–411.

Kamenica, E. and M. Gentzkow (2011). Bayesian persuasion. *The American Economic Review 101*(6), 2590–2615.

Legros, P. and A. F. Newman (2007). Beauty is a beast, frog is a prince: assortative matching with nontransferabilities. *Econometrica 75*(4), 1073–1102.

Linnainmaa, J. T., B. T. Melzer, A. Previtero, and C. Grace (2015). *Costly Financial Advice: Conflicts of Interest or Misguided Beliefs?* December.

Malenko, A. and N. Malenko (2016). Proxy advisory firms: The economics of selling information to voters.

Mullainathan, S., M. Noeth, and A. Schoar (2012). The market for financial advice: An audit study. Technical report, National Bureau of Economic Research.

Ostrovsky, M. and M. Schwarz (2010). Information disclosure and unraveling in matching markets. *American Economic Journal: Microeconomics 2*(2), 34–63.

Shapley, L. S. and M. Shubik (1971). The assignment game I: The core. *International Journal of game theory 1*(1), 111–130.

Stoughton, N. M., Y. Wu, and J. Zechner (2011). Intermediated investment management. *The Journal of Finance 66*(3), 947–980.

Terviö, M. (2008). The difference that CEOs make: An assignment model approach. *The American Economic Review 98*(3), 642–668.

A Proofs

A.1 Proof of Lemma 1

The result follows from the law of iterated expectations. Consider an information policy P. The values of adviser and customer are

$$\begin{aligned}\tilde{V}(b,l,P) &= \alpha(1-F(b\lambda))E_P[i(\mu)]\\ \tilde{U}(b,l,P) &= (1-F(b\lambda))E_P\left[\max\left\{r\int_{\lambda}^{1}d\mu(s)-l\int_{b\lambda}^{\lambda}d\mu(s)-\alpha_c,0\right\}\right].\end{aligned}$$

Let $\mathcal{M}_0=\{\mu:i(\mu)=0\}$ denote the subset of the space of posterior beliefs where the customer does not invest and let $\mathcal{M}_1=\{\mu:i(\mu)=1\}$ denote the subset where he does. Since $U(p)>u^0(b,a)$, both $\mathcal{M}_0$ and $\mathcal{M}_1$ are non-empty.[38] We have

$$\tilde{V}(b,l,P)=\alpha(1-F(b\lambda))p(\mathcal{M}_1)$$

and

$$\tilde{U}(b,l,P)=(1-F(b\lambda))p(\mathcal{M}_1)E_p\left[r\int_{\lambda}^{1}d\mu(s)-l\int_{b\lambda}^{\lambda}d\mu(s)-\alpha_c|\mu\in\mathcal{M}_1\right].$$

Now define $\mu_N=E_p[\mu|\mu\in\mathcal{M}_0]$, $\mu_I=E_p[\mu|\mu\in\mathcal{M}_1]$, and $p=P(\mathcal{M}_1)$. Consider an information policy that puts mass p on μ_I and mass $1-p$ on μ_N. We have $i(\mu_h)=1$ and $i(\mu_l)=0$, and the ex-ante values under this alternative policy must be the same as under P by the law of iterated expectations.

A.2 Proof of Proposition 1

The Bayes plausibility condition (7) is now $p\mu_I+(1-p)\mu_N=\mu_1$,which immediately implies that μ_I and μ_N are absolutely continuous with respect to μ_1. Therefore, they both must admit densities, which we denote with $\mu_I(s)$ and $\mu_N(s)$. We denote the corresponding cdfs with F_I and F_N.

[38] If $\mathcal{M}_1$ were empty, then $U(p)=0$, which is below the customer's no-information value $u^0(b,a)$. If $\mathcal{M}_0$ were empty, then the customer always invests. But then necessarily $U(p)=u^0(b,a)$.

The adviser's problem in (9) can therefore be written as

$$\begin{aligned} v(b,l,\bar{u}) &= \max_{p,\mu_I,\mu_N} \alpha\left(1-F(b\lambda)\right)p \\ s.t.\ & \left(1-F(b\lambda)\right)p\left(r\left(1-F_I(\lambda)\right)-l\left(F_I(\lambda)-F_I(b\lambda)\right)-\alpha_c\right)-u^0(b,l)\geq\bar{u} \\ & p\mu_I(s)+(1-p)\mu_N(s)=\frac{f(s)}{1-F(b\lambda)}\ \forall s\geq b\lambda \\ & \mu_I(s),\mu_N(s)\geq 0\ \forall s\geq b\lambda \\ & \int_{b\lambda}^{1}\mu_I(s)\,ds=\int_{b\lambda}^{1}\mu_N(s)\,ds=1 \end{aligned} \tag{28}$$

Using Bayes plausibility, we can solve for μ_N as a function of p and μ_I. We get

$$\mu_N(s)=\frac{\frac{f(s)}{1-F(b\lambda)}-p\mu_I(s)}{1-p}$$

and we can optimize over p and $\mu_I(s)$ only subject to $\mu_N(s)$ being non-negative. We now show that the solution must be bang-bang, i.e. either $\mu_I(s)=0$ or $\mu_N(s)=0$ and $\mu_I(s)=\frac{1}{p}\frac{f(s)}{1-F(b\lambda)}$.

Lemma 6. *If $s\geq\lambda$, then $\mu_I(s)=\frac{1}{p}\frac{f(s)}{1-F(b\lambda)}$ at the optimal solution.*

Proof. Take a candidate optimal policy $\{p,\mu_I(.)\}$ and suppose the set

$$\mathcal{S}=\left\{s\geq\lambda:\mu_I(s)<\frac{1}{p}\frac{f(s)}{1-F(b\lambda)}\right\}$$

has positive Lebesgue measure. We show that there exists an alternative policy where we set $\mu_N(s)=0$ on $\mathcal{S}$ and where we can increase p without violating any constraints. Specifically, we choose $\tilde{p}>p$ and

$$\tilde{\mu}_I(s)=\begin{cases}\frac{1}{\tilde{p}}\frac{f(s)}{1-F(b\lambda)} & \text{for}\quad s\in\mathcal{S}\\ \min\left\{\mu_I(s),\frac{1}{\tilde{p}}\frac{f(s)}{1-F(b\lambda)}\right\} & \text{for}\quad s\notin\mathcal{S}\end{cases}$$

such that $\int_{b\lambda}^{1} \tilde{\mu}_I(s) = 1$.[39] The value to the customer under this new policy is

$$\begin{aligned}
&(1 - F(b\lambda))\,\tilde{p}\left(r\left(1 - \int_{b\lambda}^{\lambda} \tilde{\mu}_I(s)\,ds\right) - l\int_{b\lambda}^{\lambda} \tilde{\mu}_I(s)\,ds - \alpha_c\right) \\
&= (1 - F(b\lambda))\,\tilde{p}\left(r - (r+l)\int_{b\lambda}^{\lambda} \tilde{\mu}_I(s)\,ds - \alpha_c\right).
\end{aligned}$$

His value under the old policy is $(1 - F(b\lambda))\,p\left(r - (r+l)\int_{b\lambda}^{\lambda} \mu_I(s)\,ds - \alpha_c\right)$, which is smaller because $\tilde{p} > p$ and $\tilde{\mu}_I(s) \leq \mu_I(s)$ for $s \in [b\lambda, \lambda]$. □

Lemma 7. *If $s \in [b\lambda, \lambda]$, then without loss of generality $\mu_I(s) \in \left\{0, \frac{1}{p}\frac{f(s)}{1-F(b\lambda)}\right\}$ at the optimal solution.*

Proof. Take again a candidate optimal policy such that

$$\mathcal{S} = \left\{s \in [b\lambda, \lambda] : 0 < \mu_I(s) < \frac{1}{p}\frac{f(s)}{1 - F(b\lambda)}\right\}$$

has positive Lebesgue measure. This policy can be replaced by one where μ_I puts weight on 0 and $\frac{1}{p}\frac{f(s)}{1-F(b\lambda)}$ only without altering the payoffs to adviser and customer. We choose a subset $\mathcal{S}_1 \subset \mathcal{S}$ and define $\tilde{\mu}_I$ as

$$\tilde{\mu}_I(s) = \begin{cases} \frac{1}{p}\frac{f(s)}{1-F(b\lambda)} & \text{for} \quad s \in \mathcal{S}_1 \\ 0 & \text{for} \quad s \in \mathcal{S} \setminus \mathcal{S}_1 \\ \mu_I(s) & \text{for} \quad s \notin \mathcal{S} \end{cases}$$

[39] This is always possible: The integral

$$\int_{b\lambda}^{1} \tilde{\mu}_I(s)\,ds = \int_{\mathcal{S}} \frac{1}{\tilde{p}}\frac{f(s)}{1 - F(b\lambda)}ds + \int_{[b\lambda,1]\setminus\mathcal{S}} \min\left\{\mu_I(s), \frac{1}{\tilde{p}}\frac{f(s)}{1 - F(b\lambda)}\right\} ds$$

is decreasing in $\tilde{p}$ and as $\tilde{p} \to 1$, it is bounded by $\int_{b\lambda}^{1} \frac{f(s)}{1-F(b\lambda)}ds$ from above, which integrates to one. Since we are restricting attention to the case when $U(p) > u^0(b,l)$, our initial p is strictly below one, so increasing $\tilde{p} > p$ until $\tilde{\mu}_I$ integrates to one is always feasible.

while p stays the same. $\tilde{\mu}_I$ satisfies

$$\int_{b\lambda}^{1} \tilde{\mu}_I(s)\, ds = \int_{\mathcal{S}_1} \frac{1}{p}\frac{f(s)}{1-F(b\lambda)} ds + \int_{[b\lambda,1]\setminus\mathcal{S}} \mu_I(s)\, ds.$$

For any μ_I and p, we can choose $\mathcal{S}_1 \subset \mathcal{S}$ such that $\tilde{\mu}_I$ integrates to one, which happens whenever[40]

$$\int_{\mathcal{S}_1} \frac{f(s)}{1-F(b\lambda)} ds - p\int_{\mathcal{S}} \mu_I(s)\, ds = 0.$$

The customer's value under this policy is

$$\begin{aligned}
&(1-F(b\lambda))\, p\left(r-(r+l)\int_{b\lambda}^{\lambda} \tilde{\mu}_I(s)\, ds - \alpha_c\right) = \\
&(1-F(b\lambda))\, p\left(r-(r+l)\left(\int_{\mathcal{S}_1} \frac{1}{p}\frac{f(s)}{1-F(b\lambda)} ds + \int_{[b\lambda,\lambda]\setminus\mathcal{S}} \mu_I(s)\, ds\right) - \alpha_c\right) = \\
&(1-F(b\lambda))\, p\left(r-(r+l)\left(\int_{\mathcal{S}} \mu_I(s)\, ds + \int_{[b\lambda,\lambda]\setminus\mathcal{S}} \mu_I(s)\, ds\right) - \alpha_c\right) = \\
&(1-F(b\lambda))\, p\left(r-(r+l)\int_{b\lambda}^{\lambda} \mu_I(s)\, ds - \alpha_c\right),
\end{aligned}$$

so his value remains unchanged. $\square$

Finding the optimal μ_I reduces to finding the optimal set $\mathcal{S}^* \subset [b\lambda, 1]$ so that

$$\mu_I(s) = \begin{cases} \frac{1}{p}\frac{f(s)}{1-F(b\lambda)} & \text{for} \quad s \in \mathcal{S}^* \\ 0 & \text{for} \quad s \notin \mathcal{S}^*. \end{cases}$$

[40] Since $\frac{1}{p}\frac{f(s)}{1-F(b\lambda)} > \mu_I(s)$ on $\mathcal{S}$, for $\mathcal{S}_1 = \mathcal{S}$, the expression is strictly positive, while for $\mathcal{S}_1 = \emptyset$, it is negative. The first integral in the above equation is simply $\mu_1(\mathcal{S}_1)$, i.e. the measure of $\mathcal{S}_1$ under μ_1. Since $\mathcal{S}$ has positive Lebesgue measure, Lebesgue measure is atomless, and μ_1 is absolutely continuous with respect to the Lebesgue measure, Aliprantis and Border (2006), Th. 10.52, p. 395, implies that for any number $0 < q < \mu_1(\mathcal{S})$, there exists a measurable subset $\mathcal{S}_1 \subset \mathcal{S}$ such that $\mu_1(\mathcal{S}_1) = q$. This establishes that we can always pick a set $\mathcal{S}_1$ to make $\tilde{\mu}_I$ integrate to one.

For such a policy, the ex-ante value of the customer becomes

$$
\begin{aligned}
U &= (1-F(b\lambda))\,p\left(r\int_{[\lambda,1]\cap\mathcal{S}^*}\frac{1}{p}\frac{f(s)}{1-F(b\lambda)}ds - l\int_{[b\lambda,\lambda]\cap\mathcal{S}^*}\frac{1}{p}\frac{f(s)}{1-F(b\lambda)}ds-\alpha_c\right)\\
&= r\int_{[\lambda,1]\cap\mathcal{S}^*} f(s)\,ds - l\int_{[b\lambda,\lambda]\cap\mathcal{S}^*} f(s)\,ds - \alpha_c\,(1-F(b\lambda))\\
&= \int_{\mathcal{S}^*}(y(s,l)-\alpha_c)\,f(s)\,ds.
\end{aligned}
$$

Since μ_I integrates to one, we have $p=\frac{1}{1-F(b\lambda)}\int_{\mathcal{S}^*} f(s)\,ds$. With these expressions, we can reduce the adviser's problem in (28) to

$$
\begin{aligned}
v(b,l,\bar{u}) &= \max_{\mathcal{S}^*\subset[b\lambda,1]}\int_{\mathcal{S}^*} f(s)\,ds \\
& \quad s.t. \quad \int_{\mathcal{S}^*}(y(s,l)-\alpha_c)\,f(s)\,ds - u^0(b,l)\geq\bar{u}.
\end{aligned}
\tag{29}
$$

The problem above is analogous to choosing the acceptance region of a statistical test to maximize its power. We can use the Neyman-Peason Lemma,[41] which guarantees that the optimal solution is of the form

$$
\begin{aligned}
1 &\geq k y(s,l) \quad \text{for} \quad s\in\mathcal{S}^*\\
1 &\leq k y(s,l) \quad \text{for} \quad s\notin\mathcal{S}^*
\end{aligned}
$$

for some constant $k\in\mathbb{R}$. For $s\in[b\lambda,\lambda]\cap\mathcal{S}^*$, we have $1\geq -kl$ and for $s\in[b\lambda,\lambda]\setminus\mathcal{S}^*$, we have $1\leq -kl$, so that $k=-\frac{1}{l}$.[42] Since $y(s,l)$ is piecewise constant, we can without loss of generality choose a set of the form $\mathcal{S}^*=[\hat{s},1]$, where $\hat{s}\in[b\lambda,\lambda]$. This establishes Proposition 1.

[41] See e.g. Dantzig and Wald (1951).

[42] The first condition becomes $1\geq -\frac{r}{l}$ for $s\in\mathcal{S}^*$, which holds.

A.3 Proof for Lemma 2

Consider $l' > l$, and $b' > b$. Then, we have

$$
\begin{aligned}
& v(b,l',U(b)) - v(b',l',U(b')) \\
= \; & v(b,l,U(b)) - v(b',l,U(b')) \\
& -\alpha \left\{ \left\{ \frac{F(b\lambda)l'}{l'+\alpha_c} - \frac{F(b\lambda)l}{l+\alpha_c} \right\} - \left\{ \frac{F(b'\lambda)l'}{l'+\alpha_c} - \frac{F(b'\lambda)l}{l+\alpha_c} \right\} \right. \\
& \left. + \left\{ \frac{U(b)+\alpha_c}{l'+\alpha_c} - \frac{U(b)+\alpha_c}{l+\alpha_c} \right\} - \left\{ \frac{U(b')+\alpha_c}{l'+\alpha_c} - \frac{U(b')+\alpha_c}{l+\alpha_c} \right\} \right\} \\
= \; & v(b,l,U(b)) - v(b',l,U(b')) \\
& -\alpha \left\{ (F(b\lambda) - F(b'\lambda)) \left\{ \frac{l'}{l'+\alpha_c} - \frac{l}{l+\alpha_c} \right\} + (U(b) - U(b')) \left\{ \frac{1}{l'+\alpha_c} - \frac{1}{l+\alpha_c} \right\} \right\} \\
> \; & v(b,l,U(b)) - v(b',l,U(b')).
\end{aligned}
$$

That last inequality follows from the last term is positive, given that, $l' > l$, $U(b) - U(b') > 0$, and $F(b\lambda) - F(b'\lambda) < 0$

Hence, if l weakly prefers a lower b, then l' must be strictly prefer b. That is, if $v(b,l,U(b)) \geq v(b',l,U(b'))$, then $v(b,l',U(b)) > v(b',l',U(b'))$. Similarly,

$$
\begin{aligned}
& u(b,l',V(l')) - u(b,l,V(l)) \\
= & u(b',l',V(l') - u(b',l,V(l)) + (F(b'\lambda) - F(b\lambda))(l' - l) \\
> & u(b',l',V(l') - u(b',l,V(l))
\end{aligned}
$$

Hence, if b' weakly prefer a higher l', then b must strictly prefer l'.

A.4 Proof for Proposition 3

Given the constructed $\ell(b)$ and $s^*(b)$, $U(b)$ is well defined. Advisers then choose which customer to attract, taking $U(b)$ as given.

$$
\begin{aligned}
V(l) &= \max_{\tilde{b}} v(\tilde{b}, l, U(\tilde{b})) \\
&= \max_{\tilde{b}} (\alpha_0 + \alpha_c) \left(1 - \left(\frac{F(\tilde{b}\lambda)l}{l + \alpha_c} + \frac{U(\tilde{b}) + \alpha_c}{l + \alpha_c}\right)\right).
\end{aligned}
$$

The construction of $U(b)$ guarantees the FOC of advisers is satisfied, as $U'(b) = F'(b\lambda)\lambda\ell(b)$.

$$
\frac{dv(b, l, U(b))}{db} = -(\alpha_0 + \alpha_c) \left(\frac{F'(b\lambda)\lambda l + U'(b)}{(l + \alpha_c)}\right)
$$

The SOC is also satisfied by construction:

$$
\begin{aligned}
\frac{d^2 v(b, l, U(b))}{d^2 b}|_{b=\ell^{-1}(l)} &= -\frac{(\alpha_0 + \alpha_c)}{(l + \alpha_c)} \left((F''(\ell^{-1}(l)\lambda)\lambda^2 l + U''(\ell^{-1}(l))\right) \\
&= \frac{(\alpha_0 + \alpha_c)}{(l + \alpha_c)} F'(\ell^{-1}(l)\lambda)\ell'(b) < 0
\end{aligned}
$$

where as $U'(b) = -F'(b\lambda)\lambda\ell(b)$ and $U''(b) = -F''(b\lambda)\lambda^2\ell(b) - F'(b\lambda)\lambda\ell'(b)$. This thus proves the optimality of advisers (E2).

Customers' optimization problem:

$$
\begin{aligned}
U(b) &= \max_{l} \left\{\tilde{U}(b, l, s^*(l)) - u^0(b, l)\right\} \\
&= \max_{l} \left(F(s^*(l)) - F(b\lambda)\right) l - \alpha_c(1 - F(s^*(l)))
\end{aligned}
$$

The ODE of $s^*(l)$ (20) guarantees the FOC of (13) is satisfied.

$$
\frac{d\{\tilde{U}(b, l, s^*(l)) - u^0(b, l))}{dl} = (F(s^*(l)) - F(b\lambda)) + F'(s^*(l))(l + \alpha_c)\frac{ds^*(l)}{dl}
$$

The second order condition yields

$$
\frac{d^2\{\tilde{U}(b, l, s^*(l)) - u^0(b, l))}{d^2 l}|_{l=\ell(b)} = F'(\ell^{-1}(l)\lambda)\frac{d\ell^{-1}(l)}{dl} = F'(b\lambda)\lambda\frac{1}{\ell'(b)} < 0
$$

This thus proves customers' optimality (E1).

A.5 Proof of Lemma 3

The information policy must solve the following equation:

$$\frac{ds^*(l)}{dl} = -\frac{(F(s^*(l)) - F(\beta(l)\lambda)}{F'(s^*(l))(l+\alpha)} \equiv m(l, \alpha). \tag{30}$$

When $Q(\bar{b}) < 1$, the initial condition is then given by the full disclosure: $s_0^{*\alpha'} = s_0^{*\alpha} = \lambda$. When $Q(\bar{b}) > 1$, the initial condition $s^{*\alpha}(l^*) = s_0^{*\alpha}$ increases with α. Hence, overall $s_0^{*\alpha'} \geq s_0^{*\alpha'}$, together with $m(l, \alpha') > m(l, \alpha)$, by comparison theorem: $s^*(l, \alpha') > s^*(l, \alpha)$.

B Customers with Heterogeneous Wealth

B.1 Proof of Proposition 5

To see the sorting between (l, w), it is easier to look at the gain for a customer w when matching with an adviser l with utility v, denoted by $u(w, l, v)$. The expression can be derived from the optimization problem (26), which yields

$$\begin{aligned} u(w, l, v) &= w\,(F(\hat{s}) - F(b\lambda))\,l - v \\ &= l(w - \frac{v}{\alpha_c}) - F(b\lambda)lw - v \end{aligned}$$

Since $u_{wv} = 0$ and $u_{wl} = 1 - F(b\lambda) > 0$, customers with higher w will then match with an adviser with more valuable expertise.

B.2 Characterization

Customer's utilities are given by $U^*(w) = \max_l u(w, l, V^*(l))$, and thus

$$\frac{dU^*(w)}{dw} = \ell(w)(1 - F(b\lambda)). \tag{31}$$

Given the sorting, the market clearing condition thus becomes $\int_w^{\bar{w}} dQ(\tilde{b}) = \int_{\ell(w)}^{\bar{l}} dG(\tilde{l})$, which yields differential equation:

$$\frac{d\ell(w)}{dw} = \frac{dQ(w)}{dG(\ell(w))}. \tag{32}$$

Same as before, when advisers are scarce, the marginal customers w^* must earn zero. Thus, $U^*(w)$ and $\ell(w)$ are thus given by Equations (31) and (32), with the boundary condition $\ell(\bar{w}) = \bar{l}$ and $U^*(w^*) = 0$, which then pins down the policy function within the match. Thus, the policy received by customer with wealth w is then given by $s^*(w) = \hat{s}(w, \ell(w), U(w))$, where $\hat{s}(w, l, \bar{u})$ denotes the solution to the optimization problem (26).

Note that one key difference from the environment when customers differ in their level of information is that, conditional on the threshold, customers' types are irrelevant to advisers. On the other hand, advisers do care about the wealth w of customers beyond the threshold. Therefore, the threshold posted by advisers must be contingent on customers' types.

C Customer's Problem

— FOR ONLINE PUBLICATION —

In the model setup, we have assumed that advisers are infinitesimal in that they know about only one asset and that customers choose the adviser who yields them the highest gain relative to the outside option of investing based on their own information. We have made this assumption to keep the model tractable and to have a clean characterization of the information provided by each adviser as a single threshold. In this section, we provide an explicit microfoundation for the customer's decision problem when advisers are not infinitesimal in the sense that they have information over a strictly positive interval of assets. We show that the customer indeed maximizes the gain when choosing an adviser and that the tradeoff between adviser and customer utilities which determines matching in our model approximates the one we have found in Section 3.1. That is, the economics of our model remain unchanged.

Assets are indexed by $l \in \left[\underline{l}, \bar{l}\right]$ and customers by $b \in \left[\underline{b}, \bar{b}\right]$ as before. Each adviser learns the state of assets in the interval $[l - \Delta, l + \Delta] \subset \left[\underline{l}, \bar{l}\right]$ for some l and fixed $\Delta > 0$. Intuitively, each adviser has information about assets that are similar to l. We denote a generic asset in $[l - \Delta, l + \Delta]$ with ℓ and we can again index advisers by l without loss of generality. The set of adviser types is now $\left[\underline{l}+\Delta, \bar{l}-\Delta\right]$. We assume that Δ is sufficiently small so this interval has strictly positive size. After matching with adviser l and receiving information, the customer then decides whether or not to buy each asset. To keep notation minimal we set $\alpha_c = 0$ throughout this section.

His investment decision for asset ℓ is the same as in equation (8). We can use the same argument as in Section 3.1 to show that without loss of generality, the adviser recommends that the customer invests in asset ℓ whenever s_ℓ is above a threshold $\hat{s}_\ell \in [b\lambda, \lambda]$. The only difference to the main model is thus that the adviser provides the customer with multiple thresholds $(\hat{s}_\ell)_{\ell\in[l-\Delta,l+\Delta]}$. Given this policy, the utility of customer b from receiving the adviser's information is

$$\tilde{U}(b,l) = \int_{l-\Delta}^{l+\Delta} \left[r\left(1-F(\lambda)\right) - \ell\left(F(\lambda) - F(\hat{s}_\ell)\right)\right] d\ell.$$

His utility from investing in asset ℓ without advice is given by

$$u^0(b,\ell) = r\left(1-F(\lambda)\right) - \ell\left(F(\lambda) - F(b\lambda)\right).$$

The customer's problem of choosing adviser l thus becomes

$$\max_{l\in\left[\underline{l}-\Delta,\bar{l}+\Delta\right]} \tilde{U}(b,l) + \int_{\ell\notin[l-\Delta,l+\Delta]} u^0(b,\ell)\, d\ell.$$

That is, in choosing adviser l, the customer understands that the adviser only has information about assets $\ell \in \left[\underline{l}+\Delta, \bar{l}-\Delta\right]$. The customer's utility from receiving this advice is the first term in the maximization problem. For all other assets, he must invest based on his own information, which is the second term. We can rewrite this equation as

$$\max_{l\in\left[\underline{l}-\Delta,\bar{l}+\Delta\right]} \tilde{U}(b,l) - \int_{l-\Delta}^{l+\Delta} u^0(b,\ell)\, d\ell + \int_{\underline{l}}^{\bar{l}} u^0(b,\ell)\, d\ell,$$

which shows that the customer's utility from choosing adviser l is the gain from receiving information about assets in $[l-\Delta, l+\Delta]$, which is represented by the first two terms, plus the utility he would receive by simply investing in all assets by himself. This problem is equivalent to solving

$$\begin{aligned} & \max\nolimits_{l\in\left[\underline{l}-\Delta,\bar{l}+\Delta\right]} \tilde{U}\left(b,l\right)-\int_{l-\Delta}^{l+\Delta} u^0\left(b,\ell\right)d\ell \\ = \quad & \max\nolimits_{l\in\left[\underline{l}-\Delta,\bar{l}+\Delta\right]} \int_{l-\Delta}^{l+\Delta}\left[\ell\left(F\left(\hat{s}_\ell\right)-F\left(b\lambda\right)\right)\right]d\ell, \end{aligned}$$

which is the analog of equation (10) in Section 3.1.

In choosing advisers with given information policies, the customer thus chooses the one with the highest gain. This is analogous to how we have modeled the customer's behavior in the main sections of the paper. Extending the model in this way does not qualitatively alter the equilibrium, except that we now have a vector of thresholds for each adviser. To see this, we can write the problem of adviser l conditional on having to promise a certain gain $\Delta\bar{u}$ to attract customer b, which is the analog of the adviser's problem in equation (11),

$$\begin{aligned} v\left(b,l,\bar{u}\right) \quad = \quad & \max_{\left(\hat{s}_\ell\right)_{\ell\in\left[l-\Delta,l+\Delta\right]}} \alpha\int_{l-\Delta}^{l+\Delta}\left(1-F\left(\hat{s}_l\right)\right)d\ell \\ & \int_{l-\Delta}^{l+\Delta}\left[\ell\left(F\left(\hat{s}_\ell\right)-F\left(b\lambda\right)\right)\right]d\ell \geq \Delta\bar{u}. \end{aligned}$$

Substituting the constraint into the objective yields the expression

$$v\left(b,l,\bar{u}\right)=2\Delta\alpha\left(1-\left(\frac{u}{l}+F\left(bl\right)\right)+o\left(\Delta\right)\right).$$

That is, for sufficiently small Δ, the adviser's value function within each match has approximately the same shape as the one we have found in equation (**??**).