Arbitration with Uninformed Consumers

Mark Egan, Gregor Matvos, and Amit Seru*

January 2025

Abstract

This paper studies the impact of the arbitrator selection process on consumer outcomes. Using data from consumer arbitration cases in the securities industry over the past two decades, where we observe detailed information on case characteristics, the randomly generated list of potential arbitrators presented to both parties, the selected arbitrator, and case outcomes, we establish several motivating facts. These facts suggest that firms hold an informational advantage over consumers in selecting arbitrators, resulting in industry-friendly arbitration outcomes. We then develop and calibrate a quantitative model of arbitrator selection in which firms hold an informational advantage in selecting arbitrators. Arbitrators, who are compensated only if chosen, compete with each other to be selected. The model allows us to decompose the firms' advantage into two components: the advantage of choosing pro-industry arbitrators from a given pool and the equilibrium pro-industry tilt in the arbitration pool that arises because of arbitrator competition. Selecting arbitrators without the input of firms and consumers would increase consumer awards by $60,000 on average relative to the current system. Forty percent of this effect arises because the pool of arbitrators skews pro-industry due to competition. Even an informed consumer cannot avoid this pro-industry equilibrium effect. Counterfactuals suggest that redesigning the arbitrator selection mechanism for the benefit of consumers hinges on whether consumers are informed. Policies intended to benefit consumers, such as increasing arbitrator compensation or giving parties more choice, would benefit informed consumers but hurt the uninformed.

Keywords: Arbitration, financial advisers, brokers, and fraud

*We especially thank Pedro Bordalo, Eric Budish, John Campbell, Bentley MacLeod, Jesse Shapiro and Johannes Stroebel. We also thank Sumit Agarwal, Judith Chevalier, Lauren Cohen, Alexander Dyck, Glenn Ellison, Allesandro Gavazza, Nicola Gennaioli, Will Gerken, Paul Goldsmith-Pinkham, Shane Greenstein, Colleen Honigsberg, Ben Iverson, Terrance Odean, Giorgia Piacentino, Andrea Pozzi, Andrew Tuch, and the seminar participants at Bocconi University, Columbia Business School, the Consumer Financial Protection Bureau, Einaudi Institute for Economics and Finance, FINRA, Harvard Business School, the Kellogg School of Management at Northwestern University, NYU Stern School of Business, Penn State, Stanford Graduate School of Business, Tilburg University, the University of California at Los Angeles, University of Texas at Austin, the University of Wisconsin School of Business, and the Yale School of Management. We also thank participants at the ASSA meetings, Financial Conduct Authority and Imperial Business School Conference on Household Finance and Consumer Decision-Making, workshop on Financial Intermediation and Regulation at Queens University, NBER Household Meetings, NBER Industrial Organization Meetings, NBER Law and Economics Meetings, Stanford-Berkeley joint workshop, and the Yale School of Management Jr. Finance Conference.

I Introduction

A large share of potential disputes between consumers and firms in the United States are settled through mandatory arbitration rather than the court system. Arbitration is a private dispute resolution mechanism in which an arbitrator issues a legally binding resolution. When consumers purchase a product or service, the purchase often contains a pre-dispute arbitration provision. This provision prohibits the consumer from suing the seller in court and mandates resolving any dispute via arbitration. In the United States, such clauses are currently used by all brokerage firms and the largest insurance and financial firms, as well as non-financial firms such as online retailers, wireless providers, and sharing economy firms, covering trillions of dollars of transactions.[1][2]

A central feature of arbitration is the ability of both parties to explicitly exert control over arbitrator selection. For example, in securities arbitration, each party is presented with a randomly generated list of arbitrators and can strike a limited number of arbitrators from the list. Practitioners believe that choosing an arbitrator can significantly affect the case outcome: "The selection of an appropriate arbitrator or arbitration tribunal is nearly always the single most important choice confronting parties in arbitration" (Stipanowich et al., 2010). The existing literature has mainly focused on arbitration in which both parties are equally informed when choosing arbitrators.[3] This paper studies consumer arbitration, examining the consequences of firms holding an informational advantage over consumers in selecting arbitrators. More broadly, our results apply to dispute resolution settings in which parties can influence the choice of the decision maker and differ in their ability to do so, such as venue shopping in corporate bankruptcy (Gennaioli and Rossi, 2010).

This paper has two goals. The first is to establish motivating facts that support the notion that firms hold an informational advantage over consumers in selecting arbitrators, which in turn results in industry-friendly arbitration outcomes. The second goal is to develop and calibrate a model of arbitrator selection in which firms hold an informational advantage in selecting arbitrators. Arbitrators, who are compensated only if chosen, compete with each other to be selected. The model allows us to decompose the firms' advantage into two components: the ability to choose pro-industry arbitrators from a given pool and the equilibrium pro-industry tilt in the arbitration pool that arises because of competition between arbitrators. The model reveals that accounting for this advantage is critical in assessing changes to arbitration design. Policies that would benefit consumers if they were informed, such as increased competition or compensation for arbitrators, could in fact be harmful to consumers, illustrating the importance of investor sophistication in consumer financial markets. Our study contributes to the existing literature by evaluating and quantifying

[1]Examples include the largest insurance companies (AIG, Aetna Inc., Blue Cross and Blue Shield, and Travelers); the largest financial firms (American Express, Bank of America, Chase Bank, and Citigroup); the largest fintech firms (PayPal, Venmo, and Square); the largest online retailers (Amazon, eBay, and Walmart.com); and the largest wireless providers (Verizon, AT&T, T-Mobile, and Sprint), among others.

[2]The American Arbitration Association alone oversaw 68,000 consumer cases in 2020, about five times as many as the 12,922 consumer protection lawsuits filed in courts in 2020. Sources: https://www.adr.org/sites/default/files/document_repository/AAA_2020_AnnualReport_and_Financial_Statements.pdf, https://lexmachina.com/media/press/lex-machina-releases-2021-consumer-protection-litigation-report/.

[3]For example, arbitrations between unions and employers, or in an experimental setting (Farber and Bazerman, 1986; Bloom, 1986; Ashenfelter et al., 1992).

these economic forces and recent policy proposals.

We study arbitration in the securities industry using a new data set of roughly 5,000 disputes between consumers and financial advisers from 1998 to 2019. The securities industry lends itself to studying arbitration because of the data availability and the institutional setting. Our data on securities arbitration come from the Financial Industry Regulatory Authority's (FINRA's) Arbitration Awards Database, which we merge with FINRA's BrokerCheck data. We observe detailed information on the claimant (consumer), respondent (firm), arbitrators, dispute details, and awards. In addition, the institutional environment has several useful features. Pre-dispute arbitration agreements (PDAAs) are required in virtually all broker-dealer contracts, implying that there is no selection of firms or consumers in arbitration clauses. FINRA provides a uniform pool of arbitrators, as well as rules governing arbitration, so the choice of venue is also fixed. Moreover, the selection system used by FINRA is similar to those of the largest consumer arbitration forums such as the American Arbitration Association (AAA) and JAMS. Most important for the research design, FINRA randomizes the list of potential arbitrators from which the parties select the arbitration tribunal, which we exploit in our research design.[4] For a significant subset of cases, we also observe the randomly generated list of potential arbitrators presented to both parties in addition to the arbitrators who are selected to the case.

Arbitration in the brokerage industry is also interesting per se. Roughly 20 million U.S. households hold a brokerage account, amounting to $20 trillion of assets (2016 Survey of Consumer Finances). The cases involve significant amounts: mean and median damages requested are $760,000 and $240,000, respectively, providing substantial incentives for the parties in arbitration. The regulator, FINRA, established the Dispute Resolution Task Force in 2014 to investigate concerns that the arbitration procedures lead to outcomes favoring the industry. More recently, the Consumer Financial Protection Bureau (CFPB) proposed a new rule regulating mandatory arbitration clauses in certain financial products (Arbitration Agreements, 12 C.F.R. ï¿œ 1040 2017). Understanding arbitration design in the financial industry therefore has direct policy relevance.

We demonstrate in two steps that firms have an informational advantage in selecting arbitrators. We use these facts to motivate the assumptions behind our quantitative model. First, we confirm practitioner intuition that some arbitrators are systematically more industry-friendly and that others are more consumer-friendly (Stipanowich et al., 2010). Controlling for the arbitrator overseeing the case explains an additional 36% of the variation in arbitration awards in excess of case characteristics. An arbitrator who is more industry-friendly by one standard deviation awards 14 percentage points (pp) less in damages relative to the amount claimed. For a median case ($240,000), this translates to a $33,600 smaller award for the consumer. It is therefore not surprising that, anecdotally, brokerage firms maintain proprietary internal arbitrator rankings, or arbitrator "strike lists," to guide their arbitrator selection process.

Second, we find that firms have an informational advantage over consumers in choosing arbi-

[4]Ernst & Young has verified the randomization process. See https://www.finra.org/arbitration-mediation/arbitrator-selection.

trators who are favorable to them. The list from which the parties select arbitrators is randomly generated by FINRA. If both parties were equally well-informed, they would strike arbitrators who favor the opposing side, and the median arbitrator from the list would be chosen. By this logic, being industry- or consumer-friendly should not increase selection chances. Instead, we find that industry-friendly arbitrators are 50% more likely to be chosen from the list than their consumer-friendly counterparts. Several tests exploiting within geography $\times$ time variation, among other factors, suggest that these patterns are not due to unobserved case characteristics. The strongest evidence comes from the subset of cases in which we observe all arbitrators on the randomly generated list. Exploiting *within case* differences in arbitrators' industry friendliness on these cases, we show that arbitrator selection arises because firms are better at eliminating consumer-friendly arbitrators. Additional evidence reveals that a firm's advantage increases with experience in arbitration. Conversely, a firm's informational advantages are diminished when consumers use an attorney specialized in arbitration.

In the second part of the paper, we present a quantitative model of arbitrator selection. The analysis in the first part of the paper shows that firms have an informational advantage when choosing arbitrators from a given pool of arbitrators. The model highlights a second advantage of the informed party: because arbitrators compete to be selected by the informed party (the firm), the whole pool of arbitrators becomes more industry-friendly in equilibrium. We use the calibrated model to quantify these two components and evaluate different proposed arbitration designs from the perspective of consumer gains and losses.

The model mirrors the institutional setting: firms and consumers strike arbitrators from a randomly generated list. Arbitrators differ in their underlying beliefs of fair awards. They can depart from these beliefs and choose how consumer- or industry-friendly they are (i.e., their "slant"). This concept of slant is similar to the choice of judicial discretion in court decisions (Gennaioli and Rossi, 2019). Arbitrators are compensated only if they are selected to arbitrate a case. They compete with other arbitrators to be selected on the arbitration panel. In doing so, they trade off their preferences for a fair award with monetary compensation from arbitration. Sophisticated firms observe arbitrators' slants, while consumers are uninformed.

A key result of the model is that because arbitrators compete to be selected, the whole pool of arbitrators becomes industry-friendly, increasing the informational advantage of firms. This competition between arbitrators exacerbates the informational advantage of firms in equilibrium. Even though the underlying beliefs of arbitrators are unbiased, competition among arbitrators drives all arbitrators to slant their case decisions in favor of firms. Intuitively, when consumers are uninformed, arbitrators compete to avoid being eliminated by firms. This effect can best be seen in the special case when arbitrators only maximize their monetary payoffs and have no fairness concerns. To avoid being eliminated, every arbitrator wants to be a bit more industry-friendly than other arbitrators. Because all arbitrators want to do the same, this results in a "race to the bottom." The only equilibrium is for all arbitrators to be as industry-friendly as possible. In our full model, the competition effect does not unravel to the same degree because it is costly for arbitrators to deviate

from what they believe is fair. The extent of this trade-off is pinned down by the data. The fact that the industry friendliness of the overall pool of arbitrators changes in response to firms being informed has several implications.

First, the model illustrates that the analysis in the first part of the paper measures only one part of firms' information advantage because it cannot measure the extent of the competition effect. This is by design: to eliminate as much variation as possible, we compare how industry-friendly arbitrators are *relative to each other*. The analysis then measures the informational advantage for a given pool of arbitrators. The fixed effect analysis cannot detect whether the whole pool is industry-friendly relative to arbitrators' beliefs of fair awards. We use the calibrated model to back out arbitrator beliefs and decompose the informational advantage into its two components.

Second, the competition result stands in stark contrast to arbitration in which both parties are informed. When both parties are informed, competition between arbitrators is desirable because it leads to less biased outcomes and statistical exchangeability of arbitrators. The idea behind statistical exchangeability is that "Since the parties play a role in the selection of the arbitrator who will decide their dispute, arbitrators who are known to favor one of the parties will be eliminated. This selection process created incentives for arbitrators to maintain characteristics that make them 'statistically exchangeable' with other arbitrators" (Ashenfelter et al., 1992, p. 1408). This argument is very powerful when both parties are informed about which arbitrators to eliminate (for example, in the setting of employer/union arbitration). The competitive forces that lead to statistical exchangeability when both parties are informed lead to biased outcomes when one party holds an informational advantage. Our counterfactuals illustrate the impact of this result considering policy changes to the arbitrator system.

We calibrate the model to quantify firms' informational advantage. Calibrating the model, we obtain the underlying distribution of arbitrators' beliefs, i.e., the awards that arbitrators would have chosen absent incentives provided by the arbitration selection mechanism. We use the estimates to decompose the informational advantage into the advantage of being better at striking arbitrators from a fixed arbitrator pool, as well as the equilibrium effect on the pool itself. Firms' informational advantage is substantial: on average, selecting arbitrators without input from parties would increase average consumer awards by 8pp, or $60,000 (relative to the amount requested), relative to the current system. Approximately 60% of that effect can be attributed to firms' informational advantage when striking arbitrators. Competition between arbitrators accounts for the remaining 40%.

We use the calibrated model to evaluate alternative arbitrator selection schemes, such as those proposed by the regulator (FINRA). Policy proposals that aim to improve arbitration outcomes are frequently designed without considering the informational advantage of firms. The counterfactuals suggest that several proposals, which would be "consumer-friendly" if consumers were as informed as firms, yield industry-friendly outcomes once one accounts for firms' informational advantage. For example, increasing arbitrator compensation has been touted as potentially improving arbitration outcomes for consumers (FINRA Notice 14-49, 2014). Our estimates suggest that doubling arbi-

trator compensation would decrease awards by $45,000, on average, because increasing arbitrator compensation further incentivizes arbitrators to act industry-friendly if firms hold an informational advantage. One implication of our model is that lower-powered incentives for arbitrators, potentially coupled with a flat wage, could decrease the pro-industry slant in arbitration. Similarly, increasing the number of strikes, which was also proposed to benefit consumers, would instead substantially lower awards to consumers.

We study several extensions of our model. The industry often argues that increasing consumer friendliness of the arbitration system would result in an increased propensity for "frivolous" cases. In such cases, consumers hope to win a large award on a meritless case, such as by drawing a consumer-friendly arbitrator. Our estimates imply a $2,150 expected payoff to bringing a low-merit case under the current system. Therefore, filing frivolous cases has a limited upside given the fixed costs of arbitration. Additionally, contrary to industry concerns, more industry-friendly mechanisms not only decrease the payoffs to frivolous cases, but also decrease incentives to file cases with merit. Finally, our results suggest that mechanisms that reduce the variance of arbitration outcomes could be effective in decreasing the number of frivolous arbitration cases.

Overall, our model highlights the qualitative and quantitative importance of accounting for firms' informational advantage when designing arbitration in the securities industry and consumer arbitration more broadly. Policies that are consumer-friendly when both parties are informed can end up being industry-friendly when only firms are informed. We provide a workhorse quantitative model for policy evaluation, which can be extended to account for risk aversion, frivolous cases, and different welfare criteria relevant to evaluating arbitration mechanisms.

One implication of our results is that the pro-industry bias in arbitration increases incentives to engage in socially inefficient adviser misconduct.[5] These issues loom larger because the market for financial advice is not perfectly competitive, and less sophisticated consumers are unaware of adviser misconduct and biased arbitration. Consistent with this view, Egan, Matvos, and Seru (2019) show that adviser misconduct is relatively common (1 in 10 advisers in the United States have a past record of misconduct) and concentrated in areas with more elderly and less educated households.

Related Literature: Our paper relates to the existing literature on arbitration. One strand of the literature provides empirical evidence that arbitrators are statistically exchangeable (Farber and Bazerman, 1986; Bloom, 1986; Ashenfelter et al., 1992). This result stands in contrast to the large differences among arbitrators we document. These studies mainly focus on arbitration in which both parties are equally informed, such as arbitrations between unions and employers, or arbitration in an experimental setting. In contrast, we study consumer arbitration where potential differences in parties' information loom large.

The focus on the information gap in consumer arbitration also distinguishes our work from existing work on arbitrator selection. Our findings are consistent with Bloom and Cavanagh's

[5] For example, roughly 16% of arbitration cases involve churning, where the advisers excessively trade in clients' accounts to generate additional commissions.

(1986a), who find that arbitration parties tend to select arbitrators based on their preferences in arbitrations operated by the New Jersey Public Employment Relations Commission. Our work suggests that parties can only do so when informed. Kondo (2006) examines securities arbitration by the National Association of Securities Dealers (NASD) from 1991 to 2004. In contrast to the standard arbitrator selection process in which the parties choose arbitrators, the regulator actively participated in selecting arbitrators during this period. Choi, Fisch, and Pritchard (2014) also study NASD arbitration from 1998 to 2000 and find that arbitrators with industry experience and those selected more frequently tend to give lower awards. In contrast, our work focuses on understanding the impact of consumer sophistication on arbitrator selection and outcomes in FINRA arbitration over the last two decades. During this time, FINRA's arbitration selection process resembled those followed in consumer arbitration disputes more generally, making our findings potentially widely applicable (see Appendix E for some suggestive evidence). More importantly, we develop and calibrate a workhorse quantitative model of arbitrator selection and use it to decompose the firms' advantage into two components: the advantage of choosing pro-industry arbitrators from a given pool and the equilibrium pro-industry tilt in the arbitration pool that arises because of arbitrator competition. We show that both are quantitatively important when evaluating alternative policies.

Our paper is related to the theoretical literature on designing arbitration mechanisms. A large part of this literature has focused on the difference between conventional arbitration and final-offer arbitration (Stevens, 1966).[6] De Clippel et al. (2014) develop a framework for understanding the selection of arbitrators from the perspective of implementation theory and test their theoretical framework in an experimental setting. The existing literature maintains the assumption that both parties are equally informed and have complete information when selecting arbitrators, which is reasonable in many settings (e.g., arbitration between firms, unions, or countries). We depart from the literature by finding evidence that one party holds an informational advantage and by studying the associated consequences of arbitration design.

Our model and setting closely relates to the market for corporate judges and judicial discretion (Giammarino and Nosal, 1994; Bernhardt and Nosal, 2004; Gennaioli and Shleifer, 2008; Gennaioli, 2013). Our quantitative model closely builds on the theoretical work of Gennaioli and Rossi (2010), who develop a model of judicial bankruptcy with biased corporate judges. Because debtors can forum shop, they will find it optimal to file in pro-debtor courts, creating a systematic pro-debtor bias in bankruptcy filings. This is analogous to what we refer to as the "striking effect" in our setting, in which firms use their information advantage to systematically select more industry-friendly arbitrators. When corporate judges have career concerns, Gennaioli and Rossi (2010) show

[6] Crawford (1979, 1982) studies the effect of conventional and final-offer arbitration on negotiated settlements. Farber (1980) and Farber and Katz (1979) explore the case where the parties are uncertain about the arbitrator's preferences and find that the outcomes under conventional and final-offer arbitration generally differ. Brams and Merrill (1983, 1986) model arbitration as a zero-sum game of imperfect information. Gibbons (1988) analyzes strategic communication in equilibrium models of conventional and final-offer arbitration and emphasizes the role of learning by the arbitrator from the parties' offers about the state of the employment relationship. Rosenthal (1978), Samuelson (1991), Farmer and Pecorino (1998, 2003), Deck and Farmer (2007), and Olszewski (2011) compare different arbitration procedures under incomplete, asymmetric information.

that the pro-debtor bias becomes stronger, as all judges have an incentive to establish a pro-debtor reputation to attract more cases in the future. This is analogous to our "competition effect," in which all arbitrators slant their decisions in favor of the industry to increase their probability of being selected in the future to earn fees. The results from Gennaioli and Rossi (2010) suggest that the effects we quantify are important in any setting in which parties can influence the choice of the decision-maker and in which decision-makers want to be selected.

Our study is related to work in behavioral finance, which highlights the importance of trust and investor sophistication in consumer financial markets (Campbell, 2006; Guiso et al., 2008; Gennaioli et al., 2015; Agarwal et al., 2015; Argyle et al., 2019; Andersen et al., 2020ab; Goldsmith-Pinkham and Shue, 2020; Gomes et al., 2020). Our conclusion that consumers fail to select friendly arbitrators is consistent with evidence that individual investors underperform in financial markets due to a lack of consumer sophistication (Barber and Odean, 2000; Egan, 2019) and due to inattention (Giglio et al., 2020).[7]

More broadly, our study relates to the literature using quantitative models to study the effect of competition in financial markets (Koijen and Yogo, 2016; Benetton, 2018; Crawford, Pavanini, and Schivardi, 2018; Allen, Clark, and Houde, 2019; Bhattacharya, Illanes, and Padi, 2019; Allen et al., 2020; Gilbukh and Goldsmith-Pinkham, 2020). We depart from this literature by focusing on competition between arbitrators in the area of dispute resolution, rather than on competition in consumer financial products. Findings suggesting that competition can sometimes be "undesirable" is related to the larger literature of market design in financial markets (Budish, Cramton, and Shim, 2015; Aquilina, Budish, and O'Neil, 2020; Budish, Lee, and Shim, 2019).

II Institutional Details: Consumer Arbitration

II.A Consumer Arbitration in the United States

Arbitration is a private dispute resolution alternative to civil courts. It differs from the civil court system along several important dimensions. Arbitration is typically binding without appeals, and courts have had limited ability to vacate or modify arbitration awards (*Hall Street Associate, LLC vs. Mattel, Inc.*, 552 US 576, 2008). Advocates of arbitration argue that this feature means arbitration is usually quicker and less expensive than litigation (U.S. Chamber of Commerce Institute for Legal Reform, 2005). Second, as described below, the parties involved in a given dispute exert significant control in selecting arbitrators, while courts select judges. Third, while judges are frequently paid a fixed salary, arbitrators are compensated only if they are selected for a case.

Consumer arbitration is ubiquitous in the United States. The CFPB's Arbitration Study (2015) estimates that 50% of credit card loans ($500 billion) and 44% of insured deposits ($3.1 trillion) are subject to mandatory arbitration. Arbitration is common in most consumer financial products, such as automobile loans, brokerage accounts, payday loans, etc., and in many other non-financial products, such as cable TV, cell phone, internet, and car rental contracts (Silver-Greenberg and

[7] Our paper also relates to the literature on misconduct among financial advisers, including Dimmock et al. (2018); Qureshi and Sokobin (2015); Egan et al. (2017, 2019); and Charoenwong et al. (2019).

Gebeloff, 2015). Arbitration is also prominent in employment contracts. More than half (54%) of non-union private-sector employers have mandatory arbitration procedures, affecting an estimated 60 million American workers (Colvin, 2018).

Arbitration proceedings are governed by an administrator/forum who determines the procedural rules. Administrators often provide a list of potential arbitrators and govern the selection process. Our analysis focuses on securities arbitration between consumers and brokerage firms, which is exclusively administered by FINRA. The two other dominant forums for consumer arbitration are AAA and JAMS.[8]

A central feature of arbitration is the parties' control over the arbitrator selection process. This selection process is based on the premise that arbitrators differ in terms of how favorable they might be to either party. Although the specifics vary across forums, the process typically involves ranking and striking potential arbitrators by the consumer (claimant) and firm (respondent). For example, in FINRA and JAMS arbitration, the administrator sends a list of potential arbitrators to the consumer and firm. Each party can remove/strike a fixed number of arbitrators from the list and then must rank the remaining arbitrators. Among arbitrators who are not struck, the one with the lowest joint rank is selected.

II.B FINRA (NASD) Arbitration

Here, we briefly discuss the institutional details of the arbitration proceedings and the arbitrator selection process used by FINRA or, prior to 2007, NASD.[9] We focus on consumer arbitration in which consumers file a claim against a brokerage firm. Consumer arbitration mechanisms differ from mechanisms used to arbitrate union contracts, international business, or country treaties, which are not the focus of this paper. As we discuss in Appendix E, FINRA's arbitrator selection mechanism and arbitrator incentives are similar to other *consumer* arbitration settings.

Consumers initiate arbitration by filing a Statement of Claim with FINRA, in which they provide details of the dispute and the type of relief requested. Consumers can modify these claims until an arbitration panel is appointed; afterwards, consumers can only modify their claim if they are granted a formal motion to amend the claim (FINRA 12309).

Next, consumers and brokerage firms select arbitrators. FINRA (formerly NASD) maintains a roster of more than 7,000 eligible arbitrators. Generally, arbitrators must have at least 5 years of any paid work experience and at least 2 years of college. FINRA describes the pool of arbitrators as ranging from "freelancers to retirees to stay-at-home parents" ("Become an Arbitrator Frequently Asked Questions," 2018). As we document in Section III, arbitrators are often current or former financial advisers. Prior to hearing cases, an arbitrator must have completed FINRA's 12-hour Basic Arbitrator Training Program. Arbitrators are classified as public, non-public, and public chairpersons. Public arbitrators are those that have not worked in the financial industry in the past 5

[8]For example, AAA is listed as a potential forum in over 80% of credit card, checking account, prepaid card, and mobile wireless arbitration clauses studied by the CFPB (2015). The National Arbitration Forum previously administered consumer arbitrations but ceased administering consumer arbitration in 2009.

[9]Full details on the arbitration proceeding can be found on the FINRA website: https://www.finra.org/arbitration-and-mediation/code-arbitration-procedure.

years.[10] Conversely, non-public arbitrators either currently work in the financial services industry or have done so within the past 5 years. Public arbitrators can qualify as chairpersons if they: (i) have served on at least three arbitration panels or (ii) have served on at least one arbitration panel, have a law degree, and are members of the bar.

Arbitrators are selected using the Neutral List Selection System (NLSS). For each case, an automated process generates a list of arbitrators on a rotational basis based on the geographic location of the hearing site (FINRA 10308(b)(4)(A)). In general, an arbitration panel consists of one or three arbitrators. The composition of the arbitration panel depends on the claim amount. Under the current guidelines, claims under $50,000 generally have one chairperson arbitrator, claims between $50,000-$100,000 have one chairperson arbitrator but can have up to three arbitrators, and claims over $100,000 generally have three arbitrators. Cases with three arbitrators have one chairperson, one public arbitrator, and one non-public arbitrator.[11] For cases with one arbitrator, the NLSS randomly generates a list of 10 public arbitrators from the FINRA public chairperson roster. For cases with three arbitrators, the NLSS randomly generates three separate lists: a list of 10 arbitrators from the non-public arbitrator roster, a list of 10 arbitrators from the FINRA public arbitrator roster, and a list of 10 arbitrators from the FINRA public chairperson roster (FINRA 12400). [12]

Two aspects are critical to the process. First, to generate the list, NLSS randomly selects arbitrators. According to FINRA, "The randomized process has been verified by an Ernst & Young audit in a report that confirmed that a 'random pool management algorithm [is] used to ensure that each arbitrator in the pool has the same opportunity to appear on a list as all other arbitrators in that pool.'"[13] Second, each party then reviews and ranks the list of arbitrators according to the following rules. A party may strike one or more arbitrators from a list for any particular reason. Prior to a 2007 rule change, parties were allowed to strike an unlimited number of arbitrators from each list. Starting in 2007, the number of strikes was limited to four on each list. The struck arbitrators are immediately deemed ineligible to preside over the arbitration hearings. The parties then rank the remaining arbitrators. Arbitrators are then appointed based on their cumulative ranking, which is constructed by adding the rankings of both parties.

Arbitrators are compensated only for the cases they arbitrate. The minimal compensation for an arbitrator is $75 per hour, and can be substantially larger for shorter hearings.[14] In addition, arbitrators are entitled to reasonable local expenses. Their compensation is almost twice the me-

[10] Since 2015, this definition was expanded to exclude all individuals with any experience in the financial industry (SR-FINRA-2014-028 eff. June 26, 2015).

[11] Starting in 2011, FINRA allowed customers to choose an all-public arbitration panel (FINRA Regulatory Notice - 11-05).

[12] In 2017, FINRA increased the size of the list of public arbitrators from 10 to 15 and allowed both parties to strike 6 arbitrators from the list (FINRA Regulatory Notice 16-44).

[13] https://www.finra.org/arbitration-mediation/arbitrator-selection.

[14] Compensation comprises $300 per hearing (chairpersons earn an additional $125 per day), which can last at most 4 hours, with at most two hearings per day (hearings can be from the same case). The typical case lasts 4 days, which means arbitrators could expect to earn $1,200-$2,900 on the average case, depending on whether the arbitrator serves as the chairperson and on the number of hearings per day.

dian hourly compensation of $39.8 for financial analysts and financial advisers, who comprise a substantial amount of the arbitration pool, and is comparable to the average compensation of federal district judges ($79 per hour).[15] [16] Becoming an arbitrator also offers non-pecuniary benefits. FINRA advertises that arbitrators have the opportunity to "build networks," "gain professional experience," and "acquire knowledge of the securities industry" ("Become an Arbitrator Frequently Asked Questions," 2018). Given the sizable compensation, it is not surprising that FINRA has a large roster of potential arbitrators at its disposal. Critically, arbitrators are paid only if they are selected onto a panel; they do not receive benefits or other payments simply for being on the roster.

III Data

III.A Data Construction

We construct a novel data set containing the details and awards of roughly 5,000 securities arbitration cases involving consumer disputes with financial advisers occurring from 1998 to 2019. The claimant is always a consumer and the respondent is always a financial adviser. We observe the details of each arbitration case, including the parties involved (claimant, respondent, and arbitrator), the nature of the allegations being arbitrated, detailed information on the respondent, and the outcome of the proceedings. We construct the data set primarily from two sources: FINRA's Arbitration Awards Online and FINRA's BrokerCheck website. We collect additional data, which we describe in the body of the paper.

FINRA's Arbitration Awards Online Data: FINRA's Arbitration Awards Online contains the details of the FINRA and NASD universe of arbitration hearings. For each arbitration case, we collect the case/award documents and systematically parse each document for information regarding the consumer (claimant), financial adviser (respondent), and arbitrator. The arbitration documents also contain detailed accounts of the nature of the disputes and awards, providing us with a detailed picture of the cases and the similarities between cases. The data cover consumer arbitration, as well as other arbitration disputes such as employment disputes between advisers and their respective employers. We match the arbitration awards data with FINRA's BrokerCheck data, which provide additional granular details on each consumer arbitration case. We also obtain detailed data on defendants' employment and misconduct history and obtain the same information for any arbitrator who was employed in the financial industry.

Random List of Arbitrators: For a subset of cases (536) studied in Honigsberg and Jacob (2020), we also observe the lists of arbitrators that were presented to the litigants. These lists are randomly generated through the NLSS. The lists are de-identified but have been merged with our estimates of individual arbitrator industry friendliness. For each case, FINRA typically generates three separate lists of 10 arbitrators, and one arbitrator is selected from each list. In our empirical analysis, we

[15] Adviser compensation data are from the Bureau of Labor Statistics: https://www.bls.gov/oes/current/oes_nat.htm#13-0000. Average compensation of federal district judges is the annual salary as of 2006 divided by the number of annual work hours ($52 \times 40 = 2,080$). Judicial compensation data are from https://www.uscourts.gov/judges-judgeships/judicial-compensation.

[16] https://www.finra.org/sites/default/files/Education/p117487_0_0.pdf.

examine the probability an arbitrator is selected for a case conditional on the arbitrator appearing on the list. Thus, even though we only observe the list for a subset of our cases, the effective sample size is quite large (5,000+) because the appropriate unit of observation is at the case-by-list-by-arbitrator level. We also exploit exogenous variation in which arbitrators appear on the randomly generated list to estimate arbitrator fixed effects.

FINRA's BrokerCheck Data: We use BrokerCheck to obtain additional data on the respondent, as well as case details such as specific allegations that triggered the arbitration, requested damages, and arbitration award, all of which we discuss in more detail below. These data contain the employment, registration, and disclosure history for all individuals registered with FINRA. We collect the details of each financial adviser to construct a data set of the universe of financial advisers as described in Egan et al. (2019). Using these data, we also construct employment histories of arbitrators who have been employed in the financial industry in the past.

If a financial adviser is involved in an arbitration proceeding, the proceeding is reported on their disclosure record. Using unique case identifiers, we perfectly match the arbitration records reported in BrokerCheck to the arbitration case details reported in the Arbitration Awards Online database. We match 4,699 consumer arbitration disputes, which is the *universe* of arbitration disputes reported in BrokerCheck. These matched data represent our main data set. We describe the information we can observe in detail in the next section.[17]

III.B Summary Statistics: Cases and Arbitrators

Our primary unit of observation is at the case-by-arbitrator level. Roughly 10% of consumer complaints in arbitration involve multiple advisers and claims, and arbitrators separately assess damages across advisers. Consequently, we define an arbitration case at the case-by-adviser level but account for the potential correlation when computing standard errors. Our baseline data set consists of 4,699 consumer arbitration cases and 11,756 arbitrator-by-case observations. These cases involve substantial monetary amounts: mean and median damages requested are $758,648 and $240,000, respectively. The median award granted is 35% of the requested amount, with large differences in arbitration outcomes: the standard deviation is 60% (Figure 1; Table 1). The distribution is skewed to the right, with a mean award of 53% of damages, partially because awarded claims can exceed damages requested if punitive damages are awarded.

We report the summary statistics for our main outcome variables of interest in Table 1 and report summary statistics for the remaining variables in the Appendix. As discussed in the Appendix, we observe detailed information on the nature of the financial products, advisers involved in the dispute, and the specific allegations. In our baseline analysis, we classify allegations into 11 different categories across six different types of financial products. In general, the cases involve advisers consciously engaging in misconduct for their own financial benefit over their clients'. Common

[17]We observe the names of the arbitrators selected for 10,000 additional non-consumer arbitration cases, which are not reported BrokerCheck. Because non-consumer arbitration cases are not reported in BrokerCheck, we do not readily observe any information regarding them beyond the name of the arbitrator. These additional non-consumer arbitration cases allow us to examine an arbitrator's total case load in addition to the arbitrator's consumer arbitration case load.

allegations include misrepresentation (34% of cases), fraud (30% of cases), unauthorized activity (20% of cases), and churning (16% of cases). In the next section, we show that the specific allegations are highly correlated with arbitration awards. We use bag-of-words, a common natural language processing method (Gentzkow, Kelly, and Taddy, 2019; Bodoh-Creed, Boehnke, and Hickman, 2018), to further distill case allegations into 500 dummy variables for the most common words in case documents. Our data set also contains detailed information on the employment, registration, and disclosure history of each respondent–the financial adviser named in the consumer dispute. Because the securities industry is highly regulated, financial advisers must be licensed in order to engage in certain business activities, such as providing advice and selling mutual funds, insurance, and other products. Advisers can hold up to 61 different types of licenses, which helps us control for potential differences across arbitration cases.

We observe 4,992 unique arbitrators and have multiple observations (cases) for 2,714 arbitrators. The arbitration panel size consists of one or three arbitrators. Even though many arbitrators have little case experience, an arbitrator for an average case has overseen 11 total arbitrations (including non-consumer cases) and 4 consumer arbitrations. Therefore, most cases in our sample are overseen by arbitrators with extensive arbitration experience.

IV Motivating Evidence: Arbitrator Heterogeneity

The arbitrator selection process is based on the premise that arbitrators differ in how favorable they are to either party. These differences are why parties are allowed to eliminate arbitrators in the first place. We measure systematic differences between arbitrators in their awards. As our model in Section V highlights, differences between arbitrators arise in equilibrium as arbitrators compete to be selected. They therefore reflect both underlying differences in arbitrator beliefs and arbitrators' strategic responses to competition. Here, we present simple evidence that motivates the underlying assumptions in our structural model: (i) there are systematic differences in how arbitrators award cases (some arbitrators are more industry-friendly than others), and (ii) firms have an informational advantage in selecting favorable arbitrators.

IV.A Arbitrator Industry Friendliness

Ideally, we would observe two arbitrators ruling on identical cases, in the same location, and at the same point in time. An arbitrator who grants the lower award would be more industry-friendly, and the magnitude of the difference in awards would measure the extent of the difference in arbitrators' industry friendliness. We construct an empirical equivalent of this thought experiment to construct our measure of arbitrator industry friendliness. In this section, we construct our baseline estimates of arbitrator industry friendliness, which are transparent and easy to understand. We discuss additional estimates with richer text-based data, as well as robustness, in the Appendix.

We estimate a model of awards granted as a function of case characteristics, location and time of the arbitration, and, critically, the identity of the arbitrator:

$$Awarded_{ijlt} = \beta X_i + \mu_j + \mu_l + \mu_t + \epsilon_{ijlt}. \quad (1)$$

Observations are at the arbitrator-by-case level, where i indexes the arbitration case, j indexes the location, l indexes the arbitrator, and t indexes time. The dependent variable $Awarded_{ijlt}$ reflects the award granted divided by the award requested.[18]

We condition on case characteristics, X_i, to control for potential differences in the type of claim that is arbitrated and the merit of the claim. In the baseline specification, we control for the 11 different allegations and six different financial products covered in the case. We proxy for complexity of the case by the length of the case in sentences and words, as well as by the size of the reward requested. We control for the size of the case and the composition of the arbitration panel by including the number of arbitrators. To further control for case merit in our baseline specifications, we include the characteristics of the defendant financial adviser: the adviser's experience, the six most popular adviser qualifications/licenses, the adviser's total number of qualifications, and any past record of misconduct. In the Appendix, we do extensive robustness testing using natural language processing to control for more case characteristics.

We also include time (μ_t) and location (μ_j) fixed effects. In other words, we compare how industry-friendly an arbitrator is relative to other arbitrators in the same location and at the same point in time. Location fixed effects help control for possible geographic differences in claims. Because arbitrators are drawn from a pool based on the hearing location, these fixed effects allow us to compare an arbitrator to the pool of other arbitrators who could be potentially assigned to the case. Time fixed effects help account for aggregate differences in claims and institutional changes in the arbitration proceedings. As discussed in the Appendix, we find that these observable adviser and case characteristics explain 19% of the variation in awards. The case characteristics also predict awards in meaningful and intuitive ways. For example, cases in which the adviser's guilt is verifiable such as those involving fraud, churning, and selling unregistered securities tend to have higher awards on average.

The objects of interest are arbitrator fixed effects, μ_l, which measure whether an arbitrator, conditional on case characteristics, awards higher claims to consumers than other arbitrators at the same location and arbitration forum. These fixed effects account for persistent differences in how arbitrators award cases. An arbitrator l who is more industry-friendly than arbitrator l' will have a lower associated fixed effect $\mu_l < \mu_{l'}$. This measure is relative. We do not measure whether arbitrators awarded too much or too little relative to some "correct" amount. We measure if arbitrators awarded more or less relative to other arbitrators. In Section VI.B, we use a model to estimate the arbitrators' beliefs over what the fair or correct award would have been.

We report the estimated distribution of arbitrator fixed effects in Figure 2. We normalize the mean of fixed effects to match the average percent of awards granted in the data, 53%. As such, we can categorize arbitrators with a fixed effect below 53% to be, on average, more industry-friendly than other arbitrators. Including arbitrator fixed effects explains a substantial amount of the variation awards, increasing the R^2 from 19% to 55% in Eq. (1).[19] The F-test for differences among

[18]For robustness, we examine in the Appendix the log of awards granted through arbitration as a function of case observable characteristics and the log of awards requested.

[19]When computing the R^2, we only include arbitrators with multiple observations. The adjusted R^2 with arbitrator

arbitrators is jointly statistically significant at 1%. In other words, arbitrators play a significant role in determining arbitration awards.

Because individual arbitrator fixed effects are estimated with noise, the estimated differences in industry friendliness among arbitrators will be larger than the true underlying differences between them. We account for noise by constructing empirical Bayes estimates of arbitrator industry/consumer friendliness, $\widehat{\mu_l^{EB}}$ (e.g., Chetty, Friedman, and Rockhoff, 2014), which scales the distribution of ordinary least squares (OLS) fixed effects by a constant factor.[20] Although the variation in the empirical Bayes estimated fixed effects is 39% of the variation in OLS estimated fixed effects, the results indicate substantial differences across arbitrators. If an arbitrator who is more industry-friendly by one standard deviation is chosen to arbitrate the case, the damages awarded to the consumer will be 14pp smaller relative to the amount requested, holding other attributes of the case fixed. For example, if a consumer claims damages equal to the median amount,$240,000, such an arbitrator would award $33,600 less than the claimed amount. Overall, our results are consistent with the idea that the choice of arbitrator can have a meaningful impact on case outcomes.

Accounting for Selection: One important consideration in estimating arbitrators' industry friendliness is the presence of unobserved case characteristics, which determine the extent of the award and simultaneously affect which arbitrators are chosen. Such selection could bias our estimates of arbitrator fixed effects. We find limited evidence of selection on observable case characteristics. In the Appendix we re-estimate arbitrator fixed effects while accounting for selection. We construct an instrument exploiting the idea that, all else equal, arbitrators who live farther away are less likely to accept a case because of the travel involved. It is unlikely that arbitrator industry friendliness to a case is related to how far they are from the case venue. We also use the fact that we can observe the full randomly generated list of arbitrators for a subset of our data. The first instrument is the arbitrator's own distance to the case venue, where we also control for both the home location of the arbitrator and location of the case. The second instrument is constructed based on the average distance of all other arbitrators, and is thus limited to cases for which we can observe the entire randomly generated list of arbitrators for a case.[21]

Because we control for the location of the arbitrator and the case venue, the instrument is not picking up the idea that some venues have different cases, or that arbitrators' locations are correlated with their quality. In Appendix A, we show that distance affects an arbitrator's probability of presiding over a case, so our instruments are relevant. Accounting for potential selection has little effect on the estimated arbitrator fixed effects. This is not necessarily surprising given the way

fixed effects is 35%. We conduct a placebo test in which we randomly allocate arbitrators to cases and recompute arbitrator fixed effects. The R^2 in the placebo regression is substantially lower at 37% and the adjusted R^2 is 18%.

[20] $\widehat{\mu_l^{EB}} = \alpha(\widehat{\mu}_l - \bar{\mu})$, where $\bar{\mu}$ is the average OLS estimated fixed effect and $\alpha = \frac{F-1-\frac{2}{k-1}}{F}$, where F is the F-test statistic of a joint test of statistical significance of the fixed effects and k is the number of fixed effects under the assumption that the variance of the estimation error is homoskedastic (Cassella, 1992).

[21] For cases where the list is not observed, we set the value of the instrument to zero and include a dummy variable for such cases. Not observing the list for all cases does not invalidate our instrument but does reduce its statistical power. See Appendix A for further details.

arbitrators are quasi-randomly assigned to cases.

IV.B Arbitrator Appointments and Industry Friendliness

The previous section documents that some arbitrators are relatively more friendly to firms, while others are more friendly to consumers. The idea behind striking and ranking is that parties can reduce favoritism in awards by eliminating arbitrators most favorable to the other party. Here, we show that firms are better than consumers at choosing arbitrators because they eliminate those favoring the other side. As noted earlier, the list from which arbitrators are selected is randomly generated and audited for randomness by external auditors. If both sides were equally good at eliminating arbitrators, then neither side would have an advantage, and arbitrators' favoritism of a side would not help their selection. Alternatively, if firms were better than consumers at eliminating unfriendly arbitrators, then industry-friendly arbitrators would be chosen with a higher probability. Below we show that the latter is indeed the case, and that industry-friendly arbitrators are more likely to be selected.

IV.B.1 How Does Industry Friendliness Impact the Probability of Being Selected?

In this section, we show that industry-friendly arbitrators are more likely to be selected from a randomly generated list of potential arbitrators. For a subset of cases, we observe the entire list of arbitrators generated by the NLSS from which the parties can strike arbitrators. Observing the list allows us to exploit *within-case variation* in arbitrators' industry friendliness, and therefore implicitly control for case characteristics. We do so by estimating a multinomial logit regression model:

$$\Pr\left(D_{il}\right) = \frac{\exp\left(X_l \beta + \gamma \widehat{\mu_l^{EB}}\right)}{\sum_{n \in \mathcal{N}_i} \exp\left(X_n \beta + \gamma \widehat{\mu_n^{EB}}\right)}. \quad (2)$$

The probability that arbitrator l is selected to case i, indicated by the dummy variable $D_{il} = 1$, is a function of her consumer friendliness $\widehat{\mu_l^{EB}}$ and her characteristics $X_l \beta$ relative to the characteristics and consumer friendliness of all other arbitrators on the list. We measure consumer friendliness using our estimated empirical Bayes fixed effects ($\widehat{\mu_l^{EB}}$) from the previous section (Figure 2), where a higher fixed effect indicates higher consumer friendliness.[22] The set of arbitrators on the list for case i is denoted $\mathcal{N}_i$. The above probability highlights the importance of relative differences in arbitrator consumer friendliness. What matters for selection is not how consumer-friendly an arbitrator is in absolute terms $\widehat{\mu_l^{EB}}$, but how consumer-friendly an arbitrator is relative to the other arbitrators on the list $\widehat{\mu_n^{EB}}$. The parameter of interest is γ, which measures how consumer friendliness translates into the probability an arbitrator is selected. Under the hypothesis that firms are more informed than consumers, we would expect $\gamma < 0$ such that arbitrators who are more consumer-friendly (higher $\widehat{\mu_n^{EB}}$) are less likely to be selected.

We estimate the model on the subset of cases from Honigsberg and Jacob (2020) for which we observe the entire lists of potential arbitrators. We estimate a negative and significant relationship

[22] As the empirical Bayes adjustment only re-scales OLS fixed effects, it aids in interpreting the magnitudes but does not otherwise affect the estimates.

between how consumer-friendly an arbitrator is and her probability of being selected (Table 2). For cases involving multiple arbitrators and consequently multiple arbitration lists, we observe all arbitrator names that appear on the lists, but we do not observe whether they appeared on the chairperson list, public arbitrator list, or non-public arbitrator list. This introduces potential measurement error in the set $\mathcal{N}_i$. As a robustness check, we re-estimate our multinomial logit model restricting our attention to those arbitrators l who have ever been chairpersons in our data—under the assumption that a chairperson is picked from such a list—which reduces measurement error in the set $\mathcal{N}_i$, and present the estimates in column (2). Arbitrator fixed effects are measured with noise, especially for arbitrators with few awards. To reduce the extent of measurement error, in columns (3)-(4), we restrict our analysis to arbitrators who grant at least five awards in our baseline data and estimate the corresponding conditional likelihood (conditioning on the arbitrator observing at least five awards). In each specification, we estimate a negative and significant relationship between how consumer-friendly an arbitrator is and her probability of being selected. The results in column (4) indicate that, conditional on being on the list, a one standard deviation decrease in consumer friendliness (a 14pp decline in $\widehat{\mu_n^{EB}}$) is associated with a 33% (7pp) increase in the probability of being selected.[23] In other words, the results in column (4) indicate that an arbitrator who grants awards that are below average $\left(\widehat{\mu_n^{EB}} < \overline{\widehat{\mu_n^{EB}}}\right)$ is roughly 50% more likely to be selected than an arbitrator who grants above-average awards $\left(\widehat{\mu_n^{EB}} > \overline{\widehat{\mu_n^{EB}}}\right)$.

IV.B.2 Extension to the Full Sample

We do not observe the list of arbitrators from which the parties can strike arbitrators for the full sample of cases. Nevertheless, we want to illustrate the idea that the more frequent selection of industry-friendly arbitrators is robust to the extended sample. We examine how an arbitrator's estimated fixed effect $\widehat{\mu_l^{EB}}$ impacts her probability of being selected in the following regression:

$$Number\,of\,Cases_l = \beta X_l + \gamma \widehat{\mu_l^{EB}} + \mu_g + \epsilon_l. \tag{3}$$

Our observations are at the arbitrator level, where l indexes the arbitrator.[24] $Number\,of\,Cases_l$ measures the number of cases an arbitrator oversees in her career. The key independent variable of interest is again the arbitrator's fixed effect $\widehat{\mu_l^{EB}}$. The term X_l is a vector of arbitrator controls and includes the number of years she has been active as an arbitrator in our sample.

We also include geographic region fixed effects (μ_g) to account for the fact that FINRA randomly selects a list of arbitrators based on the geographic region of the case. This fixed effect captures regional differences in case load and arbitrator fixed effects. FINRA randomly constructs its list of arbitrators based on the arbitrator's location and public arbitrator status. This determines the pool of arbitrators who compete to be selected. We already control for the location of the case

[23] We compute the marginal effect as $0.07 = -0.42 * \bar{p} * (1 - \bar{p})$ where $\bar{p} = 0.21$ is the average probability of being selected in the sample corresponding to column (3).

[24] An arbitrator enters our data as soon as she oversees her first case and remains in the data until 2015. We control for the number of years she's been active and the number of cases in the data set she has overseen in order to adjust different attrition rates among arbitrators.

when estimating arbitrator fixed effects. We also interact geographic region fixed effects with whether the arbitrator is a public arbitrator (i.e., non-industry-affiliated), non-public arbitrator, or eligible chairperson, which allows us to compare differences in how industry- or consumer-friendly the arbitrator is within the pool of eligible arbitrators. We estimate Eq. (3) using linear regression in our baseline specification, but we also estimate the model using Poisson and negative binomial regressions to account for the fact that the dependent variable is a count variable. The estimates presented in Table 3, column (3) indicate that an increase of one standard deviation in an arbitrator's industry friendliness is associated with the arbitrator being chosen at an 8.5% higher rate.

We find that, controlling for case characteristics, arbitrators who grant larger awards to consumers are less likely to be selected. This is despite the fact that they have an equal chance of making it on the list, which is randomly generated. These results suggest that consumer-friendly arbitrators face higher chances of elimination than industry-friendly arbitrators and that firms have an informational advantage in striking arbitrators. In Appendices A and B, we show the result is robust to alternative measures of industry friendliness such as backward-looking measures, measures constructed using machine learning, and accounting for selection.

IV.C Interpretation

IV.C.1 Sophistication

Our central hypothesis is that firms are more sophisticated than consumers in striking arbitrators, which results in industry-friendly arbitrators being chosen more often. One source of firms' advantage is that they have more experience in arbitration than the average consumer (see Nichols, 1999; Black and Gross, 2002; Barr, 2014; Silver-Greenberg and Gebeloff, 2015). Anecdotal evidence suggests that brokerage firms often maintain internal arbitrator "strike lists." Additional experience in arbitration should allow firms to design a better "strike list," as well as other information on selecting arbitrators. Similarly, consumers can acquire expertise by hiring attorneys from the Public Investors Advocate Bar Association (PIABA) who specialize in arbitration.

Here, we provide more direct evidence that parties' sophistication in arbitration helps them choose more favorable arbitrators. We examine the fixed effect of the arbitrator l selected to case i as a function of firm and consumer sophistication:

$$\widehat{\mu_{il}} = \phi_0 + \phi_1 Attorney_i + \phi_2 PIABA_i + \phi_4 Trust_i + \phi_5 Firm_Experience_i + \varepsilon_{il}. \quad (4)$$

Observations are at the arbitrator-by-case level, where i indexes the arbitration case and l indexes the arbitrator. The dependent variable $\widehat{\mu_{il}}$ measures the fixed effect of the arbitrator l selected for case i. The estimated arbitrator fixed effects correspond to Eq. (1), where a higher fixed effect implies a more consumer-friendly arbitrator. The independent variable $Attorney_i$ indicates whether the consumer used an attorney, $PIABA_i$ indicates whether the consumer used a PIABA attorney, and $Trust_i$ indicates whether the consumer is part of a trust. $Firm_Experience_i$ indicates whether the firm has above-median arbitration case experience in terms of the number of consumer arbitration cases in which a firm is involved. The omitted category comprises consumers who are

self-represented/do not use an attorney, are not part of a trust, and face a less experienced firm.

We find that consumers who use attorneys specialized in arbitration, PIABA attorneys, have their cases overseen by arbitrators whose awards are 8.14pp higher than in cases where consumers are self-represented (Table 4, column 1).[25] Moreover, our results also suggest that more sophisticated consumers choose more consumer-friendly arbitrators: the arbitrators chosen when consumers are a part of a trust grant 5.83pp higher awards on average (column 1). In other words, parties' expertise in arbitration allows them to select more favorable arbitrators. We also find some modest evidence that firms with more experience in arbitration select more industry-friendly arbitrators.

IV.C.2 Alternative Explanations

The main alternative hypothesis to our findings is that firms and consumers are equally sophisticated. Some arbitrators appear industry-friendly in the data because they are systematically selected to cases with unobservable characteristics that merit lower awards. These arbitrators are selected more often—not eliminated by either consumers or firms—because their skills make them a good fit for future cases. If we controlled for these case characteristics, there would be no pro-industry bent. As noted in Section IV.A, we exploit the randomly generated list of arbitrators and the distance of arbitrators from the arbitrator case to account for such selection in computing arbitrator fixed effects. Moreover, in Section IV.B, we show that even holding the case fixed, i.e., using *within-case* variation, the more industry-friendly arbitrators from the list are chosen more frequently. Here, we discuss additional evidence that selection is not an issue, which we present in detail in Appendices A and B. We exploit the 2007 change in FINRA rules governing arbitration, which reduced the number of arbitrators that each party could strike. If firms' advantages in arbitration are indeed driven by their ability to choose which arbitrators to eliminate, then restricting the number of arbitrators that each party can eliminate should reduce the impact of firms' informational advantages. We find that the effect of firms' informational advantages declines after the reform by more than half. We do not find evidence that the composition of cases overseen by consumer- or industry-friendly arbitrators changes (which would be required to generate the results surrounding the 2007 rule change) if both parties were equally informed.

We also perform Altonji et al. (2005) and Oster (2019) style tests and compute arbitrator fixed effects iteratively by conditioning on richer case characteristics. Using these case characteristics, we can explain approximately 85% of variation in case awards. Despite large changes in R^2, we find little effect on the estimated arbitrator fixed effects (see Appendix B, Figure A2, Table A6).

Lastly, we find that an arbitrator's personal experiences shape their bias. Roughly 40% of arbitrators have work experience in the securities industry. We find that arbitrators who have been involved in customer disputes themselves tend to grant lower awards. Conversely, arbitrators who were fired from the industry (have an employment termination flag on their regulatory file) tend to grant higher awards to consumers. These results suggest that the estimated fixed effects are driven

[25] One might ask why so many consumers choose non-PIABA attorneys. We speculate that finding specialized attorneys already requires some information/sophistication from consumers, which is also the reason why they need a specialized attorney in the first place.

by inherent differences across arbitrators rather than by differences in case characteristics.

V A Model of Arbitrator Competition

Our empirical analysis suggests that firms possess an informational advantage in choosing arbitrators from a given pool of arbitrators. Arbitrators who are more industry-friendly than other arbitrators are more likely to be selected to a case. Because arbitrators can choose how they rule on a case, our analysis suggests that arbitrators have an incentive to be more industry-friendly than other arbitrators. In other words, arbitrators compete with each other in industry friendliness. Intuitively, this competition between arbitrators should change the overall industry and consumer friendliness of the arbitrator pool as a whole. By design, the analysis we perform in the first part of the paper cannot measure the extent of the competition effect. To eliminate as much variation as possible, we compare how industry-friendly arbitrators are *relative to each other*. Fixed effects sweep away the average level of pro-industry or pro-consumer tilt of the arbitration pool. To isolate the tilt in the arbitration pool as a whole, we next develop a stylized but quantitative model of consumer arbitration. The model is informed by our empirical findings and the institutional details laid out in Section II, but explicitly models how arbitrators compete on industry friendliness.

We use the model for several purposes. First, the model allows us to recover the beliefs of arbitrators, i.e., the awards that arbitrators would have chosen absent incentives provided by the arbitration selection mechanism. We use the estimates to quantify the full extent of firms' informational advantage. We then decompose firms' informational advantage into two distinct components, i.e., advantage of striking from a fixed arbitrator pool, and the equilibrium effect on the pool itself. Second, we use the calibrated model to show that these insights are critical in assessing changes to arbitration design both qualitatively and quantitatively. While competition between arbitrators can be desirable if both parties are informed, it helps firms and hurts consumers when firms are better informed. We use the calibrated model to compute gains or losses to consumers from alternative arbitrator selection schemes proposed by regulators, which have been touted to improve arbitration outcomes for consumers (FINRA Notice 14-49, 2014). Policy changes that aim to help consumers but ignore the informational advantage of firms end up hurting consumers instead. As we discuss in Appendix E, these insights are applicable to consumer arbitration proceedings more generally and to other arbitrator selection mechanisms.

V.A Setup

The consumer (claimant) and firm (respondent) are arbitrating a claim that will be overseen by one of the available arbitrators, who determines the award. The timing is as follows. First, arbitrators choose their slant, i.e., how industry- or consumer-friendly they are going to be. In choosing slant, arbitrators commit to how they will award a case to the participants. Second, following the institutional design for arbitrator selection, a list of arbitrators is randomly constructed using the pool of all available arbitrators. The consumer and firm can simultaneously strike a limited number of arbitrators from the list. Among the remaining arbitrators, one is selected randomly. Lastly, the selected arbitrator is paid a fee for arbitrating the case, and awards are paid to the parties. Next,

we describe the incentives and information structure of the problem in more detail.

V.A.1 Consumers, Firms, and Arbitrators

Consumers and Firms: The award is the share of the requested damages $a_G \epsilon \left[\underline{a}, \bar{a}\right]$ that is granted to the consumer. Because the award is just a transfer from the firm to the consumer, it is a zero-sum game. We denote the payoff to the consumer as $U_C = a_G$ and the payoff to the firm as $U_R = -a_G$. For simplicity of exposition, we assume both parties are risk-neutral. Risk aversion does not change the parties' strategies for selecting arbitrators or the resulting equilibrium. It would affect parties' preferences over alternative arbitration mechanisms (policies), which we incorporate in the Appendix.

Arbitrators: Arbitrators trade off monetary incentives of being selected on a case against the psychological costs of departing from their view of a "fair" award. This allows us to nest the extreme cases of arbitrators who are motivated purely by monetary incentives, as well as arbitrators who are motivated only by fairness concerns. As we discuss below, both features are important in order to capture arbitrator behavior in the data.

Conditional on the observable case characteristics, each arbitrator has an inherent belief $b_i \epsilon \left[\underline{b}, \bar{b}\right]$ regarding the fair award for the arbitration case.[26] We can think of these beliefs as innate characteristics that arbitrators bring to the case that determine how an arbitrator would rule on the case in the absence of monetary incentives. These could be formed based on their prior work experience, education, upbringing, or personal interaction with the industry. For example, in the Appendix, we show that arbitrators who were fired from the financial advisory industry tend to be more consumer-friendly and grant higher awards. The distribution of beliefs among arbitrators in the population is $F(\cdot)$; the density $f(\cdot) = F'(\cdot)$ is continuous and strictly positive everywhere.

Arbitrators earn a fee (fee) if they are selected to arbitrate a case. The probability that a given arbitrator i will be selected depends on the firm's and consumer's expectations of the award a_i that the arbitrator would grant were she selected, or the arbitrator's "slant." For simplicity, we assume that arbitrators can pre-commit to what they would award for a case a_i before being selected on the panel. The idea is that, just as in the data, arbitrators can choose their slant, i.e., how industry-friendly they want to be. Instead of modeling the reputation-building process, which is not the focus of this paper, we assume that arbitrators choose their slant before arbitrating a case. To keep the notation simple, an arbitrator's slant directly commits them to an award rather than to a noisy unbiased signal of the award, which would not alter the analysis.

Arbitrators can have a sense of fairness. When their decisions depart from their beliefs of fair awards, $a_i \neq b_i$, they suffer a disutility of $\theta |a_i - b_i|$. The parameter θ measures the weight that an arbitrator places on fairness relative to the monetary payoffs from arbitration. A lower θ implies

[26] The idea that arbitrators have an inherent notion of a "fair" outcome is standard (Crawford, 1979; Farber 1979, 1980; Farber and Katz, 1979; Ashenfelter and Bloom, 1984; De Clippel et al., 2014). That arbitrators only differ with respect to their belief over a fair award is consistent with the baseline assumptions in the judge literature (e.g., Gennaioli and Shleifer, 2008; Gennaioli and Rossi, 2010; and Gennaioli, 2013). Allowing arbitrators to be multidimensional (differ in terms of their preference for fairness or the variance of beliefs) would not change the underlying economics of the model as long as firms are informed and consumers are uninformed.

that arbitrators care more about monetary payoffs. In the extreme case that arbitrators care only about monetary payoffs, $\theta = 0$. As $\theta \to \infty$, arbitrators are motivated only by their fairness beliefs and do not respond to monetary incentives.

Let $G(\cdot)$ be the equilibrium distribution of arbitrators' chosen slants, and denote the equilibrium probability that an arbitrator with slant a_i is chosen as $\Gamma(a_i, G(\cdot))$. An arbitrator's expected utility depends on her expected probability of being selected on the case, Γ, the fee she earns from arbitrating, fee, and the award she grants relative to her beliefs:

$$U(b_i, a_i) = \Gamma(a_i, G(\cdot))(fee - \theta|a_i - b_i|). \tag{5}$$

Sophistication: Firms are frequently large institutions that engage in arbitration repeatedly, while consumers typically engage in arbitration once. Consistent with our empirical setting and analysis, we assume that firms are the informed party. They recognize arbitrators' slants and can therefore predict their awards when choosing among them. Consumers, however, are uninformed and do not observe or anticipate how a specific arbitrator will award a case.[27] While consumers can be very unsophisticated, arbitrators need to understand the equilibrium distribution of arbitrator slant as well as incentives facing them. Unlike consumers, arbitrators engage in arbitration repeatedly. Furthermore, roughly 40% of the arbitrators in our sample have work experience in the financial services industry and may have had personal experience with arbitration.[28]

V.A.2 Arbitration Selection Process and Uninformed Consumers

N risk-neutral arbitrators are randomly drawn from the population of arbitrators $A = \{a_1, a_2, ..., a_n\}$, and the "list" is presented to the parties. Without any loss in generality, arbitrators are indexed such that the most industry-friendly arbitrator who grants the lowest awards is indexed by 1 and the least industry-friendly arbitrator who grants the highest awards is indexed by n such that $a_1 < a_2 < ... < a_n$. Both the consumer and firm simultaneously submit $k < \frac{n}{2}$ arbitrators to be struck from the list of available arbitrators. Among the remaining arbitrators, one is chosen randomly. The chosen arbitrator j grants the award according to their chosen slant $a_G = a_j$. Firms observe the slant $a_1, a_2, ..., a_n$ of each arbitrator appearing on the randomly generated list. Consumers, being uninformed, do not observe the slant. Given the equilibrium distribution of slant $G(\cdot)$, we denote $\widetilde{G}(\cdot)$ the distribution of awards granted through arbitration, $a_G \sim \widetilde{G}(\cdot)$.

V.A.3 Equilibrium Definition

We study a pure monotone strategy symmetric Bayesian Nash equilibrium, which is characterized by the optimal behavior of consumers, firms, and arbitrators. Firms and consumers optimally strike arbitrators from the arbitration pool A to maximize their utility, holding the strategy of the opposing

[27] Consumers can understand all other aspects of the arbitration process, including the distribution of arbitrators' beliefs and equilibrium slant.

[28] Consumers are sent an Arbitration Disclosure Report for each arbitrator that appears on the list of potential arbitrators. It lists each arbitrator's employment, education, any conflicts, and a list of the prior case IDs. Using this information, a consumer could in principle replicate our analysis to estimate arbitrator fixed effects. We think this is unlikely, given that even modest search costs prevent consumers from optimizing in much simpler settings (e.g., Hortacsu and Syverson, 2004; Woodford and Hall, 2011; Gurun et al., 2016; Egan, 2019; Agarwal et al., 2020).

party fixed. Arbitrators maximize their expected utility (Eq. 5) by choosing their slant and taking the strategies of firms, consumers, and other arbitrators as given.

V.B Equilibrium: Arbitrator Selection and the Arbitrator Pool

V.B.1 Arbitrator Selection from a Fixed Pool

We first analyze the advantage of informed firms over uninformed consumers, taking the arbitrator pool as given, i.e., given the equilibrium distribution of slant, $G(\cdot)$. The incentives of firms and consumers are straightforward. The firm, being informed, will find it optimal to always strike the arbitrators with the k highest (most consumer-friendly) slant. By contrast, uninformed consumers randomly strike k arbitrators. An arbitrator is randomly selected from the pool of eligible (non-stricken) arbitrators. Then the equilibrium probability that an arbitrator with slant a_i will be selected, given the distribution of other arbitrator slant in the population, is:

$$\Gamma(a, G(.)) = \frac{1}{n-k} P(a_i; 1, n-k, n) \tag{6}$$

where $P(a_i; l, m, n) = \sum_{j=l}^{m} \frac{(n-1)!}{(j-1)!(n-j)!} G(a_i)^{j-1}(1-G(a_i))^{n-j}$ denotes the probability that the arbitrator is between the $l'th$- and $m'th$- order statistics among a sample of n arbitrators.

This expression highlights the role of different information structures in the selection of arbitrators for a given arbitrator pool. Firms strike the k most consumer-friendly arbitrators with the highest slant. Thus, an arbitrator is selected only if her slant is one of the $n-k$ lowest-order statistics among the n arbitrators. The probability an arbitrator is selected is then decreasing in her slant a, and arbitrators who are more industry-friendly are more likely to be selected. We illustrate the striking effect in Figure 3. The distribution of awards granted $\widetilde{G}(\cdot)$ is smaller (in terms of first-order stochastic dominance) relative to the equilibrium distribution of arbitrator slant, $G(\cdot)$. Intuitively, this is the effect documented by our reduced form analysis in Section IV.B.

It is useful to also provide the benchmark if both consumers and firms are informed about arbitrator slant, which is frequently used to guide policy discussion. If consumers are informed, the arbitrators in either tail of the distribution face elimination, and the probability that an arbitrator is selected becomes:

$$\Gamma(a, G(.)) = \frac{1}{n-2k} P(a_i; k+1, n-k, n).$$

Informed consumers remove k arbitrators with the most pro-industry (lowest) slant, and firms remove the k arbitrators with the highest slant. Thus, an arbitrator is selected only if she is one of the $k+1 : n-k$ middle-order statistics of the distribution of slant among the set of N arbitrators appearing on the list. The striking mechanism helps eliminate extreme outcomes, as the closer an arbitrator's slant (a) is to the median, the higher the probability she is selected. This discussion illustrates that assuming the extent to which consumers are informed in arbitration has important consequences for arbitration outcomes.

V.B.2 Arbitrator Pool: Equilibrium Choice of Slant

Our discussion above holds the distribution of arbitrator slant fixed. In other words, it does not account for arbitrators' incentives to be selected on the panel. Arbitrators, however, can choose how they rule on cases, and can therefore choose how consumer- or industry-friendly they want to be. We show that competition among arbitrators can be desirable if both parties are equally informed. But in the presence of an information gap, competition leads to the whole pool of arbitrators becoming industry-friendly. Next, we characterize how the primitives of the model affect the severity of the equilibrium shift in the pool.

When arbitrators choose slant, they trade off two forces. On the one hand, they want to be selected on the arbitration panel (increase $\Gamma(a_i, G(\cdot))$) to earn the arbitration fee fee. To do so, they want to choose a slant that will minimize their chance of being struck from the arbitrator panel by an informed firm or consumer. This probability is determined by their slant *relative* to other arbitrators. On the other hand, choosing awards that depart from their convictions, $a_i - b_i$, causes disutility. Arbitrator i with inherent belief b_i chooses slant a_i to maximize her expected utility given the choices of other arbitrators:

$$max_{a_i}\Gamma(a_i, G(\cdot))\left(fee - \theta\left|a_i - b_i\right|\right). \quad (7)$$

We look for a monotone equilibrium: arbitrators with more consumer-friendly beliefs choose a more consumer-friendly slant. For ease of intuition, assume that $\Gamma\left(a_i; G\left(\cdot\right)\right)$ is differentiable. The corresponding first-order condition can be written as:

$$|a_i - b_i| = \frac{fee}{\theta} - sgn(a_i - b_i) \times \frac{\Gamma\left(a_i; G\left(\cdot\right)\right)}{\gamma\left(a_i; G\left(\cdot\right)\right)} \forall a_i \neq b_i \quad (8)$$

where $\gamma(a_i; G(\cdot) = \frac{\partial\Gamma(a_i;G(\cdot))}{\partial a}$. An arbitrator's choice of slant relative to their underlying beliefs b_i depends on the trade-off between the costs and benefits of slant. Firms eliminate the k most consumer-friendly arbitrators from the pool. Therefore, the probability an arbitrator is selected is decreasing in her slant a, $\gamma(a, G(\cdot)) < 0$. This implies that $a_i \leq b_i$. The choice in slant becomes:

$$a_i = \min\left\{b_i - \frac{fee}{\theta} - \frac{\Gamma\left(a_i; G\left(\cdot\right)\right)}{\gamma\left(a_i; G\left(\cdot\right)\right)}, b_i\right\}. \quad (9)$$

This expression shows the extent of an individual arbitrator's pro-industry slant. All arbitrators choose their slant to be more industry-friendly than their underlying belief, $a_i < b_i$, as long as $\frac{fee}{\theta} + \frac{\Gamma(a_i;G(\cdot))}{\gamma(a_i;G(\cdot))} > 0$. The term $\frac{\Gamma(a_i;G(\cdot))}{\gamma(a_i;G(\cdot))}$ measures the inverse of the relative change in the probability of being selected for a marginal change in an arbitrator's slant, holding other arbitrators' slant choices fixed. The term $\frac{fee}{\theta}$ is the fee that the arbitrator earns in utility terms if she is selected. Arbitrators who choose their slant equal to their beliefs ($a_i = b_i$) will award what they think is fair if the marginal benefit of slanting their award is less than the marginal cost when $a_i = b_i$ such that $\frac{fee}{\theta} + \frac{\Gamma(b_i;G(\cdot))}{\gamma(b_i;G(\cdot))} \leq 0$. In other words, arbitrators will find it optimal to skew pro-industry and grant lower awards relative to their true beliefs. We can express the distribution of equilibrium

probabilities as a function of the equilibrium distribution of slant:

$$a_i = \min\left\{b_i - \frac{fee}{\theta} - \sum_{j=1}^{n-k} \binom{n-1}{j-1} \frac{(n-1)!}{(j-1)!(n-j)!} \frac{G(a_i)(1-G(a_i)}{g(a_i)\,(j-1-(n-1)G(a_i)}, b_i\right\}$$

This equation is at the center of our estimation approach in Section VI. We also compute a closed-form expression for the equilibrium distribution of arbitrator slant as a function of model primitives: the distribution of beliefs, the size of the list from which arbitrators are chosen, and the number of strikes from the list (see derivation in Appendix D):

$$a_i = \min\left\{b_i - \frac{fee}{\theta} + \frac{\int_{b_i}^{\bar{b}} \Gamma(\tilde{b}, F(\cdot))d\tilde{b}}{\Gamma(b, F(\cdot))}, b_i\right\}. \tag{10}$$

We use the closed-form expression (Eq. 10) when computing counterfactual equilibria, which link the model with actual policy proposals in Section VI.C. We illustrate the effect of competition among arbitrators in Figure 3, which corresponds to our model estimates (discussed in Section VI): the distribution of arbitrator slant $G(\cdot)$ is industry-friendly relative to the distribution of arbitrator beliefs based on what a "fair" award should be, $F(\cdot)$, in the sense of first-order stochastic dominance. Together, the striking effect and the competition effect result in a distribution of awards for consumers that are lower than what arbitrators believe is fair: the distribution of awards granted $\widetilde{G}(\cdot)$ is stochastically dominated by the underlying distribution of arbitrator beliefs $F(\cdot)$.

V.B.3 Special Case: Race to the Bottom and Statistical Exchangeability

Here we present a special case in which arbitrators care only about monetary incentives ($\lim \theta \to 0$). This case highlights why the competition effect can be very large. In fact, if arbitrators only care about monetary incentives, the competition effect results in a race to the bottom: all arbitrators choose the most industry-friendly arbitrator slant possible $a_i = \underline{a}$. To see why, imagine that in equilibrium some arbitrators are more industry-friendly than others – $G(\cdot)$ features different arbitrator slants. Then there is an arbitrator with the most pro-consumer slant, $\tilde{a}$. This arbitrator will certainly be eliminated by the informed firm, so she will never be selected on an arbitration panel. If she instead chooses a slant that is more industry-friendly than that of other arbitrators, then she will be selected with a positive probability, increasing her expected monetary payoff. She has no fairness concerns, so there is no utility cost to changing her slant. Choosing the most industry-friendly slant is therefore clearly a profitable deviation. Because every arbitrator wants to be the most industry-friendly, there is a "race to the bottom." The only equilibrium is one in which all arbitrators, regardless of their beliefs of what is fair, choose a slant that is as industry-friendly as possible, $a_i = \underline{a}$.

This example also highlights why "statistical exchangeability" is not a good fairness criterion when one party holds an information advantage. Statistical exchangeability of arbitrators implies that the identity of the arbitrator (their pro-industry or pro-consumer slant) does not affect arbitration outcomes. As our example above illustrates, with purely monetary incentives, the competition effect results in statistical exchangeability whether consumers are informed or not: all arbitrators reach the same decision. If only firms are informed, the resulting decision can be quite "unfair,"

since all arbitrators are as industry-friendly as possible. In other words, statistical exchangeability could be a necessary but insufficient condition for fairness.

VI Informational Advantage and Policy Analysis

In this section, we calibrate the model to quantify the advantage of the informed party in the current system and decompose the advantage into two components: the striking advantage from a fixed arbitrator pool and the competition effect that shapes the pro-industry tilt of the arbitrator pool. In Section VI.C we use this model to study whether and to what extent different arbitrator selection schemes benefit the industry over consumers. Rather than considering a complete redesign of the system, we examine changes to the features of the existing system of choosing and compensating arbitrators, quantitatively linking the model with current and past policy proposals.

VI.A Calibration

We calibrate the model using an approach that resembles the methodology developed in the auction literature by Guerre, Perrigne, and Vuong (2000). We use the observed distribution of arbitrator fixed effects to recover the underlying distribution of slant $G(\cdot)$ and the underlying distribution of arbitrator beliefs $F(\cdot)$. The idea is that an arbitrator's choice of slant in equilibrium is a best response to other arbitrators' choices of slant. From the data, we can measure other arbitrators' equilibrium choices of slant a_i, as we describe below. Given the other arbitrators' equilibrium choices of slant, we can infer every arbitrator's true beliefs b_i from her own choice of slant a_i as follows:

$$b_i = \max\left\{a_i + \frac{fee}{\theta} + \sum_{j=1}^{n-k}\binom{n-1}{j-1}\frac{(n-1)!}{(j-1)!(n-j)!}\frac{G(a_i)(1-G(a_i))}{g(a_i)\left(j-1-(n-1)G(a_i)\right)}, a_i\right\}. \quad (11)$$

In order to recover the true beliefs, b_i, for an arbitrator with slant a_i, we need to observe the arbitrator fee, disutility from deviating from one's beliefs θ (which we have to calibrate), and the unconditional density and distribution of arbitrator slant $G(\cdot)$ and $g(\cdot)$. We parameterize and calibrate the model as follows.

We estimate the distribution $G(\cdot)$ and density $g(\cdot)$ of slant nonparametrically in the data. We use the empirical Bayes estimates of arbitrator fixed effects to estimate the equilibrium distribution of slant, where we restrict the data to the post-2007 reform when FINRA limited the number of strikes to four. In the data, we observe the distribution of slant, conditional on arbitrators being chosen, $\tilde{G}(\cdot)$, rather than the distribution of arbitrators in the population $G(\cdot)$. Because k consumer-friendliest arbitrators are removed from the randomly generated list of n arbitrators, we observe the distribution of slant a_i conditional on a_i not being one of the k highest-order statistics. Formally, the distribution $\widetilde{G}(\cdot)$ represents a weighted average of the $n-k$ first-order statistics of $G(\cdot)$. To obtain the unconditional distribution of slant $G(\cdot)$, we proceed in two steps. We first estimate $\widetilde{G}(\cdot)$ from the data nonparametrically using the empirical distribution function. Then, we use the model

to invert the underlying distribution given firms' striking behavior:

$$\widetilde{G}(a_i) = \sum_{i=k}^{n-1} \left(\sum_{j=n-i}^{n} \frac{n!}{j!(n-j)!} G(a_i)^j (1 - G(a_i))^{n-j} \right) \tag{12}$$

numerically solving for $G(\cdot)$. We also estimate the density of the slant distribution $g(\cdot)$. The density of arbitrator slant among selected arbitrators $\tilde{g}(a)$ is equal to the unconditional density $g(a)$ multiplied by the probability of being selected $n \times \Gamma(a, G(\cdot))$:$\tilde{g}(a) = g(a) \times n \times \Gamma(a, G(\cdot))$. We estimate $g(\cdot)$ nonparametrically using kernel density estimation, where we weight each observation by our estimates of the inverse probability of being selected $\frac{1}{\Gamma(\widehat{a, G}(\cdot))}$.[29]

We next calibrate the parameters fee and θ, which measure the trade-off between monetary incentives and the cost of deviating from arbitrators' beliefs. Only their relative trade-off $\frac{fee}{\theta}$ matters in equilibrium (Eq. 11). Arbitrators earn $\$300$ per hearing, and the typical case lasts four days (FINRA Rule 12214), so we set the per-case fee equal to $fee = \$1,200$. We calibrate θ, which reflects the cost of deviating from an arbitrator's true beliefs using the 2007 rule change (the change decreased the number of strikes from nine to four). Reducing the number of strikes curtails an arbitrator's incentive to slant their decisions in favor of the industry. Consistent with this intuition, we estimate in Appendix B that after the 2007 rule change, arbitrators increased awards by roughly 5pp. We calibrate the model to match this average change in awards when the number of strikes shifts from nine to four in the model. This calibration yields $\theta = 10,000$. This estimate implies that arbitrators are willing to lower their awards below what they think is fair by 1pp for an extra $100 increase in income. For example, suppose the arbitrator believed that a fair award was to grant 100% of the amount requested. The arbitrator would be willing to grant an award of 90% in exchange for a $1,000 increase in income.

Because potential non-pecuniary benefits of being an arbitrator are difficult to measure (see Section II.B), we experiment with alternative calibrations in which we scale $\frac{fee}{\theta}$ by 50% and 150% in the Appendix. The alternative parameterizations yield similar inferences in Section VI.C. Once we have obtained the magnitudes of disutility from deviating from one's beliefs θ, arbitrator compensation fee, and the unconditional density and distribution of arbitrator slant $\widehat{G(a)}$ and $\widehat{g(a)}$, we use Eq. (11) to compute the density of arbitrators' beliefs of what a fair award would be, $(\widehat{f(b)})$.

VI.B The Cost of "Industry-Friendly" Arbitration for Consumers

VI.B.1 Overall Cost of "Industry-Friendly" Arbitration for Consumers

We use the calibrated model to evaluate the cost to consumers because firms have an informational advantage. Figure 4 displays our nonparametric estimation results. We also present a parametric model in which the underlying distribution of beliefs follows a gamma distribution, which is computationally tractable for computing counterfactual equilibria. We estimate the parametric distribution of beliefs via maximum likelihood to match the nonparametrically estimated distribution of beliefs, and present them in Figure 3. We consider two ways to quantify firms' informational

[29]Specifically, we use a Gaussian kernel and a smoothing parameter of 5%, which is in line with Silverman's Rule of Thumb (1986).

advantage.

We first compare arbitration outcomes under the current system to how consumers would fare if arbitrators would be selected to the cases randomly, like judges in some courts. Under the latter scenario, parties have no input in the selection of arbitrators, and the distribution of awards granted simply reflects the distribution of arbitrator beliefs, as arbitrators no longer have any incentive to slant awards. Figure 4 shows that the density of arbitrator beliefs $f(\cdot)$ is shifted to the right of the density of arbitration awards $\widetilde{g}(\cdot)$. Thus, if arbitrators were to be randomly assigned, the distribution of awards would shift from $\widetilde{g}(\cdot)$ to $f(\cdot)$ and become more consumer-friendly. The average award under the current system is 53% of the amount requested. If neither party had any input into the selection process, our estimates imply that the mean award would be 61%. Given that the average award is on the order of $750,000, the model estimates suggest that the current arbitrator selection scheme costs consumers roughly 8pp, or $60,000. The shift in the distribution of awards affects the top half of the distribution more: the 10^{th} percentile award increases from 32% to 33%, while the 90^{th} percentile increases from 73% to 89%. In other words, the current arbitration system especially decreases the propensity of large awards to consumers relative to the system in which arbitrators are randomly chosen. The current mechanism does reduce the variance of outcomes, which is an often touted benefit of the arbitration selection process. The standard deviation of outcomes is reduced by 24% relative to a system in which parties have no input into arbitrator selection. These results show the extent to which the current arbitration scheme results in a biased distribution of arbitration awards relative to the underlying distribution of beliefs of fair awards.

The second method to benchmark the effect of firms' informational advantage in arbitration outcomes is to estimate outcomes under the assumption that consumers are as informed as firms (Figure 5). In Appendix D, we show that when both parties are informed, the arbitrator selection mechanism results in a distribution of arbitration awards that is a median preserving contraction of arbitrators' underlying beliefs, such that the median award equals the median belief. The mean awards would be roughly 60%, similar to the scenario where parties have no input in arbitrator selection. Importantly, with informed consumers, the selection mechanism reduces the variance of arbitration awards by 68% relative to the variance of beliefs. In that scenario, the current arbitration selection mechanism would be effective in reducing the variance of arbitration awards and would be fair in the sense that the median award reflects the median belief. This is in sharp contrast to the scenario where only firms are informed, resulting in a lower mean and median award relative to the underlying distribution of arbitrator beliefs.

VI.B.2 Decomposition: Striking and Competition

The current arbitrator selection scheme costs consumers roughly 8pp of awards, or $60,000. Figure 4 decomposes firms' advantage into two components. The first is the advantage that firms derive in striking arbitrators from a given pool. This is measured as the shift between the awards granted $\widetilde{g}(\cdot)$ and the density of equilibrium slant $g(\cdot)$. Because firms strike the most consumer-friendly arbitrators, the mean award is roughly 4.5pp lower than the equilibrium density of slant $g(\cdot)$.

The striking advantage of firms therefore accounts for approximately 60% of firms' informational advantage.

In response to incentives provided by selection, arbitrators compete to be selected by choosing a pro-industry slant a that is biased relative to their beliefs b, which generates the competition effect. Intuitively, we compare how individual arbitrators are ruling to how they would rule in absence of incentives provided by selection, i.e., their belief of a fair ruling. Formally, the magnitude of the competition effect, roughly 3.5pp, is illustrated by comparing the distribution of slant $g(a)$ to the distribution of beliefs. Intuitively, the average arbitrator gives out an award that is 3.5pp lower than what she believes is fair because doing so increases her probability of being selected for arbitration. Forty percent of firms' total informational advantage, therefore, comes from changes in the arbitrator pool as a whole. Recall that we cannot measure this aspect using the reduced-form fixed effects approach in the first part of the paper.

Informed Individual Consumers and Information Externalities: Another interpretation of the competition effect is that it is the advantage that the industry holds even over *an individual consumer who is as informed as the industry*. Formally, consider a situation in which only a measure zero of consumers are informed—for example, because they purchase expertise by hiring PIABA attorneys (see Appendix C for a formal treatment). This consumer would be as good as firms in striking arbitrators from a given pool. Nevertheless, she would be at a disadvantage. The whole pool would still have a pro-industry tilt because the ex ante chances of arbitrators being struck by an informed consumer are essentially zero. On the other hand, if all consumers are informed, the competition effect is eliminated. This result implies that the consumer benefits from being informed as a group are larger than the sum of informed individuals: being informed generates positive externalities for other consumers. This externality opens the door for potential regulation. One example of such regulation is the prohibition of arbitration clauses that rule out class-action claims, such as the one proposed by the CFPB and later overturned by Congress ("New protections against mandatory arbitration," 2017).

VI.C Changing the Arbitrator Selection System

We use our model to quantitatively investigate different arbitrator selection schemes on dimensions that were considered by FINRA. Therefore, rather than considering a complete redesign of the system, we examine changes to the features of the existing system of choosing and compensating arbitrators. These policy changes were proposed with the idea that the arbitration process might lead to "fairer" outcomes for the consumer. We show that instead of achieving the intended objective, the outcomes are, by and large, more industry-friendly once one considers the informational advantage that firms hold in the arbitration process. To estimate the counterfactuals, we numerically solve for the updated slant strategies given the change in the arbitrator selection scheme and underlying arbitrator beliefs (see Appendix D). For computational convenience, we use the parametrically estimated belief distribution displayed in Figure 3 rather than our nonparametric estimates.

***Changing the Number of Strikes and Arbitration List Size*:** In 2017, FINRA proposed increasing the number of arbitrators on the list from 10 to 15 and simultaneously increasing the number of strikes from 4 to 6. We separately study how each of these dimensions changes the distribution of awards, and then examine the plan as a whole. We first present the changes in awards as the number of strikes increases from one to seven in Figure 6a. As the number of strikes increases, the awards distribution becomes more favorable to the industry. Consider the concrete example of FINRA's proposed changes of increasing strikes from four to six. The average award we observe in the data when both parties are allowed four strikes, $k = 4$, is 53%. As the number of strikes increases to six, $k = 6$, our estimates suggest that the average award will decline by 6pp. This change partially occurs because firms are able to select more favorable arbitrators from the list, and also because arbitrators are incentivized to act more industry-friendly. This counterfactual illustrates that increasing the control that the parties have over the process increases the slant in arbitration outcomes when consumers are uninformed. This result stands in stark contrast to consequences of this policy if consumers were informed. In such situations, increasing the number of strikes would indeed shrink the distribution of awards toward the "fairer" median outcome.

We next show in Figure 6b that increasing the size of the list from which arbitrators are chosen would benefit consumers. With the increased list size, arbitrators are less likely to be selected in general. Critically, all else equal, a given pro-consumer arbitrator is less likely to be one of the k most consumer-friendly arbitrators on the list, and thus is less likely to be eliminated. Figure 6b indicates that arbitrators would also be slightly less biased relative to their beliefs if they were chosen from a larger list. Holding the number of strikes fixed, increasing the number of arbitrators from 10 to 15 increases the average award by 3.7pp.

Overall, FINRA's proposal therefore introduced two changes that work in opposite directions. Increasing the number of strikes increases the pro-industry slant, but increasing the list size decreases it. Figure 6c illustrates that the proposed policy change would further increase the pro-industry slant, but the effects are modest. The average award decreases by 0.7pp.

Changing Arbitrator Compensation: Another policy proposal that has been frequently considered is to increase the fees paid to arbitrators (FINRA Notice 14-49, 2014). The idea is that higher fees will provide arbitrators with higher-powered incentives to set aside their biases, and instead work toward reaching a fair outcome; i.e., that awards will be closer to the median. This intuition is correct if consumers are as informed as firms. If consumers are uninformed, the policy change favors firms instead. Doubling the fee paid to the arbitrator would decrease the average award by 6pp (Figure 6d). The intuition is simple: increasing the fee paid makes the arbitrator more responsive to the informed party's preferences. With higher-powered incentives, arbitrators are more willing to be industry-friendly in order to increase the probability of being selected. Conversely, if FINRA reduced arbitrator compensation by 25%, the average awards would increase by roughly 1pp (Figure 6e). This counterfactual again illustrates that policies that would potentially improve arbitration outcomes if consumers were informed, worsen the pro-industry slant in arbitration outcomes when consumers are uninformed. These results also suggest that lower-powered incentives

for arbitrators, coupled with a flat wage, could decrease the pro-industry slant in arbitration.

Eliminating Career Concerns of Arbitrators: We also consider the counterfactual where arbitrators no longer have career concerns or other pecuniary motives. In this counterfactual, arbitrators no longer slant their decisions in favor of the industry; instead, they choose the award they consider fair. This differs from the mechanism under which arbitrators are allocated randomly in Section VI.B.1, because it still gives the informed party an advantage when choosing arbitrators from a given pool (the striking effect) but eliminates the competition effect. In the baseline arbitrator selection mechanism, the competition effect is responsible for lowering awards by 3pp. Figure 6f indicates that arbitration awards would be roughly 1pp higher (more consumer-friendly) if arbitrators did not have future career concerns.

VI.D Model Extensions

In the Appendix, we consider several model extensions. First, we explore how differences in the arbitration selection mechanism impact consumers' incentives to bring a case to arbitration. On one hand, because awards to consumers cannot be negative (awarded to the firm), the consumer may be willing to pay the legal costs of filing a meritless ("frivolous") case in hope of winning a large award, such as by drawing a consumer-friendly arbitrator. Making the arbitration system more consumer-friendly could increase the probability a consumer files a frivolous case. On the other hand, the arbitration selection mechanism also impacts a consumer's willingness to file cases with merit. For example, if arbitration becomes more biased in favor of the industry, consumers may be less willing to bring small cases to arbitration, even though the case potentially has merit. In the Appendix, we explore how the arbitration selection mechanism impacts a consumer's incentive to file frivolous as well as high-merit cases and the associated costs. We find that the incentives to file frivolous cases under the current system are limited, but there is a significant chilling effect on consumers' propensity to file cases.

Second, we consider the case where consumers and/or firms are risk-averse. As discussed in the Appendix, incorporating consumer/firm risk aversion does not impact the equilibrium distribution of awards or arbitrator slant. It can, however, alter preferences of the parties *across* different arbitration mechanisms. For example, risk-averse consumers may prefer a system that is more pro-industry on average, but has a lower variance of awards. We quantify these effects in the Appendix.

VII Conclusion

We argue that firms have an informational advantage over consumers in selecting arbitrators in consumer arbitration and document how the selection process impacts arbitration outcomes. We use securities disputes as a laboratory for our study. Securities disputes present a good laboratory for arbitration: arbitration is mandatory for all disputes, so it eliminates selection concerns, the parties choose arbitrators from a randomly generated list, and this selection mechanism is similar to those in other major arbitration forums.

Here, we want to highlight some more speculative implications of our findings. The estimates from our model suggest a substantial pro-industry tilt in the arbitration pool that, because of arbi-

trator competition, accounts for 40% of the informational advantage. Individual consumers, even if they are fully informed, cannot avoid this equilibrium effect. Being informed generates positive externalities for other consumers because the presence of informed consumers incentivizes arbitrators to develop a reputation for being consumer-friendly. Because individual consumers do not internalize the benefits of every consumer being informed, this externality opens the door for potential regulation. While analyzing such regulations is beyond the scope of the current paper, one example of such regulation is the prohibition of arbitration clauses that rule out class-action claims. For example, the CFPB proposed a rule preventing companies from using mandatory arbitration clauses, which was overturned by Congress in 2017.

Our counterfactuals suggest that redesigning incentive compensation and arbitrator selection design can ameliorate the pro-industry tilt, but only if the design accounts for uninformed consumers. We show that certain policies, such as increasing arbitrator compensation or giving parties more choice, can benefit consumers if they are informed, but hurt them if they are uninformed. One avenue for future research is to examine the extent to which this result is generic. More broadly, our findings suggest that limiting the firm's and consumer's inputs over the arbitrator selection process could significantly improve outcomes for consumers. We hope that future work can extend our workhorse quantitative model to evaluate such alternative arbitration selection mechanisms.

Data Availability Statement

The code and publicly available data underlying this research are available on Zenodo: `https://doi.org/10.5281/zenodo.10869588`

References

Acquilina, M., E. Budish, and P. O'Neill. "Quantifying the High-Frequency Trading 'Arms Race': A Simple New Methodology and Estimates." Working Paper.

Agarwal, S., S. Chomsisengphet, N. Mahoney, and J. Stroebel. 2015 "Regulating Consumer Financial Products: Evidence from Credit Cards." *The Quarterly Journal of Economics*, 130 (1): 111–164.

Agarwal, S., J. Grigsby, A. Hortacsu, G. Matvos, A. Seru, and V. Yao. 2020 "Searching for Approval." Working Paper.

Allen, J., R. Clark, B. Hickman, and E. Richert. 2020. "Resolving Failed Banks: Uncertainty, Multiple Bidding & Auction Design." Working Paper.

Allen, J., R. Clark and J.F. Houde. 2019. "Search Frictions and Market Power in Negotiated-Price Markets. *Journal of Political Economy*, 127(4): 1550-1598.

Altonji, J. G., T.E. Elder, and C.R. Taber. 2005. "Selection on observed and unobserved variables: Assessing the effectiveness of Catholic schools." *Journal of Political Economy*, 113(1): 151-184.

Andersen, S., C. Badarinza, L. Liu, J. Marx, and T. Ramadorai. 2020a. "Reference Dependence in the Housing Market." Working Paper.

Andersen, S., J.Y. Campbell, K.M. Nielsen, and T. Ramadorai. 2020b. "Sources of Inaction in Household Finance: Evidence from the Danish Mortgage Market." *American Economic Review*, Forthcoming.

Argyle, B., C. Palmer, and T. Nadauld. 2019." Monthly Payment Targeting and the Demand for Maturity. *Review of Financal Studies*, Forthcoming.

Ashenfelter, O. 1987. "Arbitrator behavior." *The American Economic Review P&P*, 77(2): 342-346.

Ashenfelter, O, and D. E. Bloom. 1984. "Models of Arbitrator Behavior: Theory and Evidence." *The American Economic Review*, 74(1): 111-124.

Ashenfelter, Orley, Janet Currie, Henry S. Farber, and Matthew Spiegel. 1992. "An experimental comparison of dispute rates in alternative arbitration systems." *Econometrica*, 60(6): 1407-1433.

Barr, Michael S. 2014 "Mandatory arbitration in consumer finance and investor contracts." New York University Journal of Law and Business, 11(4): 793-817.

Barber, Brad M., and Terrance Odean. 2000."Trading is hazardous to your wealth: The common stock investment performance of individual investors." *Journal of Finance,* 55(2): 773-806.

Benetton, M. 2018. "Leverage Regulation and Market Structure: An Empirical Model of the UK Mortgage Market." *Journal of Finance,* Forthcoming.

B., Vivek, G. Illanes, and M. Padi. 2019. "Fiduciary duty and the market for financial advice." NBER Working Paper No. w25861.

Black, B. and J.I. Gross. 2002. "Making it Up as They Go Along: The Role of Law in Securities Arbitration." Cardozo Law Review, 23(3): 991-1047.

Bloom, D.E. 1986. "Empirical Models of Arbitrator Behavior under Conventional Arbitration." *Review of Economics and Statistics,* 68(4): 578-585.

Bloom, D.E., and C.L. Cavanagh. 1986a. "An Analysis of the Selection of Arbitrators." *American Economic Review* 76 (3): 408-22.

Block, R.N., and J.Stieber. 1987. "The impact of attorneys and arbitrators on arbitration awards." ILR Review 40, no. 4: 543-555.

Bodoh-Creed, A., J. Boehnke. and B. Hickman. 2018. "Using Machine Learning to predict price dispersion." Working Paper.

Bordalo, P., N. Gennaioli, and A. Shleifer. 2015. "Salience Theory of Judicial Decisions."

Journal of Legal Studies, 44(S1): S7-S33

Brams, S. and S. Merill. 1983. “Equilibrium Strategies of Final-Offer Arbitration: There is No Median Convergence,” *Management Science*, 29(8): 927-941.

Brams, S. and S. Merill. 1986. "Binding versus Final-Offer Arbitration: A Combination Is Best." *Management Science,* 32(10): 1346-55.

Budish, E., P. Crampton, and J. Shim. 2015. “The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response.” *The Quarterly Journal of Economics*, 130(4): 1547–1621.

Budish, E., R. Lee, and J. Shim. 2019. “A Theory of Stock Exchange Competition and Innovation: Will the Market Fix the Market?” National Bureau of Economic Research. NBER Working Paper No. 25855

Campbell, J.Y. (2006). “Household Finance.” *Journal of Finance,* 61(4): 1553–1604.

Casella, G. 1992. “Illustrating empirical Bayes methods.” *Chemometrics and Intelligent Laboratory Systems* 16(2): 107-125.

Charoenwong, B., A. Kwan, and T. Umar. 2019. “Does Regulatory Jurisdiction Affect the Quality of Investment-Adviser Regulation?” *American Economic Review,* 109(10): 3681-3712.

Chetty, R., J. Friedman, and J. Rockhoff. 2014. “Measuring the Impacts of Teachers I: Evaluating Bias in Teacher Value-Added Estimates.” *American Economic Review* 104(9): 2593-2632

Choi, S.J., J.E. Fisch, and A.C. Pritchard. 2014. "The influence of arbitrator background and representation on arbitration outcomes." *Virginia Law & Business Review,* 9: 43.

Colvin, A.J.S. 2018. “The growing use of mandatory abitration.” Economic Policy Institute Report.

Crawford, G.S., N. Pavanini,and F. Schivardi. 2018. “Asymmetric Information and Imperfect Competition in Lending Markets. *American Economic Review*, 108(7): 1659-1701.

Crawford, V. P. 1979. “On compulsory-arbitration schemes.” *Journal of Political Economy*, 87(1): 131-159.

Crawford, V. P. 1982. “"Compulsory Arbitration, Arbitral Risk and Negotiated Settlements: A Case Study in Bargaining Under Imperfect Information.” *Reveiw of Economic Studies*, 69: 69-82.

De Clippel, G., K. Eliaz, and B. Knight. 2014. “On the selection of arbitrators.” *The American Economic Review*, 104(11): 3434-3458.

Deck, C., and A. Farmer. 2007. "Bargaining over an Uncertain Value: Arbitration Mechanisms Compared." *Journal of Law, Economics, and Organization*, 23(3): 547-79.

Dimmock, S.G., W.C. Gerken and N.P. Grahm. 2018. “Is Fraud Contagious? Career Networks and Fraud by Financial advisers.” *Journal of Finance,* 73(3): 1417-1450.

Egan, M.. 2019. “Brokers vs. Retail Investors: Conflicting Interests and Dominated Products.” *Journal of Finance,* 74(3): 1217-1260.

Egan, M., A. MacKay, and H. Yang. 2020. “Recovering Investor Expectations from Demand for Index Funds." NBER Working Paper No 26608.

Egan, M., G. Matvos, and A. Seru. 2017. “When Harry Fired Sally: The Double Standard in Punishing Misconduct.” NBER Working Paper No. w23242.

Egan, M., G. Matvos, and A. Seru. 2019. “The Market for Financial Adviser Misconduct.” *Journal of Political Economy,* 127(1): 233-295.

Farber, H. S. 1980b. "An Analysis of Final-Offer Arbitration." *Journal of Conflict Resolution,* 24(4): 683-705.

Farber, H. S., and Harry C Katz. 1979. "Interest Arbitration, Outcomes, and the Incentive to Bargain." *Industrial and Labor Relations Review,* 33(1): 55-63.

Farber, H.S., and M.H. Bazerman. 1986, "The General Basis of Arbitrator Behavior: An Em-

pirical Analysis of Conventional and Final-Offer Arbitration." *Econometrica,* 54(6): 1503-1528.

Farmer, A., and P. Pecorino. 1998. "Bargaining with Informative Offers: An Analysis of FinalOffer Arbitration." *Journal of Legal Studies*, 27(2): 415-32.

Farmer, A., and P. Pecorino. 2003. "Bargaining with Voluntary Transmission of Private Information: Does the Use of Final Offer Arbitration Impede Settlement?" *Journal of Law, Economics, and Organization,* 19(1): 64-82.

Gennaioli, N. 2013. "Optimal contracts with enforcement risk." *Journal of the European Economic Association,* 11(1): 59-82.

Gennaioli, N., and S. Rossi. 2010. "Judicial Discretion in Corporate Bankruptcy." *The Review of Financial Studies*, 23(11): 4078-4114.

Gennaioli, N., and A. Shleifer. 2008. "Judicial Fact Discretion." *Journal of Legal Studies* 37(1): 1-35.

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

Gentzkow, M., B. Kelly, B. and M. Taddy. 2019. "Text as Data." *Journal of Economic Literature*, 57(3): 535-74.

Gentzkow, M., and J.M. Shapiro. 2006. "Media Bias and Reputation." *Journal of Political Economy,* 114(2): 280–316.

Gibbons, Robert. 1988. "Learning in Equilibrium Models of Arbitration." *American Economic Review,* 78(5): 896-912.

Giglio, S., Maggiori, M., Stroebel, J. and Utkus, S.,. 2019. "Five Facts about Beliefs and Portfolios." NBER Working Paper (No. w25744). National Bureau of Economic Research.

Goldsmith-Pinkham, P. and K. Shue. 2020. "The Gender Gap in Housing Returns" *Journal of Finance*, Forthcoming.

Golbukh, S. and P. Goldsmith-Pinkham. 2020. "Heterogeneous Real Estate Agents and the Housing Cycle." Working Paper.

Guerre, E., I. Perrigne, and Q. Vuong. 2000. "Optimal Nonparametric Estimation of First-Price Auctions." *Econometrica*, 68(3): 525-574.

Gomes, F., M. Haliassos, and T. Ramadorai. "Household Finance." *Journal of Economic Literature*, Forthcoming.

Guiso, L., P. Sapienza, and L. Zingales. (2008). "Trusting the stock market." *Journal of Finance*, 63(6): 2557-2600.

Gurun, U., G. Matvos, and A. Seru. (2016). "Advertising Expensive Mortgages." *Journal of Finance*, 71(5): 2371-2416.

Honigsberg, C. and M. Jacob. 2020. "Deleting Misconduct: The Expungement of BrokerCheck Records. *Journal of Financial Economics,* Forthcoming.

Koijen, R.S.J., and M. Yogo. 2016. "Shadow insurance." *Econometrica*, 84(3): 1265-1287.

Kondo, J. 2006. "Self-Regulation and Enforcement in Financial Markets: Evidence from Investor-Broker Disputes at the NASD." Working Paper.

Nichols, C. 1999. "Arbitrator Selection at the NASD: Investor Perception of a Pro-Securities Industry Bias." Ohio State Journal on Dispute Resolution. 15(1): 63-134.

Olszewski, W. 2011. "A welfare analysis of arbitration." *American Economic Journal: Microeconomics*, 3(1), pp.174-213.

Oster, E. 2019. "Unobservable Selection and Coefficient Stability." *Journal of Business Economics and Statistics,* 37(2): 187-204.

Qureshi, H. and J. Sokobin. 2015. "Do Investors Have Valuable Information About Brokers?" FINRA Office of the Chief Economist Working Paper.

Rosenthal, R.W. 1978. "Arbitration of two-party disputes under uncertainty." *The Review of*

Economic Studies, 45(3): 595-604.

Samuelson, William F. 1991. "Final-Offer Arbitration under Incomplete Information." Management Science, 37(10): 1234-47.

Silver-Greenberg, Je. and R. Gebeloff. 2015. "Arbitration Everywhere, Stacking the Deck of Justice." *NYT.*

Stevens, C. 1966. "Is Cumpulsory Arbitration Compatible with Bargaining?" *Industrial Relations*, 5: 38-52

Stipanowich, T., C. von Kannn and D. Rothman. 2010. *Protocols for Expeditious, Cost-Effective Commercial Arbitration: Key Action Steps for Business Users, Counsel, Arbitrators & Arbitration Provider Institutions.* The College of Commercial Arbitrators.

Figure 1: Award Distribution

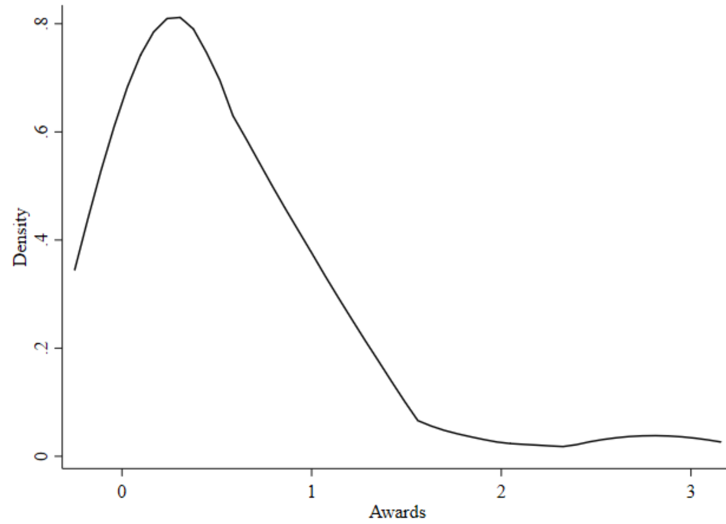

Note: Figure 1 displays the distribution of arbitration Awards. We calculate Awards as the percentage of awards granted through arbitration divided by the awards initially requested by the consumer. The distribution of Awards is winsorized at the 1% level.

Figure 2: Arbitrator Fixed Effects

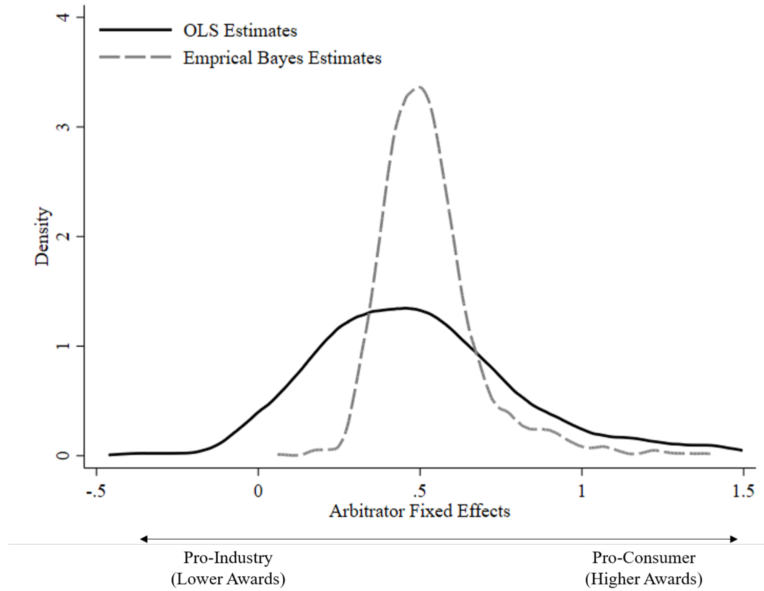

Note: Figure 2 displays the estimated distribution of arbitrator fixed effects corresponding to eq. (1). The black empirical density reflects the distribution of OLS estimated arbitrator fixed effects estimated and the gray empirical density reflects the distribution of empirical Bayes estimated arbitrator fixed effects. We normalize the mean of the distribution of fixed effects to 53% which is the average Award in our sample.

Figure 3: Estimated Distribution of Arbitrator Beliefs, Slant, and Awards

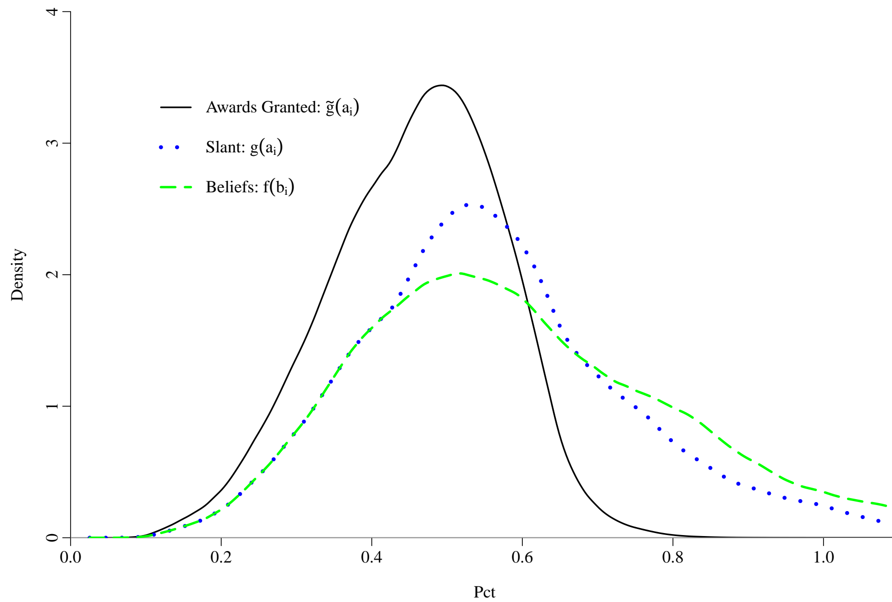

Note: Figure 3 displays the estimated density of awards among the conditional distribution of selected arbitrators $\tilde{g}(a)$, the density of slant among the unconditional (entire) population of arbitrators $g(a)$, and the distribution of true beliefs among the unconditional (entire) population of arbitrators $f(b)$. The distributions of awards, slant, and beliefs correspond to our parametric estimates as described in Section VI.

Figure 4: Estimated Distribution of Arbitrator Beliefs, Slant, and Awards

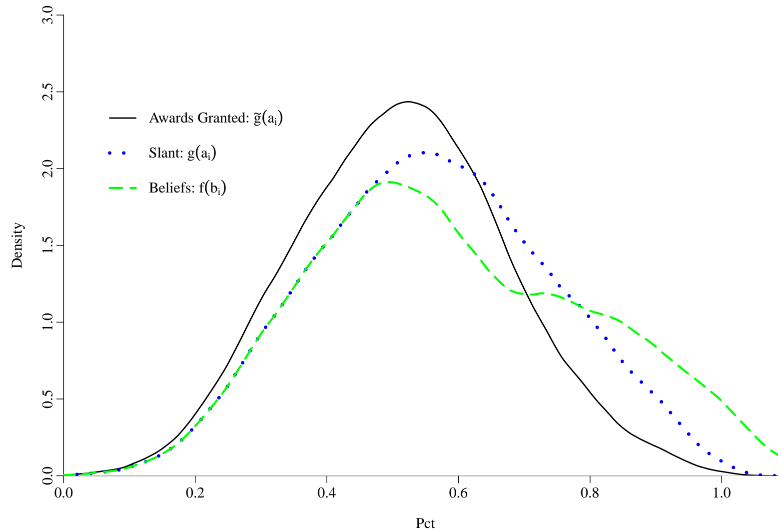

Note: Figure 4 displays the estimated density of awards among the conditional distribution of selected arbitrators $\tilde{g}(a)$, the density of slant among the unconditional (entire) population of arbitrators $g(a)$, and the distribution of true beliefs among the unconditional (entire) population of arbitrators $f(b)$. The distributions of awards, slant, and beliefs correspond to our nonparametric estimates.

Figure 5: Distribution of Arbitrator Beliefs, Slant, and Awards—Informed Consumers

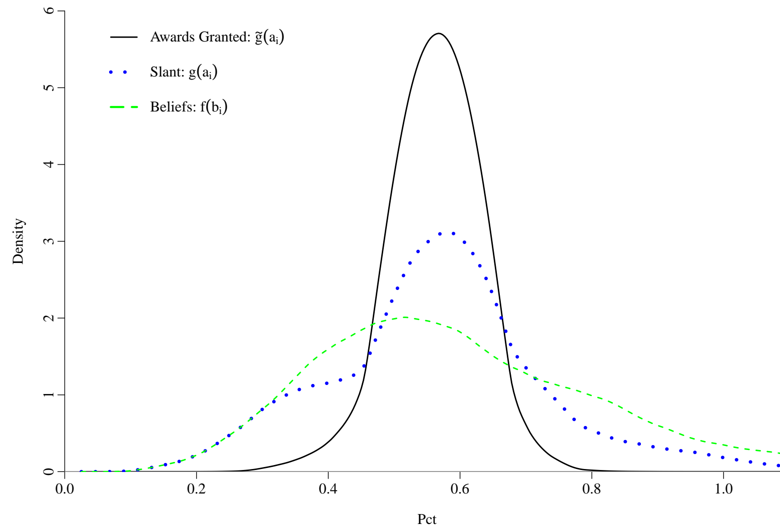

Note: Figure 5 displays the model implied density of awards among the conditional distribution of selected arbitrators $\tilde{g}(a)$, the density of slant among the unconditional (entire) population of arbitrators $g(a)$, and the distribution of true beliefs among the unconditional (entire) population of arbitrators $f(b)$ if both parties are informed.

Figure 6: Arbitration Awards Under Alternative Selection Mechanisms

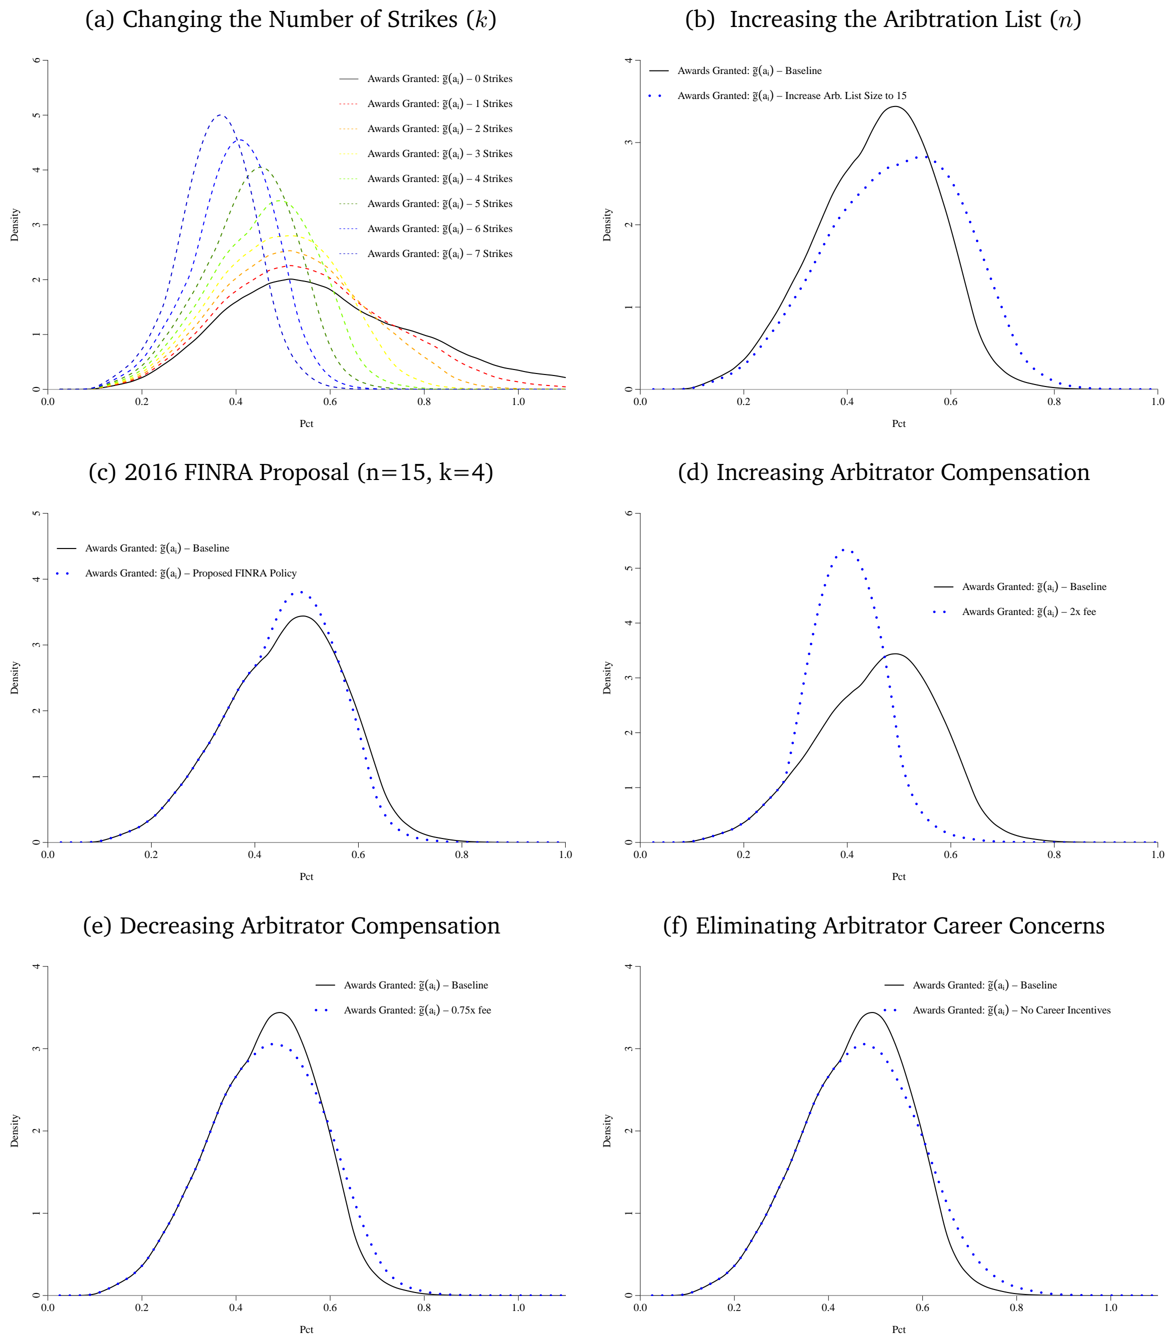

Note: Panels (a)-(f) of Figure 6 display the distribution of arbitration awards if regulators were to (a) change the number of strikes, (b) increase the number of arbitrators on the list from ten to fifteen, (c) increase the number of arbitrators on the list to fifteen and increase the number of strikes to six (FINRA's recent proposal), (d) double the fee paid to arbitrators, (e) decrease the fee paid to arbitrators by 25%, and (f) eliminate career concerns of arbitrators.

Table 1: Arbitration Summary Statistics

Variable	Obs	Mean	Std. Dev.	Median
Requested Awards	11,756	758,648	1,826,864	240,000
Percent of Requested Awards Granted	11,756	53%	60%	35%
Arbitrator Characteristics				
Total Number of Arbitration Cases	11,756	11.29	13.02	7.00
Total Number of Consumer Dispute Cases	11,756	3.64	3.40	3.00
Former/Current Financial Adviser	11,756	41%		

Note: Table 1 corresponds to the consumer dispute arbitration case characteristics. Observations are at the consumer arbitration case-by-arbitrator. Additional summary statistics are presented in Table A1.

Table 2: Arbitrator Selection From the List

	(1)	(2)	(3)	(4)	(5)	(6)
Consumer Friendliness	-0.17*	-0.32***	-0.30*	-0.42**	-1.24***	-1.26***
	(0.097)	(0.12)	(0.17)	(0.20)	(0.38)	(0.41)
Observations	7,059	4,281	2,534	1,920	788	719
Arbitrator Controls	X	X	X	X	X	X
Sample:						
Chairperson Sample		X		X		X
Greater than 5 or More Obs.			X	X		
Greater than 7 or More Obs.					X	X

Note: Table 2 displays the results corresponding to a multinomial logit model (eq. 2). Observations are at the case by eligible arbitrator level. Consumer Friendliness is measured using the empirical Bayes estimated arbitrator fixed effects as described in Section IV.A, which are standardized to ease in interpretation and winsorized at the 1% level to account for outliers. More consumer friendly (a higher arbitrator fixed effect) indicates that, all else equal, the arbitrator gives out higher awards and is therefore more consumer friendly. Arbitrator controls include whether the arbitrator is a public arbitrator and eligible to serve as a chairperson. Robust standard errors are in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.10$.

Table 3: Are Consumer Friendly Arbitrators Selected Less Often?

	(1)	(2)	(3)	(4)	(5)	(6)
Consumer Friendliness	-0.26***	-0.20***	-0.085***	-0.071***	-0.088***	-0.072***
	(0.077)	(0.074)	(0.026)	(0.025)	(0.027)	(0.025)
Observations	2,714	2,714	2,714	2,714	2,714	2,714
R-squared	0.002	0.139				
Arbitrator Controls		X		X		X
Geographic × Arbitrator Type F.E.		X		X		X
Model						
OLS	X	X				
Poisson			X	X		
Negative Binomial					X	X

Note: Table 3 displays the results corresponding to a regression model (eq. 3). Observations are at the arbitrator level. The dependent variable is the total number of consumer arbitration cases an arbitrator has overseen over the period 1998-2019. Consumer Friendliness is measured using the empirical Bayes estimated arbitrator fixed effects as described in Section IV.A, which are standardized to ease in interpretation and winsorized at the 1% level to account for outliers. A higher Arbitrator Fixed Effect indicates that, all else equal, the arbitrator gives out higher awards and is therefore more consumer friendly. Arbitrator controls include the number of years the arbitrator has been active as an arbitrator. We control for geographic region fixed effects and geographic region fixed effects interacted with the type of arbitrator (public, non-public, chairperson eligible). Robust standard errors are in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.10$.

Table 4: Industry Friendly Arbitrator Selection and Consumer Sophisitication

	(1)	(2)	(3)
Attorney	-1.51	0.23	-0.12
	(2.57)	(2.74)	(2.94)
PIABA Attorney	8.05***	8.14***	7.94***
	(1.60)	(1.60)	(1.75)
Trust	4.60**	5.83***	6.21***
	(2.00)	(2.03)	(2.29)
Firm Experience (Above Median)	-1.77	-0.28	-0.43
	(1.09)	(1.23)	(1.37)
Observations	9,427	9,427	9,294
R-squared	0.010	0.030	0.112
Other Controls		X	X
Year F.E.			X
County F.E.			X

Note: Table 4 displays the regression results for a linear regression model (eq. 4). Observations are at the arbitrator-by-case level over the period 1998-2019 and come from our consumer dispute arbitration data set. The dependent variable is the selected Arbitrator's Fixed Effect as calculated in column (3) of Table A2. A higher Arbitrator Fixed Effect indicates that, all else equal, the arbitrator gives out higher awards and is therefore more consumer friendly. Attorney is a dummy variable indicating whether the consumer used an Attorney. PIABA Attorney indicates whether the consumer used a attorney who specializes in arbitration and is a member of the of the Public Investors Arbitration Bar Association. Trust indicates whether the consumer claimant is a trust. Firm Experience is a dummy variable indicating whether the firm has above median experience in terms of the number of arbitration cases it has been involved in. The omitted category is consumers who are self-represented, not part of a trust, and are facing firms with below average experience. Coefficients are in percentage points. Other Controls include case size, the arbitration panel size, the case length in terms of the number of words, and other adviser characteristics. Other controls also include the corresponding adviser's qualifications: Series 6, Series 7, Series 24, Series 63, Series 65/66, and number of other qualifications. Standard errors are clustered at the case level. *** p<0.01, ** p<0.05, * p<0.10.