Fiduciary Duty and the Market for Financial Advice

VIVEK BHATTACHARYA[1], GASTÓN ILLANES[1], and MANISHA PADI[2]

[1]Department of Economics, Northwestern University and NBER
[2]University of California, Berkeley, School of Law

May 19, 2020

ABSTRACT. Fiduciary duty aims to solve principal-agent problems, and the United States is in the middle of a protracted debate surrounding the merits of extending it to all financial advisers. Leveraging a transaction-level dataset of deferred annuities and state-level variation in common law fiduciary duty, we find that it raises risk-adjusted returns by 25 bp. Through the lens of a model of entry and advice provision, we argue that this effect can be due to both an increase in compliance costs (a fixed cost channel) and a direct constraint on low-quality advice (an advice channel), and we show how to disentangle these two effects. Model estimates indicate that the advice channel is the dominant force in explaining the observed results, and counterfactual simulations suggest that further increases in the stringency of fiduciary duty, such as a federal fiduciary standard, will continue to improve advice.

KEYWORDS. fiduciary duty, financial regulation, financial advice, retirement markets, annuities.

JEL CODES. G23, G28, K15, L51, L84.

CONTACT. vivek.bhattacharya@northwestern.edu, gaston.illanes@northwestern.edu, mpadi@berkeley.edu. We are grateful to the Wharton Consumer Analytics Initiative for organizing the collaboration with the data sponsor, as well as to many people at the data sponsor for assistance in understanding the dataset and the industry. We would like to thank Kenneth Ayotte, Iván Canay, Ignacio Cuesta, George Deltas, Mark Egan, Daniel Garrett, Dan Green, Igal Hendel, Ali Hortaçsu, Robert Jackson, S.P. Kothari, Neale Mahoney, Gregor Matvos, Dina Mayzlin, Frank Partnoy, Rob Porter, Jim Poterba, Mar Reguant, Jimmy Roberts, Nancy Rose, Brad Shapiro, and Paulo Somaini for helpful comments, as well as participants at the 4th Rome Junior Conference on Applied Microeconomics, the American Law and Economics Association Meetings, Berkeley Law School, Duke Law, the International Industrial Organization Conference, Michigan, Michigan State, the Summer Institute on Competitive Strategy at UC Berkeley, the SEC Conference on Financial Regulation, the SITE session on Financial Regulation, University of Illinois College of Law, and University of Toronto Faculty of Law. Matias Escudero and David Stillerman provided excellent research assistance. Bhattacharya and Illanes acknowledge financial support from the National Science Foundation under Grant No. SES 1824463, the Center for the Study of Industrial Organization at Northwestern, and the Northwestern Institute on Complex Systems. Padi acknowledges support from NSF Social, Behavioral and Economic Sciences Fellowship, Award Number 1715010, and the Center for Equitable Growth. All errors are our own.

1. Introduction

Informed agents working on behalf of uninformed principals are subject to fundamental conflicts of interest. The primary legal mechanism for bridging this principal-agent problem has historically been *fiduciary duty*. Agents subject to fiduciary duty must act in the best interest of their principals, including a duty of care that requires agents to exert effort on behalf of them, and a duty of loyalty that requires agents to put aside any opportunities for private benefit. If agents fail to satisfy their fiduciary duty, they can be liable for any losses the principals incur.

This paper sheds light on the effect and mechanisms of fiduciary duty in a setting currently undergoing significant policy upheaval in the United States: the regulation of financial advisers. Americans save almost $30 trillion for retirement, much of which is in complex financial products sold through advisers. A patchwork of state and federal law has resulted in many advisers not being classified as fiduciaries, and the past decade has seen various regulators—including the Department of Labor, the Securities and Exchange Commission, and many state legislatures—propose to bridge this gap. Consumer and industry groups have spent millions lobbying on this issue, with the former alleging serious financial losses in vulnerable older populations and the latter arguing that fiduciary duties simply place undue burden on advisers without affecting outcomes.

Supporters of expanding fiduciary duty to all advisers argue that it directly alleviates conflicts of interest and thus makes it more costly to offer low-quality advice. We call this potential mechanism the *advice channel*. Opponents argue for a second mechanism: fiduciary duty does not have an impact on product choice directly—perhaps because investors already know which product to buy or because conflicts of interest are minimal—and instead simply raises the cost of doing business regardless of the quality of advice. This may lead to fewer advisers in the market and perhaps to even worse advice in equilibrium. We term this potential mechanism the *fixed cost channel*, where fiduciary duty increases fixed costs and shifts the equilibrium set of entrants.

We develop a model of entry and advice provision that captures these two forces: fiduciary duty may increase both the cost of providing low-quality advice and fixed costs. Each mechanism will change observed advice, directly in the case of the advice channel and indirectly through entry incentives for both. The model allows us to unpack the relative contributions of both channels and simulate the impact of alternative regulatory regimes while taking into account entry and exit responses to regulation. The potential for the entry margin to undo thedirect effect of a regulation is a concern in any intervention that affects the profitabilityof advice quality. A main contribution of this paper is totake into account both changes in advice and changes in entry decisions when evaluating policy interventions, which do by quantifying the effects of regulation both on fixed costs and on the costs of providing low quality advice.

To estimate these effects, we leverage a new dataset of transaction-level data for deferred annuity sales from an anonymous financial services provider ("FSP"). FSP is among the top-five companies by market share of annuities and representative of other large companies in this industry. This dataset contains information about every contract sold by FSP from 2013–2015, detailed data about the product and adviser, and some limited data on the client. For each transaction we observe the fiduciary status of the adviser and granular geographic information about the transacting parties. We supplement this data with hand-coded information about contract characteristics from SEC filings and open records requests as well as data from Morningstar and CRSP about investment options within annuities. We develop a dynamic model of the execution of these contracts to aggregate these multidimensional characteristics into a single valuation for each annuity.

The key variation we exploit is differences in fiduciary duty across types of advisers and across state borders. Advisers licensed as *registered investment advisers* (RIAs) have a fiduciary duty towards their clients at the national level, while those licensed as *broker-dealers* (BDs) do not. BDs are excluded from fiduciary duty because they historically have been considered order takers without a significant advisory function. Today, however, they do similar work with respect to retail investors (SEC, 2011, 2013a,b) and largely carry the same annuities at the same "prices" (fees, contract characteristics, etc.). Crucially, however, state courts in several states have ruled that BDs are fiduciaries within their borders, setting up common law variation in fiduciary standards. We compare behavior of BDs in states in which they have fiduciary duty to states in which they do not, using the difference in behavior of RIAs as a control. To control for differences across states, we restrict to counties along state borders at which there is a change in common law fiduciary standards.

We leverage this variation and the aforementioned product valuations to document in the reduced form that fiduciary duty improves the quality of transacted products in equilibrium. In particular, BDs facing fiduciary duty sell products with risk-adjusted returns that are 25 basis points higher. The increase in returns arises from a change in the set of transacted products. We find a shift towards fixed indexed annuities and away from variable annuities. Within variable annuities, sales shift towards those with more investment options, a larger variety of highly-rated investment options, and options with higher historical returns.

These results are a novel quantification of the causal effect of fiduciary duty on treated markets. However, they are not informative of the effect of fiduciary duty outside the markets under study, or of the effects a federal fiduciary standard—which is likely to be more stringent than a common law standard—would have, as they do not directly inform the mechanisms through which the regulation operates. To quantify these mechanisms, we develop and estimate a model of entry into the provision of financial advice with heterogenous adviser qualities (or types) and differentially regulated firms. To capture the advice channel, the model is flexible regarding the extent to which different adviser types vary their advice when facing fiduciary standards. To capture the fixed cost channel, the model does not restrict the relationship between profitability and firm type, so that changes in fixed costs can drive high or low quality firms out of the market.

Using the model, we show that one can identify the presence of an advice channel by examining how the distribution of advice, rather than simply the mean, changes with the imposition of fiduciary duty. If fiduciary duty were to only increase fixed costs, then the set of advisers in the market as well as the set of observed advice would contract. On the other hand, if the advice channel were substantial, then we might observe the emergence of new, especially high-quality, advice—both because existing advisers adjust their advice and because entry and exit would skew the composition of advisers towards those who do not find it costly to offer higher-quality advice. This implication of the model leads to a nonparametric reduced-form test that we take to the data, proxying advice with risk-adjusted returns, and we find evidence for the presence of an advice channel.

This observation also feeds into the quantification of the two channels through the lens of the structural model. We allow for unobserved types across firms which dictate both their latent propensity to offer high-quality advice and their profitability with and without fiduciary duty. We estimate the model using a two-step procedure. In the first step we recover beliefs about entry probabilities from observed entry decisions in the data à la Sweeting (2009). In the second step, we impose that observed entry must be profitable given these beliefs and recover the remaining parameters. As in the reduced form, we flexibly control for differences across borders and use comparisons across borders to inform the structural parameters. To do so in a tractable manner, we develop a computational Bayes approach.

Model estimates show that fiduciary duty operates both by increasing the cost of offering distorted advice and by increasing fixed costs. Moreover, the increase in fixed costs induced by fiduciary duty drives out high-quality advisers from the market, reducing average returns. On net, however, the advice channel significantly outweighs the fixed cost channel; almost all the observed effect on advice is due to the advice channel. We use the model to simulate increasing the stringency of fiduciary duty while allowing for the composition of advisers to endogenously change. We find that tripling the stringency of fiduciary standards relative to common law does induce exit of broker-dealers, and these broker-dealers tend to offer higher-quality advice. However, this exit is small relative to the full market for financial advice, and it is insufficient to counteract the direct effect of the advice channel substantially. Taking into account both the direct effect through the advice channel and the indirect effect through endogeneous exit, average returns provided by brokers increase by 20 bp relative to common-law fiduciary standards. Taken together, these results suggest that increasing stringency of fiduciary standards may continue to benefit retirees.

Related Literature. This paper contributes to a growing literature on the industrial organization of financial markets. Like this paper, this literature uses structural econometric methods to study market structure and consumer behavior in settings such as car loans (Einav et al., 2012; Grunewald et al., 2019), credit cards (Nelson, 2020; Gavazza and Galenianos, 2020), insurance (Koijen and

Yogo, 2015, 2016, 2018), mortgages (Allen et al., 2014, 2019; Benetton, 2019; Robles-Garcia, 2020; Grigsby et al., 2019), municipal bonds (Brancaccio et al., 2020), pensions (Hastings et al., 2017; Illanes, 2017; Illanes and Padi, 2019), personal loans (Cuesta and Sepúlveda, 2019; Liberman et al., 2019; Xin, 2020), small business lending (Bachas and Liu, 2019), and student loans (Bachas, 2019).

More narrowly, this paper relates to the literature on expert advice in financial decision-making. Theoretical work on financial advice has a long tradition (Inderst and Ottaviani, 2012a,b), and there is a growing body of recent empirical work on this issue. A number of papers have documented advisers responding to commissions and other incentives rather than offering clients appropriate advice (Anagol et al., 2017; Bergstresser et al., 2009; Christoffersen et al., 2013; del Guercio and Reuter, 2014; Guiso et al., 2018; Mullainathan et al., 2012; Robles-Garcia, 2020; Garrett, 2019). Focusing specifically on financial advisers, Egan et al. (2019) study the prevalance and geographic concentration of misconduct in this industry, and Charoenwong et al. (2019) show that the agency in charge of enforcement affects quality, as proxied by complaints. Our contribution to this literature is to study how fiduciary duty, the main policy lever to constrain poor advice, affects adviser behavior.

Despite the policy importance of fiduciary duty, there has been limited empirical analysis of it. Finke and Langdon (2012) identify cross-state common law variation and show that advisers do report that fiduciary duty constrains their advice. Kozora (2013) considers a temporary change in the fiduciary standard for the municipal bond market and finds that stricter standards led to more sales of investment-grade bonds. Finally, Egan (2019) documents a high likelihood of purchase of dominated products in the reverse convertible bond market. Through the lens of a search model, he estimates that extending fiduciary duty would increase risk-adjusted returns by 5–21 bp. We contribute to this literature both by identifying the effect of fiduciary duty in the reduced form and by taking into account the entry margin when considering the counterfactual effect of extending fiduciary duty to all financial advisers. This allows us to simulate the impact of different levels of stringency on returns without assuming that the set of advisers is held fixed. Given that federal fiduciary standards have not been formulated, and several approaches have been proposed, it is critical to build predictions that consider alternative stringency levels in order to inform policy.

2. Institutional Details

In this section, we introduce the institutional setting. Section 2.1 discusses financial advisers in the US and how fiduciary standards have evolved. Section 2.2 discusses details of variable and fixed indexed annuities, the products we study in this paper.

2.1. Financial Advisers and Fiduciary Duty

The United States has two types of financial advisers, which evolved separately for historical reasons but now largely serve similar functions. The first are registered investment advisers (RIAs), who are regulated at the federal level by the SEC under the Investment Advisers Act of 1940. The second are broker-dealers (BDs), who are subject to the Securities Exchange Act of 1934 and regulated by state law and by FINRA, a private industry regulator. BDs are not regulated under the Investment Advisers Act as they were initially conceived as mere brokers. Since then, however, they have grown into the role of providing financial advice as well. RIAs must be affiliated with a brokerage firm to sell certain products, including annuities, and thus many such advisers are *dually registered* as broker-dealers and investment advisers. They are subject to fiduciary duty at the federal level on their advisory accounts. In our sample, all transacting advisers are either broker-dealers or dual registrants—as they are selling annuities—but we refer to them as BDs and RIAs nevertheless.

All financial advisers perform many of the same functions when working with individuals. Their primary role is to recommend and facilitate the purchase of investment vehicles, which are issued by upstream financial services providers. Broker-dealers are typically paid by commission, receiving a payment from the upstream supplier from every sale while charging nothing directly to clients. Compensation schemes for RIAs tend to be a combination of commissions and a percentage of assets under management. Advisers who are compensated, even in part, on the basis of commissions have a conflict of interest: they have an incentive to recommend high-commission products over ones that may be cheaper for their clients. Moreover, informed advisers with uninformed clients may have no incentive to exert effort to maximize their client's value if clients cannot verify the quality of advice ex post.

The patchwork of federal, state, and private regulation overseeing adviser behavior attempts to combat this conflict of interest by imposing legal duties on advisers. All BDs nationwide have a federal duty to deal fairly with their client and must recommend products that are "suitable." This requirement does not specify that BDs must prioritize the client's best interest over their own, as long as the product they recommend satisfies FINRA's suitability rules.[1] BDs are also required to provide clients with each product's prospectus, which includes all technical details about the investment vehicle but is not easily understood by a layperson. Any dispute that arises over a BD's regulatory compliance is arbitrated through FINRA's private dispute resolution process. Other claims may be brought under state or federal law. Nationwide regulation of RIAs is more stringent. RIAs have fiduciary duty imposed on them by the SEC, which requires that they entirely disregard their own interest and work in the best interest of their client. RIAs may still take commissions, but must disclose the resulting conflict of interest to their client.[2] If a client has a dispute with an RIA, he may sue in state or federal court, or enter into FINRA arbitration or external private arbitration.[3]

[1] See http://www.finra.org/industry/suitability.

[2] RIAs that recommend higher commission products must justify that recommendation by using SEC-approved software that validates recommendations and by drafting disclosures to clients, among other costly compliance measures.

[3] Arbitrability varies across claims and states, although, to our knowledge, not across adviser types. Some states will allow tort claims to be brought that are very similar in nature to arbitrable claims even when there are mandatory

Consumer groups and the SEC have long been troubled by the difference in regulatory standards across BDs and RIAs. Studies by the SEC (SEC, 2011, 2013a,b) have suggested that consumers often do not realize that BDs have an incentive to sell high commission products. They also are unable to tell whether their financial adviser is technically classified as a BD or a RIA, and many assume that all advisers are fiduciaries. Motivated by these concerns, the SEC recommended that standards be harmonized, requiring all advisers dealing with retail investors to offer the best possible contract in the investor's interest. The DOL promulgated a rule in 2016 largely following the SEC recommendation.[4] The rule would place a fiduciary duty on BDs that handle retirement savings for retail investors and require all advisers to sell clients the best available contract for them. In addition, the DOL rule requires contracts between advisers and consumers that specify the fiduciary duty and allows consumers to bring class action lawsuits to enforce it. The financial adviser industry pushed back on this rule, claiming it would significantly increase compliance costs for BDs and raise the spectre of expensive class action litigation, potentially putting some BDs out of business (Kelly, 2017). Litigation ultimately caused the DOL rule to be delayed indefinitely.[5] In June 2019, the SEC passed a final rule clarifying the duties placed on both RIAs and BDs, called "Regulation Best Interest". This rule harmonizes the standards to which BDs and RIAs are held, and requires all advisers to act in the best interest of their consumers.[6] Debates continue regarding the effect of this rule, relative to a more traditional fiduciary duty approach (Bernard, 2019; Marsh, 2019).

This project estimates the impact of imposing fiduciary duties on BDs by leveraging cross-state variation in state common law. In some states, court rulings have imposed a common law duty of care that rises to the level of a fiduciary duty—a higher standard than required of BDs at the federal level. Finke and Langdon (2012) classify states into ones with no common law fiduciary duty on advisers and ones with some level of fiduciary duty; Figure 1 plots this classification.[7] These duties allow clients to sue their financial advisers for low quality advice.[8] Since all RIAs already comply

arbitration clauses in the contract between client and adviser.

[4] See https://www.dol.gov/agencies/ebsa/laws-and-regulations/rules-and-regulations/completed-rulemaking/1210-AB32-2.

[5] The Fifth Circuit Court of Appeals vacated the DOL Rule in March 2018, stating the DOL had overstepped its authority, and it currently seems unlikely the DOL Rule will be resurrected. States have responded by imposing fiduciary duty through legislation, rather than common law.

[6] Clarifying guidance includes disclosure requirements and other documentation intended to ensure that consumers receive high quality advice. See https://www.sec.gov/rules/final/2019/34-86031.pdf.

[7] Finke and Langdon (2012) develop this classification based on legal research involving careful readings of case law. In Appendix G, we outline the procedure we use to validate their legal research and arrive at the same classification. We also discuss two alternate decisions pertaining to treatment of federal cases and case law for insurance providers that yield a modified classification. We show the main results of this paper are stronger under this alternate classification.

[8] Advisers who lie to their clients in a way that causes them material loss can always be sued for fraud or misrepresentation, under standard principles of tort law. Additional duties of care, including fiduciary duty, allow clients to recover losses sustained even when advisers have told clients the truth. This can occur when advisers suggest risky investments, "churn" across assets to increase their commissions, and otherwise do not tailor their advice to the needs of their client. For further discussion, see the Joint SEC/NASD Report (https://www.sec.gov/news/studies/secnasdvip.htm).

Figure 1: Common law fiduciary duty on broker-dealers by state

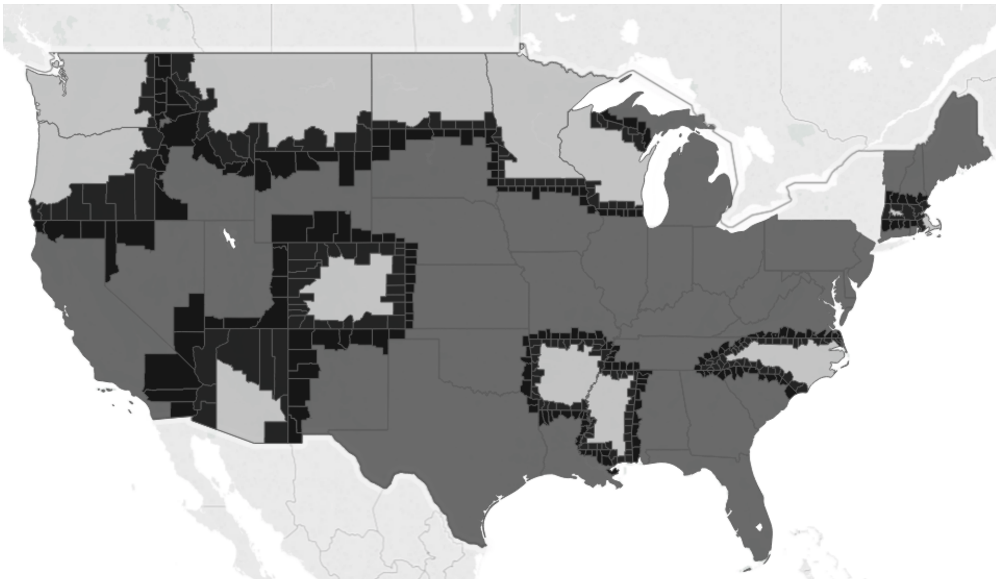

States with some degree of fiduciary duty (dark grey) and none (light grey), per Finke and Langdon (2012). Counties in black are ones at borders between states with different fiduciary standards and consitute our main sample. New York, which does not impose common law fiduciary duty on its broker-dealers, and its surrounding counties are omitted, as New York has different suites of products.

with uniform federal fiduciary duty standards, they provide a control against which to compare treated BDs (facing a fiduciary duty) relative to control BDs (facing only FINRA suitability rules). If fiduciary duty is effective, BDs will modify their behavior and their compliance programs, resulting in changes to their recommendations and to the investments made by their clients. Additionally, competitor behavior and market structure may be affected. Of course, states may not always be able to enforce these duties and common law may be less salient than legislation, suggesting that any estimate obtained by comparing state law regimes will likely be an underestimate of the impact of a federal rule.[9] A benefit of our approach, which combines a reduced form estimate of the effect of fiduciary duty and a structural entry and advice model, is that the model allows us to address this issue by predicting the effect of this legislation under counterfactual stringency levels.

2.2. Fixed and Variable Annuities

We restrict attention to annuities, one of the most common retirement vehicles, with over $3 trillion in reserves. In addition to the size and importance of the annuity market, the DOL directly mentioned concerns about annuities as the impetus for their 2016 rule.[10] Most annuity contracts sold in the US

[9]Most state law fiduciary duty claims are brought by individual litigants, while statutory fiduciary duty claims could allow for more state enforcement actions and class actions.

[10]The DOL stated that "[m]any other products, including various annuity products, among others, involve similar or larger adviser conflicts, and these conflicts are often equally or more opaque." It went on that the "greater

are *deferred annuities*.[11] These products involve an accumulation phase, during which money is contributed to an account and invested, and a payout phase, during which payments are made from the account to the annuitant. Fixed indexed (FIA) and variable annuities (VA) are the most popular deferred annuity products. They share the structure of an accumulation and a payout phase, but differ in how the account grows during accumulation, in the ways money can be withdrawn during both phases, in fee structure, and in the *riders*, or options, that can be added to the contract.

Investors in FIAs distribute their funds during the accumulation phase between a series of *crediting strategies*. Crediting strategies include fixed rates of return and the performance of the S&P 500, with a cap and a floor. All crediting strategies fully protect the investor from downside risk. In most cases, fees are not directly charged, so the client does not need to understand any further features of the product.[12] The main exception to this statement are *surrender charges*, which tax withdrawals taken in the first years of the accumulation period if they exceed a free withdrawal amount (typically 10% of contract value). Fixed indexed annuities can be converted into a fixed annuity once investors are sufficiently old, transitioning the contract into the payout phase; alternatively, they can be withdrawn. In the case of death during the accumulation period, beneficiaries receive the contract amount.

Variable annuities replace the small set of crediting strategies in FIAs with a pool of investment funds, with a wide range of asset allocations, risk profiles, and fees. The most basic VA contract resembles an FIA, with contract values accruing interest according to the performance of the set of funds chosen, and investors receiving the option of an annuity upon entering the payout phase. For this contract, investors pay an annual percentage fee, the expense ratios of the funds they invest in, and potentially surrender charges. Often, VAs are sold with *living benefit riders*, which establish a separate account called an *income base*, which for a fixed period of time grows by the maximum of the realizations of the fund return and a fixed rate. During the payout phase, clients choose between drawing down the account value, annuitizing it, or receiving a percentage of the income base in perpetuity. These riders essentially convert the VA into an option (Koijen and Yogo, 2018). This structure incentivizes risk-taking in fund selection. To mitigate this incentive, companies impose restrictions on an annuitant's investment portfolio. Optimal execution of VAs requires choosing appropriately from the pool of investment options, and if the contract is coupled with a living benefits rider, it further requires making correct decisions about when to take withdrawals. As a result, these contracts are more complex and difficult to value than a fixed indexed annuity. They

degrees of complexity, magnif[ies] both investors' need for good advice and their vulnerability to biased advice." See https://www.federalregister.gov/documents/2016/04/08/2016-07924/definition-of-the-term-fiduciary-conflict-of-interest-rule-retirement-investment-advice.

[11] Fixed immediate annuities, in which investors turn over a lump sum in exchange for fixed periodic payments until death, are a very small fraction of the US annuity market.

[12] The margin comes from the the realized return of the index less the amount accrued.

also expose the annuitant to relatively more risk than FIAs do.

For annuities sold by FSP, there is no difference between BDs and RIAs in terms of the characteristics of the products they can choose to recommend. This implies both types of advisers can offer the same product with the same investment options and fees. A client choosing a particular product would have the same payout stream regardless of the adviser. What differs is how advisers are compensated by FSP.

3. Data

In Section 3.1, we describe the data provided to us by FSP about its transactions and the advisers that sell its products. Section 3.2 discusses data for the individual products in the dataset. Section 3.3 presents our calculations for returns.

3.1. Transactions, Advisers, and Clients

We have transaction-level data from a major financial services provider, FSP, which sells a mix of annuities and insurance products in all fifty states, has household name recognition, and is publicly traded. Our main dataset consists of information about all transactions associated with financial products offered by FSP in the United States between 2008 and 2015. For each transaction, we observe the specific FSP product transacted, the date, the adviser selling the product, and the dollar amount. If a contract involves multiple transactions, such as recurring payments, then they can be grouped together, and we report the sum of the transaction amounts. The only client-level information we have is the client's zipcode and age. Although clients can also be linked across contracts, clients purchasing multiple contracts is rare.

Additionally, FSP has provided us data from Discovery Data for all advisers who could potentially sell annuities or life insurance in 2015, regardless of whether they transact with the company. This dataset allows us to observe basic demographics of the adviser as well as regulatory information such as licensing and whether the adviser is registered as a BD, an RIA, or both. While advisers cannot be matched externally, we are able to match them to FSP transactions. Discovery also includes information about the firms, including the firm footprint (e.g., local or national). A drawback of Discovery is that since we only have a snapshot in 2015, we have to restrict our analysis to a window of time around this period to ensure the accuracy of each adviser's licensing information; we restrict to 2013–2015. Additional sample selection decisions are reported in Appendix F.2.

Table 1 provides summary statistics for FSP contracts sold in the border counties highlighted in Figure 1 and for the advisers associated with them. About 21% of advisers are BDs. BDs and RIAs each sell about 5.7 FSP contracts on average over the sample, with some advisers selling significantly more. Conditional on selling an FSP annuity, BDs and RIAs sell VAs 79% and 90% of

Table 1: Summary statistics for border counties

				Percentiles				
	N	Mean	Std.	10%	25%	50%	75%	90%
Adviser-Level Quantities								
Is Broker-Dealer								
FSP Advisers	3,936	0.207						
Contracts per Adviser								
BD	814	5.7	9.2	1	1	2	6	14
RIA	3,122	5.7	9	1	1	3	6	14
Contract-Level Quantities								
Is Variable Annuity								
BD	4,678	0.793						
RIA	17,794	0.900						
Contract Amounts ($K, 2015)								
BD	4,678	119.4	139.8	24.2	42.6	79.9	148.6	251.5
RIA	17,794	153.0	179.7	34.3	54.4	100.9	188.2	304.1
Client Age								
BD	4,678	61.3	10.3	49	55	62	68	74
RIA	17,794	64.5	9.5	54	59	65	71	77

the time, respectively. Contract amounts are about $34,000 larger for RIAs. Finally, the average client is around retirement age, with a difference of about 3 years between BD and RIA clients. Summary statistics for the entire nation are broadly similar; see Appendix B.1.

3.2. Product Characteristics

We match the transaction dataset to external data sources containing information about the products. Beacon Research has provided historical data about the all fees and investment options available to annuitants; this data is sourced from quarterly prospectuses that VAs are required to file with the SEC. We also hand collected information about restrictions on investments and rider rules from prospectuses stored in EDGAR, the SEC's online database. We match investment options to the Morningstar Investment Research Center to collect information about fund ratings and investment styles, and we match them to the CRSP US Mutual Fund database for historical returns.

Contract characteristics for transacted annuities are summarized in Table 2, separated by whether the adviser is a BD or an RIA. Panel (A) shows historical undiscounted returns (net of expense ratios) of the underlying investment options, assuming either the return-maximizing allocation (subject to investment restrictions) or an equal allocation across funds (Benartzi and Thaler, 2001). Panel (B) shows the minimum and average expense ratio of all potential investments. Panel (C) shows the mortality and expense fee, an annual percentage fee that must be paid on all products, along with the average surrender charge over the surrender schedule—which must be paid only if

Table 2: Summary statistics for annuities sold by BDs and RIAs, border counties

	BD		RIA	
Characteristic	Mean	Std.	Mean	Std.
(A) Fund Return (%)				
Return-Maximizing	0.152	0.088	0.160	0.087
Equal	0.011	0.011	0.012	0.010
(B) Fund Expense Ratios (%)				
Minimum	0.503	0.022	0.500	0.020
Average	1.270	0.213	1.261	0.198
(C) Fees				
M&E Fee (%)	1.189	0.215	1.064	0.305
Surrender Charge (%)	3.737	1.197	2.963	1.436
(D) # Funds				
All	97.52	37.56	96.65	33.49
High Quality	27.39	12.63	33.12	14.09
Low Quality	34.74	17.24	30.57	19.06
(E) # Equity Styles				
Some High Quality	6.85	2.05	7.30	1.94
Only Low Quality	1.03	1.75	0.83	1.62
(F) # FI Styles				
Some High Quality	4.05	1.05	4.49	1.57
Only Low Quality	3.05	0.30	3.02	0.25
(G) Contract Return (all products)				
Risk-adjusted	0.031	0.012	0.027	0.010
Unadjusted	0.064	0.021	0.064	0.023

Panels (A)–(F) summarize characteristics of transacted VAs. Panel (G) summarizes returns of all transacted annuities. In Panels (E) and (F), "Some High Quality" refers to styles covered at least by one high quality fund, and "Only Low Quality" refers to styles convered only by low quality funds.

money is withdrawn early.[13]

Panels (D)–(F) measure the potential for diversification together with Morningstar's quality metrics for the underlying funds. Morningstar rates each fund on a scale of 1–5 stars based on its historical risk-adjusted return (net of expenses) relative to a peer group of funds. A fund is labeled *high-quality* if it receives at least 4 stars and *low-quality* if it receives 2 or fewer. Second, Morningstar categorizes the *style* of both the equity and fixed-income investment of each fund into nine potential styles. Panel (D) counts the number of distinct investment options available per product, unconditionally and across quality levels. Panels (E) and (F) report the number of equity and fixed-income styles that are covered by at least one high-quality fund, as well as the number only covered by low-quality funds.

Table 2 shows that the variation across BDs and RIAs is small relative to the variation within

[13]The surrender charge varies by year since the purchase of the contract, and it declines to zero within ten years. We average the surrender charges over this period (averaging in zeros if needed).

adviser category. Given this heterogeneity, there is scope for advice to materially affect client outcomes and thus for regulation that shifts advice to have an impact. These characteristics may affect the return of the annuity, which we report in Panel (G). We discuss the procedure to calculate this return in Section 3.3.

While FIAs do not have to file product characteristics with the SEC, we collected archived rate sheets through a series of open records requests to state insurance agencies. Beacon Research provides further information about them. Unfortunately, rates depend on the crediting strategies available in an FIA, so we do not have simple summary characteristics for FIAs like we do for VAs. However, we fold these rates into the return calculations.

3.3. Calculating Net Returns

We aggregate contract characteristics into returns using two methods. Our preferred metric computes risk-adjusted returns, using a stochastic discount factor corresponding to a three-factor model (Cochrane, 2009). We also compute unadjusted returns, as they may align more closely with the information given to retail investors; del Guercio and Reuter (2014) shows that unsophisticated investors are sensitive to unadjusted returns of mutual funds.

We compute returns of each annuity in an environment where the annual risk-free rate is 3%, for an individual who values money left to heirs equally as her own consumption. Computing the expected net present value of these products requires (i) information about the fees of the basic contract and all riders, (ii) expectations over the distribution of returns for all underlying funds in which the annuitant can invest, (iii) a stance on the discount rates, and (iv) an understanding of portfolio allocations (for a VA) or crediting strategies (for an FIA) and how the annuitant chooses whether and when to take the rider. This information, together with age and contract amount, generates an NPV for each transaction. For interpretation, we present values as the annualized returns necessary in a fixed account to achieve the same NPV by the terminal age of the contract.[14]

As discussed above, we have fees and rate sheets, which directly deals with (i). We proxy (ii) using a Fama-French three-factor model for the underlying mutual funds, estimated using the historical distribution of returns from CRSP. We deal with (iii) discounting in two ways: for adjusted returns, we compute the stochastic discount factor that prices the factors and use this quantity to discount various states of the world. Alternatively, we compute returns for an individual who discounts all states of the future at 3%. Finally, given that a limitation of our dataset is that we do not see portfolio allocations of clients or execution of the riders, we tackle (iv) by formulating

[14] If A is age, $\beta = 3\%$ is the discount rate, and T is the contract's terminal age, we find the return R such that

$$(1+\beta)^{T-A} \cdot (\text{Net Present Value}) = (1+R)^{T-A} \cdot (\text{Transaction Amount})\,.$$

Figure 2: Distribution of returns, border counties

(a) Risk-adjusted

(b) Unadjusted

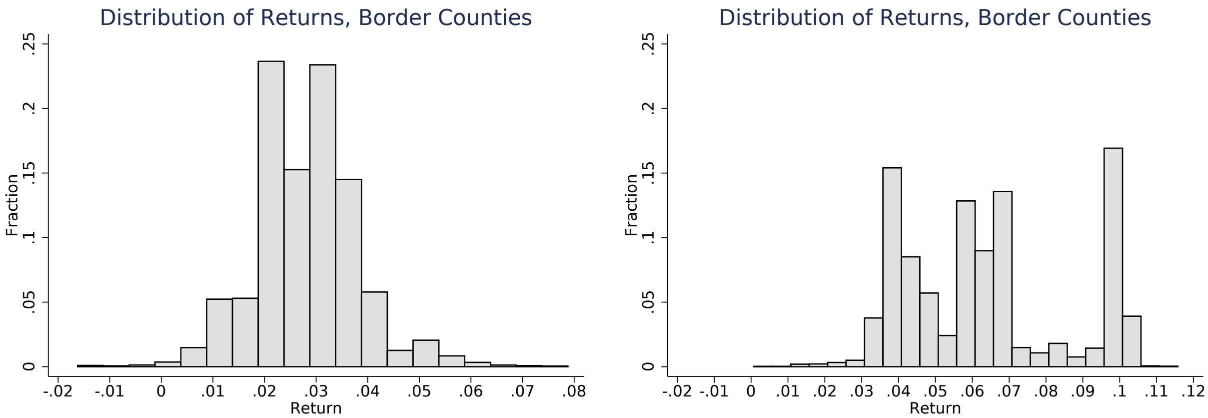

and solving the dynamic programming problem to find optimal execution of portfolio allocation or crediting strategy decisions, withdrawal decisions, and rider execution. Details of the factor model and discounting are in Appendix C, and an exposition of the dynamic program is in Appendix D.

Panel (G) of Table 2 shows that average returns of transacted products are slightly higher for BDs than RIAs. Figure 3 shows the full distribution of returns, which vary highly across products. Risk adjusted returns for VAs and RIAs range largely between 0 and 6%, with long tails in either direction. Products in the mean of the distribution have risk adjusted returns of about 2.5%, meaning that client returns could potentially double if they were advised to invest in a different product. Similar observations apply to the distribution of unadjusted returns.

4. Does Fiduciary Duty Affect Outcomes?

This section presents reduced-form estimates of the effect of fiduciary duty on advice and entry. These effects are the total impact of the advice and the fixed cost channels discussed in the introduction, and the model in Section 5 provides a roadmap for disentangling them.

4.1. Empirical Strategy

The simple comparison of product sales across legal regimes is tainted by the fact that fiduciary standards are not randomly assigned. For example, if preferences for financial instruments have influenced the adoption of fiduciary standards, then differences in product sales across states confounds the effect of fiduciary standards with differences in preferences. Instead, we think of fiduciary duty as an endogenous object that is the result of each state's judicial process. We address this issue in two steps. First, we restrict the analysis to counties on either sides of a border between states that differ in fiduciary status, since we expect that—and subsequently provide corroborating evidence for the fact that—border counties are similar to each other. Second, we compare the difference across the border for BDs to that for RIAs, leading to a difference-in-differences strategy. In particular, for a variety of outcomes Y_{ist}, we run the regression

$$\begin{aligned} Y_{ist} = \alpha_0 + \alpha_1 \cdot \mathbb{1}[\text{State has FD for BDs}]_s \cdot \mathbb{1}[\text{Adviser is BD}]_i \\ + \alpha_2 \cdot \mathbb{1}[\text{State has FD for BDs}]_s \cdot \mathbb{1}[\text{Adviser is RIA}]_i \\ + \alpha_3 \cdot \mathbb{1}[\text{Adviser is BD}]_i + \text{Border FE} + \text{Month FE} + \text{Age FE} + \epsilon_{ist}, \quad (1) \end{aligned}$$

where i represents an adviser, s a state, and t a transaction. We include border fixed effects to use only within-border variation, month-of-contract fixed effects to address any changes in product offerings and rates over time, and client age fixed effects.

Within (1), there are three objects of interest. First is the straightforward difference-in-differences estimator, $\alpha_1 - \alpha_2$ in this formulation. Under the null hypothesis that fiduciary duty has no equilibrium impact on market outcomes, we should estimate $\alpha_1 - \alpha_2$ to be zero. One may worry that counties on either side of a state border differ from each other, either in the underlying demand for financial products or the supply of financial advice. However, the difference-in-differences estimator should alleviate this concern: as long as market differences across state borders are equal for BDs and RIAs, we would still expect $\alpha_1 - \alpha_2$ to be 0. In the results below, we will reject that $\alpha_1 - \alpha_2 = 0$ for most outcomes of interest, suggesting that fiduciary duty has an equilibrium impact. Under the assumption that there are no spillover effects onto RIAs one can interpret this difference-in-difference estimate as the causal effect of fiduciary duty on BDs.

We also interpret α_1 and α_2 separately. Under the assumption that market conditions do not change sharply across the state border, α_1 alone is the causal impact of fiduciary duty on BDs, and α_2 can be interpreted as the spillover effect of BDs fiduciary duty onto RIAs. That is, interpreting both α_1 and α_2 as separate causal effects requires no shift in underlying market characteristics at the border.

The results show an effect of fiduciary duty on BDs, with α_1 significantly different than zero for a variety of outcomes. However, we find no evidence of spillover effects on RIAs, with α_2 economically and statistically zero for most outcomes. Moreover, we find limited evidence throughout for within-firm changes in the behavior of RIAs and on RIA entry.

We provide four arguments in favor of the assumption that underlying market characteristics do not change sharply at the state border. First, demographic characteristics are balanced across the border (Appendix B.2). Second, even with covariate balance, one may be worried about differential selection of consumers to advisers as a function of the fiduciary status of the state. However, there is extensive survey evidence (SEC, 2011, 2013a,b; Hung et al., 2008) suggesting that consumers

Figure 3: Returns for border counties, by adviser type and fiduciary status

(a) Risk-adjusted returns

(b) Unadjusted returns

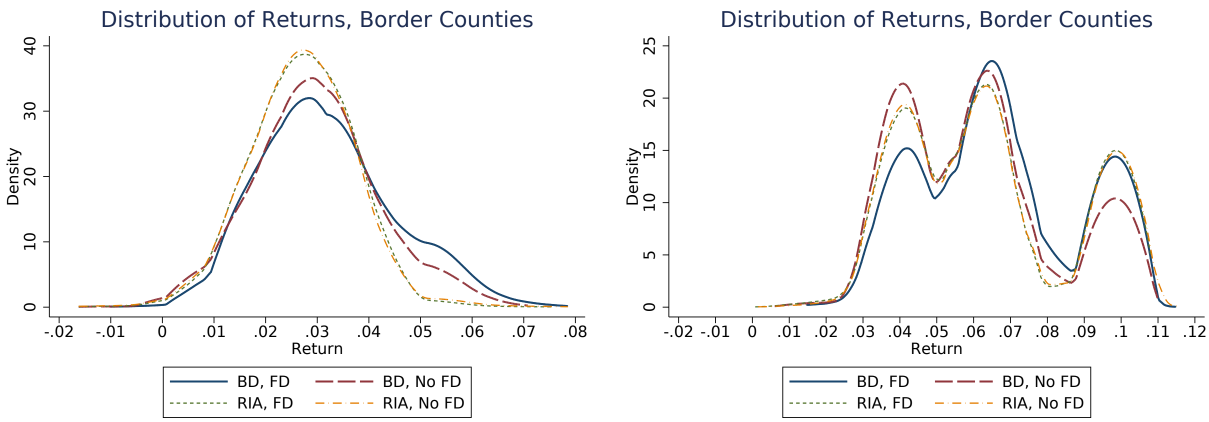

have very little information about which type of adviser they visit. Of course, there can still be selection on observables—certain consumers may choose to visit large companies, which are more likely to have RIAs—but the extent of this selection would have to vary significantly across borders for this to be a legitimate concern. Third, one can test for differential selection by using client and contract characteristics as outcomes in (1). Table B.4 in Appendix B.2 shows no significant effects on transaction amount, client age, or incidence of cross-state shopping (i.e., whether the adviser and client are from the same state), providing more suggestive evidence against differential selection. Finally, if there were significant differences across borders, we would have expected differences in RIA behavior as well.

To understand whether investors are better off from the imposition of fiduciary duty, we look at three sets of outcomes. First, in Section 4.2 we ask whether fiduciary duty increases investor returns. Second, in Section 4.3 we study how the characteristics of transacted products change with fiduciary status. Finally, in Section 4.4 we check whether improvements in returns are negated by a contraction in the size of the market.

4.2. Effects on Returns

Figure 3 shows the distribution of returns, both risk-adjusted and not, of products sold by advisers in border counties, conditional on adviser type and fiduciary status. The distribution of returns for BDs in states with fiduciary duty is shifted rightward relative to states without it, for both risk-adjusted and unadjusted returns. The distributions for RIAs are almost identical for states with and without fiduciary duty, lending credence to our strategy.

The behavior of BDs with fiduciary duty does not mimic that of RIAs. Indeed, we do not expect it to. Broker-dealers and RIAs may work at firms that negotiate different contracts with FSP, may attract different clienteles, or may have different business models. Our identification strategy allows

Table 3: Returns on variable annuity products

	(1) Risk Adjusted Returns	(2) Unadjusted Returns
DID	0.0025**	0.0047*
	(0.0011)	(0.0023)
FD on BD	0.0020**	0.0034
	(0.0009)	(0.0021)
FD on RIA	-0.0006	-0.0013*
	(0.0010)	(0.0007)
Mean of Dep. Var	0.028	0.064
N	22,472	22,472

Annualized returns for variable annuities sold. Contracts are restricted to borders, specifications include border fixed, contract month, and age fixed effects. Standard errors are clustered at the state. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

for this heterogeneity across types, as long as it is independent of the fiduciary status of the state.

Table 3 reports estimates of (1). Even controlling for compositional differences underlying Figure 3, we find a statistically and economically significant effect of fiduciary status on returns. Risk-adjusted returns increase by about 25 bp, which corresponds to approximately 9% of the base mean. This difference is due almost entirely to the effect on BDs, and—consistent with the figure—the effect on RIAs is negligible. Results are similar for unadjusted returns. The results are robust to heterogeneity in discounting across the population: in Appendix B.3, we let clients be a mix of those evaluating products in a risk-adjusted vs. an unadjusted manner. Over the space of all possible mixtures, we find that fiduciary duty improves returns by at least 18 bp.

4.3. Effects on Product Characteristics

What are the changes in the characteristics of the underlying products transacted that lead to this change in returns?[15] Answering this question not only helps unpack the return effect, but also yields evidence on the behavior of financial advisers under different regulatory regimes. After all, these characteristics are usually salient in prospectuses and brochures. Thus, they may well be the avenue through which steering towards higher-quality products happens: advisers may be more upfront about fees and expenses, or highlight that certain products have more restrictive investment options.

We estimate (1) with the raw properties of annuities mentioned in Section 3 on the left-hand side. The most salient characteristic is the type of annuity: variable or fixed indexed. Given that variable and fixed annuities serve similar purposes, the type of annuity is a salient characteristic of a product that an adviser can influence. Column (1) of Table 4 uses a dummy for whether the annuity is a variable annuity as the outcome variable, and we find a difference-in-differences estimate of a

[15] Recall that products characteristics, and thus payout streams, do not vary across states; what varies is the probability they are transacted.

drop in the probability that the annuity is a variable annuity of 11 pp, or 12.5% of the base mean. Once again, the RIA effect is small (2.1 pp) compared to the BD difference (-8.9 pp), consistent with the fact that RIAs face the same regulatory regime and with the assumption that there are no changes in market characteristics at the border.

An adviser with fiduciary duty may be drawn to fixed annuities for a variety of reasons. First, FIAs tend to have higher (risk-adjusted) returns according to our calculations, and advisers may be aware that such annuities tend to be "better deals" and thus less willing to push variable annuities if they have fiduciary duty. Second, FIAs are simpler to explain to clients, because they do not include income and contract bases, or the complex riders that come with variable annuities. A shift to simpler products may limit the likelihood of the adviser being brought to the courtroom or arbitration by a client who claims fees and terms had not been properly explained. It would also be consistent with advisers using complexity as a proxy for (worse) quality; there is evidence that such a correlation exists in other settings (Célérier and Vallée, 2017). Finally, given that FIAs cannot generate negative unadjusted returns while VAs can, the shift to FIAs would also be consistent with a shift towards products that limit complaints from downside realizations.[16] Column (2) provides evidence of a shift towards products with lower downside risk, using the 10^{th} percentile of the total growth of a product as a measure.[17] Broker-dealers with fiduciary duty sell products with higher 10^{th} percentile returns.

The remainder of Table 4 studies shifts within the VA market. A salient property of the investment menu is the expense ratio of the funds. Column (3) shows that the minimum expense ratio decreases by about 0.6 bp off the baseline of 50 bp, showing that clients have access to a (slightly) lower fee option. However, Column (4) shows that the average expense ratio increases by about 5.4 bp, which may be relevant if one is concerned about naive allocation methods. Column (5) documents a shift towards VAs that have funds with higher mean returns, net of expense ratio, assuming a return-maximizing allocation; the effect is substantial, amounting to about 13% of the base mean. Column (6) shows a similar result assuming a naive equal allocation rule, which allays concerns about the increase in the average expense ratio.

Columns (7) and (8) documents noisy effects on the two most salient fees associated with the product: the M&E fee and the surrender charge. We find a small and statistically insignificant decrease of 5.5 bp in the M&E fee and a noisy increase of about 21 bp in the surrender charge. We should highlight that unlike M&E ratios and expense ratios, the surrender charge is not necessarily

[16] Only the income base of a VA is guaranteed to not have a negative return. The actual account value is not. Since the income base cannot be withdrawn, only annuitized, and the NPV of this annuity is lower than the dollar value of the income base, this implies that individuals with sufficiently low returns will receive lower payments than the value of their investment amount.

[17] An outcome where at the terminal age of the product, the client can withdraw K times the initial principal of the contract will be recorded as K. See Appendix D for details.

Table 4: Characteristics of products transacted

			Expense Ratio		Fund Returns		Fees	
	$\mathbb{1}$[VA]	10^{th} Perc.	Minimum	Average	Optimal	Equal	M&E	Surr. Chg.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
DID	-0.110***	0.704**	-0.006*	0.053**	0.0197*	0.0023**	-0.055	0.213
	(0.039)	(0.342)	(0.003)	(0.023)	(0.0107)	(0.0010)	(0.038)	(0.153)
FD on BD	-0.089**	0.568	-0.007**	0.062***	0.0195*	0.0023**	-0.047	0.123
	(0.035)	(0.354)	(0.003)	(0.020)	(0.0098)	(0.0010)	(0.035)	(0.158)
FD on RIA	0.021	-0.135	-0.001	0.009	-0.0002	-0.0000	0.009	-0.089
	(0.027)	(0.186)	(0.002)	(0.010)	(0.0036)	(0.0006)	(0.020)	(0.078)
Base Mean	0.878	2.609	0.501	1.263	0.159	0.012	1.088	3.109
N	22,472	22,472	19,730	19,730	19,730	19,730	19,730	19,730

	# Funds			# Equity Styles		# FI Styles	
	All	$\geq$ 4 Stars	$\leq$ 2 Stars	High Q.	Only Low Q.	High Q.	Only Low Q.
	(9)	(10)	(11)	(12)	(13)	(14)	(15)
DID	8.44*	3.84**	1.88	0.748**	-0.501*	0.277	-0.079**
	(4.28)	(1.89)	(2.01)	(0.330)	(0.251)	(0.186)	(0.034)
FD on BD	10.87***	3.59**	3.51	0.764***	-0.564**	0.169	-0.092***
	(3.91)	(1.57)	(2.15)	(0.262)	(0.214)	(0.171)	(0.029)
FD on RIA	2.43	-0.25	1.63	0.016	-0.063	-0.108	-0.013
	(2.20)	(0.86)	(1.31)	(0.148)	(0.127)	(0.091)	(0.008)
Base Mean	96.81	32.04	31.35	7.215	0.865	4.408	3.028
N	19,730	19,730	19,730	19,730	19,730	19,730	19,730

Estimates of (1) for various product characteristics. Columns (1) and (2) use the set of all annuities transacted in the border, while the other columns restrict to variable annuities. Standard errors are clustered at the state level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

paid. Additionally, lower fee FSP products always come with higher surrender charges, so advisers who are unconcerned about their clients needing to withdraw early should steer them towards higher surrender charge products.

Another characteristic of interest is the number of funds available to investors. Column (9) estimates that fiduciary duty leads BDs to sell products with about 8.4 more funds, relative to the difference in RIA sales. Column (10) shows an increase of about 12% in the number of "high-quality" funds, as measured by Morningstar ratings of 4 or 5 stars. However, Column (11) reports a positive but less precisely estimated increase of about 6% in low-quality funds as well—as proxied by 2 or fewer stars. The increase in high-quality (or low-quality) funds is not a mechanical consequence of having a larger set of funds: the set of options offered is an active product design decision by FSP, and when it chooses to offer a product with more options it could only add low-quality funds.

A second relevant metric is the diversity of funds available. Using the categorization into equity and fixed income styles discussed in Section 3, Columns (12) and (13) document an economically and statistically significant increase in the number of equity styles covered by at least one high-quality fund and a decrease in the number of equity styles covered by only low-quality funds. Columns (14) and (15) repeat the analysis for fixed income styles, but the effects are noisier and of smaller magnitude. While many of these characteristics feed into the previously discussed returns, not all are directly tied to them. However, they are salient to clients and advisers, and responsiveness of such observable dimensions provides further credence that fiduciary duty is having an effect. Moreover, these characteristics are interesting since they are tied, at least heuristically, to higher quality. Historical returns of investment options are publicized in prospectuses and marketing brochures, and advisers with fiduciary duty may be hesitant to recommend products with low investment returns—even if risk-adjusted returns are aligned with the market. An adviser and a client who have a more-choice-is-better mindset may find products with a large number and variety of investment options more attractive. In the process of following these quality heuristics, advisers may well steer clients to products that indeed have higher returns on net.

Another reason these characteristics are interesting is that they may be related to recourse. Litigation about fiduciary duty in other settings, including ERISA, has cited higher numbers of investment options, higher quality funds, lower expense ratios, higher returns, and lower fees as supporting the conclusion that fiduciaries are performing their function. FINRA arbitration sometimes also cites similar characteristics as complaints against advisers. We are unable to say whether advisers are operating on heuristics they truly believe to be correlated with higher quality, or whether they are responding to other incentives such as a desire to avoid litigation; nevertheless, regardless of the underlying mechanism, we find evidence that characteristics of transacted products change when fiduciary duty is introduced.

Table 5: Market size and structure

	All Products	FSP Products		Entry		
	VA Sales (1)	# of Contracts (2)	Total Sales (3)	Total Firms (4)	BD Firms (5)	RIA Firms (6)
$\mathbb{1}$[FD]	0.001 (0.049)	-0.023 (0.064)	0.043 (0.046)	-0.092 (0.069)	-0.157** (0.076)	-0.037 (0.068)
Mean	$51.1 M	55.5	$8.1 M	10.99	3.23	7.75
N	411	411	411	411	411	411

Regression of various metrics for total sales and number of firms on the fiduciary status of the county, controlling for log population, log median household income, and median age. Column (1) shows total sales of variable annuities across all firms. Columns (2) and (3) restrict to FSP and show number of annuity contracts (both fixed and variable) and total dollar sales of these contracts. Columns (4)–(6) show regressions of the number of firms of each type. All specifications use the $\log(x+1)$ transformation of the left-hand side, although means are presented without taking logs. Specifications include border fixed effects and standard errors are clustered at the border level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

4.4. Effects on Market Size and Structure

While the previous sections document increases in returns, conditional on purchasing a financial product, critics of fiduciary standards often claim that the net impact of such standards may be to decrease the number of firms and advisers in the market, thus limiting access to financial products for clients. To analyze this concern in the reduced form, we study whether the market size and the number of firms in the market changes.

First, we regress measures of market size on a fiduciary dummy, county controls, and border fixed-effects. We use three measures of market size: (i) total dollar sales of VAs at the county, which FSP has provided us through its membership in a consortium of annuity providers; (ii) total number of FSP contracts sold; and (iii) total dollar sales of FSP annuities. Table 5 provides results of these regressions. We find no statistically significant effects on market size. We estimate a zero effect of fiduciary status on dollar sales of VAs (across all providers). The standard errors allow us to rule out shifts of 10% in either direction with 95% confidence. We do not have data on sales of FIAs outside FSP, so Columns (2) and (3) focus on total FSP sales. We estimate a negative impact of fiduciary status on the number of annuity contracts sold by FSP and positive impact on total dollar sales of FSP annuities, but these effects are statistically indistinguishable from zero.

Second, we regress the (log of one plus the) total number of firms in a county on fiduciary status, controlling for border fixed effects and county covariates. We say a firm has entered a county if it employs at least one adviser in that county who is marked as actively selling financial products in Discovery, regardless of whether it transacts with FSP. We find evidence of both a level and a compositional effect of fiduciary duty on market structure. Column (4) shows that imposing fiduciary duty reduces the total number of firms in the market by about 9%, although we cannot rule out a zero effect at the 10% level. Columns (5) and (6) suggest that this effect comes primarily from a drop in the number of BD firms, which are affected by the regulation. The number of such firms drops by 16% in counties with fiduciary duty, a number that is significant at the 5% level. We do not estimate a statistically (or economically) significant effect on the number of RIA firms.[18]

These results suggest that the concern the detractors have about fiduciary duty inducing exit of BDs has merit. However, there are some reasons to believe that this effect may have limited import: the effect on the total number of firms is potentially small, and there is limited evidence for a significant drop in the total quantity transacted in the market. On net, an improvement in advice without a large contraction in the size of the market may suggest that fiduciary duty is beneficial for clients. While one may thus conclude that further increases in the stringency of fiduciary duty would continue to benefit clients, we argue in the subsequent sections that this statement depends on the mechanisms by which fiduciary duty induces these patterns. To uncover these mechanisms, we turn our attention to a structural model of the market for financial advice. The shifts in entry and advice documented in this section are also incorporated into the estimation of the parameters.

5. A Model of Fiduciary Duty

The previous sections have estimated the causal effect of extending common law fiduciary duty to BDs. However, they cannot speak to the mechanisms through which fiduciary duty operates. That is, it may be the case that fiduciary duty constrains low-quality advice, but it could also be the case that fiduciary duty solely increases fixed costs and that in the markets under study advisers that provide low-quality advice are also less profitable.[19] Determining which mechanism dominates is critical to understanding whether we can extend these results to speak to the effects of extending fiduciary duty to BDs at the federal level.

To make headway, we build a model of entry and advice provision. We first derive testable implications of the presence of an advice channel that depend only on the underlying economic structure: we can nonparametrically test whether fiduciary duty directly constrains low-quality advice. The model also provides a structure to quantify the channels. In Section 6, we implement both the nonparametric test and the structural quantification.

The intuition for the nonparametric test is simple: say firms earn profits as a function of the advice they give and of competition, and that there is heterogeneity across firms in both their profit-maximizing advice and their actual profits. In equilibrium, firms enter in decreasing order

[18] In Appendix B.4, we study whether fiduciary duty induced a compositional shift even within BD firms, and we divide firms into natural categories based on their footprint—e.g., whether they are local or national. We find evidence that local firms are most strongly affected by common law fiduciary status. Moreover, while results are noisy, we do not find any evidence of an increase in the number of firms of any footprint.

[19] One cannot assume that the advisers who offer the worst advice are also the most profitable: there is substantial heterogeneity across firms in commission schedules negotiated with FSP, scale, reputational considerations, and exposure to legal liability, among other issues.

of profitability until the marginal firm breaks even. If fiduciary duty only raises the fixed cost of doing business, the marginal firm would have to be more profitable, but the profit ordering would not change. This implies that the set of entering firms is contained by the set of entrants in the baseline. However, if fiduciary duty increases the cost of providing low-quality advice, this will alter the relative profitability of firms, potentially leading to a different set of advice in the market. Moreover, each entering firm may also change their advice. Thus, we might see the emergence of especially high-quality advice.

In this section, we introduce the model, formalize the intuition above, study its robustness to several extensions, and deliver a set of testable implications we can take to the data.

5.1. Elements of the Model

To begin, assume that all firms are BDs; we add RIA firms to the model in Appendix A.2. Suppose there are M *categories* of firms indexed by m. This is meant to capture that the effect of fiduciary duty can vary across local, regional, and national firms. Each firm j has a *type* $\theta_j \in [0,1]$ and can choose advice $a \in [0,1]$. We adopt the convention that higher values of a correspond to worse, or more distorted, advice. The distribution of types within category m is $H_m(\cdot)$. We assume $H_m(\cdot)$ is continuous, and we abuse notation by letting $H_m(S)$ denote the mass of types in set S. A firm of type θ and category m has a *base profit* function $\pi_m(a + g_m(\boldsymbol{\mu}); \theta)$ that we assume is single-peaked. As a normalization, we say that the maximum is attained at $a = \theta$ for some known value $\bar{\boldsymbol{\mu}}$. The actual profit of a firm of category m and type θ who enters and gives advice a when the equilibrium mass of entrants is $\boldsymbol{\mu} = (\mu_1, \ldots, \mu_M)$ is

$$f_m(\boldsymbol{\mu}) \cdot \pi_m \left(a + g_m(\boldsymbol{\mu}); \theta\right) - K_m,$$

where $f_m(\cdot)$ is decreasing in every component of $\boldsymbol{\mu}$, $g_m(\cdot)$ is increasing in each component of $\boldsymbol{\mu}$, and both are independent of θ. We conceptualize $f_m(\cdot)$ as the number of clients a firm receives if there are $\boldsymbol{\mu}$ entrants, $g_m(\cdot)$ as the direct effect of competition on advice, and K_m as the fixed cost of entry.

In equilibrium, the firms enter if and only if they make positive profits. Denote by $\mathcal{E}_m(\boldsymbol{\mu}, K_m)$ the set of types θ_j of category m who would enter if they believe that a mass $\boldsymbol{\mu}$ of firms of each category would enter and the fixed cost of entry is K_m. Then, for a fixed cost vector $\boldsymbol{K} \equiv (K_1, \ldots, K_M)$, an equilibrium consists of a mass $\boldsymbol{\mu}^*(\boldsymbol{K})$ such that

$$H_m\left(\mathcal{E}_m(\boldsymbol{\mu}^*(\boldsymbol{K}), K_m)\right) = \mu_m^*(\boldsymbol{K}).$$

It is instructive to discuss the elements of this model. First, θ captures the latent propensity to offer distorted advice. We remain agnostic about the sources of differences in θ. Firms may have negotiated different commission schedules with wholesalers and may also provide different splits of the commissions to individual advisers. They may also place different levels of emphasis on reputational considerations, or have different beliefs about the probability or cost of litigation. A key aspect of θ is that the costs of fiduciary duty—which we will model in detail below—may vary depending on the advice given and on firm category, but do not *directly* depend on θ. This is meant to capture that the effects of regulation can vary as a function of the actual advice given and the firm category (for example, local or national), but not on the latent profitability of giving worse advice.

Second, $f_m(\cdot)$ and $g_m(\cdot)$ capture the two ways in which competition can affect advice: by shifting the quantity of consumers a firm receives ($f_m(\cdot)$) and by directly changing advice ($g_m(\cdot)$). Since $f_m(\cdot)$ changes how total profits scale with competition, it is natural to assume that it decreases with each component of $\boldsymbol{\mu}$. Note that we are excluding a direct effect of θ on $f_m(\cdot)$, essentially ruling out that the mass of consumers received by a firm (conditional on their category) is a function of their advice quality. We find this assumption realistic for a number of reasons. First, given the previous evidence on the lack of consumer information in this market (SEC, 2011, 2013a,b; Egan et al., 2019), it seems unlikely that consumers are sorting to advisers based on unobserved profitability differences that remain after conditioning on firm observables captured by m; sorting that depends on characteristics like whether the firm is nationally recognized are captured through the dependence on m. Second, this assumption is analogous to assuming that θ enters into $f_m(\cdot)$ in a multiplicatively separable fashion, so that we can envelope the effect of θ on $f_m(\cdot)$ into π, which does depend flexibly on θ. Thus, the restriction that $f_m(\cdot)$ is independent of θ is saying that the effect of the type on profits does not *differentially* change with competition.

Next, consider $g_m(\cdot)$. We introduce this function to allow for competitive effects on advice—in particular, for the possibility that increased competition directly improves advice. Upon entry, a firm will choose advice a to maximize $\pi(a + g_m(\boldsymbol{\mu}); \theta)$. Thus, $g_m(\cdot)$ shifts the location of optimal advice without directly affecting profits. As discussed before, we will assume that $g_m(\boldsymbol{\mu})$ is increases in each component of $\boldsymbol{\mu}$, so that increasing competition improves advice by shifting the optimal advice $a^*(\theta; \boldsymbol{\mu}) \equiv \arg\max_a \pi(a + g_m(\boldsymbol{\mu}); \theta)$ to the left. We believe that this monotonicity assumption is justifiable for a number of reasons. Tougher competition makes it easier for consumers to visit multiple financial advisers and identify questionable advice, as in some credence goods models (Dulleck and Kerschbamer, 2006). Furthermore, evidence from Egan et al. (2019) suggests that financial advisers with misconduct records are more likely to survive in markets with lower competition. Third, given that the "price" of the product is the same regardless of which adviser the client visits, concerns like showrooming effects—in which competition decreases the incentive to provide effort in advising clients—are not present in this market. Finally, firm strategies that depend on the distribution of θ likely also rely on consumers' knowledge of θ for each firm, which is unlikely in this setting. As with $f_m(\cdot)$, we still let $g_m(\cdot)$ depend directly on m so that consumers can be influenced by more salient aspects, like whether the firm is national.

Finally, we do not let $f_m(\cdot)$ or $g_m(\cdot)$ depend directly on whether the market has fiduciary duty. Arguing that $f_m(\cdot)$ and $g_m(\cdot)$ changes due to demand side factors induced by fiduciary standards suggests that imposing common law fiduciary duty changes how many people go to various firms, what type of firms they go to, or what sort of products they ask for when they arrive at these firms. Given the substantial survey evidence cited above that clients are not even aware of the fiduciary status of their advisers, we find it a priori implausible that consumers are making decisions about which advisers to talk to based on the common law fiduciary status of the state.

To illustrate the model, consider the case with $M = 1$ category and $g(\cdot) = 0$. Define $\pi^*(\cdot) \equiv \max_a \pi(a; \theta)$. Given that we do not take a stance on the source of heterogeneity, we also cannot take a stance on the behavior of $\pi(\cdot; \theta)$, and thus $\pi^*(\theta)$, with θ. Figure 4(a)–(c) illustrates three possibilities for $\pi^*(\cdot)$ and sample graphs of $\pi(\cdot; \cdot)$. Panel (a) illustrates the case where worse advice corresponds to highest profits. As discussed above, however, higher θ firms may in fact have lower profits so that cases such as (b) and (c) are also possible. Below, we develop predictions that hold over any shape of $\pi^*(\cdot)$.

5.2. The Fixed Cost Channel

We return to the general model. We say that fiduciary duty operates through a *pure fixed cost channel* if imposing fiduciary duty on a market increases fixed costs of entry from K_m to $K'_m \geq K_m$ for all θ but does not alter $\pi(\cdot; \cdot)$ or the distribution $H_m(\cdot)$ of types in any way. This increase in fixed costs could correspond to compliance software or insurance, increased paperwork, increased overhead time required to deal with regulation, increased effort dedicated to oversight, etc.[20] In Appendix A.1.2, we prove the following.

Proposition 1. *Suppose $K'_m \geq K_m$ and that $\mu^*_m(\boldsymbol{K}') \leq \mu^*_m(\boldsymbol{K})$. Define $\mathcal{E}_m \equiv \mathcal{E}_m(\boldsymbol{\mu}^*(\boldsymbol{K}), K_m)$ and $\mathcal{E}'_m$ analogously. Then, $\mathcal{E}'_m \subseteq \mathcal{E}_m$.*

Proposition 1 states that if only the fixed cost increases, and if this leads to weak decreases in the mass of each category of firm, then the new set of firms who enters is a subset of the original set of firms. The assumption that $\mu^*_m(\boldsymbol{K}') \leq \mu^*_m(\boldsymbol{K})$ is a not one on primitives. However, it is natural to expect that an increase in fixed costs leads to a decrease in entry is a natural one. To formalize this intuition, Lemmas 1 and 2 in Appendix A.1.1 consider the simpler model with $M = 1$

[20] In this section, we write the change in fixed costs as a change to the fixed costs of entry. We can instead have a constant fixed cost of entry and say that the effect of the fixed cost channel is to change the base profit function from $\pi(\cdot; \cdot, \cdot)$ to $\pi(\cdot; \cdot, \cdot) - c$. This would correspond to an increased per-transaction cost due to fiduciary duty. The key similarity, as discussed later, is that c is independent of advice and the ordering of profitability of types does not change with the imposition of fiduciary duty. Essentially, one should think of the "fixed" cost as fixed across types.

Figure 4: Illustration of $\pi(\cdot;\cdot)$ and $\pi^*(\cdot)$, and the effects of a pure fixed cost channel

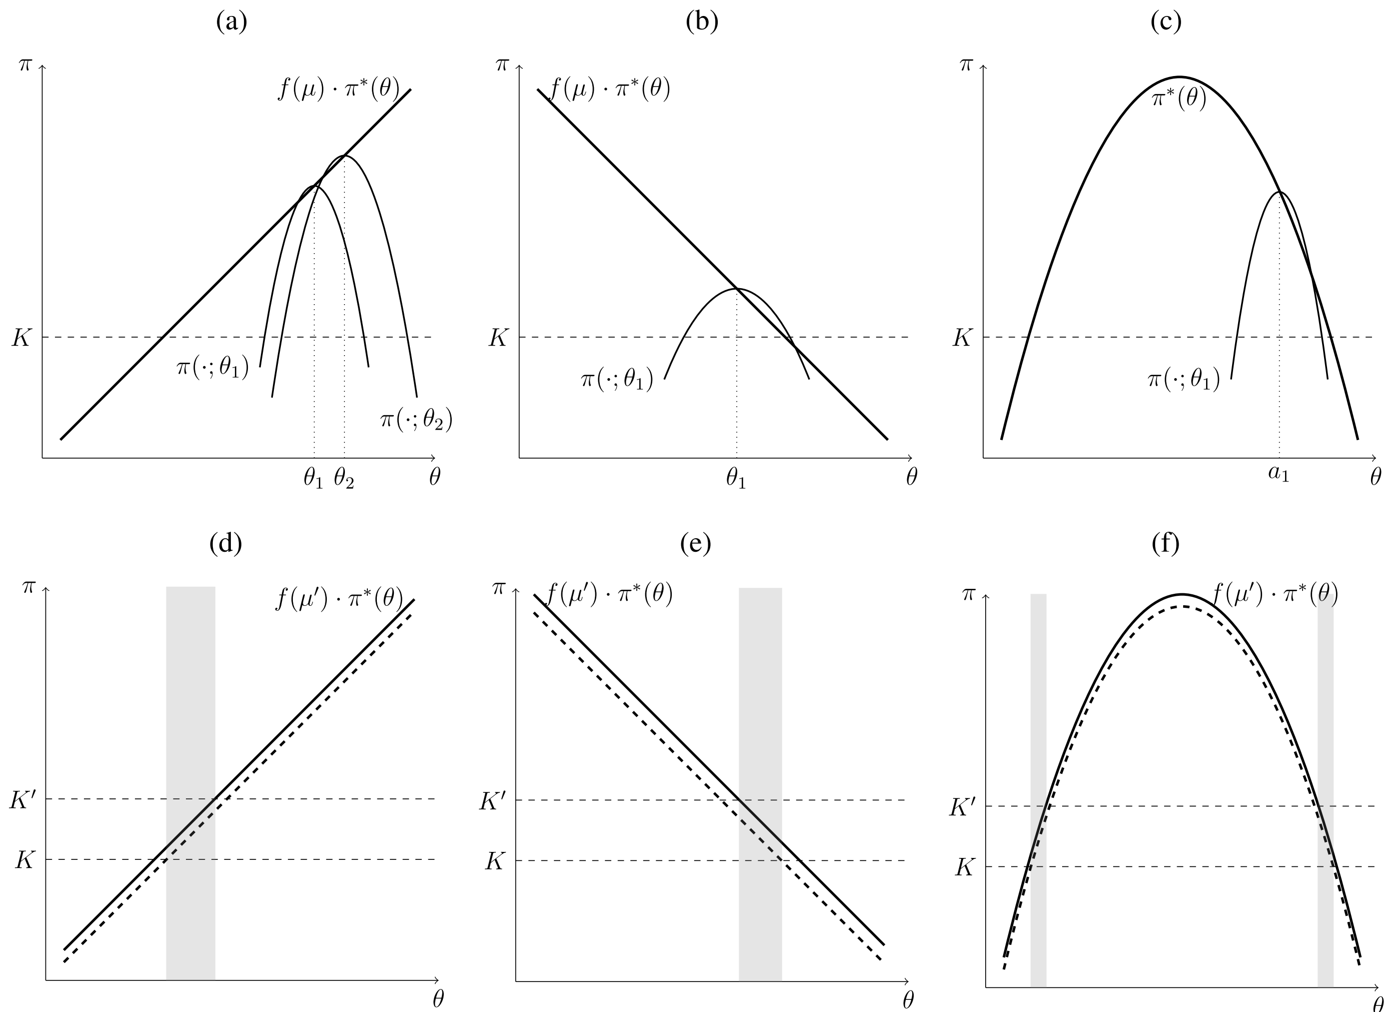

Different possible profit envelopes $\pi^*(\cdot)$, along with plots of the underlying $\pi_i(\cdot;a)$ that generate them. The fixed cost K is presented, and the fixed cost channel involves increasing this value. Panels (d)–(f) illustrate the effects of a pure fixed cost channel, by increasing the fixed cost from K to K'. The shaded types are the ones who exit the market. Note that types map directly to advice (in the same way) in each panel, but we do not show the underlying density $H(\cdot)$ of types.

and verify that the equilibrium is unique and the comparative statics with fixed costs imply that the number of entrants decreases with fixed cost increases.[21] We impose this assumption for two reasons. With $M > 1$ categories it is in principle possible to have the mass of one category increase due to decreased competition from another. Furthermore, given a partition of firms into categories, the mass of firms that enters is observable. Thus, this condition is testable and empirically useful.

Note also that the type θ can be multidimensional, to incorporate effects like provision of different advice to different groups of clients. Appendix A.1.3 provides some examples and argues that the testable predictions below do not change. The key connection between these generalizations—as discussed at the start of this section—is that the above inclusion holds as long as fiduciary duty does not change the relative profitability of different types of firms. Thus, it simply shrinks the set of types who enter rather than rearranging them.

Since θ is not observable to the econometrician, to take Proposition 1 to the data we look for predictions on advice. In the following observation, we denote by $\underline{a}(\boldsymbol{K})$ and $\bar{a}(\boldsymbol{K})$ the least and most distorted advice observed among any entrants of any category in the market, as a function of the fixed costs.

Proposition 2. *Suppose $K'_m \geq K_m$ and that $\mu^*_m(\boldsymbol{K}') \leq \mu^*_m(\boldsymbol{K})$. If $g_m(\boldsymbol{\mu}) = 0$ for all m, then $\underline{a}(\boldsymbol{K}') \geq \underline{a}(\boldsymbol{K})$ and $\bar{a}(\boldsymbol{K}') \leq \bar{a}(\boldsymbol{K})$. If $g_m(\boldsymbol{\mu})$ is increasing in every component of its argument, $\underline{a}(\boldsymbol{K}') \geq \underline{a}(\boldsymbol{K})$.*

We prove this proposition in Appendix A.1.2. Under the pure fixed cost channel, the set of types that enter the market under fiduciary duty is a subset of the set that enters without. If competition does not have a direct impact on advice, then it must be that the advice we observe is also a subset. This would imply that the best advice in the market must (weakly) worsen and the worst advice should (weakly) improve. If competition improves advice, exit induced by the fixed cost increase would worsen all advice; thus, the prediction on best advice remains while the prediction on worse advice is now ambiguous. Thus, one testable prediction is that under the fixed cost channel the best observed advice does not improve when imposing fiduciary duty.

Importantly, there are no analogous predictions for how fiduciary duty affects moments such as the mean of the advice distribution, even if it operates purely through a fixed cost channel. This is because we are not taking any stance on the shape of $\pi^*(\cdot)$ or $H(\cdot)$. Panels (d)–(f) of Figure 4 illustrate the effects of increasing the fixed cost in panels (a) through (c), restricting to $M = 1$ and $g(\cdot) = 0$. In each situation, K increases to K', but the effective profit function $f(\mu) \cdot \pi^*(\cdot)$ also increases slightly due to exit of firms, from the dashed to the solid lines. On net, however, firms exit, as denoted by the shaded areas. In panel (d), fiduciary duty operating through a fixed cost channel

[21] We can in fact go further and say that even if there are firms who are not directly impacted by fiduciary duty, as long as competition between different firm categories is "not too strong"—in a manner that can be formalized—then the aforementioned comparative statics hold.

increases the mean a since $\pi^*(\cdot)$ increases in θ and increasing the fixed cost simply excludes low-θ firms from the market. In panel (e), the argument is reversed. In panel (f), the effect on the mean depends on $H(\cdot)$. In all three panels, however, the extremes of advice (weakly) decrease.

A second prediction relates to how a particular firm changes the advice it provides as a function of fiduciary duty. Suppose first that competition does not directly impact advice. Then, if a firm is able to cover the fixed cost of entry, the advice it provides does *not* depend on the fixed cost. If instead competition directly improves advice, then if the imposition of fiduciary duty increases fixed costs, the advice a firm provides will (weakly) worsen. We formalize these observations in the following.

Proposition 3. *Suppose $K'_m \geq K_m$ and that $\mu^*_m(\boldsymbol{K}') \leq \mu^*_m(\boldsymbol{K})$. Let $a^*_m(\theta; \boldsymbol{K})$ be the advice provided by a type θ firm of category m who enters when costs of entry are $\boldsymbol{K}$. Then $a^*_m(\theta; \boldsymbol{K}) \leq a^*_m(\theta; \boldsymbol{K}')$, with equality if $g_m(\cdot) = 0$.*

The proof, which we omit, notes that $a^*_m(\theta; \boldsymbol{K}) \equiv \arg\max_a \pi_m(a + g_m(\boldsymbol{\mu}); \theta)$ does not depend on $\boldsymbol{K}$ directly, and the direct effect of competition simply shifts the location of the maximum of the profit function. The testable implication is that under a pure fixed cost channel we should not see the advice of a firm improving upon imposition of fiduciary duty.

5.3. The Advice Channel

Alternatively, fiduciary duty could make it differentially more costly to offer low-quality advice. We call this effect the *advice channel*. To model this channel, we say that the imposition of fiduciary duty introduces a cost function $c(a)$ with $c'(a) > 0$. The profit to type θ from giving advice a is then $\pi_m(a + g_m(\boldsymbol{\mu}); \theta) - c(a)$. In this section, we will show that the predictions outlined in the previous section need not hold under an advice channel.

As an illustration, set $g_m(\cdot) = 0$ and suppose $c(\cdot)$ is such that fiduciary duty places a cap on advice: $c(a) = 0$ for $a \leq \bar{a}$ and $c(a)$ is infinite for $a > \bar{a}$. This leads to two effects not present in a fixed cost channel. First, Figure 5(a) illustrates that firms with especially high values of θ, such as θ_2, cannot profitably offer any level of advice, and will be forced to exit. If there is exit of high θ firms, this makes it profitable for very low-θ firms to now enter, leading to the appearance of previously unprofitable high-quality advice. That is, the lowest type $\underline{\theta}$ that enters decreases, and thus the highest-quality advice observed improves as well. Second, a firm that remains in the market after the imposition of fiduciary duty can actually improve its advice. Firms with moderately high values of θ, such as θ_1, will still profitably operate but will adjust their advice to $\bar{a} < \theta_1$. Neither of these observations could be rationalized through a pure fixed cost channel.

These observations are robust to any increasing $c(\cdot)$ and not a consequence of the stark assumption that fiduciary duty places a cap on advice. If $c(\cdot)$ is increasing, then it effectively acts as a

Figure 5: Further illustration of the model

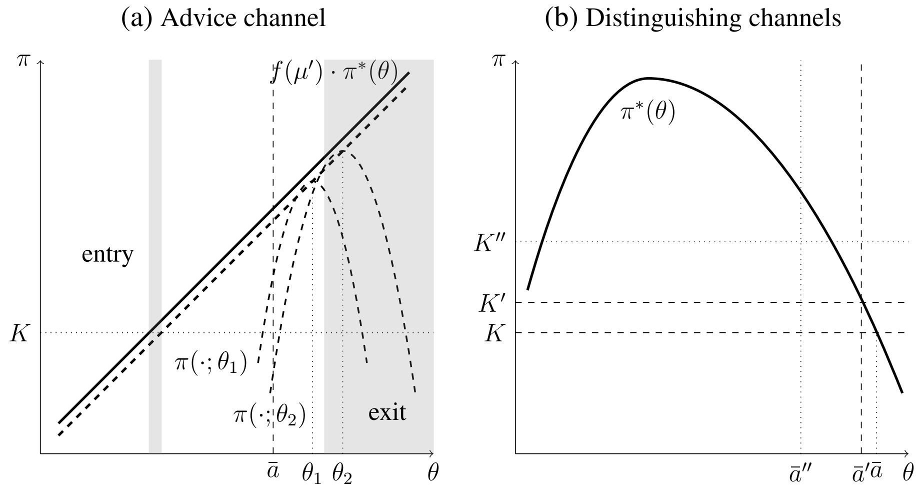

(a) Moving from the baseline (thick, dashed lines) to a fiduciary standard in which advice can be no larger than $\bar{a}$. The shaded area to the right illustrates types who exit due to the regulation since they cannot profitably adjust their advice. The shaded area to the left illustrates types offering previously unprofitably good advice who enter since the effective profit function increases due to the aforementioned exit. (b) A profit envelope under which strengthening fiduciary standards will lead to different results under a pure fixed cost channel and an advice channel (proxied by a cap)

handicap for higher-θ firms and can induce them to exit the market, leading to entry of lower-θ firms. Also, it is not necessarily the case that only high θ firms will improve their advice. Indeed, in the absence of a competitive effect on advice, all firms will have an incentive to improve their advice.[22] This also implies that in general, the emergence of high quality advice upon imposing fiduciary duty can come both from firms who only enter under fiduciary duty and from firms who enter in both regulatory regimes improving their advice.

One should not interpret the previous observations as necessary conditions for an advice channel. It is still possible for both extremes of the advice distribution to contract and for firms who enter both with and without fiduciary duty to offer worse advice under the more stringent standard, just like in a pure fixed cost channel. For example, if competition improves advice, then exit of low quality firms might lead surviving firms to worsen the advice they give. This would happen if the effect of competition is stronger than the effect of the cost of providing distorted advice, and could lead to a contraction of the best observed advice. Moreover, note that if an advice channel is present, then the worst advice could also worsen upon imposing fiduciary duty: in the case where firm types are multidimensional (see Appendix A.1.2), it is possible for the advice channel to induce entry of firms who give low a to most types of consumers but especially high a to a small set of them. The key observation, however, is that in an advice channel—unlike in a fixed cost channel—it is not necessary that both extremes of the advice distribution contract or for within-firm advice to worsen.

[22]See Appendix A.1.4 for a simple argument with monotone comparative statics.

5.4. The Importance of Distinguishing These Channels

We have argued that distinguishing whether common law fiduciary duty operates through the advice channel or through the fixed cost channel offers insights into the effects of extending fiduciary duty at the federal level, and that quantifying the effect of fiduciary duty on the mean of observed advice is not sufficient to identify the channel through which it operates. We can now use the model to formalize these statements. First, consider the situation in Figure 5(b), and suppose that in the baseline market without any fiduciary standards, the worst observed advice is $\bar{a}$, and that imposing fiduciary standards moves the worst observed advice to $\bar{a}'$. This shift could be rationalized by either fixed costs moving to K' or a cap of $\bar{a}'$ being imposed through fiduciary standards. Second, assume that the regulator is considering making the policy more stringent.[23] In an advice channel, tightening the cap to $\bar{a}'' < \bar{a}'$ would push low-quality advice out of the market. However, tightening a fixed cost channel to $K'' > K'$ would induce exit of both high and low quality advice.

This figure also highlights that the external validity of the causal effect depends critically on whether fiduciary duty operates through the advice channel or the fixed cost channel. In the former, every surviving firm will distort their advice weakly less, leading to an overall improvement of average advice. In the latter, whether average advice increases or decreases depends on whether more low-quality or high-quality advice firms are displaced.

These two channels are neither mutually exclusive nor exhaustive: fiduciary duty could both increase fixed costs and constrain advice, and it could be the case that it affects neither. In what follows, we first test the hypothesis that there is no advice channel. We then make decompose the effect of fiduciary duty into its effect on advice and its effect on fixed costs by leveraging the full structure of the model as well as parametric assumptions for estimation.

6. Quantifying Mechanisms and Effects of Counterfactual Regulation

In this section, we decompose the effect of fiduciary duty observed in Section 4 into a part that is due to the advice channel and one that is due to an increase in fixed costs. Beyond the inherent interest in understanding the mechanisms through which this regulation operates, doing so also allows us to study the effect of more stringent versions of fiduciary standards. To begin, in Section 6.1 we implement the nonparametric tests proposed in Section 5. This allows us to establish the presence of an advice channel before imposing further structure. Section 6.2 introduces a parameterization that allows us to take the full model to the data. Importantly, the parameterization we impose is based directly off the quasi-experimental variation explored in Section 4. Section 6.3 discusses

[23] Stringency of fiduciary duty regulations is a matter of current policy debate. Advocates of the defunct DOL Rule argue that the SEC's Best Interest Regulation does not live up the same standards. Proposed state legislation (rather than common law) is also anecdotally of different stringencies, especially since enforcement methods will be different.

identification and estimation of this model, and Section 6.4 presents estimation results. Using these results, Section 6.5 unpacks the roles of the advice channel and the fixed cost channel in changing observed advice, and Section 6.6 simulates the effects of counterfactual stringency levels for fiduciary standards.

6.1. Nonparametric Tests for the Presence of an Advice Channel

Consider two identical markets, but where only one imposes fiduciary duty on BDs. We aim to test whether an advice channel exists, i.e., whether fiduciary duty engenders a direct constraint on low-quality advice. The primary test derived in Section 5 is at the market level. Under a pure fixed cost channel, the highest quality advice offered by any BD in the market with fiduciary duty is weakly worse than the highest quality advice offered in the market without. Under the advice channel, this highest quality advice can improve.[24]

We use our preferred metric of risk-adjusted returns as the measure of the quality of advice, partialling out border, contract month, and age fixed effects, to arrive at a "normalized" risk-adjusted return that is comparable across all transactions. The test is based on the support of the distribution of this advice across adviser types, and we proxy for the support with the mass in the tails, i.e., the proportion of normalized returns that are above x for large values of x.[25]

The row marked "BD Proportion" of Table 6 shows the proportion of normalized returns above various cutoffs for BDs without fiduciary duty; "BD Difference" shows the change in this proportion when moving to border counties with fiduciary duty. For extreme cases, we find an economically and statistically significant increase in this proportion, consistent with an expansion of high-quality advice when imposing fiduciary duty. For RIAs, we find that changes in the shares in the tails are economically and statistically zero, which lends further credence to the fact that the changes in the distribution for BDs are not spurious. Again, the expansion in high-quality advice cannot be explained by a pure fixed cost channel but is consistent with the presence of an advice channel.

The model also provides a firm-level test. In a pure fixed cost channel, if a BD firm enters both markets, it offers weakly worse advice in the market with fiduciary duty. Under an advice channel, this firm may improve its advice under fiduciary duty. This test, however, is likely to be underpowered: if fiduciary duty does not greatly affect the cost of providing high-quality advice,

[24] The tests in this section are predicated on a decrease in the number of BD firms in the market, which Section 4.4 supports. Moreover, Appendix B.4 suggests that there is no evidence of increases in the number of BD firms of any geographic footprint—a proxy for "categories."

[25] Suppose we have two distributions A and B (with continuous and strictly increasing cdfs on their support) with the maximum M_A of the support of A strictly less than the maximum M_B of the support of B. We know that $F_A(M_A) = 1$ and $F_B(M_A) < F_B(M_B) = 1$, where $F_T(\cdot)$ is the cdf of T. Thus, $F_A(M_A) > F_B(M_A)$, so for x sufficiently close to M_A, $1 - F_A(x) < 1 - F_B(x)$ as well. For similar reasons, we could look at the effect on extreme quantiles; results are similar and available upon request. Mass in tails or quantiles are less sensitive to single observations than estimates for the support.

Table 6: Effects on tails of risk-adjusted return distribution

Cutoff	0.010 (1)	0.015 (2)	0.020 (3)	0.025 (4)	0.030 (5)
BD Proportion	0.063	0.008	0.006	0.003	0.003
	(0.008)	(0.002)	(0.002)	(0.002)	(0.002)
BD Difference	0.003	0.012***	0.010***	0.010***	0.006**
	(0.011)	(0.004)	(0.004)	(0.004)	(0.003)
RIA Proportion	0.116	0.048	0.030	0.015	0.009
	(0.004)	(0.002)	(0.002)	(0.001)	(0.001)
RIA Difference	-0.002	0.002	0.002	-0.001	-0.002
	(0.005)	(0.003)	(0.002)	(0.001)	(0.001)

Proportion of normalized risk-adjusted returns above various cutoffs as a function of adviser type and fiduciary duty. "BD Proportion" refers to the mass of advice above each cutoff for BDs in states without fiduciary duty. "BD Difference" is the difference in this quantity for BDs with and without fiduciary duty. The rows for RIAs are analogous. Standard errors are computed through the bootstrap. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

then most firms entering both markets will not shift their recommendations. Nevertheless, we estimate (1) for all outcomes considered in this paper but also add firm fixed effects. Table B.6 in Appendix B.5 shows results of this analysis. While the results are noisy, as expected, the sign of the within-firm effect is broadly consistent with an increase in quality. This would not happen under a pure fixed cost channel.

6.2. Parameterization of the Model

We now leverage the structure of the model to not just test for the presence of an advice channel but to quantify the relative contribution of the fixed cost and advice channels. In doing so, we mirror the decisions made in the reduced-form analysis as closely as possible.

We follow the model presented in Section 5 but add a parsimonious but flexible parametric structure to the elements to take the model to data. In particular, we parameterize the profit of a firm f of status $T \in \{BD, RIA\}$ with type θ_f that operates in market m as:

$$\pi^T_{fm}(\theta_f) \equiv \max_a f^T\left(N_{BD}, N_{RIA}\right) \cdot \pi^T(\theta_f, a) - K_{mf}. \tag{2}$$

Let $-T$ denote the opposite of T. We let $f^T\left(N_{BD}, N_{RIA}\right) \equiv \left[(N_T+1)^\gamma + \alpha \cdot (N_{-T}+1)^\gamma\right]^{-1}$ parameterize the effect of competition. This functional form allows the entry of firms to expand the total market (if $\gamma < 1$), and it allows for the number of broker-dealer and RIA firms to both affect profits and to enter asymmetrically. The other two terms in the max operator are parameterized as

$$\pi^T(\theta_f, a) \equiv \delta^T_0 - \delta^T_1 \cdot \left(\theta_f - \bar{\theta}^T\right)^2 - \tilde{\lambda} \cdot (\theta_f - a)^2 - \tilde{c} \cdot a^2 \cdot \mathbb{1}[FD]_m \cdot \mathbb{1}[T = BD]. \tag{3}$$

The function $\pi^T(\cdot)$ parameterizes the profits that a type receives for offering advice a. The functional form is deliberately parsimonious, but it is such that base profits can be increasing, decreasing, or non-monotone in the firm's type θ_f, as δ_1 and $\bar{\theta}$ flexibly govern which of the cases in Figure 4 are empirically relevant. The final terms of (3) govern the advice channel. In markets without fiduciary duty, firms set $a = \theta_f$. In markets with fiduciary duty for BDs, we impose an additional cost of distorting advice, which we parameterize as $\tilde{c} \cdot a^2$. To maximize profits the firm will set $a = \theta_f/(1+\tilde{c}/\tilde{\lambda}) \equiv \theta_f/(1+c)$. When c is positive, this induces a deviation from the firm's optimal advice, which lowers profits. Substituting into (3) and defining λ appropriately, we find that at the optimal advice level

$$\pi^{T*}(\theta_f) = \delta_0^T - \delta_1^T \cdot \left(\theta_f - \bar{\theta}^T\right)^2 - \lambda \cdot \theta_f^2 \cdot \mathbb{1}[FD]_m \cdot \mathbb{1}[T = BD].$$

Thus, the advice channel is parameterized by c, which governs how much distortion is affected by fiduciary duty, and λ, which governs how much distortion affects profits.[26]

The final term K_{mf} of (2) is the fixed cost of entry, which we parameterize as

$$K_{mf} = \kappa_0 \cdot \mathbb{1}[FD]_m + \kappa_1 \cdot \mathbb{1}[BD]_f + \cdot \mathbb{1}[FD]_m \cdot \mathbb{1}[BD]_f \cdot \big(\kappa_{2L} \cdot \mathbb{1}[\text{Local}]_f + \kappa_{2R} \cdot \mathbb{1}[\text{Regional}]_f + \kappa_{2N} \cdot \mathbb{1}[\text{National}]_f\big) + X_{mf}\beta + \xi_{b(m)} + \epsilon_{mf}. \quad (4)$$

The first three terms allow for fixed cost differences across markets with and without fiduciary duty, differences between BDs and RIAs, and an interaction in these differences as well. The coefficients $\kappa_{2\times}$ parameterizes the magnitude of the fixed cost channel for local, regional, and national firms, as categorized by Discovery. We allow this to vary by firm footprint, since one may believe that for firms with larger footprints any changes in fixed costs may be concentrated at the central level rather than spread out over branches. We also control for firm- and county-level covariates: log population, log median household income, and log median home price. Interactions of these county covariates with BD status and footprints are all encapsulated in X_{mf} for notational compactness.[27] Importantly, we include a full set of border fixed effects $\xi_{b(m)}$, allowing for the possibility that fixed costs vary arbitrarily at the border level, and estimate this model on the same sample of border counties as in previous sections. Finally, we include an unobserved firm-market-specific profit shifter $\epsilon_{fm} \sim N(0,1)$, which also provides the scale normalization in this model.

When making its entry decision, the firm knows its own θ_f and ϵ_{fm} draw. We assume that it does

[26]Note that $\tilde{c} = c = 0$ for RIAs, consistent with the fact that they do not face differential fiduciary standards across markets. The extent to which their national fiduciary standard does penalize distortion would be captured by a different type distribution and profit function for RIAs.

[27]For instance, we include terms such as $\mathbb{1}[BD]_f$ interacted with log population and whether the firm is regional. The main effects for national and regional, as well as the interactions with fiduciary status and broker-dealer status are also all in X_{mf} and omitted from the notation.

not know the realizations of θ_f of other potential entrants, but it does know the distribution from which they are drawn. A firm enters if and only if it expects to make positive profits conditional on entry, given its beliefs over the entry probabilities of all other potential entrants in the market. An equilibrium is such that beliefs are consistent with true entry probabilities. For instance, in an equilibrium where each firm f has a probability p_f^* of entry when integrating out over the realizations θ_f and ϵ_f, it must be for a BD firm f that

$$\int \Pr\left(\mathbb{E}_{p^*_{-f}}\left[f(N_{BD}+1, N_{RIA})\right] \cdot \pi^{BD}(\theta_i) - K_{mf} + \epsilon_f \geq 0\right) \, dH^{BD}(\theta_i) = p_f^*, \tag{5}$$

where the $\Pr(\cdot)$ is taken over realizations of ϵ_f and the inner expectations is taken over realizations of N_{BD} and N_{RIA} given the equilibrium entry probabilities p^*_{-f} of all other firms. The system specified by (5) and analogous equations for DRs define an equilibrium.[28]

Since in this model a firm would always issue the same advice were it not for fiduciary duty, we incoporate a degree of "measurement error" into the framework. In particular, when a firm gives advice a_{ft} on a transaction t, we observe $\tilde{a}_{ft} \equiv a_{ft} + \epsilon^a_{ft}$ where $\epsilon^a_t \sim N(0, \sigma_a^2)$.[29] In the data, we take our preferred quality metric—the risk-adjusted returns of the product sold—as the backbone of our measure of $\tilde{a}$. We say that $\tilde{a}_t = \bar{r} - r_t$, where r_t is the residualized risk-adjusted return—after partialling out border, contract-month, and client age fixed effects—and $\bar{r}$ is the 99.5$^{\text{th}}$ percentile of this distribution. In line with the convention in Section 5, larger values of a_t correspond to more distorted advice. We parameterize the distribution $H^T(\cdot)$ from which firms' types are drawn as normal with mean μ_θ^T and standard deviation σ_θ. This distribution does not depend on the market, as we assume that all cross-market differences in advice are controlled by the fixed effects.

To take the model to the data, we need to take a stance on the set of potential entrants. We follow a common approach in this literature of using "nearby" firms as potential entrants.[30] In particular, we assume that (i) national firms are potential entrants in all markets, (ii) any regional firm operating at a border is a potential entrant everywhere at that border, and (iii) the number of local potential entrants in every county equals one more than the maximum number of local entrants in any county within that border. When we generate local potential entrants, we assume their types

[28] Lemma 4 in Appendix A.2 shows that in a simpler case with two types (BD and RIA, say), and no unobserved shocks, the equilibrium is unique as long as cross-type competitive effects are not too strong. We do not have a proof that equilibria are unique in our richer empirical model. This is not an issue for estimation of the model, as we follow a two-step approach and recover beliefs regarding entry probabilities before estimating the model. For counterfactuals, we follow one common approach in the literature (e.g., Seim (2006)) and look for multiplicity by testing different starting points. Across the board, starting our solver from different initial guesses results in the same fixed point.

[29] We can interpret ϵ^a as the result of tailoring advice to different clients, as long as this tailoring is not relevant to the cost of distortion.

[30] Berry (1992) uses airlines operating at both points of a route as potential entrants. Roberts and Sweeting (2013) use nearby bidders as potential bidders in an auction. Jia (2008) uses (essentially) the maximum the number of small entrants in similar markets, which is similar to our choice for local potential entrants.

θ_f are independent of other firms, but national and regional potential entrants retain their identity and thus their θ_f.

The decisions in this parameterization mirror the ones made in the reduced-form section as closely as possible. Counties at the same border share a fixed effect that affects entry profitability (and advice through the normalization). We allow the entry cost to depend on a full interaction of BD status and fiduciary duty, which mirrors the difference-in-difference specification throughout the reduced-form analysis and allows the entry pattern of RIAs to serve as a control for BDs. The distribution of types of potential entrants is common on both sides of the border, allowing the model to exploit an implicit comparison between counties. Finally, we force θ_f to be fixed within firm, which causes the model to use within-firm comparisons to inform the advice channel, like the fixed-effects regressions from Section 6.1.

This model is reminiscent of entry models with heterogeneous competitors. Unlike these models, we do not allow the type of a competitor (quality in Mazzeo (2002), location in Seim (2006), and θ with BD/RIA status in this paper) to be a choice. We believe this is a realistic assumption in this setting: a firm will likely not change its licensing status or any internal policies that cause it to distort advice more or less than competitors for every market that it enters. However, we still allow for the types that choose to enter to be a selected subset of the latent distribution. Our model also incorporates firm-level unobserved heterogeneity in ϵ_{fm} and market-level unobserved heterogeneity in ξ_m, like Berry (1992) or Seim (2006). A difference is that other papers incorporate market-level heterogeneity as a random effect; instead, we continue with the reduced-form strategy of comparing similar markets along a border and incorporate a fixed effect common to a subset of markets.

6.3. Identification and Estimation

Before discussing the estimation procedure, we provide some intuition for how the parameters are identified. While the tests in Section 6.1 are implemented only to test for the presence of an advice channel, the magnitudes of these effects—the prevalence of especially high-quality advice in markets with fiduciary standards relative to other similar markets as well as the extent of within-firm changes in advice—are informative of the magnitude of c. A stronger advice channel reduces profits and thus reduces entry in markets with fiduciary duty, but the magnitude of the fixed cost channel $\kappa_{2\times}$ can move to match the entry rate by footprints in these markets. As in canonical entry models, variation in the number of potential entrants informs the competitive effect. In our setting, however, we also have quasi-exogenous policy variation that can help pin down cross-type competitive effects: since fiduciary duty differentially affects the profits of BDs without directly affect those of RIAs, the responsiveness of RIA entry to fiduciary duty will inform how strongly the two types of advisers compete with each other. Finally, fixed cost shifters (such as fiduciary duty but also variation in covariates) will change the distribution of advice in the market differently depending on the shape of the profit function, which is governed by $\bar{\theta}$. For instance, as illustrated in Figure 4, if changes in market covariates that lead to increases in fixed costs lead to a more contracted distribution of advice, this is evidence that the profit function may be an inverted-U (like in panel (c) of Figure 4).

The intuition behind identification suggests that the distribution of advice, rather than just its mean, provides important information about the parameters. As such, we employ a likelihood approach that can leverage the full distribution we observe. If θ_f were observed for each f, estimation would amount to a probit. Since θ_f is unobserved, the standard option would be to compute the likelihood of observing the (i) entry decisions and the (ii) advice provided by each firm, integrating out over θ_f. Optimizing this likelihood is cumbersome, especially with a moderately large number of border fixed effects. We instead take a computational Bayesian approach of a Metropolis-in-Gibbs sampler with data augmentation, built off the Gibbs sampler developed by McCulloch and Rossi (1994) for a probit.

An ingredient into the Gibbs sampler is the equilibrium beliefs that firms have over their opponents' entry probabilities. Instead of computing an equilibrium for each candidate set of parameters in the estimation procedure, we use a two-step approach, as in Sweeting (2009). In the first stage, we use the observed probabilities of entry across markets to predict an empirical probability of entry.[31] At the market-potential entrant level, we estimate a linear probability model of whether a firm enters on the same set of covariates as in our fixed cost parameterization in (4). From this regression, we arrive at an estimated probability of entry $\hat{p}_{fm}$ for each potential entrant, from which we derive beliefs that a firm has over the distribution of competitor BDs and RIAs conditional on entry. This allows us to compute $\mathbb{E}[f^T(\cdot)]$ for each firm, conditional on entry and given γ and α. A benefit of this approach is that it is robust to multiplicity of equilibria. In the second stage of estimation, we follow the steps below.

0. *Initialize.* Pick a guess for all parameters. Augment the parameters with a guess for $\theta_f^{(0)}$ for all firms and draws $\epsilon_{fm}^{(0)}$ for each firm f in each market m such that with these shocks, f makes positive profits in m if and only if it enters in the data. Set $i = 0$.

1. *Metropolis Step for θ_f.* For each f, draw a θ'_f for each $\theta_f^{(i)}$ from a proposal distribution $Q_\theta(\cdot|\theta_f^{(i)})$. If given all other parameters and the maintained draws for $\epsilon_{fm}^{(i)}$, the implied entry decisions are consistent with observed ones, then compute

$$L_f(\theta; c) \equiv \prod_m \phi\left(\tilde{a}_{ft} - \frac{\theta_f}{1 + c \cdot \mathbb{1}[BD]_f \cdot \mathbb{1}[FD]_m}; 0, \sigma_a^2\right),$$

[31] While Sweeting (2009) uses multiple observations of entry into the same market, we use observations of entry into similar markets. We can use this procedure since we omit market-level random effects from the model in favor of fixed effects at the border level, which groups together multiple markets.

where $\phi(\cdot;\mu,\sigma^2)$ is the pdf of $N(\mu,\sigma^2)$. We set $\theta_f^{(i+1)} = \theta'_f$ with probability $\min\left[L_f(\theta'_f)/L_f(\theta_f^{(i)}), 1\right]$ and to $\theta_f^{(i)}$ otherwise.

2. *Update μ_θ, σ_θ, and σ_a.* We draw $\mu_\theta^{(i+1)}$ and $\sigma_\theta^{(i+1)}$ from a Bayesian OLS of $\theta_f^{(i+1)}$ on dummies for BD and RIA.[32] We draw $\sigma_a^{(i+1)}$ by using the observations $a_{ft} - \theta_f/(1+c_t)$ (where c_t is shorthand for c if fiduciary duty is relevant for that transaction) to update the standard deviation of a normal with mean 0.

3. *Metropolis Step for c.* Draw $c' \sim Q_c\left(\cdot|c^{(i)}\right)$. Since this changes expected profits conditional on entry, we first check whether all entry decisions are consistent with observed ones given this new draw of the strength of the advice channel. If so, we compute $L(c) \equiv \psi_c(c)\prod_f L_f\left(\theta_f^{(i+1)};c\right)$, where $\psi_c(\cdot)$ is the prior on c, and update $c^{(i+1)}$ to c' with probability $\min\left[L(c')/L(c^{(i)}),1\right]$.

4. *Metropolis Step for (γ,α).* Draw $(\gamma',\alpha') \sim Q(\cdot|\gamma,\alpha)$ and thus new beliefs over the competitive effect $\mathbb{E}[f^T(\cdot,\cdot)]$. If all entry conditions are satisifed with this draw, we update $(\gamma^{(i+1)},\alpha^{(i+1)})$ to this new draw with probability $\min\left[\psi_{\gamma,\alpha}(\gamma',\alpha')/\psi_{\gamma,\alpha}(\gamma^{(i)},\alpha^{(i)}),1\right]$.

5. *Gibbs Sampler for the Probit.* Given the draws of θ_f and (c, γ, α), we can update all other parameters in (2) via the Gibbs sampler for the probit in McCulloch and Rossi (1994). This involves a Bayesian OLS of profits on covariates with the current draw of $\epsilon^{(i)}$ followed by a new draw of $\epsilon^{(i+1)}$. We increment i and then loop to Step 1.

We place diffuse normal inverse-gamma priors on all quantities that are updated by Bayesian OLS, to take advantage of conjugacy. We place "uninformative" uniform priors on (c,γ,α) so that $\psi_\times(\cdot)$ is cancelled from all expressions. While sufficiently long chains are guaranteed to sample from the posterior distribution, convergence may be slow, especially if the chain is not initalized well. To guard against this concern, we run a Sequential Monte Carlo proposed by del Moral et al. (2006) and used by Chen et al. (2019). This runs a large number of chains in parallel and mixes them based on their likelihood—allowing for exploration while eventually killing off chains that are in suboptimal regions. Details are in Appendix E.

6.4. Parameter Estimates

Table 7 shows results from this estimation procedure.[33] Panel (A) shows the latent distribution of types. The mean type for BDs corresponds to a distortion of approximately 3.3, and the distribution has standard deviation 0.54. This number can be compared to the parameters of the profit function

[32] See Section 2.8 of Rossi et al. (2005) for details.
[33] See Appendix B.6 for a discussion of model fit.

Table 7: Parameter estimates from structural model

Parameter	Mean	Standard Error	95% Credible Interval
(A) Distribution of Types			
Mean Type for BD (μ_θ^{BD})	3.257	0.045	[3.169, 3.345]
Mean Type for RIA (μ_θ^{RIA})	3.935	0.042	[3.853, 4.019]
Standard Deviation of Types (σ_θ)	0.537	0.018	[0.502, 0.574]
(B) Profit Function			
Minimizing Type for BD ($\bar{\theta}^{BD}$)	2.098	0.060	[1.981, 2.216]
Minimizing Type for RIA ($\bar{\theta}^{RIA}$)	2.706	0.014	[2.680, 2.735]
(C) Competition			
Strength of Competition (γ)	0.0965	0.001	[0.0949, 0.0987]
Cross-Type Competition (α)	0.292	0.004	[0.284, 0.298]
(D) Advice Channel			
Effect on Advice (c)	0.0312	0.0062	[0.020, 0.0451]
Effect on Profits ($\lambda \times 10$)	0.0402	0.0337	[0.0013, 0.1247]
(E) Fixed Cost Channel			
Effect for Local Firms (κ_{2L})	0.144	0.105	[0.006, 0.393]
Effect for Regional Firms (κ_{2R})	0.078	0.060	[0.003, 0.223]
Effect for National Firms (κ_{2N})	0.060	0.041	[0.003, 0.155]

shown in panel (B). For the estimated parameters, $\bar{\theta}$ corresponds to a profit minimizing type since $\delta_1^T < 0$. Given $\bar{\theta}^{BD}$ is slightly more than two standard deviations lower than the mean type for BDs, we can think of the profit function as increasing over the relevant range of types, as in panel (a) of Figure 4. The solid black line in Figure 6 illustrates this profit function for a representative market at the observed level of competition, with zero normalized to be the fixed cost of entry without fiduciary duty.[34]

Panel (C) shows the effect of competition. We estimate γ to be especially low: the point estimate suggests that doubling the number of competitors corresponds to a reduction in variable profits by about 7%. We also find a low value of α, suggesting that cross-type competition is not too strong. This is consistent with the observation that the number of RIAs does not respond noticeably to fiduciary duty, which is a quasi-random policy change that would affect the number of BDs. Given the parameterization of the fixed costs, this policy change is encapsulated in the model and informs the estimate of α.

Panel (D) shows the parameters most relevant for the advice channel. We estimate c to be 0.031, suggesting that the advice channel leads to reductions in distortion by about 3%. The estimate of $\lambda = 0.0097$ governs the effect this distortion has on profits. To put this number into perspective, we can compare it to the numbers in panel (E), which show the parameters related to the fixed cost channel. The increase in fixed cost for local BDs (κ_{2L}) due to fiduciary duty is slightly more than twice the reduction in profits due to an advice channel for a firm whose optimal distortion θ is 4.

[34]Recall that the normalization for profits is that the variance of the firm-level idiosyncratic shock to entry is 1.

Figure 6: Profit function and fixed costs for BDs in a typical market

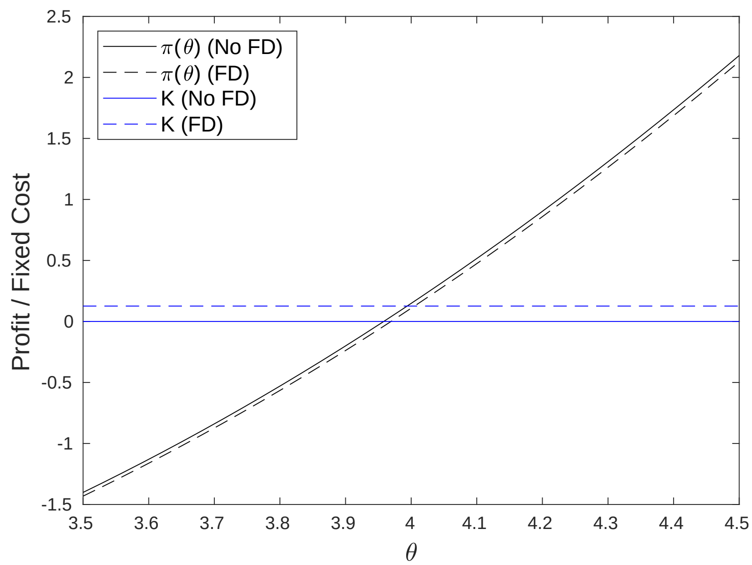

Regional and national firms have smaller increases in fixed costs: local firms bear the brunt of the fixed cost channel of fiduciary duty.

Figure 6 illustrates the advice and fixed costs channels in the representative market as well. Upon imposition of the advice channel, the profit function drops to the dashed black line. While this effect is in fact larger for higher θ, the graph illustrates that the effect on profits is smaller (across values of θ) than the average effect of the fixed cost channel—illustrated by moving from the solid blue line to the dashed blue line. Given the slope of the profit function, an increase in this fixed cost actually harms lower-distortion firms, and it would by itself lead to an increase in distortion. That is, concerns about fiduciary duty driving out high-quality advisers from the market through an increase in fixed costs are warranted. However, the advice channel outweighs this effect so that on net fiduciary duty improves observed advice. We quantify these underlying forces in the subsequent section, and then discuss their role under counterfactual stringency levels.

6.5. Quantifying the Advice Channel and Fixed Cost Channels

With the estimated parameters in hand, we can compute the counterfactual effect of fiduciary duty had it simply operated through an advice channel and compare how it relates to the total effect of fiduciary duty. We loop through all markets and compute the expected distortion $\mathbb{E}[a_m^N]$ provided by a BD conditional on entry—using that market's covariates and potential entrant distribution—assuming there is no fiduciary duty. That is, we set $\lambda = c = 0$ and $\kappa_{2\times} = 0$ as well as all interactions with regional and national dummies to 0 but keep all other parameters at their estimated values. We then allow fiduciary duty to operate only through the advice channel, setting λ and c to their

Figure 7: Decomposition of the total effect of fiduciary duty

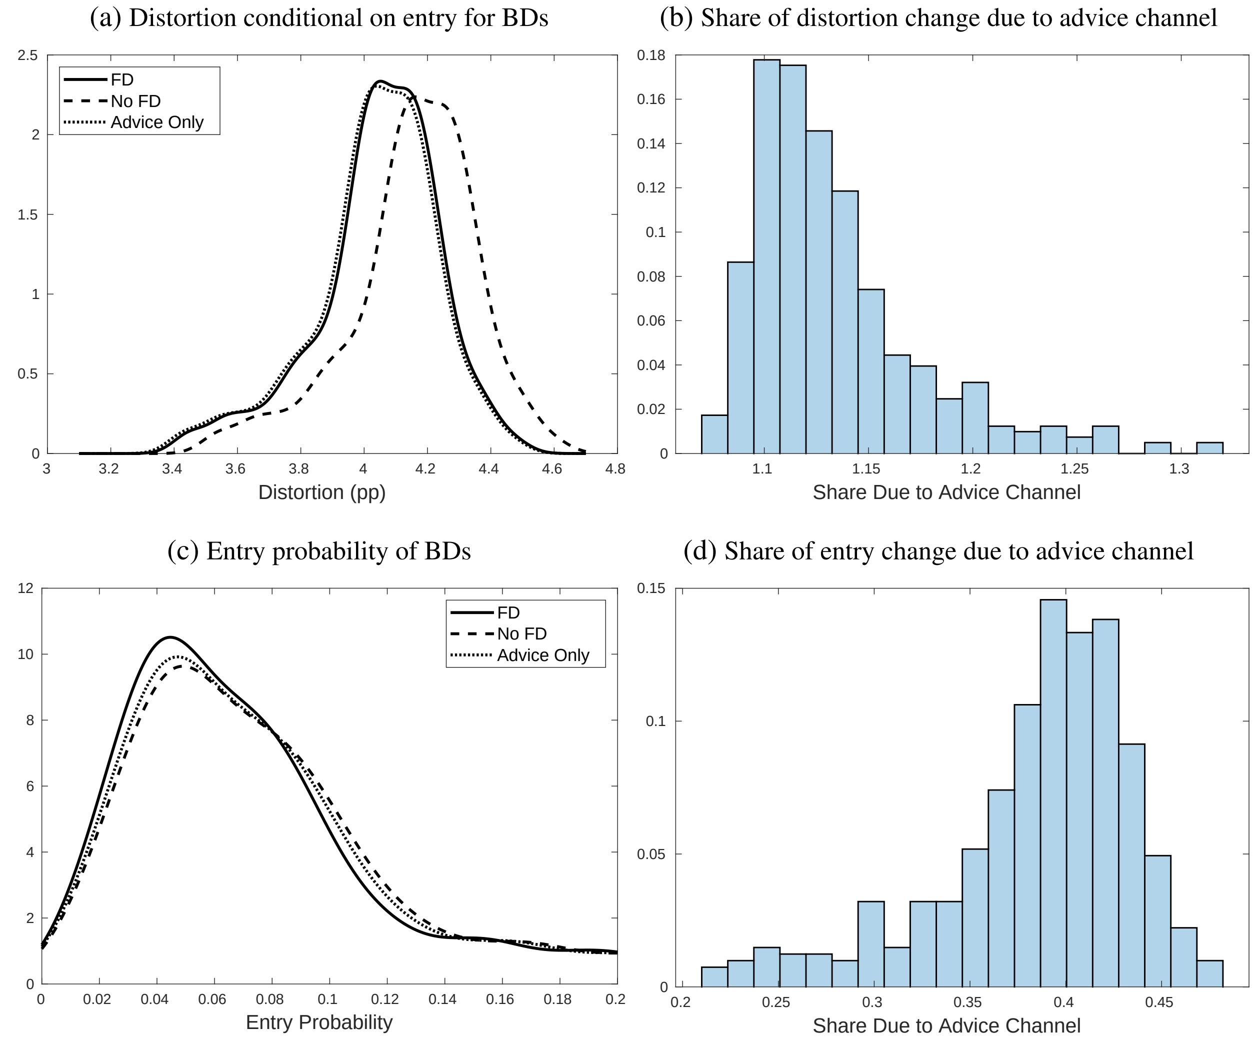

estimated values but keeping $\kappa_{2\times}$ at 0 and computing expected advice. This gives us the expected distortion $\mathbb{E}[a_m^A]$. Finally, we set λ, c, and $\kappa_{2\times}$ to their estimated values, simulating the total effect of fiduciary duty $\mathbb{E}[a_m^{FD}]$. We do this exercise for 100 draws of parameters selected uniformly at random from the chain post burn-in, and all outputs are averaged within market across draws.

Figure 7(a) shows the distribution of expected distortion conditional on entry for BDs for the three cases: no fiduciary duty, simply the advice channel, and full fiduciary duty. Fiduciary duty improves advice, seen as a leftward shift in the distribution in this normalization. Notably, however, the distribution induced by simply the advice channel (dotted line) is rather close to the distribution with full fiduciary duty—and in fact slightly farther to the left of the distribution with full fiduciary duty. Distortion is in general lower under the pure advice channel than under the full effect of fiduciary duty. This observation ties back to the parameter estimates in Section 6.4: since the profit function increases in distortion, the increase in fixed costs from the fixed cost channel harms lower-distortion firms. Panel (b) computes the number at the heart of this quantification. For each market, we compute the share of the effect of fiduciary duty due to the advice channel: $(\mathbb{E}[a_m^A]-\mathbb{E}[a_m^N])/(\mathbb{E}[a_m^{FD}]-\mathbb{E}[a_m^N])$. As before, we average within market across draws of parameter estimates andand arrive at the distribution in Figure 7(b). The average of this distribution is 1.13, suggesting that were fiduciary duty to act purely through the advice channel, the effect on distortion would be 13% larger than with both the advice and fixed cost channels.

Panels (c) and (d) repeat this exercise with the entry probability of BDs (averaged across local, regional, and national ones) as the outcome. Panel (c) shows that imposing fiduciary duty reduces entry probabilities. However, unlike with distortion, the distribution when simply imposing the advice channel is between the ones with and without fiduciary duty: the advice channel reduces entry, and the fixed cost channel reduces it further. Panel (d) shows that the advice channel accounts for between 1/4 and 1/2 of the total effect of fiduciary duty on entry, with a mean of 0.39.

The takeaway from this decomposition is that both proponents and detractors of expanding fiduciary duty to all financial advisers have identified empirically relevant mechanisms through which fiduciary duty affects advice. Like proponents argue, fiduciary duty does increase the cost of offering distorted advice, so that extending it would improve advice quality. However, like detractors argue, fiduciary duty also increases the fixed cost of offering advice, and this fixed cost increase drives out high-quality advisers from the market. On net, the advice mechanism is stronger than the fixed cost mechanism, so that fiduciary duty improves advice quality.

6.6. Changing the Stringency of Fiduciary Duty

Does increasing stringency continue to improve advice? The previous results leverage variation across common-law fiduciary standards, but a federal standard may be significantly stronger. That is, it could be the case that expanding fiduciary duty at the federal level yields a larger increase in the cost of offering distorted advice and in fixed costs than what we observe in our data. This could shift entry and it is therefore important to understand how market equilibria would shift under more stringent regimes.

Given the parameter estimates, it is not a foregone conclusion that increasing the stringency of fiduciary standards would continue to improve advice. In particular, if we conceptualize an increase in stringency as an increase in both the advice and the fixed cost channels, then we have two competing effects. While the advice channel induces a reduction in distortion, both the advice channel and the fixed cost channels induce exit.[35] Given the shape of the profit function shown in Figure 6, would lead to exit of low-distortion firms. Thus, the exit effect could counteract any improvement in advice, and the model provides a way to quantify this tradeoff.

[35] The advice channel has a larger effect on profits for high-distortion firms, although this dimension of heterogeneity is estimated to be small.

Figure 8: Effect of stringency

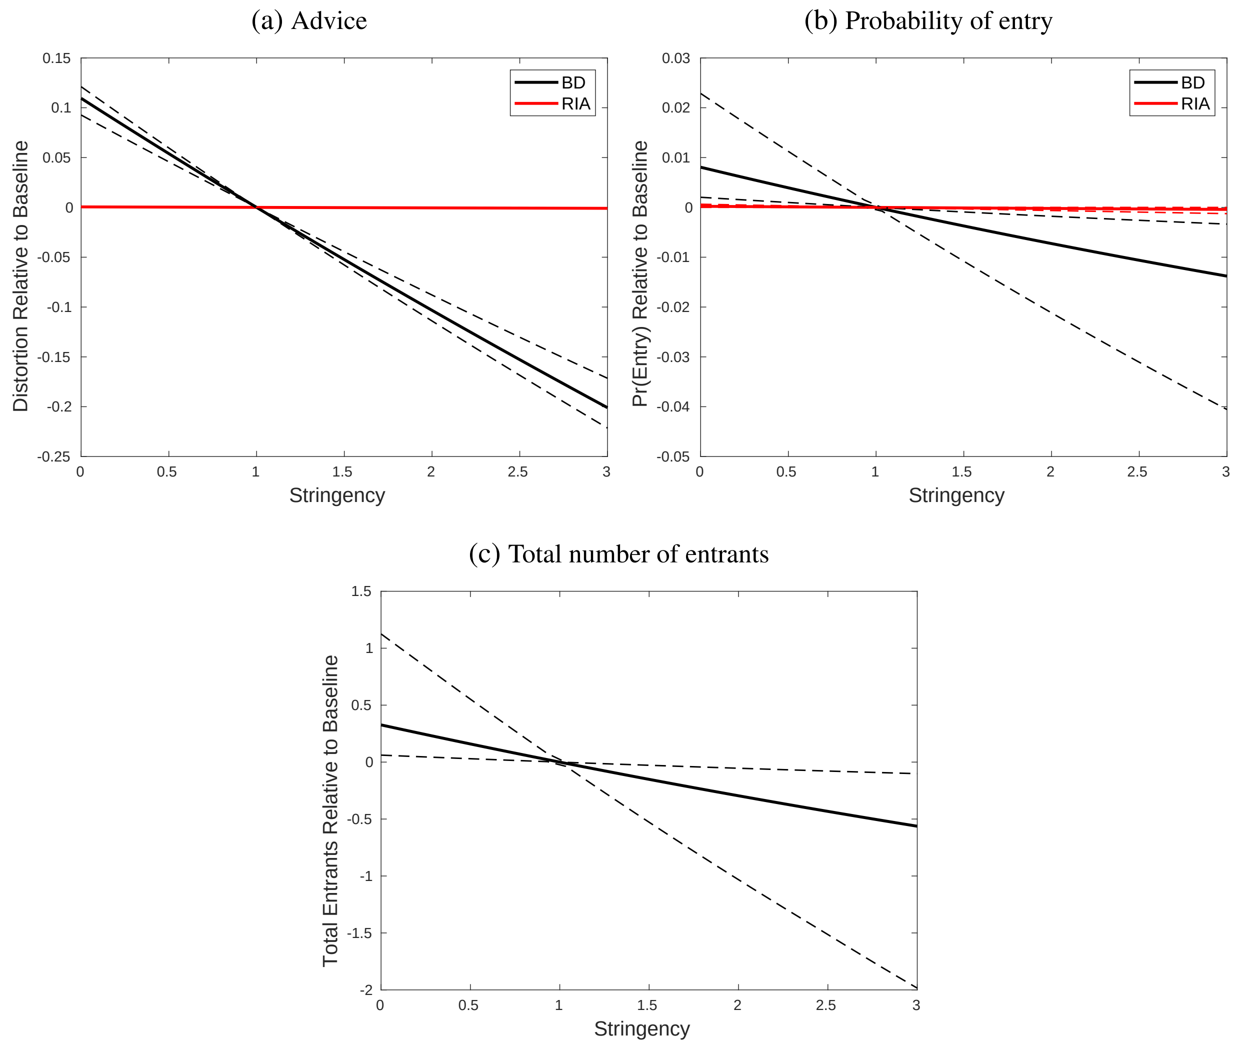

Effect of stringency on advice, probability of entry, and total number of entrants, relative a baseline that corresponds to stringency of 1. The solid lines show the mean effect across markets and the dotted lines show the middle 90% of effects across markets. Larger values of advice correspond to more distorted advice.

To simulate counterfactual stringencies, we keep the ratio of the advice and fixed cost channels fixed. We scale c and $\kappa_{2\times}$ by k, for $k \in [0, 3]$.[36] Note that $k = 0$ corresponds to no fiduciary duty and $k = 1$ corresponds to the common law standard. One may expect that national fiduciary standards enforced by the SEC could correspond to $k > 1$.

We draw 100 parameter estimates at random from the chains. For each market and parameter estimate, we simulate outcomes for a wide range of k. We then subtract the value at $k = 1$ to show changes relative to the baseline common law fiduciary standard—so that the value is mechanically 0 at $k = 1$—and average the results within-market. Figure 8 shows the results of these computations.

[36]In doing so, we multiply λ by $[kc/(1 + kc)]/[c/(1 + c)]$ as well. In principle, our methodology corresponds to scaling $\tilde{c}$ by k, using the original notation in (3).

Panel (a) shows that as fiduciary standards become more stringent, advice keeps improving. Recall that this result is not mechanical, as the set of entrants varies with stringency. At $k = 3$, advice improves over baseline by about 20 bp on average, off a baseline distortion at $k = 1$ of 4.05.[37] For comparison, moving from common-law duty at $k = 1$ to the stringent standard at $k = 3$ without changing entry at all improves advice by 23.2 bp, so that not accounting for endogenous entry in the counterfactual simulations overpredicts the improvement in advice by about 16%. The dashed lines in the figure show the heterogeneity in this effect across markets in our sample, and they indicate that the qualitative result that stringency improves advice holds market-by-market. While we hesitate to draw conclusions about markets outside our sample, to the extent that firms have similar profit functions in such markets, we may expect that stringent fiduciary standards would improve advice across the US.

We show the effect of advice on RIAs in red: fiduciary duty on BDs may have a spillover competitive effect on RIAs, thus adjusting the advice they provide. In panel (a), we see no evidence that advice by RIAs is affected appreciably. These results are consistent with low estimates of γ and α in the model and also with the reduced-form results from Section 4 that behavior of RIAs is not appreciably different across the border.

A natural concern is that this is accompanied by a reduction in the number of entering firms. Panel (b) repeats this exercise for the probability of entry. We do see a noticeable decrease in the probability of entry: at $k = 3$, the probability of entry decreases by about 1.4 pp for BDs relative to a baseline of about 9 pp at $k = 1$. That is, advice continues to improve despite the reduction in entry and thus the shift towards higher-θ firms: that the net effect still improves advice is an empirical result. The effect on RIAs is again minimal: we see a very small increase in the number of RIAs operating in the market (since γ and α, while small, are nevertheless positive), but the quantitative magnitude is economically insignificant. Panel (c) puts these observations together and reports the effect on the total number of entrants, off a baseline of 13.1 firms at $k = 1$. Since RIA firms are more common, and the empirical estimates suggest they are not affected much by the imposition of fiduciary duty, even the effect of tripling stringency on total number of firms in the market is a drop is 0.5 firms, which is somewhat small relative to the baseline.

This exercise indicates that further increases to stringency of fiduciary duty beyond common law, such as a federal standard, will continue to have a significant positive impact on advice provided by BDs. While we do predict the negative impact on the entry of BDs that detractors highlight, the

[37] In this analysis we maintain the assumption that the relative contribution of the two channels will be the same in any fiduciary duty regulation, but the magnitude will be influenced by stringency. One may wonder, however, how robust this result of improved advice is with respect to this maintained assumption. To investigate, we scale the advice channel by $k = 3$ and scale the fixed cost channel by as much as $k = 10$, thus instead assuming that stringent fiduciary duty has a much stronger effect on the fixed cost channel. As expected, this does increase distortion relative to the equal scaling case. However, even at $k = 10$ it does so by less than 10 bp, suggesting that at least half the effect of increased stringency on advice survives this robustness check.

total effect on entry is somewhat small in comparison to the entire market for financial advice. This may alleviate concerns that clients may not longer have access to financial products, and together the results suggest that strengthening fiduciary standards may continue to benefit clients.

7. Conclusion

This paper evaluates the effects of extending fiduciary duty to broker-dealers on returns, market structure, and observable characteristics of the the set of products consumers purchase. This question is motivated by recent regulatory discussion around expanding fiduciary duty to all broker-dealers. Supporters of the expansion argue that imposing fiduciary duty on all advisers will alleviate the conflict of interest and ensure that retirees choose products that are better suited to their needs. Opponents argue that fiduciary duty does not have a noticeable impact on product choice—perhaps because competition already disciplines financial advisers or perhaps because the conflict-of-interest was overblown to begin with—but will instead simply increase the cost of doing business, which will lead to fewer advisers in the market and fewer retirees purchasing beneficial products.

We evaluate these claims empirically by leveraging transactions-level data from a major financial services provider and a comprehensive dataset on the set of practicing financial advisers. We find that in the market for annuities, fiduciary duty increases risk-adjusted returns by 25 bp and induces an reduction of 16% in the number of BD firms without a change in the total sales of annuities. Unpacking this change in risk-adjusted returns we find that BDs with fiduciary duty are less likely to sell variable annuities; when selling a variable annuity, they are more likely to steer clients towards products with more and higher-quality investment options. These results offer a extensive picture of the different effects of fiduciary duty in the market for financial advice.

These results on the mean causal impact of fiduciary duty present credible reduced-form evidence that common-law fiduciary duty improves financial advice in the markets under study. However, they are silent about its effects in other markets or about the effects of stronger fiduciary standards. This last point is especially important, as federal fiduciary standards may be significantly more stringent than those imposed by common law. We show that to understand how fiduciary duty would operate in these counterfactual settings, one needs to unpack how much of its effects operate by increasing fixed costs and how much of its effects operate by constraining low-quality advice. We then develop a model of firms entering a market and selecting their advice that identifies properties of the distribution of advice that allow us to unpack these mechanisms. We find evidence in favor of the presence of a constraint on low-quality advice; that is, fiduciary duty does not simply increase fixed costs. Moreover, taking the model to the data, we find that not only is the advice channel present, but it is also an especially dominant force underlying the observed effect. Even though fiduciary duty increases fixed costs and drives out high quality advisers from the market, as detractors of extending fiduciary duty argue, its effect on low-quality advice more than compensates. The counterfactual analysis guided by our model suggests that implementing a federal standard that is more stringent than common law fiduciary duty would deliver increased returns for retirees.

References

ALLEN, J., R. CLARK, AND J.-F. HOUDE (2014): "The Effect of Mergers in Search Markets: Evidence from the Canadian Mortgage Industry," *American Economic Review*, 104, 3365–96.

——— (2019): "Search Frictions and Market Power in Negotiated-Price Markets," *Journal of Political Economy*, 127, 1550–1598.

ANAGOL, S., S. COLE, AND S. SANKAR (2017): "Understanding the Advice of Commissions-Motivated Agents: Evidence from the Indian Life Insurance Market," *Review of Economics and Statistics*, 99, 1–15.

BACHAS, N. (2019): "The Impact of Risk-Based Pricing and Refinancing on the Student Loan Market," Working Paper.

BACHAS, N. AND E. LIU (2019): "Market Power in Small Business Lending: A Two Dimensional Bunching Approach," Working Paper.

BENARTZI, S. AND R. H. THALER (2001): "Naive Diversification Strategies in Defined Contribution Saving Plans," *American Economic Review*, 91, 79–98.

BENETTON, M. (2019): "Leverage Regulation and Market Structure: A Structural Model of the UK Mortgage Market," Working Paper.

BERGSTRESSER, D., J. M. R. CHALMERS, AND P. TUFANO (2009): "Assessing the Costs and Benefits of Brokers in the Mutual Fund Industry," *Review of Financial Studies*, 22, 4129–4156.

BERNARD, T. S. (2019): "S.E.C. Adopts New Broker Rules That Consumer Advocates Say Are Toothless," *New York Times*.

BERRY, S. T. (1992): "Estimation of a Model of Entry in the Airline Industry," *Econometrica*, 889–917.

BRANCACCIO, G., D. LI, AND N. SCHÜROFF (2020): "Learning by Trading: The Case of the US Market for Municipal Bonds," Working Paper.

BULOW, J. I., J. D. GEANAKOPLOS, AND P. D. KLEMPERER (1985): "Multimarket Oligopoly: Strategic Substitutes and Complements," *Journal of Political Economy*, 93, 488–511.

CÉLÉRIER, C. AND B. VALLÉE (2017): "Catering to Investors Through Security Design: Headline Rate and Complexity," *Quarterly Journal of Economics*, 132, 1469–1508.

CHAROENWONG, B., A. KWAN, AND T. UMAR (2019): "Does Regulatory Jurisdiction Affect the Quality of Investment-Adviser Regulation?" *American Economic Review*, 109, 3681–3712.

CHEN, X., T. CHRISTENSEN, AND E. TAMER (2019): "Monte Carlo Confidence Sets for Identified Sets," *Econometrica*, 86, 1965–2018.

CHRISTOFFERSEN, S. E. K., R. EVANS, AND D. K. MUSTO (2013): "What Do Consumers' Fund Flows Maximize? Evidence from Their Brokers' Incentives," *Journal of Finance*, 68, 201–235.

COCHRANE, J. H. (2009): *Asset Pricing*, Princeton University Press.

CUESTA, I. AND A. SEPÚLVEDA (2019): "Price Regulation in Credit Markets: A Trade-off between Consumer Protection and Credit Access," Working Paper.

DEL GUERCIO, D. AND J. REUTER (2014): "Mutual Fund Performance and the Incentive to Generate Alpha," *Journal of Finance*, 69, 1673–1704.

DEL MORAL, P., A. DOUCET, AND A. JASRA (2006): "Sequential Monte Carlo Samplers," *Journal of the Royal Statistical Society B*, 68, 411–436.

DULLECK, U. AND R. KERSCHBAMER (2006): "On Doctors, Mechanics, and Computer Specialists: The Economics of Credence Goods," *Journal of Economic Literature*, 44, 5–42.

EGAN, M. (2019): "Brokers vs. Retail Investors: Conflicting Interests and Dominated Products," *Journal of Finance*, 74, 1217–60.

EGAN, M., G. MATVOS, AND A. SERU (2019): "The Market for Financial Adviser Misconduct," *Journal of Political Economy*, 127, 233–295.

EINAV, L., M. JENKINS, AND J. LEVIN (2012): "Contract Pricing in Consumer Credit Markets," *Econometrica*, 80, 1387–1432.

FINKE, M. AND T. LANGDON (2012): "The Impact of the Broker-Dealer Fiduciary Standard on Financial Advice," Working Paper.

GARRETT, D. (2019): "Conflicts of Interest in Municipal Bond Advising and Underwriting," Working Paper.

GAVAZZA, A. AND M. GALENIANOS (2020): "Regulatory Interventions in Consumer Financial Markets: The Case of Credit Cards," Working Paper.

GRIGSBY, J., A. HORTACSU, A. SERU, AND V. YAO (2019): "Searching for Approvals," Working Paper.

GRUNEWALD, A., J. LANNING, D. LOW, AND T. SALZ (2019): "Auto Dealer Loan Intermediation: Consumer Behavior and Competitive Effects," Working Paper.

GUISO, L., A. POZZI, A. TSOY, L. GAMBACORTA, AND P. E. MISTRULLI (2018): "The Cost of Steering in Financial Markets: Evidence from the Mortgage Market," *Working Paper*.

HASTINGS, J., A. HORTASU, AND C. SYVERSON (2017): "Sales Force and Competition in Financial Product Markets: The Case of Mexico's Social Security Privatization," *Econometrica*, 85, 1723–1761.

HUNG, A. A., N. CLANCY, J. DOMINITZ, E. TALLEY, C. BERREBI, AND F. SUVANKULOV (2008): "Investor and Industry Perspectives on Investment Advisers and Broker-Dealers," .

ILLANES, G. (2017): "Switching Costs in Pension Plan Choice," Working Paper.

ILLANES, G. AND M. PADI (2019): "Retirement Policy and Annuity Market Equilibria: Evidence from Chile," Working Paper.

INDERST, R. AND M. OTTAVIANI (2012a): "Competition Through Commissions and Kickbacks," *American Economic Review*, 102, 780–809.

——— (2012b): "Financial Advice," *Journal of Economic Literature*, 50, 494–512.

JIA, P. (2008): "What Happens When Wal-Mart Comes to Town: An Empirical Analysis of the Discount Retailing Industry," *Econometrica*, 76, 1263–1316.

KELLY, B. (2017): "DOL fiduciary rule compliance costs exceed 4.7 billion: SIFMA study," *InvestmentNews*.

KOIJEN, R. AND M. YOGO (2018): "The Fragility of Market Risk Insurance," Working Paper.

KOIJEN, R. S. J. AND M. YOGO (2015): "The Cost of Financial Frictions for Life Insurers," *American Economic Review*, 105, 445–75.

——— (2016): "Shadow Insurance," *Econometrica*, 84, 1265–1287.

KOZORA, M. L. (2013): "The Effect of Regulatory Regimes on the Provision of Retail Investment Advice," *Working Paper*.

LIBERMAN, A., C. NEILSON, L. OPAZO, AND S. ZIMMERMAN (2019): "The Equilibrium Eects of Information Deletion: Evidence from Consumer Credit Markets," Working Paper.

MARSH, A. (2019): "SEC guts RIA industry with a footnote, degrading fiduciary duty," *Financial Planning*.

MAZZEO, M. J. (2002): "Product Choice and Oligopoly Market Structure," *RAND Journal of Economics*, 221–242.

MCCULLOCH, R. AND P. E. ROSSI (1994): "An Exact Likelihood Analysis of the Multinomial Probit," *Journal of Econometrics*, 64, 207–240.

MULLAINATHAN, S., M. NOETH, AND A. SCHOAR (2012): "The Market for Financial Advice: An Audit Study," Working Paper.

NELSON, S. (2020): "Private Information and Price Regulation in the US Credit Card Market," Working Paper.

ROBERTS, J. W. AND A. SWEETING (2013): "When Should Sellers Use Auctions?" *American Economic Review*, 103, 1830–61.

ROBLES-GARCIA, C. (2020): "Competition and Incentives in Mortgage Markets: The Role of Brokers," Working Paper.

ROSSI, P. E., G. M. ALLENBY, AND R. MCCULLOCH (2005): *Bayesian Statistics and Marketing*, Wiley.

SEC (2011): "Study on Investment Advisers and Broker-Dealers," Tech. rep., US Securities and Exchange Commission.

——— (2013a): "Duties of Brokers, Dealers, and Investment Advisers," Tech. rep., US Securities and Exchange Commission.

——— (2013b): "Study Regarding Financial Literacy Among Investors," Tech. rep., US Securities and Exchange Commission.

SEIM, K. (2006): "An Empirical Model of Firm Entry with Endogenous Product-Type Choices," *RAND Journal of Economics*, 619–640.

SWEETING, A. (2009): "The Strategic Timing Incentives of Commercial Radio Stations: An Empirical Analysis Using Multiple Equilibria," *RAND Journal of Economics*, 40, 710–742.

XIN, Y. (2020): "Asymmetric Information, Reputation, and Welfare in Online Credit Markets," Working Paper.

Supplemental Material for Fiduciary Duty and the Market for Financial Advice

VIVEK BHATTACHARYA[1], GASTÓN ILLANES[1], and MANISHA PADI[2]

[1]Department of Economics, Northwestern University and NBER
[2]University of California, Berkeley, School of Law

May 19, 2020

A. Further Analysis of the Model

In this appendix, we formalize the statements in Section 5, provide proofs of the propositions presented in that section, and provide further results.

A.1. Only Broker-Dealers

A.1.1. Settings with a Single Category

First consider a simple version of the model in Section 5.1, setting $M = 1$. There is a continuous distribution of types $\theta_j \sim H(\cdot)$ on compact support. Each type has a base profit function $\pi(a - g(\mu); \theta)$ maximized at $a = \theta$, and we define $\pi^*(\theta) \equiv \max_a \pi(a - g(\mu); \theta)$. Note that since the effect of competition is modeled as shifting the optimal advice, $\pi^*(\cdot)$ does not depend on μ. The actual profit a type-θ firm earns upon entering is $f(\mu) \cdot \pi^*(\theta) - K$, where K is the entry cost and $f(\cdot)$ is a strictly decreasing function of the mass μ of entrants capturing competitive effects. While we do not place much structure on π in general, suppose that $H(\cdot)$ and $\pi(\cdot)$ are jointly such that the distribution of $\pi^*(\theta)$ does not have any mass points; in the following, we will essentially consider the distribution of $\pi^*(\theta)$.

While the ordering of θ has an interpretation in Section 5.1, we strip it of its interpretation as the quality of advice in this appendix. Instead, relabel and rescale types $\tilde{\theta}$ be to be one-to-one with base profits $\pi^*(\theta)$ so that $\tilde{\theta}' > \tilde{\theta}$ if and only if $\tilde{\theta}'$ earns lower profits $\tilde{\pi}(\theta')$ than does $\tilde{\theta}$. Moreover, rescale types so that they are uniform on the unit interval. Let $\tilde{\Theta} : \theta \mapsto \tilde{\theta}$ be this function. Then, an equilibrium is such that $f(\mu) \cdot \tilde{\pi}(\mu) = K$, where μ is the marginal type who enters, as long as $\mu \in (0, 1)$. If $f(0) \cdot \tilde{\pi}(0) < K$ then no one enters, and if $f(1) \cdot \tilde{\pi}(1) > K$ then everyone enters.

CONTACT. vivek.bhattacharya@northwestern.edu, gaston.illanes@northwestern.edu, mpadi@berkeley.edu.

Lemma 1. *There is a unique equilibrium.*

Proof. Note that $f(\mu)\cdot\tilde{\pi}(\mu)$ is strictly decreasing in μ. Thus, either $f(0)\cdot\tilde{\pi}(0) < K$ or $f(1)\cdot\tilde{\pi}(1) > K$, or it can take on a value of K at most once in $(0, 1)$. □

Lemma 2. *The set of types θ_j who enter at an entry cost of $K' > K$ is a subset of the set of types who enter at an entry cost of K.*

Proof. Let $\mu^*(K)$ be such that $f(\mu^*(K)) \cdot \tilde{\pi}(\mu^*(K)) = K$. It is easy to see that $\mu^*(\cdot)$ is decreasing in its argument. The set of types who enters is simply $\tilde{\Theta}^{-1}([0, \mu^*(K)])$, where $\tilde{\Theta}^{-1}(\cdot)$ is the inverse map of the function above. Thus, the set of types who enters under K' is the image of a smaller set, which means it is a subset of those who enter under K. □

Lemma 2 shows that the nonprimitive condition in Propositions 1 and 2 is indeed an implication of $K' > K$ for $M = 1$ type. This result verifies that the conditions in these propositions are not mutually inconsistent: it is a potential implication at least in certain cases, and one that is testable.

A.1.2. Multiple Categories of Broker-Dealers

The model in Section 5.1 allows for $M > 1$ categories. A natural concern is that even if fiduciary duty operates through a pure fixed cost channel, national BDs might experience a smaller increase in fixed cost than local BDs. We might imagine that $K'_{\text{local}} - K_{\text{local}} > K'_{\text{national}} - K_{\text{national}}$, and we may also expect these categories have different profit functions.

In this situation, it is *not* necessarily true that the advice observed in the market without fiduciary duty is a superset of advice observed with. One can construct an example in which $K'_1 > K_1$, $K'_2 = K_2$, and the support of the advice provided by Category 2 firms is strictly to the right of the support of that provided by Category 1—in the absence of fiduciary duty. Under reasonable conditions on $f(\cdot)$ (such as the ones in Appendix A.2), fiduciary duty will lead to a decrease in the number of Category 1 firms in the market and an increase in the Category 2 firms. Then, the advice under fiduciary duty will not be a subset of that without.[1] By itself, this possibility poses a difficulty for the testable restrictions discussed in Section 5, as expansion of advice could *still* be possible under a pure fixed cost channel with heterogeneous changes in fixed cost. However, note that this example required an expansion of the number of Category 2 broker-dealers. Indeed, (the contrapositive of) Proposition 2 is a general requirement for us to see an expansion of advice upon imposing of fiduciary duty, in a pure fixed cost channel. Here we simply provide proofs of the argument in Section 5.2.

[1] One can essentially go through Appendix A.2 and label the broker-dealers as "local broker-dealers" and the investment advisers as "national broker-dealers."

Proof of Proposition 1. First, just as in Appendix A.1.1, let $\pi^*_m(\theta) \equiv \max_a \pi_m(a - g(\boldsymbol{\mu}); \theta)$. This does not depend on $\boldsymbol{\mu}$ since it just shifts the optimal advice. Let μ_m denote the equilibrium mass of type-m firms in a world with fixed costs K_m, and let μ'_m denote this mass in a world with fixed costs $K'_m < K_m$. Suppose $\mu'_m < \mu_m$. Then, a Category m firm with type θ enters at K'_m if $f_m(\boldsymbol{\mu}') \cdot \pi^*_m(\theta) \geq K'_m$, or $\pi^*_m(\theta) \geq K'_m / f_m(\boldsymbol{\mu}')$. Similarly, (θ, m) enters with costs K_m if $\pi^*_m(\theta) \geq K_m / f_m(\boldsymbol{\mu})$. Since $\mu'_m < \mu_m$, it must be that $K'_m / f_m(\boldsymbol{\mu}') > K_m / f_m(\boldsymbol{\mu})$, meaning if a type θ firm enters with fiduciary duty, it must enter without fiduciary duty as well. □

Proof of Proposition 2. If $g(\boldsymbol{\mu}) = 0$, then $a^*(\theta) = \theta$; i.e., the advice offered by a type θ firm upon entry is θ. Proposition 1 shows $\mu'_m < \mu_m$ implies that the set of all firms who enter contracts. Thus, the set of advice offered by these entrants must shrink as well. This means the highest a observed in the market decreases, and the lowest a in the market increases.

If $g(\boldsymbol{\mu}) \neq 0$, then $a^*(\theta; \mu)$ depends on $\boldsymbol{\mu}$. If $g(\boldsymbol{\mu})$ is increasing in all its arguments, then $a^*(\theta; \boldsymbol{\mu})$ increases upon an increase in fixed costs. Proposition 1 still implies the minimum θ among all entrants would increase. The advice a that this entrant would provide would also increase. Thus, the lowest quality advice would worsen as $K'_m > K_m$ if $\mu'_m < \mu_m$. □

To reiterate Proposition 2, we can reject a pure fixed cost channel with potential heterogeneity in the impact on fixed costs if we observe a *decrease* in the mass of a particular type of broker-dealers with a corresponding *introduction* of previously unseen advice.

A.1.3. Extending the Type

Note that these arguments just depend on the fact that there is a unidimensional ordering of types in terms of their base profits, and the base profits are the only component of these types that matter for who enters. Moreover, an increase in fixed costs of entry does not impact the ordering of these base profits; i.e., if $\pi^*_m(\theta_1; \boldsymbol{\mu}) < \pi^*_m(\theta_2; \boldsymbol{\mu})$ when entry costs are $\boldsymbol{K}$, then $\pi^*_m(\theta_1; \boldsymbol{\mu}) < \pi^*_m(\theta_2; \boldsymbol{\mu})$ when entry costs are $\boldsymbol{K}'$ as well. We show below that some natural extensions satisfy these conditions.

Idiosyncratic Entry Costs. Suppose that each potential entrant is now categorized by an ordered pair (θ_j, ϵ_j) and a category m, where $\epsilon_j \sim G(\cdot|\theta_j)$. A firm of type (θ_j, ϵ_j) has a base profit function $\pi_m(a; \theta_j) + \epsilon_j$. This extension allows firms who would offer the same profit conditional on entry to be differentially profitable. As before, let $\mathcal{E}^*_m(\boldsymbol{K})$ denote the set of types of category m who would enter with a fixed cost of $\boldsymbol{K}$. Thus, if we define

$$\underline{\theta}(K) \equiv \min \{\theta : \text{ there exists } m \text{ and } \epsilon \in \operatorname{supp} G(\cdot|\theta) \text{ such that } (\theta, \epsilon) \in \mathcal{E}^*_m(\boldsymbol{K})\}$$

and $\bar{\theta}(\boldsymbol{K})$ analogous with the min replaced by the max, we would again have $\underline{\theta}(\boldsymbol{K}) \leq \underline{\theta}(\boldsymbol{K}')$ and $\bar{\theta}(\boldsymbol{K}) \geq \bar{\theta}(\boldsymbol{K}')$. Since θ is the component of the type that is one-to-one with advice, the prediction that the extremes of advice weakly contract remains. If the profit function depended on $\boldsymbol{\mu}$ directly, it is easy to check that the second part of Proposition 2 would hold as well.

Heterogeneous Consumers. So far, we have allowed for one dimension of heterogeneity in advice among firms. In reality, firms face a variety of consumers and the advice that the firm offers could be specific to the type of consumer. To accommodate this possibility, let a firm's type be denoted by a vector $\boldsymbol{\theta}_j$ such that the profit of offering a consumer of type i advice a is $\pi(a; \theta_{ij})$, maximized at $a = \theta_{ij}$.[2] Thus, firms are now categorized by the advice they give to each type of consumer. We assume *random sorting* of consumers to firms so that each consumer receives a mass ν_i of consumers of type i. Then, the profit of a type $\boldsymbol{\theta}_j$ firm if a mass μ firms enter is

$$f(\mu) \cdot \sum_i \pi(\theta_{ij}; \theta_{ij})\nu_i - K.$$

Again, Proposition 1 applies, so that $\mathcal{E}^*(K') \subseteq \mathcal{E}^*(K)$. Denote

$$\underline{\theta}(K) \equiv \min \{\theta : \theta = \min \boldsymbol{\theta}_j \text{ such that } \boldsymbol{\theta}_j \in \mathcal{E}^*(K)\}$$

as the minimum advice given to some consumer in the market, and define $\bar{\theta}(K)$ analogously. Then, once again, $\underline{\theta}(K) \leq \underline{\theta}(K')$ and $\bar{\theta}(K) \geq \bar{\theta}(K')$ purely from the fact that the set of firms who enter shrinks if fiduciary duty operates through a pure fixed cost framework.

A.1.4. A "Smooth" Advice Channel

The example in Section 5.3 uses a stark advice channel where advice above a level is infintely costly. Here, we simply record the straightforward result that we can relax this assumption.

Proposition 1. *Suppose the cost $c(\cdot)$ of advice is weakly increasing. Then, holding the entry rate $\boldsymbol{\mu}$ fixed, advice of a firm weakly improves when moving from a market with fiduciary duty to a market without.*

Proof. Fix a type θ and a entry rate $\boldsymbol{\mu}$; suppress the dependence on $\boldsymbol{\mu}$. Let $a^*_{NFD}(\theta) \equiv \arg\max_a \pi(a; \theta)$ be the advice given by this type without fiduciary duty. θ is the advice given by this type in the absence of fiduciary duty. Suppress dependence on Note that the advice with fiduciary duty is

$$a^*_{FD}(\theta) \equiv \arg\max_a \pi(a; \theta) - c(a).$$

[2] We drop categories to limit the number of subscripts we must carry in the notation, but the arguments apply with multiple categories as well.

Consider the function $s(a, \lambda) \equiv \pi(a; \theta) - c(a)$, and let $a^*(\lambda)$ be the maximizer of this function. Note that $s(a, \lambda)$ has weakly decreasing differences in (a, λ) since $c(\cdot)$ is weakly increasing. Then, it must be that $a^*(\lambda)$ is weakly decreasing in λ. Since $a^*_{FD}(\theta) = a^*(1)$ and $a^*_{NFD}(\theta) = a^*(0)$, it must be that $a^*_{NFD}(\theta) \geq a^*_{FD}(\theta)$. Thus, advice weakly improves upon imposition of fiduciary duty, as long as the cost $c(\cdot)$ is increasing in its argument. □

A.2. Adding Registered Investment Advisers

Now suppose that in additional to broker-dealers, there are registered investment advisers in the market as well. These RIAs will not be impacted by fiduciary duty in any way. We should first note that in a model with $M > 1$ categories of broker-dealers, we could think of an RIA as one of the categories—e.g., one for whom K_m never changes with policy. Indeed, in this section, we will effectively treat RIAs in this manner. In Section A.1.2, we noted that having $M > 1$ may not necessarily lead to comparative statics in which the set of broker-dealers drops. Thus, we show in this section that with one category of broker-dealer and one RIA type, there are natural conditions under which the set of broker-dealers who enters the market would shrink under an increase in fixed costs.

Both broker-dealers and RIA firms have a type θ_j, and the latent distribution of types for broker-dealers and RIAs is given by $H_{BD}(\cdot; \theta_j)$ and $H_{IA}(\cdot; \theta_j)$ respectively. We do not take a stance on how $H_{BD}(\cdot; \cdot)$ and $H_{IA}(\cdot; \cdot)$ relate to each other. A type θ_j firm has profit function $\pi_T(\cdot; \theta_j)$ and pays entry cost K_T to enter, where $T \in \{BD, IA\}$. While we will use the notation θ_j throughout, note that type can be replaced by any of the extended types from before, e.g., (θ_j, ϵ_j) or $\boldsymbol{\theta}_j$. A firm who enters will earn profits (net of entry costs)

$$f_T\left(\mu_{BD}, \mu_{IA}\right) \cdot \pi^*_T(\theta_j) - K_T,$$

where $\pi^*_T(\theta_j) = \max_a \pi_T(a; \theta_j)$ and f_T is a share function that is decreasing in both the proportion of broker-dealers who enter and the proportion of RIA firms who enter. An equilibrium is defined to be a pair (μ^*_{BD}, μ^*_{IA}) such that

$$H_T\left(\mathcal{E}_T(\mu^*_{BD}(K_{BD}, K_{IA}), \mu^*_{IA}(K_{BD}, K_{IA}), K_T)\right) = \mu^*_T(K_{BD}, K_{IA})$$

for $T \in \{BD, IA\}$, where $\mathcal{E}_T\left(\mu_{BD}, \mu_{IA}, K_T\right)$ is the set of firms of type T who would enter if they believe the share of broker-dealers who enter to be μ_{BD}, the share of RIA firms who enter is μ_{IA}, and the entry cost of type T is K_T.[3] As before, let the equilibrium set of entrants of type T be

[3]The entry decision for broker-dealers does not directly depend on the entry cost for RIA firms, say, but does indirectly depend on it in equilibrium through the entry decision of RIAs.

$\mathcal{E}_T^*(K_{BD}, K_{IA})$. Fiduciary duty influences neither $\pi_{IA}(\cdot;\theta_j)$ nor K_{IA}. If fiduciary duty operates through a pure fixed cost channel, then K_{BD} increases to K'_{BD}.

Rearrange the types of these firms in decreasing order of profits so that the distribution of types is $[0,1]$. Then, an equilibium consists of

$$(\mu_{BD}^*(K_{BD}, K_{IA}), \mu_{IA}^*(K_{BD}, K_{IA}))$$

such that

$$\begin{aligned} \hat{\pi}_{BD}\left(\mu_{BD}^*, \mu_{IA}^*\right) \equiv f_{BD}\left(\mu_{BD}^*, \mu_{IA}^*\right) \cdot \tilde{\pi}_{BD}\left(\mu_{BD}^*\right) = K_{BD} \\ \hat{\pi}_{IA}\left(\mu_{BD}^*, \mu_{IA}^*\right) \equiv f_{IA}\left(\mu_{BD}^*, \mu_{IA}^*\right) \cdot \tilde{\pi}_{IA}\left(\mu_{IA}^*\right) = K_{IA}, \end{aligned} \tag{A.1}$$

where $f_T(\cdot;\cdot)$ is strictly decreasing in both of its terms and captures the competitive effects. Accordingly, the effective profit functions $\hat{\pi}_T(\cdot;\cdot)$ are decreasing in both its arguments.

We impose that cross-price competitive effects are not too strong.[4]

Assumption 1. *Assume*

$$\frac{\partial \hat{\pi}_{BD}}{\partial \mu_{BD}} \cdot \frac{\partial \hat{\pi}_{IA}}{\partial \mu_{IA}} > \frac{\partial \hat{\pi}_{BD}}{\partial \mu_{IA}} \cdot \frac{\partial \hat{\pi}_{IA}}{\partial \mu_{BD}}. \tag{A.2}$$

The left-hand side of (A.2) is the product of the sensitivities of effective profits to the own-type competition, and the right-hand side is the sensitivity of profits to cross-type competition. The following example provides some intuition on Assumption 1.

Lemma 3. *Suppose*

$$f_{BD}^{-1}\left(\mu_{BD}, \mu_{IA}\right) = \gamma_{11}\mu_{BD} + \gamma_{12}\mu_{IA} \textit{ and } f_{IA}^{-1}\left(\mu_{BD}, \mu_{IA}\right) = \gamma_{21}\mu_{BD} + \gamma_{22}\mu_{IA}.$$

Then, if $\gamma_{11}\gamma_{22} > \gamma_{12}\gamma_{21}$, then Assumption 1 is satisfied.

Proof. Direct computations show that the left-hand side of (A.2) is

$$L \equiv \left[\pi'_{BD}\left(\gamma_{11}\mu_{BD} + \gamma_{12}\mu_{IA}\right) - \pi_{BD} \cdot \gamma_{11}\right] \cdot \left[\pi'_{IA}\left(\gamma_{21}\mu_{BD} + \gamma_{22}\mu_{IA}\right) - \pi_{IA} \cdot \gamma_{22}\right],$$

times a positive constant. Both terms in parentheses are negative, so we can say

$$L > \pi_{BD}\gamma_{11} \cdot \pi_{IA}\gamma_{22}.$$

The right-hand side is

$$\pi_{BD}\gamma_{12} \cdot \pi_{IA}\gamma_{21},$$

[4] See Bulow et al. (1985) for an example of a paper where similar conditions are used to impose stability of equilibria in a pricing game.

times the same positive constant. If $\gamma_{11}\gamma_{22} > \gamma_{12}\gamma_{21}$, we thus have the result. $\square$

Similar calculations show that a sufficient condition for Assumption 1 under more general f involves replacing $\hat{\pi}_T$ by f_T in (A.2). Under Assumption 1, we can prove both uniqueness and intuitive comparative statics.

Lemma 4. *If Assumption 1 holds, then (i) there is a unique solution to* (A.1)*; (ii) holding K_{IA} fixed, the set of broker-dealers who enter under at K_{BD} is a superset of those who enter at $K'_{BD} > K_{BD}$, and (iii) holding K_{IA} fixed, the set of RIA firms who enter under at K_{BD} is a subset of those who enter at $K'_{BD} > K_{BD}$.*

Proof. According to the Gale-Nikaido Theorem, the solution to (A.1) is unique if the matrix

$$\begin{pmatrix} -\frac{\partial \hat{\pi}_{BD}}{\partial \mu_{BD}} & -\frac{\partial \hat{\pi}_{BD}}{\partial \mu_{IA}} \\ -\frac{\partial \hat{\pi}_{IA}}{\partial \mu_{BD}} & -\frac{\partial \hat{\pi}_{IA}}{\partial \mu_{IA}} \end{pmatrix}$$

is a P-matrix. This conditions means all principal minors must be positive. Both diagonal elements are positive since the effective profit is decreasing in the number of entrants of either type. Under Assumption 1, the determinant is positive as well.

To prove (ii) and (iii), take the total derivative of (A.1) with respect to K_{BD}. Then,

$$\begin{pmatrix} \frac{\partial \hat{\pi}_{BD}}{\partial \mu_{BD}} & \frac{\partial \hat{\pi}_{BD}}{\partial \mu_{IA}} \\ \frac{\partial \hat{\pi}_{IA}}{\partial \mu_{BD}} & \frac{\partial \hat{\pi}_{IA}}{\partial \mu_{IA}} \end{pmatrix} \begin{pmatrix} \frac{d\mu_{BD}}{dK_{BD}} \\ \frac{d\mu_{IA}}{dK_{BD}} \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}. \tag{A.3}$$

Solving (A.3) for the derivatives gives

$$\begin{pmatrix} \frac{d\mu_{BD}}{dK_{BD}} \\ \frac{d\mu_{IA}}{dK_{BD}} \end{pmatrix} = \left(\frac{\partial \hat{\pi}_{BD}}{\partial \mu_{BD}} \cdot \frac{\partial \hat{\pi}_{IA}}{\partial \mu_{IA}} - \frac{\partial \hat{\pi}_{BD}}{\partial \mu_{IA}} \cdot \frac{\partial \hat{\pi}_{IA}}{\partial \mu_{BD}} \right)^{-1} \begin{pmatrix} \frac{\partial \hat{\pi}_{IA}}{\partial \mu_{IA}} & -\frac{\partial \hat{\pi}_{BD}}{\partial \mu_{IA}} \\ -\frac{\partial \hat{\pi}_{IA}}{\partial \mu_{BD}} & \frac{\partial \hat{\pi}_{BD}}{\partial \mu_{BD}} \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix}. \tag{A.4}$$

Assumption 1 ensures the first term in (A.4) is positive. The elements of the first column are negative and positive, respectively, which completes the argument. $\square$

Thus, as long as cross-type competitive effects are not too strong, we have

$$\mathcal{E}^*_{BD}(K'_{BD}, K_{IA}) \subseteq \mathcal{E}^*_{BD}(K_{BD}, K_{IA}) \text{ and } \mathcal{E}^*_{IA}(K_{BD}, K_{IA}) \subseteq \mathcal{E}^*_{IA}(K'_{BD}, K_{IA}). \tag{A.5}$$

The result in (A.5) is important for two reasons. First, it shows that even in the presence of a set of firms unaffected by the regulation, the prediction that a pure fixed cost channel must shrink the set of broker-dealers remains robust—at least with a reasonable condition on how strongly these firms compete with one another. Accordingly, the predictions on the extrema of advice discussed above will still bear out. Second, it provides predictions about *spillover* effects onto RIAs. In particular, since the set of RIA firms expands (weakly), it must be the case that the best advice offered by them improves and the worst advice becomes worse.

An example similar to the cap from Section 5.3 shows that if fiduciary duty operates through an advice channel as well, then it is still possible for the best advice given by broker-dealers to improve. However, as long as the mass of broker-dealers who enters decreases, the mass of RIA firms would weakly increase. Since the base profit functions of the RIA firms do not change, we would still have an expansion in the set of RIAs, meaning that the predictions on the support of the advice will be isomorphic in both channels.

B. Additional Empirical Results

B.1. Nationwide Summary Statistics

While the body of the paper focuses on relevant border counties, we provide further summary statistics on all advisers and transactions in the dataset. Table B.1 shows summary statistics for all advisers in the US between 2013 and 2015 who sell at least one FSP contract. About 19% of advisers are broker-dealers. BDs tend to sell slightly fewer FSP contracts over this time period, amounting to about 5.3 on average compared to 5.5 for RIAs. Half of advisers sell fewer than three contracts in this time period, although there is a sizable tail of advisers selling many more. Conditional on selling an FSP annuity, BDs sell VAs about 77% of the time, while the proportion is somewhat larger for RIAs. Contract amounts are indeed significantly larger for RIAs than BDs, by about $40,000 off a baseline of about $120,000 for BDs. Finally, most of the clients are nearing or slightly past retirement, as would be expected in a market for retirement products. BDs and RIAs tend to have similar clientele, although the average age of clients in RIAs is higher by about 3 years.

Comparing Tables 1 and B.1 shows that restricting to the border limits us to about 10% of the sample in terms of advisers and about 11% in terms of contracts. However, the characteristics of financial advisors and financial transactions are rather representative of the broader US. The proportion of broker-dealers is about 2 pp lower nationally than in the border. Advisers at the border sell a slightly larger number of contracts on average than the typical adviser in the US, although inspection of the quantiles of this distribution suggests that this result may be driven by a longer upper tail of advisers. The probability of a transaction corresponding to a variable rather than a fixed annuity is similar for advisers at the border relative to advisers overall. Contract amounts tend to be slightly lower at the border, a result driven once again by the tail of contracts, and the ages of the client are not appreciably different from the population of clients in the US.

Table B.2 shows summary statistics for characteristics and returns of all transactions. Comparing the means to Table 2 suggests that the products transacted at the border are also comparable to ones

Table B.1: Summary statistics for all counties

				Percentiles				
	N	Mean	Std.	10%	25%	50%	75%	90%
Adviser-Level Quantities								
Is Broker-Dealer								
FSP Advisers	39,013	0.186						
Contracts per Adviser								
BD	7,244	5.3	8.7	1	1	2	6	12
RIA	31,769	5.5	8.5	1	1	3	6	13
Contract-Level Quantities								
Is Variable Annuity								
BD	38,041	0.774						
RIA	174,479	0.912						
Contract Amounts ($K, 2015)								
BD	38,041	118.8	146.8	23.2	40.2	77.3	143.8	252.3
RIA	174,479	157.7	197.7	34.4	56.1	101.4	197.5	314.3
Client Age								
BD	38.041	61.7	10.4	49	56	62	68	75
RIA	174,479	64.6	9.8	53	59	65	71	77

transacted nationwide, which may further allay some concerns about whether the products in the main sample are representative.

B.2. Covariate Balance

Our identifying assumption rests on the argument that even though common law fiduciary status of a state may be correlated with average demand in the state, there are no demand discontinuities at the border. For corroborating evidence on this point, we run covariate balance checks for a variety of demographic and economic characteristics. To run these checks, we estimate regressions at the county level of the demographic quantity on a dummy for whether the county has fiduciary duty. We estimate specifications with and without fixed effects and sometimes dropping counties that do not have any transactions from FSP. In all specifications, we restrict to the relevant border. Standard errors are clustered by state.

Table B.3 shows the results of these regressions. Each row corresponds to an outcome, and each column (except for the mean columns (3) and (6)) corresponds to a regression. Columns (1) and (2) restrict to counties with at least one transaction from FSP, and run the regression with and without border fixed effects. Column (3) represents the mean of the outcome variable on this sample. Columns (4)–(6) repeat this on the set of all counties in the Discovery dataset, restricted to the border. The takeaway from Table B.3 is that on almost all covariates, we estimate fairly tight zeros on the difference between means for counties with and without fiduciary duty.

Table B.2: Summary statistics for annuities sold by BDs and RIAs, all counties

	BD		RIA	
Characteristic	Mean	Std.	Mean	Std.
(A) Fund Return (%)				
Return-Maximizing	0.160	0.087	0.159	0.088
Equal	0.012	0.011	0.012	0.010
(B) Fund Expense Ratios (%)				
Minimum	0.501	0.021	0.501	0.022
Average	1.279	0.256	1.262	0.239
(C) Fees				
M&E Fee (%)	1.195	0.206	1.109	0.302
Surrender Charge (%)	3.780	1.199	3.072	1.440
(D) # Funds				
All	99.56	36.09	96.79	35.57
High Quality	27.48	11.97	31.59	14.56
Low Quality	35.98	16.64	31.88	19.02
(E) # Equity Styles				
Some High Quality	6.85	2.05	7.30	1.94
Only Low Quality	1.03	1.75	0.83	1.62
(F) # FI Styles				
Some High Quality	4.00	1.02	4.32	1.53
Only Low Quality	3.05	0.28	3.05	0.30
(G) Contract Return (all products)				
Risk-adjusted	0.031	0.013	0.026	0.010
Unadjusted	0.065	0.022	0.064	0.023

Panels (A)–(F) summarize characteristics of transacted VAs. Panel (G) summarizes characteristics of all transacted annuities. In Panel (D), "High Quality" refers to funds rated by Morningstar as 4 or 5 stars, and "Low Quality" refers to funds rated as 1 or 2 stars. In Panels (E) and (F), "Some High Quality" refers to styles covered at least by one high quality fund, and "Only Low Quality" refers to styles convered only by low quality funds.

Table B.4 shows evidence that there is no differential selection at the border into broker-dealers and registered investment advisers on some limited client dimensions we do observe. In particular, we view the age of the contract holder (at the time of purchase) and whether the client is a cross-border shopper—i.e., the client state is different from the adviser's state of business. We estimate the same regression as in (1), excluding client age fixed effects, with these as the left-hand side variambles. We find no evidence that there is differential selection by age induced by fiduciary duty. One may also wonder that clients would be willing to travel across the border to a state with fiduciary standards to purchase an annuity from a broker-dealer. This does have difficulties associated with it: for instance, the adviser would have to be licensed in the client's home state (although this is not an especially binding constraint in our dataset, since many advisers are licensed in all states). Columns (3) and (4) show that there is no differential cross-border shopping that

Table B.3: Covariate balance

	Transactions			Discovery		
	No Border FE (1)	Border FE (2)	Mean (3)	No Border FE (4)	Border FE (5)	Mean (6)
Population (K)	168.61	-105.45	134.03	35.66	28.46	102.55
	(230.00)	(97.68)		(42.48)	(26.25)	
Median Age	-0.33	0.29	40.69	-0.57	-0.60	41.37
	(0.80)	(0.45)		(0.87)	(0.43)	
Pop Black (K)	27.37	-17.52	16.17	7.72	7.13**	12.57
	(38.16)	(25.16)		(5.04)	(2.92)	
Pop Hispanic (K)	130.82	0.31	21.96	15.85	12.83	16.48
	(97.45)	(20.29)		(14.57)	(9.84)	
Median HH Income (K)	0.06	0.70	45.74	1.99	1.23*	44.45
	(6.11)	(1.97)		(2.61)	(0.68)	
Mean HH Income (K)	-1.36	-1.00	59.97	2.26	1.28	58.38
	(7.65)	(2.88)		(3.04)	(0.86)	
Pct. Unemployment	0.61	-0.55***	9.32	-0.16	-0.08	9.30
	(0.81)	(0.20)		(1.06)	(0.31)	
Pct. Poverty	-0.17	-1.00	17.34	-0.68	-0.36	17.72
	(1.81)	(0.71)		(1.67)	(0.50)	
Pct. HH with less than $25k	-0.89	-1.18	28.38	-0.99	-0.52	29.14
	(2.09)	(1.11)		(1.96)	(0.52)	
Pct. HH with less than $50k	-0.94	-1.33	54.86	-1.82	-1.10*	56.11
	(4.10)	(1.49)		(2.40)	(0.64)	
Pct. HH with less than $75k	-0.28	-0.56	73.15	-1.52	-0.77	74.31
	(4.66)	(1.48)		(2.09)	(0.61)	
Pct. HH with less than $100k	0.29	0.03	84.46	-1.26	-0.68	85.45
	(4.26)	(1.33)		(1.56)	(0.48)	
Pct. Pop less than HS	1.53	-0.44	14.50	-0.03	0.36	14.97
	(1.45)	(0.62)		(1.61)	(0.39)	
Pct. Pop HS	2.31**	1.81**	32.85	1.66	1.73***	33.68
	(0.87)	(0.87)		(1.39)	(0.52)	
Pct. Pop BA or Higher	-4.19	-1.99	19.75	-0.35	-0.71	18.65
	(3.07)	(1.42)		(1.64)	(0.57)	

Covariate balance for various economic and demographic characteristics. Each pair of columns, for each row, corresponds to the results of one regression. The first column in each pair gives the coefficient on the fiduciary duty dummy. All specifications cluster at the state level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

induces excess shopping onto the side with fiduciary duty: even if we believe that unobservably different (on sophistication, say) shoppers are the ones engaging in cross-border shopping, this effect is the same across the border. We also see from Columns (5) and (6) that running the same regression with transaction amount of the left-hand side returns statistically insignificant, albeit slightly noisier, coefficients. To the extent that transaction amount is a proxy for consumer income or wealth, this would indicate a lack of differential selection on this consumer characteristic as well. However, we interpret this result with some caution: one might worry that advisers influence the transaction amount, and fiduciary duty might affect how much they try.

Table B.4: Client covariates

	Age of Contract Holder		Cross-Border Shopper		Trans. Amount ($K)	
	(1)	(2)	(3)	(4)	(5)	(6)
DID	-0.197	0.680	-0.013	0.003	4.20	9.23
	(0.833)	(0.521)	(0.028)	(0.029)	(16.71)	(9.95)
FD on BD	-0.200	0.519	0.005	0.021	0.81	4.40
	(0.762)	(0.499)	(0.034)	(0.035)	(15.19)	(9.37)
FD on RIA	-0.003	-0.161	0.018	0.018	-3.39	-4.83
	(0.299)	(0.166)	(0.025)	(0.017)	(5.48)	(3.37)
Firm FE	No	Yes	No	Yes	No	Yes
Mean of Dep. Var	63.8	63.8	0.320	0.320	146.0	146.0
N	22,472	22,451	22,472	22,451	22,472	22,451

Contract-level regression using (1), with age of the contract holder, whether the contract is due to cross-border shopping (client state is different from adviser state), and transaction amount on the left-hand side. All specifications include border fixed effects and contract-month fixed effects but exclude age fixed effects, and Columns (2), (4), and (6) also include firm fixed effects. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

B.3. Combining Risk-Adjusted and Unadjusted Returns

One interpretation of risk-adjusted returns is that they correspond to how an individual whose SDF prices the factors in the economy would value the annuity. This individual is risk-averse, with a particular risk aversion. An interpretation of unadjusted returns is that they correspond to how much a risk-neutral individual would value the annuity. In the body of this paper, we have estimated returns using one valuation method at a time, with the risk-adjusted valuation being our preferred one given its prevalence in the finance literature.

However, a natural concern may be that valuation methods are heterogeneous. In particular, perhaps an individual who is risk-averse is more likely to buy an FIA rather than a VA. To investigate this, we estimate (1) and allow for heterogenous valuations in the population that may depend on the product purchased. In particular, we assume that a client can value each annuity either using the risk-adjusted method ("is risk-averse") or the unadjusted method ("is risk-neutral"). On the side without fiduciary duty, we assume that a proportion $\eta_{VA} \in [0, 1]$ of the clients who purchase VAs are risk-averse and the remainder are risk-neutral; a proportion $\eta_{FA} \in [0, 1]$ of clients who purchase FIAs are risk-averse. Then, we value each VA on the side without fiduciary duty as a convex combination of the risk-adjusted and unadjusted returns, with a weight η_{VA} times the risk-adjusted return; we value FIAs analogously. Given the assumption that populations do not change on either side of the border, we compute the proportions η'_{VA} and η'_{FA} on the side with fiduciary duty so that the total proportions of risk-averse individuals is constant on both sides of the border,[5] and we use

[5]That is, we impose that $\eta'_{VA} \cdot \Pr(\text{purchase VA with FD}) + \eta'_{FA} \cdot \Pr(\text{purchase FIA with FD}) = \eta_{VA} \cdot \Pr(\text{purchase VA without FD}) + \eta_{FA} \cdot \Pr(\text{purchase FIA without FD})$. We find (η'_{VA}, η'_{FA}) to minimize the distance

Table B.5: Number of firms, by footprint

	(1) Local	(2) Multistate	(3) Regional	(4) National
All Firms	-0.133*	-0.0657	0.0036	-0.0398
	(0.0702)	(0.0495)	(0.0577)	(0.0580)
BD Firms	-0.115*	-0.0277	-0.0190	-0.0645
	(0.0681)	(0.0324)	(0.0485)	(0.0679)
RIA Firms	-0.0225	-0.0483	0.0173	-0.0296
	(0.0175)	(0.0485)	(0.0483)	(0.0639)

Regressions of the number of each type of firm (using the $\log(x+1)$ transformation) on fiduciary status, county controls (log population, log median household income, and median age), border fixed effects, and standard errors clustered at the border. Each coefficent shown comes from a separate regression, and the number in the table is the coefficient on the fiduciary dummy. All regressions have $N = 411$ observations. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

these proportions on the side with fiduciary duty.

We allow η_{VA} and η_{FA} to independently vary over a fine grid on $[0, 1]$ and compute the difference-in-differences estimate from (1). Of course, $\eta_{VA} = \eta_{FA} = 1$ corresponds to the risk-adjusted result and $\eta_{VA} = \eta_{FA} = 0$ corresponds to the unadjusted one, but note that using other combinations of these parameters does not necessarily imply that the estimate lies between the ones in Table 3. Nevertheless, even with this flexibility, the difference-in-differences estimates are robustly positive. Indeed, over the entire range of parameters (η_{VA}, η_{FA}), the lowest estimate that we find is 18 bp, and the 10^{th} percentile is 24 bp, both of which are statistically significant at the 5% level. This exercise provides credence that our main results are robust to some degree of heterogeneity in valuation methodologies.

B.4. Entry Rates by Firm Categories

We next study whether fiduciary duty induced a compositional shift even within broker-dealer firms, focusing on firm footprint. We use Discovery Data's classification into local, multistate, regional, and national firms. The rationale behind this investigation is two-fold. First, a natural concern is that local broker-dealers may be more susceptible to increases in costs induced by fiduciary duty—perhaps because they lack the legal and compliance departments to deal with the regulatory costs of such laws. Second, if different groups of broker-dealer firms sustain different increases in fixed costs, then even under a pure fixed cost channel we may see an expansion in advice from broker-dealers. However, Section 5.2 shows that this expansion cannot happen without an expansion in at least of the groups. As such, the effect of fiduciary duty on entry for a natural grouping of broker-dealer firms is a relevant robustness check for the testable predictions of the model.

to (η_{VA}, η_{FA}) subject to satisfying the aforementioned equality.

Table B.5 presents results of regressions where the left-hand side is (the log of one plus) the count of the number of firms of each footprint, and the right-hand side has the same set of variables the regressions in Table 5. The numbers presented in the table are the coefficient of the fiduciary dummy in separate regressions. The first row shows that among all firms, the ones that are affected most strongly by regulation are the ones with a local footprint, with the number of local firms dropping by about 13%. Consistent with the notion that the direct incidence falls on broker-dealers, the second row shows that local broker-dealers are affected strongly. The third row suggests no strong compostional effect among RIA firms. We should note, however, that the compositional shift we identify among broker-dealers is due to "exit" of firms: we do not see any evidence that the decrease in the number of local broker-dealers induces *more* regional or national broker-dealers to enter.

B.5. Estimates with Firm Fixed Effects

Table B.6 reports estimates of (1), but adding firm fixed effects, for all outcomes in this paper. A prediction of the fixed cost channel is that within-firm behavior should not change as a function of fiduciary duty. While results are underpowered, we broadly find that point estimates of within-BD changes are large—1/2 to 2/3 of the change within firm fixed effects—and often statistically significant. They are (almost) uniformly in the same direction as the total effect. These results provide suggestive evidence in favor of an advice channel.

B.6. Evaluating Model Fit

In this subsection, we evaluate the fit of the structural model with regards to advice and entry. To do so, we compute predicted probabilities of entry for BDs and RIAs at the market level, averaging across all footprints and realizations of θ; that is, we use the specific market's covariates and set of potential entrants. Note that this involves explicitly using the fixed point procedure to solve for the equilibrium, which is not done at any point in estimation. As with all counterfactuals in this paper, we draw 100 points from the Markov Chain at random post burn-in and estimate the entry probabilities for each draw, and we then average across draws within-market. Figure B.1(a) compares these predictions with the observed probabilities of entry at the market level, breaking the predictions up by fiduciary status of the market and BD/RIA status. The points are roughly in line with the 45° line, although the model tends to overpredict entry in markets with especially low entry. However, note that given sampling error in predicted and observed entry rates, the data points still lie close to the 45° line. Additionally, observed entry probabilities lie within the confidence region of predicted entry probabilities for both BDs and DRs. More precisely, the average observed entry probability for BDs is 8.35%, and the prediction of the model is 9.61%, with a 95% confidence

Table B.6: Characteristics of products transacted, with firm fixed effects

	Returns				Expense Ratio		Fund Returns		Fees	
	Risk-Adj. (3.1)	Unadj. (3.2)	$\mathbb{1}$[VA] (1)	10^{th} Perc. (2)	Minimum (3)	Average (4)	Optimal (5)	Equal (6)	M&E (7)	Surr. Chg. (8)
DID	0.0005	0.0029	-0.041	0.411	-0.004	0.042**	0.0143	0.0015*	-0.021**	-0.016
	(0.0010)	(0.0019)	(0.031)	(0.334)	(0.003)	(0.019)	(0.0087)	(0.0008)	(0.010)	(0.167)
FD on BD	0.0004	0.0022	-0.027	0.288	-0.004*	0.044**	0.0137*	0.0013	-0.014*	-0.066
	(0.0009)	(0.0017)	(0.034)	(0.305)	(0.002)	(0.018)	(0.0076)	(0.0008)	(0.008)	(0.130)
FD on RIA	-0.0001	-0.0001	0.015	-0.123	-0.000	0.002	-0.0005	-0.0002	0.007	-0.050
	(0.0005)	(0.0005)	(0.013)	(0.161)	(0.00)	(0.009)	(0.0034)	(0.0005)	(0.008)	(0.055)
Base Mean	0.028	0.064	0.878	2.610	0.501	1.263	0.159	0.012	1.088	3.108
N	22,451	22,451	22,451	22,451	19,711	19,711	19,711	19,711	19,711	19,711

	# Funds			# Equity Styles		# FI Styles	
	All (9)	≥ 4 Stars (10)	≤ 2 Stars (11)	High Q. (12)	Only Low Q. (13)	High Q. (14)	Only Low Q. (15)
DID	5.75	2.45*	1.74	0.393*	-0.321	0.050	-0.066***
	(3.48)	(1.36)	(1.45)	(0.224)	(0.215)	(0.091)	(0.017)
FD on BD	6.81**	2.14*	2.60*	0.367*	-0.331*	-0.008	-0.067***
	(3.28)	(1.27)	(1.34)	(0.206)	(0.194)	(0.060)	(0.017)
FD on RIA	1.06	-0.31	0.86	-0.026	-0.009	-0.058	-0.002
	(1.63)	(0.41)	(0.93)	(0.083)	(0.100)	(0.055)	(0.009)
Base Mean	96.81	32.05	31.34	7.216	0.864	4.408	3.028
N	19,711	19,711	19,711	19,711	19,711	19,711	19,711

Estimates of (1) for various product characteristics. Columns (3.1), (3.2), (1), and (2) use the set of all annuities transacted in the border, while the other columns restrict to variable annuities. All specifications include firm fixed effects. Columns (3.1) and (3.2) should be compared to Columns (1) and (2) of Table 3, and all other columns should be compared to the corresponding columns in Table 4. Standard errors are clustered at the state level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Figure B.1: Model fit

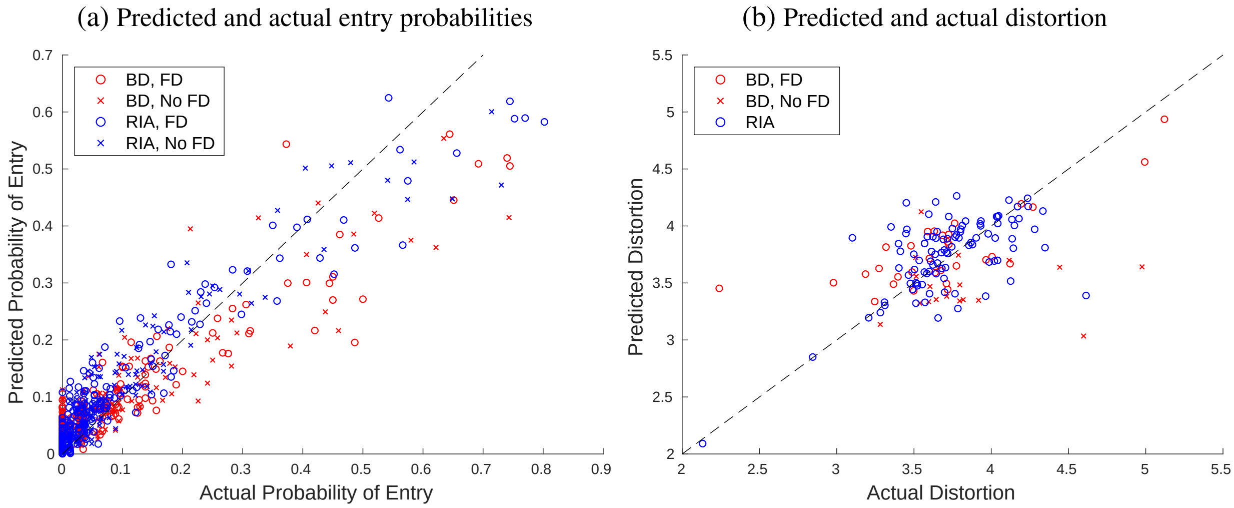

region of $[7.45\%, 11.99\%]$. For DRs, the average observed entry probability is 8.77%, and the prediction of the model is 10.93%, with a 95% confidence region of $[8.74\%, 13.20\%]$.

We consider another check on the entry probabilities, as they are an important component to the counterfactuals and are also simulated post-estimation. In the data, we can run a regression at the market-potential entrant level of whether the potential entrant enters on a dummy for BD, a dummy for fiduciary status of the market, their interaction, controlling for border fixed effects. Doing so leads to a coefficient of -0.89 pp for the interaction (standard error of 0.55 pp). We can repeat this exercise in the simulations, with the probability that a potential entrant enters on the left-hand side. Over 100 parameter draws, we arrive at a mean of -1.10 pp as the coefficient on the interaction, with a 95% confidence interval of between -1.76 pp and -0.52 pp over draws. Thus, the model is able to match this difference-in-difference coefficient well.

Finally, we consider fit for advice. Figure B.1(b) computes predicted and average advice at the firm level—for firms for whom we observe advice (i.e., transact FSP products). While there is still noticeable variation from the 45° line, the pattern is still clear.

C. Computing Investment Returns

In this section, we detail how we compute investment returns for the investment options (often called *subaccounts*) available to the clients and decide on the set of investment allocations from which the clients can choose. We also discuss how we aggregate historical information from FIA rate sheets. These inputs feed into the calculation of the net present values computed in Appendix D.

C.1. Computing Returns for Variable Annuities

For each investment option in the variable annuity dataset, we can match by name to CRSP Survivorship-Bias-Free US Mutual Fund Database. CRSP provides a permanent fund number, which is invariant to name changes, which we then track to find montly net asset values dating from January 1, 1990. We compute monthly returns from changes in this net asset value instead of using CRSP's monthly return, since variable annuity subaccounts do not reinvest dividends on behalf on the annuitants: reinvested dividends accrue to the firm. (One can check that the computed number is identical to the CRSP monthly return less dividends reinvested.) From CRSP, we also collect historical monthly risk-free rates (proxied by the one-month treasury), the excess return of the market, and the Fama-French factors, at the monthly level from January 1990.

We compute returns and covariances using two main methods. These estimates then feed into the computation of the optimal portfolios.

Stochastic Discount Factor. The first is employing a linear factor model for both the stochastic discount factor and the returns of the annuity. In this process, we first need an estimate of the stochastic discount factor m_t. We model $m_t = a - \sum_i b_i f_i$ where f_i consists of just the excess return of the S&P index (over the risk-free rate) and the size premium (small minus big) and the value premium (high minus low) in the three-factor case. In the one-factor case, we simply use the excess return of the S&P index. We then posit a risk-free rate r^* that we will use to value the variable annuity. Then, we use the restrictions $\mathbb{E}[m(1 + r^*)] = 1$ and $\mathbb{E}[mf_i] = 0$ for all i to estimate a and b_i, by replacing the expectations with their empirical counterparts. We convert the monthly returns to quarterly ones to compute a quarterly discount factor. In practice, we use all groups of three consecutive months as a separate observation of the quarter.

We then must then value the funds. We use a factor model for the returns as well, positing that for fund j in quarter t

$$r_{jt} - r_t = \alpha_j + \sum_i \beta_{ji} f_{it} + \epsilon_{jt}, \tag{C.1}$$

where r_t is the observed risk-free rate in quarter t. We can estimate α_j and β_{ji} through OLS, and we also recover a distribution of abnormal returns ϵ_{jt} for the quarters where we observe returns of the fund. While almost all estimates α_j are negative—consistent with these funds having higher than normal expense ratios and sometimes withholding dividends—we estimate some funds to have positive (but especially small) α.

Using these estimates, we can compute an (i) expected discounted mean for each fund and (ii) covariance matrix for all funds that are options. We estimate the mean as simply its empirical counterpart

$$\frac{1}{T}\sum_t \hat{m}_t \left(r^* + \hat{\alpha}_j + \sum_i \hat{\beta}_{ji} f_{it} \right), \tag{C.2}$$

where the sum ranges over all T quarters starting from 1990, r^* is the *posited* discount rate to be used for the value calculations, and the hats denote the estimates computed from above. In this version of the computation, $\hat{\beta}$ do not play a role in this calculation by construction, and $\hat{m}_t$ was chosen so that their product with the discount factor averaged to 0.

The covariance matrix is computed in two steps. We first compute the empirical covariance matrix of the distribution of the terms in the summand in (C.2) across funds j. Call this $\hat{V}_1$. We then compute the empirical covariance matrix of the abnormal returns, and we denote this $\hat{V}_\epsilon$. Since funds may not have full overlap (as they enter into the market at different times), we compute the elements of the covariance matrix pairwise, which means that $\hat{V}_\epsilon$ is not guaranteed to be positive semidefinite. Direct expansion of the terms for the covariance of the discounted returns shows that the total covariance matrix is $\hat{V} \equiv \hat{V}_1 + \mathbb{E}[\hat{m}^2]\hat{V}_\epsilon$. Since this expression need not be positive semidefinite in finite samples (though it often is), our final step involves finding the closest positive semidefinite matrix to it, to convert it to a valid covariance matrix. Letting $QUQ' \equiv \hat{V}$ denote the Schur decomposition of $\hat{V}$, we generate the matrix U^+, which replaces all negative elements of U (which will be a diagonal matrix in this case) with zeros. We then use $\hat{V}^+ \equiv QU^+Q'$ as the estimated covariance matrix.[6]

Risk-Free Rate. We run another version of the computations in which the agent discount returns not via the stochastic discount factor but via the posited risk-free rate. In this situation, we follow all the above steps but simply impose $m = 1/(1+r^*)$. In particular, we still model the returns using the factor structure: given that some funds were only introduced after the crisis and others have endured periods of downturns as well, the raw means and variances would introduce substantial bias.

C.2. Optimal Portfolio Allocation for Variable Annuities

Investment restrictions partition the set of funds available into groups and place minimums and maximums on the shares of assets that can be placed in each group. If s is the vector of shares of each fund, this effectively amounts to a linear restriction $Ms \geq m$. The only portfolios a client can choose are ones that satisfy this restriction. If r is the vector of estimated returns, the maximum

[6]We have checked for numerical issues by using a semidefinite solver, which achieves the same solution through a different algorithm. Furthermore, the norm of $\hat{V}^+ - \hat{V}$ is usually very small, suggesting this procedure does not change the matrix appreciably—as one would hope.

possible return is simply the linear program

$$\max_{s} r \cdot s \text{ s.t. } Ms \geq m \text{ and } s \cdot \mathbb{1} = 1, \tag{C.3}$$

if $\mathbb{1}$ is a vector of ones. This program can be solved efficiently; we use Gurobi.

However, the client will not necessarily pick the mean-maximizing return. Moreover, the set of possible allocations is still infinite, so we cannot solve the dynamic programming problem over this entire set. Instead, we allow the client to choose portfolios on the mean-variance frontier. The intuition is simple: facing two portfolios with the same volatility, the client should pick the one with the higher mean. Thus, for a fixed variance, we could find the highest mean attainable and thus compute an "extended" efficient frontier. Alternatively, for each mean, we can compute the lowest and higher variance attainable. Due to the convexity of the contract, the client may prefer higher variance. However, due to different funds having different returns, high variance may come at a cost, just as low variance comes with a cost.

We can solve for the typical variance-minimizing portfolios as

$$\min_{s} s'\hat{V}^{+}s \text{ s.t. } Ms \geq m,\ r \cdot s \geq \bar{r}, \text{ and } s \cdot \mathbb{1} = 1, \tag{C.4}$$

for a fine grid of minimum returns $\bar{r}$ from the minimum possible return to the maximum one (i.e., the solution to (C.3)). This is a convex quadratic program and can also be solved efficiently by Gurobi. The analoguous variance-maximizing program is identical but with the min replaced by a max. This problem is non-convex, but we find using KNITRO's multistart that we can reliably and efficiently find a solution.

In the case where α is set to a constant and we use a stochastic discount factor, all funds return the same mean. In these cases, we simply find the minimum and maximum variance attainable and allocations that attain them. Since the set of attainable portfolios is convex,[7] all variances between the extremes can be attained. We then use nine equally spaced allocations between the two extremes as additional elements of the choice set.

C.3. Computing Rates for Fixed Indexed Annuities

In this paper, we compute returns in a world with a one-year risk-free return of 3%. The factor models for returns account for this risk-free rate directly. However, rates for fixed indexed annuities

[7] The caveat to this statement is that some products have two possible investment restrictions: clients can choose funds that satisfy one set or the other. In such situations, the set of possible portfolios need not be convex. However, we have checked that we do not have any situations where the set of attainable variances do not overlap, i.e., the minimum in one set is never larger than the maximum in the other. Thus, the same spanning property holds. In other situations, we simply take account of these two sets by solving the minimization or maximization problem separately for each set of restrictions.

are set for different crediting strategies, and they are changed monthly as the interest rate changes. To impute rates for these crediting strategies in the return calculations, we interpolate based on the relationship between the historical rates for different crediting strategies for a particular annuity and treasury rates. Fixing a product, the procedure follows.

1. Rates for different crediting strategies for a product are strongly (linearly) correlated, and we wish to use this relationship to improve the accuracy of our predictions of rates. To do so, we "normalized" rates to the rate that would be provided by the fixed crediting strategy, as all products in our dataset have a fixed crediting strategy. That is, for each crediting strategy c and month m, we regress the fixed rate r_m^x on r_m^c. We then compute $\hat{r}_m^{xc}$, the predicted value of the fixed rate implied by the crediting strategy c in month m.

2. We regress r_m^{xc} on the five-year treasury rate: observations are at the month level, and we stack the regression across all crediting strategies c provided by the product. This regression then lets us predict the rate provided by the fixed crediting strategy for any value of the five-year treasury rate.

3. We compute the five-year rate implied by a one-year rate of 3%, averaging across historical realizations of the yield curve. To do so, we regress the five-year rate r_m^5 on the one-year rate r_m^1, where each observation is a month (starting at 1990). We estimate an implied rate of 3.67%. We then plug this estimate into the regression from Step 2 to impute the rate provided by the fixed crediting strategy for this product.

4. To compute the rates for other strategies, we run the reverse of the regression from Step 1, i.e., r_m^c on r_m^x. We then use the predicted value at the imputed rate for the fixed strategy from Step 3.

We have experimented with variations of this procedure. The results in this paper are robust to modifications such as dropping Step 3 (so that the rates are predicted at a five-year rate of 3%) or using a ten-year treasury rather than the five-year rate.

D. Computations of Net Present Values

This appendix section presents the detailed explanation of how variable and fixed income annuities are valued. It is divided into three subsections. The first introduces notation and presents relevant definitions. The second derives how to value a variable annuity contract with a minimum withdrawal living benefit and an account value death benefit, the most prevalent contract in our dataset. The third modifies this derivation for variable annuities and fixed indexed annuities.

D.1. Definitions and Contract Rules

When a VA or an FIA contract is signed, the invested amount becomes the contract value at period 0, c_0. Contracts with living benefit riders also generate an income base b_0, which is equal to c_0 at this moment, but will typically diverge over time. Let $c_t \in \mathbb{R}^+$ denote the contract value in period t and $b_t \in [c_0, \bar{b}]$ denote the income base in period t. Contract values are bounded below by zero, as annuitants cannot go into debt with the insurance company, and income bases are bounded above by an amount set by the insurance company (in our data, $10 million dollars) and below by the original contract value.

Let $\mathcal{I}_t$ denote the set of feasible asset allocations available to the annuitant in period t. For variable annuities, this is restricted both by the set of funds available given the chosen contract and rider, and by the investment restrictions imposed by the contract-rider combination. For fixed indexed annuities, this corresponds to the set of crediting strategies the annuitant can choose from. Let $i_t \in \mathcal{I}_t$ denote a vector of chosen allocations in period t, and let $r_{t+1}(i_t)$ denote the return of that asset allocation, which is realized in period $t+1$. In some cases, crediting strategies for fixed indexed annuities are realized in longer horizons. For expositional clarity, we will ignore this for now and return to this issue below.

Variable and fixed indexed annuity contracts may have a fixed fee f_t, which for some contracts is waived for contract values above $\bar{f}$ and for all contracts is waived after 15 years, and a variable fee on the income base v^b. Variable annuity contracts also have a variable fee v^c on the contract value. In what follows, $v^c = 0$ for all fixed indexed annuity contracts, let $\bar{f} = \infty$ if the contract does not waive the annual fee for high contract values, and let $f_t = 0$ after fifteen contract years.

Contracts with a minimum withdrawal living benefit rider have two additional features that affect transitions of the income base and of the contract value must be introduced. First, after a given age annuitants have the option of withdrawing the Guaranteed Annual Income (GAI) amount, which is equal to the income base times the relevant GAI rate for the period, $g_t \in \{g_1, ..., g_G\}$. We detail which GAI rate is available to the annuitant in each period below, as it is a complicated function of the sequence of choices made in the past. Let $w_t \in \{0, 1\}$ denote whether the annuitant decides to withdraw the GAI amount in period t, so that the GAI withdrawal amount is $w_t \cdot g_t \cdot b_t$. Second, for the first E years of the contract, known as the enhancement period, the income base is guaranteed to grow at least by the enhancement rate e. Moreover, if certain conditions are met, an additional E years of enhancement rate eligibility can be earned. We denote the enhancement rate in period t by $e_t \in \{0, e\}$. Typical values of the enhancement period and enhancement rate during our sample period are 10 years and 5%, respectively.

Transitions of the contract value and the income base are governed by

$$\tilde{c}_t = c_t - \left(w_t g_t + v^b\right) b_t - f_t \cdot 1[c_t < \bar{f}]$$

$$c_{t+1} = \max[(1 + r_{t+1}(i_t) - v^c(i_t))\tilde{c}_t, 0]$$

$$b_{t+1} = \begin{cases} \min\left[\max\left[(1+e_t)\, b_t, \tilde{c}_t\right], \bar{b}\right] \text{ if } a_t < \bar{a} \\ b_t \text{ if } a_t \geq \bar{a} \end{cases}.$$

Define $\tilde{c}_t$ as the end-of-period contract value, equal to the contract value minus the annual fee, the fee on the income base, and the GAI withdrawal amount. In an abuse of notation, we set $w_t g_t = 0$ in years where GAI withdrawals are not available. The next period contract value is equal to the end of period contract value times the net rate of return, or the difference between the realized return on investments and the contract fee. As mentioned earlier, contract value is bounded below by zero. Finally, in every period where the annuitant's age (a_t) is less than the contract's maximum purchase age, $\bar{a}$, the income base is equal to the maximum of the contract value and the enhanced income base, provided this amount is below the maximum income base. Because of this transition rule, the income base cannot fall below the initial investment amount. After the contract's maximum purchase age, the income base is locked in and cannot change. Note that GAI withdrawals decrease the contract value but do not decrease the income base, and that they continue even when contract value equals zero.

On a period where contract value exceeds the value of the enhanced income base and no GAI withdrawals take place, the contract is said to have "stepped up." After a step up, the contract is eligible for E more years of enhancement. Let s_t denote the number of years since the last step up. Then $s_0 = 0$, $s_{t+1} = s_t \cdot 1\left[b_{t+1} \neq \tilde{c}_t \text{ or } w_t = 1\right] + 1$, and $e_t = e \cdot 1\left[s_t \leq E\right] \cdot 1\left[a_t < \bar{a}\right]$.

The GAI rate available in period t is a function of the age at which the first GAI withdrawal occurs, a^{first}. GAI withdrawals cannot be taken before a certain age a_0, typically 55, and they are increasing in the age of first withdrawal, until either 70 or 75. The contract specifies a map $G\left(a^{first}\right) : \{a_0, ..., \bar{a}\} \rightarrow \{g_1, ..., g_G\}$ from all possible ages at first withdrawal to GAI rates. For example, a contract might specify that an annuitant who takes a GAI withdrawal for the first time at age 60 receives a 3% GAI rate, while they would receive a 5% rate if they wait until age 75. Annuitants are locked in to the GAI rate at the age of first withdrawal, unless a step up takes place at a later age with a higher GAI rate. Then the GAI rate available in period t is

$$g_t = \begin{cases} \emptyset & \text{if } a_t < a_0 \\ g_{G(a_t)} & \text{if } a_t \leq a^{first} \\ g_{G(a_{t-1})} & \text{if } a_t > a^{first} \text{ and } \tilde{b}_{t-1} = \tilde{c}_{t-1} \\ g_{t-1} & \text{if } a_t > a^{first} \text{ and } \tilde{b}_{t-1} \neq \tilde{c}_{t-1} \end{cases}.$$

In summary, the set of relevant state variables in period t is (c_t, b_t, s_t, g_t), and the annuitant's control variables are whether to take a GAI withdrawal w_t and the investment allocation i_t. Finally, annuitants can withdraw the contract value at any time, receiving $c_t \cdot (1 - d_t)$, where d_t is the surrender charge in period t, or they can annuitize the contract value, receiving an expected present discounted value of the annuity stream $z(a_t, c_t)$. Note that both full withdrawal of the contract value and annuitization induces the loss of the guaranteed annual income.

Defining μ_t as the probability of being alive in period t conditional having lived to period $t-1$, the value of a contract in period t is equal to

$$V_t(c_t, b_t, s_t, g_t) = \max\left[\max_{(w_t, i_t)} w_t \cdot g_t \cdot b_t + E\left[\delta\left(\mu_{t+1} E\left[V_{t+1}(c_{t+1}, b_{t+1}, s_{t+1}, g_{t+1})\right]\right)\right.\right.$$
$$\left.\left. + (1 - \mu_{t+1})\beta E[c_{t+1}]\right], (1-d_t)c_t, E[PDV(z(a_t, c_t))]\right]$$

D.2. Solving for the Value of Variable and Fixed Indexed Annuity Contracts with a Minimum Withdrawal Living Benefit Rider

Assume that the probability of death in period T is 1, and that annuitants value a dollar left after their death by β. In our calculations, we set $\beta = 1$. Then in period $T-1$ the continuation value of the contract is $\beta E[c_T]$. Since $a_{T-1} > \bar{a}$, the income base and GAI rate are locked in (at $b_{\bar{t}}$ and $g_{\bar{t}}$, respectively), so the years since last step up are irrelevant. Then the problem in period $T-1$ is

$$V_{T-1}(c_{T-1}, b_{\bar{t}}, g_{\bar{t}}) = \max\left[\left(\max_{(w_{T-1}, i_{T-1})} w_{T-1} \cdot g_{\bar{t}} \cdot b_{\bar{t}} + \beta \cdot E[\delta \cdot c_T]\right),\right.$$
$$\left. z(a_{T-1}, c_{T-1}), (1 - d_{T-1}) \cdot c_{T-1}\right] \quad \text{(D.1)}$$

subject to

$$E[\delta c_T] = E\left[\delta \max\left[(1 + r_T(i_{T-1}) - v_T^c)\,\tilde{c}_{T-1}, 0\right]\right]$$
$$\tilde{c}_{T-1} = c_{T-1} - \left(w_{T-1} g_{\bar{t}} + v_{T-1}^b\right) b_{\bar{t}} - f_{T-1} \cdot 1[c_{T-1} < \bar{f}].$$

We use the 2012 Individual Annuity Mortality Basic Table, from the Society of Actuaries, for death probabilities. This sets $T = 121$. Additionally, contracts cannot be annuitized after age 99, so annuitization is not an option in $T-1$. Rather than introducing notation to keep track of when annuitization is available, we will always include it as an option, and implicitly set $z(a_{T-1}, c_{T-1}) = 0$ whenever it is not. Furthermore, since the maximum purchase age is 85 for variable annuities and 96 for fixed indexed annuities, and surrender periods are never more than 10 years long, in practice $d_{T-1} = 0$. We will also keep surrender charges in the notation and set them to 0 when the surrender period has expired.

To solve for the value of continuing with the contract, we discretize both the set of feasible investments $\mathcal{I}_t$, and the space of $(c_{T-1}, b_{\bar{t}})$. For every element in the contract value - income base grid, (c^k, b^k), and conditional on the GAI rate, we find the asset allocation that yields the highest expected present discounted value for both the case where the annuitant decides to take GAI withdrawals and where they do not. Taking the maximum over the utilities under both withdrawal strategies and over annuitization and full surrender yields $V^*_{T-1}(c^k, b^k, g_{\bar{t}})$, the value of following the optimal withdrawal and investment strategy after arriving at period $T-1$ with contract value c^k and income base b^k. We interpolate linearly over the (c_{T-1}, b_{T-1}) space to obtain $\hat{V}^*_{T-1}(c_{T-1}, b_{\bar{t}}, g_{\bar{t}})$, the value function in period $T-1$ for all possible combinations of contract value, income base, and GAI rate. In period $T-2$, we then solve

$$\begin{aligned} V_{T-2}\left(c_{T-2}, b_{\bar{t}}, g_{\bar{t}}\right) = \max \Bigg[& \max_{(w_{T-2}, i_{T-2})} w_{T-2} \cdot g_{\bar{t}} \cdot b_{\bar{t}} \\ & + \left(\mu_{T-1} \cdot E\left[\delta \hat{V}^*_{T-1}(c_{T-1}, b_{\bar{t}}, g_{\bar{t}}) \right] + (1-\mu_{T-1}) \cdot E[\delta c_{T-1}] \right), \\ & z(a_{T-2}, c_{T-2}), (1-d_{T-2}) \cdot c_{T-2} \Bigg] \end{aligned} \tag{D.2}$$

subject to

$$\begin{aligned} E\left[\delta c_{T-1}\right] &= E\left[\delta \max\left[\left(1 + r_{T-1}\left(i_{T-2}\right) - v^c_{T-1}\right) \tilde{c}_{T-2}, 0\right]\right] \\ \tilde{c}_{T-2} &= c_{T-2} - \left(w_{T-2} g_{\bar{t}} + v^b_{T-2}\right) b_{\bar{t}} - f_{T-2} \cdot 1[c_{T-2} < \bar{f}]. \end{aligned}$$

Again, discretizing over $(c_{T-1}, b_{\bar{t}})$ and over the set of feasible investments allows us to find $V^*_{T-2}(c^k, b^k, g_{\bar{t}})$, the value of following the optimal withdrawal and investment strategy after arriving at period $T-2$ with contract value c^k and income base b^k, and linear interpolation yields $\hat{V}^*_{T-2}(c_{T-2}, b_{\bar{t}}, g_{\bar{t}})$. We continue this process recursively until we reach the maximum purchase age in period $\bar{t}$, where we obtain $\hat{V}^*_{\bar{t}}(c_{\bar{t}}, b_{\bar{t}}, g_{\bar{t}})$.[8]

In period $\bar{t}-1$, the annuitant can still step up or enhance the income base. A step up increases the GAI rate to its highest possible level, if the annuitant is not there already. Moreover, having one or more remaining enhancement years is irrelevant. The problem is

$$V_{\bar{t}-1}\left(c_{\bar{t}-1}, b_{\bar{t}-1}, s_{\bar{t}-1}, g_{\bar{t}-1}\right) = \max \Bigg[\max_{\left(w_{\bar{t}-1}, i_{\bar{t}-1}\right)} w_{\bar{t}-1} \cdot g_{\bar{t}-1} \cdot b_{\bar{t}-1}$$

[8] Note that when contract value equals zero, we can obtain the value of the problem analytically, as annuitization and withdrawal are not available and the income base is fixed. As a result, $V^*_{\bar{t}}(0, b_{\bar{t}}, g_{\bar{t}}) = g_{\bar{t}} \cdot b_{\bar{t}} \cdot \left(1 + \sum_{\tau=\bar{t}+1}^{T} \delta^{\tau-\bar{t}} \prod_{\tau'=\bar{t}+1}^{\tau} \mu_{\tau'}\right)$.

$$+\left[\mu_{\bar{t}} \cdot E\left[\delta \hat{V}^{*}_{\bar{t}}\left(c_{\bar{t}}, b_{\bar{t}}, g_{\bar{t}}\right)\right] + (1-\mu_{\bar{t}}) \cdot \beta \cdot E\left[\delta c_{\bar{t}}\right]\right], z(a_{\bar{t}-1}, c_{\bar{t}-1}), (1-d_{\bar{t}-1}) \cdot c_{\bar{t}-1})\Big] \quad \text{(D.3)}$$

subject to

$$\begin{aligned}
E\left[\delta c_{\bar{t}}\right] &= E\left[\delta \max\left[\left(1+r_{\bar{t}}\left(i_{\bar{t}}\right)-v^{c}_{\bar{t}}\right)\tilde{c}_{\bar{t}-1}, 0\right]\right] \\
\tilde{c}_{\bar{t}-1} &= c_{\bar{t}-1} - \left(w_{\bar{t}-1} g_{\bar{t}-1} + v^{b}_{\bar{t}-1}\right) b_{\bar{t}-1} - f_{\bar{t}-1} \cdot 1[c_{\bar{t}-1} < \bar{f}] \\
b_{\bar{t}} &= \min\left[\max\left[\left(1+e_{\bar{t}-1}\right) b_{\bar{t}-1}, \tilde{c}_{\bar{t}}\right], \bar{b}\right] \\
g_{\bar{t}} &= \begin{cases} g_{A\left(a_{\bar{t}-1}\right)} \text{ if } b_{\bar{t}} = \tilde{c}_{\bar{t}-1} \text{ or } a^{first} = a_{\bar{t}} \\ g_{\bar{t}-1} \text{ otherwise} \end{cases}.
\end{aligned}$$

To increase numerical precision, we transform the state space into a single dimension by working with $\frac{CV_{\bar{t}-1}}{IB_{\bar{t}-1}}$ as the state variable. Note that an individual who continues receiving GAI withdrawals at age $\bar{t}$ receives $CV_{\bar{t}} \cdot \frac{CV_{\bar{t}}}{IB_{\bar{t}}}^{-1} \cdot NPV(1, g_{\bar{t}})$, where $NPV(1, g_{\bar{t}})$ is the NPV of receiving $g_{\bar{t}} \cdot 1$ dollars as an annuity, while an individual who withdraws the contract value receives $\frac{CV_{\bar{t}}}{IB_{\bar{t}}} \cdot IB_{\bar{t}}$. Therefore, $\hat{V}^{*}_{\bar{t}}(\frac{CV_{\bar{t}}}{IB_{\bar{t}}}, g_{\bar{t}}) = \max\left[NPV(1, g_{\bar{t}}), \frac{CV_{\bar{t}}}{IB_{\bar{t}}}\right]$, and $\hat{V}^{*}_{\bar{t}}(c_{\bar{t}}, b_{\bar{t}}, g_{\bar{t}}) = IB_{\bar{t}} \cdot \hat{V}^{*}_{\bar{t}}(\frac{CV_{\bar{t}}}{IB_{\bar{t}}}, g_{\bar{t}})$.

We discretize the $\frac{CV}{IB}$ space and solve for the optimal asset allocation for every combination of GAI rate-enhancement availability-withdrawal decision. Taking the maximum over withdrawal decisions, and comparing to the value of both annuitization and full withdrawal yields $V^{*}_{T-2}(\frac{CV}{IB}^{k}, s_{\bar{t}-1}, g_{\bar{t}})$, the value at each grid point for all combinations of GAI rates and years since the last step up. As argued earlier, in this period $V^{*}_{T-2}(\frac{CV}{IB}^{k}, 1, g_{\bar{t}}) = V^{*}_{T-2}(\frac{CV}{IB}^{k}, y, g_{\bar{t}})\, \forall y \in \{2, ..., E\}$, as the income base is locked in period $\bar{t}$. Linear interpolation yields $\hat{V}^{*}_{\bar{t}-1}\left(\frac{CV}{IB}^{k}, s_{\bar{t}-1}, g_{\bar{t}-1}\right)$.

The general recursive formulation for earlier periods is

$$\begin{aligned}
V_t\left(c_t, b_t, s_t, g_t\right) = \max\Bigg[\max_{(w_t, i_t)} w_t \cdot g_t \cdot b_t + \cdot \Big[\mu_t \cdot E\left[\delta \hat{V}^{*}_{t+1}\left(c_{t+1}, b_{t+1}, g_{t+1}\right)\right] \\
+ (1-\mu_{t\bar{+}1}) \cdot \beta \cdot E\left[\delta c_{t+1}\right]\Big], z(a_t, c_t), (1-d_t) \cdot c_t)\Bigg] \quad \text{(D.4)}
\end{aligned}$$

subject to

$$\begin{aligned}
E\left[\delta c_{t+1}\right] &= E\left[\delta \max\left[\left(1+r_{t+1}\left(i_t\right)-v^{c}_{t}\right)\tilde{c}_{t}, 0\right]\right] \\
\tilde{c}_{t} &= c_{t} - \left(w_{t} g_{t} + v^{b}_{t}\right) b_{t} - f_{t} \cdot 1[c_{t} < \bar{f}] \\
b_{t} &= \min\left[\max\left[\left(1+e_{t}\right) b_{t}, \tilde{c}_{t}\right], \bar{b}\right] \\
g_{\bar{t}} &= \begin{cases} g_{A(a_t)} \text{ if } b_t = \tilde{c}_t \text{ or } a^{first} = a_t \\ g_{t-1} \text{ otherwise.} \end{cases}
\end{aligned}$$

Since we work in $\frac{CV}{IB}$ space, we must show that the obtained values are equivalent. Note that

$$V_t(CV_t, IB_t, Y, g) = \max_{w,i} g \cdot IB_t \cdot w + E[\delta V_{t+1}(CV_{t+1}, IB_{t+1}, Y, g)]. \tag{D.5}$$

Expanding the second term, we have

$$\begin{aligned} IB_t \cdot E\Bigg\{ \delta \cdot \Bigg[1\left[\frac{CV_t}{IB_t}(1-v^c) - (g \cdot w + v^b)R \geq e_t\right] \left(\frac{CV_t}{IB_t}(1-v^c) - v^b - g \cdot w\right) \\ \cdot V_{t+1}(1, 1, \bar{Y}, g_{t+1})\Bigg] + \Bigg[1\left[\frac{CV_t}{IB_t}(1-v^c) - (g \cdot w + v^b)R < e_t\right] e_t \\ \cdot V_{t+1}\left(\frac{CV_t}{IB_t}(1-v^c) - (g \cdot w + v^b)\frac{R}{e_t}, 1, \bar{Y}, g_{t+1}\right)\Bigg]\Bigg\}, \end{aligned} \tag{D.6}$$

where $\bar{Y} \equiv \min[E, \bar{t} - t - 1]$. Grouping (D.5) and (D.6), we see that the net expression is

$$IB_t \cdot V\left(\frac{CV_t}{IB_t}, 1, Y, g\right).$$

Backward induction until the initial period yields the value of the contract, $\hat{V}_0^*(c_0, c_0, E, g_0)$. Note that as the periods decrease the set of possible GAI rates decreases, as one need not solve for the value function at age 70 for GAI rates that are only available if the first withdrawal is at age 75. Moreover, the problem is initialized with 0 years since the last step up, and the annuitant is guaranteed E enhancement years, so one need not solve for the value function for infeasible values of years since last step up during the first E years of the contract. Finally, some asset allocation alternatives for fixed indexed annuities lock in funds for more than one period. When that happens, we value that alternative using the continuation value for the appropriate horizon, rather than the continuation value for the next period.

D.3. Solving for the Value of a Variable Annuity and Fixed Indexed Annuity Contracts without a Living Benefit Rider

The problem is significantly simpler in this case, as there is no income base, no enhancement, and no step up. The problem in period $T-1$ is

$$\begin{aligned} V_{T-1}(c_{T-1}) &= \max\left[\beta \cdot E[\delta c_T], z(a_{T-1}, c_{T-1}), (1-d_{T-1}) \cdot c_{T-1})\right] \\ \text{subject to} \quad E[\delta c_T] &= E\left[\delta \max\left[(1 + r_T(i_{T-1}) - v_T^c)\,\tilde{c}_{T-1}, 0\right]\right] \\ \tilde{c}_{T-1} &= c_{T-1} - f_{T-1} \cdot 1[c_{T-1} < \bar{f}]. \end{aligned}$$

Discretizing the space of contract value allows us to solve for the optimal asset allocation if the contract is continued, and comparing this value to that of annuitization or full withdrawal yields the optimal strategy in this period for a grid of contract values. Interpolation yields $\hat{V}^*_{T-1}(c_{T-1})$, the value of following the optimal strategy in period $T-1$ if landing on that period with contract value c_{T-1}. In this setting, the only difference between a variable annuity contract and a fixed indexed annuity contract will come from the menu of investment strategies available and the value of the fees.

The recursive formulation for previous periods is

$$V_t(c_t) = \max\left[\mu_{t+1} \cdot E[\delta \hat{V}^*_{t+1}(c_{t+1})] + (1-\mu_{t+1}) \cdot \beta \cdot E[\delta c_{t+1}], z(a_t, c_t), (1-d_t) \cdot c_t\right]$$

$$\text{s.t. } E[\delta c_{t+1}] = E[\delta \max[(1 + r_{t+1}(i_t) - v^c_t)\tilde{c}_t, 0]]$$
$$\tilde{c}_t = c_t - f_t \cdot 1[c_t < \bar{f}].$$

Solving this problem by backward induction yields the value of the contract, $\hat{V}^*_0(c_0)$.

D.4. Forward Simulations

In Table 4, we report results of the effect of extending fiduciary duty to BDs on the 10th percentile of the distribution of returns of the products they sell. This requires moving beyond the mean return of each asset, the object of interest in the previous subsections, and obtaining instead the distribution of returns.

To do so, we save the optimal policies from the aforementioned problems, and draw 100 paths of returns from the time of purchase to the maturity date of the contract. The optimal policies give us the set of actions an individual would take for a grid of realizations of contract value and income base (if pertinent) for every age between contract purchase and maturity. For each draw of the path of returns, we start at contract purchase, execute the optimal action, observe the transition to the next period, and execute the optimal action again. We repeat this process until maturity. Since we only have optimal policies for a grid of contract value and income base, we interpolate them whenever necessary.

This process yields the contract value and income base available to the client at maturity for each draw of returns paths. We calculate the NPV of the optimal action at this stage, retirement or withdrawal, and add to this the NPV of any flows received prior to maturity. For example, the NPV of all GAI withdrawals taken prior to that age. For each draw, this yields the value of the contract at maturity, which we then transform to a return. The vector of return draws is our approximation to the return distribution.

E. Details of Structural Model

In this appendix, we discuss details of the sequential Monte Carlo algorithm used to sample from the posterior. As discussed in Section 6.3, we use a computational Bayesian approach and sample from the posterior implied by the likelihood. In order to make sure our chain samples a large part of the parameter space, and does not get stuck in local modes, we use a sequential Monte Carlo (SMC) algorithm. This method is proposed by del Moral et al. (2006) and used in Chen et al. (2019), and we follow the guidance in those papers in deciding on the specifications of the chain. The steps below follow the discussion in Chen et al. (2019) closely.

0. *Initialization*. We first initalize C different parameter values. Since Steps 2 and 5 of the Gibbs sampler estimate μ_θ, σ_θ, and the parameters of (2) via OLS, they are updated extremely quickly from initial guesses. As such, we focus on exploring a large set of initial guesses for γ, α, c, and θ_f. Set $j = 1$ and let $w_0^c = 1$ for all $c \in \{1, \ldots, C\}$.

1. *Inner Loop*. For each of the C parameter values, we run K steps of the chain described in Section 6.3 and compute a log likelihood conditional on the draws of θ_f and all other parameters at the tenth step of each chain c. Call this L_{jc}.

2. *Selection*. We let $v_{jc} \equiv \exp(\phi_j L_{jc})$ and let $w_{jc} \propto v_{jc} w_{j-1,c}$, normalized so that $\sum_j w_{jc} = 1$. Let the effective sample size $ESS_j \equiv \left[\frac{1}{C}\sum_c w_{jc}^2\right]^{-1}$.

   - If $ESS_j \geq \bar{s}$, then we restart every one of the C chains at its final draw of the parameter values.
   - If $ESS_j < \bar{s}$, then we start each chain c at the parameters from the end of chain c' with probability $w_{jc'}$. We then set $w_{jc} = 1$ for all c.

   Using these new starting points for the inner chain, we increment j by 1 and loop to Step 1 if $j < J$. If $j = J$, then stop.

We set $C = 2000$, $K = 10$, $\bar{s} = 2/3$, $J = 2000$, and use half the chain as a burn-in. Furthermore, $\phi_j = [j/200]^2 - [(j-1)/200]^2$ for $j \leq 200$ and $\phi_j = \phi_{200}$ for $j > 200$. This functional form allows chains with lower likelihood to survive earlier in the process. Furthemore, note for interpretation that the effective sample size is low if weights w_{jc} are especially concentrated in only a few chains, at which point the sampler is designed to (probabilistically) drop chains with low likelihoods.

The decision of $J = 2000$ is larger than the recommendations in the examples of del Moral et al. (2006) and Chen et al. (2019), who usually use $J = 200$. In practice, it seems that the chains collapse to the point where they effective mix with each other within 200 iterations. We use a longer chain to ensure convergence. Figure E.1 shows trace plots of selected parameters, overlaying all chains, with different colors used for different chains c. We show all draws for all chains so that the horizontal axis reaches $K \cdot C = 20000$. While there is a substantial amount of variance initially due to heterogeneity in initial conditions of chains, there is evidence of convergence both within and across chains.

F. Dataset Details

The analysis relies on seven sources of data: Transactions, Discovery, Beacon Annuity Nexus, Morningstar, CRSP, VA prospectuses, and FIA rate sheets. Below, we describe the data in detail, including the collection process and methods used to map across sources. We also discuss the sample selection criteria.

F.1. Data Sources

Transactions. The Transaction dataset contains information on each of FSP's transactions of annuity, deferred-contribution, and insurance products sold between January 1, 2008 and February, 2016. We restrict attention to deferred annuity (variable and fixed indexed) contracts initiated between 2013 and 2015. The unit of observation is an individual payment, including lump sum and periodic payments, but we aggregate to the contract level. In our final dataset, each observation is a unique contract, and we observe the contract amount at purchase, age of the contract holder, adviser(s) associated with the sale, as well as information on the financial product, importantly the product type and share class, and codes indicating any supplemental rider purchases.

Discovery. The Discovery dataset serves two purposes. First, we rely on it to augment the Transaction dataset with detailed information about advisers. The Discovery dataset contains information on advisers and the firms with which they were employed on December 31, 2015. We observe adviser characteristics, such as an indicator of whether the adviser is a BD or DR, the adviser's age, gender, and the location of the branch office. We use this branch location to define the adviser's fiduciary standard. Additionally, the Discovery dataset provides unique identifiers of the adviser's BD firm and RIA firm (if applicable) and includes characteristics such as firm footprint, number of employees, and primary business line. We map information from the Discovery dataset to the Transaction dataset using a unique adviser ID provided by FSP and restrict to advisers and firms available in Discovery. We cannot use this adviser ID to map externally, however.

We also leverage the Discovery dataset for the market structure analysis. We observe the universe of registered financial advisers who are able to sell annuities as of December 31, 2015. For our main specifications, the outcomes of interest are the aggregate number of advisers and associated firm branches at the county level. We also explore heterogeneity by firm footprint. Discovery defines the

Figure E.1: Trace plots for selected parameters

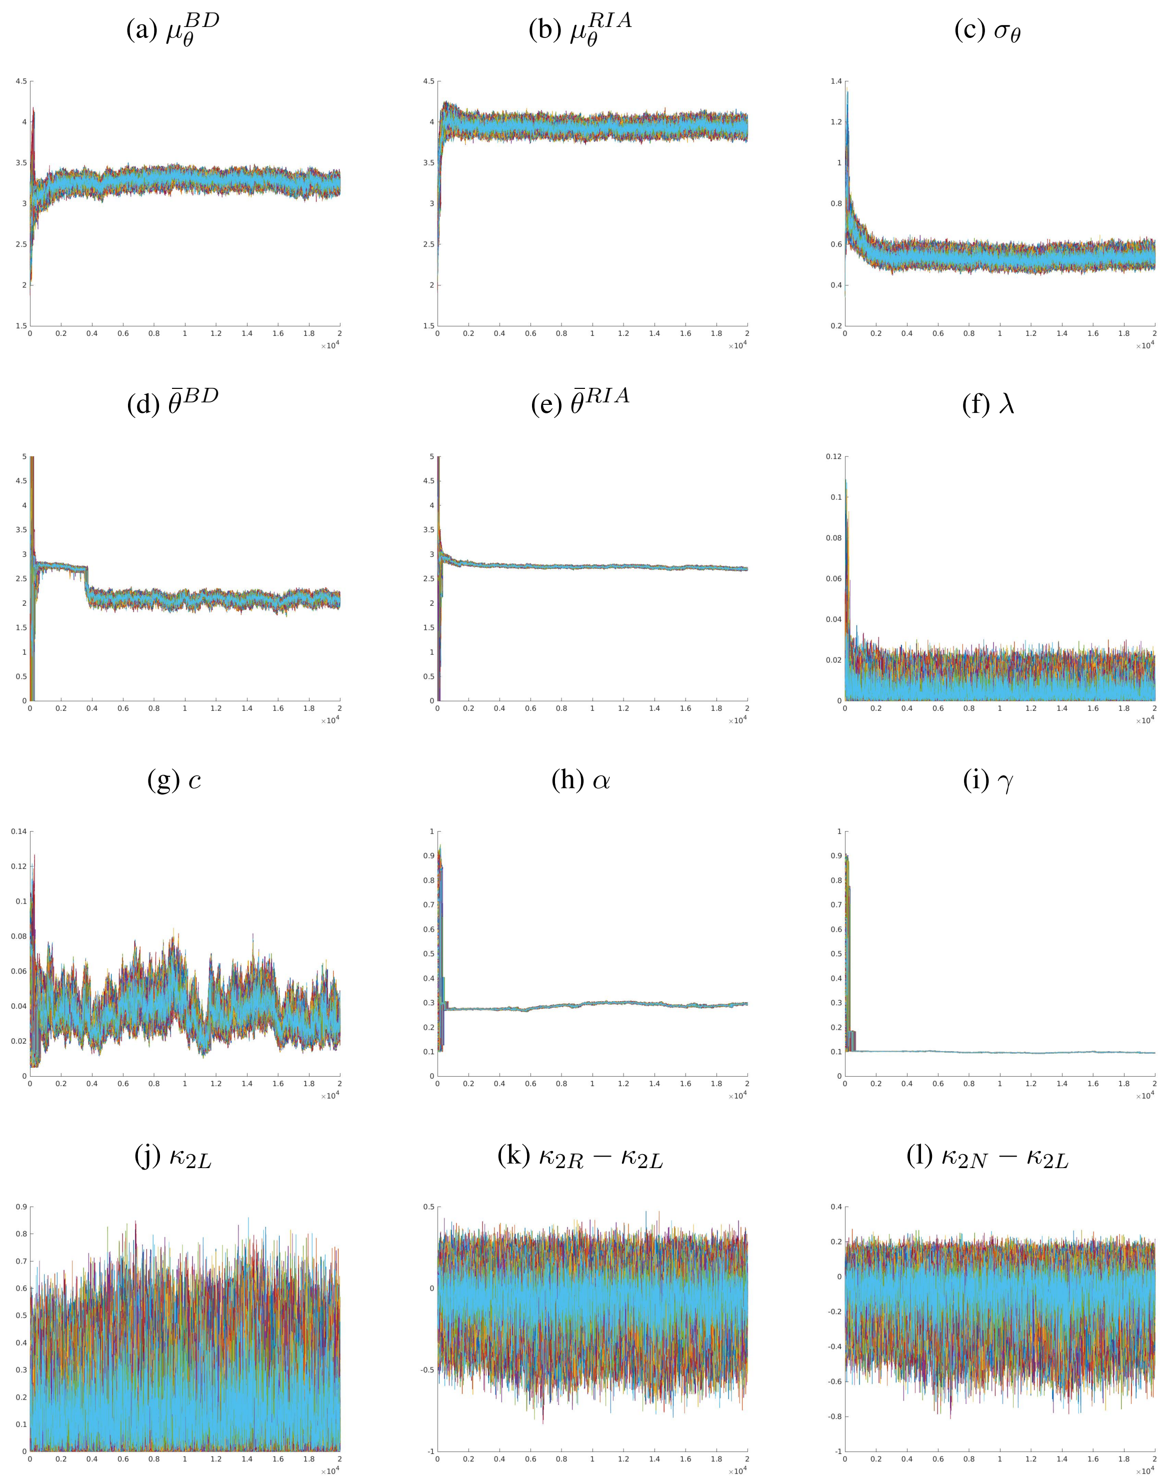

firm footprints as follows:

- Local: located in no more than a few offices in one state or close proximity
- Multistate: located in multiple states but not large or concentrated enough to be categorized as a regional firm
- Regional: substantial office and adviser coverage across a region, e.g., the Midwest
- National: substantial office and adviser coverage across the U.S.

Beacon Research. For detailed product information, we rely on Beacon Research's Annuity Nexus. This dataset provides historical information on annuity fees and characteristics, as well as changes in availability and characteristics of supplemental riders.

We manually map product names and share classes from Beacon to the detailed descriptions provided in the Transaction dataset. This mapping is straightforward because a high level of detail is provided in the Transaction dataset. The mapping of rider selections is more difficult. The Transaction dataset provides a unique code for each rider selection but does not include a description. Instead, we rely on temporal restrictions on rider availability to match the codes with Beacon. The process is as follows:

- *Rider Availability Restrictions*: Create a crosswalk that lists each rider code combination and any potential corresponding rider name in Beacon. In this step, we rely on rider availability restrictions. Specifically, if a rider is not available for a given product, then it is eliminated as a potential mapping for all rider code combinations associated with that product in the Transaction dataset. Note that, after implementing the availability restrictions, there are certain combinations of rider codes that could only correspond to a single Beacon name, while others could correspond to more than one.
- *Temporal Restrictions*: For the rider code combinations that may correspond to more than one Beacon name, we implement temporal restrictions in an attempt to obtain a unique mapping. We compare the first and last transaction dates (from the Transaction dataset) for a given product and set of rider codes with the Beacon introduction and closing dates. We eliminate a rider as a potential Beacon mapping if the first transaction date is before the introduction date or if the last transaction date is after the closing date. Note again that temporal restrictions are only used if there are multiple potential Beacon mappings.

After implementing the above restrictions, we obtain unique rider mappings for approximately 68% of contracts issued between 2008 and 2016.

Morningstar. Morningstar provides data on the subaccounts underlying annuity products, and we use a number of measures contained in Morningstar's data, including subaccount fees, investment styles, and the number of "high quality" funds, as measures of investment quality. We manually map annuity product names from Morningstar to the product descriptions provided in the Transaction dataset.

CRSP. CRSP provides returns net of expense ratios for each subaccount. We manually match fund names in the CRSP database with those provided in VA prospectuses (described in Section VI below). The fund names do change over time for the same fund, so we use CRSP's permanent fund number to aggregate historical returns for the fund. Finally, we use historical Fama-French factors from CRSP.

VA Prospectuses. For the NPV calculations, we rely on data obtained from VA prospectuses stored in the SEC's EDGAR database. We manually collect information on investment restrictions that contract holders must follow when they elect supplemental riders. Additionally, we obtain the number of accumulation units in the subaccounts for each product, which measure aggregate investment choices. We map this information to the transaction dataset using the Beacon product names and riders obtained through the process described in Appendix F.1.

FIA Rate Sheets. Historical data on formulas and rates from crediting strategies available in each FIA product come from *rate sheets*, which are issued monthly by FSP and distributed to advisers. While these rate sheets, unlike VA prospectuses, are not consistently filed in any publicly available database, we collect them through two means. First, some advisory firms have posted historical rate sheets for FIAs online, and we develop a large archive of such sheets through extensive web searches. Second, some states require FSP to file rate changes to FIA products with the state insurance agency. Through a series of Open Records Requests with the Texas Department of Insurance and the Florida Office of Insurance Regulation, we have collected further rate sheets to complete the historical database and corroborate the sheets obtained from advisory firms. As expected, since rates and crediting strategies do not depend on the state or the adviser who sells the product, rate sheets for the same month obtained through two different sources always agree.

F.2. Sample Selection

Note that the transactions dataset contains all transactions from 2008–15, and Discovery contains licensing information in 2015. To arrive at the final sample for analysis, we make a number of restrictions. First, we restrict to contracts sold in 2013–15 so that licensing and regulatory information is likely to be correct; this takes us to 248,103 transactions (from 689,454 annuity in the full dataset). Second, we keep transactions in which geography, (masked) adviser identity, and adviser type are all identified (234,135 observations after the restriction). These restrictions ensure that we know the fiduciary standard of the adviser who sells the product, if we can map to Discovery. Third, we drop all contracts sold in New York; there is substantially different financial regulation in that state—to the point where advisers in New York carry a different line of FSP products than those in other states. Indeed, most financial services providers sell a different suite of products in New York through advisers. We have 221,547 contracts after the restriction. Fourth, we restrict to firms and advisers with a record in Discovery, which takes us to 215,967 contracts. Fifth, we restrict to deferred annuities (variable and fixed indexed); we only drop about 1% of contracts (2,392) with this restriction, consistent with the fact that immediate annuities are especially rare in the United States. Finally, we restrict to contracts sold to individuals age 85 and younger, as variable annuities are not available to individuals over age 85; this drops 1,055 contracts.

After these restrictions, the sample contains 22,472 contracts sold in border counties, 19,730 of which are VAs. Nationwide, the sample consists of 212,520 contracts, 188,542 of which are VAs.

G. Robustness to Alternative Fiduciary Duty Classifications

G.1. Fiduciary Duty by State

Under state common law, any private party that has experienced some harm during a relationship with a financial advisor can bring a tort claim claiming breach of fiduciary duty. States differ in whether they recognize that a particular advisor has a fiduciary duty towards their client. In general, states agree that the fiduciary relationship must be assessed on a case by case basis. Some states, such as Massachusetts, have repeatedly stated that brokers in general owe no special duty to their clients past an ordinary duty of care. Others, including Texas, have been willing to find a fiduciary relationship in special circumstances, such as when the broker repeatedly recommends particular products or the broker's behavior would cause a reasonable person to believe they had an advisory role. Finally, some states, like Missouri, have held that stockbrokers generally do have a fiduciary duty towards their clients. In Missouri, for instance, fiduciary duties of brokers include "to manage the account as dictated by the client's needs and objectives, to inform of risks in particular investments, to refrain from self-dealing, to follow order instructions, to disclose any self-interest, to stay abreast of market changes, and to explain strategies." *State ex rel. PaineWebber, Inc. v. Voorhees*, 891 S.W.2d 126, 129 (1995).

Most states fall in the middle of this range of possibilities. These states do not specify the particularly duties placed on brokers, but instead rely on general principles of agency to guide their imposition of fiduciary duty on a subset of brokers. Compliance with these duties are therefore up to the interpretation of individual firms and their legal counsel. Firms may decide that compliance with

FINRA suitability largely satisfies their duties to their clients, or they may choose to supplement those requirements with more complete disclosure or higher standards for recommendations.[9] Note that taking fees on an assets under management basis is considered as evidence of fiduciary duty being satisfied, but is not necessary or sufficient to satisfy a fiduciary duty. Furthermore, there are only a handful of opinions in each state that evaluate the specifics of the fiduciary relationship, and even fewer of these that result in a monetary disposition of the case. Brokers have a strong incentive to settle these cases, so this may not reflect the efficacy of the common law standard. On the other hand, several cases seem to include the language that brokers are subject to fiduciary duty in dicta, meaning that the statement may not be reflected by other parties adopting that standard as binding. Finally, state laws are most likely effective in regulating only state registered brokers, since federally registered brokers can be more effectively sued for breach of federal duties in federal court.

Finke and Langdon (2012) have classified states according to their common law cases into three categories: those with fiduciary standards on broker-dealers, those with quasi-fiduciary standards, and those without a fiduciary standard for broker-dealers. We validate this categorization as follows.

First, we restrict our attention to state appellate court opinions that mention fiduciary relationships. Within that case law, we search for cases discussing the application of fiduciary duty to broker-dealers handling non-discretionary accounts. If there is a case with unequivocal language, either extending fiduciary duty to most such transactions or denying the possibility of a fiduciary relationship being found in an arms-length transaction, the state is categorized as within the first or third group. If there is equivocal language or there is no case addressing the application of fiduciary principles to brokers, we move to the second stage.

In these states, we read through any cases that describe the requirements needed for a finding of a fiduciary relationship, focusing on cases involving clients and brokers in arms length relationships. We exclude cases where there is a statutory fiduciary duty or with an unusually close relationship such as family and business partners. These cases usually arise between debtors and creditors, or with clients of real estate and insurance brokers. If a case exists that shows that the state's court is willing to expand the reach of fiduciary duty when clients face losses due to seller's poor guidance, the state is coded as a quasi-fiduciary state. If no cases exist in a state or if the cases define fiduciary duty very narrowly, the state is classified as having no fiduciary duty on broker-dealers.

Following this procedure, we are able to largely replicate Finke and Langdon (2012)'s classification of states. It is important to note that this classification is missing two key features. First, the classification ignores federal cases where state law is applied to the fiduciary duty question. A brief look at these excluded cases shows that federal courts tend to heighten the duties placed on broker

[9] Note that FINRA suitability tends to require full documentation about the client rather than about the interests of the broker. Complying with fiduciary duty is likely to involve more significant disclosures on the side of the adviser. In some ways, therefore, they are separate requirements and can be layered on top of each other.

dealers. Second, the classification doesn't account for the fact that advisors in our sample are often registered insurance producers, subject to heightened duties under state insurance law. Since both of these omissions would underestimate the strength of the duty placed on brokers, we assume that any evidence of fiduciary duty placed on broker dealers qualifies a state as imposing a fiduciary duty, pooling the fiduciary and quasi-fiduciary states under our fiduciary classification.

For robustness, we take an alternate stance on the decisions above and generate a "modified" classification of states by common law fiduciary duty. This classification accounts for the challenges laid out above and codes states as imposing a fiduciary duty anytime a state or federal case mentions that brokers have a fiduciary duty placed on them, excluding fiduciary duties placed on analogous principal-agent relationships. This classification accounts for agency relationships' duties diverging significantly from each other (eg. stockbrokers facing different duties than real estate brokers). Moreover, this accounts for the increasing role of federal cases in interpreting state law. States where we find a different outcome have new cases that signal a change in the court's attitude. Our classification differs from Finke and Langdon's in the following states: Maine, Nebraska, New Jersey, Rhode Island, Vermont, Wyoming. In each of these states, Finke and Langdon find a quasi-fiduciary duty on brokers, but our research have not uncovered cases that are directly analogous to the retail investor/financial advisor relationship. In sum, this classification effectively refines the decisions by Finke and Langdon that lead to the quasi-fiduciary category.

Tables G.1 and G.2 show the results of the modified classification. The results are very similar to the original classification, with some effects being stronger. Risk adjusted returns increase by 33 bp in states with fiduciary duties on broker dealers, while raw returns increase by 54 bp. These increases are larger and more precisely estimated than with the original classification, consistent with the theory that Finke and Langdon's classification measures the true state common law stringency with some measurement error.

The results are driven by largely similar changes in product characteristics. The probability of selling a variable annuity, relative to a fixed indexed annuity, drops by 14% with the revised classification relative to 12% in the baseline classification, and the lowest 10th percentile return on a product sold increases more significantly with the revised classification. Minimum expense ratios decrease, just as in the original classification, but average expense ratios do not increase as much, with the coefficient not reaching statistical significance. Fund returns within variable annuities sold increase regardless of how investments are allocated across funds. In addition, the modified classification shows that mortality and expense ratios on products sold drop significantly, by nearly 7%.

As robustness checks, we consider whether the results are largely the same under two other potential definitions of a fiduciary standard:

1. Continuous classification

Table G.1: Returns on variable annuity products using modified classification

	(1) Risk Adjusted Returns	(2) Unadjusted Returns
DID	0.0033***	0.0054***
	(0.0011)	(0.0019)
FD on BD	0.0030***	0.0042***
	(0.0008)	(0.0017)
FD on RIA	-0.0002	-0.0012*
	(0.0007)	(0.0006)
Mean of Dep. Var	0.028	0.063
N	32,115	32,115

Annualized returns for variable annuities sold. Contracts are restricted to borders, specifications include border fixed, contract month, and age fixed effects. Standard errors are clustered at the state. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

This classification uses the Finke and Langdon (2012) classification and quantifies the strength of each category, assigning value 0 to non-fiduciary duty states, .5 to quasi-fiduciary states, 1 to fiduciary states where the relevant court opinions appear to mention the application fiduciary duty to brokers in that state but not lay out detailed guidance on compliance to future parties, and 2 to fiduciary states with at least one state court opinion providing compliance guidance. This continuous metric is intended to reflect the *strength* of the fiduciary duty placed on advisors within that state, particularly in its stringency being significantly beyond FINRA suitability.

2. Full fiduciary only classification

   This conservative metric codes states Finke and Langdon (2012) designate as "full fiduciary" states as including a fiduciary duty, while quasi-fiduciary states are classified as not being subject to any heightened duty beyond FINRA suitability. This robustness check allows for the possibility that quasi-fiduciary states rarely impose duties beyond suitability on advisor-client relationships, or that enforcement is more lax in quasi-fiduciary states. By classifying quasi-fiduciary states as imposing "no duty," this robustness check specifically estimates the marginal effect on advisor behavior of broader language regarding fiduciary duties in state court opinions. If this language itself is not the only mechanism through which standards are imposed, we would expect an underestimate of the true effect of fiduciary duty.

In each robustness check, we replace the Finke and Langdon classification with the alternative classification. Each classification is associated with a different sample, because the analysis includes only border counties where the neighboring state has a different standard imposed on BDs. The results are qualitatively similar to the baseline classification, but there are some significant differences. The continuous classification results in very similar outcomes for the difference-in-difference in term of magnitude, but some outcomes are no longer statistically significant. The

Table G.2: Characteristics of products transacted using modified classification

			Expense Ratio		Fund Returns		Fees	
	$\mathbb{1}$[VA] (1)	10^{th} Perc. (2)	Minimum (3)	Average (4)	Optimal (5)	Equal (6)	M&E (7)	Surr. Chg. (8)
DID	-0.123***	2.078***	-0.006*	0.048	0.0191*	0.0028**	-0.075**	0.167
	(0.033)	(0.676)	(0.003)	(0.031)	(0.0094)	(0.0011)	(0.030)	(0.139)
FD on BD	-0.129**	1.814***	-0.008**	0.066**	0.0187*	0.0027**	-0.080**	0.098
	(0.029)	(0.514)	(0.003)	(0.025)	(0.0079)	(0.0010)	(0.030)	(0.150)
FD on RIA	0.006	-0.264	-0.002	0.018	-0.0004	-0.0001	-0.004	-0.069
	(0.026)	(0.283)	(0.002)	(0.013)	(0.0032)	(0.0005)	(0.015)	(0.067)
Base Mean	0.872	2.827	0.501	1.267	0.159	0.012	1.124	3.209
N	32,115	32,115	27,998	27,998	27,998	27,998	27,998	27,998

	# Funds			# Equity Styles		# FI Styles	
	All (9)	≥ 4 Stars (10)	≤ 2 Stars (11)	High Q. (12)	Only Low Q. (13)	High Q. (14)	Only Low Q. (15)
DID	9.21*	4.35**	2.07	0.837**	-0.525*	0.293**	-0.116*
	(5.34)	(1.66)	(2.32)	(0.332)	(0.246)	(0.134)	(0.058)
FD on BD	11.15***	4.75**	2.90	0.931***	-0.607**	0.270**	-0.142***
	(4.37)	(1.50)	(1.96)	(0.293)	(0.209)	(0.125)	(0.055)
FD on RIA	1.94	0.39	0.83	0.094	-0.083	-0.024	-0.026*
	(2.26)	(0.74)	(1.16)	(0.145)	(0.118)	(0.065)	(0.015)
Base Mean	97.49	30.88	32.77	7.067	0.911	4.250	3.040
N	27,998	27,998	27,998	27,998	27,998	27,998	27,998

Estimates of (1) for various product characteristics. Columns (1) and (2) use the set of all annuities transacted in the border, while the other columns restrict to variable annuities. Standard errors are clustered at the state level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

loss of statistical significance signifies that the results are not driven heavily by those states which provide written guidance for future parties on how to comply with state fiduciary duties. This is reasonable given that states often cite each other, and parties facing a common law fiduciary duty can look to other states to guide their compliance.

The second robustness check uses only states bordering those which Finke and Langdon classify as "full fiduciary" states. These results look qualitatively similar to the baseline classification in sign, but magnitudes are approximately 1/3 as large and many effects are not statistically significant. The results suggest that financial advisor behavior is impacted by the threat of courts scrutinizing the particulars of a transaction to determine whether a fiduciary duty exists, as well as by broadly imposed fiduciary duties on all broker-dealers without scrutiny of a particular set of facts. This is consistent with our evidence that advisors comply with fiduciary duty by trying to avoid litigation altogether. Of course, to the extent that this classification groups many states that effectively impose strong duties on broker-dealers with states that do not, we would expect an underestimate of the effect.

References

BULOW, J. I., J. D. GEANAKOPLOS, AND P. D. KLEMPERER (1985): "Multimarket Oligopoly: Strategic Substitutes and Complements," *Journal of Political Economy*, 93, 488–511.

CHEN, X., T. CHRISTENSEN, AND E. TAMER (2019): "Monte Carlo Confidence Sets for Identified Sets," *Econometrica*, 86, 1965–2018.

DEL MORAL, P., A. DOUCET, AND A. JASRA (2006): "Sequential Monte Carlo Samplers," *Journal of the Royal Statistical Society B*, 68, 411–436.

FINKE, M. AND T. LANGDON (2012): "The Impact of the Broker-Dealer Fiduciary Standard on Financial Advice," Working Paper.