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. $\square$

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 $\mu'_m < \mu_m$. $\square$

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. $\square$

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 $f_T\left(\cdot;\cdot\right)$ 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 varianbles. 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 (230.00)	-105.45 (97.68)	134.03	35.66 (42.48)	28.46 (26.25)	102.55
Median Age	-0.33 (0.80)	0.29 (0.45)	40.69	-0.57 (0.87)	-0.60 (0.43)	41.37
Pop Black (K)	27.37 (38.16)	-17.52 (25.16)	16.17	7.72 (5.04)	7.13** (2.92)	12.57
Pop Hispanic (K)	130.82 (97.45)	0.31 (20.29)	21.96	15.85 (14.57)	12.83 (9.84)	16.48
Median HH Income (K)	0.06 (6.11)	0.70 (1.97)	45.74	1.99 (2.61)	1.23* (0.68)	44.45
Mean HH Income (K)	-1.36 (7.65)	-1.00 (2.88)	59.97	2.26 (3.04)	1.28 (0.86)	58.38
Pct. Unemployment	0.61 (0.81)	-0.55*** (0.20)	9.32	-0.16 (1.06)	-0.08 (0.31)	9.30
Pct. Poverty	-0.17 (1.81)	-1.00 (0.71)	17.34	-0.68 (1.67)	-0.36 (0.50)	17.72
Pct. HH with less than $25k	-0.89 (2.09)	-1.18 (1.11)	28.38	-0.99 (1.96)	-0.52 (0.52)	29.14
Pct. HH with less than $50k	-0.94 (4.10)	-1.33 (1.49)	54.86	-1.82 (2.40)	-1.10* (0.64)	56.11
Pct. HH with less than $75k	-0.28 (4.66)	-0.56 (1.48)	73.15	-1.52 (2.09)	-0.77 (0.61)	74.31
Pct. HH with less than $100k	0.29 (4.26)	0.03 (1.33)	84.46	-1.26 (1.56)	-0.68 (0.48)	85.45
Pct. Pop less than HS	1.53 (1.45)	-0.44 (0.62)	14.50	-0.03 (1.61)	0.36 (0.39)	14.97
Pct. Pop HS	2.31** (0.87)	1.81** (0.87)	32.85	1.66 (1.39)	1.73*** (0.52)	33.68
Pct. Pop BA or Higher	-4.19 (3.07)	-1.99 (1.42)	19.75	-0.35 (1.64)	-0.71 (0.57)	18.65

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

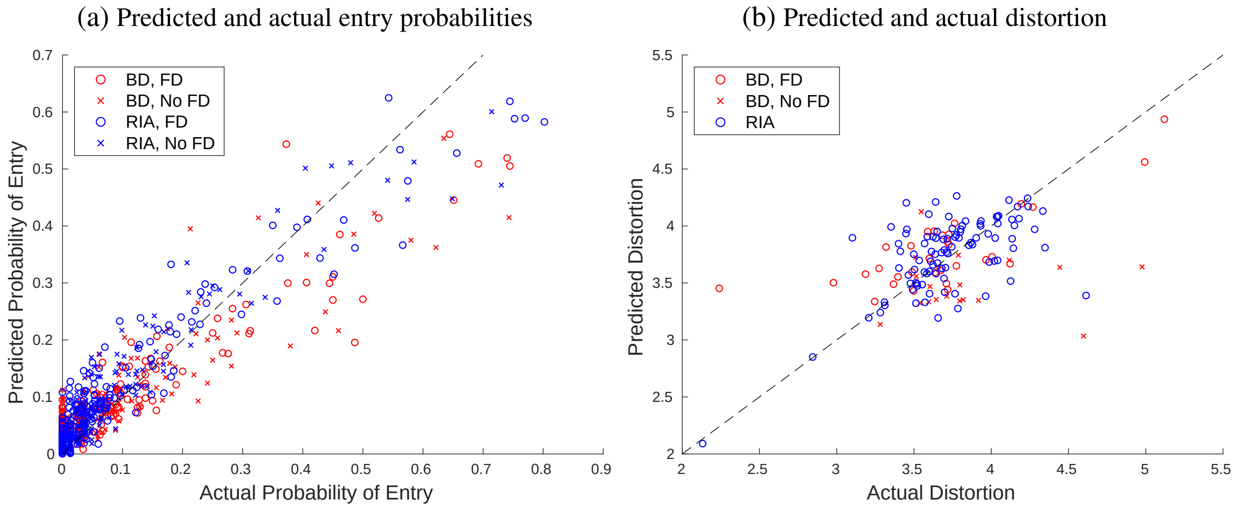

region of $[7.45\%, 11.99\%]$. For DRs, the average observed entry probability is 8.77%, and the prediction of the model is 10.93%, with a 95% confidence region of $[8.74\%, 13.20\%]$.

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 $G\left(a^{first}\right) : \{a_0, ..., \bar{a}\} \to \{g_1, ..., g_G\}$ 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

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

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

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

subject to

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

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, $V^*_{\bar{t}}(0, b_{\bar{t}}, g_{\bar{t}}) = g_{\bar{t}} \cdot b_{\bar{t}} \cdot \left(1 + \sum_{\tau=\bar{t}+1}^{T} \delta^{\tau-\bar{t}} \prod_{\tau'=\bar{t}+1}^{\tau} \mu_{\tau'}\right)$.

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

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

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

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

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

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

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 $c \in \{1, \dots, C\}$.

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

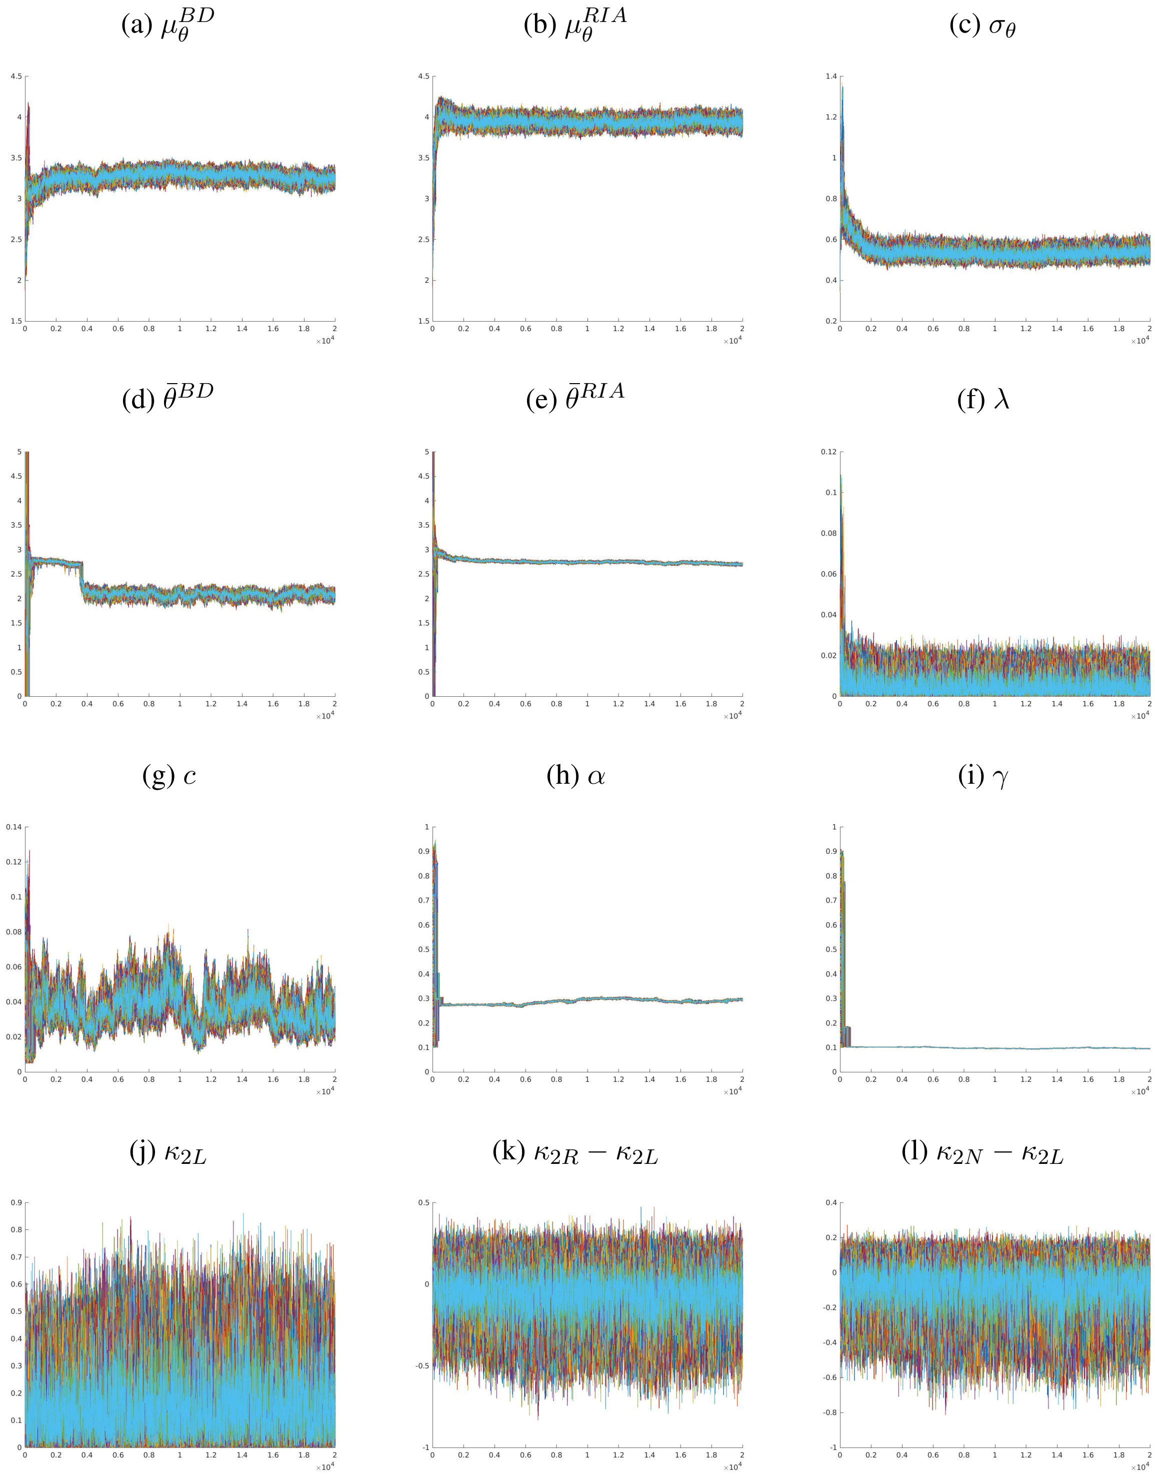

(a) μ_θ^{BD} (b) μ_θ^{RIA} (c) σ_θ

(d) $\bar{\theta}^{BD}$ (e) $\bar{\theta}^{RIA}$ (f) λ

(g) c (h) α (i) γ

(j) κ_{2L} (k) $\kappa_{2R} - \kappa_{2L}$ (l) $\kappa_{2N} - \kappa_{2L}$

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.