

Signaling, Random Assignment, and Causal Effect Estimation*

Christopher A. Hennessy
LBS, CEPR, and ECGI.

Gilles Chemla
Imperial College, DRM/CNRS, and CEPR.

March 2021

Abstract

Evidence from quasi-random assignment (e.g. natural experiments, IV, and RDD) has been labeled “the most credible.” We argue such causal evidence is often misleading in finance/economics, omitting a key component of the true empirical causal effect. Random assignment, in eliminating self-selection, simultaneously precludes signaling via treatment choice. However, outside experiments, agents enjoy discretion to signal, thereby *causing* changes in beliefs and outcomes. Therefore, if the goal is informing discretionary decisions, rather than predicting outcomes after forced/mistaken actions, randomization is problematic. As shown, signaling amplifies, attenuates, or reverses signs of causal effects. Thus, traditional empirical finance methods, e.g. event studies, are often more credible/useful.

1 Introduction

It is hard to overstate the influence of Angrist and Pischke’s (2009) *Mostly Harmless Econometrics* on modern-day empirical finance. Even the language of the profession has shifted, as evidenced by frequent use of the terms “identification” and “causal effect.” As one indicator, Bowen, Frésard, and Taillard (2017) find that in the top-three finance journals, the share of empirical corporate finance papers using what they term “identification technologies” rose from roughly 0 percent in the late 1980’s to over 50 percent by 2012. As an example of the converse, consider that in 1986 the *Journal of Financial Economics* devoted one-half of a double issue to five event studies analyzing endogenous

*Contact chennessy@london.edu (corresponding author) and g.chemla@imperial.ac.uk. We thank seminar participants at Université Paris Dauphine, Mannheim, Imperial College, and ESCP Europe.

capital structure changes, an unlikely journal configuration given contemporary methodological norms.

The crowding-out of such quasi-observational studies is understandable, and indeed a good thing for the progress of finance as a science *provided* one accepts the key initial premise put forward by Angrist and Pischke (2009) that, “The goal of most empirical research is to overcome selection bias, and therefore to have something to say about the causal effect of a variable.” If their causal effect concept is accepted, then it must be agreed that, as they assert, “The most credible and influential research designs use random assignment.” Apparently, traditional observational studies, especially those making no attempt to “overcome selection”, e.g. the path-breaking event study by Fama, Fisher, Jensen and Roll (1969), are somehow lacking in “credibility.”

The methodological stance of Angrist and Pischke, amongst others, draws much of its inspiration from the notion that economists should strive to utilize the same causal measures and methods as medical science, an appealing metaphor at face value. For example, Duflo (2004) argues, “Creating a culture in which rigorous randomized evaluations are promoted, encouraged, and financed has the potential to revolutionize social policy during the 21st century, just as randomized trials revolutionized medicine during the 20th.” Indeed, the two textbooks by Angrist and Pischke (2009, 2015) open with motivating examples of causal effects of hospitalization and insurance on health outcomes. Similarly, the influential textbook of Imbens and Rubin (2015) uses the same causal effect definition and methodological tool-kit across biomedical and social sciences.

The objective of this paper is to explore the merits and limitations of the causal effect definition and estimators as advocated by Angrist and Pischke (2009, 2015), amongst others, within the specific context of financial economics. To begin, we note that, to the best of our knowledge, absent from the methodological discussions amongst financial economists is the fact that medical outcome variables differ in-kind from most of the financial outcome variables we study. In general, pure physical responses to medical treatments (e.g. pills and surgeries) are not mediated by the beliefs of other agents.¹ In stark contrast, for the majority of economic and financial variables, equilibrium

¹Health may be mediated by own-beliefs. Philipson and Desimone (1997) and Chemla and Hennessy (2019) show

outcomes are heavily influenced by the beliefs of other market participants.

Should financial economists strive to attain the same type of evidence as medical scientists? Alternatively, are the qualitative differences between health outcome variables and economic outcome variables sufficiently important for financial economists to reconsider how we *define and measure* causal effect? In order to address these important questions, or at least initiate a more transparent and constructive dialogue on these questions, this paper analyzes the meaning and utilization of alternative causal effect measures in the types of settings commonplace in finance and economics, those where individual agents have private information and outcome variables are mediated by the beliefs of other agents in the economy.

As we show, the narrow causal effect definition utilized by Angrist and Pischke (2009, 2015), while appropriate for health outcomes, is problematic from the perspective of financial economists. This is because random assignment, in eliminating self-selection, simultaneously prevents signaling via treatment choice. After all, if a treatment is randomly forced upon an agent, the act of taking it cannot signal private information. In causal biomedical research, shutting down the signaling channel through random assignment is of no concern since an individual cannot cause their health quality to improve by publicly and voluntarily taking a particular pill that signals something to onlookers. In sharp contrast, in financial markets, individuals and firms can and do improve their outcomes by publicly and voluntarily taking particular actions.

If the goal is guiding discretionary decisions, rather than predicting outcomes after forced actions or mistakes, random assignment is problematic precisely because it strips out the signaling component of the overall empirical causal chain. As we show, taken in isolation, estimates derived from random assignment are often faulty guides regarding *discretionary* policies—the very types of policies that empirical evidence is often intended to inform. In fact, as shown, causal effects derived from random assignment can undershoot, overshoot, and even have signs opposite to causal effects cum signaling effects.

We suggest utilization of two distinct causal effect definitions. *Partial causal effects* are to be

RCTs may deliver biased estimates of physiological effects due to such placebo effects.

understood as changes in the outcome variable arising from changes in the forcing variable holding fixed the beliefs of other agents. This type of causal effect is recovered under random assignment. *Total causal effects* are changes in the outcome variable resulting from discretionary changes in the forcing variable allowing for endogenous equilibrium changes in beliefs and payoffs. The total causal effect can be viewed as the sum of the partial causal effect and signaling effect. This causal effect measure is recovered by observational evidence such as event studies, and perhaps by ordinary least-squares in some instances. As we show, the total causal effect is often a sufficient statistic for optimal discretionary corporate financial decisions.

We clarify the issues by way of three examples. To begin, we revisit Molina (2005), who examines the effect of debt on observed bond ratings with firm tax rates serving as an instrument for debt. Importantly, Molina finds that the effect of leverage on ratings is three times stronger under IV estimation than if debt endogeneity is ignored. In the spirit of Molina, we consider a parable economy in which firms trade-off tax benefits of debt against reputational costs of bankruptcy, with firm quality being a latent variable unobservable to rating agencies, investors, and the econometrician. Here tax rates represent an ideal instrument, and the IV coefficient is large while the OLS coefficient is zero. Nevertheless, in this parable economy, the OLS coefficient actually captures the relevant causal effect from the perspective of CFOs making leverage choices: In equilibrium, ratings are invariant to discretionary changes in leverage conditional upon observables. Intuitively, the very same latent variables problem that troubles the econometrician, implies that debt serves as a credible signal of latent quality to outsiders, including rating agencies. Phrased differently, the very same endogeneity effect that the econometrician struggles to eliminate via instrumentation is actually part of the overall causal chain that CFOs must account for if they are to make optimal decisions.

In the second example, we revisit the recent paper by Dittmar, Duchin and Zhang (DDZ below) (2020) who claim to “provide some of the cleanest estimates, to date, of the timing and causal effects of SEOs.” In particular, we develop a simple parable economy that mimics key aspects of the fuzzy regression discontinuity design (RDD) they employ. In this economy, firms are ex ante identical and

start life at the same stock price. At an interim date, firms are hit with an observable idiosyncratic shock to their investment opportunity set, with the shock having the unit interval as its support. Consistent with the empirical evidence provided by DDZ, we hard-wire those firms with a shock above (below) one-half to receive board approval for some type of SEO with a relatively high (low) probability. In this economy, evidence from the fuzzy RDD is consistent with the notion that equity issuance has a positive effect on stock price. In particular, the interim-date stock price exhibits a discontinuous upward jump at the threshold where boards approve some SEO with relatively high probability. Further, stock prices react favorably if the board randomly approves some type of SEO. Finally, as in DDZ, the stock price reaction for board SEO approval is higher for firms above the cutoff. However, in this parable economy, all of the preceding evidence is of limited practical use to the ultimate decision-makers, the CFOs and investment bankers who must decide on the actual number of shares to issue. In particular, these decision-maker still confront a downward sloping stock price reaction function conditional upon observable investment fundamentals, due to the negative signal greater stock issuance sends about the value of assets in place.

In the final example, we revisit the influential argument of Romer and Romer (2010) who argue for using legislative histories and the like in order to isolate quasi-random changes in fiscal policy. In this spirit, we consider a government with private information about economic fundamentals that is contemplating changing some policy, e.g. regulation or tax policy, in response to this information. Here we show that evidence from random assignment can identify potentially stimulative policies—in the sense of having positive partial causal effects. But the actual impact of discretionary government policy is correctly measured by the total causal effect. That is, the way to ultimately estimate the response to discretionary policy is to observe the response to discretionary policy. We conclude that in some applied settings both forms of causal effect estimates have a place in the decision-maker’s tool-kit.

The present paper shares with Gomes (2001), Alti (2003), and Moyen (2004) the idea of using variations on canonical models to shed light on the meaning and interpretation of empirical evidence.

However, none of these papers comments on random assignment, causal effect measures, or the signaling channel. Keane (2010) and Rust (2010) argue for the need to isolate distinct causal channels and identify deep technological parameters, recommending the use of structural models for this purpose. Instead, our paper discusses how different forms of reduced-form evidence can be used to measure partial and total causal effects. There is nothing in our argument that has direct bearing on the structural versus reduced-form debate.

The central argument in our paper is related to a forthcoming paper by Fudenberg and Levine (2020). They show that Bayesian learning by agents can drive a wedge between partial causal effects and causal estimates derived from regression discontinuity designs. This is because uninformed agents on opposing sides of regression discontinuity boundaries endogenously form sharply different beliefs about effort returns. The central difference between the papers concerns the methodological message. In particular, Fudenberg and Levine (2020) do not challenge the primacy of partial causal effects, but rather show that regression discontinuity may fail to recover them. In contrast, we show that when a privately informed agent moves first, partial causal effects are often a faulty basis for decision-making.

Our paper is related to the literature on signaling, with Spence (1973) being the pioneering paper. Ross (1977) and Leland and Pyle (1977) were early applications of signaling theory to corporate finance. For general surveys, see Riley (2001) and Löfgren, Persson, and Weibull (2002). Our paper complements this literature by flushing out the implications of signaling for applied econometric work that seeks to inform decision-making by individuals, firms, and governments.

The remainder of the paper proceeds as follows. Section 2 revisits the data in Bowen, Frésard, and Taillard (2017) in order to give the reader a rough sense of the applicability of our argument within empirical corporate finance research using some of the most popular identification strategies. Section 3 presents the instrumental variables example and Section 4 presents the fuzzy RDD example. Finally, Section 5 considers a more complex setting that illustrates the complementary roles the alternative causal effect estimators can play in government policy setting.

2 Revisiting Bowen, Frésard and Taillard (2017)

In the final sections of this paper we present detailed examples of how our critique applies to specific papers. We do not intend to single out any particular papers, but have simply chosen individual papers because they were perhaps the easiest to capture in simple tractable models. But this then might raise the question of just how widely applicable is our critique. The objective of this section is to provide the reader with a general sense of the scope of the problem. Of course, such assessments are necessarily subjective. Thus, greater detail on each of the 253 papers we read, and the basis for classifications is provided in the Online Appendix.

Our first task was to assemble a reasonably comprehensive data set, preferably one compiled by other researchers in order to mitigate concerns about cherry picking. Conveniently, Bowen, Frésard and Taillard (BFT below) (2017) were kind enough to provide us with their data.² Importantly, BFT offer perhaps the most comprehensive analysis of the development of identification strategies in the empirical corporate finance literature over recent decades.

BFT rely on keyword searches in order to pin down empirical corporate finance papers mentioning/using some form of “identification strategy.” They document that within the top three finance journals, the share of empirical corporate finance papers using at least one identification strategy increased from 0% in 1980 to 10% in 2000 before skyrocketing to over 50% in 2012. For this reason, we focused our analysis on the more recent subset of the BFT sample papers published in the top three finance journals during the period from 2000 to 2012 using three identification strategies: instrumental variables (IV), regression discontinuity design, and controlled experiments. As discussed below, IV is by far the most popular of these three methods.

We excluded from our analysis the papers in the BFT sample that rely upon difference-in-differences (DiD) estimation. This is because, in our view, our critique does not generally apply to DiD estimation as it is implemented in practice. After all, most authors applying DiD rely primarily on the assumption that the treatment and control groups would have had parallel trends had it not

²We thank these authors for the generosity.

been for the treatment. Although this parallel trends assumption could potentially be defended by invoking some notion that the treatment arrived randomly, random assignment is not a necessary condition for the parallel trends assumption to be satisfied. Rather, random assignment can in some instances be viewed as akin to a sufficient condition for satisfaction of the parallel trends assumption in DiD estimation. Further, most DiD papers study policy changes, and it is hard to take seriously most claims that a deliberate policy change arrived randomly.

As shown below in Table 1, the BFT sample contains 253 empirical corporate finance papers in the top three finance journals from 2000-2012 using/mentioning either IV, RDD or an experiment: 243 used IV, 10 used RDD (including three using both RDD and IV), and 3 used experiments. After this first cut, we set about reading each of the 253 papers to come up with a determination as to whether the respective paper should be included in our final sample, and whether our critique applied.

In order to obtain our final sample, we eliminated those papers that placed little/no emphasis on the respective identification strategy. For example, if a paper was included in the BFT sample as employing IV, but then upon reading the paper it was discovered that IV was only mentioned in a footnote, the paper was excluded. As another common example, if a paper mentioned that “results are robust to instrumentation” but the robustness checks were not even tabulated, the paper was excluded from our sample. We also eliminated from our sample those papers that did not set about measuring any causal effect but instead set about estimating a structural parameter via GMM or structural model moment matching, for example. We also eliminated from the sample papers that performed estimation using simulated rather than real data. After carefully reading the individual papers in the BFT sample, we arrived at a final sample of 185 out of the 253 BFT papers (73%) that either stress or rely heavily upon IV, RDD or Experiments.

Our next task was to determine whether our critique was applicable to each individual paper. This determination depended upon the nature of the causal variable studied and the respective dependent variable. In particular, in order for our critique to apply, a theoretically plausible signaling

channel would need to be operative—with the identification strategy eliminating from the causal effect measurement that component of the causal chain arising from signaling effects.

In order for signaling to be operative, the causal variable would need to be the choice variable of an individual/unitary decision-maker, such as a CFO or board of directors. For example, leverage is a choice of the CFO, and this can signal firm quality, and executive compensation is chosen by a board, and this can signal their view of executive quality. By way of contrast, causal variables such as an industry Herfindahl index, network density, number of analysts covering a firm, or cash flow shocks are not unilateral decisions that convey signals.

Further, in order for signaling to be operative, one or more of the dependent variables studied must be mediated by counterparty beliefs. For example, measures of credit spreads, credit ratings, debt values, equity values, firm enterprise values, and market-to-book ratios are directly mediated by market beliefs. In some instances real outcomes will also be mediated by counterparty beliefs, such as the beliefs of consumers or suppliers having an effect on corporate earnings or other performance measures.

As detailed in Table 1, using these criteria, we assessed that our critique applied to 130 out of the 185 papers in the final sample, or 70%. While one might disagree with the final tally, the takeaway point here is that our critique applies to a large number of papers in the empirical corporate finance literature. And this is hardly surprising. After all, asymmetric information and signaling motives are ubiquitous in financial markets. Moreover, many commonly-studied outcome variables, e.g. valuations, prices, and returns, are mediated by market beliefs.

Of course, some caveats are in order. First, just how strong the signaling effect is likely to be in a given context is a subjective judgment. Second, to say that our critique applies should not be construed as implying that the respective papers do something wrong. Rather, we simply argue that the papers subject to our critique focus on partial causal effects whereas total causal effects may well be of greater operational importance for real-world decision-makers such as CFOs. Nevertheless, our argument does cast doubt on a reflexive assumption that so-called “cleaner” identification methods

are better identification methods.

3 Instrumental Variables

This section considers the interpretation and utilization of IV estimates. To this end, we revisit Molina (2005) who examines the effect of leverage on bond ratings. Importantly, Molina questions whether existing estimates are biased downwards due to the endogeneity of leverage choice. Thus, Molina uses firm tax rates as an instrument for debt. Importantly, Molina finds that the effect of leverage on ratings is three times stronger under IV estimation than if debt endogeneity is ignored.

With this in mind, consider the following economy. There are two dates $t \in \{0, 1\}$. All agents are risk-neutral, and there is no discounting. At $t = 0$ there is a large finite number of private equity (PE) investors who will be packaging up and selling off claims to the future cash flows of the respective firms they currently own. They can package cash flows as equity or zero coupon bonds. The bond face value chosen for company j is denoted $B_j \geq 0$. As in Gourio (2013), packaging cash flow as debt is assumed to be tax advantaged, with the government providing firm j with an up-front tax rebate at time $t = 0$ equal to $\tau_j B_j$.

Each firm's debt tax shield parameter τ_j is observable to all agents in the economy, including the econometrician. Each debt tax shield parameter τ_j represents the realization of an i.i.d. random variable $\tilde{\tau} \in \{\tau^l, \tau^m, \tau^h\}$ where $0 < \tau^l < \tau^m < \tau^h < 1$. In reality, firms may have different debt tax shield values due to different non-debt tax shields and/or differential exposures to state and international taxation. For simplicity, we capture such effects in reduced-form, so as to make the instrumentation strategy ideal.

The after-tax cash flow (\tilde{c}) accruing to firm j at $t = 1$ is a random drawn from a uniform distribution with support $[0, \Theta_j]$. The upper bound on each firm's cash flow Θ_j represents the realization of an independent random variable with known support $[0, \bar{\Theta}]$. If realized cash flow at $t = 1$ is insufficient to repay B , the firm's original private equity sponsor will incur a reputational cost LB , with the loss parameter $L > 0$ being public information and homogeneous. In the spirit

of Ross (1977) and Leland (1994), the objective of the private equity shop in packaging up cash flows, that is in choosing B_j , is to maximize the value of marketable claims on the firm less expected reputational costs arising from bankruptcy.

The econometrician is interested in empirically estimating the relationship between debt levels and credit ratings as assigned by “the rating agencies.” To streamline exposition, assume the credit rating scale is continuous, with rating agencies delivering their best estimate of each firm’s log default probability. Let the observed rating (log default probability) for firm j be denoted by Λ_j .

Given her understanding of the technologies in this economy, including terminal cash flows being drawn from uniform distributions, the econometrician considers that a reasonable empirical specification for log default probability is:

$$\ln [\Pr (\tilde{c} \leq B_j)] = \ln \left[\int_0^{B_j} \frac{1}{\Theta_j} dc \right] = \underbrace{\beta_1}_{\beta_1=1} \times \ln B_j + \underbrace{\beta_2}_{\beta_2=-1} \times \ln \Theta_j. \quad (1)$$

Firm debt levels (B_j) and credit ratings (Λ_j) are public information, but, unfortunately for the econometrician, firm quality (Θ_j) is not. Of particular concern for the econometrician is her understanding that firms’ private equity sponsors have the ability to condition their choice of debt levels on their private knowledge of the quality (Θ_j) of the firms they have been owning and operating. This leads the econometrician to propose instrumenting $\ln B_j$ with the debt tax shield parameter τ_j .

To motivate her chosen instrument, the econometrician presents the following stylized model of firm capital structure, in the spirit of Leland (1994). She considers that observed capital structures are represented by

$$\begin{aligned} B_j^* &\in \arg \max_B \frac{\Theta_j}{2} + \tau_j B - LB \times \underbrace{\Theta_j^{-1} B}_{=\Pr(\tilde{c} \leq B_j)}. \\ \Rightarrow B_j^* &= \left(\frac{\tau_j}{2L} \right) \Theta_j. \end{aligned} \quad (2)$$

That is, the econometrician considers private equity shops to be maximizing their expectation of cash flow plus tax shield value less reputational costs of bankruptcy. The implied first-order condition

equates marginal tax benefits of debt with marginal reputational costs of bankruptcy:

$$\tau_j = 2LB_j^*\Theta_j^{-1} \Rightarrow B_j^* = \frac{\tau_j\Theta_j}{2L}.$$

By construction, the tax shield parameter represents an ideal instrument here. After all, firms with higher tax shield parameters can be expected to take on more debt. Moreover, the exclusion restriction is satisfied since the debt tax shield parameter only changes default risk through its effect on B_j . Finally, the debt tax shield parameter is randomly-assigned.

We consider a simple parameterized economy broadly consistent with Molina (2005). As in Molina, the sample size is 2678 simulated firm observations and the mean value of the debt tax shield parameter is 33%, with three potential values: $\tau^l=3\%$, $\tau^m=33\%$ and $\tau^h=43\%$. The first two tax rates have probability 20% and the latter has probability 60%. The bankruptcy reputational loss is set to $L=1/2$. Finally, the firm quality parameter is uniformly distributed on $[0, 1/2]$.

Panel A of Table 2 presents regression output when the econometrician employs OLS and regresses credit ratings Λ_j directly on $\ln B_j$, with no attempt at instrumentation. As shown, the slope coefficient is negligible. The econometrician argues that, “Naive reliance on OLS output here would lead one to conclude that a CFO’s decision to increase his debt level will have no effect on his firm’s bond rating. But the OLS output is likely subject to endogeneity bias.”

Panel B of Table 2 presents the econometrician’s preferred regression output, that relying on the tax shield parameter as an instrument for $\ln B_j$. A number of points are worthy of note. To begin, the first-stage regression results are broadly consistent with the econometrician’s formulation of the private equity shop’s decision problem (equation (2)), with the tax shield parameter apparently having a strong effect on observed debt levels. Second, when observed bond ratings Λ_j are regressed on instrumented $\ln B_j$, the slope coefficient is near unity, consistent with the econometrician’s empirical specification of default probability in equation (1). The econometrician is thus likely to conclude that “increases in leverage will actually lead to sharp deterioration in bond ratings, an effect captured by IV regression, but not OLS.”

But what will be the actual effect on bond rating should a CFO choose to issue more debt?

Figure 1 provides the answer. As shown, there are three distinct horizontal lines capturing observed bond ratings as a function of debt level. The underlying difference between the firms along the three different lines is that they have been randomly assigned different debt tax shield parameters τ_j . The firms with the highest log default probability were those assigned with the highest tax shield parameter of 43%. The firms with the lowest log default probability were those publicly randomly assigned to the lowest tax shield parameter of 3%. For a real-world CFO, the *operational* significance of the figure is simple: “Conditional upon variables observable to investors, equilibrium bond ratings are actually insensitive to changes in debt levels.” Or, “If I change my company’s leverage ratio, its bond rating actually will not change.”

Where then has the econometrician gone wrong? To begin, note that in her attempt to describe the decision problem of the CFO (equation (2)), the econometrician has violated rational expectations in implicitly treating agents outside the firm as being privy to the very same latent variable (firm quality Θ_j) that generated her original concern about omitted variables bias and motivated her IV strategy. That is, she has implicitly treated outside investors as having access to the private information held by the firm’s current owner. In reality, the initial owners of a firm will have some amount of private information and will attempt to signal this information to equity investors, debt investors, and bond rating agencies.

A correct specification of the capital structure optimization problem here reflects the fact that, as in Ross (1977), the original firm sponsor will use their own private knowledge of Θ_j in computing the probability of incurring reputational costs in the event of subsequent vehicle default, but must use the beliefs of the rating agencies and investors $\hat{\Theta}_\tau(B)$ (given an observable tax shield parameter τ) to determine what the market will be willing to pay for the rights to the cash flows accruing at $t = 1$. In turn, rating agency and investor beliefs will be influenced by the debt chosen by the firm. Thus, the true optimum financial policy for CFOs in this latent variables economy is:

$$B_j^{**} \in \arg \max_B \frac{\hat{\Theta}_\tau(B)}{2} + \tau_j B - LB \times \underbrace{\Theta_j^{-1} B}_{=\Pr(\tilde{c} \leq B_j)} . \quad (3)$$

The first-order condition for the CFO can be written as:

$$2\tau_j\Theta_j + \Theta_j \frac{\partial}{\partial B} \widehat{\Theta}_\tau(B^{**}) = 4LB^{**}. \quad (4)$$

Imposing the equilibrium condition that investors draw correct inferences, with $\widehat{\Theta}_\tau(B_j^{**}) = \Theta_j$, we obtain the following differential equation for the investor inference function:

$$2\tau_j \widehat{\Theta}_\tau(B) + \widehat{\Theta}_\tau(B) \frac{\partial}{\partial B} \widehat{\Theta}_\tau(B) = 4LB. \quad (5)$$

It is readily verified that the preceding ODE has a simple linear solution.

$$\widehat{\Theta}_\tau(B) = \left(\sqrt{\tau^2 + 4L} - \tau \right) B. \quad (6)$$

That is, outsiders, including bond rating agencies, will rationally infer that firm quality increases linearly with the debt level, with the slope of the investor inference function varying with the observed tax shield parameter. Substituting the preceding quality inference equation back into the first-order condition (4), the optimal debt schedule is:

$$B_j^{**}(\Theta) = \left[\frac{\tau_j + \frac{1}{2} \left(\sqrt{\tau_j^2 + 4L} - \tau_j \right)}{2L} \right] \Theta \geq B_j^*(\Theta). \quad (7)$$

The preceding equation informs us that firms in this economy will issue more debt than is predicted by the econometrician, since the very same latent variable problem confronting the econometrician here gives rise to a signaling benefit from debt which augments tax shield benefits.³

In estimating log default probabilities, the rating agencies must use their belief regarding firm type, with

$$\Lambda_j = \ln B_j - \ln \widehat{\Theta}_\tau(B_j) = \ln B_j - \left[\ln B_j + \ln \left(\sqrt{\tau_j^2 + 4L} - \tau_j \right) \right] = -\ln \left(\sqrt{\tau_j^2 + 4L} - \tau_j \right) \quad (8)$$

Notice, the critical implication of the preceding equation is that, consistent with Figure 1, conditional upon the observable tax shield value τ_j , a firm's bond rating will not vary with the level of debt it chooses—a prediction consistent with the OLS output. This is because the natural increase

³Notice, in this LCSE the lowest type issues zero debt, just as under the perfect information B^* .

in default probability coming from an increase in debt, holding beliefs fixed, is here just offset by the equilibrium increase in beliefs regarding firm quality generated by discretionary increases in debt.

In this economy, as in the real-world where econometricians and investors confront latent variables, CFOs must incorporate signaling effects if they are to arrive at optimal policies—and here OLS does so. By way of contrast, IV here informs the CFO of the effect on default probability of a *forced random change* in leverage induced by exogenous variation τ_j . It is hard to see why the latter quantity, delivering the answer to a contrived thought-experiment, would be of practical value to CFOs when making discretionary decisions.

4 Fuzzy Regression Discontinuity

4.1 Technology and Timing

In this section, we revisit the recent paper by Dittmar, Duchin and Zhang (DDZ below) (2020) who use fuzzy RDD in an attempt to provide “clean” estimates of the causal effect of SEOs. With this in mind, we develop a simple parable economy that mimics key aspects of their study.

Consider a discrete-time stock market economy in which investors are risk-neutral and stocks are priced at their expected terminal payoff. Within this economy there is a large yet finite number of equity-financed firms with initial shares outstanding at time $t = 0$ normalized at 1.⁴ Stock prices at each point in time are denoted p_t . Firms are identical ex ante, with identically distributed real assets-in-place and growth option technologies, implying the same initial price p_0 . The growth option technology delivers a sure payoff of two dollars for each dollar invested, so that each unit of investment has net present value equal to 1. However, investment scale is limited. In particular, maximum investment for each firm is an idiosyncratic i.i.d. random variable \tilde{I} uniformly distributed on $[0, 1]$, with realized value denoted I .

Timing is as follows. At $t = 1$, the realization of each firm’s maximum investment scale is publicly observed, causing stock prices to move to $p_1(I)$, with p_1 strictly monotone increasing in

⁴For simplicity, debt finance is ignored. One could assume firms face covenants prohibiting additional debt.

I , as shown below. That is, interim-date stock prices are increasing in feasible investment. At $t = 2$ each firm's board of directors meets and randomizes over whether to allow the firm's CFO to issue some amount of stock in an SEO. In the spirit of the the fuzzy regression discontinuity design employed by DDZ (2021), the board's SEO approval probability $\rho(I)$ is assumed to be an increasing step function.⁵ In particular for all $I < 1/2$, board approval occurs with probability $\underline{\rho}$ and for all $I \geq 1/2$ board approval occurs with probability $\bar{\rho}$. It is assumed that

$$1 > \bar{\rho} > \underline{\rho} > 0.$$

Intuitively, since investment has positive net present value, stock price will move upwards to $p_2(I) > p_1(I)$ if (and only if) the board gives the green light for doing some form of SEO.

Before proceeding, it is worth noting that this parable economy is designed to mimic the reference point identification strategy employed by DDZ. In particular, by construction, the price ratio $p_1(I)/p_0$ exceeds unity for those firms with $I \geq 1/2$ and falls below unity for those firms with $I < 1/2$. Thus, one could equally well think of the econometrician here as exploiting a discontinuity in SEO probability at the price ratio $p_1(I)/p_0 = 1$. However, we prefer to think of the econometrician as using firm-level maximum investment scale I as the conditioning information, since I is an exogenous random variable whereas price ratios are determined endogenously, perhaps complicating interpretation.

If given approval for an SEO, then at $t = 3$ the firm's CFO will work with an investment bank to perform due diligence and then decide on the exact number of shares of stock (s) to issue in the SEO. Asymmetric information arises at this point in time because the due diligence process privately reveals to the CFO and the investment bank the value of the firm's assets-in-place α . Investors simply know that α represents a draw from a uniform distribution with support $[0,2]$.

⁵In reality, private benefits can make insiders reluctant to dilute control. Of course, behavioral explanations are abundant.

4.2 Stock Market Equilibrium

The corporation will receive investment funding i equal to shares issued times the equilibrium share price:

$$i(s; I) \equiv sp_3(s; I). \quad (9)$$

The CFO's objective is to maximize the value of the claim held by the firm's current shareholders, subject to the constraint that she no more than double shares outstanding, since this would risk transferring corporate control to outsiders.

Letting α denote the value of assets-in-place, the CFO's program can be written as:

$$\max_{0 \leq s \leq 1} \frac{\alpha + 2i(s, I)}{1 + s} = \frac{\alpha + 2sp_3(s, I)}{1 + s}. \quad (10)$$

The equity market equilibrium from this point on is in the spirit of Myers and Majluf (1984) and Krasker (1986). However, here there will be a continuum of equilibrium pricing functions $p_3(\cdot, I)$. That is, given the observed feasible investment scale I for a firm, prices will vary endogenously with the CFO's announced value of s .

The first-order condition for the CFO's program (10) is:

$$2i_s(s, I) = (1 + s)^{-1}[\alpha + 2i(s, I)]. \quad (11)$$

Conjecturing a least-costly separating equilibrium (LCSE) in which the amount of stock issued fully reveals α , equilibrium entails new equity investors providing funding just equal to their expected payoff, or

$$i(s, I) = \left(\frac{s}{1 + s} \right) [\alpha + 2i(s, I)]. \quad (12)$$

Substituting the preceding equity market equilibrium condition into the CFO's first-order condition we obtain the following differential equation

$$2si_s(s, I) = i(s) \Rightarrow i(s, I) = A(I)\sqrt{s}. \quad (13)$$

where $A(I)$ is to be determined.

In the LCSE the worst type, with $\alpha = 0$, implements their symmetric information allocation, issuing the maximum number of shares ($s = 1$) and funding the maximum investment scale I . We then have

$$i(1, I) = I \Rightarrow i(s, I) = I\sqrt{s} \Rightarrow p_3(s, I) = \frac{i(s, I)}{s} = \frac{I}{\sqrt{s}}. \quad (14)$$

It follows from the preceding equation that the first-order condition (11) pinning down the optimal discretionary choice of s can be written as

$$Is^{-1/2} = (1 + s)^{-1}[\alpha + 2Is^{1/2}]. \quad (15)$$

From the preceding equation it follows that companies with more valuable assets-in-place will choose to issue fewer shares. In particular:⁶

$$s^*(\alpha, I) = \left[-\frac{\alpha}{2I} + \frac{1}{2} \sqrt{\left(\frac{\alpha}{I}\right)^2 + 4} \right]^2. \quad (16)$$

4.3 A Look at the Empirical Evidence

We consider now simulated stock price data in this economy, with the simulation parameterization featuring $\bar{\rho} = 7\%$ and $\underline{\rho} = 3\%$. These parameters are chosen to be consistent with the findings of DDZ (2021) that empirical SEO probabilities exhibit a similar jump at a current to historical price ratio threshold $p_1(I)/p_0 = 1$. Indeed, the simulated setting considered might seem to lend itself naturally to a fuzzy regression discontinuity research design exploiting the discrete upward jump in the probability of the SEO treatment for firms at the investment scale threshold $I = 1/2$.

With this in mind, consider first the behavior of the simulated stock price $p_1(I)$. As shown in Figure 2, the stock price is strictly monotone increasing in firms' respective maximum investment scale I . More importantly, the interim-date stock price exhibits a sharp upward jump at precisely the point where the SEO treatment probability jumps upward from 3% to 7%. This would seem to provide “clean” empirical evidence of a positive causal effect of stock issuance on stock prices.

⁶Let $\Omega = s^{1/2}$ and solve the quadratic in Ω , then compute $s = \Omega^2$.

In order to better understand the behavior of simulated stock prices in this economy, it is worth noting that the break-even condition for new shareholders, equation (12), implies that the intrinsic value of the original share outstanding can be rewritten as:

$$\frac{\alpha + 2i(s, I)}{1 + s} + \left(\frac{s}{1 + s} \right) [\alpha + 2i(s, I)] - i(s, I) = \alpha + i(s, I). \quad (17)$$

It follows that for all times $t \leq 2$:

$$p_t = \mathbb{E}_t[\alpha + i] = 1 + \mathbb{E}_t[i].$$

Intuitively, given that new investors just break even in the LCSE, the original shareholders rationally expect to capture the value of assets-in-place plus the net present value of new investment, with each unit of investment here having NPV=1. Since maximum investment scale is roughly the same in a small neighborhood about $I = 1/2$, the upward jump in stock price at this threshold is properly understood as arising from the higher probability of new investment funding at this threshold. Thus, one can think of the positive stock price reaction to the higher treatment probability as capturing the partial causal effect of stock issuance: If beliefs about firm type are held fixed, a higher probability of stock issuance leads to a higher stock price.

Consider next the behavior of simulated stock prices at $t=2$ when boards randomize over allowing the SEO process to move forward. Employing the same methodology as in DDZ (2021), we consider the simulated announcement effect arising from boards announcing that approval has been given for an SEO, focusing on stock price reactions for simulated firms with pre-announcement price ratios $p_1(I)/p_0$ in small neighborhoods just below and just above their posited anchoring ratio of 1. Again, we recall that by construction in the simulated economy there is indeed a discrete jump upward in SEO treatment probability at the threshold of unity. Here too, the simulated data would seem to support the general notion that stock issuance has a positive causal effect on stock price. For all firms, stock prices at $t=2$ react positively to board approval of SEOs. Moreover, consistent with the findings in DDZ, the average abnormal return for firms above the price ratio=1 threshold is higher (39.78%) than for firms below the threshold (9.29%). This is because simulated firms above the threshold have better growth options on average due to higher maximum investment scales.

All this “credible” evidence notwithstanding, it is of limited value to the CFO who must decide at $t=3$ on just how much stock to issue—after consulting with the investment bank. In fact, the preceding evidence is actually misleading. After all, if stock issuance has a positive causal effect on stock price, this would seem to imply that the CFO should issue the maximum feasible number of shares, setting $s = 1$. Of course, this reasoning neglects the fact that maximum share issuance sends the signal to the market that the firm has the lowest possible assets-in-place value of $\alpha = 0$, which would send the stock price tumbling downward. More generally, the partial causal effects captured by the regression discontinuity design completely strip out the signaling channel of stock issuance—an effect of first-order importance to CFOs.

What would be useful to a CFO standing at $t = 3$ in this parable economy would be evidence from event studies examining stock price reactions to *endogenous* decisions made by prior CFOs who found themselves at the endogenous decision margin at which the CFO finds herself—with stock price reactions being conditioned upon variables observable to all agents in the economy, investors and econometricians. More specifically, arriving at an optimal solution to the program in equation (10) requires the CFO to understand much more than the general notion that stock price responds positively to higher SEO approval probabilities. Rather, the CFO must know the stock price reaction function to discretionary stock issuance volumes $p_3(\cdot, I)$ with the second argument I representing observable conditioning information, here maximum feasible investment scale. To this end, a traditional corporate finance event study, or even a conversation with an investment banker, would seem to provide more practically useful information than that derived from exogenous variation in SEO propensities.

Figure 3 illustrates, depicting equilibrium stock price reaction functions $p_3(\cdot, I)$ (equation (14)) for two sets of firms—those with observables just below ($I = .45$) and just above ($I = .55$) the SEO treatment cutoff threshold of $I = .50$. The figure reveals two important pieces of information to CFOs making discretionary decisions in this economy, information not provided by the RDD identification strategy. First, stock price responds negatively to marginal increases in the number of

shares issued—a causal effect that would be captured using a traditional event study approach, e.g. Asquith and Mullins (1986). Second, apparently the CFO will face a steeper stock price reaction function if maximum investment scale is higher.

5 Evidence-Based Government Policy

Much of the motivation given for random assignment is based upon the notion that the resulting causal estimates are particularly valuable in terms of setting optimal government policies. Indeed, this is the argument made by Romer and Romer (2010) who argue that effects of tax changes are better gauged by examining what appear to be exogenous shocks to tax policy.

With the Romer and Romer (2010) study in mind, this section considers a more detailed and complex example of how the two forms of causal effects, partial and total causal effects, can both be used by a government policymaker in a complementary fashion. The government’s decision problem is a common one, how to use empirical evidence to decide on regulatory and tax policies.

Time is continuous and horizons are infinite. Agents are risk-neutral and discount cash flows at the risk-free rate r . Firms (or other agents) accumulate a stock (say capital) according to the law of motion

$$dK = (I - \delta K)dt \tag{18}$$

with the price of a unit of capital being 1 and adjustment costs being γI^2 .

A government (e.g. state or national) has discretion to choose the state of its policy variable. We will call this economy, the “endogenous policy economy.” The policy state is binary, $S \in \{S1, S2\}$. The policy variable influences marginal product, and with it, investment. In particular, in policy state S the marginal product is $\Pi_S X K$, where X is a geometric Brownian motion evolving according to

$$\frac{dX_t}{X_t} = \mu_t dt + \sigma dW \tag{19}$$

where W is a standard Wiener process. We shall think of X as representing an aggregate shock hitting firms in the endogenous policy economy. Notice, the drift μ_t is time-varying. In particular,

as described in greater detail below, we assume μ_t is a binary stochastic process. The realization of this process is private information to the government.

For the sake of the illustration, assume

$$\Pi_{S2} > \Pi_{S1}.$$

The first causal inference problem is that, by assumption, the government does not know the preceding inequality, nor the magnitude of either Π_S . That is, the government does not know which policy variable state is *technologically* more stimulative.

Suppose now that there is a neighboring economy (the “experimental policy economy”) ex ante identical in all respects to the endogenous policy economy, but with the exception that this neighbor randomizes its policy variable, alternating between Π_{S1} and Π_{S2} . In particular, over any infinitesimal time interval dt , with probability λdt the policy variable will switch states. This stochastic process is independent of any other random variable including the aggregate shock hitting the experimental economy which has the following law of motion

$$\frac{d\tilde{X}}{\tilde{X}} = \tilde{\mu}_t dt + \sigma d\tilde{W}$$

where \tilde{W} is also a standard Wiener process and $\tilde{\mu}_t$ is also a binary stochastic process that is unobservable to firms, with identical probability law as μ_t . The fact that the experimental policy economy is endowed with the same probability law for the aggregate shock as the endogenous policy economy makes it a convenient benchmark.

The government of the endogenous policy economy will first attempt to use evidence from the experimental policy economy’s shock responses to infer the (relative) magnitudes of Π_{S1} and Π_{S2} . Since the neighbor randomizes its policy variable, this first step will be an exercise in estimating partial causal effect (signs).

Assuming the optimizing government is successful in determining which policy state is more stimulative in terms of underlying latent technological parameters, it faces a second challenge: determining the magnitude of total causal effects. In particular, we assume the government in the

endogenous policy economy will adopt as a policy rule switching to the more (less) stimulative policy if the current instantaneous aggregate drift is low (high). However, as we show, since the aggregate drift is private information to the government, the response to policy variable changes will be dampened (and potentially reversed) due to the opposing signal content. Here the econometrician must estimate total causal effects in order to correctly predict how firms will respond to discretionary policy interventions.

5.1 Experimental Policy

We consider first the experimental policy economy in which the policy variable is an exogenous stochastic process. Following Veronesi (2000), assume that the instantaneous drift can take on two potential values, $\mu_1 > \mu_2$. This holds true in both economies.

Recall, the drift is unobservable to all parties except the government. Over any infinitesimal time interval dt with probability pdt a drift will be randomly drawn according to the probability distribution $\mathbf{f} = (f_1, f_2)$. Let \mathbf{Z} be the two-dimensional vector of probability weights agents place on each potential drift and let

$$\mu(\mathbf{Z}) \equiv Z_1\mu_1 + Z_2\mu_2. \quad (20)$$

In the experimental economy, the government provides no signals, so agents must instead form beliefs based upon the realized path of \tilde{X} . From Lemma 1 in Veronesi (2000) it follows beliefs evolve as diffusions, with:

$$dZ_n = \underbrace{p(f_n - Z_n)dt}_{\equiv \mu_{z_n}} + \underbrace{\frac{Z_n[\mu_n - \mu(\mathbf{Z})]}{\sigma}}_{\equiv \sigma_{z_n}} d\tilde{W}. \quad (21)$$

The Hamilton-Jacobi-Bellman (HJB) equation for the firm is:

$$\begin{aligned} & rV(K, X, S, \mathbf{Z}) \\ = & \max_I \Pi_S KX - I - \gamma I^2 + V_k(I - \delta K) + V_x \mu(\mathbf{Z})X + \mu_{z_1} V_{z_1} + \mu_{z_2} V_{z_2} \\ & + \lambda[V(K, X, S', \mathbf{Z}) - V(K, X, S, \mathbf{Z})] + \frac{1}{2} V_{xx} \sigma^2 X^2 + \frac{1}{2} V_{z_1 z_1} \sigma_{z_1}^2 + \frac{1}{2} V_{z_2 z_2} \sigma_{z_2}^2 \\ & + V_{z_1 z_2} \sigma_{z_1} \sigma_{z_2} + V_{x z_1} X \sigma \sigma_{z_1} + V_{x z_2} X \sigma \sigma_{z_2}. \end{aligned} \quad (22)$$

The HJB equation can be understood as an equilibrium condition demanding that the expected holding return on the firm's stock must be just equal to the opportunity cost. The holding return consists of dividends plus expected capital gains. In turn, the capital gains can be understood as a second-order Taylor expansion using the rules of Ito calculus.

We conjecture and verify that the value function takes the following separable form:

$$rV(K, X, S, \mathbf{Z}) = KQ(X, S, \mathbf{Z}) + G(X, S, \mathbf{Z}). \quad (23)$$

Isolating those terms in the HJB equation involving the instantaneous investment control we find that the optimal investment policy solves

$$\max_I IQ(X, S, \mathbf{Z}) - I - \gamma I^2 \Rightarrow I^*(X, S, \mathbf{Z}) = \frac{Q(X, S, \mathbf{Z}) - 1}{2\gamma}. \quad (24)$$

That is, investment is linear in the shadow value of capital Q .

Next we note that since the HJB equation must hold pointwise, the terms scaled by K must equate. Using this fact, we obtain an equation for pinning down the shadow value of capital Q :

$$\begin{aligned} & (r + \delta + \lambda)Q(X, S, \mathbf{Z}) \\ = & \Pi_S X + \mu(\mathbf{Z})XQ_x + \mu_{z_1}Q_{z_1} + \mu_{z_2}Q_{z_2} + \lambda Q(X, S', \mathbf{Z}) \\ & + \frac{1}{2}\sigma^2 X^2 Q_{xx} + \frac{1}{2}\sigma_{z_1}^2 Q_{z_1 z_1} + \frac{1}{2}\sigma_{z_2}^2 Q_{z_2 z_2} \\ & + \sigma_{z_1} \sigma_{z_2} Q_{z_1 z_2} + X \sigma \sigma_{z_1} Q_{xz_1} + X \sigma \sigma_{z_2} Q_{xz_2}. \end{aligned} \quad (25)$$

Now let $X\Psi_S^n$ denote the shadow value of capital in policy state S given drift rate μ_n . As shown in the Online Appendix, we have the following lemma pinning down the solution to the preceding shadow value equation for the experimental economy.

LEMMA 2. *In the experimental economy, the shadow value of capital is*

$$Q(X, S, \mathbf{Z}) = X[Z_1\Psi_S^1 + Z_2\Psi_S^2]$$

where the shadow value constants solve the following linear system

$$\begin{aligned}
[r + \delta - \mu_1 + \lambda + pf_2]\Psi_{S1}^1 &= \Pi_{S1} + pf_2\Psi_{S1}^2 + \lambda\Psi_{S2}^1 & (26) \\
[r + \delta - \mu_2 + \lambda + pf_1]\Psi_{S1}^2 &= \Pi_{S1} + pf_1\Psi_{S1}^1 + \lambda\Psi_{S2}^2 \\
[r + \delta - \mu_1 + \lambda + pf_2]\Psi_{S2}^1 &= \Pi_{S2} + pf_2\Psi_{S2}^2 + \lambda\Psi_{S1}^1 \\
[r + \delta - \mu_2 + \lambda + pf_1]\Psi_{S2}^2 &= \Pi_{S2} + pf_1\Psi_{S2}^1 + \lambda\Psi_{S1}^2.
\end{aligned}$$

Subtracting the first equation listed in the lemma from the third and the second equation from the fourth, the following inequalities are readily verified:

$$\Pi_{S2} > \Pi_{S1} \Rightarrow \Psi_{S2}^1 > \Psi_{S1}^1, \Psi_{S2}^2 > \Psi_{S1}^2. \quad (27)$$

We then have the following important proposition showing the utility of natural policy experiments in determining the relative magnitude of deep technological parameters. Of course, inferring technological parameters is often a natural pre-requisite for setting policy optimally, and this is the case here.

PROPOSITION 1. *Observing any shock response in the experimental economy allows for correct measurement of the signs of partial causal effects, with the investment response to an exogenous transition of the policy variable from state S to S' being*

$$SR_{SS'} = \frac{X}{2\gamma} \times [Z_1(\Psi_{S'}^1 - \Psi_S^1) + Z_2(\Psi_{S'}^2 - \Psi_S^2)]. \quad (28)$$

Notwithstanding the positive conclusion of the proposition, that natural experiments here allow for a correct ranking of relative stimulus provided by the alternative policies (sans-signaling), it is also clear that latent time-varying beliefs (\mathbf{Z}) will make it hard for the government to correctly infer the absolute magnitudes of the technological parameters. Anticipating, this will be problematic since, once policy discretion is introduced, there will be a signaling effect working in the opposite direction of the partial causal effect.

5.2 Endogenous Policy

Suppose now that, based upon the experimental evidence (Proposition 1), the government of the endogenous policy economy tries to lean against the wind, implementing policy Π_{S2} whenever it privately observes that the drift rate is low ($\mu_t = \mu_2$) and Π_{S1} whenever it privately observes that the drift rate is high ($\mu_t = \mu_1$), recalling $\mu_1 > \mu_2$. What will be the observed total causal effect?

The HJB equation for the firm here is:

$$\begin{aligned} & rV(K, X, S) \\ &= \max_I \Pi_S K X - I - \gamma I^2 + V_k(I - \delta K) + V_x \mu_S X + \\ & \quad + p(1 - f_S)[V(K, X, S')] - V(K, X, S)] + \frac{1}{2} V_{xx} \sigma^2 X^2. \end{aligned} \tag{29}$$

We conjecture and verify that the value function takes the following separable form:

$$V(K, X, S) = Kq(X, S) + g(X, S). \tag{30}$$

Isolating those terms in the HJB equation involving the instantaneous investment control we find that the optimal investment policy solves

$$\max_I Iq(X, S) - I - \gamma I^2 \Rightarrow I^*(X, S) = \frac{q(X, S) - 1}{2\gamma}. \tag{31}$$

Next we note that since the HJB equation must hold pointwise, the terms scaled by K must equate. Using this fact and rearranging terms we obtain an equation for pinning down the shadow value of capital:

$$[r + \delta + p(1 - f_S)]q(X, S) = \Pi_S X + \mu_S X q_x(X, S) + \frac{1}{2} \sigma^2 X^2 q_{xx}(X, S) + p(1 - f_S)q(X, S'). \tag{32}$$

Since the dividend is linear in X we conjecture the shadow value is also linear in X taking the form

$$q(X, S) = X\psi_S.$$

Substituting this into the preceding equation and rearranging terms we obtain the following two

equations pinning down the shadow values in the endogenous policy economy

$$\begin{aligned} [r + \delta - \mu_1 + p(1 - f_1)]\psi_1 &= \Pi_{S1} + p(1 - f_1)\psi_2 \\ [r + \delta - \mu_2 + p(1 - f_2)]\psi_2 &= \Pi_{S2} + p(1 - f_2)\psi_1. \end{aligned} \quad (33)$$

Solving this system we find

$$\begin{aligned} \psi_1 &= \frac{\Pi_{S1}}{r + \delta - \mu_1} \\ &+ \frac{p(1 - f_1)(r + \delta - \mu_2) [\Pi_{S2}/(r + \delta - \mu_2) - \Pi_{S1}/(r + \delta - \mu_1)]}{(r + \delta - \mu_1)(r + \delta - \mu_2)[1 + p(1 - f_1)/(r + \delta - \mu_1) + p(1 - f_2)/(r + \delta - \mu_2)]} \end{aligned} \quad (34)$$

with the symmetric expression for ψ_2 .

Since investment is increasing in $q \equiv X\psi_S$, implementation of Π_{S2} will be followed by an increase in investment iff $\psi_2 > \psi_1$. Comparing the shadow value constants from equation (34) we have the following proposition.

PROPOSITION 2. *Investment will increase after discretionary government implementation of the technologically stimulative policy $\Pi_{S2} > \Pi_{S1}$ if and only if the partial causal effect is sufficiently large to imply*

$$\frac{\Pi_{S2}}{\Pi_{S1}} > \frac{r + \delta - \mu_2}{r + \delta - \mu_1}. \quad (35)$$

Essentially, the preceding proposition shows that the total causal effect will have the same sign as the partial causal effect if and only if the latter causal effect is sufficiently large to offset the negative signal that the government's discretionary stimulus intervention sends regarding aggregate drift. To see this, note that the condition in the proposition can be stated in terms of shadow values under constant policies and drifts:

$$\frac{\Pi_{S2}}{r + \delta - \mu_2} > \frac{\Pi_{S1}}{r + \delta - \mu_1} \Rightarrow q(X, S2) > q(X, S1). \quad (36)$$

The point of this example is to illustrate in a concrete way how *both* forms of causal effect estimates may be necessary operationally. Perhaps ironically, here partial causal effects were shown

to be especially helpful in terms of inferring deep technological parameters. But once these technological parameters are assessed, it seems that (perhaps) only through actually implementing the policy rule in a discretionary fashion can the government get a better sense of how things will work in reality. Moreover, in contrast to the present stylized example, in reality it will be hard to predict the magnitude of signaling effects since the true nature and quality of agent information is hard for a government to know with a high degree of precision. In fact, in reality, information quality and beliefs will vary over time, giving rise to time-varying signal content in many applied settings. This might well be an interesting direction for future applied theory and empirical work.

6 Conclusion

This paper questions the notion that evidence from random assignment is somehow more “credible” than more traditional forms of evidence in finance such as event studies. As we show, if a prospective decision-maker is privately informed and is indeed attempting to make optimal *discretionary* decisions, rather than attempting to understand the consequences of random mistakes, then the signal content of her decisions is payoff-relevant and decision-relevant. Causal effect estimates derived from random assignment strip out signaling effects, potentially leaving the decision-maker ignorant of the true implications of alternative decisions. Nevertheless, we show how partial causal effect estimates can be used in conjunction with total causal effect estimates to decompose real and information signaling channels. As shown, in practice, both forms of empirical estimates, partial and total causal effects, may be necessary inputs in order to pin down optimal policies. If this fact is appreciated, then a broader definition of “identification” should be adopted, and a broader set of evidence should be accepted as “credible.”

References

- [1] Alti, Aydogan, 2003, How Sensitive is Investment to Cash Flow when Financing is Frictionless? *The Journal of Finance* 58, 707-722.
- [2] Angrist, Joshua D. and Jorn-Steffen Pischke, 2009, *Mostly Harmless Econometrics: An Empiricist's Companion*, Princeton University Press.
- [3] Angrist, Joshua D. and Jorn-Steffen Pischke, 2015, *Mastering Metrics: The Path from Cause to Effect*, Princeton University Press.
- [4] Asquith, Paul and David Mullins, 1986, Equity Issues and Offering Dilution, *Journal of Financial Economics* 15 (1), 61-89.
- [5] Bowen, Donald E., Laurent Fresard, and Jerome Taillard, 2017, What's Your Identification Strategy? Innovation in Corporate Finance Research, *Management Science*, 2529-2548.
- [6] Chemla, Gilles and Christopher A. Hennessy, 2019, Controls, Belief Updating, and Bias in Medical RCTs, *Journal of Economic Theory*.
- [7] Dittmar, Amy, Ran Duchin, and Shuran Zhang, 2020, The Timing and Consequences of Seasoned Equity Offerings: A Regression Discontinuity Approach, *Journal of Financial Economics* 138 (1), 254-276.
- [8] Duflo, Esther, 2004, Scaling Up and Evaluation, in *Annual World Bank Conference on Development Economics: Accelerating Development*, ed. François Bourguignon and Boris Pleskovic, 341-69. Washington, D.C.: World Bank; Oxford and New York: Oxford University Press.
- [9] Fama, Eugene, Lawrence Fisher, Michael Jensen, and Richard Roll, 1969, The Adjustment of Stock Prices to New Information, *International Economic Review* 10, 1-28.
- [10] Fudenberg, Drew and David K. Levine, 2020, Learning in Games and the Interpretation of Natural Experiments, forthcoming in *American Economic Journal: Microeconomics*.

- [11] Gomes, Joao, 2001, Financing Investment, *American Economic Review* 91, 1263-1285.
- [12] Gourio, Francois, 2013, Credit Risk and Disaster Risk, *American Economic Journal: Macroeconomics* 5 (3), 1-34.
- [13] Hara, Chiaki, Ilya Segal, and Steve Tadelis, 1997, *Solutions Manual for Microeconomic Theory*, Oxford University Press.
- [14] Imbens, Guido W. and Donald B. Rubin, 2015, *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*, Cambridge University Press.
- [15] Keane, Michael P., Structural vs. Atheoretic Approaches to Econometrics, 2010, *Journal of Econometrics* 156.
- [16] Krasker, William, 1986, Stock Price Movements in Response to Stock Issues under Asymmetric Information, *Journal of Finance*.
- [17] Leland, Hayne E., 1994, Corporate Debt Value, Bond Covenants, and Optimal Capital Structure, *Journal of Finance*.
- [18] Leland, Hayne E. and David Pyle, 1977, Information Asymmetry, Financial Structure, and Financial Intermediation, *Journal of Finance* 32, 371-387.
- [19] Löfgren, Karl-Gustaf, Torsten Persson, and Jorgen Weibull, 2002, Markets with Asymmetric Information: The Contributions of George Akerlof, Michael Spence and Joseph Stiglitz, *Scandinavian Journal of Economics* 104(2), 195-211.
- [20] Molina, Carlos A., 2005, "Are Firms Underleveraged? An Examination of the Effect of Leverage on Default Probabilities, *Journal of Finance* 60 (3), 1427-1459.
- [21] Moyen, Nathalie, 2005, Investment-Cash Flow Sensitivities: Constrained vs Unconstrained Firms. *Journal of Finance* 59, 2061-2092.

- [22] Myers, Stewart, and Nicholas S. Majluf, 1984, Corporate Financing and Investment Decisions when Firms have Information that Investors do not Have, *Journal of Financial Economics* 13, 187-221.
- [23] Philipson, Tomas, and Jeffrey Desimone, 1997, Experiments and Subject Sampling, *Biometrika*, 619-630.
- [24] Riley, John G., 2001, Silver Signals: Twenty-Five Years of Screening and Signaling, *Journal of Economic Literature*, 432-478.
- [25] Romer, Christina D. and David H. Romer, 2010, The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks, *American Economic Review* 100, 763-801.
- [26] Ross, Stephen, 1977, The Determination of Financial Structure: The Incentive Signaling Approach, *Bell Journal of Economics*, 23-40.
- [27] Rubin, Donald B., 1974, Estimating the Causal Effects of Treatments in Randomized and Non-Randomized Studies, *Journal of Educational Psychology* 66, 688-701.
- [28] Rust, John, 2014, The Limits of Inference with Theory: A Review of Wolpin 2013, *Journal of Economic Literature*.
- [29] Spence, Michael, 1973, Job Market Signaling, *Quarterly Journal of Economics* 87(3), 355-374.
- [30] Veronesi, Pietro, 2000, How Does Information Quality Affect Stock Returns? *Journal of Finance*, 807-837.

Table 1: Survey of Empirical Corporate Finance in Top Finance Journals, 2000-2012

	Total	IV	RDD	Experiments
Bowen, Fresard and Tallard Sample using IV, RDD, Experiments	253	243	10	3
Number of Papers in Our Sample	185	177	9	2
Papers Subject to Signaling Critique	130	125	4	2
Papers where the main explanatory variable is				
Corporate governance - ownership	30	30	0	0
Corporate governance - board	12	11	1	1
Corporate governance - other	22	21	1	0
Executive compensation	10	10	0	0
Investment	12	11	0	1
Capital structure/financing/payout policy	35	33	2	0
Other	9	9	0	0
Papers where the main dependent variable is				
Valuation/performance measure	77	75	1	1
Investment	28	26	2	0
Financing	22	21	1	0
Ownership Structure	7	6	0	1
Other	1	1	0	0

Table 2: OLS and IV Estimates of Bond Rating Determinants

PANEL A: OLS			
	Coefficient	Stand. Error	t statistic
LN DEBT	0.0113122	0.0020809	5.44
CONSTANT	-0.0947435	0.0043199	-21.93
PANEL B: IV			
(First Stage)	Coefficient	Stand. Error	t statistic
TAX SHIELD	0.6716006	0.123597	5.43
CONSTANT	-2.0453	0.0451979	-45.25
	Coefficient	Stand. Error	z statistic
LN DEBT	1.0361	0.190605	5.44
CONSTANT	1.773156	0.3479821	5.11

Figure 1: Leverage and Bond Ratings

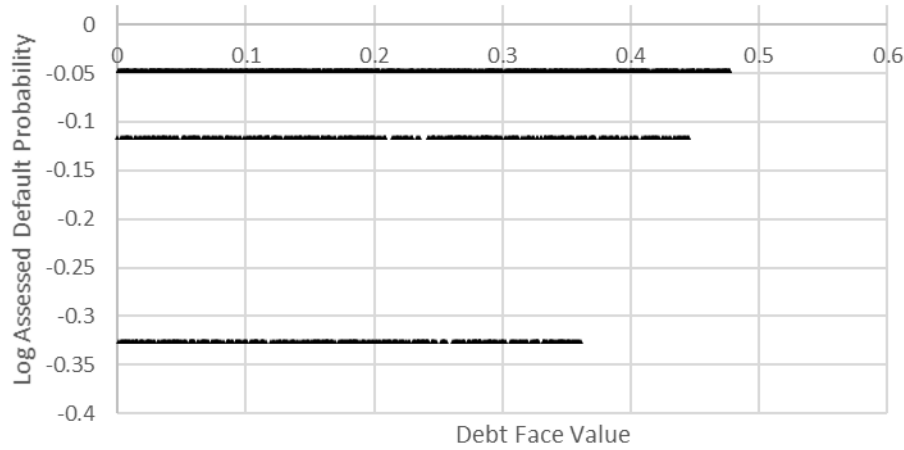


Figure 2: Causal Effect of SEOs using Fuzzy RDD

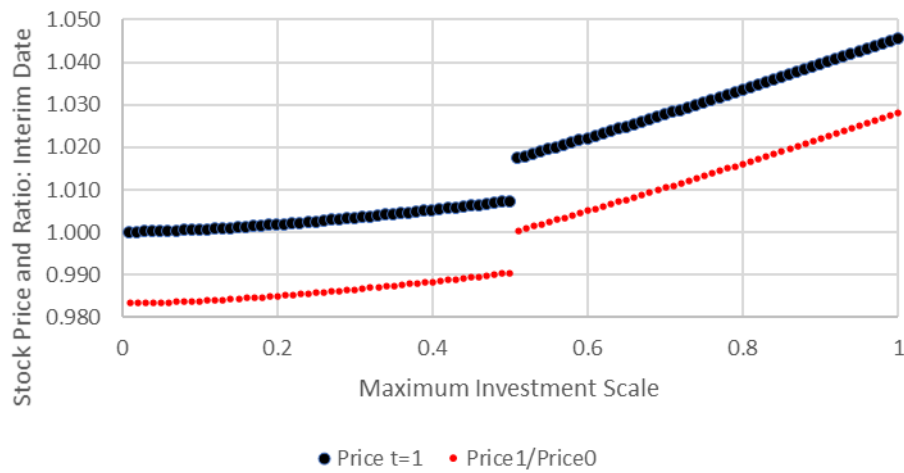


Figure 3: Issuance Volume and Stock Price

