

LBS Research Online

[C Hennessy](#) and G Chemla

Signaling, instrumentation, and CFO decision-making

Article

This version is available in the LBS Research Online repository: <https://lbsresearch.london.edu/id/eprint/2214/>

[Hennessy, C](#) and Chemla, G

(2022)

Signaling, instrumentation, and CFO decision-making.

Journal of Financial Economics, 144 (3). pp. 849-863. ISSN 0304-405X

DOI: <https://doi.org/10.1016/j.jfineco.2021.07.014>

Elsevier

<https://www.sciencedirect.com/science/article/pii/...>

Users may download and/or print one copy of any article(s) in LBS Research Online for purposes of research and/or private study. Further distribution of the material, or use for any commercial gain, is not permitted.

Signaling, Instrumentation, and CFO Decision-Making*

Christopher A. Hennessy
LBS, CEPR, and ECGI.

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

April 2021

Abstract

Building parable economies embedding econometricians, we view alternative estimators (IV, fuzzy RD, natural experiments, OLS, event studies) from the perspective of privately-informed decision-makers, e.g. CFOs. IV estimates can be misleading since randomization through observable instruments eliminates signal content arising from discretion. If the goal is informing discretionary decisions, rather than predicting outcomes after forced/mistaken actions, instrumentation is problematic, whereas OLS or event studies can be sufficient. The analysis shows the utility of alternative estimators hinges upon oft-neglected assumptions about agent/econometrician information sets, as distinct from exclusion restrictions. We recommend parable economy estimation as precursor to real-world IV estimation.

1 Introduction

In their textbook, *Mostly Harmless Econometrics: An Empiricist's Companion* (MHE below), Angrist and Pischke (2009) argue that “The most credible and influential research designs use random assignment.” This view has gained traction across a wide-range of fields including empirical corporate finance. As one indicator, Bowen, Frésard, and Taillard (2017) find that in the top-three finance journals, the share of empirical corporate finance papers using “identification technologies” rose from roughly 0 percent in the late 1980’s to over 50 percent by 2012. By way of contrast, consider that in 1986 the *Journal of Financial Economics* published in a single issue five event

*Contact chennessy@london.edu (corresponding author) and g.chemla@imperial.ac.uk. We thank seminar participants at UC Berkeley, LBS, Université Paris Dauphine, Mannheim, Imperial College, and ESCP Europe. Thanks also to Zhengming Li for research assistance.

studies analyzing market responses to endogenous capital structure changes, an unlikely journal configuration today.

Duflo (2017) recommends that development economists adopt the mindset of “plumbers”, using “credible” experimental evidence to help governments make informed decisions about policy. Nevertheless, as argued by Banerjee, Chassang and Snowberg (2017), there has been little in the way of formal attempts to connect experimental economics with decision theory. Moreover, Banerjee, Chassang, Montero and Snowberg (2020) (BCMS below) show that deterministic experiments dominate randomized controlled trials (RCTs) under some conditions.

To our knowledge, there has been no attempt to model which forms of evidence should be preferred within empirical corporate finance taking into account the role the evidence will play in decision-making. This paper does so. Of course, the fact that empirical corporate finance is a subfield of applied microeconomics would allow one to gain some insights directly from the aforementioned decision-theoretic literature. Nevertheless, a number of factors limit extrapolation. First, a primary objective of empirical corporate finance is to help inform the decisions of atomistic CFOs making unilateral decisions as distinct from governments imposing common tax/regulation across all firms/households. Second, while development economists use RCTs to help decision-making governments learn about the private information of households, decision-making CFOs instead generally enjoy superior information relative to outside investors, as stressed by Myers and Majluf (1984). Finally, whereas development economists have successfully marshaled resources for RCTs, researchers in empirical corporate finance have instead relied upon instrumental variables (IV) and variations upon IV, e.g. quasi-natural experiments and fuzzy regression discontinuity.

Angrist and Pischke (2009) open their textbook by arguing that “the estimators in common use almost always have a simple interpretation that is not heavily model dependent.” Our paper takes the opposite perspective, viewing alternative forms of empirical evidence through the lens of models or parable economies intended to approximate key features of the environment occupied by the econometrician, the decision-making CFO, and outside investors. Specifically, taking the perspective

of a privately-informed decision-maker (e.g. CFO), we explore the merits and limitations of IV-type estimators relative to traditional methods such as OLS and event studies.

A first important implication of the analysis is that evidence derived from standard IV methods is often problematic from the perspective of privately-informed decision-makers such as CFOs.¹ This is because instrumentation, in achieving quasi-randomization, simultaneously strips out the signal content arising from discretionary actions. Intuitively, since instrumentation achieves orthogonal randomization in a way observable to the econometrician and other agents in the economy, concomitant changes in the causal variable do not signal private information. However, when discretion is ultimately exercised, signaling effects form part of the causal chain the decision-maker must account for if she is to make optimal decisions.

Anticipating, we offer and make use of two alternative causal effect definitions. *Partial causal effects* are changes in the outcome variable arising from changes in the forcing variable holding fixed beliefs. As shown, this type of causal effect is recovered by standard IV methods. *Total causal effects* are changes in the outcome variable resulting from discretionary changes in the forcing variable allowing for endogenous changes in beliefs. A second important implication of our analysis is that, as shown, the total causal effect is often a sufficient statistic for optimal decisions by privately-informed CFOs. A third closely related point is that total causal effects can often be recovered by traditional methods that currently enjoy less favor, specifically conditional event studies and even (“biased”) OLS with clearly endogenous regressors. Conversely, pursuing valid instruments with the aim of recovering partial causal effects may be unnecessary if the goal is informing discretionary decisions.

A final broad conclusion that emerges from our analysis is that interpretation and utilization of empirical evidence depends upon granular assumptions about the distinct information sets of decision-makers, investors, *and* econometricians. Thus, we advocate empiricists develop and present simple stylized models mimicking the exploited real-world data generating process, inclusive of information sets, when trying to decide between competing estimators and estimates. Indeed, the

¹See Young (2019) for an alternative argument that even biased OLS may dominate IV in terms of MSE.

three stylized models we develop in the paper illustrate how this can be done in practice, without resorting to complex structural modeling.

We first revisit Molina (2005) who finds the effect of leverage on bond ratings is three times stronger under IV than if debt endogeneity is not eliminated. With this in mind, we consider a parable economy in which CFOs face tax benefits of debt and reputational costs of bankruptcy, with firm quality being unobservable to rating agencies, investors, and the econometrician. By construction tax rates are an ideal instrument in this economy, and the simulated IV coefficient is large while the OLS coefficient is negligible. Nevertheless, the OLS coefficient actually captures the relevant causal effect from the perspective of CFOs: Ratings are invariant to discretionary leverage changes conditional upon observables. Intuitively, rating agencies and investors face the same latent variables problem as the econometrician. Consequently, the endogeneity that instrumentation eliminates actually forms part of the causal chain CFOs face. Nevertheless, total causal effects could be measured while maintaining the benefits of IV if the econometrician herself has private information, observing a “hidden instrument” whose variation goes undetected by the market.

We next revisit Dittmar, Duchin and Zhang (DDZ below) (2020) who “provide some of the cleanest estimates, to date, of the timing and causal effects of SEOs.” We develop a parable economy that mimics key aspects of their fuzzy regression discontinuity (RD) design. At an interim date, ex ante identical firms face 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 exogenous board approval for moving forward with the SEO process with a relatively high (low) probability. In this economy, evidence from the fuzzy RD supports the notion that equity issuance has a positive causal effect on stock price. In particular, the interim-date stock price exhibits a discontinuous upward jump at the threshold where boards green light SEOs with relatively high probability. Further, stock prices react favorably if the board randomly approves the SEO process. However, the preceding evidence is of limited use to the decision-making CFOs who decide on the actual number of shares to

issue following the acquisition of private information during due diligence. These CFOs still confront a downward sloping stock price function conditional upon observables, and conditional event studies reveal this. By way of contrast, the fuzzy RD evidence would have been useful to boards making green light decisions prior to the arrival of private information, revealing the SEO process to be profitable in expectation.

Although the primary focus is CFOs, our arguments apply to other privately-informed decision-makers, e.g. governments. To illustrate, our final example revisits Romer and Romer (2010) who argue for using legislative histories in order to isolate quasi-random fiscal shocks. In this spirit, we consider a government that privately observes binary macroeconomic drift, contemplating changing a binary policy (tax/regulation) variable to lean against the wind. Here we show responses to exogenous shocks can identify potentially stimulative policies—in the sense of having positive partial causal effects. However, the actual impact of discretionary stimulus is correctly measured by the total causal effect which can be directly observed once the government does exercise policy discretion. In addition, the government’s own macroeconomic signaling is critical in allowing it to infer deep causal parameters since the signaling eliminates the need for the government to infer private sector beliefs about macroeconomic drift. Once again, modeling the data generating process reveals the subtle role played by information sets in the interpretation and utilization of empirical evidence.

In addition to the papers cited above that stress decision-making, our paper is in the spirit of Chassang, Snowberg, Seymour and Bowles (2015) who show that double-blind RCTs may understate the causal impact of medical treatments by failing to account for interactions between a patient’s beliefs and moral hazard.² We stress that IV methods in corporate finance generally fail to capture causal effects arising from endogenous changes in investor beliefs, whereas event studies and OLS may do so, potentially making them more useful in decision-making. Our paper is related to work by Fudenberg and Levine (2021) who show agent Bayesian learning can drive a wedge between true partial causal effects and causal estimates obtained from regression discontinuity. Intuitively, agents on opposing sides of regression discontinuity boundaries may form different beliefs about

²Chemla and Hennessy (2019) model bias from Bayesian patient placebo effects.

effort returns. In contrast to our paper, Fudenberg and Levine do not question the utility of partial causal effects in decision-making, but instead show that regression discontinuity may fail to recover them.

Signaling effects are related to selection effects but may differ from a decision-maker perspective. For example, in the motivating health-hospitalization example in MHE, it is clear that an ill person cannot cause health improvements by deviating and avoiding hospitalization. Here observational evidence is indeed subject to “selection bias” as a guide to decisions. In contrast, all firms, including weak firms, can cause stock price increases by signaling and mimicking strong firm policies. Here, as we show, observational evidence can deliver an unbiased estimate of causal implications of actions. Similarly, signaling and screening are predicted to differ in terms of equilibrium distortions/contracts which future granular empirical work can examine.³ Despite such differences, our paper argues for renewed interest in estimation of signaling effects with an eye toward informing CFO decisions, just as recent years have seen a resurgence in estimation of selection effects in insurance and lending markets (e.g. Einav and Finkelstein (2011), Hertzberg, Liberman and Paravisini (2018)) with an eye toward informing public policy decisions.

Finally, our paper is similar in spirit to Prabhala (1997) in using the rational expectations paradigm to interpret competing estimators. Prabhala analyzes the merits and limitations of standard event studies versus conditional event studies. In contrast, we focus on the merits and limitations of IV-type estimators versus OLS, event studies and conditional event studies, again stressing decision-making CFOs. The papers share an underlying theme in advocating making explicit assumptions about information sets.

The remainder of the paper proceeds as follows. Section 2 presents the instrumental variables example and Section 3 presents the fuzzy RD example. Section 4 considers a more complex setting with dynamic policy shocks. Finally, Section 5 revisits the data in Bowen, Frésard, and Taillard (2017) in order to give the reader a sense of the applicability of our arguments within the empirical corporate finance literature.

³See Salanie (1997) for a concise treatment.

2 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 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 . Further, let $b_j \equiv \ln B_j$ and let $\theta_j \equiv \ln \Theta_j$.

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:

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

Firm log 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 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

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

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:

$$\tau_j = 2LB_j^* \Theta_j^{-1} \Rightarrow B_j^* = \frac{\tau_j \Theta_j}{2L} \Rightarrow b_j^* = \theta_j + \ln(\tau_j) - \ln(2L).$$

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, with firm's tax shield parameters being i.i.d. draws from a uniform distribution with support $[0, 1/2]$. The bankruptcy reputational loss is set to $L = 1/2$. Finally, the firm quality parameter is uniformly distributed on $[0, 1]$.

Panel A of Table 1 presents regression output when the econometrician employs OLS and regresses credit ratings λ_j directly on 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 1 presents the econometrician's preferred regression output, that relying on the tax shield parameter as an instrument for 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 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 in this economy should a CFO choose to issue more debt? Figure 1 suggests an answer, plotting bond ratings as a function of debt level for firms in the bottom and top tax shield parameter quintiles. The top (bottom) cluster captures firms in the top (bottom) tax shield quintile. Apparently, bond ratings are invariant to debt levels within each quintile. Panel C of Table 1 confirms this visual impression, regressing bond ratings directly on b_j along with dummy variables for tax shield parameter quintiles. Notice, the R-squared

is 0.96, a near perfect fit in this stylized economy. Importantly, the coefficient on b_j is negligible and statistically insignificant. Given the high R-squared here, it is apparent that firms within each tax shield quintile have similar bond ratings, with the bond ratings being insensitive to debt levels within each quintile. For the CFOs in this economy, this OLS output would have direct relevance for decision-making: “Conditional upon variables observable to investors, equilibrium bond ratings are actually insensitive to changes in debt levels.” Or, more simply, “If I change my company’s leverage ratio, its bond rating actually will not change.”

To understand the empirical evidence here, note that in her attempt to describe the decision problem of the CFO (equation (2)), the econometrician has actually 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. A correct specification of the capital structure optimization problem here reflects the fact that the latent variables problem confronting the econometrician also confronts investors. Thus, as in Ross (1977), the original firm sponsor will use their private knowledge of Θ_j in computing the default probability, but must use rating agency and investor beliefs $\widehat{\Theta}_\tau(B)$ (given an observable tax shield parameter τ) to determine the market’s valuation of the cash flows accruing at $t = 1$. Thus, the true optimum financial policy for CFOs in this latent variables economy is:

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

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 infer firm quality increases linearly in debt, with the slope of the 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.⁴

In estimating log default probabilities in this latent variables economy, the rating agencies must use their inference regarding firm type (6) rather than the actual firm type, so that observed ratings are properly understood as taking the form:

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

For the CFO, the significance of the preceding equation is that conditional upon observables (the tax shield value τ_j), the firm's bond rating will not vary with the level of debt it chooses. But note, this was the exactly the message delivered by the OLS regression output. To understand this, consider that the true empirical relationship is $\lambda_j = b_j - \widehat{\theta}$. Using the fact that the IV coefficient delivers an estimate of $\partial\lambda/\partial b = 1$, the OLS coefficient from a regression of λ on b can be written as:

$$\beta_{\lambda b}^{OLS} = \frac{cov(b - \widehat{\theta}, b)}{var(b)} = 1 - \beta_{\widehat{\theta} b}^{OLS} \approx \beta_{\lambda b}^{IV} - \beta_{\widehat{\theta} b}^{OLS}. \quad (9)$$

As illustrated by the preceding equation, the OLS coefficient captures the total causal effect of debt on bond rating, consisting of the partial causal effect and the effect of debt on inferred firm quality ($\beta_{\widehat{\theta} b}^{OLS}$). Stated differently, the wedge between the OLS and IV coefficients here delivers an estimate

⁴Notice, in this LCSE the lowest type issues zero debt, just as under the perfect information B^* .

of the magnitude of signaling effects, or how inferred firm type varies with the choice of debt in equilibrium.

Returning to equation (3), we see that the essential piece of outside information required for the CFO to choose optimal policies in this latent variables economy is the market inference function $\hat{\Theta}_\tau$. In other words, as in reality, the CFO must correctly anticipate the debt and equity valuations the market will assign, as functions of chosen debt levels. The OLS output in Panel C of Table 1 delivers just that information, showing how pricing schedules will vary according to the observables here, the tax shield parameter. In particular, using the respective intercept term for the firm's tax shield parameter quintile, the CFO can use equation (6) to infer the following expressions for market assessed enterprise value:

$$EV = B \exp(-INTERCEPT). \quad (10)$$

The point of the preceding analysis is to illustrate in a stark way potential shortcomings of standard IV estimates from the perspective of decision-making CFOs. In contrast, OLS output delivered to the CFO the information needed to make optimal decisions here, e.g. firms face flat rating-debt schedules conditional upon observables. Intuitively, the CFOs in this economy need to know what consequences *en toto* will follow from their discretionary actions, and OLS with an endogenous regressor reveals this. In contrast, IV strips out the signaling component of discretionary decisions.

Another subtle point illustrated by the example is that omitted variables, per se, are not necessarily a problem for OLS. Here firm quality was an omitted variable, yet the OLS coefficient captured the total causal effect. What this stylized example did not consider was the potential existence of other variables observable to all agents in the economy other than the econometrician. Such variables would generate an omitted variables bias in the OLS estimate of total causal effects. In this case, one would want to capture the benefit of IV strategies (robustness to unwanted omitted variables) while somehow capturing the total causal effect estimate needed by decisionmakers (here CFOs). A solution suggests itself when we consider that the problem with the instrument in this

parable economy, the tax rate proxy, was that it was observable to all parties. Thus, changes in leverage induced by changes in tax shields would be understood to be quasi-exogenous rather than discretionary. If a more opaque parameter, say reputational costs of bankruptcy L , varied in a way perceived by CFOs/econometricians but somehow unperceived by rating agencies, then this instrument L would capture total causal effects since rating agencies would misinterpret quasi-exogenous leverage variation as being discretionary.

Thus, we suggest that searching for such “invisible instruments” is a potentially fruitful and novel direction for future empirical work seeking to inform atomistic decision-makers such as CFOs making discretionary decisions.⁵ Indeed, in the future one could envision firms experimenting on their own accounts to infer total causal effects. For example, a firm could randomize its net debt by randomizing the schedule at which it distributed excess cash. Similarly, a firm could randomize its schedule of share repurchases. Of course, such activity would need to be perceived as probability zero or the inference procedure would be complicated.

The more general message delivered by the analysis is that the nature and interpretation of IV estimates is contingent upon subtle and granular descriptions of information sets. In addition to exclusion restrictions, observability assumptions merit debate and discussion in empirical work.

3 Fuzzy Regression Discontinuity

3.1 Technology and Timing

In this section, we revisit the recent paper by Dittmar, Duchin and Zhang (DDZ below) (2020) who use fuzzy RD 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

⁵We thank Gustavo Manso and Vikrant Vig for suggesting this discussion.

of equity-financed firms with initial shares outstanding at time $t = 0$ normalized at 1.⁶ 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 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 any new stock. In the spirit of the the fuzzy regression discontinuity design employed by DDZ, the board's SEO approval probability $\rho(I)$ is assumed to be an increasing step function.⁷ 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

⁶For simplicity, debt finance is ignored. One could assume firms face covenants prohibiting additional debt.

⁷Private benefits can make insiders reluctant to dilute control. Behavioral explanations are abundant.

exogenous random variable whereas price ratios are determined endogenously, perhaps complicating interpretation.

If the board approved new stock issuance, 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 $[0,2]$.

3.2 Stock Market Equilibrium

In an SEO, the corporation will receive investment funding i equal to shares issued times the equilibrium share price:

$$i(s; I) \equiv sp_3(s; I). \quad (11)$$

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:

$$\max_{0 \leq s \leq 1} \frac{\alpha + 2i(s, I)}{1 + s} = \frac{\alpha + 2sp_3(s, I)}{1 + s}. \quad (12)$$

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 (12) is:

$$2i_s(s, I) = (1 + s)^{-1}[\alpha + 2i(s, I)]. \quad (13)$$

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

$$i(s, I) = \left(\frac{s}{1+s} \right) [\alpha + 2i(s, I)]. \quad (14)$$

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

$$2si_s(s, I) = i(s) \Rightarrow i(s, I) = A(I)\sqrt{s}. \quad (15)$$

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

$$i(1, I) = I \Rightarrow i(s, I) = I\sqrt{s} \Rightarrow p_3(s, I) = \frac{i(s, I)}{s} = \frac{I}{\sqrt{s}}. \quad (16)$$

It follows from the preceding equation that the first-order condition (13) pinning down the optimal discretionary choice of s can be written as

$$Is^{-1/2} = (1+s)^{-1}[\alpha + 2Is^{1/2}]. \quad (17)$$

From the preceding equation it follows that companies with more valuable assets-in-place will choose to issue fewer shares. In particular:⁸

$$s^*(\alpha, I) = \left[-\frac{\alpha}{2I} + \frac{1}{2} \sqrt{\left(\frac{\alpha}{I} \right)^2 + 4} \right]^2. \quad (18)$$

3.3 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 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

⁸Let $\Omega = s^{1/2}$ and solve the quadratic in Ω , then compute $s = \Omega^2$.

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.

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 (14), implies that the intrinsic value of the original share outstanding can be rewritten as:

$$\frac{\alpha + 2i(s, I)}{1 + s} + \left(\frac{s}{1 + s} \right) [\alpha + 2i(s, I)] - i(s, I) = \alpha + i(s, I). \quad (19)$$

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, 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. In particular, simulated stock prices at $t=2$ react positively to board allowance of an SEO. Moreover, consistent with the findings in DDZ, the average abnormal return for simulated firms above the price ratio=1 threshold is higher (39.78%) than for the simulated 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.

But note, the preceding clean evidence is of no value to the CFOs in this parable economy who must decide at $t=3$ on just how much stock to issue—after consulting with their investment banks. Rather, the preceding evidence could be misleading if interpreted naively.⁹ After all, if stock issuance has a positive causal effect on stock price, this might be seen to imply that the CFO should issue the maximum feasible number of shares, setting $s = 1$. Of course, this reasoning neglects the fact that, as shown, 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 strip out the signaling channel of stock issuance.

What would be useful to a CFO standing at $t = 3$ in this parable economy would be evidence from conditional event studies examining stock price reactions to *endogenous* decisions made by prior CFOs who found themselves at the same 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 (12) 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

⁹We consider equilibria, and so treat such mistakes as non-existent or measure zero.

representing observable conditioning information, here maximum feasible investment scale. To this end, a traditional conditional corporate finance event study would be informative.

Figure 3 illustrates, depicting equilibrium stock prices (equation (16)) from simulated event studies (e.g. Asquith and Mullins (1986)) for two sets of firms issuing shares of stock—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 RD identification strategy. First, stock price responds negatively to marginal increases in the number of shares issued. Second, apparently the CFO will face a steeper stock price reaction function if maximum investment scale is higher.

Before closing it is worth noting that this analysis should not be interpreted as implying that the evidence provided by DDZ has low value. Rather, we have presented a stylized example in which a conditional corporate finance event study would have value to the decision-maker and the RD study would not. In reality, both the SEO event study and RD evidence could have independent value to decision-makers. To see this, recall that in order to capture the spirit of the DDZ fuzzy RD approach, we assumed that boards randomized the SEO approval process, with a sharp discontinuity in the approval probability at $I = .50$. If at future dates boards had the ability to optimize their SEO approval process, the evidence coming from the simulated fuzzy RD would reveal that boards were not behaving optimally since stock price is increasing in SEO approval probability, a fact revealed by Figure 2.

Thus, in practice the conjunction of alternative forms of evidence may be needed to make sequentially optimal decisions. In particular, evidence stripped of signaling effects (e.g. RD) can here be seen as useful to those making decisions before the arrival of private information while evidence cum signaling effects (event studies) can here be seen as useful to those making decisions after the arrival of private information.

4 Dynamic Natural Experiments

Natural experiments can be viewed as a special case of IV estimation in which the causal variable is discrete. Indeed, when available, natural experiments are viewed as an attractive source of causal inference. An inherent feature of many natural experiments is their time dimension—shocks are taking place over time. With this in mind, the present section considers the interpretation and utilization of evidence from dynamic natural experiments. In particular, we revisit 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.

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

$$dK = (I - \delta K)dt \tag{20}$$

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

$$\frac{dX_t}{X_t} = \mu_t dt + \sigma dW \tag{21}$$

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.

4.1 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

$$\mu(\mathbf{Z}) \equiv Z_1\mu_1 + Z_2\mu_2. \quad (22)$$

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:

$$dZ_n = \underbrace{p(f_n - Z_n)dt}_{\equiv \mu_{z_n}} + \underbrace{\frac{Z_n[\mu_n - \mu(\mathbf{Z})]}{\sigma}}_{\equiv \sigma_{z_n}} d\tilde{W}. \quad (23)$$

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

$$\begin{aligned} & rV(K, X, S, \mathbf{Z}) \\ = & \max_I \Pi_S K X - I - \gamma I^2 + V_k(I - \delta K) + V_x \mu(\mathbf{Z}) X + \mu_{z_1} V_{z_1} + \mu_{z_2} V_{z_2} \\ & + \lambda[V(K, X, S', \mathbf{Z}) - V(K, X, S, \mathbf{Z}) + \frac{1}{2} V_{xx} \sigma^2 X^2 + \frac{1}{2} V_{z_1 z_1} \sigma_{z_1}^2 + \frac{1}{2} V_{z_2 z_2} \sigma_{z_2}^2 \\ & + V_{z_1 z_2} \sigma_{z_1} \sigma_{z_2} + V_{xz_1} X \sigma \sigma_{z_1} + V_{xz_2} X \sigma \sigma_{z_2}. \end{aligned} \quad (24)$$

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:

$$rV(K, X, S, \mathbf{Z}) = KQ(X, S, \mathbf{Z}) + G(X, S, \mathbf{Z}). \quad (25)$$

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

$$\max_I IQ(X, S, \mathbf{Z}) - I - \gamma I^2 \Rightarrow I^*(X, S, \mathbf{Z}) = \frac{Q(X, S, \mathbf{Z}) - 1}{2\gamma}. \quad (26)$$

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 :

$$\begin{aligned} & (r + \delta + \lambda)Q(X, S, \mathbf{Z}) \\ &= \Pi_S X + \mu(\mathbf{Z})XQ_x + \mu_{z_1}Q_{z_1} + \mu_{z_2}Q_{z_2} + \lambda Q(X, S', \mathbf{Z}) \\ & \quad + \frac{1}{2}\sigma^2 X^2 Q_{xx} + \frac{1}{2}\sigma_{z_1}^2 Q_{z_1 z_1} + \frac{1}{2}\sigma_{z_2}^2 Q_{z_2 z_2} \\ & \quad + \sigma_{z_1} \sigma_{z_2} Q_{z_1 z_2} + X\sigma\sigma_{z_1} Q_{xz_1} + X\sigma\sigma_{z_2} Q_{xz_2}. \end{aligned} \quad (27)$$

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

$$\begin{aligned} [r + \delta - \mu_1 + \lambda + pf_2]\Psi_{S1}^1 &= \Pi_{S1} + pf_2\Psi_{S1}^2 + \lambda\Psi_{S2}^1 \\ [r + \delta - \mu_2 + \lambda + pf_1]\Psi_{S1}^2 &= \Pi_{S1} + pf_1\Psi_{S1}^1 + \lambda\Psi_{S2}^2 \\ [r + \delta - \mu_1 + \lambda + pf_2]\Psi_{S2}^1 &= \Pi_{S2} + pf_2\Psi_{S2}^2 + \lambda\Psi_{S1}^1 \\ [r + \delta - \mu_2 + \lambda + pf_1]\Psi_{S2}^2 &= \Pi_{S2} + pf_1\Psi_{S2}^1 + \lambda\Psi_{S1}^2. \end{aligned} \quad (28)$$

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

$$\Pi_{S2} > \Pi_{S1} \Rightarrow \Psi_{S2}^1 > \Psi_{S1}^1, \Psi_{S2}^2 > \Psi_{S1}^2. \quad (29)$$

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*

$$SR_{SS'} = \frac{X}{2\gamma} \times [Z_1(\Psi_{S'}^1 - \Psi_S^1) + Z_2(\Psi_{S'}^2 - \Psi_S^2)]. \quad (30)$$

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.

4.2 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:

$$\begin{aligned}
& rV(K, X, S) \\
&= \max_I \quad \Pi_S K X - I - \gamma I^2 + V_k(I - \delta K) + V_x \mu_S X + \\
&\quad + p(1 - f_S)[V(K, X, S')] - V(K, X, S)] + \frac{1}{2} V_{xx} \sigma^2 X^2.
\end{aligned} \tag{31}$$

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

$$V(K, X, S) = Kq(X, S) + g(X, S). \tag{32}$$

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

$$\max_I \quad Iq(X, S) - I - \gamma I^2 \Rightarrow I^*(X, S) = \frac{q(X, S) - 1}{2\gamma}. \tag{33}$$

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:

$$[r + \delta + p(1 - f_S)]q(X, S) = \Pi_S X + \mu_S X q_x(X, S) + \frac{1}{2} \sigma^2 X^2 q_{xx}(X, S) + p(1 - f_S)q(X, S'). \tag{34}$$

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:

$$\begin{aligned}
[r + \delta - \mu_1 + p(1 - f_1)]\psi_1 &= \Pi_{S1} + p(1 - f_1)\psi_2 \\
[r + \delta - \mu_2 + p(1 - f_2)]\psi_2 &= \Pi_{S2} + p(1 - f_2)\psi_1.
\end{aligned} \tag{35}$$

Solving this system we find:

$$\begin{aligned}
\psi_1 &= \frac{\Pi_{S1}}{r + \delta - \mu_1} \\
&+ \frac{p(1 - f_1)(r + \delta - \mu_2) [\Pi_{S2}/(r + \delta - \mu_2) - \Pi_{S1}/(r + \delta - \mu_1)]}{(r + \delta - \mu_1)(r + \delta - \mu_2)[1 + p(1 - f_1)/(r + \delta - \mu_1) + p(1 - f_2)/(r + \delta - \mu_2)]}
\end{aligned} \tag{36}$$

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 (36) 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*

$$\frac{\Pi_{S2}}{\Pi_{S1}} > \frac{r + \delta - \mu_2}{r + \delta - \mu_1}. \quad (37)$$

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 partial 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:

$$\frac{\Pi_{S2}}{r + \delta - \mu_2} > \frac{\Pi_{S1}}{r + \delta - \mu_1} \Rightarrow q(X, S2) > q(X, S1). \quad (38)$$

This example illustrates how, in yet another setting, both forms of causal effect estimates may be necessary operationally. In particular, Proposition 1 showed how the response of investment to exogenous random policy shocks could be used to determine the relative magnitudes of policy parameters Π_{S1} and Π_{S2} . However, inference regarding the absolute magnitudes of these two policy parameters was clouded by the fact that investment shock response magnitudes depended upon the latent beliefs (\mathbf{Z}) of firms regarding aggregate drift. After all, with an exogenous policy process, firms remain ignorant of latent macroeconomic drift. By way of contrast, once the government in the endogenous policy economy implements its strategy of leaning against the wind, the signal sent by the government reveals the drift to firms, removing beliefs \mathbf{Z} as a source of opacity. That is, the signaling government knows what firms believe about drift. This would then allow the signaling government to infer the absolute magnitudes of Π_{S1} and Π_{S2} from observed responses of investment

to its countercyclical policy changes—revealing both partial and total causal effects. Thus, here we have a stark example in which the government’s signaling activity can actually be seen as aiding causal inference.

5 Revisiting Bowen, Frésard and Taillard (2017)

The preceding three sections of this paper offered a detailed deconstruction of three high-quality papers published in top-tier journals, using parable economies to flush out subtle questions of interpretation and practical utility. The objective of this section is to provide the reader with a more general sense of the relevance of our arguments. Of course, such assessments are necessarily subjective. Thus, greater detail on each of the 253 papers included in our final sample, and the basis for classifications is provided in the Online Appendix.

Our first task was to assemble a reliable data set reflective of the literature. Conveniently, Bowen, Frésard and Taillard (BFT below) (2017) were kind enough to provide us with their data.¹⁰ 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

¹⁰We thank these authors for the generosity.

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 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 2, the BFT sample contains 253 empirical corporate finance papers in the top three finance journals from 2000-2012 using/mentioning either IV, RD or an experiment: 243 used IV, 10 used RD (including three using both RD 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, RD 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 our signaling critique to apply, 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. Further, in order for our signaling critique to apply, 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 2, applying these criteria, our critique applied to 130 out of the 185 papers in our final sample, or 70%. As shown in Table 2, our critique most commonly applies to papers in which the dependent variable concerns valuation/performance and the explanatory variable concerns corporate governance or financial structure.

By way of contrast, in many of the remaining 55 sample papers in which our critique does not apply, the main explanatory variable was not a unilateral choice variable or the dependent variable was not mediated by investor or counterparty beliefs. For example, causal variables such as an industry Herfindahl index, network density, number of analysts covering a firm, or cash flow shocks are not unilateral decisions. In these papers too, many dependent variables use performance or pricing measures that are mediated by investor beliefs, as is common in finance research, but the nature of the explanatory variable suggested our critique did not apply. In other papers, the main dependent variable had no obvious first-order link with investor or counterparty beliefs, e.g. measures of country-level financial development or legal protections.

Two caveats are in order. First, just how strong the signaling effect is likely to be 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.

To be more specific, consider the commonly-studied case (Table 2) in which the dependent variable is Tobin's Q and the explanatory variable concerns corporate governance. From the perspective of an individual firm's board interested in maximizing share price, the total causal effect would seem to be of most interest, since signaling effects will be part of the total causal chain following a unilateral discretionary change in governance structure. On the other hand, from the perspective of a regulator interested in cleaning up markets after a wave of corporate scandals, the partial causal effect recovered by an IV strategy would be of interest, since a government policy imposed on all firms will have zero signal content, leaving only the moral hazard channel. Similarly, individual CFOs making unilateral capital structure decisions will be most interested in total causal effects, whereas a bank regulator imposing a leverage cap might be most interested in estimates of partial causal effects.

6 Conclusion

This paper questions the notion that evidence from instrumentation is, as a general matter, more credible than traditional forms of evidence in corporate finance such as event studies, conditional event studies, or even endogeneity-plagued OLS. Rather, we demonstrate the practical utility of alternative estimators by building parable economies embedding econometricians, investors, and privately-informed decision-making CFOs. As shown, IV estimators have their place, but can also be misleading since randomization through observable instruments eliminates signal content arising from discretion. In contrast, if the goal is informing discretionary decisions, OLS or event studies may well be sufficient. We recommend, by way of illustration, estimation in parable economies as

a precursor to real-world IV estimation as a means of clarifying the interpretation and utility of competing estimators.

References

- [1] Angrist, Joshua D. and Jorn-Steffen Pischke, 2009, *Mostly Harmless Econometrics: An Empiricist's Companion*, Princeton University Press.
- [2] Asquith, Paul and David Mullins, 1986, Equity Issues and Offering Dilution, *Journal of Financial Economics* 15 (1), 61-89.
- [3] Banerjee, Abhijit V., Sylvain Chassang and Erik Snowberg, 2017, Decision Theoretic Approaches to Experiment Design and External Validity, *Handbook of Field Experiments*, Vol. 1, Edited by A. Banerjee and E. Duflo, New York, Elsevier.
- [4] Banerjee, Abhijit V., Sylvain Chassang, Sergio Montero, and Erik Snowberg, 2020, A Theory of Experimenters: Robustness, Randomization, and Balance, *American Economic Review* 110 (4), 1206-1230.
- [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] Chassang, Sylvain, Erik Snowberg, Ben Seymour and Cayley Bowles, 2015, Accounting for Behavior in Treatment Effects: New Applications for Blind Trials, *PLOS One* (June 10).
- [7] Chemla, Gilles, and Christopher A. Hennessy, 2019, Controls, Belief Updating, and Bias in Medical RCTs, *Journal of Economic Theory*.
- [8] 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.
- [9] 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.

- [10] Einav, Liran and Amy Finkelstein, 2011, Selection in Insurance Markets: Theory and Empirics in Pictures, *Journal of Economic Perspectives* 25 (1), 115-138.
- [11] Fama, Eugene, Lawrence Fisher, Michael Jensen, and Richard Roll, 1969, The Adjustment of Stock Prices to New Information, *International Economic Review* 10, 1-28.
- [12] Fudenberg, Drew and David K. Levine, 2020, Learning in Games and the Interpretation of Natural Experiments, forthcoming in *American Economic Journal: Microeconomics*.
- [13] Gourio, Francois, 2013, Credit Risk and Disaster Risk, *American Economic Journal: Macroeconomics* 5 (3), 1-34.
- [14] Hertzberg, Andrew, Andres Liberman and Daniel Paravisini, 2018, Screening on Loan Terms: Evidence from Maturity Choice in Consumer Credit, *Review of Financial Studies*, 3533-3567.
- [15] Krasker, William, 1986, Stock Price Movements in Response to Stock Issues under Asymmetric Information, *Journal of Finance*.
- [16] Leland, Hayne E., 1994, Corporate Debt Value, Bond Covenants, and Optimal Capital Structure, *Journal of Finance*.
- [17] Molina, Carlos A., 2005, "Are Firms Underleveraged? An Examination of the Effect of Leverage on Default Probabilities, *Journal of Finance* 60 (3), 1427-1459.
- [18] 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.
- [19] Prabhala, N.R., 1997, Conditional Methods in Event Studies and an Equilibrium Justification for Standard Event-Study Procedures, *Review of Financial Studies* 10 (1), 1-38.
- [20] 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.

- [21] Ross, Stephen, 1977, The Determination of Financial Structure: The Incentive Signaling Approach, *Bell Journal of Economics*, 23-40.
- [22] Salanie, Bernard, 1997, *The Economics of Contracts: A Primer*, MIT Press.
- [23] Veronesi, Pietro, 2000, How Does Information Quality Affect Stock Returns? *Journal of Finance*, 807-837.
- [24] Young, Alwyn, 2019, Consistency without Inference: Instrumental Variables in Practical Application, working paper.

Table 1: 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

Table 2: 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

Figure 1: Leverage and Bond Ratings



Figure 2: Causal Effect of SEOs using Fuzzy RDD




Figure 3: Issuance Volume and Stock Price

