UNIVERSITA' COMMERCIALE "LUIGI BOCCONI"

PhD SCHOOL

PhD program in ECONOMICS AND FINANCE

Cycle: XXXVII

Disciplinary Field (code): SECS-P/07

Essays in Disclosure

Advisor: PETER POPE

Co-Advisor (if any): WANLI ZHAO

PhD Thesis by

NICOLA MARIA FIORE

ID number: 1828029

Year 2026

Abstract

This dissertation contains three self-contained papers on disclosure, contracting, and economic behavior. The first two papers study rebate contracting in the U.S. pharmaceutical supply chain. The third paper examines how countries' fiscal capacity affects firms' risk perceptions and investment. The chapters are independent; the first two are directly connected, while the third addresses a separate macro-finance question.

The first chapter provides an empirical evaluation of state laws requiring pharmacy benefit managers (PBMs) to disclose rebates to health plans. Using administrative Medicare Part D data covering all prescription-drug plans from 2016–2025 and exploiting the staggered adoption of transparency laws between 2019 and 2022, the analysis tests whether increased visibility over rebate flows reduces consumer costs. Across premiums, deductibles, negotiated prices, and out-of-pocket payments, I find no statistically or economically significant effects. Plan formularies, benchmark acquisition costs, and coinsurance structures also remain unchanged. The only systematic response is a cross-state spillover among multi-state plans, which raise prices in states not subject to the law. The evidence indicates that these transparency mandates did not meaningfully alter PBM–plan contracting or reduce drug spending.

The second chapter develops a theoretical explanation for why greater rebate transparency may fail to improve outcomes. I study a contracting environment in which a PBM privately observes realized rebate revenue and exerts costly effort to generate it on behalf of a health plan. Under secrecy, the PBM can retain part of negotiated rebates, limiting pass-through but preserving incentives to exert effort and generate surplus. Under disclosure, realized rebates become verifiable and retained amounts can be recovered through costly enforcement. While transparency increases pass-through conditional on high rebates, it also compresses the PBM's upside and weakens incentives to negotiate aggressively. The model identifies a "transparency hold-up" effect: by strengthening ex post surplus extraction, disclosure can reduce ex ante surplus creation. As a result, mandatory transparency improves welfare only when enforcement is sufficiently effective and effort is not highly sensitive to retained surplus.

The third chapter studies how fiscal capacity shapes firms' sensitivity to disaster risk. Using a global panel of firms and constructing a firm-level measure of exposure to fiscal constraints based on the sales-weighted fiscal space of countries, the chapter examines the onset of the COVID-19 pandemic as an exogenous shock. Firms more exposed to fiscally constrained countries experience significantly larger increases in fiscal-constraint-related risk in earnings-call transcripts, with no comparable increase for other categories of risk. In the pre-pandemic period, the same firms exhibit higher perceived risk, higher internal discount rates used for investment, and lower investment in both capital and R&D. The results document a direct link between countries' fiscal limitations and firm-level disaster vulnerability, providing micro-evidence on how fiscal constraints can depress investment and, ultimately, growth.

Overall, the dissertation provides empirical evidence and theoretical results on three distinct questions: whether rebate transparency lowers drug costs (it does not), why transparency fails when verification is costly, and how sovereign fiscal constraints affect firms' perceived risk and investment behavior.

Can Transparency Reduce Drug Costs?

Nicola Maria Fiore [*]

May 18, 2026

Abstract

High drug prices and opaque middle-man contracting have driven calls for transparency. I examine whether mandates to disclose rebate flows lower prescription drug costs. I find no meaningful reduction in premiums, deductibles, or out-of-pocket payments across different market structures and incentive conditions. Instead, I find that costs can rise in markets not directly regulated—consistent with intermediaries shifting costs across jurisdictions. Contrary to regulators' expectations, transparency alone may not reduce drug spending—and may even redirect it—when contracts span multiple markets.

Keywords: Transparency; Information Disclosure; Vertical Contracting; Intermediation; Regulation.

JEL Codes: D82, L15, L22, L51, I11, I18.

[*]Bocconi University, Department of Accounting. Correspondence: via personal webpage nicola-maria-fiore.gitlab.io/project/. I am especially grateful to my advisors, Peter Pope and Wanli Zhao, for their guidance and support. I thank Joseph Ross for providing the Quarterly Prescription Drug Plan Formulary, Pharmacy Network, and Pricing Information (Medicare Part D) files used in this analysis. I thank Oriana Ciani, Simone Ghislandi, John Kepler, Maureen McNichols, Joseph Ross, Delphine Samuels, Regina Wittenberg-Moerman and Youchang Wu for their valuable feedback and thoughtful suggestions on earlier drafts of this paper. Portions of this paper were developed while I was visiting the University of Oregon and Stanford GSB, whose faculty and researchers provided invaluable feedback. I also thank seminar participants at University of Oregon, Stanford GSB, Yale School of Medicine, the EuHEA annual conference, EuHEA PhD conference, Università della Svizzera italiana, Bocconi University, and EAA Annual Congress for constructive comments and stimulating discussions.

1 Introduction

Prescription drug spending places a persistent burden on U.S. households, yet the prices patients face are shaped as much by opaque contracting as by manufacturers' list prices. In particular, rebates negotiated between drugmakers, pharmacy benefit managers (PBMs), and health plans determine who pays what, but these flows are largely hidden from buyers and regulators. Policymakers have therefore turned to transparency mandates, expecting disclosure to discipline intermediaries and lower prices. Economic theory, however, is ambiguous: making information observable can sharpen competition by reducing asymmetries, or it can shift incentives and facilitate coordination in concentrated, vertically integrated markets. This paper asks whether greater transparency in rebate contracting reduces what patients pay. The answer matters both for the design of disclosure policy and for understanding pricing in complex supply chains.

I examine whether transparency in pharmaceutical rebate contracting lowers drug costs for patients. I study a series of state laws that, beginning in 2019, required pharmacy benefit managers (PBMs) to disclose detailed information on rebates and administrative fees to health plans. These laws create a natural experiment: they targeted a specific layer of the supply chain, changed information flows without directly regulating prices, and were introduced at different times across states. Using comprehensive administrative data on prescription drug plans, I trace how prices, premiums, and patient cost sharing evolved before and after adoption. This design isolates how greater visibility into rebate flows affects the costs ultimately borne by consumers. By linking policy-driven changes in information to equilibrium outcomes, the analysis provides direct evidence on whether transparency alone can realign pricing incentives in one of the most opaque markets in the U.S. economy.

The rebate-transparency setting provides a powerful laboratory to study how information disclosure affects competition and welfare. Economic theory offers clear but contrasting predictions: when disclosure reduces information asymmetry, it can sharpen price competition and improve efficiency; when it alters strategic behavior or facilitates coordination, it can instead raise prices or reallocate rents. Rebate transparency laws isolate this tension in a market where both forces are present—high concentration and complex contracting on one hand, but large potential gains from competition on the other. Because the reforms changed what plans and intermediaries know without constraining prices or quantities, they provide a clean test of how greater information flows shape equilibrium outcomes in a major consumer market.

The analysis uses comprehensive administrative data covering nearly the entire U.S. prescription drug plan market to examine how greater transparency affects prices. The data link detailed plan design—premiums, deductibles, negotiated drug prices, and patient cost sharing—to the timing of state rebate-transparency mandates adopted between 2019 and 2022. Because these laws required pharmacy benefit managers to disclose rebate information but did not constrain prices or coverage, they provide a clean test of how information disclosure alone reshapes competition and bargaining. The empirical strategy compares the same plans before and after adoption, relative to otherwise similar plans in states that had not yet implemented a law, and exploits the staggered rollout to separate the effect of transparency from broader market trends. Event-study estimates show flat pre-trends and precise post-adoption coefficients centered around zero, indicating that the identifying assumption of parallel trends is well supported.

The results show that greater transparency did not reduce what patients pay. Premiums, deductibles, and out-of-pocket payments are unchanged following adoption, and the estimated average effect on total annual spending is less than one percent of the mean. Point estimates are close to zero and confidence intervals narrow, indicating that transparency did not produce economically meaningful changes in patient costs. Mechanism tests reveal why. Negotiated prices between plans and manufacturers, benchmark acquisition costs, and the structure of formularies all remain stable, implying that disclosure neither intensified price competition nor altered the division of surplus along the supply chain. Plans with stronger incentives to pass through savings—those facing binding medical loss ratio constraints—or with greater bargaining power do not react differently, nor do those in more competitive markets. Across every dimension of market structure and incentives, behavior remains unchanged. These patterns suggest that transparency, while increasing information, left the underlying bargaining environment intact.

The one exception occurs across state borders. Multi-state plans that were indirectly exposed to transparency through operations in adopting states raised prices or cost sharing in untreated states by roughly seven percent. This spillover points to a broader equilibrium response: when transparency reduces flexibility in one market, firms adjust prices elsewhere to preserve margins. Together, the evidence shows that disclosure did not strengthen competition or improve pass-through; instead, it prompted cost reallocation consistent with market power in a vertically integrated system. Transparency increased oversight but not affordability—an outcome that clarifies how information interacts with incentives in complex contracting environments.

Literature and contribution. I contribute to the literature on transparency and competition. Transparency is increasingly promoted as a market-based alternative to price regulation (Weil et al., 2013), yet economic theory remains divided on whether disclosure disciplines or distorts competition. Classic models show that information sharing among firms can either enhance efficiency or sustain coordination, depending on market structure and strategic interaction (Hwang and Kirby, 2000). When disclosure reduces uncertainty about rivals' costs or demand, it may sharpen price competition, but when firms produce complements or compete repeatedly, it can facilitate tacit collusion or rent alignment (Darrough, 1993; Bourveau et al., 2020b; Goncharov and Peter, 2019). Work on vertical contracting similarly shows that revealing wholesale or rebate terms can alter rent distribution and coordination along supply chains (Arya and Mittendorf, 2011). Across these theories, transparency is not uniformly pro-competitive: it can lower prices by reducing asymmetry or raise them by changing strategic behavior. This paper provides new evidence on this central question in a concentrated, vertically integrated industry where both mechanisms are plausible.

Second, the paper contributes to the accounting and information economics literature on the real effects of disclosure regulation. Much of this work studies how disclosure influences capital markets, yet growing research emphasizes its broader consequences for real decisions and welfare in the economy (Beyer et al., 2010; Leuz and Wysocki, 2016; Kanodia and Sapra, 2016; Christensen et al., 2017; Ball, 2024). Healthcare offers a rich context for these questions because disclosure is frequent, measurable, and linked to both financial and social outcomes. Evidence from this setting shows that disclosure can meaningfully alter behavior but in heterogeneous directions: firms adjust R&D and trial reporting to peer actions (Capkun et al., 2023; Zhang, 2024; Bourveau et al., 2020a), conflict-of-interest disclosure can reduce perceived credibility (Leuz et al., 2023), and public reporting of prices or quality metrics has produced mixed effects on hospital pricing and service quality (Christensen et al., 2020; Eyring, 2020; Kepler et al., 2024). This paper extends that literature in two ways. First, it examines a new form of disclosure—rebate transparency—that targets financial flows within a vertically structured supply chain rather than firm performance or product quality. Second, it connects the accounting perspective on real effects to research on supply-chain transparency and vertical contracting, which studies how disclosure affects coordination, bargaining, and rent allocation across firms (Arya and Mittendorf, 2011; Ellis et al., 2012; Feng et al., 2022; Nathan, 2024). In pharmaceutical markets, intermediaries such as pharmacy benefit managers (PBMs) negotiate rebates with manufacturers and determine how savings are shared with insurers and patients. Analyzing this setting shows how transparency about intermediaries' margins influences pricing and the division of surplus, providing evidence on the allocation of rents and efficiency of intermediation in complex vertical relationships.

The closest related paper is Scanlon (2025), who studies transparency laws across private health plans. My paper differs in two key respects. First, I focus on rebate-specific disclosure mandates—the most stringent form of transparency regulation—rather than general price reporting. Second, I study Medicare Part D, a large public insurance program covering tens of millions of beneficiaries, which provides a standardized and nationally relevant setting. Together with recent calls for research on vertical relationships in drug supply chains (Brot-Goldberg et al., 2022), this analysis extends the evidence on how transparency affects pricing dynamics in regulated pharmaceutical markets. By documenting null effects in this context, the paper shows that disclosure rules targeting rebates—unlike transparency about product quality or access—may fail to lower patient costs when intermediaries retain bargaining power and when insurance design insulates consumers from underlying price changes.

This paper thus speaks to a central question in economics: when does more information change behavior, and when does it simply reveal it? Evidence from rebate-transparency laws shows that disclosure alone may not realign incentives or reduce costs when bargaining power remains concentrated. The findings clarify how market structure mediates the effects of transparency and help frame the limits of information-based regulation.

2 Institutional Setting and Conceptual Framework

2.1 Pharmacy Benefit Managers and the Drug Pricing Chain

Drug prices in the United States are shaped by intermediaries rather than manufacturers or pharmacies. Pharmacy Benefit Managers (PBMs) negotiate rebates with drug producers on behalf of health plans and employers, design formularies that determine coverage, and set patient cost sharing. By doing so, PBMs mediate the flow of money and information across the pharmaceutical supply chain (see Figure 1).

Whether PBMs reduce or raise overall drug spending remains contested. They can lower prices by bargaining with manufacturers but may also extract rents through opaque contracts. Recent research highlights this tension: PBMs can discipline monopolist suppliers yet inflate list prices and patient payments when their incentives diverge from those of insurers and consumers (Che, 2025). Some studies find efficiency gains and improved access (Feng and Maini, 2024); others emphasize rent capture and weaker innovation incentives. Conti et al. (2021) describe PBMs as "double agents" that withhold savings, Ho and Lee (2024) show that tiered formularies amplify rebates without clarifying who benefits, and Agha et al. (2022) document reduced R&D in vulnerable drug classes. Rising rebates have widened the gap between list and net prices (Kakani et al., 2020), increasing patient out-of-pocket costs. In Medicare Part D, rebates in protected classes have grown far less than in comparable markets (Kakani et al., 2024). Vertical integration among PBMs, insurers, and pharmacies further limits competition (Gray et al., 2023), while selective contracting reduces retail prices only where consumers respond to cost sharing (Starc and Swanson, 2021).

Rebates are the primary mechanism through which PBMs capture and redistribute value. Suppose a drug lists for \$100. Without a PBM, a health plan would pay the manufacturer the full price. By pooling purchases, a PBM might negotiate a \$30 rebate in exchange for favorable formulary placement, lowering the plan's effective cost to \$70. The rebate is paid retroactively, and patient copays are typically tied to the list price, not the net cost. Over time, PBMs have retained a growing share of these rebates. Tiered rebate schedules tied to market share often steer coverage toward high-rebate drugs even when cheaper substitutes exist.

The economics of PBM contracting resemble those of financial advisors compensated through supplier commissions. In both markets, intermediaries are paid by producers rather than buyers, creating kickback incentives that distort recommendations. Models of financial advice show that such contingent payments arise endogenously when buyers are uninformed (Inderst and Ottaviani, 2012a,b,c). Disclosure can curb exploitation but may reduce efficiency by weakening screening and information acquisition. Nonlinear and volume-based bonus schemes persist even under stricter regulation (Honda et al., 2024). PBMs operate under the same logic: manufacturers reward placement, while patients and insurers rely on their recommendations. Transparency could improve oversight but need not reduce costs.

2.2 State Rebate-Transparency Laws

Between 2019 and 2022, four states—California (AB 315), Indiana (SB 241), Minnesota (SF 278), and Montana (SB 395)—enacted rebate-transparency statutes that introduced unusually detailed reporting requirements for pharmacy benefit managers (Figure 2, Table A1). Each law mandates purchaser-specific disclosure of rebate and fee information rather than broad, market-wide aggregates. PBMs must provide health plans, on a confidential basis and at regular intervals, with data linking rebate payments and administrative fees to the plan's own utilization. The required disclosures typically include three elements: (i) the aggregate wholesale acquisition cost by therapeutic category, (ii) the aggregate amount of rebates received from manufacturers, explicitly including utilization discounts, and (iii) any administrative or other fees retained by the PBM. These provisions apply across retail, mail-order, and specialty drugs and are delivered under nondisclosure agreements that protect proprietary terms. By making rebate flows observable to plans with this level of granularity, the statutes create a credible change in rebate visibility and, in principle, shift bargaining power toward plan sponsors.

The statutory provisions are enforceable and carry credible compliance mechanisms. Minnesota and Montana authorize daily fines of up to $1,000 per violation, Indiana classifies noncompliance as insurance fraud under SB 241, and California's AB 315 embeds a duty of good faith and fair dealing enforceable through contract law. These provisions make adherence to disclosure requirements a binding condition of PBM operations, ensuring that the resulting increase in rebate transparency represents a substantive regulatory change rather than a nominal reporting exercise.

2.3 Conceptual Framework and Predicted Effects

Rebate transparency aims to lower drug costs by shifting bargaining power from PBMs to health plans. By revealing rebate magnitudes, disclosure should help plans benchmark PBM performance and negotiate better terms, reducing net prices and patient payments. Because rebate contracts are confidential, voluntary disclosure is rare, and PBMs argue that revealing terms would weaken their bargaining position (Shepherd, 2014; Benecurv, 2024). State mandates therefore generate exogenous variation in rebate visibility that would not otherwise occur.

Several institutional features may blunt the effect. PBM concentration limits competitive pressure; dominant intermediaries can sustain margins even under disclosure. The design of Medicare Part D also weakens pass-through incentives: its bidding and reconciliation rules reward plans for minimizing net costs to the federal government rather than lowering beneficiaries' out-of-pocket spending (Medicare and Services, 2023a,b,c). In addition, PBMs can adjust other contract terms—raising administrative fees or narrowing rebate eligibility—to preserve profits.

Within this environment, transparency could still reduce costs if greater disclosure strengthens plan bargaining or intensifies competition among PBMs. Such effects should be most visible where financial or regulatory constraints give plans strong incentives to pass savings through to consumers. Whether these conditions hold, and to what extent transparency changes contracting outcomes, are the questions examined in the empirical analysis that follows.

3 Data and Descriptive Statistics

3.1 Data Sources and Variable Construction

The analysis uses comprehensive administrative data from the Centers for Medicare and Medicaid Services (CMS) covering all Medicare Part D prescription drug plans from 2016 through 2025. The data combine CMS's *Quarterly Prescription Drug Plan Formulary, Pharmacy Network, and Pricing Information* files with linked information on plan characteristics, enrollment, and medical loss ratios (MLRs). Details of data cleaning, identifier harmonization, and variable construction are reported in Appendix B.

The CMS quarterly files report the universe of plan contracts, monthly premiums, annual deductibles, and the full menu of negotiated drug prices and cost-sharing terms offered to beneficiaries. These records describe, for each plan and county, the prices faced by enrollees and the reimbursement amounts paid to pharmacies. The original data are observed at the drug–plan–county–year level. For analysis, I aggregate them to the plan–county–year level using two complementary approaches.

First, I identify the minimum-drug within each plan and county—the National Drug Code (NDC) with the lowest cost sharing within the relevant therapeutic class—and carry its corresponding values for negotiated price, acquisition cost, coinsurance, branded status, and specialty tier. This captures the lowest-cost option available to enrollees in that class. Second, I compute Jevons price indices for negotiated prices, acquisition costs, cost sharing, and coinsurance to summarize average price dynamics across all drugs within a plan. These two aggregation strategies provide complementary measures of plan-level pricing: one reflecting the least costly option available, and the other summarizing overall pricing patterns.

The analysis focuses on the initial coverage phase of Part D—after the deductible is met but before the Initial Coverage Limit—where most prescriptions occur and where cross-plan differences in pricing and cost sharing are most meaningful for consumer incentives. The three primary outcomes are the annual premium (*premiumyearly*), the annual deductible (*deductible*), and the annualized average patient payment (*costsharing*). I also define total annual beneficiary cost,

$$\textit{costyearly} = \textit{premiumyearly} + \textit{deductible} + \textit{costsharing},$$

as a summary measure of plan generosity.

Each plan's formulary is linked to standardized product data from the FDA, RxNorm, and the National Average Drug Acquisition Cost (NADAC) database. These sources identify whether a product is branded or generic and provide benchmark acquisition costs at the pharmacy level. A drug is classified as branded ($\textit{branded} = 1$) if RxNorm lists it as a brand term type, the FDA reports any active patent or exclusivity, or NADAC designates it as brand. Benchmark acquisition costs are annualized as *nadacyearly* and compared with the plan's negotiated prices (*negotiatedpriceyearly*) and patient cost sharing.

The analysis is restricted to drugs in WHO ATC class A10BJ, which includes glucagon-like peptide-1 (GLP-1) receptor agonists such as semaglutide and liraglutide. This class provides a clean setting for studying rebates: it is dominated by branded products but still features active competition among manufacturers, creating scope for large rebates and variation in net prices across plans.

Plan identifiers are normalized and concatenated to form a persistent key (*contractplansegmentid*). Indicators for PBM affiliation (*planpbm*), nonprofit status (*nonprofit*), and single-state operation (*singlestate*) are coded manually from plan names and service areas. Enrollment counts come from the CMS *Medicare Advantage/Part D Contract and Enrollment Data*, and MLRs are drawn from CMS Part D MLR public-use files; *mlrbelow85* equals one when a plan's reported MLR is 85 percent or lower. Treatment indicators are based on the staggered adoption of state rebate-transparency laws documented by the National Academy for State Health Policy. The variables *treatedg*, *cohort*, and *dummygt* capture direct exposure to a law, while *treatedgspill* and *dummygtspill* identify contracts operating in untreated states that also serve treated states in the same year.

The key variables are premiums, deductibles, and patient cost sharing, complemented by negotiated and acquisition prices, coinsurance rates, and indicators for branded and specialty status. These measures, together with plan characteristics and treatment timing, form the basis for the descriptive evidence and regressions that follow. Appendix B provides complete definitions.

3.2 Summary Statistics

Table 1 summarizes the plan–county–year data for 2016–2025, comprising more than 1.2 million observations across nearly all Medicare Part D contracts. The market is large and heterogeneous, offering substantial variation in prices and plan design. Premiums and deductibles vary widely: the mean annual premium is about \$380, but one quarter of plans charge no premium and another quarter exceed \$550. Deductibles average \$285 with comparable dispersion. Average annual patient cost sharing is roughly \$535, producing total annual beneficiary costs of about \$1,200 once premiums and deductibles are included. These magnitudes imply that even modest changes in plan pricing can meaningfully alter what beneficiaries pay.

Negotiated drug prices average \$1,700 per year, while benchmark acquisition costs from NADAC are higher because manufacturers price at bulk package levels. Nearly all covered drugs are branded, and specialty-tier products are rare. The data therefore capture a branded, high-rebate segment of the market—GLP-1 therapies—where transparency could plausibly influence negotiations and pass-through to consumers.

Table 1, Panel B describes plan structure. About 30 percent of plans are affiliated with a PBM, 5 percent are nonprofit, and one quarter operate in only one state. Seven percent report medical loss ratios below the 85 percent threshold. Enrollment ranges from a few dozen beneficiaries in small regional contracts to hundreds of thousands in national ones, and county-level concentration (HHI) averages 0.56. These figures show substantial variation across plans and markets—the variation exploited in the empirical analysis of transparency's effects.

3.3 Pre-Policy Balance and Sample Comparability

Before the adoption of rebate-transparency laws, plans in treated and untreated states were highly similar along observable characteristics. Table 2 reports mean differences in plan costs, pricing variables, and market attributes during the pre-policy period. Panel A shows that premiums, deductibles, and total beneficiary costs are of comparable magnitude across groups, with differences that are small relative to their standard deviations. Although treated states exhibit slightly higher premiums and lower deductibles, these gaps are economically minor. Panel B indicates that the composition of plans is also similar: the shares of PBM-affiliated, nonprofit, and low-medical-loss-ratio plans differ only modestly between groups, and both sets of markets display comparable levels of concentration and enrollment size. Panel C documents the scope of exposure to the policy. Four states—California, Minnesota, Indiana, and Montana—implemented rebate-transparency laws between 2019 and 2022, together accounting for roughly 10% of all plan-county-year observations. Despite the limited number of adopting states, their plans represent large and diverse markets within Medicare Part D. Overall, the evidence in Table 2 indicates that adoption timing was not systematically related to baseline plan generosity or market structure. Treated and untreated states entered the policy period on similar footing, supporting the identifying assumption that, absent the laws, trends in prices and beneficiary costs would have evolved in parallel.

4 Empirical Strategy

4.1 Identification and Estimation

The staggered adoption of state rebate-transparency laws provides a natural setting to estimate how greater disclosure affects plan pricing and patient costs. States enacted these laws at different times between 2016 and 2025, creating variation in treatment timing across otherwise comparable markets. The analysis focuses on plans unaffiliated with pharmacy benefit managers (PBMs), for which information on rebates is likely to change most when transparency increases. PBM-affiliated plans serve as a falsification group in Section 5.

Each observation corresponds to a plan-state-county-year cell (i, s, c, t), where i indexes plans, s states, c counties, and t calendar years. The outcome variables are the main cost components described in Section 3.1: the annual premium, deductible, cost sharing, and their sum (*costyearly*). The treatment indicator $\textit{Transparency}_{st}$ equals one when state s has a rebate-transparency law in effect in year t. Adoption is permanent, so treatment status is weakly increasing over time, and never-adopting states serve as contemporaneous controls. Multi-state contracts are assigned according to the county's state of operation.

The baseline estimating equation is

$$Y_{isct} = \alpha_{isc} + \lambda_t + \beta \, \textit{Transparency}_{st} + \varepsilon_{isct}, \tag{1}$$

where Y_{isct} denotes a pricing or cost outcome for plan i in state s, county c, and year t; α_{isc} are plan-county fixed effects that absorb persistent differences in pricing and generosity; and λ_t are year fixed effects that capture aggregate shocks to drug markets and Medicare regulation. The coefficient β measures the mean contemporaneous change in outcomes when a state implements a transparency law, relative to not-yet-treated states in the same year. Identification therefore comes from within-plan comparisons over time, contrasting the evolution of outcomes in adopting and non-adopting states.

To address heterogeneous treatment effects across cohorts, I implement the estimator of De Chaisemartin and d'Haultfoeuille (2020), which aggregates two-period comparisons between newly treated and not-yet-treated states. This estimator remains unbiased under treatment-effect heterogeneity and avoids the negative-weight and contamination problems that affect conventional two-way fixed-effects regressions. The specification in equation (1) thus recovers the average treatment effect across all adopting states, weighted by enrollment within plan-county cells.

4.2 Empirical Implementation

Identification relies on a standard parallel-trends assumption: absent transparency laws, average changes in premiums, deductibles, and cost sharing would have evolved similarly in adopting and non-adopting states. Because no state repeals its law, only the untreated-trend condition is required. I assess its plausibility using event-time plots and placebo regressions one year before adoption. Consistent with these checks, pre-policy trends are flat and differences between treated and control states are small, as reported in Section 3.3.

Equation (1) includes only plan-county and year fixed effects, matching the structure of a canonical difference-in-differences design. Standard errors are clustered at the *state level*, the level of policy variation, following the guidance in de Chaisemartin and D'Haultfœuille (2023). Clustering at this level accounts for serial correlation and unobserved shocks common to all counties and plans within a state.

The empirical analyses proceed in several steps. The main specification estimates equation (1) for unaffiliated plans to test whether rebate-transparency laws reduce patient drug costs. Section 5 first reports effects on the four core cost outcomes—premiums, deductibles, cost sharing, and total annual costs—providing a direct test of whether transparency lowers patient expenditures.

Subsequent analyses explore mechanisms and heterogeneity in these effects. Mechanism tests re-estimate equation (1) for negotiated drug prices, benchmark NADAC acquisition costs, and coinsurance rates to distinguish whether transparency affects underlying acquisition prices or benefit design. Additional outcomes—the branded and specialty indicators—capture whether plans substitute toward lower-cost drugs, while Jevons price indices track a fixed basket of molecules to isolate negotiation effects on identical drugs rather than substitution across products.

Heterogeneity analyses examine how market incentives shape policy impacts. I test whether effects are stronger among plans facing binding medical-loss-ratio constraints (greater pass-through pressure), among large contracts with greater bargaining power, and in counties with more competitive insurance markets (lower Herfindahl-Hirschman Index). More competition should amplify cost pass-through and strengthen any transparency effect.

Robustness checks restrict the sample to single-state plans to rule out cross-state contamination, and placebo tests examine PBM-affiliated plans, for which no effect is expected. Finally, spillover analyses assess whether contracts operating in both treated and untreated states transmit effects across borders. Together, these analyses provide a comprehensive assessment of whether, how, and for whom rebate-transparency laws influence prescription-drug costs.

5 Results

This section presents the empirical results. I begin by estimating the overall effect of rebate-transparency laws on plan-level pricing and patient costs using the difference-in-differences specification in equation (1). I then examine whether the effects differ across PBM-affiliated and unaffiliated plans, test potential mechanisms related to negotiated prices and formulary design, and explore heterogeneity across plans with different financial incentives and market structures. The final subsection reports robustness checks and spillover analyses. Together, these results provide a comprehensive assessment of how rebate disclosure shapes the cost structure of Medicare Part D plans and beneficiaries' out-of-pocket spending.

5.1 Main Effects: Rebate Disclosure and Patient Costs

Table 3 and Figure 3 report the estimated dynamic effects of state rebate-transparency laws on plan-level costs for non–PBM-affiliated Medicare Part D plans. Event-time coefficients show no discernible changes in premiums, deductibles, or patient cost sharing following the enactment of these laws. Point estimates in the first three years after adoption are economically small and statistically indistinguishable from zero. The average post-adoption effect on total annual beneficiary cost is approximately \$83 below baseline (column 4), corresponding to less than one percent of the sample mean (\$1,227).

Pre-policy coefficients are flat, and joint placebo tests fail to reject equality across pre-treatment periods ($p > 0.3$), indicating no differential trends prior to adoption. Post-policy estimates are also jointly insignificant ($p > 0.4$). Standard errors are clustered at the state level, and regressions are weighted by plan–county–year enrollment. Taken together, these findings suggest that rebate-transparency laws did not measurably affect plan pricing or patients' total out-of-pocket spending during the study period.

The absence of systematic changes across all cost components indicates that increased rebate disclosure neither lowered premiums or deductibles nor altered the cost sharing borne by beneficiaries. This pattern is consistent with the institutional design of the transparency statutes themselves: each law—in California (AB 315, 2017), Indiana (SB 241, 2020), Minnesota (SF 278, 2019), and Montana (SB 395, 2021)—requires disclosure only of aggregate rebates and manufacturer fees, while leaving other revenue channels such as plan-side administrative fees, spread pricing, and affiliated-pharmacy margins unregulated. In effect, the policies increased reporting requirements without necessarily changing the contractual or financial incentives governing PBMs. As discussed in Fiore (2025), such partial or aggregated transparency may not eliminate leakage opportunities and thus might have limited ability to alter equilibrium prices or pass-through behavior. The null estimates here are consistent with this possibility: transparency appears to have illuminated one category of transfers while leaving others open, potentially allowing intermediaries to reallocate rents without reducing total patient costs. The next subsection examines whether this null effect persists when comparing PBM-affiliated and unaffiliated plans in a triple-difference design.

5.2 Falsification and Triple-Difference: PBM vs. Non–PBM Plans

Table A2 compares the estimated effects of rebate-transparency laws for non–PBM and PBM-affiliated plans. Panel A reproduces the average post-adoption estimates for non–PBM plans from Table 3; Panel B reports the corresponding effects for PBM-affiliated plans; and Panel C presents the triple-difference estimates capturing the difference between the two groups. The design serves both as a falsification and heterogeneity test: because PBM-affiliated plans directly manage rebate contracts, they should be less affected by transparency laws that primarily increase information available to unaffiliated plans.

Across all four cost components—premiums, deductibles, cost sharing, and total annual costs—the estimated differences between non–PBM and PBM-affiliated plans are small and statistically insignificant. The triple-difference coefficients are close to zero and imprecisely estimated ($p > 0.15$ in all cases). For example, the post-policy change in total annual beneficiary cost for non–PBM plans is only about $170 lower than that for PBM-affiliated plans (column 4), an economically negligible difference relative to the mean of $1,227.

These findings reinforce the main result: transparency laws did not materially alter plan pricing or patient costs. The absence of differential effects between PBM-affiliated and unaffiliated plans suggests that greater rebate disclosure neither shifted bargaining dynamics nor changed incentives for pass-through within vertically integrated intermediaries. The next subsection explores potential mechanisms by examining whether transparency affected negotiated prices or formulary composition.

5.3 Mechanisms: Formulary and Pricing Channels

Table 5 examines potential mechanisms through which rebate-transparency laws could influence plan and patient costs. The analysis focuses on pricing and formulary outcomes for non–PBM-affiliated plans, testing whether transparency affected negotiated prices, benchmark acquisition costs (NADAC), cost-sharing rates, or formulary composition. Panel A reports level outcomes for the minimum-cost drug within each plan–county–year, while Panel B presents results using Jevons price indices that track a fixed basket of molecules over time.

Across all outcomes, the coefficients are small and statistically indistinguishable from zero. The average post-adoption effects on negotiated prices and benchmark NADAC costs are positive but economically negligible—about \$480 and \$20,700 respectively, corresponding to less than two percent of the sample means. Coinsurance rates decline by only six percentage points, and there is no measurable change in the probability that a plan's lowest-cost option is branded or specialty. Similarly, Jevons price indices show no systematic movement in negotiated or acquisition prices, with estimated changes below half a percentage point. Pre-policy coefficients are flat, and post-policy estimates remain well within confidence intervals spanning zero.

Taken together, the results indicate that rebate-transparency laws did not alter the underlying pricing or formulary channels through which savings could reach patients. The absence of effects on both list and acquisition prices suggests that PBMs and manufacturers maintained preexisting rebate structures. Likewise, the stability of formulary composition implies that disclosure did not shift coverage toward lower-cost or generic alternatives. Overall, transparency increased information but did not translate into changes in contracting, pricing, or benefit design. The next section examines whether these findings differ across plans with varying financial incentives or market competitiveness.

5.4 Heterogeneity Across Plan Incentives and Market Structure

Table 6 investigates whether the effect of rebate-transparency laws varies with plans' financial incentives and competitive environments. I examine three dimensions of heterogeneity: (i) whether plans face binding medical-loss-ratio (MLR) constraints that limit retained profits (Panel A), (ii) whether plans have greater bargaining power as proxied by pre-policy enrollment (Panel B), and (iii) whether they operate in more or less competitive county markets as measured by the Herfindahl-Hirschman Index (Panel C). Each subgroup estimate reflects the average post-adoption effect on total annual beneficiary cost among non–PBM-affiliated plans.

Across all dimensions, the estimated differences are small and statistically insignificant. In Panel A, transparency laws reduce total annual costs by roughly $45 among MLR-constrained plans and $104 among unconstrained ones, with no significant difference between groups ($p = 0.77$). In Panel B, plans with greater bargaining power experience slightly smaller cost changes (-$83) than less powerful plans (-$895), but the difference is again imprecisely estimated ($p = 0.41$). In Panel C, effects are virtually identical in more and less concentrated markets ($p = 0.98$). All confidence intervals are wide and include zero, indicating no systematic heterogeneity across financial or structural conditions.

The absence of heterogeneous effects suggests that transparency did not interact meaningfully with either regulatory incentives or market competition. Plans subject to binding pass-through requirements did not translate disclosure into lower patient costs, and larger or more competitive plans did not leverage additional information to negotiate better terms. These results reinforce the main conclusion: transparency mandates, on their own, are insufficient to alter contracting outcomes even where incentives to pass through savings are strongest. The next subsection reports robustness checks and spillover analyses.

5.5 Robustness: Single-State v. Multi-State Plans

Table 7 examines the robustness of the main findings by comparing two alternative specifications that differ in sample construction. Panel A restricts the estimation to single-state, non–PBM-affiliated plans, thereby removing potential cross-state contamination that could arise when multi-state contracts operate concurrently in treated and untreated markets. Panel B reports corresponding estimates for multi-state plans, and Panel C presents the difference between the two sets of coefficients.

The estimated post-adoption effects are small in magnitude and statistically indistinguishable from zero in both samples. For single-state plans, the average cumulative change in total annual beneficiary cost is approximately $15, or about one percent of the sample mean, while the estimates for multi-state plans are similarly modest. Differences across specifications (Panel C) are economically minor and not statistically significant ($p > 0.25$).

Overall, Table 7 suggests that the conclusions drawn from the main analysis are not sensitive to the inclusion of multi-state contracts. Whether the sample is limited to single-state plans or includes the full set of contracts, there is no evidence that rebate-transparency laws meaningfully affected plan premiums, cost sharing, or total beneficiary costs.

5.6 Spillovers to Untreated States

Table 8 and Figure 4 examine whether rebate-transparency laws generated cross-market spillovers through multi-state plan operations. Because many Medicare Part D contracts span both adopting and non-adopting states, exposure in one jurisdiction may prompt portfolio-wide adjustments in pricing or formulary design. To test for such indirect effects, I re-estimate equation (1) for non–PBM-affiliated plans in untreated states whose contracts also operate in states implementing transparency laws (*treatedgspill*=1, *dummygtspill*=1).

Event-time coefficients reveal positive post-adoption effects concentrated in cost sharing and total annual beneficiary costs. The average spillover effect implies an increase of approximately $830 per enrollee—about 7 percent of the sample mean ($1,212). Although estimates are imprecise, the direction and magnitude suggest that multi-state plans may have raised prices or adjusted cost sharing upward in response to transparency requirements elsewhere. Pre-policy coefficients show no systematic trends, and placebo tests yield mixed results, indicating that the observed pattern emerges primarily after adoption.

These findings suggest that partial transparency may shift, rather than reduce, drug costs when intermediaries operate across states. Because the laws require disclosure only of aggregate rebates and manufacturer fees—while leaving other revenue channels such as administrative or data fees unregulated—PBMs and plans could have reoptimized their portfolios by adjusting unregulated fees in non-adopting states. This potentially reflects fee-shifting behavior: when disclosure constrains one source of revenue, intermediaries may respond by expanding or repricing other, non-disclosed fees. In equilibrium, this mechanism could help explain the higher cost sharing observed for patients in untreated states, even though transparency was intended to lower costs overall.

6 Discussion and Policy Implications

The evidence shows that rebate-transparency laws did not measurably reduce drug costs for Medicare Part D beneficiaries and, if anything, modestly increased costs in markets indirectly exposed through multi-state plan operations. These findings imply that transparency alone—without complementary changes to incentives or market structure—is unlikely to realign pricing behavior in vertically integrated pharmaceutical supply chains.

From an economic perspective, the null results are informative. In theory, transparency can discipline markets by reducing information asymmetry, but it can also have offsetting effects when intermediaries retain market power or face weak pass-through incentives. The stability of plan premiums, deductibles, and negotiated drug prices suggests that PBMs and insurers internalized the disclosure requirements without changing underlying contract terms. In this setting, the marginal information revealed by transparency was insufficient to alter bargaining outcomes or competitive dynamics. This pattern is consistent with models in which disclosure redistributes rents without affecting total surplus when intermediaries can adjust along unregulated margins.

The positive spillover effects observed in untreated states underscore this mechanism. When only a subset of states imposed transparency mandates, multi-state PBMs and insurers appear to have rebalanced pricing across their national portfolios, spreading compliance or profit-adjustment costs to unregulated markets. Such behavior highlights the challenge of implementing partial transparency in a nationally integrated industry: when regulatory exposure is uneven, intermediaries may respond at the portfolio level rather than the jurisdictional level. Similar cross-market adjustments have been documented in financial, energy, and environmental regulation, where firms internalize policy shocks by reallocating prices or margins across markets.

Taken together, the findings suggest that transparency mandates may redistribute rather than reduce drug costs. By making rebate flows more visible but leaving the underlying contracting incentives unchanged, these policies improved oversight but not affordability. A growing number of federal and state initiatives aim to curb drug spending by mandating rebate disclosure, enforcing pass-through requirements, or prohibiting spread pricing altogether (Myshko, 2025; Mintz, Levin, Cohn, Ferris, Glovsky and Popeo, P.C., 2025). The results here indicate that such transparency measures, on their own, may have limited effectiveness in vertically integrated pharmaceutical markets. When intermediaries retain market power and insurers have weak incentives to pass through savings, greater disclosure of rebate information does not translate into lower out-of-pocket costs for patients. Effective cost containment may therefore require complementary instruments—such as minimum pass-through standards, limits on spread pricing, or uniform federal disclosure—that directly alter how intermediaries allocate rebates and negotiate prices. More broadly, the results caution that transparency in complex, vertically structured markets can have unintended general-equilibrium effects, particularly when intermediaries operate across regulatory boundaries.

Overall, this evidence contributes to a broader understanding of when and how disclosure regulation affects real outcomes. In markets where pricing reflects strategic contracting rather than consumer search, transparency can enhance accountability but is unlikely to substitute for incentive alignment. The experience of state rebate-transparency laws illustrates both the promise and the limits of transparency as a policy tool for controlling drug spending in the United States.

These null findings are themselves informative. As emphasized by Abadie (2020), in empirical settings with large samples and diffuse priors—typical of applied economics—failure to reject a null hypothesis conveys substantial information about the absence of economically meaningful effects. In this context, the lack of measurable price responses to rebate-transparency laws should not be interpreted as inconclusive, but rather as evidence that transparency alone is unlikely to alter equilibrium pricing or pass-through behavior in pharmaceutical markets. The precision of the estimates rules out large cost reductions and suggests that any potential effects are economically minor. In Abadie's terms, the nonsignificant estimates meaningfully update beliefs about the likely impact of disclosure-based regulation, providing information about what policies do not achieve as well as what they do.

7 Conclusion

This paper examines whether rebate-transparency laws reduced prescription drug costs in Medicare Part D. Using a difference-in-differences design that exploits the staggered adoption of state transparency mandates, I find no evidence that disclosure lowered premiums, deductibles, or patient cost sharing. Instead, costs rose modestly in untreated states linked to multi-state contracts, suggesting that intermediaries reallocated prices across their portfolios in response to uneven regulation.

These results highlight an important boundary condition for disclosure-based policy. Transparency can improve oversight but may not change outcomes when intermediaries retain market power or can adjust along unregulated margins. In such settings, information alone is insufficient to discipline pricing behavior, and partial regulation can induce offsetting effects across markets.

More broadly, the findings illustrate that the real effects of disclosure depend on how transparency interacts with the structure of contracting and incentives in the underlying market. In the case of prescription drugs, effective cost containment may require policies that directly target the allocation of rebates and the bargaining power of intermediaries, rather than transparency alone. Understanding when disclosure alters—not merely reveals—economic behavior remains a central question for research on regulation and information design.

References

Abadie, A. (2020): "Statistical nonsignificance in empirical economics," *American Economic Review: Insights*, 2, 193–208.

Agha, L., S. Kim, and D. Li (2022): "Insurance design and pharmaceutical innovation," *American Economic Review: Insights*, 4, 191–208.

Arya, A. and B. Mittendorf (2011): "Disclosure standards for vertical contracts," *The RAND Journal of Economics*, 42, 595–617.

Ball, R. (2024): "By what criteria do we evaluate accounting? Some thoughts on economic welfare and the archival literature," *Journal of Accounting Research*, 62, 7–54.

Benecurv (2024): "The Transparency of Rebates among Pharmacy Benefit Managers (PBMs)," .

Beyer, A., D. A. Cohen, T. Z. Lys, and B. R. Walther (2010): "The financial reporting environment: Review of the recent literature," *Journal of accounting and economics*, 50, 296–343.

Bourveau, T., V. Capkun, and Y. Wang (2020a): "Consequences of Disclosing Clinical Trial Results: Evidence from the Food and Drug Administration Amendments Act," Working paper. Available at SSRN.

Bourveau, T., G. She, and A. Žaldokas (2020b): "Corporate disclosure as a tacit coordination mechanism: Evidence from cartel enforcement regulations," *Journal of Accounting Research*, 58, 295–332.

Brot-Goldberg, Z. C., C. Che, and B. R. Handel (2022): "Pharmacy benefit managers and vertical relationships in drug supply: state of current research," Tech. rep., National Bureau of Economic Research.

Capkun, V., Y. Lou, C. A. Otto, and Y. Wang (2023): "Do firms respond to peer disclosures? Evidence from disclosures of clinical trial results," *The Accounting Review*, 98, 71–108.

Che, C. (2025): "Do Middlemen Raise Drug Prices? Countervailing Market Power Meets Agency Frictions," *Publisher: University of California, Berkeley.*

Christensen, H. B., E. Floyd, L. Y. Liu, and M. Maffett (2017): "The real effects of mandated information on social responsibility in financial reports: Evidence from mine-safety records," *Journal of Accounting and Economics*, 64, 284–304.

Christensen, H. B., E. Floyd, and M. Maffett (2020): "The only prescription is transparency: The effect of charge-price-transparency regulation on healthcare prices," *Management Science*, 66, 2861–2882.

Conti, R. M., B. Frandsen, M. L. Powell, and J. B. Rebitzer (2021): "Common agent or double agent? Pharmacy benefit managers in the prescription drug market," Tech. rep., National Bureau of Economic Research.

Darrough, M. N. (1993): "Disclosure policy and competition: Cournot vs. Bertrand," *Accounting review*, 534–561.

de Chaisemartin, C. and X. D'Haultfœuille (2023): "Credible answers to hard questions: differences-in-differences for natural experiments," *Available at SSRN 4487202*.

De Chaisemartin, C. and X. d'Haultfoeuille (2020): "Two-way fixed effects estimators with heterogeneous treatment effects," *American economic review*, 110, 2964–2996.

Ellis, J. A., C. E. Fee, and S. E. Thomas (2012): "Proprietary costs and the disclosure of information about customers," *Journal of Accounting Research*, 50, 685–727.

Eyring, H. (2020): "Disclosing physician ratings: Performance effects and the difficulty of altering ratings consensus," *Journal of Accounting Research*, 58, 1023–1067.

Feng, J. and L. Maini (2024): "Demand inertia and the hidden impact of pharmacy benefit managers," *Management Science*, 70, 8940–8961.

Feng, R., M. D. Kimbrough, and S. Wei (2022): "The role of information transparency in the product market: an examination of the sustainability of profitability differences," *Review of Accounting Studies*, 27, 668–705.

Fiore, N. M. (2025): "The Transparency Hold-up: Pharmacy Benefit Managers and Rebate Disclosure," *Available at SSRN 5682442*.

Goncharov, I. and C. D. Peter (2019): "Does reporting transparency affect industry coordination? Evidence from the duration of international cartels," *The Accounting Review*, 94, 149–175.

Gray, C., A. Alpert, and N. Sood (2023): "Disadvantaging Rivals: Vertical Integration in the Pharmaceutical Market," Tech. Rep. w31536, National Bureau of Economic Research, Cambridge, MA.

Ho, K. and R. S. Lee (2024): "Contracting over pharmaceutical formularies and rebates," Tech. rep., National Bureau of Economic Research.

Honda, J., R. Inderst, and M. Ottaviani (2024): "When liability is not enough: Regulating bonus payments in markets with advice," *Management Science*, 70, 1301–1314.

Hwang, Y. and A. J. Kirby (2000): "Competitive effects of disclosure in a strategic entry model," *Review of Accounting Studies*, 5, 57–85.

Inderst, R. and M. Ottaviani (2012a): "Financial advice," *Journal of Economic Literature*, 50, 494–512.

——— (2012b): "How (not) to pay for advice: A framework for consumer financial protection," *Journal of Financial Economics*, 105, 393–411.

——— (2012c): "Regulating financial advice," *European Business Organization Law Review (EBOR)*, 13, 237–246.

Kakani, P., M. Chernew, and A. Chandra (2020): "Rebates in the pharmaceutical industry: Evidence from medicines sold in retail pharmacies in the us," Tech. rep., National Bureau of Economic Research.

Kakani, P., M. A. Kyle, A. Chandra, and L. Maini (2024): "Medicare Part D Protected-Class Policy Is Associated With Lower Drug Rebates: Article examines Medicare Part D rebates," *Health Affairs*, 43, 1420–1427.

Kanodia, C. and H. Sapra (2016): "A real effects perspective to accounting measurement and disclosure: Implications and insights for future research," *Journal of Accounting Research*, 54, 623–676.

Kepler, J. D., V. V. Nikolaev, N. Scott-Hearn, and C. R. Stewart (2024): "Quality transparency and healthcare competition," *Available at SSRN 3963418*.

Leuz, C., A. Malani, M. Muhn, and L. Jakab (2023): "Do conflict of interests disclosures work? Evidence from citations in medical journals," Tech. rep., National Bureau of Economic Research.

Leuz, C. and P. D. Wysocki (2016): "The economics of disclosure and financial reporting regulation: Evidence and suggestions for future research," *Journal of accounting research*, 54, 525–622.

Medicare, C. f. and M. Services (2023a): "Final Contract Year (CY) 2024 Part D Bidding Instructions," Tech. rep., U.S. Department of Health and Human Services, Centers for Medicare and Medicaid Services, Baltimore, MD.

——— (2023b): "Final Medicare Part D DIR Reporting Guidance for 2022," Tech. rep., U.S. Department of Health and Human Services, Centers for Medicare and Medicaid Services, Baltimore, MD.

——— (2023c): "Instructions for Completing the Part D Bid Pricing Tool (BPT) - CY 2024," Tech. rep., U.S. Department of Health and Human Services, Centers for Medicare and Medicaid Services, Baltimore, MD.

Mintz, Levin, Cohn, Ferris, Glovsky and Popeo, P.C. (2025): "PBM Policy and Legislative Update - Spring 2025 Edition," .

Myshko, D. (2025): "States Become More Aggressive with PBM Reform," .

Nathan, B. (2024): "Show Your Hand: The Impacts of Fair Pricing Requirements in Procurement Contracting," *Journal of Accounting Research*, 62, 1405–1448.

Scanlon, G. (2025): "Prescription for Savings? Disclosure in the Drug Market," Ph.D. thesis, The University of North Carolina at Chapel Hill.

Shepherd, J. (2014): "Written Testimony for the ERISA Advisory Council Hearing on PBM Compensation and Fee Disclosure," .

Sood, N., T. Shih, K. Van Nuys, and D. P. Goldman (2017): "Follow The Money: The Flow Of Funds In The Pharmaceutical Distribution System," .

STARC, A. AND A. SWANSON (2021): "Preferred Pharmacy Networks and Drug Costs," *American Economic Journal: Economic Policy*, 13, 406–446.

WEIL, D., M. GRAHAM, AND A. FUNG (2013): "Targeting transparency," *Science*, 340, 1410–1411.

ZHANG, Y. (2024): "Corporate R&D investments following competitors' voluntary disclosures: Evidence from the drug development process," *Journal of Accounting Research*, 62, 335–373.

Figures

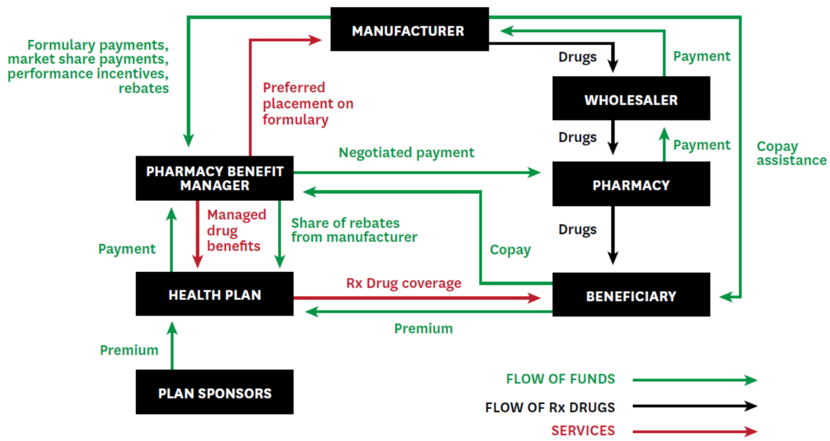

(a) Pharmaceutical supply chain

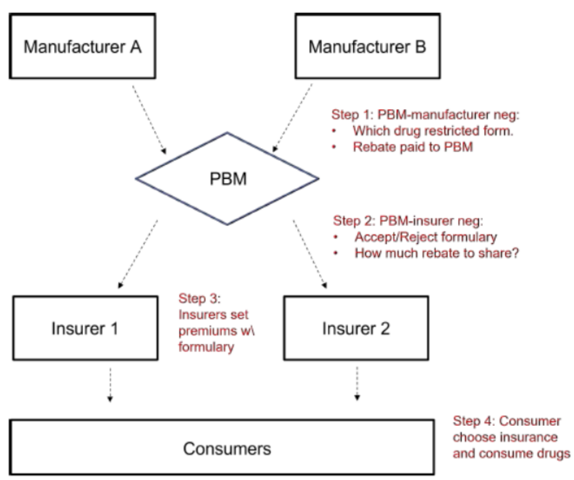

(b) Contracting over rebates and formularies ...

Figure 1: Overview of the pharmaceutical distribution and contracting system

Notes: Panel (a) illustrates the relationships among manufacturers, wholesalers, pharmacies, pharmacy benefit managers (PBMs), health plans, and consumers, showing the flow of products, payments, and rebates across the pharmaceutical supply chain. Panel (b) depicts the contracting process between drugmakers, PBMs, and insurers, emphasizing how rebate negotiations and formulary placement affect pricing and incentives throughout the system. *Sources:* Panel (a) adapted from Sood et al. (2017); Panel (b) adapted from Brot-Goldberg et al. (2022).

(e) The pharmacy benefit manager shall, on a quarterly basis, disclose, upon the request of the purchaser, the following information with respect to prescription product benefits specific to the purchaser:

(1) The aggregate wholesale acquisition costs from a pharmaceutical manufacturer or labeler for each therapeutic category of drugs containing three or more drugs, as outlined in the state's essential health benefits benchmark plan pursuant to Section 1367.005 of the Health and Safety Code.

(2) The aggregate amount of rebates received by the pharmacy benefit manager by therapeutic category of drugs containing three or more drugs, as outlined in the state's essential health benefits benchmark plan pursuant to Section 1367.005 of the Health and Safety Code. The aggregate amount of rebates shall include any utilization discounts the pharmacy benefit manager receives from a pharmaceutical manufacturer or labeler.

(a) California (AB 315, 2017)

Sec. 6. **[62W.06] PHARMACY BENEFIT MANAGER TRANSPARENCY.**

Subdivision 1. **Transparency to plan sponsors.** (a) Beginning in the second quarter after the effective date of a contract between a pharmacy benefit manager and a plan sponsor, the pharmacy benefit manager must disclose, upon the request of the plan sponsor, the following information with respect to prescription drug benefits specific to the plan sponsor:

(1) the aggregate wholesale acquisition costs from a drug manufacturer or wholesale drug distributor for each therapeutic category of prescription drugs;

(2) the aggregate wholesale acquisition costs from a drug manufacturer or wholesale drug distributor for each therapeutic category of prescription drugs available to the plan sponsor's enrollees;

(3) the aggregate amount of rebates received by the pharmacy benefit manager by therapeutic category of prescription drugs. The aggregate amount of rebates must include any utilization discounts the pharmacy benefit manager receives from a drug manufacturer or wholesale drug distributor;

(b) Minnesota (SF 278, 2019)

Sec. 25. (a) A party that has contracted with a pharmacy benefit manager to provide services may, at least one (1) time in a calendar year, request an audit of compliance with the contract. The audit may include full disclosure of rebate amounts secured on prescription drugs, whether product specific or general rebates, that were provided by a pharmaceutical manufacturer and any other revenue and fees derived by the pharmacy benefit manager from the contract. A contract may not contain provisions that impose unreasonable fees or conditions that would severely restrict a party's right to conduct an audit under this subsection.

(b) A pharmacy benefit manager shall disclose, upon request from a party that has contracted with a pharmacy benefit manager, to the party the actual amounts paid by the pharmacy benefit manager to any pharmacy.

(c) A pharmacy benefit manager shall provide notice to a party contracting with the pharmacy benefit manager any consideration that the pharmacy benefit manager receives from a pharmacy manufacturer for any name brand dispensing of a prescription when a generic or biologically similar product is available for the prescription.

(c) Indiana (SB 241, 2020)

Section 6. Pharmacy benefit manager transparency to carriers and plan sponsors. (1) Beginning in the second quarter after the effective date of a contract between a pharmacy benefit manager and a health carrier, plan sponsor, or workers' compensation insurance carrier, the pharmacy benefit manager shall disclose, within 45 days of a request of the health carrier, plan sponsor, or workers' compensation insurance carrier, the following information regarding prescription drug benefits specific to the health carrier, ~~or~~ plan sponsor, or workers' compensation insurance carrier:

(a) the aggregate wholesale acquisition costs from a manufacturer or wholesale distributor for each therapeutic category of prescription drugs;

(b) the aggregate wholesale acquisition costs from a manufacturer or wholesale distributor for each therapeutic category of prescription drugs available to enrollees of the health carrier or plan sponsor or injured workers of the workers' compensation insurance carrier;

(c) the aggregate amount of rebates received by the pharmacy benefit manager by therapeutic category of prescription drugs;

(d) Montana (SB 395, 2021)

Figure 2: State pharmacy benefit manager (PBM) rebate-disclosure laws

Notes: Panels (a)–(d) reproduce key statutory excerpts from state rebate-disclosure laws. Each panel shows the core language requiring pharmacy benefit managers to disclose manufacturer rebates and fees to plan sponsors or regulators. See Table A1 for details and sources.

Figure 3: Dynamic Effects of Rebate-Transparency Laws on Total Annual Beneficiary Cost

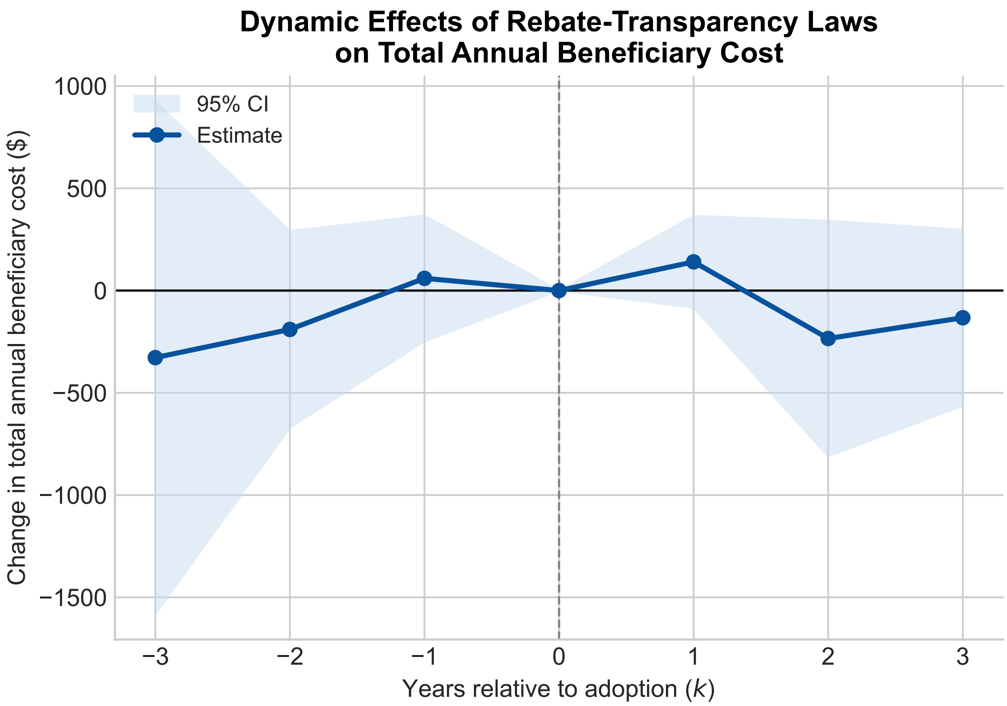

Notes: This figure plots event-time coefficients from equation (1) using plan–county–year data for non–PBM-affiliated plans. Each point represents the estimated change in total annual beneficiary cost (*costyearly*) in year k relative to the year before adoption ($k = -1$). Vertical line at $k = 0$ denotes the year of adoption of a rebate-transparency law. Shaded bands show 95% confidence intervals. Regressions include state and year fixed effects, are weighted by plan–county–year enrollment, and cluster standard errors at the state level. All monetary values are in nominal U.S. dollars.

Figure 4: Spillover Effects of Rebate-Transparency Laws on Total Annual Beneficiary Cost

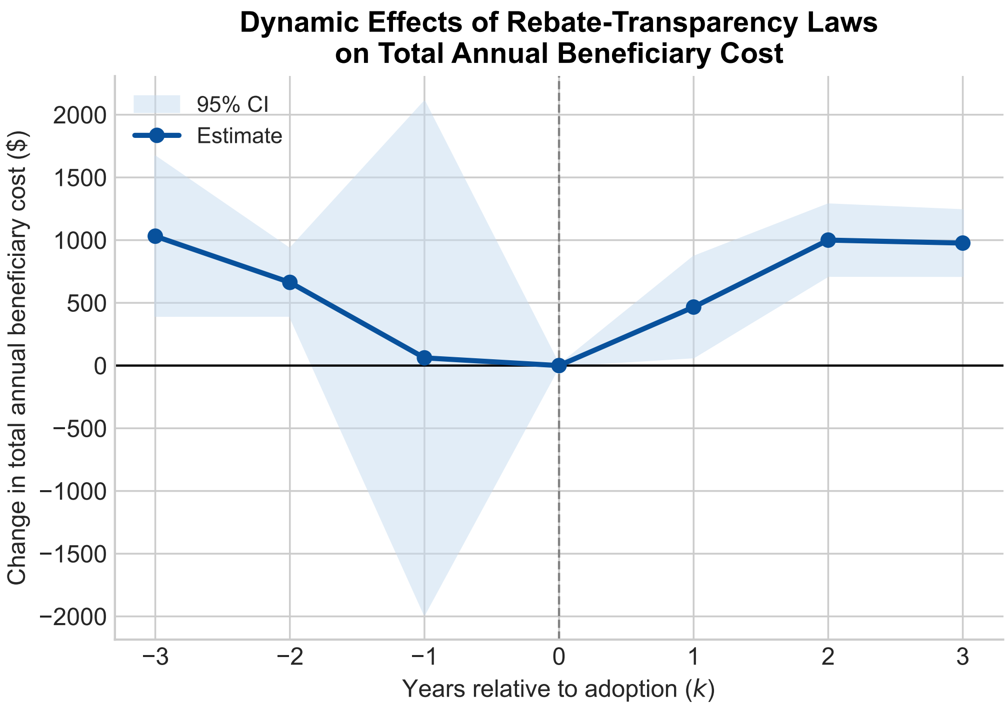

Notes: This figure plots event-time coefficients from equation (1) using plan–county–year data for *non–PBM-affiliated* plans in *untreated states* whose contracts also operate in treated states. Each point represents the estimated change in total annual beneficiary cost (*costyearly*) in year k relative to the year before adoption ($k = -1$). Vertical line at $k = 0$ denotes the year in which connected treated states adopted a rebate-transparency law. Shaded bands show 95% confidence intervals. Regressions include state and year fixed effects, are weighted by plan–county–year enrollment, and cluster standard errors at the plan (*contractid*) level. All monetary values are expressed in nominal U.S. dollars.

Tables

Table 1: Summary Statistics (Plan-County-Year, 2016–2025)

	Mean	SD	P25	P75	N
Panel A. Plan costs and prices					
Premium	378.01	373.31	0.00	556.80	1,275,731
Deductible	284.67	212.91	0.00	445.00	1,275,731
Cost sharing	534.62	1,268.46	116.77	450.17	1,257,784
Total cost	1,197.44	1,311.05	635.55	1,301.80	1,257,784
Negotiated price	1,731.13	2,741.85	124.19	2,076.42	1,257,784
NADAC acquisition cost	30,977.59	35,069.68	10,431.62	45,641.90	1,092,388
Coinsurance rate	0.68	0.40	0.20	1.00	1,257,784
Boxed warning (=1)	1.00	0.06	1.00	1.00	921,607
Branded drug (=1)	1.00	0.04	1.00	1.00	1,257,784
Specialty drug (=1)	0.00	0.06	0.00	0.00	1,257,784
Panel B. Plan and market characteristics					
PBM-affiliated plan (=1)	0.29	0.45	0.00	1.00	1,275,731
Nonprofit plan (=1)	0.05	0.21	0.00	0.00	1,275,731
Single-state plan (=1)	0.24	0.43	0.00	0.00	1,275,731
MLR $\leq$ 85% (=1)	0.07	0.25	0.00	0.00	770,791
Enrollment (count)	266.76	1,330.04	0.00	158.00	1,275,731
County HHI below median (=1)	0.56	0.50	0.00	1.00	1,275,236
Observations	1,275,731 plan-county-year observations				
Years	2016–2025				

Notes: The unit of observation is plan-county-year. All monetary variables (premium, deductible, negotiated price, NADAC acquisition cost, cost sharing, and total cost) are expressed in U.S. dollars. "Boxed warning," "Branded drug," and "Specialty drug" indicate whether the minimum-drug observation in a plan-county-year carries an FDA boxed warning, is branded, or is specialty tier, respectively. "County HHI below median" equals one if the Herfindahl–Hirschman Index (HHI), computed from plan enrollment shares within a county–year, is below the sample median. All statistics are unweighted.

Table 2: Pre-Policy Balance of Treated and Untreated States

	Treated states (pre-policy)	Untreated states	Treated − Untreated	p-value
Panel A. Plan costs and prices				
Premium	511.7	371.5	140.2	0.00
Deductible	241.5	284.2	−42.7	0.00
Cost sharing	1,195.3	530.7	664.6	0.00
Total cost	1,949.1	1,186.5	762.6	0.00
Panel B. Plan and market characteristics				
PBM-affiliated plan (=1)	0.32	0.29	0.03	0.00
Nonprofit plan (=1)	0.08	0.04	0.03	0.00
Single-state plan (=1)	0.14	0.24	−0.11	0.00
MLR $\leq$ 85% (=1)	0.04	0.07	−0.03	0.00
Enrollment (count)	394.0	243.3	150.8	0.00
County HHI below median (=1)	0.41	0.58	−0.16	0.00
Panel C. Policy coverage				
Number of adopting states		4		
Share of plan-county-year observations in adopting states		9.9%		
Observations		1,197,497 plan-county-year observations		
Unique plans		10,739 contracts		
Counties / States / Years		3,220 counties, 56 states, 10 years (2016–2025)		

Notes: The table reports mean values for states that ever adopt a rebate-transparency law (treated) and those that never adopt (untreated). The treated sample includes only years before each state's adoption year; all years are included for untreated states. Column (3) shows differences in means (treated − untreated), and column (4) reports p-values from regressions of each variable on a treatment indicator, with standard errors clustered at the state level. Panel C summarizes policy adoption and coverage. All monetary variables are expressed in U.S. dollars. All variables are defined in Table 1.

Table 3: Dynamic Effects of Rebate-Transparency Laws on Plan Costs

	Dependent variable: plan-level cost component			
	(1)	(2)	(3)	(4)
	Premium	Deductible	Cost sharing	Total cost
Event-time coefficients				
$k = +1$	10.21	-3.50	133.58	140.27
	[0.914]	[-0.796]	[1.220]	[1.226]
$k = +2$	28.91	-0.56	-261.92	-234.36
	[1.301]	[-0.098]	[-0.846]	[-0.798]
$k = +3$	34.89	-0.76	-167.35	-132.92
	[1.163]	[-0.083]	[-0.694]	[-0.606]
Average post-adoption effect	25.24	-1.54	-106.10	-82.58
	[1.181]	[-0.281]	[-0.576]	[-0.482]
Joint p-value (post-policy)	0.366	0.862	0.555	0.593
State fixed effects	Y	Y	Y	Y
Year fixed effects	Y	Y	Y	Y
Observations (plan-county-year)	892,158	892,158	880,627	880,627
Mean of dependent variable	356.25	286.08	584.88	1,226.57

Notes: Each column reports event-time estimates of equation (1) using plan-county-year data for non-PBM-affiliated plans. Coefficients correspond to years relative to each state's adoption of a rebate-transparency law ($k = 0$ is the adoption year). The highlighted row shows the average post-adoption (cumulative) effect per treatment unit. The final line reports the p-value from a joint test that all post-adoption coefficients are equal to zero. t-statistics are shown in brackets. Estimates are weighted by plan-county-year enrollment, and standard errors are clustered at the state level. Monetary values are in U.S. dollars.

Table 4: Precision and Minimum Detectable Effects (MDE): Plan Costs

	Dependent variable: plan-level cost component			
	(1)	(2)	(3)	(4)
	Premium	Deductible	Cost sharing	Total cost
Mean of dependent variable	236.98	242.29	657.63	1,136.28
Average post-adoption effect	25.24	−1.54	−106.10	−82.58
	(21.36)	(5.48)	(184.12)	(171.43)
95% confidence interval	[−16.63, 67.11]	[−12.28, 9.20]	[−466.98, 254.78]	[−418.57, 253.42]
MDE (80% power, 5% size)	59.85	15.35	515.84	480.27
MDE as % of mean	25.25%	6.34%	78.44%	42.27%
Fixed effects	Plan-county and year			
Weights	Enrollment			
Observations (plan-county-year)	892,158	892,158	880,627	880,627

Notes: This table summarizes the precision of the average post-adoption estimates from Table 3. Standard errors are clustered at the state level. The minimum detectable effect (MDE) is the smallest absolute change in the outcome that would be detected with 80% power at 5% size under the study design and inference procedure. Monetary values are in U.S. dollars.

Table 5: Effects of Rebate-Transparency Laws on Pricing and Formulary Mechanisms (Non-PBM Plans)

	Dependent variable: pricing or formulary outcome				
	(1)	(2)	(3)	(4)	(5)
	Negotiated price	NADAC	Coinsurance	Branded drug	Specialty drug
Panel A. Levels (minimum-drug outcomes)					
$k = +1$	−15.33	14,844.94	0.03	0.00	−0.00
	[−0.07]	[1.25]	[0.95]	[.]	[−0.18]
$k = +2$	657.24	18,468.99	−0.10	0.00	0.16
	[1.33]	[1.19]	[−1.04]	[.]	[1.18]
$k = +3$	755.31	25,006.82	−0.11	0.00	0.00
	[1.34]	[1.20]	[−1.02]	[.]	[0.47]
Average post-adoption effect	483.58	20,660.29	−0.06	0.00	0.06
	[1.39]	[1.20]	[−1.00]	[.]	[1.21]
Joint p-value (post-policy)	0.523	0.645	0.710	1.000	0.359
State fixed effects	Y	Y	Y	Y	Y
Year fixed effects	Y	Y	Y	Y	Y
Observations (plan-county-year)	880,627	785,789	880,627	880,627	880,627
Mean of dependent variable	1,782.47	30,883.72	0.71	1.00	0.00

	Dependent variable: Jevons price index			
	(1)	(2)	(3)	(4)
	Negotiated price	NADAC	Cost sharing	Coinsurance
Panel B. Jevons indices				
$k = +1$	−0.48	−0.07	1.84	1.62
	[−1.24]	[−0.72]	[1.01]	[0.76]
$k = +2$	−0.01	−0.38	−7.52	−22.76
	[−0.02]	[−1.22]	[−1.09]	[−1.18]
$k = +3$	0.01	−0.45	−7.05	−21.55
	[0.02]	[−1.34]	[−0.96]	[−1.04]
Average post-adoption effect	−0.16	−0.30	−4.17	−14.02
	[−0.52]	[−1.23]	[−0.99]	[−1.11]
Joint p-value (post-policy)	0.619	0.465	0.631	0.581
State fixed effects	Y	Y	Y	Y
Year fixed effects	Y	Y	Y	Y
Observations (plan-county-year)	451,771	451,537	448,597	448,597
Mean of dependent variable	68.26	104.25	85.37	138.92

Notes: Each column reports event-time estimates of equation (1) for distinct pricing or formulary outcomes among non-PBM-affiliated plans. Panel A reports level outcomes based on the *minimum-drug* definition, which selects in each plan-county-year the National Drug Code (NDC) with the lowest cost sharing and carries its associated values for negotiated price, acquisition cost (NADAC), coinsurance rate, branded status, and specialty status. Panel B presents results for the corresponding Jevons price indices of negotiated prices, acquisition costs (NADAC), cost sharing, and coinsurance. The final line reports the p-value from a joint test that all post-adoption coefficients are equal to zero. t-statistics are shown in brackets. Estimates are weighted by plan-county-year enrollment, and standard errors are clustered at the state level. Monetary values are in U.S. dollars.

Table 6: Heterogeneity in the Effect of Rebate-Transparency Laws on Plan-Level Total Annual Costs (Non-PBM Plans)

	Dependent variable: plan-level total annual cost		
	(1) Low/Below	(2) High/Above	Diff. p-value
Panel A. By plan's MLR constraint (binding if MLR $\leq$ 85% pre-policy)			
Average post-adoption effect	-45.43 $[-0.279]$	-104.41 $[-0.336]$	p = 0.87
Observations	820,371	21,542	
Mean of dependent variable	1,235.67	1,077.91	
Panel B. By plan's bargaining power (above-median pre-policy enrollment)			
Average post-adoption effect	-894.86 $[-1.391]$	-83.31 $[-0.461]$	p = 0.23
Observations	91,805	788,822	
Mean of dependent variable	1,235.27	1,225.56	
Panel C. By local competition (county HHI below median pre-policy)			
Average post-adoption effect	-115.37 $[-0.774]$	-112.10 $[-0.535]$	p = 0.99
Observations	371,067	509,177	
Mean of dependent variable	1,293.23	1,177.98	

Notes: Each panel reports subgroup-specific estimates of the *average post-adoption* effect of rebate-transparency laws on plan-level total annual costs among non-PBM-affiliated plans. Columns (1) and (2) show the average cumulative (total) effect per treatment unit. t-statistics are shown in brackets, and p-values refer to two-sided tests of the null that the difference between the subgroup coefficients equals zero. Estimates are weighted by plan-county-year enrollment, and standard errors are clustered at the state level. "Low/Below" and "High/Above" correspond respectively to: (A) non-binding versus binding MLR (pre-policy), (B) below- versus above-median pre-policy plan enrollment, and (C) above- versus below-median pre-policy county-level HHI (so column (2) represents more competitive markets). All monetary values are in U.S. dollars.

Table 7: Dynamic Effects of Rebate-Transparency Laws on Plan Costs: Single-State vs. Multi-State Plans

	Dependent variable: plan-level cost component			
	(1)	(2)	(3)	(4)
	Premium	Deductible	Cost sharing	Total cost
Panel A. Single-state plans				
Average post-adoption effect	61.31	11.40	−56.46	15.45
	[1.253]	[0.588]	[−0.206]	[0.065]
Panel B. Multi-state plans				
Average post-adoption effect	4.53	6.73	56.89	68.42
	[0.308]	[0.990]	[0.973]	[1.003]
Panel C. Difference (Single-state − Multi-state)				
Difference estimate	56.78	4.67	−113.35	−52.97
t-statistic (z)	[1.11]	[0.23]	[−0.40]	[−0.21]
p-value	0.266	0.820	0.686	0.830
State fixed effects	Y	Y	Y	Y
Year fixed effects	Y	Y	Y	Y
Observations (plan-county-year)	283,034	283,034	281,353	281,353
Mean of dependent variable	184.97	244.78	741.67	1,171.49

Notes: Each column reports the average post-adoption (cumulative) effect from equation (1) estimated separately for single-state (Panel A) and multi-state (Panel B) plans. Panel C reports the difference between these two estimates (Single-state − Multi-state). t-statistics are shown in brackets, and p-values correspond to two-sided tests of the null that the coefficient equals zero. Estimates are weighted by plan-county-year enrollment, and standard errors are clustered at the state level. Monetary values are in U.S. dollars.

Table 8: Spillover Effects of Rebate-Transparency Laws on Plan Costs (Contracts in Untreated States Serving Treated Markets)

	Dependent variable: plan-level cost component			
	(1)	(2)	(3)	(4)
	Premium	Deductible	Cost sharing	Total cost
Event-time coefficients				
$k = +1$	27.32	17.50	421.62	466.48
	[1.292]	[2.284]	[2.237]	[2.287]
$k = +2$	11.65	53.25	936.68	1,000.32
	[0.310]	[4.447]	[5.703]	[6.904]
$k = +3$	17.56	54.79	908.16	976.64
	[0.433]	[4.193]	[6.403]	[7.332]
Average post-adoption (spillover) effect	18.47	42.97	770.02	829.64
	[0.558]	[4.368]	[7.943]	[8.087]
Joint p-value (post-policy)	0.024	0.0002	0.0000	0.0000
State fixed effects	Y	Y	Y	Y
Year fixed effects	Y	Y	Y	Y
Observations (plan-county-year)	805,827	805,827	795,362	795,362
Mean of dependent variable	348.91	285.08	579.08	1,212.39

Notes: Each column reports event-time estimates of equation (1) using plan-county-year data for non-PBM-affiliated plans in untreated states whose contracts also operate in treated states. The treatment indicator equals one when a contract is exposed to a rebate-transparency law through multi-state operations (*treatedgspill*=1, *dummygtspill*=1). Coefficients correspond to years relative to each treated state's adoption year ($k = 0$ is the adoption year). The highlighted row shows the average post-adoption (spillover) effect per treatment unit. The final line reports the p-value from a joint test that all post-policy coefficients ($k > 0$) are equal to zero. t-statistics are shown in brackets. Estimates are weighted by plan-county-year enrollment, and standard errors are clustered at the contract level (*contractid*). Monetary values are in U.S. dollars.

A Additional Tables and Figures

Table A1: State legislation requiring pharmacy benefit manager (PBM) rebate disclosure, 2018–2022

Enacted Year	State	Bill (Session Year)	Link
2019	California	Assembly Bill 315 (2017–18)	[link]
2020	Minnesota	Senate File 278 (2019)	[link]
2021	Indiana	Senate Bill 241 (2019)	[link]
2022	Montana	Senate Bill 395 (2021)	[link]

Notes: The table lists state statutes requiring pharmacy benefit managers to disclose rebate or fee information to plan sponsors or regulators. All links direct to official legislative sources.

Table A2: Rebate-Transparency Laws: Non-PBM vs. PBM Plans (Average Post-Adoption Effects)

	Dependent variable: plan-level cost component			
	(1)	(2)	(3)	(4)
	Premium	Deductible	Cost sharing	Total cost
Panel A. Non-PBM plans				
Average post-adoption effect	25.24	−1.54	−106.10	−82.58
	[1.181]	[−0.281]	[−0.576]	[−0.482]
Panel B. PBM-affiliated plans				
Average post-adoption effect	15.21	26.66	47.67	89.44
	[0.841]	[1.360]	[0.982]	[1.247]
Panel C. Difference (Non-PBM − PBM)				
Difference estimate	10.04	−28.20	−153.77	−172.02
t-statistic (z)	[0.36]	[−1.39]	[−0.81]	[−0.93]
p-value	0.720	0.166	0.419	0.355
State fixed effects	Y	Y	Y	Y
Year fixed effects	Y	Y	Y	Y
Observations (plan-county-year)	892,158	892,158	880,627	880,627
Mean of dependent variable	356.25	286.08	584.88	1,226.58

Notes: Each column reports the average post-adoption (cumulative) effect from equation (1) estimated separately for non-PBM plans (Panel A) and PBM-affiliated plans (Panel B). Panel C reports the difference between these two estimates (triple-difference). t-statistics are shown in brackets, and p-values correspond to two-sided tests of the null that the coefficient equals zero. Estimates are weighted by plan-county-year enrollment, and standard errors are clustered at the state level. Monetary values are in U.S. dollars.

B Data Construction and Sample Definition

This appendix documents how the empirical dataset is constructed from administrative and public sources.

The primary data are the CMS *Quarterly Prescription Drug Plan Formulary, Pharmacy Network, and Pricing Information* files (Q2 releases, 2016-2025). These files report plan identifiers and benefit design, including monthly premiums (*premiumyearly*), annual deductibles (*deductible*), and drug-level negotiated prices and cost-sharing terms for all covered drugs and plan-pharmacy combinations in each county. The universe includes stand-alone Prescription Drug Plans (PDPs) and Medicare Advantage Prescription Drug Plans (MA-PDs).

Each observation is identified by contract, plan, and segment codes; I normalize and concatenate these to form a persistent key, *contractplansegmentid*, enabling consistent tracking across years and geographies. Geographic identifiers are standardized to FIPS state and county codes (*statefips*, *FIPS State County Code*) using CMS crosswalks.

Plan attributes are taken from the same CMS quarterly SPUF "Plan Information" files. In addition to *premiumyearly* and *deductible*, I construct indicators for PBM affiliation (*planpbm*), nonprofit status (*nonprofit*), and single-state operation (*singlestate*). PBM affiliation and nonprofit status are inferred by manual classification of contract and labeler names; *singlestate* equals one if, in a given year, the plan's observed service area is confined to a single state. Contract-plan-segment identifiers are kept in *CONTRACT_ID*, *PLAN_ID*, and *SEGMENT_ID*, and aggregated using *contractplansegmentid*.

Enrollment counts (*Enrollment*) are obtained from the CMS Medicare Advantage/Part D Contract and Enrollment Data and merged at the contract-plan-county-year level using SSA and FIPS county codes. Medical Loss Ratio data come from the CMS Part D MLR public use files; I define *mlrbelow85* as an indicator equal to one when the reported MLR is at or below 0.85.

Drug-level attributes are constructed by linking the CMS pricing records to RxNorm, FDA Orange Book, and NADAC. RxNorm provides standardized molecule identifiers (*RXCUI*) and term types; I also harmonize package NDCs to an 11-digit format (*ndc11*). FDA Orange Book "Products," "Patents," and "Exclusivity" files identify products with active patent or exclusivity listings. NADAC supplies benchmark acquisition costs and a classification for rate setting. I classify a product as branded (*branded* $= 1$) if any of the following hold in the year: RxNorm branded term types (e.g., SBD/BN/BPCK), a patent listing, an exclusivity listing, or NADAC's brand/generic classification indicates brand; otherwise *branded* $= 0$. NADAC prices are annualized as

$$\textit{nadacyearly} = (\textit{NADAC Per Unit}) \times 365.$$

The empirical analysis focuses on the initial coverage phase of Medicare Part D—after beneficiaries satisfy their deductible but before reaching the Initial Coverage Limit—where the vast majority of prescriptions are filled and where cross-plan variation in negotiated prices and cost sharing is most pronounced. For each plan, drug, and days-supply category (30, 60, or 90 days) in this phase, negotiated prices are annualized:

$$\textit{negotiatedpriceyearly} = \frac{\textit{unit\_cost}}{\textit{days\_supply}} \times 365.$$

Patient payments (*costsharing*) reflect preferred-pharmacy terms: if a copayment is specified, the copay is divided by days-supply and multiplied by 365; if coinsurance is specified, the rate is applied to the negotiated price. The final *costsharing* is the minimum of the negotiated price, the copay-equivalent, and the coinsurance payment. The variable *coinsurance* records the implied coinsurance rate. When a drug appears with multiple days-supply options, I retain the record with the lowest *costsharing*. Specialty-tier status (*specialty*) is carried from the CMS formulary files.

From these components, I construct total annual beneficiary cost:

$$\textit{costyearly} = \textit{premiumyearly} + \textit{deductible} + \textit{costsharing}.$$

This serves as the summary measure of plan generosity in the main analysis.

The analysis sample is restricted to drugs in WHO ATC class A10BJ, which includes glucagon-like peptide-1 (GLP-1) receptor agonists. This therapeutic class is particularly well suited for studying rebate dynamics: it is dominated by branded products but still exhibits active competition among manufacturers, creating scope for substantial rebates and variation in negotiated prices and patient cost sharing. Within each plan-county-year, all covered drugs are ranked by *costsharing* and, to break ties, by *negotiatedpriceyearly*, *nadacyearly*, *specialty*, *branded*, and *ndc11*. The top-ranked (least costly) drug is selected; its values for *costsharing*, *negotiatedpriceyearly*, *nadacyearly*, *coinsurance*, *branded*, and *specialty* define the "minimum-drug" outcomes. To summarize price levels over time, I compute fixed-basket Jevons indices for each price concept $p \in \{\textit{nadacyearly}, \textit{negotiatedpriceyearly}, \textit{costsharing}, \textit{coinsurance}\}$:

$$\text{Jevons}_{p,t} = 100 \times \exp\left(\frac{1}{N} \sum_{i \in B_{2018}} \log \frac{p_{it}}{p_{i,2018}}\right),$$

where B_{2018} is the 2018 basket. The resulting variables are *nadacyearlyjevons*, *negotiatedpriceyearlyjevons*, *costsharingjevons*, and *coinsurancejevons*.

Treatment variables capture the staggered adoption of state-level rebate-transparency laws: *treatedg* equals one for states that ever adopt a disclosure law; *cohort* records the first adoption year; and *dummygt* equals one in and after the adoption year. Because many contracts operate in multiple states, a plan-county in an untreated state is coded as indirectly treated ($\textit{treatedgspill} = 1$) if the same contract operates in any treated state that year; *dummygtspill* equals one in those periods.

Pre-policy characteristics define heterogeneity splits. *mlrbelow85split* is the modal pre-policy value of *mlrbelow85* for each plan. *enrollmentcontractsplit* equals one for plans whose pre-policy average enrollment exceeds the national median. *hhicountysplit* equals one for counties whose pre-policy Herfindahl-Hirschman Index of enrollment shares lies below the median. The modal pre-policy PBM affiliation defines *planpbmsplit*.

The final dataset is an unbalanced panel of plan-county-year observations from 2016-2025. Each record contains identifiers (*CONTRACT_ID*, *PLAN_ID*, *SEGMENT_ID*, *contractplansegmentid*, *statefips*, and county FIPS), plan characteristics (*premiumyearly*, *deductible*, *Enrollment*, *planpbm*, *nonprofit*, *singlestate*, and *mlrbelow85*), pricing and cost-sharing outcomes (*negotiatedpriceyearly*, *nadacyearly*, *cost-*

sharing, *coinsurance*, and their Jevons indices), drug attributes (*ndc11*, *RXCUI*, *specialty*, *branded*), and treatment/heterogeneity indicators (*treatedg*, *cohort*, *dummygt*, *treatedgspill*, *dummygtspill*, *planpbm-split*, *mlrbelow85split*, *enrollmentcontractsplit*, *hhicountysplit*). No imputation is performed.

The Transparency Hold-up: Pharmacy Benefit Managers and Rebate Disclosure

Nicola Maria Fiore *

April 17, 2026

Abstract

This paper examines contracting when an intermediary privately observes realized revenue that is meant to be shared with a principal. Motivated by current rebate disclosure proposals in intermediary contracting—most visibly in the PBM setting linked to high drug prices—we develop a simple theory of contracting under secrecy versus disclosure. In our model, rebates are initially hidden from the principal. Under secrecy, the intermediary can retain part of the rebate, limiting pass-through but preserving upside. Disclosure makes rebates verifiable and allows costly recovery of retained amounts. While disclosure increases pass-through of realized rebates, it compresses the intermediary's high-state rents and weakens incentives to generate them. Transparency improves principal welfare only when enforcement is sufficiently credible and incentives are not highly rent-sensitive.

Keywords: Transparency, Mandatory Disclosure, Verifiability, Contracting, Real Effects.
JEL Codes: D82, D86, M41, L22.

*Bocconi University, Department of Accounting. Correspondence: via personal webpage nicola-maria-fiore.gitlab.io/project/.

1 Introduction

Transparency regulation is often justified as a way to improve contracting and discipline intermediaries. By making performance measures observable and verifiable, disclosure can tighten the link between compensation and true output, reduce agency costs, and enhance surplus. A large literature formalizes this logic: more precise and more contractible information improves incentive design and mitigates opportunistic behavior. Yet, in many settings the disclosed measure is not purely informational. Transparency can also change the environment that generates the reported outcome in the first place. When reporting requirements affect how agents allocate effort or negotiate with counterparties, the welfare effects of disclosure are no longer unambiguously positive.

Recent policy debates surrounding pharmacy benefit managers (PBMs) illustrate this tension. PBMs negotiate rebates from drug manufacturers on behalf of health plans and employers. These rebates can represent a substantial fraction of total drug spending and have become a focal point in discussions about high prescription drug prices. Policymakers have proposed greater transparency in rebate contracting, including mandatory disclosure of gross rebates and restrictions on rebate retention. The underlying premise is straightforward: if rebate flows are fully visible and verifiable, intermediaries will be less able to retain a portion of the negotiated surplus, and more savings will flow to plans and patients.

This logic reflects a standard contracting intuition. When realized revenue is observable only to the intermediary, it is difficult for the principal to verify how much surplus was created and how much was passed through. Disclosure promises to reduce this asymmetry. If gross rebates become verifiable, plans can demand greater pass-through or recover retained amounts. From this perspective, transparency strengthens enforcement and improves surplus division.

At the same time, rebate revenue is not exogenous. It depends on costly negotiation effort, expertise, and investment in manufacturer relationships. These inputs are difficult to observe and contract upon. If disclosure reduces the intermediary's ability to retain a share of the incremental surplus generated in successful negotiations, it may also reduce the incentive to generate that surplus in the first place. In other words, transparency may improve how realized rebates are divided while weakening the incentives that produce them. Whether mandatory rebate disclosure increases total surplus therefore depends on the balance between these two forces.

This paper develops a simple model to study this tradeoff. A health plan hires a PBM to negotiate manufacturer rebates. The PBM exerts costly effort that increases the likelihood of generating a high rebate. After rebates are realized, the PBM chooses how much to remit and how much to retain. Retention is feasible but costly, reflecting legal exposure, regulatory scrutiny, reputational concerns, and the risk of losing clients. We compare two regimes that differ only in what is verifiable. Under secrecy, only remitted amounts are observable and enforceable. Gross rebates remain private information, and retained amounts cannot be contractually recovered. Under disclosure, realized rebates become verifiable. The plan can attempt to claw back retained amounts, but enforcement is costly and cannot eliminate the intermediary's payoff entirely.

The model highlights a clear mechanism. Under secrecy, the PBM retains part of the high rebate when it occurs. These retained amounts are socially wasteful from a static perspective—they reflect surplus that does not flow to the plan and are partly dissipated through retention costs. However, they also provide the reward that motivates effort. The prospect of keeping some share of a high rebate increases the return to negotiation effort and raises the probability that such rebates are generated. Under disclosure, by contrast, retention becomes recoverable. Pass-through increases in the high-rebate state because the plan can enforce greater remittance. Yet this very improvement compresses the PBM's upside. By reducing the payoff difference between high and low outcomes, disclosure weakens incentives to exert effort.

The overall effect on the plan's welfare is therefore ambiguous. Disclosure unambiguously increases pass-through conditional on a high rebate. But it may reduce the probability of achieving that high rebate in the first place. If enforcement is sufficiently credible and negotiation effort is not highly sensitive to retained surplus, the gains from improved pass-through dominate. Transparency increases expected rebates received by the plan. In contrast, if effort is highly responsive to the PBM's residual claim, rent compression substantially reduces rebate generation. In this case, secrecy can sustain higher expected surplus despite lower pass-through in successful states.

The analysis yields sharp comparative statics. Disclosure is more likely to improve outcomes when enforcement costs are low and when retention costs are modest, so that secrecy allows substantial rent extraction. It is less likely to be beneficial when effort responds strongly to high-state rewards, making incentive effects quantitatively important. These results clarify that transparency does not mechanically expand the feasible contracting set in environments with hidden effort and costly enforcement. Once renegotiation and ex post recovery are possible, making realized revenue verifiable changes both the division and the creation of surplus.

Although motivated by rebate disclosure in pharmaceutical markets, the mechanism applies more broadly to intermediary relationships in which realized revenue is privately observed and partially shared. Examples include procurement agents negotiating supplier discounts, platforms collecting transaction fees, or asset managers earning trading gains. In each case, greater verifiability strengthens ex post discipline but may alter ex ante incentives to generate surplus. By isolating this mechanism, our paper contributes to the literature on disclosure and real effects by showing how transparency can simultaneously improve surplus extraction and weaken surplus creation.

The broader implication is that evaluating transparency regulation requires more than documenting changes in pass-through or observed payments. It also requires assessing how disclosure reshapes incentives to produce the underlying performance measure. In settings where effort is hidden and enforcement is costly, transparency introduces a fundamental tradeoff. Understanding that tradeoff is essential for designing disclosure policies that improve, rather than inadvertently reduce, overall surplus.

2 Literature

Our paper contributes to the literature on disclosure and its real effects by studying mandatory transparency in a contracting environment with an intermediary and renegotiation. Classic discretionary disclosure models show that managers optimally withhold information when disclosure is costly, even if disclosure improves outsiders' beliefs about firm value (Verrecchia, 1983; Dye, 1985). Subsequent work emphasizes that the quality and precision of information affect equilibrium disclosure and pricing, and that greater transparency need not be welfare improving once proprietary and strategic considerations are taken into account (Verrecchia, 1990). A related strand studies disclosure in competitive environments, demonstrating that reporting can alter product-market outcomes and entry decisions, thereby generating real effects through strategic channels (Darrough and Stoughton, 1990; Wagenhofer, 1990; Darrough, 1993). These studies primarily focus on how disclosure changes beliefs, prices, and strategic interaction. In contrast, we study how mandatory disclosure alters the allocation of ex post surplus in a bilateral contracting relationship, thereby affecting ex ante incentives.

Our analysis is also closely related to the literature comparing voluntary and mandatory disclosure regimes. Dye (1990) shows that mandatory disclosure can be desirable when private disclosure choices generate externalities, but that the welfare comparison depends on the nature of those externalities. Gigler (1994) highlights that disclosure can be self-enforcing in equilibrium, limiting opportunism even in the absence of full verifiability. More recent work formalizes how mandatory disclosure shapes reporting asymmetries and equilibrium information acquisition (Bertomeu and Magee, 2015; Bertomeu et al., 2021). We complement this literature by analyzing a setting in which mandatory disclosure changes not only the information structure but also the renegotiation environment. Once performance becomes verifiable, the principal can extract surplus ex post subject to the agent's outside option. Disclosure therefore affects not only beliefs but also bargaining power, introducing a tradeoff between governance gains and weakened production incentives.

Our paper further relates to the "real effects" literature, which studies how accounting information affects investment and contracting incentives. Periodic reporting can discipline managers but may also distort real decisions (Kanodia and Lee, 1998; Gigler and Hemmer, 1998). Transparency can improve efficiency by mitigating hidden actions, yet it can also induce over- or underinvestment depending on how prices respond to reported signals (Gigler and Hemmer, 2004). A large body of work shows that disclosure rules influence production, hedging, and investment decisions, and that more precise or more comprehensive measurement can reduce or increase surplus depending on the equilibrium response (Kanodia et al., 2000, 2004, 2005; Ewert and Wagenhofer, 2005; Kanodia and Venugopalan, 2023). In much of this literature, disclosure affects real decisions by changing market prices or monitoring intensity. In our setting, disclosure affects real decisions through a different channel: by strengthening ex post surplus extraction in the presence of noncontractible effort and unavoidable renegotiation.

Finally, our setting shares features with models of disclosure and contracting under asymmetric information. Sridhar and Magee (1996) examine how disclosure interacts with financial contracting and opportunism, while Bertomeu and Marinovic (2016) distinguish between hard (verifiable) and soft information in shaping reporting equilibria. We study a related transformation: mandatory rebate disclosure converts a privately observed gross outcome into a contractible measure after performance is realized. Because contracts cannot preclude renegotiation once new information becomes verifiable, this transformation changes the division of realized surplus. Unlike standard models in which improved verifiability weakly expands the feasible contracting set, we show that making the performance measure contractible can simultaneously weaken the agent's marginal claim on output and reduce effort. This interaction generates a non-monotone welfare effect of disclosure and links the disclosure literature's emphasis on verifiability to the real-effects literature's focus on endogenous distortions in production and investment.

3 Institutional Setting

Consider a firm that wants to lower the prices it pays but lacks the expertise or scale to negotiate effectively on its own. It hires a specialized intermediary and promises compensation tied to the financial gains the intermediary generates. The central question is straightforward: what must the firm observe in order to reward performance and limit appropriation? The market for prescription drug rebates provides a sharp and economically important instance of this general contracting problem.

Health plans finance prescription drugs for their enrollees. Rather than negotiating directly with manufacturers, most plans hire a pharmacy benefit manager (PBM). The PBM designs the drug benefit and negotiates rebate agreements with manufacturers. These rebates are payments made after drugs are dispensed, typically tied to utilization and to contractual terms secured in negotiation. The magnitude of rebates can be substantial relative to total drug spending.

The size of the rebate depends in part on how aggressively and effectively the PBM negotiates. Negotiation requires expertise, information, and effort. These inputs are costly and largely unobservable to the plan. The plan sees realized financial transfers but does not observe how hard the PBM pushed in bargaining or the full set of contractual contingencies embedded in manufacturer agreements. As in many delegated negotiation settings, effort affects surplus, but effort itself is not contractible.

A second feature is that gross rebates generated by negotiation are not always transparent to the plan. Manufacturer contracts are typically between the PBM and the manufacturer and are treated as confidential. The plan commonly observes the amounts the PBM transfers under their bilateral agreement, but it may not observe the full formula that maps utilization into gross rebate payments. Contractual terms can specify which components are passed through, which are retained, and how fees or adjustments are netted. As a result, the plan's direct visibility into total rebate generation depends on the reporting and audit rights specified in the contract or imposed by regulation.

Retention is feasible but constrained. A PBM that retains a larger share of rebate-related cash flows faces costs, including client disputes, audit exposure, regulatory scrutiny, and reputational damage. These costs tend to increase with the magnitude retained. At the same time, when negotiated rebates are larger, the scope for retention is larger.

Relationships between plans and PBMs are ongoing and subject to renegotiation. When additional information about realized rebate performance becomes available, plans can press for revised terms. Because renegotiation cannot be ruled out ex ante, improvements in verifiability affect not only how surplus is divided after the fact but also how much effort the PBM is willing to invest in negotiation in the first place. The setting thus offers a clear and economically meaningful environment in which disclosure, bargaining power, and real effort interact.

4 Model

We study a contractual relationship between a health plan (principal) and a pharmacy benefit manager (PBM, agent). The PBM exerts costly effort to generate manufacturer rebates. Effort is privately chosen and noncontractible. The sole institutional difference across regimes is whether realized rebates are verifiable and therefore enforceable.

4.1 Environment

A health plan (the principal) delegates manufacturer negotiation to a pharmacy benefit manager (PBM, the agent). The economic friction is that generating rebates requires costly expertise and effort that the plan cannot directly perform itself. The PBM can influence the rebate it obtains by exerting costly negotiation effort. Formally, the PBM chooses effort $a \geq 0$. A higher a raises the probability of obtaining a high rebate. The realized rebate is binary, $r \in \{0, \bar{r}\}$, where $\bar{r} > 0$ denotes a successful negotiation outcome. The probability of success is $q(a)$, with $q'(a) > 0$ and $q''(a) < 0$. The monotonicity of q captures that effort improves bargaining outcomes; concavity captures diminishing returns to negotiation intensity, expertise deployment, or relationship investment.

Effort is privately costly. The PBM incurs cost $c(a)$ with $c'(a) > 0$ and $c''(a) > 0$. Convexity reflects increasing marginal costs of negotiation intensity: initial improvements may be relatively inexpensive (e.g., basic analytics or routine bargaining), while pushing manufacturers further requires disproportionately greater managerial time, strategic concessions, or relationship capital. The convexity assumption ensures an interior tradeoff between effort and its expected benefit.

Once the rebate is realized, the PBM determines how much to pass through to the plan. Let $y \in [0, r]$ denote the remitted amount. The remainder is retained by the PBM and is denoted by $x \equiv r - y$. Retention is technologically feasible: the PBM controls the manufacturer contracts and rebate flows and can structure transfers in ways that allow it to keep part of the negotiated surplus.

However, retention is not costless. Retaining rebates generates private costs summarized by a function $g(x)$ satisfying $g(0) = 0$, $g'(x) > 0$, and $g''(x) > 0$. The function $g(\cdot)$ captures a broad set of frictions. First, legal and regulatory exposure typically increases with the magnitude of retained rebates. Second, greater retention increases the risk of audit disputes, client dissatisfaction, and reputational damage. Third, in repeated commercial relationships, low pass-through may lead the plan to renegotiate aggressively or switch PBMs altogether. In practice, a PBM that systematically shares little or nothing would have difficulty retaining business. These considerations imply that while some retention may be feasible, larger retention becomes increasingly costly. Convexity of g reflects that these risks escalate at the margin: modest retention may be tolerated or obscured, whereas extreme retention is likely to trigger substantial penalties or client exit.

Importantly, $g(x)$ is a real resource cost to the PBM. It does not represent a transfer to the plan but rather legal expenses, reputational losses, organizational distortions, or other privately borne losses associated with withholding funds. The assumption $g(0) = 0$ simply states that full pass-through entails no such cost.

The plan also pays the PBM a fixed transfer $w \geq 0$, interpreted as an administrative fee or contractual payment that does not depend on realized outcomes. For given choices (a, y) and realization r, payoffs are

$$\Pi_P = y - w \qquad \text{and} \qquad \Pi_A = w + x - g(x) - c(a),$$

where $x = r - y$. The plan values the amount remitted net of the fixed payment. The PBM values the fixed payment and any retained rebates, but retention is offset by the private costs $g(x)$ and effort cost $c(a)$.

4.2 Contracting and Enforcement

The relationship is modeled as a sequential game with perfect recall. In each stage, a single player makes a payoff-maximizing decision, taking as given the continuation strategies in subsequent stages. We describe the timing and optimization problems in order.

Stage 1: Contract offer.

The plan moves first and offers a contract consisting solely of a fixed transfer $w \geq 0$. The contract cannot condition on realized rebates. This restriction isolates the role of verifiability in shaping ex post enforcement, rather than allowing full state-contingent contracting at the outset.

Formally, the plan chooses w anticipating the PBM's optimal effort, retention, and (under disclosure) the plan's own optimal enforcement decision. The plan's expected payoff is

$$\mathbb{E}\big[y + t - w - k\,\mathbf{1}\{t > 0\}\big],$$

where the expectation is taken over rebate realizations induced by the PBM's effort choice. The choice variable w enters the plan's payoff directly as a transfer $-w$, and indirectly through the PBM's continuation value, affecting effort incentives, retention behavior, and the feasibility of enforcement. In addition, the contract must satisfy the PBM's ex ante participation constraint: the PBM's expected utility under the contract must be weakly greater than its outside option, which we normalize to zero.

Stage 2: Effort choice.

After observing w, the PBM chooses effort $a \geq 0$. Effort affects the distribution of rebates and enters the PBM's payoff in two places: through the success probability $q(a)$ and through the effort cost $c(a)$.

The rebate realization satisfies

$$r \in \{0, \bar{r}\}, \qquad \Pr(r = \bar{r} \mid a) = q(a),$$

with $q'(a) > 0$ and $q''(a) < 0$. Thus, the choice variable a shifts the probability of realizing the high rebate $\bar{r}$ via $q(a)$, while the PBM bears the direct cost $c(a)$, with $c'(a) > 0$ and $c''(a) > 0$.

To state the PBM's effort problem transparently, define its *ex post continuation payoff* after a rebate realization r (but before effort costs are subtracted) as

$$U(r;w) \;\equiv\; w + x^*(r) \;-\; g\big(x^*(r)\big) \;-\; t^*\big(x^*(r)\big),$$

where $x^*(r)$ denotes the PBM's optimal retention choice in Stage 4 given realization r (and anticipating enforcement), and $t^*(\cdot)$ denotes the plan's optimal recovery rule in Stage 5 under disclosure (with $t^*(\cdot) \equiv 0$ under secrecy). This definition makes clear how the fixed fee w enters additively, and how retention enters both through the direct benefit $+x^*(r)$ and the private cost $-g(x^*(r))$, as well as through any induced recovery $-t^*(x^*(r))$.

Given the binary distribution of rebates, the PBM's expected continuation payoff conditional on effort a can be written explicitly as

$$\mathbb{E}[U(r;w) \mid a] \;=\; q(a)\, U(\bar{r};w) \;+\; \big(1 - q(a)\big)\, U(0;w).$$

This expression shows exactly where the effort choice a enters: *only through* the probability weights $q(a)$ and $1 - q(a)$. The continuation payoffs $U(\bar{r};w)$ and $U(0;w)$ are determined by later-stage optimal choices given the realized rebate, and do not depend directly on a.

The PBM therefore chooses effort to maximize expected continuation value net of effort cost:

$$\max_{a \geq 0} \Big\{ q(a)\, U(\bar{r};w) \;+\; \big(1 - q(a)\big)\, U(0;w) \Big\} \;-\; c(a).$$

Equivalently, substituting the definition of $U(\cdot;w)$ and using linearity of expectation, the effort problem can be written in a compact but fully transparent form as

$$\max_{a \geq 0} \; w \;+\; \mathbb{E}\big[x^*(r) - g\big(x^*(r)\big) - t^*\big(x^*(r)\big) \;\big|\; a\big] \;-\; c(a),$$

where the conditional expectation is taken with respect to $\Pr(r = \bar{r} \mid a) = q(a)$. This form makes two points explicit: (i) w can be taken outside the expectation because it does not vary with r; and (ii) a affects expected payoffs only by shifting the distribution of r through $q(a)$.

Stage 3: Rebate realization.

Nature draws $r \in \{0, \bar{r}\}$ according to the distribution determined by a. The PBM observes r. Under disclosure, r is verifiable; under secrecy, it is not.

Stage 4: Remittance and retention.

After observing r, the PBM chooses remittance $y \in [0, r]$, equivalently retention $x = r - y \in [0, r]$. The PBM's continuation payoff at this stage is

$$w + x - g(x) - t,$$

where t is the recovery imposed in the enforcement stage (with $t = 0$ under secrecy).

The choice variable in this stage is x (or equivalently y). Retention x enters the PBM's payoff positively as $+x$, negatively through the private cost $-g(x)$, and negatively through any recovery $-t$ that may be imposed subsequently. The PBM solves

$$\max_{x \in [0,r]} \; w + x - g(x) - t.$$

Under secrecy, $t \equiv 0$. Under disclosure, $t = t^*(x)$ is determined by the plan's enforcement decision in the next stage.

Stage 5: Enforcement (disclosure regime only).

Under secrecy, the realized rebate r is not verifiable, so retention x is not contractible and no recovery is feasible. Formally, $t = 0$.

Under disclosure, r and therefore x are verifiable. After observing x, the plan chooses a recovery transfer $t \geq 0$. Enforcement is bounded by two constraints. First, recovery cannot exceed retained funds:

$$0 \leq t \leq x.$$

Second, enforcement must leave the PBM with nonnegative continuation value in this stage:

$$w + x - t - g(x) \geq 0.$$

If $t > 0$, the plan incurs a fixed enforcement cost $k > 0$.

Given x, the plan solves

$$\max_{t \geq 0} \; t - k\,\mathbf{1}\{t > 0\} \quad \text{s.t.} \quad 0 \leq t \leq x, \; w + x - t - g(x) \geq 0.$$

The choice variable t enters the plan's payoff directly as $+t$ and through the fixed cost $-k\,\mathbf{1}\{t > 0\}$. It enters the PBM's payoff as $-t$. Let $t^*(x)$ denote the solution to this problem.

Equilibrium payoffs.

For realized (r, x) and induced recovery t, payoffs are

$$\Pi_P = r - x + t - w - k\,\mathbf{1}\{t > 0\}, \qquad \Pi_A = w + x - t - g(x) - c(a).$$

The equilibrium is defined by sequential rationality: (i) the plan chooses w anticipating optimal continuation behavior; (ii) the PBM chooses a anticipating optimal retention and enforcement; (iii) the PBM chooses x anticipating enforcement; and (iv) under disclosure, the plan chooses t optimally given x.

5 Results

We characterize equilibrium behavior under each regime. In each case, we solve backward. Given the contract, we determine the PBM's remittance and effort choices. The participation constraint then pins down compensation. We focus on interior effort whenever it exists.

5.1 Secrecy

Under secrecy, the realized gross rebate r is observed by the PBM but is not verifiable to a third party. As a consequence, retention $x \equiv r - y$ is not contractible and the plan cannot impose any ex post recovery. In the notation of the previous subsection, the enforcement stage is inactive and we impose $t \equiv 0$ (so the plan never incurs the fixed enforcement cost k). We solve by backward induction.

Stage 4 (Retention and remittance). Fix a realized rebate $r \in \{0, \bar{r}\}$. The PBM chooses retention $x \in [0, r]$ (equivalently remittance $y = r - x$) to maximize its continuation payoff. Under secrecy, this problem is

$$\max_{x \in [0,r]} \; w + x - g(x).$$

Because w enters additively and does not vary with x, the PBM's retention choice is determined entirely by the tradeoff between the marginal benefit of retention and the marginal private cost $g'(x)$. Since g is increasing and strictly convex with $g(0) = 0$, the objective $x - g(x)$ is strictly concave and admits a unique unconstrained maximizer $x^\circ \geq 0$ characterized (when interior) by

$$g'(x^\circ) = 1.$$

Imposing feasibility $x \leq r$ yields the optimal retention rule

$$x^*(r) = \min\{x^\circ, r\}, \qquad y^*(r) = r - x^*(r).$$

In particular, $x^*(0) = 0$ and $x^*(\bar{r}) = \min\{x^\circ, \bar{r}\}$.

Stage 2 (Effort). After observing w, the PBM chooses effort $a \geq 0$, which affects the rebate distribution through $\Pr(r = \bar{r} \mid a) = q(a)$. Under secrecy, the PBM anticipates the Stage-4 retention rule above and obtains no enforcement-related transfers. Since r is binary, the PBM's expected payoff can be written explicitly as

$$q(a)\Big[w + x^*(\bar{r}) - g\big(x^*(\bar{r})\big)\Big] + \big(1 - q(a)\big)\Big[w + x^*(0) - g\big(x^*(0)\big)\Big] - c(a).$$

Using $x^*(0) = 0$ and $g(0) = 0$, and pulling the constant fee w outside the expectation, the effort problem is equivalently

$$\max_{a \geq 0} \quad w + q(a)\Big[x^*(\bar{r}) - g\big(x^*(\bar{r})\big)\Big] - c(a).$$

This expression makes clear that a enters payoffs only through $q(a)$ and $c(a)$: effort shifts the probability of the high rebate via $q(a)$ and is privately costly through $c(a)$, while the continuation term in brackets is determined by the PBM's optimal retention in the high-rebate state.

Any interior optimizer a^* satisfies the first-order condition

$$q'(a^*)\Big[x^*(\bar{r}) - g\big(x^*(\bar{r})\big)\Big] = c'(a^*),$$

with the usual boundary solution $a^* = 0$ when the marginal return to increasing $q(a)$ is insufficient.

Stage 1 (Contract). Given the equilibrium retention rule, the plan receives $y^*(r) = r - x^*(r)$ and pays w. The plan's expected payoff is therefore

$$\mathbb{E}[\Pi_P] = q(a^*)\big[\bar{r} - x^*(\bar{r})\big] - w.$$

The contract must also satisfy the PBM's ex ante participation constraint. Normalizing the PBM's outside option to zero, we require

$$\mathbb{E}[\Pi_A] = w + q(a^*)\big[x^*(\bar{r}) - g\big(x^*(\bar{r})\big)\big] - c(a^*) \ \geq \ 0.$$

Because the fixed fee w is a lump-sum transfer that does not depend on realized rebates, it does not affect the PBM's marginal incentives for retention in Stage 4 or effort in Stage 2. Accordingly, in equilibrium the plan sets w equal to the smallest nonnegative value satisfying the participation constraint. In particular, when the constraint binds,

$$w^* = \max\big\{0,\ c(a^*) - q(a^*)\,\big[x^*(\bar{r}) - g\big(x^*(\bar{r})\big)\big]\big\}\,.$$

provided the right-hand side is nonnegative. Since w enters additively in the PBM's objective, it does not affect the marginal conditions that determine retention or effort. The equilibrium can therefore be characterized sequentially: first solve for the retention rule $x^*(\cdot)$ and effort level a^* from the marginal conditions in Stages 4 and 2, and then determine w^* from the binding participation constraint. In particular, w^* adjusts to satisfy participation but does not feed back into the PBM's incentive tradeoffs.

Proposition 1 (Secrecy equilibrium)**.** *Suppose realized rebates are not verifiable, so that no ex post* *recovery is feasible and $t \equiv 0$. Then the game admits a (pure-strategy) equilibrium characterized as follows.*

(i) Retention and remittance. *Let x° denote the unique maximizer of $x - g(x)$ on $[0, \infty)$, characterized (when interior) by $g'(x^\circ) = 1$. After observing $r \in \{0, \bar{r}\}$, the PBM retains*

$$x^*(r) = \min\{x^\circ, r\}, \qquad \text{and remits} \qquad y^*(r) = r - x^*(r).$$

In particular, $x^*(0) = 0$ and $x^*(\bar{r}) = \min\{x^\circ, \bar{r}\}$.

(ii) Effort. *Anticipating the retention rule in (i), the PBM chooses effort $a \geq 0$ to solve*

$$\max_{a \geq 0} \;\; w + q(a)\Big[x^*(\bar{r}) - g\big(x^*(\bar{r})\big)\Big] - c(a).$$

Any interior equilibrium effort a^* satisfies

$$q'(a^*)\Big[x^*(\bar{r}) - g\big(x^*(\bar{r})\big)\Big] = c'(a^*),$$

and otherwise $a^* = 0$.

(iii) Contract. *The plan's expected payoff is*

$$\mathbb{E}[\Pi_P] = q(a^*)\big[\bar{r} - x^*(\bar{r})\big] - w.$$

Since w is not state-contingent, it does not enter the PBM's marginal conditions for either retention or effort. In equilibrium, the plan sets the fee equal to the smallest nonnegative value that satisfies the PBM's ex ante participation constraint. That is,

$$w^* = \max\big\{0,\; c(a^*) - q(a^*)\big[x^*(\bar{r}) - g\big(x^*(\bar{r})\big)\big]\big\}.$$

The equilibrium retention rule in (i) and the effort condition in (ii) remain unchanged, as the fixed fee does not affect marginal incentives.

The secrecy regime therefore yields a simple structure. Ex post, retention is disciplined only by the PBM's private cost of withholding, summarized by $g(\cdot)$. Ex ante, effort incentives are governed by the PBM's net private return from the high-rebate state,

$$x^*(\bar{r}) - g\big(x^*(\bar{r})\big),$$

rather than by any enforceable pass-through obligation.

5.2 Disclosure

Under disclosure, the realized gross rebate r is verifiable. Consequently, retention $x \equiv r - y$ is verifiable as well, and the plan can initiate an ex post recovery of retained rebates. Enforcement is costly: any strictly positive recovery $t > 0$ triggers a fixed cost $k > 0$. We solve by backward induction, beginning with the plan's enforcement decision and then moving to the PBM's retention and effort choices.

Stage 5 (Enforcement). Fix a realized retention level $x \in [0, r]$. After observing x, the plan chooses a recovery transfer $t \geq 0$ to maximize its net enforcement payoff. Recovery is constrained by (i) feasibility, $t \leq x$, and (ii) the requirement that the PBM have nonnegative continuation value in the enforcement stage:

$$w + x - t - g(x) \geq 0.$$

If the plan sets $t > 0$, it incurs the fixed cost k. Thus, conditional on x, the plan solves

$$\max_{t \geq 0} \quad t - k\,\mathbf{1}\{t > 0\} \quad \text{s.t.} \quad 0 \leq t \leq x,\ w + x - t - g(x) \geq 0.$$

Because the plan's objective is increasing in t whenever $t > 0$, it either (i) does not enforce and sets $t = 0$, or (ii) enforces and chooses the largest t consistent with the constraints. The constraints imply

$$t \leq x \qquad \text{and} \qquad t \leq w + x - g(x),$$

so the maximal feasible recovery is

$$\hat{t}(x) \equiv \min\{x,\ w + x - g(x)\}.$$

Enforcement is attractive if and only if the net gain from enforcing is nonnegative, i.e., $\hat{t}(x) - k \geq 0$. Therefore the plan's equilibrium recovery rule is

$$t^*(x) = \begin{cases} \hat{t}(x) & \text{if } \hat{t}(x) \geq k, \\ 0 & \text{if } \hat{t}(x) < k. \end{cases}$$

This rule makes explicit how the choice variable t is pinned down by x and the enforcement technology. The quantity $w + x - g(x)$ is the maximum transfer consistent with leaving the PBM nonnegative continuation value in the enforcement stage, while x is the physical upper bound on recovery.

Stage 4 (Retention and remittance). Fix a realized rebate $r \in \{0, \bar{r}\}$. After observing r, the PBM chooses retention $x \in [0, r]$ (equivalently remittance $y = r - x$), anticipating the recovery rule $t^*(x)$ derived above. The PBM's continuation payoff at this stage is

$$w + x - g(x) - t^*(x),$$

so the PBM solves

$$\max_{x \in [0,r]} \quad w + x - g(x) - t^*(x).$$

A useful implication of the enforcement rule is that the PBM never finds it optimal to choose a retention level that induces enforcement. To see this, note that if $t^*(x) > 0$, then $t^*(x) = \hat{t}(x) \geq k$ and the PBM's continuation payoff is

$$w + x - g(x) - \hat{t}(x) \leq w + x - g(x) - x = w - g(x) < w,$$

where the strict inequality follows from $g(0) = 0$ and $g'(x) > 0$ for $x > 0$. Since the PBM can always choose $x = 0$ (which yields continuation payoff w and implies $t^*(0) = 0$), any x that triggers enforcement is strictly dominated. It follows that in equilibrium the PBM chooses retention in the region where enforcement is not undertaken, i.e., where $t^*(x) = 0$.

Accordingly, the PBM's retention choice under disclosure can be written as the solution to a constrained problem that makes the avoidance of enforcement explicit:

$$\max_{x \in [0,r]} \quad w + x - g(x) \quad \text{s.t.} \quad \hat{t}(x) < k,$$

where $\hat{t}(x) = \min\{x, w + x - g(x)\}$. The constraint $\hat{t}(x) < k$ is exactly the condition under which the plan optimally sets $t^*(x) = 0$.

As in the secrecy case, w enters additively and does not affect the marginal tradeoff between x and $g(x)$ inside the objective. The unconstrained maximizer of $x - g(x)$ remains the unique $x^\circ \geq 0$ characterized (when interior) by

$$g'(x^\circ) = 1.$$

Disclosure modifies the retention rule only through the additional requirement that retention remain in the non-enforcement region. Let

$$\bar{x}(r) \equiv \sup\big\{x \in [0,r] : \hat{t}(x) < k\big\},$$

the largest retention level feasible at rebate realization r that does not induce enforcement. Then the PBM's optimal retention satisfies

$$x^*(r) = \min\{x^\circ,\, \bar{x}(r)\}, \qquad y^*(r) = r - x^*(r).$$

In particular, since $r = 0$ implies $x = 0$, we have $x^*(0) = 0$ and $y^*(0) = 0$. For $r = \bar{r}$, the equilibrium retention is $x^*(\bar{r}) = \min\{x^\circ, \bar{x}(\bar{r})\}$, where $\bar{x}(\bar{r})$ is pinned down by the enforcement threshold $\hat{t}(x) < k$.

Stage 2 (Effort). After observing w, the PBM chooses effort $a \geq 0$, which affects the rebate distribution through $\Pr(r = \bar{r} \mid a) = q(a)$. Under disclosure, the PBM anticipates the Stage-4 retention rule $x^*(r)$ and the implied enforcement outcome. As shown above, equilibrium retention satisfies $t^*(x^*(r)) = 0$ for each realized r, so the PBM's expected payoff can be written explicitly as

$$q(a)\Big[w + x^*(\bar{r}) - g\big(x^*(\bar{r})\big)\Big] + \big(1 - q(a)\big)\Big[w + x^*(0) - g\big(x^*(0)\big)\Big] - c(a).$$

Using $x^*(0) = 0$ and $g(0) = 0$, and pulling the constant fee w outside the expectation, the effort problem is equivalently

$$\max_{a \geq 0} \;\; w + q(a)\Big[x^*(\bar{r}) - g\big(x^*(\bar{r})\big)\Big] - c(a).$$

As before, a enters payoffs only through $q(a)$ and $c(a)$. The bracketed term captures the PBM's net private return in the high-rebate state under disclosure, where the retention level $x^*(\bar{r})$ is disciplined by the enforcement threshold implicit in $\bar{x}(\bar{r})$.

Any interior optimizer a^* satisfies

$$q'(a^*)\Big[x^*(\bar{r}) - g\big(x^*(\bar{r})\big)\Big] = c'(a^*),$$

with boundary solution $a^* = 0$ when the marginal return to increasing $q(a)$ is insufficient.

Stage 1 (Contract). Given the equilibrium retention rule, the plan receives $y^*(r) = r - x^*(r)$ and pays the fixed fee w. Because equilibrium retention avoids enforcement, the plan receives no recovery and incurs no enforcement cost along the equilibrium path. The plan's expected payoff is therefore

$$\mathbb{E}[\Pi_P] = q(a^*)\big[\bar{r} - x^*(\bar{r})\big] - w.$$

The contract must satisfy the PBM's ex ante participation constraint. Normalizing the PBM's outside option to zero, we require

$$w + q(a^*)\big[x^*(\bar{r}) - g\big(x^*(\bar{r})\big)\big] - c(a^*) \;\geq\; 0.$$

Since the fixed fee is not contingent on the rebate realization, it does not affect the PBM's marginal incentives for retention or effort; it shifts payoffs only through a level transfer. Accordingly, in equilibrium the plan sets

$$w^* = \max\Big\{0,\ c(a^*) - q(a^*)\big[x^*(\bar{r}) - g\big(x^*(\bar{r})\big)\big]\Big\},$$

that is, the smallest nonnegative fee that satisfies the participation constraint.

Proposition 2 (Disclosure equilibrium)**.** *Suppose realized rebates are verifiable and enforcement entails a fixed cost $k > 0$. Then the game admits a (pure-strategy) equilibrium characterized as follows.*

(i) Enforcement. *For any realized retention x, define the maximal feasible recovery*

$$\hat{t}(x) = \min\{x,\ w + x - g(x)\}.$$

The plan's optimal recovery rule is

$$t^*(x) = \begin{cases} \hat{t}(x) & \text{if } \hat{t}(x) \geq k, \\ 0 & \text{if } \hat{t}(x) < k. \end{cases}$$

(ii) Retention and remittance. *After observing $r \in \{0, \bar{r}\}$, the PBM chooses retention to solve*

$$\max_{x\in[0,r]} \ w + x - g(x) - t^*(x).$$

In equilibrium the PBM selects x so that $t^*(x) = 0$, and thus solves

$$\max_{x\in[0,r]} \ w + x - g(x) \quad \text{s.t.} \quad \hat{t}(x) < k.$$

Let x° denote the unique maximizer of $x - g(x)$ on $[0, \infty)$, characterized (when interior) by $g'(x^\circ) = 1$, and let

$$\bar{x}(r) = \sup\{x \in [0, r] : \hat{t}(x) < k\}.$$

Then equilibrium retention and remittance satisfy

$$x^*(r) = \min\{x^\circ, \bar{x}(r)\}, \qquad y^*(r) = r - x^*(r).$$

In particular, $x^*(0) = 0$.

(iii) Effort. *Anticipating the retention rule in (ii), the PBM chooses effort $a \geq 0$ to solve*

$$\max_{a\geq 0} \quad w + q(a)\Big[x^*(\bar{r}) - g\big(x^*(\bar{r})\big)\Big] - c(a).$$

Any interior equilibrium effort a^* satisfies

$$q'(a^*)\Big[x^*(\bar{r}) - g\big(x^*(\bar{r})\big)\Big] = c'(a^*),$$

and otherwise $a^* = 0$.

(iv) Contract. *The plan's expected payoff is*

$$\mathbb{E}[\Pi_P] = q(a^*)\big[\bar{r} - x^*(\bar{r})\big] - w.$$

Although the fixed fee is not state-contingent and does not enter the PBM's marginal conditions directly, it affects the enforcement feasibility constraint under disclosure through the term $w + x - g(x)$. Accordingly, the equilibrium is characterized jointly in $(a^*, x^*(\bar{r}), w^*)$: retention and effort are determined taking w as given, and the equilibrium fee w^* is then pinned down by the binding participation constraint, with the understanding that w^* feeds back into enforcement feasibility. In equilibrium, the plan sets

$$w^* = \max\left\{0,\ c(a^*) - q(a^*)\big[x^*(\bar{r}) - g\big(x^*(\bar{r})\big)\big]\right\}.$$

The retention and effort rules in (i)–(iii) characterize equilibrium given w^*, which influences retention through the enforcement constraint.

5.3 The Transparency Hold-up

We now compare the equilibrium under secrecy and disclosure, incorporating the binding participation constraint characterized above. In each regime, the fixed fee w^* is pinned down by the PBM's participation constraint and therefore adjusts to redistribute surplus but does not enter marginal incentive conditions directly. Substituting the binding participation constraint into the principal's payoff implies that, in either regime, the plan captures total expected surplus net of concealment costs.

Let $x_S(\bar{r})$ and $x_D(\bar{r})$ denote equilibrium retention in the high-rebate state under secrecy and disclosure, respectively. Let a_S and a_D denote the corresponding equilibrium effort levels. Because $r = 0$ implies $x = 0$ in both regimes, all comparisons are driven by behavior in the state $r = \bar{r}$.

Retention.

Under secrecy, retention in the high-rebate state is

$$x_S(\bar{r}) = \min\{x^\circ, \bar{r}\},$$

where x° is the unique maximizer of $x - g(x)$.

Under disclosure, retention must lie in the non-enforcement region characterized by $\hat{t}(x) < k$, where $\hat{t}(x) = \min\{x, w_D^* + x - g(x)\}$ and w_D^* denotes the equilibrium fee under disclosure. Equilibrium retention therefore satisfies

$$x_D(\bar{r}) = \min\{x^\circ, \bar{x}(\bar{r})\},$$

where $\bar{x}(\bar{r})$ is the largest retention level at which enforcement is not undertaken.

Because enforcement restricts the feasible set of retention levels, it follows that

$$x_D(\bar{r}) \leq x_S(\bar{r}),$$

with strict inequality whenever enforcement would be triggered at the secrecy retention level $x_S(\bar{r})$. Thus disclosure weakly reduces equilibrium retention in the high-rebate state.

Effort.

In both regimes, effort solves a problem of the form

$$\max_{a \geq 0} \quad w^* + q(a)\,\Delta - c(a),$$

where Δ denotes the PBM's net private return in the high-rebate state.

Under secrecy,

$$\Delta_S = x_S(\bar{r}) - g\big(x_S(\bar{r})\big),$$

while under disclosure, because equilibrium retention avoids enforcement on the equilibrium path,

$$\Delta_D = x_D(\bar{r}) - g\big(x_D(\bar{r})\big).$$

Since $x_D(\bar{r}) \leq x_S(\bar{r})$ and $x - g(x)$ is strictly concave and increasing on $[0, x^\circ]$, it follows that

$$\Delta_D \leq \Delta_S,$$

with strict inequality whenever $x_D(\bar{r}) < x_S(\bar{r})$.

Equilibrium effort in each regime satisfies

$$q'(a_S)\Delta_S = c'(a_S), \qquad q'(a_D)\Delta_D = c'(a_D),$$

whenever interior. Because $q'(\cdot) > 0$ and $c'(\cdot)$ is increasing, the reduction in the net private return implies

$$a_D \le a_S,$$

with strict inequality whenever enforcement reduces retention in the high-rebate state.

Disclosure therefore weakly reduces equilibrium effort by reducing the PBM's residual claim on negotiated surplus.

Principal payoff.

When the participation constraint binds, the equilibrium fee satisfies

$$w^* = c(a^*) - q(a^*)\big[x^*(\bar{r}) - g\big(x^*(\bar{r})\big)\big],$$

provided the right-hand side is nonnegative. Substituting this expression into the principal's payoff implies that, in either regime,

$$V_P = q(a)\bar{r} - c(a) - q(a)g\big(x(\bar{r})\big).$$

Accordingly, the plan's expected payoff under secrecy is

$$V_P^S = q(a_S)\bar{r} - c(a_S) - q(a_S)g\big(x_S(\bar{r})\big),$$

while under disclosure it is

$$V_P^D = q(a_D)\bar{r} - c(a_D) - q(a_D)g\big(x_D(\bar{r})\big).$$

Disclosure has two opposing effects on the principal's payoff. First, by reducing retention it lowers concealment costs $g(x)$ in the high-rebate state. Second, by reducing the PBM's residual claim it lowers equilibrium effort and therefore the probability of realizing the high rebate.

The net effect is therefore ambiguous.

Proposition 3 (The Transparency Hold-up)**.** *Suppose enforcement binds at the secrecy retention level, so that $x_D(\bar{r}) < x_S(\bar{r})$. Then disclosure strictly reduces equilibrium effort, $a_D < a_S$. Moreover, secrecy yields a higher expected payoff for the plan whenever*

$$q(a_S)\bar{r} - c(a_S) - q(a_S)g\big(x_S(\bar{r})\big) \;>\; q(a_D)\bar{r} - c(a_D) - q(a_D)g\big(x_D(\bar{r})\big).$$

That is, secrecy is preferred whenever the reduction in expected surplus due to lower effort under disclosure exceeds the savings in concealment costs from reduced retention.

The mechanism is straightforward. Under disclosure, the plan can appropriate a larger share of realized surplus ex post, which constrains retention and reduces the PBM's residual claim. Anticipating this, the PBM exerts less effort. Although disclosure lowers concealment costs by reducing retention, it also depresses the probability of achieving the high rebate. When effort is sufficiently responsive to the private return from success, the induced reduction in expected surplus dominates the reduction in concealment costs, and the plan is strictly worse off under disclosure.

We refer to this distortion as the *transparency hold-up*. By making realized surplus verifiable and subject to recovery, disclosure strengthens the plan's ex post position but weakens ex ante incentives to generate surplus. Transparency can therefore reduce total expected surplus and, under the condition above, lower the principal's own payoff.

Although the principal may prefer ex ante to limit recovery in order to preserve effort incentives, enforcement is chosen based on ex post profitability. Because effort and retention are sunk at the enforcement stage, the principal optimally recovers surplus whenever the net gain exceeds the fixed enforcement cost.

6 Extensions

Thus far, we have compared two exogenously imposed regimes: secrecy and mandatory disclosure. In practice, however, one may ask whether transparency could arise voluntarily. If disclosure strengthens contracting and increases pass-through, why would the intermediary not choose to disclose rebate information on its own? We now show that, within our framework, neither complete nor partial hard disclosure arises in equilibrium when disclosure is voluntary.

6.1 Complete Voluntary Disclosure

Suppose that after observing the realized rebate r, the PBM can choose whether to make r verifiable. If it does so, the relationship proceeds under the disclosure regime characterized in Proposition 2: the principal can attempt ex post recovery subject to enforcement costs and limited liability. If the PBM does not disclose, the relationship remains in the secrecy regime characterized in Proposition 1.

The key observation is that disclosure expands the principal's ex post enforcement set without relaxing any constraint faced by the PBM. Once r becomes verifiable, the principal can claw back retained rebates whenever enforcement is privately profitable. Because effort and retention are sunk at the enforcement stage, the principal's recovery decision is governed solely by ex post profitability. The PBM anticipates this behavior.

For any fixed contract w and effort choice a, the PBM's continuation payoff after disclosure is weakly lower than under secrecy. Disclosure does not create any new source of surplus for the PBM; it merely exposes a larger portion of realized surplus to potential recovery. Since the principal adjusts the fixed fee w to satisfy the PBM's participation constraint, disclosure does not increase the PBM's equilibrium expected utility. When the participation constraint binds, the PBM's ex ante expected payoff is identical across regimes. When participation is slack due to limited liability on w, secrecy weakly dominates disclosure because it preserves a larger residual claim in the high-rebate state.

Accordingly, complete voluntary disclosure does not arise in equilibrium. The PBM strictly prefers to keep realized rebates nonverifiable whenever the principal cannot commit not to exploit newly verifiable information ex post.

6.2 Partial Hard Disclosure

We next consider whether the PBM might voluntarily disclose only part of the realized rebate. Suppose that after observing r, the PBM can choose a verifiable component $s \in [0, r]$. The principal can attempt recovery only from the disclosed component s, while the residual $r - s$ remains nonverifiable and governed by the secrecy technology. There is no direct reporting cost of choosing s; the only effect of disclosure is to expand the set of enforceable transfers.

As in the case of complete disclosure, increasing s weakly enlarges the principal's enforcement opportunities without creating any compensating benefit for the PBM. For any retention choice feasible under a given $s > 0$, the PBM can replicate (and weakly improve upon) its payoff by choosing a smaller disclosed amount. In particular, the PBM can always set $s = 0$, thereby reverting to the secrecy regime.

It follows that in any subgame following rebate realization, the PBM optimally chooses

$$s^*(r) = 0.$$

Anticipating this behavior, the principal sets the fixed fee and the equilibrium coincides with the secrecy outcome characterized in Proposition 1. Partial hard disclosure therefore does not arise in equilibrium.

6.3 Implications

The analysis establishes a general principle. In environments where newly verifiable information can be exploited ex post and the principal cannot commit not to renegotiate or enforce, voluntary hard disclosure does not emerge. Making realized performance contractible shifts bargaining power toward the principal after effort and retention decisions are sunk. Anticipating this ex post appropriation, the agent optimally keeps performance measures nonverifiable.

This result provides a theoretical rationale for the absence of voluntary rebate transparency in practice. Even if disclosure increases pass-through conditional on realization, it reduces the intermediary's residual claim and exposes realized surplus to recovery. Without a commitment device or regulatory mandate, the intermediary has no incentive to harden the performance measure.

Mandatory disclosure, in contrast, alters the feasible contracting environment by eliminating the agent's ability to withhold verifiability. The welfare consequences of such mandates therefore depend not only on improved enforcement, but also on the incentive distortions analyzed in the preceding sections.

7 Conclusion

This paper examines how disclosure affects contracting when performance is created by an agent but not immediately observable to the principal. In our setting, the PBM's effort generates rebate revenue that may initially remain private. Disclosure makes realized rebates verifiable and allows the plan to recover retained amounts through costly enforcement. By strengthening enforceability, disclosure increases the share of realized rebates that are passed through. At the same time, by limiting the PBM's upside in high-rebate states, it can reduce incentives to generate those rebates.

The analysis highlights a simple tradeoff. Under secrecy, the PBM retains part of incremental rebate revenue. These retained amounts are costly and imperfectly controlled, but they provide the reward that motivates effort. Under disclosure, retention becomes recoverable. This improves pass-through of realized rebates, yet it also compresses the payoff difference between high and low outcomes. Because effort responds to that payoff difference, rebate generation may decline.

Whether disclosure improves outcomes depends on which force is stronger. When enforcement is sufficiently credible and effort is not highly sensitive to the PBM's upside, disclosure increases principal surplus by reducing retention without materially affecting production. When effort is highly responsive to high-state rewards, disclosure can weaken incentives enough to offset gains from improved recovery. In such cases, secrecy may sustain greater rebate generation despite weaker pass-through.

The broader implication is that transparency is not uniformly beneficial in environments with hidden effort and costly enforcement. Making performance verifiable changes not only how realized surplus is divided, but also how much surplus is created. Evaluating disclosure policies therefore requires attention to both redistribution and real effects.

References

Bertomeu, J. and R. P. Magee (2015). Mandatory disclosure and asymmetry in financial reporting. *Journal of Accounting and Economics 59*(2-3), 284–299.

Bertomeu, J. and I. Marinovic (2016). A theory of hard and soft information. *The accounting review 91*(1), 1–20.

Bertomeu, J., I. Vaysman, and W. Xue (2021). Voluntary versus mandatory disclosure. *Review of Accounting Studies 26*(2), 658–692.

Darrough, M. N. (1993). Disclosure policy and competition: Cournot vs. bertrand. *Accounting review*, 534–561.

Darrough, M. N. and N. M. Stoughton (1990). Financial disclosure policy in an entry game. *Journal of accounting and economics 12*(1-3), 219–243.

Dye, R. A. (1985). Disclosure of nonproprietary information. *Journal of accounting research*, 123–145.

Dye, R. A. (1990). Mandatory versus voluntary disclosures: The cases of financial and real externalities. *Accounting Review*, 1–24.

Ewert, R. and A. Wagenhofer (2005). Economic effects of tightening accounting standards to restrict earnings management. *The Accounting Review 80*(4), 1101–1124.

Gigler, F. (1994). Self-enforcing voluntary disclosures. *Journal of Accounting Research 32*(2), 224–240.

Gigler, F. and T. Hemmer (1998). On the frequency, quality, and informational role of mandatory financial reports. *Journal of Accounting Research 36*, 117–147.

Gigler, F. B. and T. Hemmer (2004). On the value of transparency in agencies with renegotiation. *Journal of Accounting Research 42*(5), 871–893.

Kanodia, C. and D. Lee (1998). Investment and disclosure: The disciplinary role of periodic performance reports. *Journal of accounting research 36*(1), 33–55.

Kanodia, C., A. Mukherji, H. Sapra, and R. Venugopalan (2000). Hedge disclosures, future prices, and production distortions. *Journal of Accounting Research*, 53–82.

Kanodia, C., H. Sapra, and R. Venugopalan (2004). Should intangibles be measured: What are the economic trade-offs? *Journal of Accounting Research 42*(1), 89–120.

Kanodia, C., R. Singh, and A. E. Spero (2005). Imprecision in accounting measurement: Can it be value enhancing? *Journal of Accounting Research 43*(3), 487–519.

Kanodia, C. and R. Venugopalan (2023). Decision facilitating information and induced volatility: A study of tradeoffs in accounting disclosure. *The Accounting Review 98*(1), 317–336.

Sridhar, S. S. and R. P. Magee (1996). Financial contracts, opportunism, and disclosure management. *Review of Accounting Studies 1*(3), 225–258.

Verrecchia, R. E. (1983). Discretionary disclosure. *Journal of accounting and economics 5*, 179–194.

Verrecchia, R. E. (1990). Information quality and discretionary disclosure. *Journal of accounting and Economics 12*(4), 365–380.

Wagenhofer, A. (1990). Voluntary disclosure with a strategic opponent. *Journal of accounting and economics 12*(4), 341–363.

Fiscal Constraints, Disaster Vulnerability, and Corporate Investment Decisions

Nicola Maria Fiore, Thorsten Martin, and Florian Nagler*

November 2024

Abstract

This paper empirically examines how government fiscal constraints influence disaster vulnerability and investment decisions in a global sample of firms. We develop a novel firm-level measure of exposure to fiscal constraints based on firms' sales distributions across countries and combine it with computational linguistics tools applied to earnings calls to measure perceived risk. Using a difference-in-differences framework, we find that firms with greater exposure to fiscal constraints experience a stronger increase in perceived risk—specifically related to fiscal constraints—during disaster. Pre-disaster, these firms exhibit higher perceived risk and discount rates, translating into reduced investment in tangible capital and R&D. Our findings suggest that fiscal constraints affect growth through a risk-based disaster vulnerability channel.

JEL-Classification: F30, G12, G14, G15, H12, H50, H63
Keywords: Government fiscal constraints, disaster, discount rates, corporate investment, textual analysis, earnings conference calls

*First Version: April 10, 2020. Nicola Maria Fiore (email: nicola.fiore@unibocconi.it) is at Bocconi University, Via Roentgen 1, 20136 Milan, Italy. Thorsten Martin (email: t.martin@fs.de) is at Frankfurt School of Finance & Management, Adickesallee 32-34, 60322 Frankfurt, Germany, Bocconi University, and IGIER. Florian Nagler (corresponding author; email: florian.nagler@unibocconi.it; Via Roentgen 1, 20136 Milan, Italy) is at Bocconi University, IGIER, and CEPR. We thank seminar participants at Aarhus University, Bank of Portugal, Cass Business School, Frankfurt School of Finance & Management, London School of Economics and Political Science, University of Bonn, University of Utrecht, University of Southern Denmark, and PMP/ISK Vienna as well as participants at the Meeting of the Society for Economic Measurement and the Meeting of the Financial Intermediation Research Society (FIRS) for helpful comments and suggestions. We further thank Filippo De Marco, Carlo Favero, Nicola Gennaioli, Stefano Giglio, Seung Lee, Silvio Petriconi, Alexander Wagner, and Christian Wagner for helpful comments and discussions. We thank Felix Gerding for outstanding research assistance. The paper previously circulated under the titles 'The Value of Fiscal Capacity in the Face of a Rare Disaster' and 'Sovereign Debt and Equity Returns in the Face of Disaster.' Any remaining errors are our own.

1 Introduction

Disasters elevate uncertainty and can lead to severe economic downturns. Consequently, government intervention to mitigate the impact of such shocks features prominently in both economic theory and public policy. However, fiscal constraints weaken a government's ability to respond effectively to crises, such as by limiting demand support.[1] Moreover, such constraints may impair a government's creditworthiness during crises, which can further tighten financing conditions for firms. We hypothesize that fiscal constraints thus increase firms' vulnerability to disasters by amplifying disaster-induced increases in perceived risk. This heightened disaster vulnerability translates into elevated risk perception even in normal times, resulting in higher discount rates used by managers in investment decisions and, ultimately, depressed investment. In essence, fiscal constraints give rise to a risk-based disaster vulnerability channel, providing a novel explanation of why fiscal constraints are associated with lower growth (see, e.g., Alesina, Favero, and Giavazzi, 2020).

Testing this hypothesis is challenging for two reasons. First, studying the uncertainty-amplifying role of fiscal constraints during disasters requires a large, unexpected exogenous shock that impacts a broad cross-section of firms. Second, it necessitates observing otherwise comparable firms with varying levels of exposure to fiscal constraints. To address the first challenge, we use the outbreak of the covid pandemic as a natural experiment. The pandemic was unforeseen, as evidenced by significant changes in asset prices and growth expectations (Gormsen and Koijen, 2020; Landier and Thesmar, 2020), and it did not arise endogenously from fiscal constraints. As a global event, the pandemic allows us to analyze a cross-section of firms from different countries experiencing the same shock simultaneously.

[1] For instance, concerns about soaring sovereign debt levels around the globe are prevalent in the media and among policymakers. High public debt levels may render governments fiscally constrained, impeding a rapid and effective crisis response, potentially exacerbating economic downturns and contributing to instability (see, e.g., Economist, 2020; Financial Times, 2023; Wall Street Journal, 2021; Barron's, 2023). Figure IA.A.1 in the Internet Appendix, based on Factiva data, shows a sharp increase in news articles discussing risk in the context of fiscal constraints during disasters.

To tackle the second challenge, we construct a novel firm-level measure of exposure to fiscal constraints based on firms' sales distribution across countries. Our analysis then proceeds in two steps. First, we examine the perceived risk dynamics of firms when disaster strikes depending on their ex-ante exposure to fiscal constraints. Second, we relate firms' fiscal constraints exposure to perceived risk, discount rates, and investment in tangible capital and R&D in the years preceding the disaster.

We start by constructing an index of fiscal constraints at the country level, combining a country's debt-to-GDP ratio, credit rating, interest payments, and fiscal balance. We show that this index predicts lower government spending in disasters, in line with the notion of limited insurance provision by constrained governments. Next, we exploit firms' sales distribution across countries and combine them with the country-level index to construct a novel firm-level measure of exposure to fiscal constraints, FCE_f. To study whether fiscal constraints amplify the impact of disasters, we then employ a difference-in-differences (DiD) framework with continuous exposure. Specifically, we compare disaster-induced changes across firms with varying (ex-ante) FCE_f. Importantly, we saturate the specifications with headquarter country fixed effects, focusing on variation in fiscal constraints exposure arising from firms' sales distributions across countries. Additionally, we include industry fixed effects, allowing us to compare firms that are headquartered in the same country and operate in the same industry, but differ in their exposure to fiscal constraints.

We use this DiD strategy to estimate how firms' sensitivity to disasters varies based on their ex-ante exposure to fiscal constraints, providing insights into how these constraints influence firms' vulnerability to disasters. Note that the DiD approach does not require randomization of FCE_f across firms. Instead, it relies on the key identifying assumption that firms with different levels of FCE_f would have followed similar trends in the absence of the disaster. Three pieces of evidence support this parallel trends assumption. First, firms with different levels of FCE_f do not differ in ex-ante characteristics that might influence their disaster response (e.g., leverage, cash holdings, tangibility, size, market-to-book, profitability). Second, our key firm outcome of interest in the disaster, perceived risk related to fiscal constraints (measured from earnings calls), exhibits parallel trends before, and begins to diverge only with the onset of the disaster. Third, other non-fiscal constraint-related risk continues to follow parallel trends even after the onset of disaster.

In the first step, we study the dynamics of perceived firm-level risk during the disaster. We adapt the methodology developed by Hassan et al. (2019) and use tools from computational linguistics to analyze earnings conference calls. Our objective is to quantify the proportion of discussions in these calls that focus on risk related to fiscal constraints. To achieve this, we construct a training library for fiscal constraints terminology using the textbook by Abbas, Pienkowski, and Rogoff (2019). We then apply this library to the content of conference calls. To associate the language with risk, we incorporate synonyms for 'risk' and 'uncertainty' alongside the fiscal constraints terminology.

Figure 1 about here

Equipped with this proxy for perceived firm-level risk associated with fiscal constraints, we examine its relationship with FCE_f in the disaster. Using headquarter country fixed effects, Panel A of Figure 1 shows that perceived risk dynamics follow similar trends before the disaster. However, as the disaster unfolds, high-FCE_f firms experience an increase in perceived firm-level fiscal constraints risk compared to low-FCE_f firms. In Panel B, we conduct a falsification test by examining perceived firm-level risk unrelated to fiscal constraints. Reassuringly, risk unrelated to fiscal constraints continues to follow parallel trends as the disaster unfolds.

Next, we employ DiD regressions controlling for a range of firm characteristics, including other sales-weighted country characteristics. Crucially, in our preferred specification, we use a narrow fixed effects strategy, exploiting variation in FCE_f within the same headquarter (HQ) country and industry (Ind). This approach effectively controls for heterogeneity across firms related to their headquarters' location and sectoral effects. In this specification, we find that a one-unit increase in FCE_f corresponds to an 82% increase in perceived fiscal constraints risk in the disaster. Additionally, we confirm that FCE_f is unrelated to changes in non-fiscal constraints-related risk.

We conduct a series of additional tests to address two concerns. First, we verify that FCE_f is uncorrelated with perceived firm-level covid risk, as measured by Hassan et al. (2023a). Second, we ensure that FCE_f captures risk specifically related to fiscal constraints rather than broader country risk. To this end, we perform two additional tests. The first exercise is a falsification test: we construct a firm-level proxy for perceived country risk and find no association with FCE_f. In the second exercise, we run a placebo test and replace FCE_f with firm-level measures of ex-ante country risk (volatility of exchange rate and GDP growth). We find that these measures are unrelated to the disaster dynamics of perceived firm-level fiscal constraints risk. Taken together, the evidence shows that high-FCE_f firms exhibit an increase in perceived fiscal constraints risk during the disaster—but not in risk unrelated to fiscal constraints. Furthermore, broader measures of country risk do not spur an increase in fiscal constraints risk, indicating that the observed effects are specific to fiscal constraints.

After documenting that fiscal constraints exposure amplifies the increase in perceived risk in the disaster, we focus on the pre-disaster period. In a panel setup, we find that, over the five years leading up to the disaster, perceived firm-level fiscal-constraints risk is higher for firms with high FCE_f. This insight remains robust in our preferred specification with HQ×Ind×Year fixed effects. In a given year, fiscal constraints risk shows a positive correlation with FCE_f, for firms sharing otherwise identical characteristics, being headquartered in the same country and operating in the same industry. We observe a 57% higher perceived fiscal constraints risk for a one-unit increase in FCE_f.

A key question is whether the higher perceived risk affects the discount rates that firm management uses in investment decisions. Using the same panel regression setup, we find that managers of high-FCE_f firms indeed use larger discount rates in their investment decisions. Specifically, using data of Gormsen and Huber (2023a,b), we observe that discount rates are higher for firms with greater exposure to fiscal constraints. In our most saturated specification, with HQ×Ind×Year fixed effects, we obtain an increase in the discount rate of 0.7% for a one-unit rise in FCE_f. These results imply that firms with greater exposure to fiscal constraints should invest less during normal times. We provide direct evidence for this conjecture and show that larger FCE_f is indeed associated with lower firm-level capital as well as R&D investment in the pre-disaster period. Specifically, we obtain a 18% (49%) reduction in capital (R&D) investment for a one-unit increase in FCE_f.

Next we explore the mechanism underlying these results. We start by dissecting the discount rate effect. Gormsen and Huber (2023a) demonstrate the existence of a wedge between the discount rates that managers use in their investment decisions (i.e., the minimum required return on capital, or hurdle rate) and the perceived cost of capital—a firm's internal estimate, typically derived using a weighted average cost of capital (WACC) approach. We find that exposure to fiscal constraints does not correlate with firms' perceived cost of capital. Instead, the higher discount rates are fully accounted for by higher wedges. One interpretation of these facts is that i) the standard WACC approach fails to account for risk associated with fiscal constraints exposure, and ii) managers recognize this and adjust their discount rates accordingly.

Managers typically implement the textbook WACC approach by estimating the cost of equity using the CAPM applied to historical data (e.g., Graham and Harvey, 2001; Dessaint et al., 2021; Graham, 2022). However, fiscal constraints become particularly relevant during severe negative events. This suggests that the market beta estimated in normal times does not fully capture the risk arising from fiscal constraints exposure. This can explain why FCE_f is unrelated to the perceived cost of capital but is associated with higher wedges: managers of high-FCE_f firms are aware that their out-of-disaster market betas are too low and thus adjust their discount rates upward.

Motivated by these patterns in perceived cost of capital and discount rates, we use a simple conditional CAPM logic to provide some evidence on the asset pricing mechanism at play. Consistent with the idea that fiscal constraints matter through a disaster-vulnerability channel, we first show that FCE_f predicts higher pre-disaster CAPM alphas, suggesting that investors require additional compensation for firms more exposed to fiscal instability. Next, using pre-disaster market betas, we find that risk-adjusted disaster drawdowns are larger for high-FCE_f firms. High-FCE_f compared to low-FCE_f firms further experience an increase in market betas in the disaster state. These observations indicate that FCE_f-related variation in disaster drawdowns could be attributed to cross-sectional differences in the dynamics of market betas in the disaster. We provide direct evidence for this conjecture: FCE_f is largely unrelated to risk-adjusted drawdowns if we use realized disaster market betas for the risk adjustment.

Finally, we provide evidence on the likely channel driving the disaster results. Two candidates are demand effects and financing conditions. First, regarding demand effects, we indeed observe a greater drop in sales among high-FCE_f firms compared to low-FCE_f firms. In our preferred specification, we find a 6% drop in sales for a one-unit increase in FCE_f. Second, we split the sample into firms borrowing domestically versus foreign. Firms borrowing only domestically are by design not exposed to foreign lending. In this sub-sample, we thus compare firms headquartered and borrowing in the same country, but with differential exposure to fiscal constraints due to their sales distribution. We find similar disaster effects in the two sub-samples, suggesting that direct financing effects do not play a crucial role for our disaster results; rather, demand effects appear to be the key driver.

Literature. We find that exposure to fiscal constraints amplifies firms' disaster vulnerability. These firms are perceived as riskier even in normal times, apply higher discount rates in investment decisions, and invest less in capital and R&D. Since investment, particularly in R&D, is a critical driver of growth (e.g., Romer, 1990; Aghion and Howitt, 1992), our findings offer micro-level evidence that helps explain why fiscal constraints are associated with lower growth in the macro literature (Reinhart and Rogoff, 2010; Reinhart, Reinhart, and Rogoff, 2012). Additionally, fiscal constraints have been linked to more severe financial crises at the aggregate level (Romer and Romer, 2018, 2019). Our results provide new insights into the mechanisms underlying this relationship: we show that fiscal constraints amplify perceived risk from normal to disaster states and exacerbate adverse demand effects, which can contribute to deeper economic downturns.

Based on novel firm-level data, Gormsen and Huber (2023a,b) examine firms' perceived cost of capital and discount rates used by management in investment decisions. Notably, Gormsen and Huber (2023a) document a wedge between discount rates and the perceived cost of capital. Our analysis shows that exposure to fiscal constraints can help explain these wedges. This aligns with a risk-based interpretation, suggesting that managers adjust their discount rates upward to account for heightened disaster vulnerability. Other studies examine discount rates using survey data (e.g., Summers, 1986; Poterba and Summers, 1995; Graham and Harvey, 2001; Meier and Tarhan, 2007; Jacobs and Shivdasani, 2012; Mukhlynina and Nyborg, 2016; Jagannathan et al., 2016; Sharpe and Suarez, 2021; Graham, 2022; Barry et al., 2023), as well as the link between the cost of capital and investment (e.g., Lamont, 2000; Krüger, Landier, and Thesmar, 2015; Van Binsbergen and Opp, 2019; Pflueger, Siriwardane, and Sunderam, 2020; Dessaint et al., 2021; Kim, 2022; He, Liao, and Wang, 2022). Exploring the role of fiscal constraints in influencing discount rates and investment distinguishes our work from these previous studies.

A central paradigm in the literature on sovereign risk is the 'doom loop' (e.g., Acharya, Drechsler, and Schnabl, 2014; Gennaioli, Martin, and Rossi, 2014; Brunnermeier et al., 2016; Arellano, Bai, and Bocola, 2017; Farhi and Tirole, 2018): banks are exposed to the financial health of the government and transmit shocks to firms by reducing the supply of credit. Andrade and Chhaochharia (2018) use stock prices in the European sovereign debt crisis to provide a market-based estimate of the cost of sovereign default.[2] Hébert and Schreger (2017) estimate the causal effect of sovereign default on asset prices by exploiting changes in legal rulings in Argentina. Jappelli, Pelizzon, and Plazzi (2021) use corporate CDS spreads and find that the covid pandemic triggered a surge in the sovereign-corporate nexus in European core countries. We innovate on three dimensions relative to this literature. First, rather than relying on the fiscal constraint of the headquarter country, we concentrate on fiscal constraints exposure derived from firms' sales distributions. This approach establishes a direct connection to customer demand effects and is arguably less susceptible to concerns related to reverse causality (e.g., low firm productivity giving rise to fiscal constraints) or omitted variable bias (e.g., cultural and structural factors influencing both fiscal constraints and firm productivity). Second, we measure perceived risk—related and unrelated to fiscal constraints—directly, allowing us to address alternative stories. Third, we investigate how fiscal constraints contribute to a risk-based disaster-vulnerability channel, impacting both firms' discount rates and their investment decisions.

A growing literature also studies the role of fiscal policies in asset prices. These papers focus on the calibration of macro-finance models in the U.S. context, do not speak to disaster-based channels, and do not measure perceived risk directly. Fiscal policy matters for aggregate stock returns (e.g., Croce et al., 2012; Croce, Nguyen, and Schmid, 2012; Pastor and Veronesi, 2012; Belo and Yu, 2013; Croce, Nguyen, and Raymond, 2021; Liu, 2022), for industry returns (Belo, Gala, and Li, 2013), and for returns of R&D intensive

[2] Corsetti et al. (2013), Bedendo and Colla (2015), and Augustin et al. (2018) also focus on the spillover effects of sovereign to corporate credit risk among Eurozone countries. De Marco (2019) provides evidence for real effects of the sovereign-corporate nexus through a bank lending channel.

firms (Croce et al., 2019). Other studies examine the dynamics of sovereign credit default swap (CDS) spreads, which can be driven by global (Pan and Singleton, 2008; Longstaff et al., 2011; Augustin and Tédongap, 2016) and local risk factors (Hilscher and Nosbusch, 2010; Augustin, 2018). Jiang et al. (2023) take an asset pricing view on fiscal capacity, arguing that governments tend to insure taxpayers against aggregate risk. This insight is in line with the idea that insurance provision in bad times is valuable. Augustin et al. (2022) use a ranking-based measure of fiscal constraints and find that sovereign CDS spreads rise more among more constrained countries, following the covid outbreak. We differ from prior studies by introducing firm-level measures of exposure to fiscal constraints and perceived firm-level fiscal constraints risk. Combined with a clear exogenous shock we can thus measure the cross-sectional dynamics of perceived risk associated with exposure to fiscal constraints from normal times to the disaster state.

Our paper also relates to contributions regarding the impact of rare disasters on asset prices. Rietz (1988) introduces a rare disaster state in a consumption-based asset pricing model to address the equity premium puzzle posed by Mehra and Prescott (1985). Barro (2006) extends the concept of rare disasters and demonstrates that calibrated disaster probabilities can explain various puzzles in asset pricing. Other studies, such as Longstaff and Piazzesi (2004), Gabaix (2012), Gourio (2012), Wachter (2013), and Tsai and Wachter (2015), explore time-variation in disaster-related risk and consider a cross-section of assets with varying degrees of resilience. Our empirical analysis suggests that exposure to fiscal constraints plays an important role for disaster resilience.

By focusing on firm-level exposure to fiscal constraints, we further advance the understanding of effects in asset prices during covid (Gormsen and Koijen, 2023, provide an overview). Ramelli and Wagner (2020) study exposure to China, Fahlenbrach, Rageth, and Stulz (2021) focus on financial flexibility, Mazur, Dang, and Vega (2021) study industry effects, and Ding et al. (2020) examine the role of corporate governance. Barrot et al. (2023)

study business closure decisions, and Bretscher et al. (2020) covid transmission channels. Arteaga-Garavito et al. (2024) study high frequency news shocks, and Pagano, Wagner, and Zechner (2023) how vulnerability to social distancing is priced in asset returns.

This article is structured as follows: Section 2 presents the hypotheses. Section 3 introduces our novel firm-level measure of exposure to fiscal constraints. Section 4 examines how fiscal constraints shape the impact of disasters on firms' perceived risk. Section 5 focuses on the effects of exposure to fiscal constraints on discount rates and investment in normal times, while Section 6 explores the mechanism. Section 7 concludes.

2 Hypotheses

We conjecture that firms exposed to fiscally constrained countries—through their sales distribution—are more vulnerable to global disasters and perceived as riskier. As a result, managers use higher discount rates in investment decisions in normal times, leading to lower investment in capital and R&D.

2.1 Fiscal constraints and disaster vulnerability

Theory suggests that exposure to fiscal constraints amplifies firms' vulnerability to global disasters through multiple channels. First, fiscal constraints can prevent governments from optimally employing fiscal policies to smooth out shocks. This can amplify downturns due to insufficient levels of support and increased uncertainty about fiscal policies aimed at i) aggregate demand (e.g., Keynes, 1936; Blinder and Solow, 1973; Auerbach and Gorodnichenko, 2012) or ii) firms' financing conditions (e.g., Kiyotaki and Moore, 1997; Holmstrom and Tirole, 1997; Bernanke and Gertler, 1989; Jermann and Quadrini, 2012).

Second, fiscally constrained countries may experience a deterioration of credit worthiness when disaster strikes, which can affect firms directly through their financing conditions or indirectly through demand effects. For instance, a decline in the value of government securities in fiscally constrained countries—due to deteriorating government creditworthiness—can lead to a 'doom-loop' (e.g., Acharya, Drechsler, and Schnabl, 2014; Gennaioli, Martin, and Rossi, 2014; Brunnermeier et al., 2016; Arellano, Bai, and Bocola, 2017; Farhi and Tirole, 2018): banks that hold government securities may cut back on lending to firms and consumers when the value of these securities declines. Additionally, increased sovereign risk premia in fiscally constrained countries can raise firms' borrowing costs by serving as a benchmark for corporate credit (Almeida and Argelina, 2009).

All these channels predict that firms exposed to fiscal constraints are more vulnerable to disasters due to a stronger and more uncertain impact on demand and credit conditions. Consequently, a general implication is that exposure to fiscal constraints amplifies the increase in risk from normal times to a disaster state.

H1: *Exposure to fiscal constraints amplifies the increase in (perceived) risk when disaster strikes.*

2.2 Discount rates and investment in normal times

Firms exposed to fiscal constraints may be perceived as riskier even in normal times due to their increased vulnerability to global shocks. Following basic corporate finance principles, managers should incorporate this higher perceived risk into the discount rates used in investment decisions during normal times (e.g., Modigliani and Miller, 1958).

H2: *Managers of firms exposed to fiscal constraints set higher discount rates in investment decisions in normal times.*

Finally, higher discount rates should lead to lower investment in physical capital (e.g., Jorgenson, 1963; Tobin, 1969; Hayashi, 1982) and R&D (e.g., Griliches, 1981; Hall, 1993).

H3: *Firms exposed to fiscal constraints invest less in physical capital and R&D in normal times.*

The main focus of the paper is to test the hypotheses outlined above and establish a link between fiscal constraints and disaster vulnerability, as well as between fiscal constraints, discount rates, and investment during normal times. In Section 6, we conduct additional tests to explore the potential mechanisms at play. To this end, we examine the behavior of stock returns both during and outside disaster periods, as well as changes in sales during the transition from normal times to the disaster state.

3 Data

In this section, we describe the data and introduce our novel firm-level measure of exposure to fiscal constraints.

3.1 Country-level data

We use the World Bank's fiscal space database (described in Kose et al., 2022) to obtain the following country-level data: debt-to-GDP ratio (sovereign debt, *sd*), foreign currency long-term debt rating (*rat*), interest payments in percent of GDP (*int*), primary fiscal balance in percent of GDP (*prim.bal*),[3] GDP per capita (*gdppc*), five-year GDP growth (*gdpg*), and government expenditure in percent of GDP (*exp*). We further obtain the history of covid cases from the data repository of John Hopkins CSSE.

[3]Primary fiscal balance is the difference between a government's revenues and its expenditures (excluding interest expenses).

3.2 Firm-level data

We obtain accounting data from the Compustat Capital IQ Global and North America databases. We exclude financials (GICS sector 40), and only keep firms with positive total assets and positive market capitalization. We require the availability of leverage $lev = (dltt + dlc)/at$; $cash = che/at$; and property, plant, and equipment $ppe = ppent/at$; as well as of market beta, β; market value of equity in USD, me; book-to-market ratio, bm; operating profitability, op; investment-to-assets, i/a; and return on equity, roe (i.e., the Fama and French (2015, FF) and Hou, Xue, and Zhang (2015, HXZ) characteristics).[4] We estimate β by regressing a stock's returns on the global market returns over the preceding 60 months, requiring a minimum of 24 monthly returns. We winsorize firm-level data at the 1% and 99% level.

We collect transcripts of earnings conference calls from Refinitiv from Q1 2015 to Q4 2020, and filter out those with fewer than 100 bigrams. Next, we use the data of Gormsen and Huber (2023a,b) on firms' discount rates, wedges, and perceived cost of capital. Stock price data and currency data are obtained from the Compustat Capital IQ Global and North America databases. We closely follow Bessembinder et al. (2019) in constructing returns. We only use common stocks (*tpci* either 0, Q, or 1) and traded assets (*exchg* either 0, 1, 13, 19, or 20), and consider only observations with a certain price status (*prcstd* either 3, 4, 5, or 10). If a stock is cross-listed, we use the price from the exchange with the highest trading volume in USD. We eliminate stock-years lacking a price observation on more than 5% of the exchange's trading days. Consistent with existing literature (e.g., Frazzini and Pedersen, 2014; Bessembinder et al., 2019), we work with stock returns based on USD prices. However, for robustness, we provide results using domestic currency returns in the Internet Appendix. The global market return is calculated as the value-weighted USD

[4] FF characteristics are β, me, bm, op, and i/a, while those of HXZ are β, me, roe, and i/a. For construction, we follow the detailed descriptions in Hou et al. (2019) and Hou, Xue, and Zhang (2020).

return across all stocks in the Compustat Capital IQ Global and North America universe.

3.3 Firm-level measure of fiscal constraint exposure

We aim to capture the multifaceted nature of countries' fiscal constraints. To do so, we build for each country an index of fiscal constraints that combines four variables (debt-to-GDP ratio, credit rating, interest expense, and primary balance) as follows:

$$FC_{c,t} = \frac{1}{4}\left(sd_{c,t} + rat_{c,t} + int_{c,t} + n.prim.bal_{c,t}\right), \tag{1}$$

where c is country and t is year.[5] The debt-to-GDP ratio is a widely used measure of a country's debt burden relative to output, while credit ratings are independent assessments of a country's creditworthiness that consider a wide range of economic and fiscal factors. Interest expense and primary balance are flow variables, with the former capturing the cost of servicing debt and the latter the current path of a government's debt. There are 21 rating notches and a higher value indicates higher credit risk. The variable $n.prim.bal$ is simply $prim.bal$ multiplied by minus one, so that deficits have positive values. We standardize the four variables (sd, rat, int, and $n.prim.bal$) in the cross-section (i.e., they all have mean zero and standard deviation one). A higher value of FC indicates more fiscal constraints.[6]

In Table IA.A.I in the Internet Appendix, we show that this measure of fiscal constraints indeed predicts lower government spending during disasters. Specifically, we demonstrate that countries entering the global financial crisis and the covid pandemic with higher values of FC spend less during these disasters. In other words, FC effectively captures the limited insurance provision by governments in times of disaster.

Next, we use Factset GeoRev to extract a firm's sales distribution across countries,

[5] An important advantage of this measure, compared to market-based alternatives such as sovereign credit default swaps, is that it can be constructed for a much larger cross-section of countries.

[6] The fiscal constraint measure is in the spirit of Augustin et al. (2022). However, we do not use rankings but construct an index from standardized variables to capture the intensity of fiscal constraints.

forming the basis for our novel firm-level measure of exposure to fiscal constraints.[7] We only use firm-year observations where a firm's sales distribution captures at least 90% of total sales, and we re-scale the shares so that they sum to one.[8] We define $w_{f\times c,t}$ as the fraction of sales of firm f in country c in fiscal year t relative to the firm's total sales across all countries in year t. We then compute firm-level exposure to fiscal constraints as:

$$FCE_{f,t} = \sum_{c} w_{f\times c,t} \cdot FC_{c,t}. \tag{2}$$

We also construct additional sales-weighted variables: $gdppc_{f,t}$, $gdpg_{f,t}$, $cases_{f,t}$, and we denote $w^{for}_{f,t}$ as a firm's share of foreign sales (i.e., sales outside the headquarter country).

4 Amplification of perceived risk in disaster

In this section, we first present our empirical design to identify the differential disaster response of firms exposed to fiscal constraints. We then apply this framework to study the role of exposure to fiscal constraints in shaping perceived firm-level risk during a disaster. We establish a link between our exposure measure, FCE_f, and perceived firm-level risk related to fiscal constraints, which we obtain by applying textual-analysis tools to earnings conference calls. This analysis allows us to demonstrate that FCE_f amplifies the increase in perceived firm-level risk—specifically related to fiscal constraints—when transitioning from normal conditions to a disaster state. Additionally, we show that FCE_f does not affect perceived firm-level risk unrelated to fiscal constraints during disasters, addressing concerns that fiscal constraints exposure may correlate with other factors influencing the

[7] Li, Richardson, and Tuna (2014) also use geographical sales distributions data to construct firm-level measures of macroeconomic forecasts. Their objective is very different from ours as they focus on the role of macroeconomic forecasts for predicting firm performance. Moreover, we study effects within headquarter countries and we rely on sales data by country rather than by geographic region.

[8] The raw sales distribution data covers typically at least 99% of a firm's total sales; hence our filter excludes only very few firms.

risk dynamics in the disaster state. In Section 6 we examine the relation between FCE_f and other disaster effects to shed more light on the mechanism at play.

4.1 Empirical design

The covid pandemic presents a unique and close-to-ideal setting to examine how fiscal constraints amplify the impact of disasters on firms. Our analysis leverages the pandemic as a large, exogenous global negative shock, allowing us to assess firms' differential responses based on their ex-ante exposure to fiscal constraints.

4.1.1 Covid pandemic as a natural experiment

The covid pandemic offers several key advantages as a laboratory for testing the hypothesis that fiscal constraints amplify firms' disaster vulnerability. First, the pandemic was unforeseen and sudden, with substantial impacts on asset prices and growth expectations (Gormsen and Koijen, 2020; Landier and Thesmar, 2020). Figure IA.A.2 in the Internet Appendix shows that stock prices across different regions of the world began to plummet simultaneously around February 21, 2020, as investors started to realize the potential global economic impact of the pandemic. The unexpected nature and clear timing of the pandemic's onset allow us to assess the plausibility of the parallel trends assumption around the outbreak, which is crucial for our difference-in-differences (DiD) analysis. Second, the pandemic was exogenous to countries' fiscal constraints and firms' financial conditions. As a global health crisis, the outbreak of the pandemic was unrelated to countries' fiscal or economic policies and the stability of the financial system. This exogeneity ensures that the shock itself is not influenced by pre-existing fiscal constraints, allowing us to isolate the amplifying role of those constraints. Third, the global nature of the pandemic allows to analyze a broad cross-section of firms across various countries.

To further illustrate the econometric appeal of the pandemic for our purposes, consider two alternative disasters: the global financial crisis (GFC) and the sovereign debt crisis. Unlike the pandemic, the GFC was deeply rooted in the financial systems and economic policies of specific countries. This endogeneity complicates any attempt to isolate the amplifying effect of fiscal constraints on firms during the GFC. Furthermore, the GFC's impact was uneven across countries, with regions differentially affected at varying points in time. Similarly, the sovereign debt crisis was closely tied to the fiscal health and economic conditions of specific countries. Moreover, it had a regional impact centered on countries with pre-existing fiscal vulnerabilities, and its timing varied across regions. The endogeneity and variability in exposure to both the GFC and the sovereign debt crisis complicate comparisons across firms, making it difficult to apply a consistent DiD framework.

4.1.2 Firm-level exposure to fiscal constraints

To address the challenge of identifying differential impacts across firms, we combine the exogenous outbreak of the pandemic with a novel firm-level measure of fiscal constraint exposure based on the distribution of firms' sales across countries. This approach offers several advantages over relying solely on the firm's headquarter country for exposure.

Most notably, using firm-level exposure allows to control for headquarter country fixed effects, thereby isolating the impact of fiscal constraints stemming from international sales rather than the domestic fiscal environment. This enables a more precise comparison between firms with similar domestic conditions but different exposure to fiscal constraints due to their sales distribution across countries. Specifically, using our firm-level exposure measure in conjunction with headquarter fixed effects reduces concerns over reverse causality, such as low firm productivity influencing the fiscal constraints of the headquarter country. Furthermore, fiscal constraints in a firm's headquarter country might be shaped by various factors—including cultural and structural aspects—that may also impact firm characteristics such as productivity. By using our sales-weighted exposure combined with headquarter fixed effects, we mitigate the risk of omitted variable bias arising from factors that influence both firm performance and national fiscal policies.

4.1.3 Difference-in-differences specifications

To estimate the impact of fiscal constraints on firms' sensitivity to disasters, we leverage the exogenous outbreak of the pandemic in combination with our ex-ante firm-level fiscal constraint exposure in a difference-in-differences (DiD) framework. By comparing firms with varying levels of fiscal constraint exposure before and after the outbreak, we aim to isolate the amplifying role of these constraints in shaping the disaster's impact on firms' perceived riskiness and performance.

Importantly, the DiD approach does not require does not require fiscal constraint exposure to be randomized across firms. Instead, the key identifying assumption is that, absent the disaster, firms with differing levels of fiscal constraint exposure would have followed similar trends over the considered timeframe.[9] Three pieces of evidence support this parallel trends assumption. First, firms with different levels of fiscal constraint exposure do not differ in ex-ante characteristics that might influence their disaster response. Table IA.A.II in the Internet Appendix shows that our measure of fiscal constraint exposure, $FCE_{f,2018}$, is uncorrelated with firms' ex-ante size, market-to-book ratio, profitability, leverage, cash holdings, and asset tangibility. Second, our key outcome of interest during the disaster—perceived fiscal constraints risk (as measured from earnings calls)—exhibits parallel trends prior to the pandemic and begins to diverge only with its onset. Third, we observe that risk unrelated to fiscal constraints continues to follow parallel trends even after the onset of the pandemic.

[9]Here, our objective is to identify the amplification effect in firms' responses to the disaster. Consequently, our approach relies on the assumption of parallel trends in the outcomes of interest surrounding the outbreak of the pandemic, while allowing for long-term level differences in perceived risk, expected returns, and investment rates—differences predicted by our disaster-vulnerability channel and studied in Section 5.

Given this evidence supporting the parallel trends assumption, we proceed to estimate the following DiD specification with continuous exposure:[10]

$$\Delta Y_{f,\mathcal{D}-\mathcal{N}} = \delta FCE_{f,2018} + \gamma_1' C_{f,2018} + \gamma_2' F_{f,2018} + FE + \epsilon_{f,\mathcal{D}-\mathcal{N}}. \tag{3}$$

$\Delta Y_{f,\mathcal{D}-\mathcal{N}}$ represents the change in firm outcomes around the outbreak of the pandemic and δ captures the impact of fiscal constraint exposure, $FCE_{f,2018}$.[11] To mitigate concerns about omitted variables (e.g., cultural and structural factors influencing both firm productivity and fiscal policy) and reverse causality (e.g., low firm productivity contributing to fiscal constraints), we employ combinations of headquarter-country fixed effects: we either use FE =HQ+Ind or FE =HQ×Ind, where Ind is based on the two-digit Global Industry Classification Standard (GICS). We add controls for other sales-weighted country characteristics: we control for the logarithm of GDP per capita in USD, and five-year GDP growth to account for differences in ex-ante economic development. This allows us to compare firms that sell to countries with the same level of development but different degrees of fiscal constraints. Additionally, we control for the logarithm of sales-weighted covid cases (measured at the end of the disaster period $\mathcal{D}$). Hence, we have that $C_f := [\log(gdppc_f), gdpg_f, \log(1 + cases_f)]$.

We add controls for the share of foreign sales (w_f^{for}), as well as debt, cash, and property, plant, and equipment (i.e., F_f contains $Other_f := [w_f^{for}, lev_f, cash_f, ppe_f]$). In addition, F_f either includes the HXZ or FF firm characteristics. These characteristics capture differences

[10]Note that this change-in-outcomes specification is akin to an interaction model in a panel setup:

$$Y_{f,t} = \delta Disaster_t \times FCE_{f,2018} + Firm\ FE + Time\ FE + \epsilon_{f,t},$$

where $Disaster_t$ is a dummy equal to one in the disaster state. We prefer the change-in-outcomes specification for simplicity of exposition, but results are robust to using the panel specification instead.

[11]We define the disaster ($\mathcal{D}$) and pre-disaster ($\mathcal{N}$) periods based on the outcome of interest. We compute our risk measures using data from the quarters immediately before and after the outbreak in Q1 2020; for stock returns, we use a window of 20 trading days around February 21, 2020; and for sales, we define fiscal years 2018 and 2019 as the pre-disaster period, and 2020 and 2021 as the disaster period.

in market risk exposures, size, growth options, or profitability. We measure $FCE_{f,2018}$ and controls (with the exception of covid cases) before the pandemic to i) reflect the information available to market participants at the time of the outbreak, and ii) avoid concerns about endogenous controls.[12] Contemporaneous values of the control variables (i.e., measured during the disaster period) are endogenous and thus considered bad controls' (Angrist and Pischke, 2009).[13] Standard errors are clustered at the industry level (six-digit GICS).

4.2 Perceived firm-level fiscal constraints risk

We begin our empirical analysis by examining the dynamics of perceived risk around the outbreak of the pandemic. First, we demonstrate that our measure of ex-ante exposure to fiscal constraints predicts an increase in perceived risk related to fiscal constraints when disaster strikes. Next, we show that our exposure measure does not predict the dynamics of risk unrelated to fiscal constraints during the disaster.

To measure firm-level risk related to fiscal constraints, we apply computational linguistics tools to earnings conference calls, following the methodology of Hassan et al. (2019). Our goal is to quantify the portion of the conversation between firm management and call participants that focuses on fiscal constraints-related risk.

To distinguish fiscal constraints-related topics from non-fiscal constraints topics, we apply a pattern-based sequence-classification method from computational linguistics (Song and Brook Wu, 2008; Schutze, Manning, and Raghavan, 2008). Specifically, we search for language patterns in earnings calls associated with fiscal constraints and identify their connection with risk using synonyms of 'risk' and 'uncertainty.'

To be precise, we construct our measure of perceived firm-level fiscal constraints risk, $FCR_{f,t}$, by defining two training libraries: one for fiscal constraints-related text, $\mathbb{F}$, and

[12] Specifically, accounting and country-level data are measured as of the end of 2018, while market data (e.g., β) are measured as of the end of 2019.

[13] The results are robust to excluding covid cases from the list of controls.

another for non-fiscal constraints-related text, $\mathbb{N}$. We utilize the textbook by Abbas, Pienkowski, and Rogoff (2019) for fiscal constraints topics and the accounting textbook[14] by Libby, Libby, and Hodge (2019), along with the Santa Barbara Corpus of Spoken American English[15], for non-fiscal constraints topics. We choose the Abbas, Pienkowski, and Rogoff (2019) textbook for its exhaustive treatment of topics related to fiscal constraints, specifically targeting practitioners. This focus aligns better with the language used in earnings calls, compared to textbooks aimed exclusively at an academic audience.

Each training library consists of all bigrams (two-word sequences) contained in $\mathbb{F}$ and $\mathbb{N}$, respectively. In Panel A of Table 1, we present the 50 most common bigrams found in the fiscal constraints training library but not in the non-fiscal constraints training library, denoted as $\mathbb{F} \setminus \mathbb{N}$. Examples include 'public debt,' 'sovereign debt,' 'government debt,' 'fiscal rules,' and 'fiscal adjustment.'

Table 1 about here

We then decompose each transcript of firm f at time t into $B_{f,t}$ bigrams (indexed by b) and count the occurrences of bigrams that indicate a discussion of fiscal constraints, within ten words of synonyms for 'risk' or 'uncertainty.' This count is then scaled by the total number of bigrams in the transcript to obtain the perceived firm-level fiscal constraints risk, calculated as follows:

$$FCR_{f,t} = \frac{1}{B_{f,t}} \sum_{b}^{B_{f,t}} \mathbb{I}\left(b \in \mathbb{F} \setminus \mathbb{N}\right) \times \mathbb{I}\left(|b - r| < 10\right) \times \frac{f_{b,\mathbb{F}}}{B_{\mathbb{F}}}, \tag{4}$$

where $\mathbb{I}$ denotes the indicator function, $\mathbb{F} \setminus \mathbb{N}$ represents the set of bigrams in $\mathbb{F}$ but not in $\mathbb{N}$, and r indicates the position of the nearest synonym for 'risk' or 'uncertainty.' We thus count the number of bigrams related to fiscal constraints (but not non-fiscal constraints)

[14] As discussed by Hassan et al. (2019), earnings calls often focus on financial disclosures and accounting information, supporting the use of an accounting textbook.

[15] Accessible at `https://www.linguistics.ucsb.edu/research/santa-barbara-corpus`.

that occur within ten words of a synonym for 'risk' or 'uncertainty,' and then scale this count by the total number of bigrams in the transcript. Each bigram is weighted based on its association with fiscal constraints topics, where $f_{b,\mathbb{F}}$ denotes the frequency of bigram b in the fiscal constraints training library, and $B_{\mathbb{F}}$ is the total number of bigrams in the training library. Examples of discussions about fiscal constraints risk in conference calls are provided in Panel B of Table 1.

In the first step, we focus on the dynamics of $FCR_{f,t}$—the risk associated with fiscal constraints as perceived by managers and call participants—around the outbreak of the pandemic. We compute a quarterly $FCR_{f,t}$ for each firm, based on the average of all conference calls in a two-quarters rolling window. We match the conference call data to Compustat, resulting in a sample of 3,043 firms from 55 countries. Panel A of Figure 1 in the Introduction illustrates the dynamics of $FCR_{f,t}$ from Q1 2019 to Q4 2020 as a function of exposure to fiscal constraints, $FCE_{f,2018}$. In each quarter $t \in$ [Q1 2019,...,Q4 2020], we run a cross-sectional regression at the individual firm level, regressing the change in perceived fiscal constraints-related risk, $\Delta FCR_{f,t}$, on ex-ante firm-level exposure to fiscal constraints, $FCE_{f,2018}$, while absorbing headquarter (HQ) country fixed effects.

As the dependent variable, we use the change in perceived firm-level fiscal constraints risk from Q4 2019 to quarter t, defined as $\Delta FCR_{f,t} = (FCR_{f,t} - FCR_{f,\text{Q4 2019}})/\overline{FCR}_{\text{Q4 2019}}$, where $\overline{FCR}_{\text{Q4 2019}}$ is the median in Q4 2019. For each quarter, we plot the coefficient of $FCE_{f,2018}$ (black line) along with 95% confidence intervals (grey lines). High-FCE_f firms experience a significant increase in perceived fiscal constraints risk compared to low-FCE_f firms when the disaster strikes, with the largest differential observed in Q3 2020. Importantly, there is no discernible difference in trends between high- and low-FCE_f firms before the disaster.

Table 2 about here

Next, we estimate equation (3) and add controls for firm and sales-weighted country characteristics, along with combinations of headquarter country and industry fixed effects. We use $\Delta FCR_{f,\text{Q3 2020}-\text{Q4 2019}} = (FCR_{f,\text{Q3 2020}} - FCR_{f,\text{Q4 2019}})/\overline{FCR}_{\text{Q4 2019}}$ as the dependent variable, that is, the change in perceived firm-level fiscal constraints risk from Q4 2019 (pre-disaster state) to Q3 2020 (disaster state), scaled by $\overline{FCR}_{\text{Q4 2019}}$, the median in Q4 2019. Summary statistics for the sample are provided in Table 2. On average, ΔFCR_f is 1.31, demonstrating that perceived firm-level fiscal constraints risk increases by 131% compared to pre-disaster levels.

Table 3 about here

We show results of equation (3) in Table 3, where we present four model specifications. In model 1 we use HQ+Ind fixed effects and control for sales-weighted country characteristics, C_f, as well as the baseline firm controls $Other_f$ (i.e., w^{for}, lev, $cash$, ppe). In model 2 we additionally use the HXZ characteristics (i.e., β, me, roe, i/a), while in model 3 we use the FF characteristics (i.e., β, me, bm, op, i/a), and model 4 then employs HQ×Ind fixed effects. We find positive and significant δ-coefficients across all specifications with $|t| > 2.0$, supporting **H1**. In our most stringent specification in column 4 we obtain a δ-estimate of 0.82 ($|t| > 2.3$). Economically, this result implies that a one-unit increase in ex-ante fiscal constraint exposure leads to an 82 percentage point increase in perceived firm-level fiscal constraints risk during the disaster, relative to the pre-disaster period.

These results show that differential exposure to fiscal constraints (FCE_f)—among otherwise identical firms headquartered in the same country and operating in the same industry—predicts increased risk related to fiscal constraints during a disaster.

4.3 Falsification and placebo tests

Next, we conduct a series of tests to provide additional support for the hypothesis that exposure to fiscal constraints amplifies disaster vulnerability. First, we demonstrate that the observed increase in perceived risk for high- relative to low-FCE firms is specific to fiscal constraints-related risk, with no corresponding rise in i) perceived risk unrelated to fiscal constraints, ii) perceived covid risk, or iii) perceived general country risk. Additionally, we show that other ex-ante country risk measures, such as the standard deviation of GDP growth and exchange rates, do not predict an increase in perceived fiscal constraints risk. These results help mitigate concerns that other factors correlated with fiscal constraints might confound the dynamics of firm-level risk during the disaster period.

We begin with a falsification test, similar to that in Hassan et al. (2019). Specifically, we create a measure of perceived firm-level risk unrelated to fiscal constraints, $NFCR_f$. This measure quantifies the share of the conversation centering around risk of non-fiscal constraints-related topics. To do so, we count $\mathbb{N} \setminus \mathbb{F}$ (instead of $\mathbb{F} \setminus \mathbb{N}$) and weight by $\frac{f_{b,\mathbb{N}}}{B_{\mathbb{N}}}$ (instead of $\frac{f_{b,\mathbb{F}}}{B_{\mathbb{F}}}$) in equation (4). In Panel B of Figure 1 in the Introduction we now use as the dependent variable the change in perceived firm-level non-fiscal constraints-related risk from Q4 2019 to quarter t, $\Delta NFCR_{f,t} = (NFCR_{f,t} - NFCR_{f,\text{Q4 2019}})/\overline{NFCR}_{\text{Q4 2019}}$, where $\overline{NFCR}_{\text{Q4 2019}}$ is the median in Q4 2019. We plot for each quarter the coefficient of $FCE_{f,2018}$ (black line) along with 95% confidence intervals (grey lines). There is no discernible difference in the dynamics of risk unrelated to fiscal constraints between high- and low-FCE_f firms in any quarter of 2019 and 2020. In other words, our exposure measure captures the dynamics of fiscal constraints-related risk, but is unrelated to risk that is not associated with fiscal constraints matters. This result further supports the parallel trends assumption: risk unrelated to fiscal constraints follows similar trends before the disaster and remains on parallel trends after the onset of disaster.

Similar to before, we then use the change in $NFCR_{f,t}$ from the pre- to the disaster period, $\Delta NFCR_{f,\text{Q3 2020}-\text{Q4 2019}} = (NFCR_{f,\text{Q3 2020}} - NFCR_{f,\text{Q4 2019}})/\overline{NFCR}_{\text{Q4 2019}}$, as the dependent variable in regression (3). The regression results in Table 4 reveal insignificant δ-coefficients across all specifications, with $|t| < 1.2$. In essence, the dynamics of risk associated with fiscal constraints matters are captured by FCE_f, while the dynamics of risk unrelated to fiscal constraints are not. This finding mitigates the concern that FCE_f captures other types of risk— in particular general country risk—as such risk would likely affect all forms of perceived risk rather than remaining specific to fiscal constraints.

Table 4 about here

Next, we conduct additional tests to address this concern. First, we examine the possibility that FCE_f correlates with perceived firm-level covid risk. To investigate this, we utilize the covid risk measure from Hassan et al. (2023a), which is a firm-level text-based measure of perceived covid risk. We run equation (3) with covid risk in Q3 2020 as the dependent variable and present the results in Table IA.A.III in the Internet Appendix. Reassuringly, we obtain negative and insignificant δ-estimates throughout all specifications with $|t| < 1.3$, showing that exposure to fiscal constraints is not related to covid risk.

Second, we further examine the possibility that FCE_f captures broader country risk dynamics rather than being specific to fiscal constraints-related risk. To address this concern, we conduct two different sets of tests. In the first exercise, we perform a further falsification test. Specifically, we use the quarterly text-based measures from Hassan et al. (2023b) on country risk. We construct a firm-level measure of perceived country risk, CR_f, by utilizing firms' sales distribution across countries from 2018, following the same approach as for FCE_f.

We then employ $\Delta CR_{f,\text{Q3 2020}-\text{Q4 2019}} = (CR_{f,\text{Q3 2020}} - CR_{f,\text{Q4 2019}})/\overline{CR}_{\text{Q4 2019}}$ as the dependent variable in equation (3), where $\overline{CR}_{\text{Q4 2019}}$ is the median in Q4 2019. The results are presented in Table IA.A.IV in the Internet Appendix, where we consistently observe negative and insignificant δ-estimates across all models, with $|t| < 1.2$. This result indicates that FCE_f does not capture the disaster dynamics of perceived country risk more broadly.

The second exercise is a placebo test. In this test, we run equation (3) with the change in perceived firm-level fiscal constraints risk, $\Delta FCR_{f,\text{Q3 2020}-\text{Q4 2019}}$, as the dependent variable. However, we replace our measure of exposure to fiscal constraints with proxies for ex-ante country risk more broadly. Specifically, in Panel A of Table IA.A.V in the Internet Appendix, we use the standard deviation of monthly currency returns (relative to USD) over the 15 years before the disaster. The coefficient δ is negative and insignificant across all models, with $|t| < 1.1$. In Panel B, we instead use the standard deviation of yearly GDP growth over the last 15 years. Here as well, we obtain negative and insignificant δ-estimates throughout all models, with $|t| < 0.8$. These analyses provide further support for our exposure measure: fiscal constraints risk increases for firms that are more exposed to fiscal-constraints ex-ante (i.e., high FCE_f-firms), but not for firms more exposed to broader measures of country risk.

Taken together, our measure of exposure to fiscal constraints predicts the disaster dynamics of perceived firm-level risk related to fiscal constraints, while it is uncorrelated with risk unrelated to fiscal constraints. Additionally, broader proxies for country risk are unrelated to fiscal constraints-related risk during the disaster. In other words, FCE_f effectively captures exposure to fiscal constraints-related concerns but is unrelated to non-fiscal constraints-related channels. Having established that exposure to fiscal constraints amplifies perceived risk in disaster, we now analyze its relationship with perceived fiscal constraints risk, discount rates, and investment during normal times.

5 Perceived risk, discount rates, and investment in normal times

Using perceived firm-level fiscal constraints risk, our previous tests show that high-FCE_f firms become riskier in a disaster compared to low-FCE_f firms. We now focus on the pre-disaster period, specifically examining the five years preceding the disaster, from 2015 to 2019.[16] We begin by studying perceived firm-level fiscal constraints risk and then turn to discount rates. Our objective is to determine whether higher levels of FCE_f are associated with heightened perceived fiscal constraints risk and higher required rates of return on investment during the pre-disaster period. In the final part of this section, we examine investment in physical capital and R&D. We find that high-FCE_f firms exhibit higher perceived fiscal constraints risk, apply higher discount rates in investment decisions, and invest less.

5.1 Pre-disaster perceived fiscal constraints risk

We now explore the relationship between fiscal constraints and perceived risk in the pre-disaster period. High-FCE_f firms exhibit an increase in perceived risk (FCR_f) relative to low-FCE_f firms in the disaster period. Next, we study if investors and analysts are concerned about fiscal constraints-related risk for high-FCE firms also in normal times. We find robust empirical support for this conjecture: Fiscal constraint exposure is associated with higher levels of perceived fiscal constraints-related risk in the pre-disaster period.

We analyze earnings conference calls from 2015 to 2019, retaining only those with

[16] We choose this time period since it follows the European sovereign debt crisis, which raised awareness of the risk associated with fiscal constraints. Figure IA.A.1 shows a spike in fiscal constraints-related news articles during the period from 2010 to 2013, with levels remaining elevated afterward. One interpretation of this fact could be that investors neglected the risk associated with fiscal constraints before (see, e.g., Bordalo, Gennaioli, and Shleifer, 2022). Moreover, the number of firms covered in the sales distribution data is stable as of 2013, whereas coverage before then is significantly lower, with the number of firms increasing by 360% (50%) from 2008 (2012) to 2013.

at least 100 bigrams. Using the methodology described in Section 4.2, we compute a measure of perceived fiscal constraints risk for each call, as defined in equation (4). We then calculate the yearly average for each firm to obtain an annual measure of firm-level perceived fiscal constraints-related risk, $FCR_{f,t}$, and match it to Compustat data. We include only firms with conference calls in at least three years and apply the same filters as described in Section 3. Our final sample consists of 2,461 firms headquartered in 57 countries. Subsequently, we conduct the following panel regression:

$$FCR_{f,t} = \delta FCE_{f,t-1} + \gamma_1' C_{f,t-1} + \gamma_2' F_{f,t-1} + FE + \epsilon_{f,t}, \tag{5}$$

where $FCR_{f,t}$ is scaled by the sample median; $C_{f,t-1} := [\log(gdppc_{f,t-1}), gdpg_{f,t-1}]$ are sales-weighted country characteristics, and $F_{f,t-1}$ are other firm characteristics, $Other_{f,t-1} := [w^{for}_{f,t-1}, lev_{f,t-1}, cash_{f,t-1}, ppe_{f,t-1}]$, as well as the HXZ or FF characteristics. In our baseline specification we set FE=HQ×Year+Ind×Year. In our most stringent specification we set FE=HQ×Ind×Year, that is, we exploit variation across firms headquartered in the same country and operating in the same industry in the same year.

Table 5 about here

We present the results in Table 5, where we find positive and significant δ-estimates throughout all specifications, with $|t| > 1.9$. In our most stringent specification in model 4, where we control for sales-weighted country characteristics, other firm characteristics, and include HQ×Ind×Year fixed effects, we obtain a δ of 0.57. This estimate implies that a one-unit increase in exposure to fiscal constraints raises perceived fiscal constraints risk by 57 percentage points.

The analysis shows that our exposure measure captures variation in perceived fiscal constraints risk in normal times, that is, high-FCE_f firms exhibit higher levels of FCR_f

compared to low-FCE_f firms outside disasters. Collectively, the results suggest that fiscal constraints influence firms' perceived risk through a disaster vulnerability mechanism: high-FCE_f firms experience an increase in perceived risk in a disaster, and investors and analysts of high-FCE_f firms are concerned about fiscal constraints risk also in normal times. Importantly, these insights are based on comparing firms with identical characteristics that are headquartered in the same country and operate in the same industry.

5.2 Pre-disaster discount rates

As demonstrated above, firms with greater exposure to fiscal constraints exhibit higher pre-disaster perceived fiscal constraints risk. Next, we study what these findings imply for a firm's pre-disaster discount rates. To do so, we turn to data from Gormsen and Huber (2023a,b), who compile firm-level discount rates from earnings conference calls. The discount rate, denoted as $r^{dr}_{f,t}$ for firm f at time t, is the required return on investment set by firm management for investment decisions. Similar to the previous section, we examine the five years leading up to the disaster (fiscal years 2015 to 2019). Our analysis includes all firms available during this period, subject to the same criteria applied in our previous filters. This procedure results in a sample of 1,885 firms headquartered in 23 countries. We run the same panel regressions as for pre-disaster perceived fiscal constraints-related risk, shown in equation (5).

Results for discount rates (i.e., $r^{dr}_{f,t}$ as dependent variable) are in Table 6. We obtain significantly positive δ-estimates throughout with $|t| > 1.9$, supporting **H2**. Model 4 is our most stringent specification, where we use the full set of firm-level controls and employ HQ×Ind×Year fixed effects. In this model, δ is around 0.7 with $|t| > 2.0$, which implies that a one-unit increase in FCE_f raises the discount rate that managers use in their investment decisions by 0.7 percentage points. This is an economically important effect, given that we draw inference about δ by comparing firms with otherwise identical characteristics that are headquartered in the same country and operate in the same industry.

Table 6 about here

Our results show that exposure to fiscal constraints is associated with higher discount rates used by firm management in investment decisions in the pre-disaster period. This suggests that exposure to fiscal constraints affects investment of firms in normal times. In the next step, we therefore investigate whether fiscal constraints indeed correlate negatively with pre-disaster investment.

5.3 Pre-disaster investment

Firms with larger exposure to fiscal constraints require a higher return on investment. We now investigate whether this translates into lower investment, and again examine the five years preceding the disaster (fiscal years 2015 to 2019) using all firms that are available over this time period after applying the same filters as before. First, we focus on capital expenditure defined as $capx_{f,t} = \text{CAPX}_{f,t}/\text{Total Assets}_{f,t-1}$. We run equation (5) with $capx_{f,t}$ as the dependent variable and present results in Panel A of Table 7. We find that δ is significantly negative across all specifications with $|t| > 2.5$, showing that firms with more exposure to fiscal constraints have lower capital expenditure. To interpret the economic magnitude, we focus on model 4, which is our most stringent specification with HQ×Ind×Year fixed effects. Here, we obtain a δ of -0.008 ($|t| > 3.9$), which implies that a one-unit increase in FCE_f reduces capital expenditure by 0.8%, which corresponds to a 18% reduction relative to the average $capx_{f,t}$ of 0.045.

Table 7 about here

In Panel B of Table 7 we repeat the exercise for R&D expenditure, that is, we define $rd_{f,t} = \text{XRD}_{f,t}/\text{Total Assets}_{f,t-1}$ and use it as dependent variable. Investment in R&D

is thought of as being particularly important for economic growth (e.g., Romer, 1990; Aghion and Howitt, 1992). In our preferred specification (model 4), we obtain a δ-estimate of -0.019 ($|t| > 2.5$), which implies a 1.9% lower investment in R&D for a one unit increase in FCE_f. This estimate corresponds to a 49% reduction in $rd_{f,t}$ relative to the average value of 0.039.[17] The larger effects are consistent with R&D investments having a longer cash-flow duration (i.e., a cash-flow profile skewed toward the long term), which makes them more sensitive to variations in discount rates. Taken together, the results support **H3**: Firms exposed to fiscal constraints invest less in normal times, with particularly strong effects on R&D investment.

Taken together, our results indicate that fiscal constraints affect a firm's investment decisions through a risk-based disaster vulnerability channel. High-FCE_f firms become riskier in a disaster, which increases required returns on investment and thus reduces investment in normal times.

6 Exploring the mechanism

In this final section, we explore potential mechanisms underlying our findings. We begin by dissecting the discount rate result and then examine whether our results align with standard asset pricing principles. Lastly, we assess the relative importance of demand and financing conditions in the disaster state in explaining our results.

6.1 Wedges and the CAPM

We start by dissecting the discount rate result. Following Gormsen and Huber (2023b), we decompose the discount rate used in investment decisions into the perceived cost of

[17] Missing values for R&D expenditure result in a smaller sample size in Panel B. Missing values are sometimes imputed as zero (e.g., Hou et al., 2022). By doing so, our most stringent specification yields a δ-estimate of -0.006 ($|t| > 2.0$), corresponding to a 33% reduction relative to the average value (0.018).

capital, $r^{fin}_{f,t}$, and a wedge, $\omega_{f,t}$, such that $r^{dr}_{f,t} = r^{fin}_{f,t} + \omega_{f,t}$. The perceived cost of capital is the firms' internal estimate of their cost of capital, typically obtained from using a weighted average cost of capital (WACC) approach. The discount rate in turn is the hurdle rate that managers actually use in investment decisions. Interestingly, Gormsen and Huber (2023b) show that this hurdle rate is on average substantially higher than the firms' internal cost of capital estimate, suggesting that firms adjust their discount rates upwards, potentially to account for risk not captured by the standard WACC approach.

6.1.1 Pre-disaster perceived cost of capital and wedges

We want to understand whether exposure to fiscal constraints is associated with higher discount rates because of higher perceived cost of capital or higher wedges. We thus run equation (5) with $r^{fin}_{f,t}$ and $\omega_{f,t}$ as the dependent variable, respectively. Results are in Table 8, where Panel A shows that firms' perceived cost of capital are largely unrelated to fiscal constraint exposure. The coefficients are positive but small—roughly 10-20% of the discount rate coefficients in Table 6—and not statistically significant. In contrast, Panel B shows that the wedge reflects 80-90% of the discount rate coefficients, significant across all models with $|t| > 1.8$. In our preferred specification (model 4) δ is 0.60, which implies that the wedge is higher by around 0.6 percentage points for a one-unit increase in FCE_f.

TABLE 8 ABOUT HERE

One interpretation of these facts is that i) the standard WACC approach fails to fully account for risk associated with fiscal constraint exposure, and ii) managers recognize this and adjust their discount rates accordingly. Managers typically implement the textbook WACC approach by estimating the cost of equity using the CAPM applied to historical data (e.g., Graham and Harvey, 2001; Dessaint et al., 2021; Graham, 2022). However, previous analyses show that perceived fiscal-constraints risk for high-FCE_f relative to low-FCE_f firms increases in a disaster. This fact suggests that market beta estimates in normal times might not fully account for the risk arising from fiscal constraint exposure. We now turn to testing this conjecture in the data.

6.1.2 Pre-disaster CAPM alphas

We adopt a conditional CAPM-style logic to understand required returns in normal times. If market beta estimated in normal times does not fully capture the risk stemming from fiscal constraint exposure in bad times, investors should require compensation for holding such stocks over and above what their beta would imply. Thus, high-FCE_f firms should have abnormal returns relative to what would be expected based on their CAPM beta in normal times. In other words, the CAPM alpha in the pre-disaster period should be increasing in FCE_f.

Table 9 about here

We thus use the yearly CAPM alpha, $\alpha_{f,t}$, as outcome variable in equation (5) and present the results in Table 9. We obtain positive and highly significant δ-estimates in all models with $|t| > 2.0$. Specifically, in the most stringent specification (model 5), where we use all firm-level controls as well as HQ×Ind×Year fixed effects and weight by the market value of equity in USD, we obtain a δ-estimate of 0.029 ($|t| > 2.0$). Hence, we observe a spread in alphas of close to 2.9% per year for a one-unit difference in FCE_f.[18] These results demonstrate that exposure to fiscal constraints predicts pre-disaster CAPM alphas, consistent with the notion that market beta estimated in normal times does not fully capture the risk associated with fiscal constraints.[19]

[18] For consistency, we interpret our results economically with a one-unit wedge in FCE_f. However, note that this does not correspond to a high-minus-low difference in FCE_f, which is used in portfolio sorts. For instance, in Table IA.A.VI the interquartile range ($q_{0.75} - q_{0.25}$) of FCE_f is approximately 0.55.

[19] We document in the Internet Appendix that our results are robust to using (raw) excess returns in Table IA.A.VII and CAPM-alphas in domestic currency in Table IA.B.I. We obtain a δ of 0.025 ($|t| > 1.9$) and 0.026 ($|t| > 1.7$) in our most stringent specification, respectively.

6.1.3 Disaster drawdowns

Next, we turn to drawdowns in the disaster state. As shown before, fiscal constraint exposure increases perceived firm-level risk in a disaster. Given this, we hypothesize that market betas measured in normal times understate the risk associated with fiscal constraints, thus not fully reflecting a firm's true exposure to systematic risk. As a consequence, when disaster strikes and the economic environment shifts, high-FCE_f firms experience sharper declines in their stock prices than would be predicted by their market beta estimated in normal times.

To test this idea, we examine a sample of 21,125 firms headquartered in 79 countries.[20] We compute for each day $t \in [-5, \ldots, 15]$ around February 21, 2020—the day when prices in equity markets around the globe started to drop—drawdowns in USD, $DD_{f,t}$. Next, we compute the CAPM-adjusted excess drawdown, $ADD_{f,t}$, for each day t. We use the pre-disaster market beta (β_f measured as of end of 2019) to compute the adjusted excess drawdown as follows: $ADD_{f,t} = \overline{DD}_{f,t} - \beta_f \overline{DD}_{m,t}$, where $\overline{DD}_{f,t}$ is the excess drawdown of firm f on trading day t, and $\overline{DD}_{m,t}$ the corresponding excess market drawdown. We use the buy-and-hold return of the one-month US Treasury Bill for the return of the risk-free asset.

Figure 2 about here

In Panel A of Figure 2 we then run on each day t around February 21, 2020, a cross-sectional regression at the individual firm level of market-beta adjusted drawdowns, $ADD_{f,t}$, against firm-level exposure to fiscal constraints, $FCE_{f,2018}$. We include headquarter-country fixed effects and weight by the market value of equity in USD. We plot for each day t the coefficient of FCE_f (black line) and 95% confidence intervals (grey lines). High- compared to low-FCE_f firms experience significantly larger drawdowns as the disaster strikes,

[20] Summary statistics are in Table IA.A.VI in the Internet Appendix.

while there are no differential effects in the days before. On day $t = 15$, the drawdown differential is close to 10% for a one-unit increase in FCE_f.

Table 10 about here

Next, we focus on trading day $t = 15$ and show results of equation (3) in Panel A of Table 10.[21] We obtain significantly negative δ-estimates in all models with $|t| > 3.5$. Our estimates reveal that high-FCE_f firms exhibit larger drawdowns adjusted for market beta estimated in normal times. We focus on model 5 to interpret the economic magnitude, where we employ all firm-level controls with HQ×Ind fixed effects and weight by market value in USD. Here, we obtain a δ-estimate of -0.09 ($|t| > 3.0$), which implies a 9% larger market-beta adjusted drawdown for a one-unit increase in FCE_f.[22] Hence, otherwise identical firms that are headquartered in the same country and operate in the same industry, but differ in their exposure to fiscal constraints, exhibit economically significant differentials in drawdowns relative to their market-risk exposure estimated in normal times.

6.1.4 Disaster betas

We now turn to the dynamics of market betas. During normal times, the market beta of firms exposed to fiscal constraints may be understated because their true sensitivity to systematic risk is not fully captured. However, in a disaster state, fiscal constraints amplify the risk associated with the economic downturn: as fiscal constraints limit government support and tighten credit conditions, the performance of firms exposed to fiscal constraints could be more dependent on overall economic conditions, leading to a higher market beta.

[21] We exclude β from the HXZ/FF characteristics in the specifications in Table 10.

[22] Robustness in domestic currency stock prices are in Table IA.B.II in the Internet Appendix. Risk-adjusted drawdowns, using pre-disaster β, are larger for firms more exposed to fiscal constraints (significantly negative δ-estimates throughout all models with $|t| > 3.3$). In our most stringent specification we obtain a δ of -0.08.

The previous analyses show that, in a disaster, perceived fiscal constraints risk increases more for high-FCE_f firms compared to low-FCE_f firms. Moreover, this section shows that risk-adjusted drawdowns (using pre-disaster market betas) around the disaster outbreak are more pronounced for firms with greater exposure to fiscal constraints. In the subsequent analysis, we investigate the dynamics of market betas to understand whether the change in risk perceptions is reflected in the dynamics of realized market betas. This analysis also helps to understand whether drawdown differentials can be understood by changes in market betas.

Panel B of Figure 2 shows the beta dynamics as a function of FCE_f. To address non-synchronous trading, we estimate daily betas for each firm using three-days overlapping returns (as in Frazzini and Pedersen, 2014) in a 180-day rolling window and calculate the average of the estimates in each quarter. Akin to previous analyses, we run in each quarter $t \in$ [Q1 2019,...,Q4 2020] a cross-sectional regression at the individual firm level of the change in market betas against firm-level exposure to fiscal constraints, $FCE_{f,2018}$, as well as headquarter-country fixed effects. Specifically, the dependent variable is the change in beta from Q4 2019 to quarter t, $\beta_{f,t} - \beta_{f,\text{Q4 2019}}$. We weight by the market value of equity in USD. We plot for each quarter the coefficient of FCE_f (black line) and 95% confidence intervals (grey lines). There is no discernible difference in trends before the disaster, but high-FCE_f firms experience a strong increase in market betas compared to low-FCE_f firms as disaster strikes, with a differential of up to 0.4 for a one-unit increase in FCE_f. Hence, the comovement with the global market portfolio increases in the disaster for firms more exposed to fiscal constraints.

Panel B in Table 10 presents the results of estimating (3) using as dependent variable the change in beta, that is $Y_f = \Delta\beta_{f,\text{Q3 2020}-\text{Q4 2019}} = \beta_{f,\text{Q3 2020}} - \beta_{f,\text{Q4 2019}}$. In all models, we obtain significantly positive δ-estimates with $|t| > 3.2$. The results confirm that during a disaster, firms with high FCE_f experience an increase in their market betas compared to firms with low FCE_f. In model 5, utilizing value-weighted regressions, we obtain a δ-estimate of 0.32, implying that a one-unit increase in FCE_f is associated with a beta wedge of 0.32 in the disaster.[23]

These results indicate that assets with higher exposure to fiscal constraints experience an increase in their comovement with the global market portfolio in a disaster state. This mechanism can explain why firms with high values of FCE_f exhibit larger market-beta adjusted drawdowns when the disaster unfolds. We offer direct evidence supporting this idea in Panel C of Table 10. Here, we compute risk-adjusted drawdowns using the realized disaster beta, $\beta_{f,\text{Q3 2020}}$, and re-run equation (3). We find substantially lower risk-adjusted drawdowns when using realized rather than pre-disaster betas: The δ-estimates in Panel C are 70-85% lower compared to Panel A and insignificant in four out of five models. Finally, we do back-of-the-envelope calculations to see if the beta increase in the disaster state is broadly consistent with the CAPM alpha in normal times (estimates in Table 9). Using VIX data to gauge market volatility, these calculations imply that an increase of beta by 0.3 in the disaster state would entail 3% higher CAPM alphas in normal times. This magnitude is close to our estimates in Table 9, which range from 2.3% to 2.9%.[24]

Overall, these dynamics in asset prices and risk align well with the notion that fiscal constraints shape the vulnerability of assets in a disaster state. Next, we provide additional tests to shed light on the possible channels underlying these disaster effects.

[23] Market betas based on domestic currency stock returns are in Table IA.B.II in the Internet Appendix. Betas increase for high-FCE_f firms compared to low-FCE_f firms in a disaster throughout all models ($|t| > 3.3$). In our preferred specification we obtain a δ of 0.34.

[24] Suppose the true expected return on a stock i in state s, either normal (n) or disaster (d), is $E(R_i^s) = R_f + \beta_i^s \times \lambda_s$, where the market risk premium in state s is $\lambda_s = \gamma \times \sigma_s^2$, with γ representing risk aversion and σ_s^2 the market variance. The expected return across states is $E(R_i) = (1 - p_d) \times E(R_i^n) + p_d \times E(R_i^d)$, with p_d denoting the disaster probability. The CAPM estimated in normal times predicts that the stock's return is $E(R_i^n)$. Hence, the observed CAPM alpha is $\alpha_i = E(R_i) - E(R_i^n) = p_d \left[\beta_d \cdot \gamma \cdot \sigma_d^2 - \beta_n \cdot \gamma \cdot \sigma_n^2\right]$. Setting $\gamma = 3$ and $p_d = 0.0355$ as in Wachter (2013), $\beta_n = 1$, $\beta_d = 1.3$, $\sigma_n = 0.15$ (average VIX 2015-2019), and $\sigma_d = 0.5$ (average VIX mid-February to mid-April 2020), we obtain $\alpha = 0.03$.

6.2 Disaster channels

If firms exposed to fiscal constraints are more vulnerable to global shocks due to demand effects, their sales should drop more during global downturns, compared to less exposed firms. Alternatively, firms exposed to fiscal constraints could be more vulnerable to global shocks due to a deterioration of their financing conditions, for instance due to the 'doom-loop' or due to sub-optimally low levels and heightened uncertainty about government intervention in lending (e.g., loan guarantees and subsidized credit).

Note that our empirical approach is more likely to pick up demand effects rather than financing effects, as we measure exposure to fiscal constraints based on firms' international sales distribution. Nevertheless, it is still possible that firms' sales distribution correlates with their financing patterns, meaning that firms selling to customers in fiscally constrained countries may also borrow more from banks in those countries. Importantly, these mechanisms are not mutually exclusive and may operate simultaneously. We do not intend to rule out one of the two channels, but rather want to gauge their relative importance in driving the results in our empirical setting.

Table IA.A.VIII in the Internet Appendix shows that high-FCE_f firms experience a greater drop in sales. Specifically, we run equation (3) using the change in sales from the pre-disaster to the disaster period as the dependent variable. We obtain negative and significant δ-estimates throughout all models ($|t| > 1.9$). In our most stringent specification (model 4), we obtain a 6% drop in sales for a one-unit increase in FCE_f.

In a second test, we examine the role of direct financing effects in driving the results. We test the relevance of this channel in our empirical setting by excluding firms which borrow from institutions outside their headquarter country, thereby eliminating direct lending exposure to foreign banks. Paired with headquarter-country fixed effects, this approach compares firms that borrow from banks in the same country but are differentially exposed to fiscal constraints through their international sales distribution. If direct financing effects play an important role in the asset price dynamics we document, fiscal constraints should amplify disaster vulnerability to a lesser extent in this sub-sample of firms that borrow solely from domestic banks, as these firms would primarily be exposed to demand effects, with the direct financing channel excluded.[25]

To explore this idea, we use Dealscan data to examine the role of domestic versus foreign borrowing. We split the sample based on whether firms borrow from foreign banks (i.e., banks outside their headquarter country), thus eliminating direct financing exposure to foreign banks for the sub-sample of firms that borrow exclusively domestically. Panel A of Table IA.A.IX in the Internet Appendix distinguishes firms that borrow from foreign banks versus those that borrow only domestically as observed in the Dealscan data. Panel B instead classifies firms based on observable characteristics indicative of foreign borrowing for the full sample. The results show no difference in risk-adjusted drawdowns or market beta dynamics between the sub-samples of domestic and foreign borrowers. This suggests that the direct financing channel does not play a significant role in the disaster effects we observe. Taken together, these additional tests suggest that fiscal constraint exposure amplifies disaster vulnerability through demand effects in our empirical setup.

7 Conclusion

We study the role of fiscal constraints for disaster vulnerability. To that end, we build a novel firm-level measure of exposure to fiscal constraints. By applying tools from computational linguistics to earnings conference calls, we first demonstrate that our exposure measure predicts higher levels of perceived fiscal constraints risk in normal times as well as

[25] One potential limitation of this test is that firms that borrow only domestically may differ systematically from those with foreign financing, potentially making them more vulnerable to economic shocks on other dimensions. To address this, we include controls for firm characteristics associated with financial vulnerability (e.g., size, leverage, cash holdings, tangibility, profitability, book-to-market). Nonetheless, we acknowledge that unobserved differences between these sub-samples may persist, and thus we view the results as suggestive rather than definitive.

an increase in this risk in a disaster state. We conduct a series of validation tests to ensure that our measure of exposure to fiscal constraints accurately captures fiscal constraints-related risks, rather than serving as a proxy for broader country risk.

We find that firms with higher exposure to fiscal constraints exhibit higher discount rates set by management and lower investment in normal times. We show that differences in discount rates stem from higher wedges rather than perceived cost of capital, suggesting that the standard WACC approach using CAPM cost-of-equity estimates does not fully capture fiscal constraints-related risk. Consistent with this notion, we find that firms exposed to fiscal constraints experience larger disaster-induced drawdowns relative to their CAPM beta estimated in normal times, while their market betas increase. We further find that their sales drop more in the disaster state, while direct financing effects do not appear to be driving the results.

It is important to note that our conclusions are based on a very stringent empirical setup. In particular, thanks to our firm-level measure, we can exploit variation in exposure to fiscal constraints for firms with identical characteristics that are headquartered in the same country and operate in the same industry. Overall, our article has important implications for understanding the effect of fiscal constraints on growth and emphasizes the role of a risk-based disaster vulnerability channel.

References

Abbas, S. Ali, Alex Pienkowski, and Kenneth Rogoff, 2019, *Sovereign debt: A guide for economists and practitioners* (Oxford University Press).

Acharya, Viral, Itamar Drechsler, and Philipp Schnabl, 2014, A pyrrhic victory? Bank bailouts and sovereign credit risk, *Journal of Finance* 69, 2689–2739.

Aghion, Philippe, and Peter Howitt, 1992, A model of growth through creative destruction, *Econometrica* 60, 323–351.

Alesina, Alberto, Carlo Favero, and Francesco Giavazzi, 2020, *Austerity: When it Works and when it Doesn't* (Princeton University Press).

Almeida, Heitor, and F. Argelina, 2009, Credit market conditions and the corporate sector: Evidence from emerging markets, *Journal of Financial Economics* 91, 143–162.

Andrade, Sandro, and Vidhi Chhaochharia, 2018, The costs of sovereign default: Evidence from the stock market, *Review of Financial Studies* 31, 1707–1751.

Angrist, Joshua D, and Jörn-Steffen Pischke, 2009, *Mostly harmless econometrics: An empiricist's companion* (Princeton University Press).

Arellano, Cristina, Yan Bai, and Luigi Bocola, 2017, Sovereign default risk and firm heterogeneity, Working paper, National Bureau of Economic Research.

Arteaga-Garavito, Maria Jose, M. Max Croce, Paulo Farroni, and Isabella Wolfskeil, 2024, When the markets get covid: Contagion, viruses, and information diffusion, *Journal of Financial Economics* 157.

Auerbach, Alan J, and Yuriy Gorodnichenko, 2012, Measuring the output responses to fiscal policy, *American Economic Journal: Economic Policy* 4, 1–27.

Augustin, Patrick, 2018, The term structure of CDS spreads and sovereign credit risk, *Journal of Monetary Economics* 96, 53–76.

Augustin, Patrick, Hamid Boustanifar, Johannes Breckenfelder, and Jan Schnitzler, 2018, Sovereign to corporate risk spillovers, *Journal of Money, Credit and Banking* 50, 857–891.

Augustin, Patrick, Valeri Sokolovski, Marti G. Subrahmanyam, and Davide Tomio, 2022, In sickness and in debt: The covid-19 impact on sovereign credit risk, *Journal of Financial Economics* 143, 1251–1274.

Augustin, Patrick, and Roméo Tédongap, 2016, Real economic shocks and sovereign credit risk, *Journal of Financial and Quantitative Analysis* 51, 541–587.

Barro, Robert J., 2006, Rare disasters and asset markets in the twentieth century, *Quarterly Journal of Economics* 121, 823–866.

Barrot, Jean-Noel, Maxime Bonelli, Basile Grassi, and Julien Sauvagnat, 2023, Causal effects of closing businesses in a pandemic, *Journal of Financial Economics* forthcoming.

Barry, John W., Bruce Carlin, Alan D. Crane, and John R. Graham, 2023, Project development with delegated bargaining: The role of elevated hurdle rates, Working paper.

Bedendo, Mascia, and Paolo Colla, 2015, Sovereign and corporate credit risk: Evidence from the Eurozone, *Journal of Corporate Finance* 33, 34–52.

Belo, Frederico, Vito D. Gala, and Jun Li, 2013, Government spending, political cycles, and the cross section of stock returns, *Journal of Financial economics* 107, 305–324.

Belo, Frederico, and Jianfeng Yu, 2013, Government investment and the stock market, *Journal of Monetary Economics* 60, 325–339.

Bernanke, Ben S., and Mark Gertler, 1989, Agency costs, net worth, and business fluctuations, *American Economic Review* 79, 14–31.

Bessembinder, Hendrik, Te-Feng Chen, Goeun Choi, and KC Wei, 2019, Do global stocks outperform US treasury bills?, Working paper, Carey School of Business.

Blinder, Alan S, and Robert M Solow, 1973, Does fiscal policy matter?, *Journal of Public Economics* 2, 319–337.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, 2022, Salience, *Annual Review of Economics* 14, 521–544.

Bretscher, Lorenzo, Alex Hsu, Peter Simasek, and Andrea Tamoni, 2020, Covid-19 and the cross-section of equity returns: Impact and transmission, *Review of Asset Pricing Studies* 10, 705–741.

Brunnermeier, Markus K., Luis Garicano, Philip R. Lane, Marco Pagano, Ricardo Reis, Tano Santos, David Thesmar, Stijn Van Nieuwerburgh, and Dimitri Vayanos, 2016, The sovereign-bank diabolic loop and ESBies, *American Economic Review* 106, 508–12.

Corsetti, Giancarlo, Keith Kuester, André Meier, and Gernot J. Müller, 2013, Sovereign risk, fiscal policy, and macroeconomic stability, *Economic Journal* 123, F99–F132.

Croce, M. Max, Howard Kung, Thien T. Nguyen, and Lukas Schmid, 2012, Fiscal policies and asset prices, *Review of Financial Studies* 25, 2635–2672.

Croce, M. Max, Thien T. Nguyen, and Steve Raymond, 2021, Persistent government debt and aggregate risk distribution, *Journal of Financial Economics* 140, 347–367.

Croce, M. Max, Thien T. Nguyen, Steve Raymond, and Lukas Schmid, 2019, Government debt and the returns to innovation, *Journal of Financial Economics* 132, 205–225.

Croce, M. Max, Thien T. Nguyen, and Lukas Schmid, 2012, The market price of fiscal uncertainty, *Journal of Monetary Economics* 59, 401–416.

De Marco, Filippo, 2019, Bank lending and the European sovereign debt crisis, *Journal of Financial and Quantitative Analysis* 54, 155–182.

Dessaint, Olivier, Jacques Olivier, Clemens A. Otto, and David Thesmar, 2021, Capm-based company (mis) valuations, *Review of Financial Studies* 34, 1–66.

Ding, Wenzhi, Ross Levine, Chen Lin, and Wensi Xie, 2020, Corporate immunity to the covid-19 pandemic, *Journal of Financial Economics* 141, 802–830.

Fahlenbrach, Rüdiger, Kevin Rageth, and Rene M. Stulz, 2021, How valuable is financial flexibility when revenue stops? Evidence from the covid-19 crisis, *Review of Financial Studies* 34, 5474–5521.

Fama, Eugene F., and Kenneth R. French, 2015, A five-factor asset pricing model, *Journal of Financial Economics* 116, 1–22.

Farhi, Emmanuel, and Jean Tirole, 2018, Deadly embrace: Sovereign and financial balance sheets doom loops, *Review of Economic Studies* 85, 1781–1823.

Frazzini, Andrea, and Lasse Heje Pedersen, 2014, Betting against beta, *Journal of Financial Economics* 111, 1–25.

Gabaix, Xavier, 2012, Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance, *Quarterly Journal of Economics* 127, 645–700.

Gennaioli, Nicola, Alberto Martin, and Stefano Rossi, 2014, Sovereign default, domestic banks, and financial institutions, *Journal of Finance* 69, 819–866.

Gormsen, Niels Joachim, and Kilian Huber, 2023a, Corporate discount rates, Working paper, University of Chicago Booth School of Business.

Gormsen, Niels Joachim, and Kilian Huber, 2023b, Firms perceived cost of capital, Working paper, University of Chicago Booth School of Business.

Gormsen, Niels Joachim, and Ralph Koijen, 2020, Coronavirus: Impact on stock prices and growth expectations, *Review of Asset Pricing Studies* 10, 574–597.

Gormsen, Niels Joachim, and Ralph Koijen, 2023, Financial markets and the covid-19 pandemic, *Annual Review of Financial Economics* 15, 69–89.

Gourio, Francois, 2012, Disaster risk and business cycles, *American Economic Review* 102, 2734–66.

Graham, John R., 2022, Presidential address: Corporate finance and reality, *Journal of Finance* 77, 1975–2049.

Graham, John R., and Campbell R. Harvey, 2001, The theory and practice of corporate finance: Evidence from the field, *Journal of Financial Economics* 60, 187–243.

Griliches, Zvi, 1981, Market value, r&d, and patents, *Economic Inquiry* 19, 22–36.

Hall, Bronwyn H., 1993, R&d tax policy during the 1980s: Successes and failures, *The American Economic Review* 83, 431–437.

Hassan, Tarek A., Stephan Hollander, Laurence Van Lent, Markus Schwedeler, and Ahmed Tahoun, 2023a, Firm-level exposure to epidemic diseases: Covid-19, SARS, and H1N1, *Review of Financial Studies* 36, 4919–4964.

Hassan, Tarek A., Stephan Hollander, Laurence Van Lent, and Ahmed Tahoun, 2019, Firm-level political risk: Measurement and effects, *Quarterly Journal of Economics* 134, 2135–2202.

Hassan, Tarek A., Jesse Schreger, Markus Schwedeler, and Ahmed Tahoun, 2023b, Sources and transmission of country risk, *Review of Economic Studies* forthcoming.

Hayashi, Fumio, 1982, Tobin's marginal q and average q: A neoclassical interpretation, *Econometrica* 50, 213–224.

He, Zhiguo, Guanmin Liao, and Baolian Wang, 2022, What gets measured gets managed: Investment and the cost of capital, National Bureau of Economic Research.

Hébert, Benjamin, and Jesse Schreger, 2017, The costs of sovereign default: Evidence from argentina, *American Economic Review* 107, 3119–3145.

Hilscher, Jens, and Yves Nosbusch, 2010, Determinants of sovereign risk: Macroeconomic fundamentals and the pricing of sovereign debt, *Review of Finance* 14, 235–262.

Holmstrom, Bengt, and Jean Tirole, 1997, Financial intermediation, loanable funds, and the real sector, *Quarterly Journal of Economics* 112, 663–691.

Hou, Kewei, Haitao Mo, Chen Xue, and Lu Zhang, 2019, Which factors?, *Review of Finance* 23, 1–35.

Hou, Kewei, Haitao Mo, Chen Xue, and Lu Zhang, 2022, The economics of security analysis, *Management Science* Advanced access.

Hou, Kewei, Chen Xue, and Lu Zhang, 2015, Digesting anomalies: An investment approach, *Review of Financial Studies* 28, 650–705.

Hou, Kewei, Chen Xue, and Lu Zhang, 2020, Replicating anomalies, *Review of Financial Studies* 33, 2019–2133.

Jacobs, Michael T., and Anil Shivdasani, 2012, Do you know your cost of capital?, *Harvard Business Review* 90, 118–124.

Jagannathan, Ravi, David A. Matsa, Iwan Meier, and Vefa Tarhan, 2016, Why do firms use high discount rates?, *Journal of Financial Economics* 120, 445–463.

Jappelli, Ruggero, Loriana Pelizzon, and Alberto Plazzi, 2021, The core, the periphery, and the disaster: Corporate-sovereign nexus in covid-19 times, Working paper, SAFE.

Jermann, Urban J, and Vincenzo Quadrini, 2012, Macroeconomic effects of financial shocks, *American Economic Review* 102, 238–271.

Jiang, Zhengyang, Hanno Lustig, Stijn Van Nieuwerburgh, and Mindy Z. Xiaolan, 2023, Fiscal capacity: An asset pricing perspective, *Annual Review of Financial Economics* 15, 197–219.

Jorgenson, Dale W., 1963, Capital theory and investment behavior, *American Economic Review* 53, 247–259.

Keynes, John Maynard, 1936, *The General Theory of Employment, Interest, and Money* (Macmillan).

Kim, Yongjin, 2022, Does corporate investment respond to the time-varying cost of capital? empirical evidence, *Journal of Financial and Quantitative Analysis* 57, 629–668.

Kiyotaki, Nobuhiro, and John Moore, 1997, Credit cycles, *Journal of Political Economy* 105, 211–248.

Kose, M. Ayhan, Sergio Kurlat, Franziska Ohnsorge, and Naotaka Sugawara, 2022, A cross-country database of fiscal space, *Journal of International Money and Finance* 128, 102682.

Krüger, Philipp, Augustin Landier, and David Thesmar, 2015, The WACC fallacy: The real effects of using a unique discount rate, *Journal of Finance* 70, 1253–1285.

Lamont, Owen A., 2000, Investment plans and stock returns, *Journal of Finance* 55, 2719–2745.

Landier, Augustin, and David Thesmar, 2020, Earnings expectations during the covid-19 crisis, *Review of Asset Pricing Studies* 10, 598–617.

Li, Ningzhong, Scott Richardson, and İrem Tuna, 2014, Macro to micro: Country exposures, firm fundamentals and stock returns, *Journal of Accounting and Economics* 58, 1–20.

Libby, Robert, Patricia Libby, and Frank Hodge, 2019, *Financial Accounting* (McGraw-Hill), 10th Edition.

Liu, Yang, 2022, Government debt and risk premia, Working paper, University of Hong Kong.

Longstaff, Francis, and Monika Piazzesi, 2004, Corporate earnings and the equity premium, *Journal of Financial Economics* 74, 401–421.

Longstaff, Francis A., Jun Pan, Lasse H. Pedersen, and Kenneth J. Singleton, 2011, How sovereign is sovereign credit risk?, *American Economic Journal: Macroeconomics* 3, 75–103.

Mazur, Mieszko, Man Dang, and Miguel Vega, 2021, Covid-19 and the march 2020 stock market crash. Evidence from S&P1500, *Finance Research Letters* 38, 101690.

Mehra, Rajnish, and Edward C. Prescott, 1985, The equity premium: A puzzle, *Journal of Monetary Economics* 15, 145–161.

Meier, Iwan, and Vefa Tarhan, 2007, Corporate investment decision practices and the hurdle rate premium puzzle, Working paper.

Modigliani, Franco, and Merton H. Miller, 1958, The cost of capital, corporation finance and the theory of investment, *American Economic Review* 48, 261–297.

Mukhlynina, Lilia, and Kjell G. Nyborg, 2016, The choice of valuation techniques in practice: education versus profession, Swiss Finance Institute Research Paper.

Pagano, Marco, Christian Wagner, and Josef Zechner, 2023, Disaster resilience and asset prices, *Journal of Financial Economics* 150, 103712.

Pan, Jun, and Kenneth J. Singleton, 2008, Default and recovery implicit in the term structure of sovereign cds spreads, *Journal of Finance* 63, 2345–2384.

Pastor, Lubos, and Pietro Veronesi, 2012, Uncertainty about government policy and stock prices, *Journal of Finance* 67, 1219–1264.

Pflueger, Carolin, Emil Siriwardane, and Adi Sunderam, 2020, Financial market risk perceptions and the macroeconomy, *Quarterly Journal of Economics* 135, 1443–1491.

Poterba, James M., and Lawrence H. Summers, 1995, A CEO survey of US companies' time horizons and hurdle rates, *MIT Sloan Management Review* 37, 43.

Ramelli, Stefano, and Alexander F. Wagner, 2020, Feverish stock price reactions to covid-19, *Review of Corporate Finance Studies* 9, 622–655.

Reinhart, Carmen M., Vincent R. Reinhart, and Kenneth S. Rogoff, 2012, Public debt overhangs: advanced-economy episodes since 1800, *Journal of Economic Perspectives* 26, 69–86.

Reinhart, Carmen M., and Kenneth S. Rogoff, 2010, Growth in a time of debt, *American Economic Review* 100, 573–78.

Rietz, Thomas A., 1988, The equity risk premium: a solution, *Journal of Monetary Economics* 22, 117–131.

Romer, Christina D., and David H. Romer, 2018, Phillips Lecture-Why some times are different: Macroeconomic policy and the aftermath of financial crises, *Economica* 85, 1–40.

Romer, Christina D., and David H. Romer, 2019, Fiscal space and the aftermath of financial crises: How it matters and why, Working paper, NBER.

Romer, Paul M., 1990, Endogenous technological change, *Journal of Political Economy* 98, S71–S102.

Schutze, Hinrich, Christopher D. Manning, and Prabhakar Raghavan, 2008, *Introduction to information retrieval* (Cambridge University Press).

Sharpe, Steven A., and Gustavo A. Suarez, 2021, Why isn't business investment more sensitive to interest rates? evidence from surveys, *Management Science* 67, 720–741.

Song, Min, and Yi-Fang Brook Wu, 2008, *Handbook of research on text and web mining technologies* (IGI global).

Summers, Lawrence H., 1986, Investment incentives and the discounting of depreciation allowances, National Bureau of Economic Research.

Tobin, James, 1969, A general equilibrium approach to monetary theory, *Journal of Money, Credit and Banking* 1, 15–29.

Tsai, Jerry, and Jessica A. Wachter, 2015, Disaster risk and its implications for asset pricing, *Annual Review of Financial Economics* 7, 219–252.

Van Binsbergen, Jules H., and Christian C. Opp, 2019, Real anomalies, *Journal of finance* 74, 1659–1706.

Wachter, Jessica A., 2013, Can time-varying risk of rare disasters explain aggregate stock market volatility?, *Journal of Finance* 68, 987–1035.

Panel A:
Perceived Fiscal Constraints Risk

Panel B:
Perceived Non-Fiscal Constraints Risk

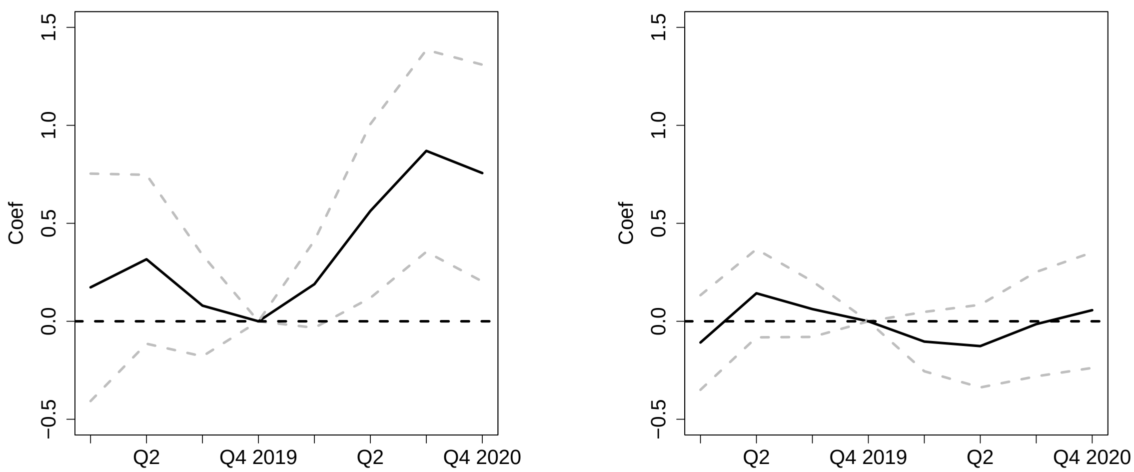

Figure 1: Exposure to fiscal constraints and perceived disaster risk dynamics. In Panel A we run in each quarter $t \in$ [Q1 2019,…,Q4 2020] a cross-sectional regression at the individual firm level of the change in perceived fiscal constraints-related risk, ΔFCR_f against firm-level exposure to fiscal constraints, $FCE_{f,2018}$, as well as headquarter-country fixed effects. Specifically, the dependent variable is the change in firm-level perceived fiscal constraints risk from Q4 2019 to quarter t, $\Delta FCR_{f,t} = (FCR_{f,t} - FCR_{f,\text{Q4 2019}})/\overline{FCR}_{\text{Q4 2019}}$, where $\overline{FCR}_{\text{Q4 2019}}$ is the cross-sectional median in Q4 2019. We plot for each quarter the coefficient of $FCE_{f,2018}$ (black line) and 95% confidence intervals (grey lines). In Panel B we repeat the analysis with the change in perceived firm-level non-fiscal constraints-related risk, $\Delta NFCR_f$, as the dependent variable. Standard errors are clustered at the industry level (six-digit GICS). The sample covers 3,043 firms headquartered in 55 countries.

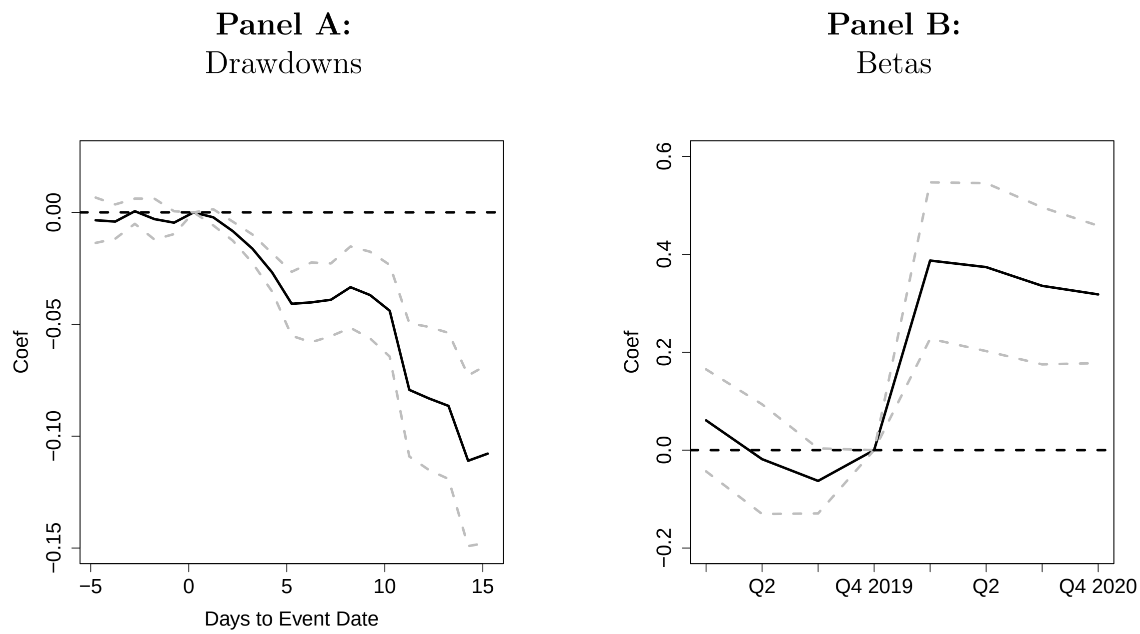

Figure 2: Exposure to fiscal constraints and disaster asset price dynamics. In Panel A we run on each day $t \in [-5, \dots, 15]$ around February 21, 2020, a cross-sectional regression at the individual firm level of the CAPM-adjusted excess drawdown in USD against firm-level exposure to fiscal constraints, $FCE_{f,2018}$, as well as headquarter-country fixed effects. We use the pre-disaster market beta (i.e. from end of 2019) to compute the adjustment, and run weighted least squares (VW), where we weight by the market value of equity in USD. We plot for each day t the coefficient of $FCE_{f,2018}$ (black line) and 95% confidence intervals (grey lines). In Panel B we run in each quarter $t \in$ [Q1 2019,. . . ,Q4 2020] a cross-sectional regression at the individual firm level of the change in market betas against firm-level exposure to fiscal constraints, $FCE_{f,2018}$, as well as headquarter-country fixed effects. Specifically, the dependent variable is the change in beta from Q4 2019 to quarter t, $\beta_{f,t} - \beta_{f,\text{Q4 2019}}$. Market betas are based on three-days overlapping returns in USD in a 180-days window, and we use for each quarter the average beta-estimate. We run weighted least squares (VW), where we weight by the market value of equity in USD. We plot for each quarter the coefficient of $FCE_{f,2018}$ (black line) and 95% confidence intervals (grey lines). Standard errors are clustered at the industry level (six-digit GICS). The sample covers 21,125 firms headquartered in 79 countries.

Table 1: The language of fiscal constraints. Panel A shows the top 50 bigrams with the highest frequency in the training library used in the construction of perceived firm-level fiscal constraints risk, FCR_f. Panel B shows examples of how participants talk about fiscal constraints in earnings conference calls.

Panel A: Bigrams

#	bigram	count	#	bigram	count
1	public debt	323	26	financial repression	55
2	sovereign debt	273	27	financial crisis	54
3	government debt	189	28	monetary policy	53
4	advanced economies	129	29	washington dc	52
5	percent gdp	129	30	debt crises	52
6	central bank	120	31	balance sheet	52
7	debt sustainability	118	32	domestic debt	51
8	international monetary	118	33	external debt	51
9	interest rates	114	34	paris club	50
10	monetary fund	108	35	real interest	48
11	fiscal policy	106	36	lowincome countries	48
12	debt restructuring	100	37	emerging market	47
13	public sector	95	38	new york	46
14	united states	94	39	foreign currency	46
15	debt management	87	40	public investment	45
16	general government	82	41	central government	43
17	interest rate	79	42	debttogdp ratio	43
18	sovereign default	77	43	local currency	42
19	economic growth	72	44	debt crisis	42
20	debt relief	68	45	fiscal rules	41
21	primary balance	64	46	contingent liabilities	41
22	debt securities	64	47	debt dynamics	40
23	world war	58	48	central banks	40
24	debt reduction	56	49	euro area	39
25	world bank	56	50	fiscal adjustment	37

Table continuous on next page.

Panel B: Examples

Firm	Call Date	Text
GRUPO CLARIN SA	9/9/2020	... During the second quarter of 2020, the Argentine economy continue to be immersed in a complex macroeconomic environment, marked by a strong fiscal and monetary imbalance, tightened controls on the foreign exchange market and high levels of uncertainty, especially derived from the results of the sovereign debt renegotiation process. ...
ITALGAS SPA	29/04/2020	... The second question is the cost of debt of the company, whether we see a risk of an increase given the Italian government yield? And the third question is related to the recent downgrade by future of the sovereign debt. And whether there could be an impact for Italgas? ...
ARCHER LIMITED	14/08/2020	... with the development of the corona pandemic, it became clear that Argentina would default on its sovereign debt. And in 2020, the government defaulted on USD 65 billion of its outstanding debt. During the following months, the country continued with its lockdown while negotiating with the creditors. During the period of negotiations, the government imposed further restriction on foreign capital transactions. In effect, it has become more difficult to transfer cash out of the country, but it has also stabilized the official foreign exchange rate and inflation for now. ...
WINNEBAGO INDUSTRIES	25/03/2020	... The next 30 days will be critical for us to better understand the magnitude of our remaining opportunity for revenue and profits in fiscal year 2020. And clarity will be provided through a variety of external actions, including legislative decisions, fiscal policy action, government social interaction directives, medical developments and so much more. Our view for these forward-looking thoughts have changed from fiscally and quarterly to monthly and weekly, if not daily. These are truly unprecedented times. ...
NORFOLK SOUTHERN CORP	29/07/2020	... A high degree of uncertainty exists. State and local government reaction to increased covid cases will have a major influence on the speed and shape of the economic recovery. The growing number of such cases and fiscal policy will impact our outlook to the extent that both influence the strength of consumer demand and labor availability, with the latter having an effect on production. ...

Table continued from previous page.

Table 2: Summary statistics: earnings conference calls. This table reports the mean, standard deviation, and 10th, 25th, 50th, 75th, and 90th percentiles for the earnings conference calls sample. We report the change in perceived firm-level fiscal constraints risk from Q4 2019 to Q3 2020, $\Delta FCR_f = (FCR_{f,\text{Q3 2020}} - FCR_{f,\text{Q4 2019}})/\overline{FCR}_{\text{Q4 2019}}$, where $\overline{FCR}_{\text{Q4 2019}}$ is the cross-sectional median in Q4 2019. We further provide statistics for changes in firm-level non-fiscal constraints-related risk, $\Delta NFCR_f$, firm-level exposure to fiscal constraints, FCE_f, the logarithm of firm-level covid cases, $\log(1 + cases_f)$, the logarithm of firm-level GDP per capita, $\log(gdppc_f)$, firm-level five-year GDP growth, $gdpg_f$, share of foreign sales, w_f^{for}, pre-disaster market beta, β, logarithm of market value of equity in USD, $\log(me)$, book-to-market ratio, bm, operating profitability, op, investment-to-assets, i/a, return on equity, roe, leverage, lev, cash holdings, $cash$, and property plant and equipment, ppe. Accounting data are from the fiscal year 2018. The sample covers 3,043 firms headquartered in 55 countries.

	obs	mean	sd	$q_{0.1}$	$q_{0.25}$	$q_{0.5}$	$q_{0.75}$	$q_{0.9}$
ΔFCR_f	3043	1.31	3.26	−1.26	−0.01	0.83	2.28	4.47
$\Delta NFCR_f$	3043	0.54	1.65	−0.79	−0.17	0.33	1.11	2.26
FCE_f	3043	−0.05	0.36	−0.47	−0.20	−0.00	0.14	0.15
$\log(1 + cases_f)$	3043	10.83	1.37	8.45	10.43	11.30	11.86	12.02
$\log(gdppc_f)$	3043	10.57	0.66	9.67	10.49	10.77	10.99	11.05
$gdpg_f$	3043	0.14	0.06	0.10	0.12	0.12	0.14	0.18
w_f^{for}	3043	0.40	0.36	0.00	0.02	0.33	0.76	0.96
β	3043	1.02	0.60	0.32	0.63	0.96	1.35	1.80
$\log(me)$	3043	21.27	1.90	18.74	20.03	21.44	22.67	23.81
bm	3043	0.68	1.14	0.07	0.18	0.39	0.76	1.38
op	3043	0.24	0.47	−0.04	0.13	0.23	0.39	0.63
i/a	3043	0.11	0.33	−0.11	−0.02	0.05	0.16	0.37
roe	3043	0.06	0.41	−0.20	0.01	0.10	0.19	0.32
lev	3043	0.25	0.18	0.00	0.09	0.24	0.36	0.49
$cash$	3043	0.16	0.18	0.01	0.04	0.09	0.20	0.43
ppe	3043	0.26	0.24	0.03	0.07	0.18	0.40	0.65

Table 3: Disaster dynamics of perceived firm-level fiscal constraints risk. This table reports cross-sectional regressions at the individual firm level. The dependent variable is the change in perceived firm-level fiscal constraints risk from Q4 2019 to Q3 2020, $\Delta FCR_f = (FCR_{f,\text{Q3 2020}} - FCR_{f,\text{Q4 2019}})/\overline{FCR}_{\text{Q4 2019}}$, where $\overline{FCR}_{\text{Q4 2019}}$ is the cross-sectional median in Q4 2019. We present several specifications, where Country$_f$ are firm-level country characteristics ($\log(1 + cases_f)$, $\log(gdppc_f)$, $gdpg_f$), Other$_f$ are firm characteristics including the share of foreign sales, w_f^{for}, leverage, lev, cash holdings, $cash$, and tangibility, ppe, HXZ refers to the Hou, Xue, and Zhang (2015) characteristics, and FF to the Fama and French (2015) characteristics. HQ+Ind and HQ$\times$Ind are fixed effects, where HQ is the headquarter country and Ind the industry classification (two-digit GICS). Accounting data are from the fiscal year 2018. t-statistics (in square brackets) are based on standard errors clustered at the industry level (six-digit GICS). The sample covers 3,043 firms headquartered in 55 countries.

	(1)	(2)	(3)	(4)
$FCE_{f,2018}$	0.793	0.799	0.801	0.823
	[2.094]	[2.095]	[2.116]	[2.317]
Country$_f$ char.	Y	Y	Y	Y
Other$_f$ char.	Y	Y	Y	Y
HXZ char.		Y		
FF char.			Y	Y
HQ + Ind	Y	Y	Y	
HQ $\times$ Ind				Y
Obs.	3,043	3,043	3,043	3,043

Table 4: Falsification test: Disaster dynamics of perceived firm-level non-fiscal constraints-related risk. This table reports cross-sectional regressions at the individual firm level. The dependent variable is the change in perceived firm-level non-fiscal constraints-related risk from Q4 2019 to Q3 2020, $\Delta NFCR_f = (NFCR_{f,\text{Q3 2020}} - NFCR_{f,\text{Q4 2019}})/\overline{NFCR}_{\text{Q4 2019}}$, where $\overline{NFCR}_{\text{Q4 2019}}$ is the cross-sectional median in Q4 2019. We present several specifications, where Country$_f$ are firm-level country characteristics ($\log(1 + cases_f)$, $\log(gdppc_f)$, $gdpg_f$), Other$_f$ are firm characteristics including the share of foreign sales, w_f^{for}, leverage, lev, cash holdings, $cash$, and tangibility, ppe, HXZ refers to the Hou, Xue, and Zhang (2015) characteristics, and FF to the Fama and French (2015) characteristics. HQ+Ind and HQ×Ind are fixed effects, where HQ is the headquarter country and Ind the industry classification (two-digit GICS). Accounting data are from the fiscal year 2018. t-statistics (in square brackets) are based on standard errors clustered at the industry level (six-digit GICS). The sample covers 3,043 firms headquartered in 55 countries.

	(1)	(2)	(3)	(4)
$FCE_{f,2018}$	0.018	0.020	0.031	0.237
	[0.094]	[0.107]	[0.160]	[1.105]
Country$_f$ char.	Y	Y	Y	Y
Other$_f$ char.	Y	Y	Y	Y
HXZ char.		Y		
FF char.			Y	Y
HQ + Ind	Y	Y	Y	
HQ × Ind				Y
Obs.	3,043	3,043	3,043	3,043

Table 5: Pre-disaster perceived firm-level fiscal constraints risk. This table reports panel regressions at yearly frequency at the individual firm level. The dependent variable is perceived firm-level fiscal constraints risk scaled by the sample median. We present several specifications, where Country$_f$ are firm-level country characteristics ($\log(gdppc_f)$, $gdpg_f$), Other$_f$ are firm characteristics including the share of foreign sales, w_f^{for}, leverage, lev, cash holdings, $cash$, and tangibility, ppe, HXZ refers to the Hou, Xue, and Zhang (2015) characteristics, and FF to the Fama and French (2015) characteristics. HQ×Year, Ind×Year, and HQ×Ind×Year are fixed effects, where HQ is the headquarter country, Ind the industry classification (two-digit GICS), and Year the fiscal year. t-statistics (in square brackets) are based on standard errors clustered at the industry level (six-digit GICS). The sample period is 2015 to 2019 and covers 2,461 firms headquartered in 57 countries.

	(1)	(2)	(3)	(4)
$FCE_{f,t-1}$	0.561	0.550	0.525	0.566
	[2.191]	[2.115]	[2.070]	[1.944]
Country$_f$ char.	Y	Y	Y	Y
Other$_f$ char.	Y	Y	Y	Y
HXZ char.		Y		
FF char.			Y	Y
HQ × Year	Y	Y	Y	
Ind × Year	Y	Y	Y	
HQ × Ind × Year				Y
Obs.	10,125	10,125	10,125	10,125

Table 6: Discount rates. This table reports panel regressions at yearly frequency at the individual firm level, where the dependent variable is the discount rate firms' use in investment decisions. The discount rate data is from Gormsen and Huber (2023a,b). We present several specifications, where Country$_f$ are firm-level country characteristics ($\log(gdppc_f)$, $gdpg_f$), Other$_f$ are firm characteristics including the share of foreign sales, w_f^{for}, leverage, lev, cash holdings, $cash$, and tangibility, ppe, HXZ refers to the Hou, Xue, and Zhang (2015) characteristics, and FF to the Fama and French (2015) characteristics. HQ×Year, Ind×Year, and HQ×Ind×Year are fixed effects, where HQ is the headquarter country, Ind the industry classification (two-digit GICS), and Year the fiscal year. t-statistics (in square brackets) are based on standard errors clustered at the industry level (six-digit GICS). The sample period is 2015 to 2019 and there are 1,885 firms headquartered in 23 countries.

	(1)	(2)	(3)	(4)
$FCE_{f,t-1}$	0.718	0.740	0.770	0.695
	[1.923]	[2.350]	[2.446]	[2.044]
Country$_f$ char.	Y	Y	Y	Y
Other$_f$ char.	Y	Y	Y	Y
HXZ char.		Y		
FF char.			Y	Y
HQ × Year	Y	Y	Y	
Ind × Year	Y	Y	Y	
HQ × Ind × Year				Y
Obs.	5,535	5,535	5,535	5,535

Table 7: Investment. This table reports panel regressions at yearly frequency at the individual firm level, where in Panel A (B) the dependent variable is CAPX (R&D expenditure) scaled by beginning of year total assets. We present several specifications, where Country$_f$ are firm-level country characteristics ($\log(gdppc_f)$, $gdpg_f$), Other$_f$ are firm characteristics including the share of foreign sales, w_f^{for}, leverage, lev, cash holdings, $cash$, and tangibility, ppe, HXZ refers to the Hou, Xue, and Zhang (2015) characteristics, and FF to the Fama and French (2015) characteristics. HQ×Year, Ind×Year, and HQ×Ind×Year are fixed effects, where HQ is the headquarter country, Ind the industry classification (two-digit GICS), and Year the fiscal year. t-statistics (in square brackets) are based on standard errors clustered at the industry level (six-digit GICS). The sample period is 2015 to 2019 and in Panel A (B) there are 22,826 (8,619) firms headquartered in 82 (63) countries.

Panel A: CAPX

	(1)	(2)	(3)	(4)
$FCE_{f,t-1}$	−0.006	−0.006	−0.006	−0.008
	[−2.755]	[−2.584]	[−2.895]	[−3.918]
Country$_f$ char.	Y	Y	Y	Y
Other$_f$ char.	Y	Y	Y	Y
HXZ char.		Y		
FF char.			Y	Y
HQ × Year	Y	Y	Y	
Ind × Year	Y	Y	Y	
HQ × Ind × Year				Y
Obs.	78,242	78,242	78,242	78,242

Panel B: R&D

	(1)	(2)	(3)	(4)
$FCE_{f,t-1}$	−0.014	−0.012	−0.014	−0.019
	[−2.398]	[−2.361]	[−2.545]	[−2.596]
Country$_f$ char.	Y	Y	Y	Y
Other$_f$ char.	Y	Y	Y	Y
HXZ char.		Y		
FF char.			Y	Y
HQ × Year	Y	Y	Y	
Ind × Year	Y	Y	Y	
HQ × Ind × Year				Y
Obs.	37,189	37,189	37,189	37,189

Table 8: Perceived cost of capital and wedges. This table reports panel regressions at yearly frequency at the individual firm level, where in Panel A the dependent variable is the perceived cost of capital and in Panel B the dependent variable is the wedge between a firm's discount rate and its perceived cost of capital. We use the data provided by Gormsen and Huber (2023a,b) for discount rates and perceived cost of capital estimates. We present several specifications, where Country$_f$ are firm-level country characteristics ($\log(gdppc_f)$, $gdpg_f$), Other$_f$ are firm characteristics including the share of foreign sales, w_f^{for}, leverage, *lev*, cash holdings, *cash*, and tangibility, *ppe*, HXZ refers to the Hou, Xue, and Zhang (2015) characteristics, and FF to the Fama and French (2015) characteristics. HQ×Year, Ind×Year, and HQ×Ind×Year are fixed effects, where HQ is the headquarter country, Ind the industry classification (two-digit GICS), and Year the fiscal year. t-statistics (in square brackets) are based on standard errors clustered at the industry level (six-digit GICS). The sample period is 2015 to 2019 and there are 1,885 firms headquartered in 23 countries.

Panel A: Perceived cost of capital

	(1)	(2)	(3)	(4)
$FCE_{f,t-1}$	0.082	0.127	0.125	0.093
	[0.728]	[1.524]	[1.522]	[1.127]
Country$_f$ char.	Y	Y	Y	Y
Other$_f$ char.	Y	Y	Y	Y
HXZ char.		Y		
FF char.			Y	Y
HQ × Year	Y	Y	Y	
Ind × Year	Y	Y	Y	
HQ × Ind × Year				Y
Obs.	5,535	5,535	5,535	5,535

Panel B: Wedge

	(1)	(2)	(3)	(4)
$FCE_{f,t-1}$	0.636	0.613	0.645	0.602
	[1.931]	[1.958]	[2.070]	[1.833]
Country$_f$ char.	Y	Y	Y	Y
Other$_f$ char.	Y	Y	Y	Y
HXZ char.		Y		
FF char.			Y	Y
HQ × Year	Y	Y	Y	
Ind × Year	Y	Y	Y	
HQ × Ind × Year				Y
Obs.	5,535	5,535	5,535	5,535

Table 9: Pre-disaster CAPM alphas. This table reports panel regressions at the individual firm level using ordinary least squares (EW) as well as weighted least squares (VW), where we weight by the market value of equity in USD. The dependent variable is the yearly β-adjusted excess return (CAPM alpha) in USD from June in year $t-1$ to June in year t. We present several specifications, where Country$_f$ are firm-level country characteristics ($\log(gdppc_f)$, $gdpg_f$), Other$_f$ are firm characteristics including the share of foreign sales, w_f^{for}, leverage, lev, cash holdings, $cash$, and tangibility, ppe, HXZ refers to the Hou, Xue, and Zhang (2015) characteristics (excluding β), and FF to the Fama and French (2015) characteristics (excluding β). HQ×Year, Ind×Year, and HQ×Ind×Year are fixed effects, where HQ is the headquarter country, Ind the industry classification (two-digit GICS), and Year the calender year. t-statistics (in square brackets) are based on standard errors clustered at the industry level (six-digit GICS). The sample period is 2014 to 2019, and covers 22,084 firms headquartered in 81 countries.

	(1)	(2)	(3)	(4)	(5)
$FCE_{f,t-1}$	0.024	0.024	0.025	0.023	0.029
	[3.004]	[3.010]	[3.216]	[3.159]	[2.033]
Country$_f$ char.	Y	Y	Y	Y	Y
Other$_f$ char.	Y	Y	Y	Y	Y
HXZ char.		Y			
FF char.			Y	Y	Y
HQ × Year	Y	Y	Y		
Ind × Year	Y	Y	Y		
HQ × Ind × Year				Y	Y
Weight	EW	EW	EW	EW	VW
Obs.	74,740	74,740	74,740	74,740	74,740

Table 10: CAPM-adjusted drawdowns and time-varying market betas. This table reports cross-sectional regressions at the individual firm. In Panel A and C, the dependent variable is the CAPM-adjusted excess drawdown in USD on trading day 15 after February 21, 2020. We use the pre-disaster beta (i.e. from end of 2019) to compute the adjustment in Panel A, and the realized disaster beta, $\beta_{f,\text{Q3 2020}}$, in Panel C. In Panel B, the dependent variable is the change in market beta from Q4 2019 to Q3 2020, $\beta_{f,\text{Q3 2020}} - \beta_{f,\text{Q4 2019}}$. Control variables and fixed effects are identical to Table 3, but we exclude β from the HXZ/FF controls. The sample covers 21,125 firms headquartered in 79 countries.

Panel A: CAPM-adjusted drawdowns using pre-disaster beta

	(1)	(2)	(3)	(4)	(5)
$FCE_{f,2018}$	−0.050	−0.050	−0.050	−0.050	−0.088
	[−4.103]	[−4.054]	[−4.112]	[−4.262]	[−3.599]
Country$_f$ char.	Y	Y	Y	Y	Y
Other$_f$ char.	Y	Y	Y	Y	Y
HXZ char.		Y			
FF char.			Y	Y	Y
HQ + Ind	Y	Y	Y		
HQ × Ind				Y	Y
Weight	EW	EW	EW	EW	VW
Obs.	21,125	21,125	21,125	21,125	21,125

Panel B: Disaster-induced change in betas

	(1)	(2)	(3)	(4)	(5)
$FCE_{f,2018}$	0.287	0.283	0.283	0.289	0.318
	[5.992]	[5.912]	[5.934]	[5.396]	[3.280]
Country$_f$ char.	Y	Y	Y	Y	Y
Other$_f$ char.	Y	Y	Y	Y	Y
HXZ char.		Y			
FF char.			Y	Y	Y
HQ + Ind	Y	Y	Y		
HQ × Ind				Y	Y
Weight	EW	EW	EW	EW	VW
Obs.	21,125	21,125	21,125	21,125	21,125

Panel C: CAPM-adjusted drawdowns using realized disaster betas

	(1)	(2)	(3)	(4)	(5)
$FCE_{f,2018}$	−0.008	−0.007	−0.007	−0.007	−0.027
	[−1.557]	[−1.486]	[−1.565]	[−1.423]	[−2.826]
Country$_f$ char.	Y	Y	Y	Y	Y
Other$_f$ char.	Y	Y	Y	Y	Y
HXZ char.		Y			
FF char.			Y	Y	Y
HQ + Ind	Y	Y	Y		
HQ × Ind				Y	Y
Weight	EW	EW	EW	EW	VW
Obs.	21,125	21,125	21,125	21,125	21,125

Internet Appendix for

"Fiscal Constraints, Disaster Vulnerability, and Corporate Investment Decisions"

NICOLA MARIA FIORE, THORSTEN MARTIN, and FLORIAN NAGLER *

This Internet Appendix contains supplemental material for the article "Fiscal Constraints, Disaster Vulnerability, and Corporate Investment Decisions." There are two sections:

- Section IA.A provides additional results that complement the main findings.
- Section IA.B provides robustness by using stock prices in domestic currencies.

*Fiore, Nicola Maria, Thorsten Martin, and Florian Nagler, Internet Appendix to "Fiscal Constraints, Disaster Vulnerability, and Corporate Investment Decisions."

IA.A Additional results

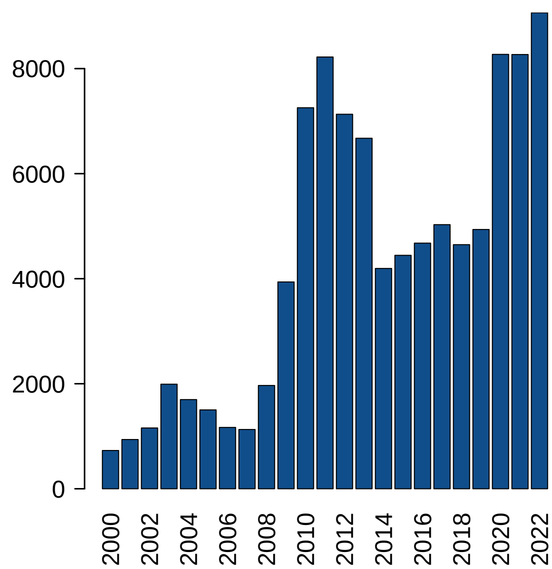

Figure IA.A.1: News articles discussing fiscal constraints-related risk. We use Factiva to search for news articles that discuss risk in the context of fiscal constraints. The search includes articles mentioning synonyms for public debt, fiscal policy, and risk. Specifically, we use the following keywords: [(*government* OR *sovereign* OR *public*) AND *debt*] AND [*fiscal* AND (*spending* OR *policy*)] AND (*risk* OR *uncertainty*). Square brackets indicate that these keywords must appear within ten words of each other. We include all sources, authors, and regions worldwide.

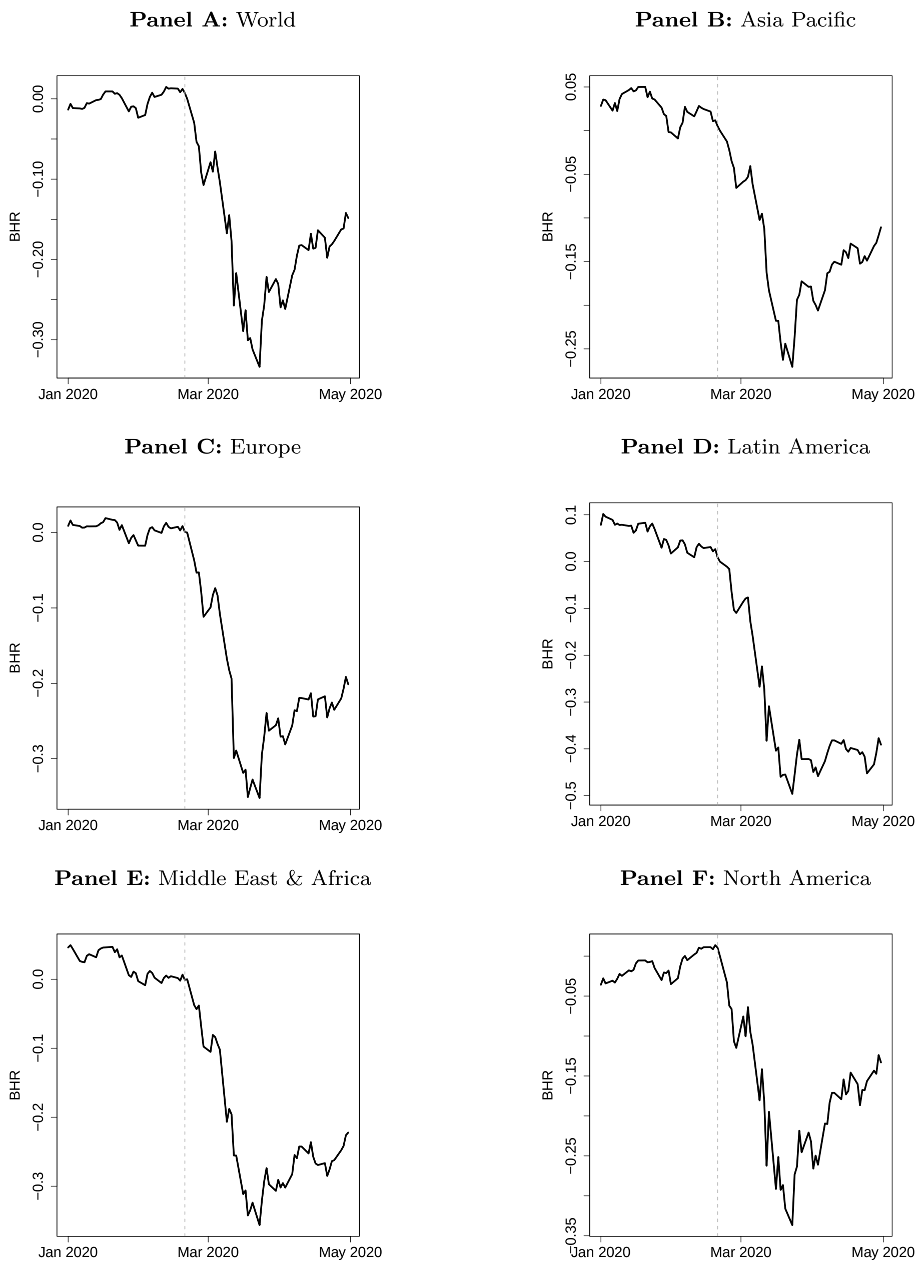

Figure IA.A.2: Disaster-induced stock market reactions around the world. We show FTSE index prices from the beginning of January to the end of April 2020 for the regions indicated above each graph: World, Asia Pacific, Europe, Latin America, Middle East and Africa, and North America. Specifically, we plot buy-and-hold returns relative to February 21, 2020 (dashed line) using daily index prices (*prccd*) from Compustat Global.

Table IA.A.I: Fiscal constraints and disaster government expenditures. This table reports cross-sectional regressions at the country level. The main explanatory variable is a country's c pre-disaster fiscal constraint's measure, FC_c. In columns 1 to 3 the dependent variable is the change in government expenditures scaled by GDP from 2007 to 2009 ($\text{expenditures}_{c,2009}/\text{GDP}_{c,2009} - \text{expenditures}_{c,2007}/\text{GDP}_{c,2007}$), and in columns 4 to 6 the dependent variable is the change in government expenditures scaled by GDP from 2018 to 2020 ($\text{expenditures}_{c,2020}/\text{GDP}_{c,2020} - \text{expenditures}_{c,2018}/\text{GDP}_{c,2018}$). We control for pre-disaster country characteristics, Country_c, that include the logarithm of GDP per capita, $\log(gdppc_c)$, and five-year GDP growth, $gdpg_c$. We further control for the contemporaneous GDP growth_c. t-statistics are shown in square brackets. The sample covers 184 countries.

	Global Financial Crisis			Covid Pandemic		
	(1)	(2)	(3)	(4)	(5)	(6)
$FC_{c,2007}$	−2.065	−2.048	−1.986			
	[−3.353]	[−2.836]	[−2.739]			
$FC_{c,2018}$				−3.380	−2.467	−2.225
				[−4.413]	[−3.081]	[−2.804]
Country_c char.		Y	Y		Y	Y
Cont. GDP growth_c			Y			Y
Obs.	184	184	184	184	184	184

Table IA.A.II: Fiscal constraints and ex-ante firm characteristics. This table reports cross-sectional regressions at the individual firm level, relating ex-ante firm characteristics to firm-level exposure to fiscal constraints, FCE_f. Firm characteristics are the logarithm of market value of equity in USD, $\log(me)$, book-to-market ratio, bm, operating profitability, op, return on equity, roe, leverage, lev, cash holdings, $cash$, and property plant and equipment, ppe. Accounting data are from the fiscal year 2018. The specifications include headquarter country fixed effects. t-statistics (in square brackets) are based on standard errors clustered at the industry level (six-digit GICS). The sample covers 21,125 firms headquartered in 79 countries.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	$\log(me)$	bm	op	roe	$cash$	lev	ppe
$FCE_{f,2018}$	−0.078	−0.017	−0.018	0.016	−0.003	0.005	0.014
	[−0.895]	[−0.245]	[−1.260]	[0.995]	[−0.438]	[0.765]	[1.401]
Obs.	21,125	21,125	21,125	21,125	21,125	21,125	21,125

Table IA.A.III: Perceived firm-level covid risk. This table reports cross-sectional regressions at the individual firm level. The dependent variable is perceived firm-level covid risk of Hassan et al. (2023a) in Q3 2020. We present several specifications, where Country$_f$ are firm-level country characteristics ($\log(1+cases_f)$, $\log(gdppc_f)$, $gdpg_f$), Other$_f$ are firm characteristics including the share of foreign sales, w_f^{for}, leverage, lev, cash holdings, $cash$, and tangibility, ppe, HXZ refers to the Hou, Xue, and Zhang (2015) characteristics, and FF to the Fama and French (2015) characteristics. HQ+Ind and HQ×Ind are fixed effects, where HQ is the headquarter country and Ind the industry classification (two-digit GICS). Accounting data are from the fiscal year 2018. t-statistics (in square brackets) are based on standard errors clustered at the industry level (six-digit GICS). The sample covers 3,035 firms headquartered in 55 countries.

	(1)	(2)	(3)	(4)
$FCE_{f,2018}$	−0.002	−0.005	−0.005	−0.020
	[−0.149]	[−0.398]	[−0.380]	[−1.280]
Country$_f$ char.	Y	Y	Y	Y
Other$_f$ char.	Y	Y	Y	Y
HXZ char.		Y		
FF char.			Y	Y
HQ + Ind	Y	Y	Y	
HQ × Ind				Y
Observations	3,035	3,035	3,035	3,035

Table IA.A.IV: Disaster dynamics of perceived firm-level country risk. This table reports cross-sectional regressions at the individual firm level. The dependent variable is the change in firm-level country risk from Q4 2019 to Q3 2020, $\Delta CR_f = (CR_{f,\text{Q3 2020}} - CR_{f,\text{Q4 2019}})/\overline{CR}_{\text{Q4 2019}}$, where $\overline{CR}_{\text{Q4 2019}}$ is the cross-sectional median in Q4 2019. CR_f is based on data from Hassan et al. (2023b). We present several specifications, where Country$_f$ are firm-level country characteristics ($\log(1 + cases_f)$, $\log(gdppc_f)$, $gdpg_f$), Other$_f$ are firm characteristics including the share of foreign sales, w_f^{for}, leverage, lev, cash holdings, $cash$, and tangibility, ppe, HXZ refers to the Hou, Xue, and Zhang (2015) characteristics, and FF to the Fama and French (2015) characteristics. HQ+Ind and HQ×Ind are fixed effects, where HQ is the headquarter country and Ind the industry classification (two-digit GICS). Accounting data are from the fiscal year 2018. t-statistics (in square brackets) are based on standard errors clustered at the industry level (six-digit GICS). The sample covers 3,012 firms headquartered in 53 countries.

	(1)	(2)	(3)	(4)
$FCE_{f,2018}$	−1.070	−1.050	−1.041	−0.777
	[−1.149]	[−1.124]	[−1.113]	[−0.675]
Country$_f$ char.	Y	Y	Y	Y
Other$_f$ char.	Y	Y	Y	Y
HXZ char.		Y		
FF char.			Y	Y
HQ + Ind	Y	Y	Y	
HQ × Ind				Y
Observations	3,012	3,012	3,012	3,012

Table IA.A.V: Placebo test. This table reports cross-sectional regressions at the individual firm level. The dependent variable is the change in perceived firm-level fiscal constraints risk from Q4 2019 to Q3 2020, $\Delta FCR_f = (FCR_{f,\text{Q3 2020}} - FCR_{f,\text{Q4 2019}})/\overline{FCR}_{\text{Q4 2019}}$, where $\overline{FCR}_{\text{Q4 2019}}$ is the cross-sectional median in Q4 2019. The main explanatory variables are firm-level measures of ex-ante country risk. In Panel A we measure country risk with the monthly exchange rate (relative to USD) volatility over the last 15 years, XR_f, and in Panel B we measure country risk using the volatility of yearly GDP growth over the last 15 years, GR_f. We present several specifications, where Country$_f$ are firm-level country characteristics ($\log(1+cases_f)$, $\log(gdppc_f)$, $gdpg_f$), Other$_f$ are firm characteristics including the share of foreign sales, w_f^{for}, leverage, lev, cash holdings, $cash$, and tangibility, ppe, HXZ refers to the Hou, Xue, and Zhang (2015) characteristics, and FF to the Fama and French (2015) characteristics. HQ+Ind and HQ×Ind are fixed effects, where HQ is the headquarter country and Ind the industry classification (two-digit GICS). Accounting data are from the fiscal year 2018. t-statistics (in square brackets) are based on standard errors clustered at the industry level (six-digit GICS). The sample covers 3,043 firms headquartered in 55 countries.

Panel A: Exchange rate risk

	(1)	(2)	(3)	(4)
$XR_{f,2018}$	−0.001	−0.001	−0.001	−0.001
	[−1.016]	[−1.083]	[−1.067]	[−0.365]
Country$_f$ char.	Y	Y	Y	Y
Other$_f$ char.	Y	Y	Y	Y
HXZ char.		Y		
FF char.			Y	Y
HQ + Ind	Y	Y	Y	
HQ × Ind				Y
Observations	3,043	3,043	3,043	3,043

Panel B: GDP growth risk

	(1)	(2)	(3)	(4)
$GR_{f,2018}$	−0.005	−0.005	−0.006	−0.013
	[−0.316]	[−0.311]	[−0.408]	[−0.743]
Country$_f$ char.	Y	Y	Y	Y
Other$_f$ char.	Y	Y	Y	Y
HXZ char.		Y		
FF char.			Y	Y
HQ + Ind	Y	Y	Y	
HQ × Ind				Y
Observations	3,043	3,043	3,043	3,043

Table IA.A.VI: Summary statistics: stock price sample. This table reports the mean, standard deviation, and 10th, 25th, 50th, 75th, and 90th percentiles of market-β adjusted excess drawdowns in USD (on trading day 15 after February 21, 2020), ADD, firm-level exposure to fiscal constraints, FCE_f, the logarithm of firm-level covid cases, $\log(1 + cases_f)$, the logarithm of firm-level GDP per capita, $\log(gdppc_f)$, firm-level five-year GDP growth, $gdpg_f$, share of foreign sales, w_f^{for}, pre-disaster beta, β, logarithm of market value of equity in USD, $\log(me)$, book-to-market ratio, bm, operating profitability, op, investment-to-assets, i/a, return on equity, roe, leverage, lev, cash holdings, $cash$, and property plant and equipment, ppe. Accounting data are from the fiscal year 2018. The sample covers 21,125 firms headquartered in 79 countries.

	obs	mean	sd	$q_{0.1}$	$q_{0.25}$	$q_{0.5}$	$q_{0.75}$	$q_{0.9}$
ADD	21125	0.084	0.256	−0.201	−0.081	0.049	0.219	0.444
FCE_f	21125	−0.097	0.499	−0.716	−0.403	−0.170	0.152	0.744
$\log(1 + cases_f)$	21125	7.565	2.579	4.210	5.926	7.754	9.305	11.233
$\log(gdppc_f)$	21125	9.848	1.014	8.065	9.320	10.243	10.604	10.902
$gdpg_f$	21125	0.190	0.101	0.060	0.120	0.156	0.292	0.337
w_f^{for}	21125	0.273	0.342	0.000	0.000	0.084	0.501	0.906
β	21125	0.987	0.707	0.153	0.507	0.921	1.426	1.949
$\log(me)$	21125	18.905	2.394	15.807	17.316	18.965	20.551	21.984
bm	21125	1.205	1.748	0.132	0.316	0.683	1.376	2.613
op	21125	0.169	0.411	−0.085	0.053	0.156	0.285	0.505
i/a	21125	0.089	0.297	−0.125	−0.025	0.042	0.140	0.314
roe	21125	0.014	0.380	−0.244	0.000	0.065	0.140	0.244
$cash$	21125	0.167	0.167	0.013	0.043	0.114	0.233	0.401
lev	21125	0.210	0.186	0.000	0.038	0.181	0.329	0.465
ppe	21125	0.276	0.224	0.024	0.088	0.228	0.414	0.617

Table IA.A.VII: Pre-disaster excess returns. This table reports panel regressions at the individual firm level using ordinary least squares (EW) as well as weighted least squares (VW), where we weight by the market value of equity in USD. The dependent variable is the yearly excess return in USD from June in year $t-1$ to June in year t. We present several specifications, where Country$_f$ are firm-level country characteristics ($\log(gdppc_f)$, $gdpg_f$), Other$_f$ are firm characteristics including the share of foreign sales, w_f^{for}, leverage, lev, cash holdings, $cash$, and tangibility, ppe, HXZ refers to the Hou, Xue, and Zhang (2015) characteristics (excluding β), and FF to the Fama and French (2015) characteristics (excluding β). HQ×Year, Ind×Year, and HQ×Ind×Year are fixed effects, where HQ is the headquarter country, Ind the industry classification (two-digit GICS), and Year the fiscal year. t-statistics (in square brackets) are based on standard errors clustered at the industry level (six-digit GICS). The sample period is 2014 to 2019, and covers 22,084 firms headquartered in 81 countries.

	(1)	(2)	(3)	(4)	(5)
$FCE_{f,t-1}$	0.022	0.022	0.024	0.021	0.025
	[2.731]	[2.757]	[2.961]	[2.914]	[1.865]
Country$_f$ char.	Y	Y	Y	Y	Y
Other$_f$ char.	Y	Y	Y	Y	Y
HXZ char.		Y			
FF char.			Y	Y	Y
HQ × Year	Y	Y	Y		
Ind × Year	Y	Y	Y		
HQ × Ind × Year				Y	Y
Weight	EW	EW	EW	EW	VW
Obs.	74,740	74,740	74,740	74,740	74,740

Table IA.A.VIII: Disaster sales. This table reports cross-sectional regressions at the individual firm level. The dependent variable is the change in sales from the pre-disaster to the disaster period, $sales_{f,\mathcal{D}}/sales_{f,\mathcal{N}} - 1$, where $sales_{f,\mathcal{N}}$ are a firm's cumulative sales in the fiscal years 2018 and 2019, and $sales_{f,\mathcal{D}}$ are a firm's cumulative sales in the fiscal years 2020 and 2021. We only use firms with fiscal year-ends in December. We present several specifications, where Country$_f$ are firm-level country characteristics ($\log(1+cases_f)$, $\log(gdppc_f)$, $gdpg_f$), Other$_f$ are firm characteristics including the share of foreign sales, w_f^{for}, leverage, lev, cash holdings, $cash$, and tangibility, ppe, HXZ refers to the Hou, Xue, and Zhang (2015) characteristics, and FF to the Fama and French (2015) characteristics. HQ+Ind and HQ×Ind are fixed effects, where HQ is the headquarter country and Ind the industry classification (two-digit GICS). Accounting data are from the fiscal year 2018. t-statistics (in square brackets) are based on standard errors clustered at the industry level (six-digit GICS). The sample covers 11,631 firms headquartered in 75 countries.

	(1)	(2)	(3)	(4)
$FCE_{f,2018}$	−0.053	−0.054	−0.055	−0.059
	[−1.912]	[−1.960]	[−2.021]	[−2.055]
Country$_f$ char.	Y	Y	Y	Y
Other$_f$ char.	Y	Y	Y	Y
HXZ char.		Y		
FF char.			Y	Y
HQ + Ind	Y	Y	Y	
HQ × Ind				Y
Observations	11,631	11,631	11,631	11,631

Table IA.A.IX: Foreign borrowing and disaster asset pricing. This table reports cross-sectional regressions at the individual firm level. In columns (1) and (2), the dependent variable is the CAPM-adjusted excess drawdown in USD on trading day 15 after February 21, 2020. In columns (3) and (4) the dependent variable is the change in beta from Q4 2019 to Q3 2020, $\beta_{f,\text{Q3 2020}} - \beta_{f,\text{Q4 2019}}$. Betas are based on three-days overlapping returns in USD in a 180-days window, and we use for each quarter the average beta-estimate. We split the sample in two groups: In Panel A we split by actual borrowing from foreign banks (foreign 1/0). In Panel B we split by high (above median) and low (below median) propensity to borrow from foreign banks. Country$_f$ are firm-level country characteristics ($\log(1+cases_f)$, $\log(gdppc_f)$, $gdpg_f$), Other$_f$ are firm characteristics including the share of foreign sales, w_f^{for}, leverage, lev, cash holdings, $cash$, and tangibility, ppe, and FF are the Fama and French (2015) characteristics (excluding β). HQ+Ind are fixed effects, where HQ is the headquarter country and Ind the industry classification (two-digit GICS). Accounting data are from the fiscal year 2018. t-statistics (in square brackets) are based on standard errors clustered at the industry level (six-digit GICS). The sample covers 21,125 firms headquartered in 79 countries.

Panel A: Actual Foreign Borrowing

	Risk-Adjusted Drawdowns		Betas	
	Foreign=1	Foreign=0	Foreign=1	Foreign=0
	(1)	(2)	(3)	(4)
$FCE_{f,2018}$	−0.058	−0.075	0.266	0.252
	[−2.000]	[−2.358]	[2.569]	[2.282]
Country$_f$ char.	Y	Y	Y	Y
Other$_f$ char.	Y	Y	Y	Y
FF char.	Y	Y	Y	Y
HQ + Ind	Y	Y	Y	Y
Obs.	1,022	636	1,022	636

Panel B: Predicted Foreign Borrowing

	Risk-Adjusted Drawdowns		Betas	
	High	Low	High	Low
	(1)	(2)	(3)	(4)
$FCE_{f,2018}$	−0.048	−0.049	0.260	0.328
	[−3.487]	[−2.847]	[5.051]	[3.703]
Country$_f$ char.	Y	Y	Y	Y
Other$_f$ char.	Y	Y	Y	Y
FF char.	Y	Y	Y	Y
HQ + Ind	Y	Y	Y	Y
Obs.	10,552	10,559	10,552	10,559

IA.B Asset prices in domestic currency

Table IA.B.I: Pre-disaster CAPM alphas in domestic currency. This table reports panel regressions at the individual firm level using ordinary least squares (EW) as well as weighted least squares (VW), where we weight by the market value of equity in USD. The dependent variable is the yearly β-adjusted excess return (CAPM alpha) in USD from June in year $t-1$ to June in year t. We present several specifications, where Country$_f$ are firm-level country characteristics ($\log(gdppc_f)$, $gdpg_f$), Other$_f$ are firm characteristics including the share of foreign sales, w_f^{for}, leverage, lev, cash holdings, $cash$, and tangibility, ppe, HXZ refers to the Hou, Xue, and Zhang (2015) characteristics (excluding β), and FF to the Fama and French (2015) characteristics (excluding β). HQ×Year, Ind×Year, and HQ×Ind×Year are fixed effects, where HQ is the headquarter country, Ind the industry classification (two-digit GICS), and Year the fiscal year. t-statistics (in square brackets) are based on standard errors clustered at the industry level (six-digit GICS). The sample period is 2014 to 2019, and covers 22,084 firms headquartered in 81 countries.

	(1)	(2)	(3)	(4)	(5)
$FCE_{f,t-1}$	0.023	0.023	0.024	0.021	0.026
	[2.852]	[2.865]	[3.065]	[2.986]	[1.738]
Country$_f$ char.	Y	Y	Y	Y	Y
Other$_f$ char.	Y	Y	Y	Y	Y
HXZ char.		Y			
FF char.			Y	Y	Y
HQ × Year	Y	Y	Y		
Ind × Year	Y	Y	Y		
HQ × Ind × Year				Y	Y
Weight	EW	EW	EW	EW	VW
Obs.	74,740	74,740	74,740	74,740	74,740

Table IA.B.II: CAPM-adjusted drawdowns and time-varying market betas in domestic currencies. This table reports cross-sectional regressions at the individual firm level using ordinary least squares (EW) as well as weighted least squares (VW), where we weight by the market value of equity in USD. In Panel A, the dependent variable is the CAPM-adjusted excess drawdown in domestic currency on trading day 15 after February 21, 2020. We use the pre-disaster beta (i.e. from end of 2019) to compute the adjustment. In Panel B, the dependent variable is the change in beta from Q4 2019 to Q3 2020, $\beta_{f,\text{Q3 2020}} - \beta_{f,\text{Q4 2019}}$. Betas are based on three-days overlapping domestic currency returns in a 180-days window, and we use for each quarter the average beta-estimate. We present several specifications, where Country_f are firm-level country characteristics ($\log(1 + cases_f)$, $\log(gdppc_f)$, $gdpg_f$), Other_f are firm characteristics including the share of foreign sales, w_f^{for}, leverage, lev, cash holdings, $cash$, and tangibility, ppe, HXZ refers to the Hou, Xue, and Zhang (2015) characteristics (excluding β), and FF to the Fama and French (2015) characteristics (excluding β). HQ+Ind and HQ×Ind are fixed effects, where HQ is the headquarter country and Ind the industry classification (two-digit GICS). Accounting data are from the fiscal year 2018. t-statistics (in square brackets) are based on standard errors clustered at the industry level (six-digit GICS). The sample covers 21,125 firms headquartered in 79 countries.

Panel A: CAPM-adjusted drawdowns using pre-disaster beta

	(1)	(2)	(3)	(4)	(5)
$FCE_{f,2018}$	−0.049	−0.048	−0.049	−0.049	−0.083
	[−3.907]	[−3.862]	[−3.901]	[−4.098]	[−3.359]
Country_f char.	Y	Y	Y	Y	Y
Other_f char.	Y	Y	Y	Y	Y
HXZ char.		Y			
FF char.			Y	Y	Y
HQ + Ind	Y	Y	Y		
HQ × Ind				Y	Y
Weight	EW	EW	EW	EW	VW
Obs.	21,125	21,125	21,125	21,125	21,125

Panel B: Disaster-induced change in betas

	(1)	(2)	(3)	(4)	(5)
$FCE_{f,2018}$	0.287	0.283	0.282	0.290	0.314
	[6.048]	[5.971]	[5.987]	[5.458]	[3.345]
Country_f char.	Y	Y	Y	Y	Y
Other_f char.	Y	Y	Y	Y	Y
HXZ char.		Y			
FF char.			Y	Y	Y
HQ + Ind	Y	Y	Y		
HQ × Ind				Y	Y
Weight	EW	EW	EW	EW	VW
Obs.	21,125	21,125	21,125	21,125	21,125