K Wasn't Built in a Day: Investment with Endogenous Time to Build

Adriano Fernandes* Rodolfo Rigato †

Job Market Paper

November 8th, 2023

Click here for the latest version

Abstract

Physical capital takes time to build. Yet, the measurement of time to build and of its response to firm behavior remain scant. We fill this gap using project-level data from India. We first document new facts about time to build. Industry heterogeneity accounts for 30% of its variation; and time to build increases on average by 0.18% for each 1% increase in project cost. We exploit quasi-experimental variation in credit supply to document that firms have control over time to build. When credit dries up, the conditional probability of completing a project over the following quarter rises by 6%, consistent with firms accelerating project development. In turn, new project starts fall by 7.5%. To rationalize our findings, we introduce a model of endogenous time to build. A credit crunch increases firm appetite for immediate cash flows relative to delayed cash flows. Firms then accelerate existing, closer to completion projects and postpone unbegun projects. Such a mechanism is borne out in the data: projects proxied to be more mature are sped up the most. We quantify endogenous time to build by calibrating our model to match our causal estimates, and the joint distribution of project-level costs and gestation lags. Moving from exogenous to endogenous time to build amplifies the response of investment to shocks, increasing investment volatility by up to 30%. Endogenous gestation lags are policy relevant. Monetary policy is more potent when the distribution of projects along their gestation cycle skews towards mature projects. Fiscal policy, in turn, can flexibly reshuffle investment expenditures over time with tax credits.

*Harvard University. Email: adrihfernandes@g.harvard.edu.

†European Central Bank. Email: rodolfo.dinis_rigato@ecb.europa.edu.

We are indebted to our advisors Gabriel Chodorow-Reich, Samuel Hanson, Jeremy Stein, and Ludwig Straub for their invaluable support. We are grateful to Robert Barro, Namrata Narain, Adi Sunderam and Elie Tamer for continuous encouragement and suggestions; as well as Pol Antràs, Patrick Bajari, Adrien Bilal, Michael Blank, Emily Breza, John Campbell, Veronica De Falco, Pedro Degiovanni, Martin Eichenbaum, Ester Faia, Xavier Gabaix, Shresth Garg, Edward Glaeser, Daniel Green, Myrto Kalouptsidi, Stephanie Kestelman, Greg Lewis, Greg Mankiw, Ross Mattheis, Max Miller, Ariel Pakes, Dev Patel, Louise Paul-Delvaux, Marco Sammon, Audrey Tiew, Boris Valleé, Davide Viviano, Christian Wolf, Lingxuan Wu, and numerous seminar participants for their helpful insights. We thank Samuel Efraim, Harold Klapper and João Menegotto for stellar research assistance. Adriano Fernandes recognizes financial support from the Molly and Domenic Ferrante Economics Research Fund on Monetary and Fiscal Policy (Harvard Economics Department), and from the Center for International Development (Harvard Kennedy School). The views expressed in this paper do not represent the views of the European Central Bank. Any errors are our own.

1 Introduction

Physical capital takes time to build. Anecdotal evidence abounds: ships, power plants, roads, factories, sewers, the list goes on. More than an anecdotal fact, time to build is a defining feature of capital formation. It is entwined with investment decisions. A firm experiencing a period of high demand may benefit from building a new plant. Whether to invest depends on whether demand will remain high by the time the plant is finished. Gestation lags shape investment. But what if the firm was already building a plant when demand picked up? Could the firm *choose* to speed up construction, even if costly, to take advantage of good times? Or is time to build an exogenously given technological constant that firms must simply accept?

The central assertion – and empirical finding – of this paper is that time to build is an *endogenous choice* of firms. Investment shapes gestation lags. We study the resulting consequences for capital budgeting, investment fluctuations and economic policy.

Despite its economic importance, there is relatively little work measuring time to build. Earlier work leveraged small samples or aggregated data. Our first contribution is to fill this gap. We measure gestation lags in a wide cross-section of projects using microdata from India and provide a battery of novel stylized facts.

Our second contribution is to provide causal evidence that time to build is a choice variable. We show that firm-level discount rate shocks affect how quickly firms complete ongoing projects. Under plausible assumptions, our findings refute an ubiquitous assumption of investment models: that gestation lags are purely exogenous, technological primitives.

Our third contribution is to revisit investment theory with endogenous time to build. We build a structural model where firms *choose* project gestation lags. The model delivers analytical insights on optimal capital budgeting, and rationalizes our empirical findings. Chiefly, it yields an additional empirical prediction: projects react more to a discount rate shock as they approach completion. We test and find support for such prediction in the data. Using our measurements and estimated casual effects, we calibrate our model. Quantitatively, investment dynamics become more responsive to shocks, and thus more volatile; and differ sharply across discount rate versus cash flow shocks.

Our fourth contribution is to study stabilization policies in a world of endogenous time to build. We show that the effects of monetary policy on investment are state-dependent. Since later-stage projects respond more to interest rate changes, the response of aggregate investment is stronger when the distribution of projects along their gestation cycle skews towards more mature projects. Second, temporary investment tax credits front-load investment flows, allowing fiscal policy to flexibly reshuffle aggregate demand over time.

We start with measurement. Our data on large capital expenditure projects comes from CapexDx (CMIE, 2023a). Projects in our data include steel plants, irrigation canals, airport terminals, office and residential buildings, among other structures. The data covers projects in India above 10 million rupees;[1] and focuses on new capacity creation. In contrast with prior work, we observe microdata on both public and private projects, across a wide cross-section of end-user industries, and over different points of project lifecycle.[2]

A project's event history allows us to measure its time to build. The distribution has long right tails: the median (average) gestation lag is 2.5 (3.5) years, the 90th percentile is 10 years, and the 95th percentile is over 15 years. The aggregate distribution masks sizable heterogeneity. We decompose the variation in time to build: industry accounts for 30% of the variation; firm heterogeneity accounts for 50% at most. We thus take industry as a parsimonious proxy for shared technological primitives across projects, aware of the sizable residual variation within industry. Project size does not account for much additional variance, but is tightly linked to gestation lags: a 1% increase in project size increases a project's time to build by 0.18%. Though the slope varies across industry and ownership cells, the log-linear relationship is consistent across and within-cells.

Saving notable exceptions (Lewis and Bajari, 2011, 2014), prior work assumed time to build was an exogenous constraint that project sponsors had to accept. A competing hypothesis is that firms partially control how fast they build projects. We causally establish that the latter view is borne out in the data. We explore a setting where a firm suffers a shock that is plausibly orthogonal to the daily affairs of project construction. Under a sharp null of exogenous gestation lags, the firm-level shock should have no bearing on project completion.

We study a shock to credit supply that arguably raised firm discount rates. Following large-scale audits by the Reserve Bank of India starting in 2015, commercial banks had to set aside more provisions for non-performing assets, potentially reducing lending capital. We use loan-level data from the Indian Ministry of Corporate Affairs to confirm that the banking shock reduced outstanding credit (Chopra et al., 2020). Controlling for demand-side confounders, firms relying on more exposed banks borrowed less. We take firm exposure to the credit supply disruption as our firm-level shock; and link projects to their sponsoring firms using data from Prowess(CMIE, 2023b).

We embed our shift-share design into a discrete-time duration model. Identification relies on exogeneity of the bank-level shocks relative to project and firm-level unobservables, conditional on controls (Borusyak et al., 2021; Chodorow-Reich et al., 2021). Previous evidence supports quasi-randomness of the shock. We further rely on a detailed set of fixed effects and firm-level controls to minimize confounding concerns. Our main finding is that firms more

[1] This amounts to 150 to 300 thousand USD depending on the foreign exchange rate. The rupee experienced a pronounced devaluation over the sample years.

[2] Fixed capital investment is particularly crucial for developing and emerging market economies, but project-level measurement exercises with such wide-coverage as ours are non-existent for such countries, to the best of our knowledge. Moreover, India has been and will continue to be a global powerhouse of capital information and economic growth (Economist 2023).

exposed to the shock speed up ongoing projects relative to firms less exposed to the shock. The intensive margin result is consistent with the view that time to build is an endogenous choice. On the extensive margin, we find that firms start fewer projects following the shock, consistent with standard investment theory. To be precise, a one-standard deviation increase in firm-level exposure raises the completion hazard by 6% and lowers project starts by 7.5% in our preferred specification.

To move beyond local elasticities, we write down a model of a firm that dynamically invests in a collection of heterogeneous, irreversible projects. Each project takes time to build, and only generates cash flows once fully completed. Crucially, gestation lags are endogenous: firms can speed up project completion, but incur ever increasing acceleration costs to do so. At each point in time, firms optimally choose how much to invest on ongoing projects; and whether to start new projects. Firms benefit from completing projects faster because they discount future cash flows. Optimal investment in ongoing projects equates the marginal returns from earlier completion to the marginal costs of acceleration. This intensive margin tradeoff pins down a project's remaining time to build. On the extensive margin, firms start new projects with positive net present values, internalizing optimal gestation lags. The model is minimalist and portable, lending itself to further applications and extensions.

How does the model match our empirical findings? We first take an analytical approach. We map the banking shock (data) to an increase in firm discount rates (model). Up to first order, we show that projects close to completion are unambiguously sped up, whereas just started projects tend to be slowed down relative to steady state.

We first show that changes in project values following the discount rate shock are given by a duration formula, a familiar fixed-income concept. Intuitively, duration gives us how long it will take to reap back the cash flows of the project. As project value changes, so does investment. We then show that this investment response is proportional to the marginal value of advancing the project, which in turn is given by how much duration falls as projects mature. Intuitively, reducing duration is akin to reaping positive cash flows earlier.

So how do incentives to invest change after an increase in discount rates? There are two opposing effects. First, investors become more impatient, pushing them to accelerate projects. Second, as discount rates rise, project values fall, since we discount future cash flows more heavily. Intuitively, however, we show that the latter effect becomes dominated as the project approaches completion: its duration is smaller, because there is "less time" for the discount rate to reduce its NPV. For mature projects, investor impatience dominates and projects are sped up. As residual time to build increases, so does project duration. Such long-lived projects suffer dramatically from increases in discount rates. The duration effect becomes meaningful, and might even overweight the impatience effect, in which case it becomes optimal to slow down the project.

Our theory provides an additional, testable prediction: projects closer to completion should react more to the shock relative to projects further away from completion. Arguably, this is the key unconditional prediction that emerges from the model. Evidence of such heterogeneous effects would provide strong support for the theory. Indeed, effect heterogeneity is borne out in the data: projects proxied to be *closer* to completion are sped up by *more* relative to earlier stage projects. These sharper reduced form tests thus back up the model's central mechanism.

The opposing intensive and extensive margin results qualify the standard intuition that capital expenditures necessarily decrease following adverse shocks to firm financial health. The project-level effect depends on how close the project is to completion. At the micro level, capital formation may very well rise. In turn, aggregation shapes the total effect, as we average the positive effects on near-completed projects and the negative effects on nascent or unbegun projects. The distribution of projects over their completion stages matters for aggregate investment.

We take the model to the data. We introduce heterogeneity in project returns, sizes, and construction technologies; and match our causal estimates by computing model counterparts to our intensive and extensive margin regressions. We combine all steps to jointly calibrate model parameters.

We use our calibrated model to dissect investment dynamics. We feed discount rate shocks into our model, and compare fluctuations under endogenous versus exogenous time to build, keeping other prices fixed (partial equilibrium). Endogenous time to build makes investment dynamics more responsive to shocks, and thus more volatile: moving from exogenous to endogenous time to build increases investment volatility by 30%. The total effect hides sizable compositional effects. When time to build is exogenous, investment volatility results exclusively from variation in project starts. When time to build is endogenous, dynamics are instead driven primarily by adjustments in gestation lags, with a more limited role for project starts.

Endogenous time to build generates *state-dependence*: the distribution of projects along completion stages affects the aggregate investment response. Since more mature projects react more to shocks, the aggregate response is stronger when the distribution is shifted towards such later-stage projects. We thus propose a novel source of state-dependence in the effects of monetary policy – one operating through the investment channel and, in particular, through the distribution of gestation lags; but with the opposite effect that we normally ascribe to discount rate shocks.

Last, we study cash flow shocks. When returns to completion temporarily rise, investment on ongoing projects spike, as firms race to reach a mean-reverting jackpot. Fiscal policy instead frequently relies on longer-lived investment tax credits. We study such policies by introducing a temporary investment subsidy. Following the introduction of a tax credit, a standard model with exogenous TTB predicts an increase in investment only due to the increase in deployment of new projects, what we call the extensive margin channel. The model with endogenous TTB, on the other hand, introduces an intensive margin channel, corresponding to the front-loading of investments before the expiration of the tax credit. From a stabilization perspective, investment tax credits shuffle expenditure intertemporally, thereby providing flexibility to the stabilization toolkit. Beyond stabilization policy, investment subsidies encourage entry, generating long-lived positive effects on the capital stock. They can thus be valuable policies to speed up structural change, such as the green transition.

Taking stock, our quantitative results suggest that models relying on ad hoc aggregate adjustment costs are subject to a Lucas critique. Discount rate and cash flow shocks generate impulses of different shapes, so that adjustment costs become *shock-dependent*. The effect of these shocks in turn vary with the underlying distribution of projects, which depends – even in steady state – on policy stances (e.g. the interest rate set by the monetary authority). Adjustment costs thus become *state-dependent*. Counterfactual policy statements arising from models with ad hoc adjustment costs may therefore yield inappropriate conclusions. In turn, endogenous time to build provides an empirically plausible microfoundation for such costs, and can be readily embedded in models of firm, industry, or aggregate dynamics.

Relation to Literature. Our paper contributes to four strands of literature. First, a small body of work measures gestation lags, but relies on either small samples (Mayer, 1960; Brooks, 2000) or aggregated data (Montgomery, 1995). We contribute to this literature by (i) relying on microdata from a wide cross-section of industries, sponsors, project sizes, and (ii) studying an emerging market.

Second, we add to the literature on the interplay between gestation lags and investment. Closest to us are (Lewis and Bajari, 2011, 2014), who document that contractors accelerate road construction to avoid procurement penalties. We are unaware of any other paper providing econometric evidence of endogenous gestation lags. Kalouptsidi (2014), Greenwood and Hanson (2014), and Oh and Yoon (2020) study the role of exogenous gestation lags in driving entry decisions, return predictability, and boom-bust dynamics in specific industries. Relative to such work, we (i) provide causal evidence on endogenous time to build in a wide cross-section of industries, thereby bolstering external validity of the mechanism; and provide a portable, microfounded model of optimal gestation lags.

Third, we relate to the vast literature on real effects of credit supply disruptions. [3] We are unaware of papers documenting a causal effect on gestation lags as we do. Further, we qualify the conventional wisdom that tighter credit necessarily reduces investment: the level of aggregation and the degree of project completion crucially influence such prediction.

[3]Methodologically, we borrow from Khwaja and Mian (2008), Chodorow-Reich (2013) and Chopra et al. (2020). Thematically, a flood of evidence documents that real activity contracts following credit supply shocks. A non-exhaustive list includes earlier work by Bernanke and Gertler (1995), Peek and Rosengren (2000), and Ashcraft (2005); and more recent work by Pinardon-Touati (2023).

We buttress econometric evidence with a theoretical model that carries practical lessons for optimal capital budgeting.

Last, we converse with a long intellectual history of time to build in theories of macroeconomic fluctuations.[4] Whether gestation lags matter for fluctuations has been debated since the Rouwenhorst (1991) critique of Kydland and Prescott (1982). Casares (2006), Edge (2007) and Lucca (2007) challenged the critique: time to build had important implications for monetary transmission.[5] Leeper et al. (2010) and Ramey (2020) argue that the potency of fiscal multipliers varies with gestation lags of public investment.[6] We push this literature in a new direction. We emphasize that endogenous gestation lags and the ensuing distribution of projects over their completion lifecycle reshape investment dynamics; and provide lessons for both fiscal and monetary policy.

The rest of the paper elaborates on our discussion so far. Section 2 describes the data; section 3 contains our measurement exercises; section 4 presents the econometric methodology and reduced-form results; section 5 details our model and analytical results; Section 6 explains our calibration; Section 7 studies investment dynamics and policy implications; Section 8 concludes.

2 Data

We obtain data on large capital expenditure projects from CapexDx, a dataset collected and maintained by the Center for Monitoring of the Indian Economy (CMIE). The data provider defines a project as a construction enterprise where "there is reasonable clarity on the capacity being created". Examples of projects include steel plants, irrigation canals, airport

[4]The hypothesis that gestation lags are endogenous to the economic environment dates back to Böhm-Bawerk (1889), and has been resurrected by Antras (2023) in the context of global value chains. Pigou (1920), Aftalion (1927) and Kalecki (1935) proposed that gestation lags were essential ingredients for theories of investment fluctuations. Hall (1977) and Taylor et al. (1982) spoke for gestation lags as microfoundations for investment theories relying on ad hoc adjustment costs or acceleration mechanisms. Asea and Zak (1999) proved that the neoclassical growth model with exogenous gestation lags predicted oscillatory fluctuations around a steady state, very much in line with early-century theories.

[5]All such papers share a common lesson: gestation lags matter for aggregate investment dynamics if they are heterogeneous across different, complementary capital stocks. Lucca (2007) further proved that a specific formulation of exogenous gestation lags was first-order equivalent to the investment adjustment costs of Christiano et al. (2005) – an ubiquitous ingredient of monetary DSGE models, essential to match the effects of monetary shocks on aggregate investment. In a similar vein, Bloesch and Weber (2023) provided causal evidence for time to onboard as a promising avenue to explain the response of R&D investment to monetary shocks. Christiano and Todd (1996), Wen (1998), Zhou (2000) and Christiano and Vigfusson (2003) also defended the role of gestation lags, but relied more on macroeconometric evidence in single-K models.

[6]Meier (2020) argues that time-series variation in gestation lags due to supply chain disruptions yields GDP and TFP losses. Liu and Tsyvinski (2023) introduce adjustment costs for intermediate inputs in a dynamic input-output network model, and show that such costs are first-order equivalent to exogenous gestation lags. Both papers measure gestation lags with the backlog ratio. Such measure is appropriate for inventories, intermediate inputs, or smaller equipments, but inadequately captures gestation lags for large structures. We interpret their work as informative on the role of backlogs in intermediate inputs rather than gestation lags in capital formation.

terminals, office and residential buildings, etc. In the parlor of National Income Account systems, projects in our data are "structures".[7]

CapexDx collects data on project name, cost, geographical location, and industry. Crucially, it tracks the lifecycle of projects within a rich event-level dataset. To quote CMIE, CapexDx "is a monitoring machinery that records the various events that a project goes through in its life-cycle." A project's event history allows us to measure time to build. We observe when the sponsor announced the project; when it started or stalled implementation; and, in most cases, when it completed or abandoned the project. For some projects, histories are censored: we do not know what becomes of the project. We will leverage econometric methods that account for right-censoring. Further, the data provider tracks its monitoring efforts and follows transparent rules when discontinuing project coverage. We elaborate on project definition rules and data collection in appendix A.1.

CMIE has collected data on projects since its inception in 1976, but has systematized its efforts from the mid-1990s onwards. Accordingly, the bulk of our projects started in the latter period. The data covers projects above 10 million rupees (150-300 thousand USD, depending on the foreign exchange rate over our sample period); and focuses on new capacity creation, with some cases of substantial expansions or modernizations of existing structures. In contrast with prior work using project-level data (Mayer, 1960; Brooks, 2000), we observe both public and private projects,[8] across a wide cross-section of end-user industries, and over different points of project lifecycle.

Another differentiating feature of CapexDx is the ability to link projects to sponsoring firms. We use CMIE's ProwessDx to obtain information on firms and their banking partners, mostly financial statements. The data is sourced from regulatory filings with the Indian government and is *not* restricted to private, publicly-traded companies. Alfaro and Chari (2014) (among other papers) provide extensive descriptions of the data.

Last, we collect loan-level data from the Indian Ministry of Corporate Affairs (MCA) website. We observe secured loans registered by lenders with the MCA. The regulatory framework incentivizes lenders to register charges against secured loans, as failing to register may result in material adverse consequences.[9] The data thus approximates a credit

[7]Systems of National Accounts usually divide "fixed assets" (the K of macroeconomic models) into structures, equipment, and, more recently, intellectual property. Our data covers structures. We do not observe replacements or maintenance of existing equipments (e.g. computers), as in the ACES data collected by the U.S. Census Bureau; nor do we observe capacity or capacity utilization, as in the ASM data collected by the U.S. Census Bureau.

[8]Sponsors in our data include government entities, foreign and Indian private firms, and (minority or majority) state-owned enterprises. Therefore, not all projects in our data are sponsored by firms. We will use the terms promoter, sponsor and firms interchangeably.

[9]In cases of borrower restructuring, a secured lender with a registered charge gets priority over a secured lender who failed to register a charge. In cases of borrower liquidation, unregistered secured lenders are treated as unsecured lenders. Failing to register a secured loan effectively nullifies the benefits of secured lending, thereby incentivizing creditors to register their charges. See Chopra et al. (2020) for further details.

registry for secured loans, and allows us to link the financial health of firms' banks to the performance of firms' projects. For each loan, we observe the amount borrowed, the date of the transaction, as well as borrower and lender name. We do no t observe interest rates, contractual maturity, covenants, or any other contractual details on the loan.

We link borrowers in the loan-level data to CMIE datasets by using the unique Company Identification Number of a firm. These identifiers are present in both datasets, and are assigned by the Indian government, facilitating the matching process. Linking lenders to the CMIE data is more challenging. We can only rely on lender names, which are often misspelled. Using a combination of fuzzy matching algorithms and extensive manual checks, we create a crosswalk between lender names in the MCA data and company names and identifiers in the CMIE data. We are able to match over 90% of our loans (both in terms of loan volume and number of charges) to a Prowess company, thereby identifying the majority of entities lending to project sponsors.

3 Deconstructing Time to Build: Novel Facts

We first describe the data. Panel A of Table 1 presents summary statistics on the cross-section of projects. We observe over 100,000 projects sponsored by over 30,000 entities. Many firms sponsor more than one project, with the distribution featuring heavy right tails. The average project was first observed in 2010, with the earliest projects entering the sample in the mid-1990s. We observe cost data for 70% of projects. The distribution of costs is right-skewed: while the median project costs 503 million rupees, the average project costs over 6 billion rupees, and projects above the 90th percentile cost at least 10 billion rupees.[10] Projects in our data are thus sizable construction enterprises.

A distinctive feature of the data is its broad coverage across industries, geographies and sponsors. We compute HHI indexes to measure whether projects seem to be concentrated in particular industries or states, but find median HHIs of 0.02 and 0.07 at the industry- and state-levels. Public sponsors account for almost 40% of the projects.[11] Importantly, projects in our sample are either new units (81%), large expansions (15%) or modernizations (4%) of existing units. They are not simply acquisitions or other forms of secondary market transactions. We thus measure actual gestation lags, not measures of shipping time or backlog.

Panel B provides an overview of projects' lifecycles. To do so, we construct project histories from event-level data.[12] All projects in our data start when they are announced. The

[10]The median (average) project costs 7 million (86 million) USD, and the right tail features projects over 140 million USD. We assume a nominal exchange rate of 70 rupees per dollar. Given the large nominal devaluation of the rupee since the GFC, these numbers likely provide a conservative estimate of project costs in USD.

[11]An important feature of the Indian economy is the presence of (partially or fully) state-owned enterprises. Public sponsors thus do not correspond exclusively to government enterprises.

[12]Our diary dataset records multiple events that happen over a project's lifecycle. Often, more than one event

Table 1: Summary Statistics: All Projects

	Mean	SD	10th	25th	50th	75th	90th
Panel A: Cross-Sectional Characteristics							
Projects per Sponsor	113.05	257.90	1.00	2.00	11.00	77.00	366.00
Year Started	2010	8	1997	2005	2011	2017	2020
Last Observed	2014	8	2001	2009	2016	2021	2022
Cost (Rs. Million)	6428.76	40583.47	46.00	125.70	503.30	2500.00	10000.00
Missing Cost	0.29	0.45					
HHI: Section	0.19	0.11	0.02	0.07	0.23	0.31	0.31
HHI: Industry	0.04	0.04	0.00	0.01	0.02	0.04	0.12
HHI: State	0.09	0.07	0.01	0.03	0.07	0.13	0.20
New Unit	0.81	0.39					
Large Expansion	0.15	0.35					
Private	0.57	0.50					
Public	0.38	0.49					
Panel B: Lifecycle Dynamics							
Ever in Status							
Under Imp	0.71	0.46					
Stalled	0.06	0.23					
Completed	0.40	0.49					
Shelved or Abandoned	0.05	0.21					
No Information	0.31	0.46					
Time in Status (Qtrs)							
Announced	6.35	9.41	0.04	0.92	3.32	8.10	15.78
Under Imp	8.74	13.13	0.00	0.00	4.51	12.12	23.62
Stalled	0.69	4.34	0.00	0.00	0.00	0.00	0.00
Number of projects	106626						
Number of firms	32456						

Notes: This table reports summary statistics on all projects available in the March 2022 vintage of CapexDx. Panel A reports cross-sectional characteristics. HHIs are computed within a particular sector, industry, or state, and then assigned to projects. Panel B reports lifecycle characteristics. To do so, we leverage our panel dataset tracking project histories. "Ever in Status" variables are dummies that take value 1 if the project was ever in that stage. They need not sum to one, as a project may transit through multiple stages over its lifecycle. "Time in Status" variables measure the total amount of time spent in a particular stage over the project's entire observable history. We only include non-terminal stages, i.e. we exclude completion, abandonment, and loss to follow-up.

median project spends 3 quarters in this phase. Most projects (70%) move to implementation, but, rarely, some are stalled or abandoned even before implementation begins.

Stalling is a temporary state: the project is put on hold, but there is the intent to continue the project after some period. Abandonments are terminal stages: the sponsor decisively states it has no intention of continuing the project.[13] We find that stallments or abandonments are rare in our data: 6% of projects are ever stalled, and 5% of projects are ever abandoned. Further, most projects do not remain stalled for over a quarter.

The median project spends 4.5 quarters under implementation. At any point during its construction, a project may be stalled, abandoned, or completed. The latter account for 40% of all projects in our data (and roughly 60% of projects that move into implementation). This figure might seem striking given the low number of abandonments. What is happening to the other projects? Some projects are just still under implementation to this day. Our data ends in March 2022, so these projects are just right-censored observations. The same can be said about projects that were recently announced, but had not moved to implementation as of March 2022.

For 30% of all projects in the data, we simply do not know what has become of them. To borrow from the epidemiology literature, they are "lost to follow-up", right-censored observations.[14] Reassuringly, we understand part of the data generating process for censoring — at least that under the discretion of CMIE. CapexDx discloses it stops tracking projects "if they are stalled or shelved or abandoned or information regarding them is not available for over 30 months in the case of manufacturing projects and for over 36 months in the case of infrastructure projects." [15]

is recorded in a given date. For instance, there might be three entries in a particular date – (i) sponsor contacted successfully, (ii) construction commenced, and (iii) foundation stone laid. Event (i) is informative that the project is being actively monitored, but does not provide us any information about the project's ongoing construction. It is simply an informational marker. Events (ii) and (iii), however, provide critical information on the project's lifecycle. We learn that the project moved into implementation. Multiple difficulties arise in providing a consistent reconstruction of a project lifecycle from event-level data, but it is essential to check that project histories are "sensible", and hence that our time to build and hazard regression estimates are not driven by measurement error. We elaborate on our extensive data cleaning in Appendix A.1.

[13]In some exceedingly rare cases, shelved or abandoned projects are "resuscitated". These zombie projects are special cases where a sponsor decisively states that it has abandoned or shelved a project. But later defaults on its previous statement and resuscitates the project. CapexDx records such events as abandonments, despite their temporary nature. For expositional purposes, we group them with stalled projects. Our results are insensitive to such grouping.

[14]The nature of the term originates from medical studies, where patients are often seen for a number of visits, but then drop out of the study. The parallel is that we always observe a project start, and often observe some intermediate events, but eventually cannot "follow-up" with the project. This right-censoring is conceptually different in nature from administrative censoring, as in the case of projects for which we don't know what happen simply because data collection stops on March 2022.

[15]Right-censoring is not a one-way street: CMIE resumes tracking previously censored projects if it manages to obtain reliable information on project status from the public domain. In short, the combination of intense monitoring and transparent censoring schemes bolsters data reliability. The 30% figure we report is after we clean project histories to exclude periods of temporary loss to follow-up.

Appendix Table A1 tabulates the frequency of transitions across different stages, excluding no transition events. Most announced projects move into implementation, and most projects under implementation are eventually completed. About 20% of announced projects, 6% of stalled projects, and 27% of projects under implementation are lost to follow-up. Appendix Table A2 instead estimates quarterly transition probabilities between stages including transitions to itself, and shows that intermediate stages are persistent.

Measuring Gestation Lags We take a broad view of time to build, and measure a project's gestation lag from its announcement to its eventual completion. For such, we restrict the analysis to projects that are eventually completed. Our results should thus be interpreted under the possibility that completed projects systematically differ from abandoned or censored projects along predictable dimensions; but this potential selection bias does not invalidate our efforts to provide a first estimate of gestation lags.

Figure 1 plots the distribution of gestation lags for all projects in our data. The median project takes 2.5 years to be completed. The distribution is notably right-skewed. While most projects are completed in a few years, a sizable share of projects take over 15 years to be completed (the plot winsorizes the histogram for visualization). To the best of our knowledge, we are the first to provide such a comprehensive, microdata-based estimate of gestation lags; and the first to provide any estimates of gestation lags for emerging economies.

The aggregate distribution hides considerable cross-sectional heterogeneity. Three key drivers of heterogeneity come to mind. First, structures vary by industry. Building a railway, a warehouse for a retailer, a manufacturing plant, or a sewage system are each very different endeavors. Second, certain projects are more likely to be undertaken by private firms, while others rely on public sponsorship. Third, projects differ in size, and larger projects presumably take longer to build.

Panel (a) of Figure 2 breaks down the distribution at the sector-level. To facilitate visualization, we only mark the median gestation lag with interquartile range bands; and exclude water and sanitation, an outlier with long gestation lags.[16] The usual suspects come up on top: transportation, construction, and other infrastructure-related sectors feature the longest gestation lags. In turn, retail, information technology, and manufacturing feature the shortest gestation lags.

Panel (b) of Figure 2 in turn summarizes time to build based on whether the sponsor is a public entity, a private Indian entity, or a foreign private entity. The differences are striking: public sponsors build projects with longer gestation lags. This is partly driven by industrial composition: the public sector frequently sponsors heavy infrastructure projects, for which social returns are higher than private returns. For example, Indian Railways is the largest

[16]Since the means are precisely estimated, we do not plot standard error bands – it's easy to pairwise reject equality of means across many sectors.

Figure 1: Distribution of Gestation Lags – Completed Projects

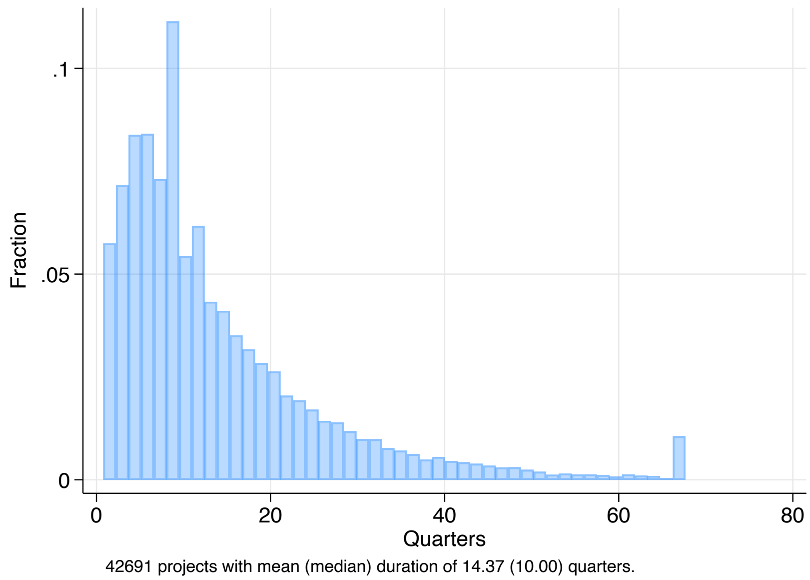

Notes: the figure plots the histogram of gestation lags for all completed projects in the Mar 2022 CapexDx vintage. We measure a gestation lag from the first date that a project appears in the data (their announcement date) to its eventual completion. We winsorize the distribution at the 1st and 99th percentile for visualization (but report the mean and median of the raw distribution).

Figure 2: Disaggregating the Time to Build Distribution

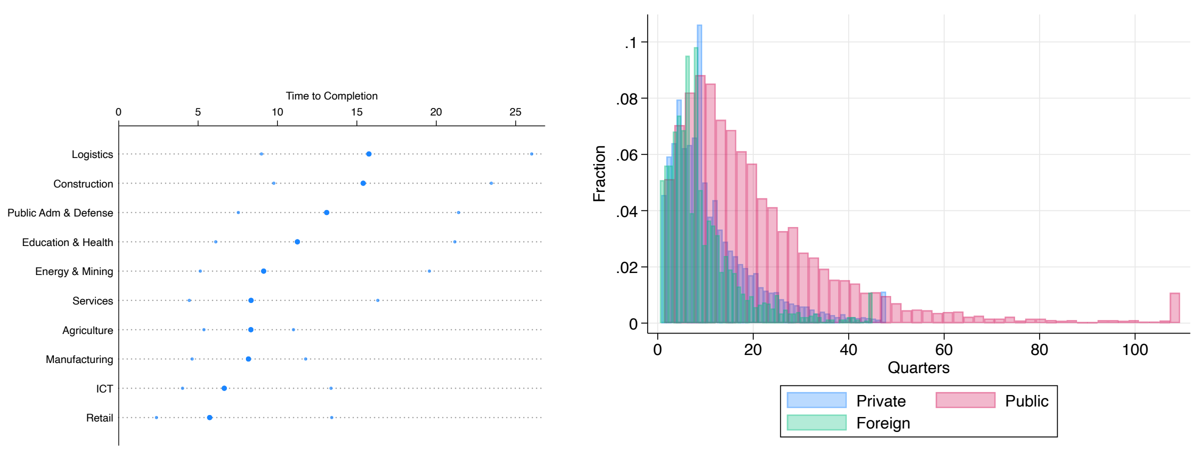

(a) Medians and IQR of Gestation Lags by Sector

(b) Distribution of Gestation Lags by Sponsors

Notes: the figures disaggregate the distribution of gestation lags across sectors (panel a) and public vs private ownership (panel b). Panel (a) plots median and interquartile ranges for gestation lags by sector. Sectors are based on a slightly aggregated version of India's NIC-defined sectors, and excludes water, sanitation and sewage (longest median gestation lags) for visualization purposes. Panel (b) plots histograms by private, foreign or public ownership. Distributions are winsorized at the 1st and 99th percentile for visualization purposes. Public ownership includes sponsorship by central, state, or local governments, as well as state-owned enterprises (i.e. public sector is majority shareholder). Private ownership includes any private Indian sponsors.

Figure 3: Cost-Time Schedules

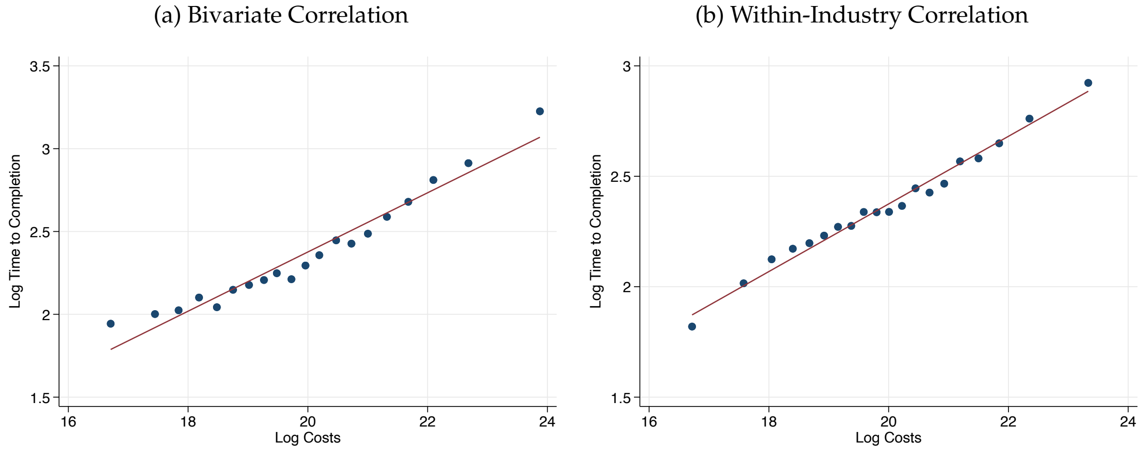

Notes: the figures plot binned scatterplots of (log) time to completion on (log) costs. The sample is restricted to all completed projects in the CapexDx data with non-missing cost data. Panel (a) plots the bivariate non-parametric relationship. Panel (b) first residualizes by industry dummies.

public enterprise in India, and railways have notoriously long gestation lags. Consistent with this view, we will see that public sponsorship explains little variation in gestation lags once we account for a project's end user industry. We cannot claim a causal effect of public sponsorship on longer gestation lags.

Last, we investigate the relationship between project size and gestation lag. We proxy for size with the cost of the project. This strategy is admittedly imperfect: projects that get delayed might be the ones incurring unexpected cost overruns; or projects built quickly might be those for which accelerating the project was crucial, regardless of the high costs incurred to ensure fast completion. Despite these limitations, monetary cost provides a common metric to compare projects across different "size" definitions. That is, while it's feasible to compare two residential buildings based on how many floors are being built, it's not as straightforward to compare railways to irrigation canals. Total cost provide a common metric across fundamentally different structures.

Figure 3 plots binned scatter plots of log time to build on log cost. Panel (a) plots a bivariate relationship, while panel (b) residualizes by industry fixed effects. There is a remarkable linear relationship in log space between gestation lags and costs. That is, the elasticity between size and gestation lags appears robustly constant. A 1% increase in costs is associated with a 0.18% increase in completion time; including industry fixed effects – and therefore averaging multiple within-industry correlations – barely moves the elasticity, from 0.18% to 0.15%. Appendix Table A3 collects estimates from the regression specifications.[17]

[17] Appendix Figure A3 documents that the constant elasticity pattern holds well within particular sectors and ownership groups.

Decomposing Variation in Gestation Lags A takeaway from our measurement is that time to build exhibits considerable variation. Towards a systematic accounting of such variance, we decompose it using fixed effects. We aim to only include "exogenous" predictors of gestation lags, as we want to explore the view that time to build is determined purely by technological factors. Table 2 collects estimates.

Column (1) shows that project sector accounts for 24% of the variation. Moving from a dozen sector bins to over a hundred industry bins in column (2) marginally increases the explained variance to 30%. Column (3) removes industry as a predictor and includes only the year the project was started. The year dummies account for 28% of the variance. In column (4), we include sector and year started fixed effects, and explain about 40% of the variation in time to build.

Admittedly, sector or industry, and particularly year started, may fail to proxy exclusively for exogenous, technological factors. Some industries might face particular contractual arrangements that shape economic incentives to complete projects faster or slower. Some years may pose particularly favorable conditions to start projects (e.g. low input costs across sectors). A stricter approach would obtain characteristics based on engineering blueprints, construction processes, and other purely technical factors. We find this approach promising, but refrain from pursuing it. Project descriptions in our data are not detailed enough to determine precisely such technical details. Further, we want to reduce the dimensionality of our predictor set. Exploiting progressively detailed information is bound to improve explanatory power, but at the cost of less interpretability.

In columns (5) to (8), we keep sector and year started fixed effects throughout. We then include other fixed effects to assess their additional explanatory power. Column (5) adds state fixed effects; column (6) adds dummies for whether the sponsor is public, private, or foreign; column (7) adds dummies for whether the project was a new unit; and column (8) includes all the preceding variables. Geography, ownership and unit type add little explanatory power relative to sector and year. Last, in column (9), we explicitly abandon dimensionality reduction and instead include sponsor and year started fixed effects. We manage to account for 58% of the variation.[18] In short, there is considerable within-sector (or even, within-firm) variation in gestation lags.

On the one hand, there is room for optimism: construction is known to be an unpredictable endeavor (Bajari and Tadelis, 2001). Yet, we managed to account for over half of the variation in gestation lags. On the other hand, the glass is not half full: better explanatory power comes from abandoning parsimony. By shackling ourselves to low-dimensionality and interpretability, we can learn at most about 30% of the variation in construction times,

[18]The increase in explanatory power is not driven solely by further granularity. In results available upon request, we show adjusted R^2. Even penalizing the number of included predictors, the specification with sponsor dummies features greater explanatory power.

Table 2: Cross-Sectional Variance Decompositions: Time to Completion

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
R^2	0.239	0.302	0.285	0.395	0.406	0.410	0.396	0.439	0.579
Unc. Var					0.011	0.015	0.001	0.044	
Partial R^2					0.018	0.024	0.002	0.072	
Ownership FE						✓		✓	
State FE					✓			✓	
Sector FE	✓			✓	☒	☒	☒	☒	
Industry FE		✓							
Year Started FE			✓	✓	☒	☒	☒	☒	✓
Type of unit FE							✓	✓	
Firm FE									✓

Notes: the table provides a variance accounting for completed projects. The first row reports the R^2 of a linear regression of gestation lags on the fixed effects indicated by ✓ in the bottom part of the table. In columns (5)-(8), all regressions include sector and year started fixed effects (marked by ☒); and the additional explanatory power of the added fixed effects (marked by ✓) is presented with uncorrelated variance shares (the difference between the R^2 to the R^2 in column 4) and the partial R^2 (the uncorrelated variance share divided by the unexplained variance share in column 4). Lower values of the uncorrelated variance share indicate that the additional predictors have low additional explanatory powers relative to the predictors included in the baseline model (column 4). Lower values of the partial R^2 indicate that the additional predictors have low explanatory power for the residual variation left from the baseline model (column 4). Type of Unit captures whether the project is a new structure, a substantial renovation, or a modernization of an existing structure; and type of ownership captures private (Indian), foreign or public sponsorship.

primarily by looking at industry, and perhaps a tad more by considering other easily observed covariates.

4 Testing for Endogenous Time to Build

We now turn to the main empirical exercise. Are gestation lags purely exogenous constraints that firms must accept? Or do firms have some control over how fast they build projects?

Testing whether firms can speed up projects requires us to devise an experiment where we do not alter the natural progression of project completion, but shock the conditions faced by the project sponsor. Consider a scenario where time to build was purely exogenous. At project onset, firms chose an idealized builder; made payments upfront; and once the contract was signed, could not make changes to the construction process. In turn, the builder consistently worked in the project. In the absence of any unforeseen disturbances, time to build would be perfectly known. Whether the firm experienced distress — weak sales, extraordinary expenses, cash shortages, or cuts in credit lines – would have no bearing on the eventual completion of a project. Ex post shocks could cause early or late deliveries by the idealized contractor, but those would be unrelated to any decisions taken by the firm. This is what we mean by “natural progression” of a project. We would *not* be able to reliably detect any relationships between shocks to firms and project completion times in the data, as only unobserved, project-level disruptions would generate variation in gestation lags.

We contrast this sharp null hypothesis against an alternative scenario where frms do have some control over how quickly to finish projects. If sales were exceptional and capacity expansion was desirable, firms could speed up completion of an ongoing manufacturing plant. Conversely, in times of weak performance, firms could stall existing projects and shelve new projects queued on their capex pipeline. In the absence of project-specific disturbances that are correlated with the firm-level disturbance, such episodes would offer encouraging evidence. Firms reacted to the shock by actively managing the portfolio of capital expenditure projects, changing how quickly they completed ongoing projects, a feature we name "endogenous time to build".

Our empirical strategy systematizes the thought experiments above. The key difficulty is to find a shock that affects the firm, but is plausibly unrelated to the daily affairs of project construction. If projects derail because their sponsors are bad firms, it is hard to argue that firms explicitly had some control over the derailment. It could be that bad firms consistently choose projects with worse technological factors, and bad outcomes are likelier to happen to bad firms. That an adverse shock hit the firm at the same time that its project derailed would thus be uninformative about whether the firm could alter the gestation lag.

We circumvent this identification challenge by considering a setting where firms face a contraction in credit supply because their banks get in trouble. Both good and bad firms have relationships with the troubled banks, and we can control for how good a firm is on average. Importantly, firms are differentially exposed to the troubled banks; and we will provide evidence that more exposed firms did not differ much from less exposed firms before the banks got in trouble. If the remaining reasons for potential project disruptions are just noise, we can compare what happens to projects of more exposed firms relative to projects of less exposed firms. If the former suffer greater disruptions than the latter, we can conclude that a firm-level shock unrelated to the natural progression of a project was responsible for the project disruption.

4.1 A Credit Supply Shock: the Indian Asset Quality Review

Our quasi-experimental setting is the Indian Asset Quality Review, a widespread audit of Indian commercial banks. This bank clean-up exercise was a preemptive measure taken by the Reserve Bank of India (RBI) during a noncrisis period: the RBI did not structure capital infusions or backstops, and relied on market forces to recapitalize troubled banks. We summarize the main features of this setting, and refer readers interested in details to previous work that shares our setting and describes the intervention at length.[19]

[19] See Chari et al. (2021), Flanagan and Purnanandam (2023), Kulkarni et al. (2021). Our summary closely follows Chopra et al. (2020).

Before the Global Financial Crisis, Indian banks had to automatically reclassify restructured loans into non-performing assets, and set aside corresponding provisions. On August 27, 2008, the RBI suspended automatic reclassification given global turmoil in financial markets. No additional provisions had to be set for restructured loans, so long as borrowers met restructuring terms. This regime lasted until April 1 2015, despite concerns of evergreening in the banking sector.[20] The Asset Quality Review was the RBI response to these concerns. It was at first a special audit of scheduled commercial banks conducted between August and December of 2015. Following this first inspection, the RBI conducted annual audits under the same guidelines. RBI auditors estimated the amount of non-performing assets in each bank's portfolio, required provisions to be set aside, and mandated that banks disclose audit results in their financial statements.[21] The audits uncovered sizable under-provisioning across systemic banks and thus had the potential to reduce bank lending.

The empirical banking literature has established exercises to identify credit supply contractions following exogenous shocks to banks. The key identification challenge is to disentangle credit demand from credit supply. That is, did borrowers want to borrow, but banks lent less because of the shock? Or did the shock also decrease borrower appetite for credit? The standard strategy is to control for credit demand (Khwaja and Mian, 2008) by including borrower or borrower-by-time fixed effects. This is possible because banks extend loans to multiple borrowers, while borrowers tend to borrow from multiple banks. The fixed effects estimator then identifies the effect of the shock on equilibrium loan amounts by comparing more and less exposed banks holding credit demand fixed. The inclusion of lender fixed effects guarantees that our estimator is not capturing time-invariant heterogeneity across banks. That is, our estimates do not compare banks that systematically make good or bad loans. Rather, the estimator asks whether banks more exposed to the shock differentially cut loans to a particular borrower in a given year, relative to its usual proclivity to extend loans.[22]

We construct a lender-borrower-year dataset from our MCA data. First, we isolate all outstanding loans between 2013 and 2015, roughly three years before the first special audit;

[20] Loan loss provisions reduce lending capital and therefore loan profits (Freixas and Rochet, 2008). Banking regulation imposes such policies to align privately optimal bank risk-taking with the ensuing externalities, i.e. increases in systemic risk (Freixas et al., 2015). Consistent with RBI concerns and with the cited theories, Chopra et al. (2020) documented sizable increases in restructurings, but no increases in non-performing assets during the forbearance period.

[21] The reporting requirement was binding either if the additions to provisions exceeded 15% of bank's net profits; or if additions to gross non-performing assets flagged by auditors exceeded 15% of yearly additions to gross non-performing assets flagged by the bank. Chopra et al. (2020) documents that these requirements were mostly binding; that additional provisions and gross non-performing assets were sizable relative to previous net profits and incremental gross NPAs; and that reporting (and thereby troubled) banks accounted for a large share of the overall secured loan volume in the MCA charge registry.

[22] Note that we use the terms bank and lender interchangeably. Similarly, we refer to sponsors as firms or borrowers. The terminology facilitates parallels with previous work in the empirical macro-finance literature.

and collect all borrower-lender pairs by summing over different charges. For each pair, we create observations for years 2013-2019. If a pair had an outstanding loan in a given year, we assign it the total outstanding loan amount in that given year. If the pair had no outstanding loans in a given year, we assign it a value of zero. The procedure gives us a balanced panel of firm-bank relationships over the sample period, fixing the composition of relationships prior to the audits.

We measure how exposed a bank b is to the audit shock at time t as the difference in provisions resulting from the yearly audit, scaled by total bank assets

$$\text{Bank Shock}_{b,t} = \left(\frac{\text{Provisions Required by the RBI Post-Audit} - \text{Provisions Pre-Audit}}{\text{Total Assets}} \right)_{b,t}$$

The measure is intuitive: banks who must post more provisions experience larger reductions in lending capital. In turn, less lending capital should translate into fewer loans. We normalize by bank assets to avoid spuriously comparing large and small banks.[23] We then regress total loans outstanding between lender b and borrower f in year t using a Poisson pseudo-maximum likelihood (PPML) estimator[24]

$$\text{Loans}_{bft} = \exp\left(\gamma \times \text{Bank Shock}_{b,t} + \lambda_{ft} + \omega_b + X_{bft}\Gamma\right) \epsilon_{bft}$$

where λ_{ft} are firm-by-year fixed effects that absorb unobservable credit demand shocks, ω_b are bank fixed effects that absorb permanent heterogeneity across banks, and X_{bft} collect other controls. We are interested in γ, the semi-elasticity of equilibrium loan amounts to the bank shock. We cluster standard errors at the lender level.

Table 3 presents the results. Column (1) provides a baseline specification with lender and borrower-by-year fixed effects. A one-standard deviation increase in bank exposure reduces outstanding credit in the average pair by 5.1%. Importantly, the inclusion of firm-year fixed effects control for time-varying demand confounders. The reduction thus represents a relative reduction in credit supply among more versus less affected banks. Column (2) relaxes the control set by including borrower and year fixed effects. Hence, we control for aggregate

[23] The shock is set to zero if the lender is not a scheduled commercial bank. We include non-SCBs in the sample as borrowers may substitute towards them following the information shock hitting SCBs. The shock also takes value zero for years prior to the start of the special audits. These years provide information on lending dynamics for the affected institutions before the information shock.

[24] A growing literature advocates for pseudo-Poisson maximum likelihood methods in settings where the dependent variable is characterized by a distribution with non-negative support and a large mass of zeros. Gourieroux et al. (1984) proved that validity of a Poisson regression extended beyond count outcomes, relying instead on a correct specification of the conditional mean. Silva and Tenreyro (2006) applied PPML to gravity equations, highlighting large discrepancies between log-linear specifications estimated by OLS with $\log(1+Y)$ as dependent variables and the corresponding Poisson specifications. Cohn et al. (2022) documented similar pitfalls in empirical corporate finance. Chen and Roth (2022) proved that $\log(1+Y)$ specifications estimated by OLS are not invariant to the scaling of the outcome variable Y, and that one can recover arbitrarily-sized treatment effects by appropriately rescaling the outcome. They explicitly recommend the use of Poisson methods instead.

Table 3: Effects of Asset Quality Review on Credit Supply

	(1)	(2)	(3)	(4)	(5)	(6)
Exposure	-0.0517***	-0.0508***	-0.0513***	-0.0497***	-0.0485***	-0.0491***
	(0.0106)	(0.0119)	(0.0120)	(0.0122)	(0.0130)	(0.0131)
Outc Mean	2497.20	2497.47	2497.47	4450.34	4450.34	4450.34
Obs	268858	268827	268827	109929	109929	109929
Full Sample	✓	✓	✓			
Hazard Sample				✓	✓	✓
Lender FE	✓	✓	✓	✓	✓	✓
Borrower by Year FE	✓			✓		
Borrower and Year FE		✓	✓		✓	✓
Firm Controls			✓			✓

Notes: the table reports PPML estimates of the semi-elasticity of outstanding loans within a borrower-lender pair in a given year to lender exposure to the Asset Quality Review. The sample consists of a balanced panel of firm-bank relationships for 2013-2019, conditional on an existing relationship (positive outstanding loans) in the 2013-2015 period. Exposure is defined as the difference between required provision after the audit and the original provisions, scaled by total bank assets. Coefficient estimates are read as the percent effect on the outcome for a one standard-deviation increase in lender exposure. For example, column (1) reports that moving from the lender with average exposure to a lender with one-standard deviation higher exposure decreases outstanding loans by 5.08%. In columns (1)-(3), the sample comprises all borrower-lender pairs in our MCA-Prowess merged sample. In columns (4)-(6), the sample is restricted to firms for which we eventually observe projects in our hazard regressions (see model (3) and table 4). Columns (1) and (4) include lender and borrower-by-year fixed effects. Columns (2) and (5) include lender, borrower, and year fixed effects. Columns (3) and (6) add firm controls. Firm controls include sales, sales growth, profitability, debt over equity, book leverage, equity issuances, cash over liquid assets, and debt over EBITDA. All firm controls are winsorized at the 1st and 99th percentile, and then standardized. All firm controls are lagged by one-period. Standard errors are clustered at the lender level and reported in parentheses. $*$, $**$, and $***$ denote statistical significance at the 10%, 5% and 1% levels, respectively.

demand shocks or systematic heterogeneity in firm credit demand, but not for time-varying credit demand shocks. Reassuringly, the point estimates are economically indistinguishable. Column (3) confirms the limited role for time-varying credit demand confounders by adding time-varying firm-level controls to the specification in column (2).

Not all firms in the merged MCA-Prowess sample sponsor projects in Capex. If the latter are better insulated from the audit shock, perhaps by banking mostly with foreign institutions, our quasi-experimental setup would be of little use. We thus restrict the sample in columns (4)-(6) to borrowers for whom we observe projects in the Capex dataset, and whose projects are included in our project-level regressions. First, the audit shock resulted in a 4.9% average contraction of credit supply among this subsample, a reassuringly similar magnitude. Second, we once again find little role for time-varying credit demand confounders relative to potential aggregate or firm-specific confounders.[25]

[25] Another sensible specification omits lender fixed effects and estimates a linear model in first-differences. Appendix Table A4 reports results from such specification. Although our results are slightly weaker in the broader sample, we still find evidence of a credit supply contraction. Moreover, the semi-elasticities obtained by dividing the coefficients by the mean of the outcome are similar in magnitude to those estimated in Table 3. This specification suggests a greater role for time-varying demand confounders, but the control set included in the regressions with only borrower and year effects approximate well the estimates from the more saturated specifications.

4.2 From Banks to Projects: Estimation and Identification

Having established that the audit shock has bite, we move to our main results. We need to map the banking shock to firms. Projects are exposed to the banking shock if so is their sponsor. If a firm is more exposed to troubled banks, it potentially faces starker credit disruptions. We model exposure for firm f by aggregating bank-level exposures to the audits, weighting each bank b by its lagged share of total borrowing (Chodorow-Reich, 2013)

$$\text{Firm Exposure}_{f,t} = \sum_{b \in B_{f,t-1}} w_{b,f,t-1} \text{Bank Shock}_{b,t} \tag{1}$$

$$w_{b,f,t-1} = \frac{\text{Outstanding Loans}_{b,f,t-1}}{\sum_{b \in B_{f,t-1}} \text{Outstanding Loans}_{b,f,t-1}} \tag{2}$$

To quantify the effect of shock exposure on project completion hazards, we estimate a generalized linear model

$$\phi(Completed_{pft}) = \beta \times \text{Exposure}_{ft} + X_{pft}\Gamma + v_{pft} \tag{3}$$

where p refers to a project, f refers to the project's sponsor (or, firm), t refers to calendar time, and $\phi(\cdot)$ is a link function (eg. affine ϕ yields linear regression). X_{pft} collects control variables, and v_{pft} denotes residuals. The effect of interest is β. We cluster standard errors at the firm level (Abadie et al., 2022).

Our specification borrows from (discrete-time) survival analysis (Singer and Willett, 1993). OLS estimation in our setting is misspecified due to right-censoring. Restricting our attention to projects for which we observe full histories – a "complete case analysis" – runs into selection issues: completed projects may systematically differ from censored projects. Survival analysis – also known as duration models in econometrics (Van Den Berg, 2001) – explicitly accounts for right-censoring, allowing us to use all projects in the data. The methodology thus circumvents these selection issues.

We organize the data into a quarterly panel of projects.[26] The outcome $Completed_{pft}$ is zero in all periods before project completion, and one in the completion period. If a project is right-censored, either because CapexDx loses track of or because the project is ongoing, $Completed_{pft}$ takes value zero in all time periods.

Different choices for the link function ϕ yield likelihood functions that match or approximate underlying continuous-time models. The choice of link function is mostly guided by the underlying data generating process for project completion. On the one hand, Allison (1982) shows that a discrete-time model for hazards yields a likelihood that coincides with

[26] We opt for a quarterly panel, despite exposure varying at the yearly level. Since completion times are small in yearly scale, we would not account well for age effects in an yearly panel; and would not disentangle well the role of project age from shocks to firms, except for firms with the largest project portfolios.

the likelihood for a logistic regression model. On the other hand, logistic regressions can run into incidental parameter problems (Neyman and Scott, 1948). This is a real concern, since we will include high-dimensional fixed effects motivated by identification concerns.

Instead, we take the view that we are grouping the underlying continuous-time process into discrete periods; and opt for a log link $\phi(x) = \log(x)$. Formally, the log link implies a semi-parametric, piece-wise exponential model for hazards. Intuitively, it assumes that conditional on surviving up to period t, completion arrives at a Poisson rate that depends on the covariates included in the RHS of (3). We find it reasonable: our outcome is binary (therefore a count) and captures the arrival of a rare event (the unconditional probability of completion is 4%). Alternatively, one can relax the the Poisson assumption, and interpret our choice as specifying a log-linear hazard function with multiplicative error structure (Silva and Tenreyro, 2006). In this case, under weak assumptions, the model can be estimated by Poisson pseudo-maximum likelihood methods, which only require the conditional expectation function to be correctly specified (Gourieroux et al., 1984).[27]

Proper identification of β requires satisfying a non-linear GMM moment condition

$$\mathbb{E}\left[v_{pft}\text{Exposure}_{ft}|X_{pft}\right] = 0 \tag{4}$$

Leveraging recent results in the econometrics of exposure designs, it is sufficient that bank shocks are (conditionally) uncorrelated with (loan-weighted) project-level unobservables. Identification leverages exogenous variation in the shocks. Correlation of the shares $w_{f,t-1}$ with unobservables and with the outcome do not invalidate the design (Borusyak et al., 2021). Instead, we should be worried if (1) there are other factors that influence project completion and are correlated with the credit supply shock; and (2) we fail to control for such confounders.

Previous research has accumulated convincing evidence on the exogeneity of the credit shock (Chopra et al., 2020, and references therein). We safeguard our analyses by intuiting potential confounders that could threaten a causal interpretation. Whenever possible, we include the appropriate controls into X_{pft}. First, some firms may systematically perform poorly, build projects slowly[28], and build relationships with unhealthy banks who are more likely to underreport non-performing assets. Firm fixed effects address this problem under

[27] Conveniently, including high-dimensional fixed effects in Poisson pseudo-maximum likelihood models with large datasets is feasible due to novel existence results and fast estimation routines (Correia et al., 2020, 2021).

[28] For instance, they may have access to better construction technologies or better relationships with bureaucrats; and banks might differentially cut credit supply due to these characteristics. For firms with a single project in the sample, inclusion of firm effects implies that the estimator compares project-level outcomes before and after exposure; for firms with multiple projects, we rely on within-firm rather than within-project variation. This allows us to strike a compromise between identification and sample size for estimation. Including project-level effects would drop too many observations due to separation concerns arising from the high-dimensional fixed effects. Firm-level fixed effects allow us to absorb substantial confounding variation, while simultaneously allowing to include other fixed effects that control for project-level confounders.

the assumption that this sorting does not vary over time.

Second, projects in our sample are immobile, which makes them susceptible to time-varying geographical shocks. Political instability may disrupt project timelines (e.g. military tensions in the Northeast relative to Southern India, or idiosyncratic changes in construction regulations). Variation in economic slack (e.g. local unemployment, construction labor supply or wages) may alter incentives to hire workers and thus change gestation lags. These factors threaten our design if banks cut credit by internalizing them. We include state-by-year fixed effects to absorb these time-varying regional confounders.

Third, aggregate (e.g. a recession) or industry-level (e.g. supply chain disruptions) shocks may delay project completion regardless of firm actions; and may lead bank loans to perform poorly, induce banks to hide losses, and thereby increase divergences in provisions. We include year fixed effects or sector-by-year fixed effects to mitigate these concerns.[29] Fourth, we include a cubic spline on project age, which removes baseline hazards and allow us to (semi-parametrically) compare projects of similar ages.

In our most conservative specifications, we further include controls related to recent firm performance in product markets (e.g. sales growth) and firm financial solvency (e.g. leverage, cash-to-assets, debt-to-gross profits, etc.) Even in the absence of confounding concerns, these controls improve precision of our estimates. For the more skeptical readers, these controls attempt to span the space of time-varying firm-level confounders that influenced both project completion likelihood and bank's propensity to hide losses. Suppose that periods of low sales growth led firms to start projects with ex ante lower expected time to build; and that low sales growth of individual firms led to defaults, impaired a sizable chunk of bank advances, and encouraged banks to hide these losses. In that case, failing to control for past sales growth could confound our effect of interest. However, we find hard to concoct stories where time-varying idiosyncratic characteristics of specific firms would generate sizable distortions in bank's propensities to hide losses. This could only be the case if either these firms were tremendously granular as to shape behavior of the entire banking sector, which we find unlikely; or if the empirical distribution of charges was characterized by dispersed lending through many small banks, which is not the case.

Table 4 presents estimates of model (3). Column (1) starts by estimating a simple bivariate specification, only controlling for baseline hazards. Taken at face value, a one-standard deviation increase in project-level exposure to the audit shock increases hazards contemporaneously by 11.5%. The effect is statistically significant at conventional levels, but potential confounders abound. Column (2) tightens identification by including firm fixed effects and

[29]Throughout, we use "financial year" fixed effects. The financial year in India runs from April 1st to March 31st of the subsequent year. All financial data, and thereby bank exposure data is measured at this frequency. Even many official macroeconomic series are released based on financial year rather than calendar year denominations.

Table 4: Effects of Firm-Level Credit Supply Shocks on Project Completion Hazards

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Exposure	0.1154***	0.0701**	0.0681**	0.0623**	0.0600**	0.0758***	0.0750***
	(0.0278)	(0.0273)	(0.0275)	(0.0265)	(0.0268)	(0.0271)	(0.0273)
DepVar Mean	0.0392	0.0392	0.0392	0.0392	0.0392	0.0409	0.0409
Observations	123857	123857	123857	123857	123857	118314	118314
Age Splines	✓	✓	✓	✓	✓	✓	✓
Firm FE		✓	✓	✓	✓	✓	✓
State x Year FE		✓	✓	✓	✓	✓	✓
Sector x Year FE				✓	✓		
Industry x Year FE						✓	✓
Firm Controls			✓		✓		✓

Notes: the table reports PPML estimates of the semi-elasticity of project-level hazards to firm exposure to the Asset Quality Review (see model (3)). The sample consists of projects observed quarterly between 2013-2019. Firm exposure is defined as the weighted-average of lender exposure, where weights correspond to outstanding loans in the previous year. Coefficient estimates are interpreted as the percent effect on the outcome for a one standard-deviation increase in firm exposure. For example, column (1) reports that moving from a firm with average exposure to a firm with one-standard deviation higher exposure increases the hazard by 11.54%. Age splines refer to cubic splines on a project's age (measured in quarters). Firm controls include sales, sales growth, profitability, debt over equity, book leverage, equity issuances, cash over liquid assets, and debt over EBITDA. All firm controls are winsorized at the 1st and 99th percentile, and then standardized. All firm controls are lagged by one-period. The reduction in sample size from columns (5) to columns (6) and (7) results from separation of additional observations by the more granular industry-by-year fixed effects relative to the sector-by-year fixed effects. Standard errors are clustered at the firm level and reported in parentheses. $*$, $**$, and $***$ denote statistical significance at the 10%, 5% and 1% levels, respectively.

project state-by-year fixed effects. A one-standard deviation in project-level exposure now raises hazards by 7%, a statistically significant effect at conventional levels.

Column (3) shows the results are unaffected by the inclusion of firm-year controls. Column (4) adds NIC sector-by-year fixed effects. The effect size drops to 6.2%, but remains significant. Column (5) is our preferred specification. The additional firm-level controls barely moves the point estimate: hazards increase by 6%, a significant effect at conventional levels. Last, columns (6) and (7) replace sector-by-year with industry-by-year fixed effects. If anything, the tighter specifications yield larger completion effects.

Perhaps lagged measures of firm real and financial performance are insufficient to control for all possible time-varying unobservable shocks. Including firm-year effects would eliminate this concern. In the bank-firm level regressions, that was our strategy to isolate the supply shock from other potential confounding shocks. However, firm-year effects are not a feasible strategy in our project-level or firm-level regressions. Exposure varies at the firm-year level, so including firm-year effects would absorb all the variation and preclude identification of our coefficient of interest.

The question then is: how well does our control set span the space of time-varying firm-level confounders? We follow Khwaja and Mian (2008) and Chodorow-Reich (2013) and use the results from our credit regressions as suggestive evidence on the absence of unobservable firm-year confounders that would violate (4). In Table 3, moving from borrower-year to borrower and year fixed effects barely changed our results. If borrower-year unobservables were the chief sources of confounding, we would have expected that relaxing the control set would change point estimates considerably. Instead, firm and year effects approximated well the information set spanned by the full suite of firm-year effects. Given we include firm and year fixed effects in our preferred specifications, we do not believe our results are driven by unobservable firm-year confounders.

We provide a battery of robustness exercises in Appendix C. Table A5 shows that our choice of weights for aggregating the bank shocks is conservative. Using pre-period weights instead of lagged, time-varying weights strengthens the results. Table A6 shows results from linear models estimated in the same sample as the PPML regressions. The semi-elasticities obtained by dividing point estimates by dependent variable means are similar to those in Table 4, which reassures our results are not driven by the choice of estimator.

4.3 Extensive Margin Results

Firms discount future cash flows, so accelerating completion increases the project's net present value. If it was costless to speed up, firms would always do so. By revealed preference, speeding up must be costly. But our results document that firms are engaging in costlier investment behavior in times of low credit supply, precisely when we would expect firms to

Table 5: Effects of Firm-Level Credit Supply Shocks on Project Starts

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Exposure	-0.0775***	-0.0698**	-0.0670**	-0.0806**	-0.0763**	-0.0584*	-0.0588**
	(0.0286)	(0.0325)	(0.0322)	(0.0316)	(0.0314)	(0.0299)	(0.0296)
DepVar Mean	0.8362	0.8362	0.8362	0.8362	0.8362	0.8326	0.8326
Observations	8994	8994	8994	8994	8994	8869	8869
Firm FE		✓	✓	✓	✓	✓	✓
Year FE		✓	✓				
Section x Year FE				✓	✓		
Division x Year FE						✓	✓
Firm Controls			✓		✓		✓

Notes: the table reports PPML estimates of the semi-elasticity of firm-level project starts to firm exposure to the Asset Quality Review (see model (5)). The sample consists of the Capex-Prowess matched sample of firms observed yearly between 2013-2019. Firm exposure is defined as the weighted-average of lender exposure, where weights correspond to outstanding loans in the previous year. Coefficient estimates are interpreted as the percent effect on the outcome for a one standard-deviation increase in firm exposure. For example, column (1) reports that moving from a firm with average exposure to a firm with one-standard deviation higher exposure reduces project starts by 7.75%. NIC Section and Division classifications are constructed from NIC codes and categorizations. Firm controls include sales, sales growth, profitability, debt over equity, book leverage, equity issuances, cash over liquid assets, and debt over EBITDA. All firm controls are winsorized at the 1st and 99th percentile, and then standardized. All firm controls are lagged by one-period. The reduction in sample size from columns (5) to columns (6) and (7) results from separation of additional observations by the more granular division-by-year fixed effects relative to the section-by-year fixed effects. Standard errors are clustered at the firm level and reported in parentheses. $*$, $**$, and $***$ denote statistical significance at the 10%, 5% and 1% levels, respectively.

cut down on capital expenditures. Though puzzling at first, our results are consistent with a potential decline in investment. Firms also start new projects. If speeding up ongoing projects is costly, and firms can cut down investment on other margins, we would expect to see a decline in new project starts.

We test for this extensive margin channel by collapsing the data to the firm-year level and estimating by PPML

$$\text{New Project Starts}_{ft} = \exp\left(\beta \times \text{Firm Exposure}_{ft} + X_{ft}\Gamma\right)\nu_{ft} \tag{5}$$

where ν_{ft} is a residual and X_{ft} collects firm fixed effects, year fixed effects, and time-varying firm-level controls. We cluster standard errors at the firm level.

Table 5 collects results. In a bivariate specification (column 1), a one-standard deviation increase in firm-level exposure reduces new project starts by 7.7%. Adding firm and year fixed effects (column 2) reduces the point estimates to 7%. Further adding time-varying firm controls barely moves the point estimates. In columns (4) and (5), we replace year fixed effects by NIC sector x year fixed effects to control for sector-specific shocks. In our preferred specification (column 5), a one-standard deviation increase in exposure reduces new project starts by 7.6%. Replacing sector-by-year with division-by-year effects slightly reduces the point estimates to 5.8%. All effects are significant at conventional levels. Reassuringly, the magnitudes are stable across specifications.

Appendix C collects a battery of robustness checks. Table A7 shows that our choice of weights for aggregating the bank shocks is conservative. Using pre-period weights instead of lagged, time-varying weights strengthens our results. Table A6 shows results from linear models estimated in the same sample as the PPML regressions. The semi-elasticities obtained by dividing point estimates by dependent variable means are similar to those in Table 5, reassuring our results are not driven by the choice of estimator.

4.4 Other Threats to Identification

Our empirical design is valid if the control set ensures that the identification condition (4) holds. Confounding would result if, conditional on the set of included observables, there existed factors that (i) influence project completion and correlate with the credit supply shock; (ii) *and* that we fail to control for.

A first concern is the absence of firm-time effects in a world where effects are instead driven by some unobservable firm-year confounder. That is, we can imagine a world where there was no causal effect of the audit shock on project completion or firm starts. Rather, the effects were driven by some unobservable characteristic that (i) banks internalized and conditioned on when reducing credit, and (ii) made firms less willing to start projects, but, somehow, more willing to complete ongoing projects. For example, loan officers could condition on some private information that we are not privy to as econometricians. They could cut loans to firms they knew were temporarily going through hard times. This scenario would threaten our results if our control set did not approximate well the information set that loan officers used when differentially reducing credit. To borrow from Chodorow-Reich (2013), the identifying assumption would be violated if there was no "'as good as random' matching of banks and borrowers conditional on observables".

We attempt to address these concerns by first presenting balance tables. If our exposure measure predicted pre-audit firm characteristics, we would be concerned that banks were differentially cutting credit on these observables. Exposure would fail to be quasi-random. Further, the effects we estimate could be driven by such observables if they were relevant for predicting our outcomes of interest. For example, some firms could be experiencing weak sales growth around the time of the audits. This weak performance could drive firms to cut on their project pipeline; and could reasonably lead banks to cut credit to these firms by more relative to healthier firms.

Appendix Table A10 ameliorates these concerns. We regress one-period lagged firm characteristics on a contemporaneous measure of exposure, including firm and year effects. We thus test for predictability in trends, since we rely on within-firm time-series variation in our design. We detect predictability for profitability and cash over assets, but fail to detect predictability among other covariates.[30]

A second, more subtle concern, is that the shock could change the ability of the data provider to collect data on projects. This possibility opens up space for a delicate form of selection bias: the sample of projects collected after the shock may differ from the sample of projects before the shock. We would not be able to make causal statements if the treatment induced such differential censoring.

In Appendix D, we provide two analyses to soothe these concerns. First, we restrict our completion analyses to projects that were in data vintages *prior* to the first audits. Even if the data collection process changed across treated and untreated units following the shocks, projects in vintages prior to the shock would not have been affected by these changes. The sample is thus immune to the selection concern. Second, we test a direct implication of differential censoring due to differential backfilling across treatment groups. If CapexDx had become less able to backfill data on treated relative to untreated projects, we should observe a downward trend in the number of treated projects relative to untreated projects before the shock. We thus test whether exposure predicts differential pre-trends in firm-level project stocks. Across the two tests, we find no evidence that our results are driven by differential censoring.

5 A Theory of Investment with Endogenous Time to Build

Our reduced form results may appear puzzling at first. Firms are experiencing an adverse shock. Expectedly, they are cutting down on new projects, consistent with standard investment theory. However, they are speeding up existing projects. By revealed preference, this must be a costly action. How can we reconcile the intensive and extensive margin results?

We rationalize our empirical findings in an investment model with endogenous time to build. By endogenous, we mean that the firm *chooses* gestation lags, rather than treating it as an exogenous primitive. Our model is minimalist. It allows us to make analytical statements that cleanly spell out the intuition for our results. Yet, it lends itself to quantitative applications. With some further enrichments, we can take the model to the data and perform counterfactuals, which we defer to section 6.

Setup Time is continuous and indexed by t. A firm manages a collection of capital investment projects.[31] Each period, a continuum of projects become available. Completing a

[30] We also run cross-sectional balance tests. In column (3) of Table A9, we regress pre-audit firm characteristics on average exposure ex post. We detect cross-sectional predictability in sales and equity issuance, but fail to detect predictability otherwise. Note that this test is more stringent that what's required by our identification assumption, but is overall encouraging.

[31] The model's interpretation is the same whether projects are managed by a single firm or by many firms sharing the same structural primitives. Alternatively, one can think of a firm being a project, and use the framework

project requires investing an amount $\bar{x}$. After the project is finished, say at time T, the firm obtains z units of capital. Project output z is drawn at project inception from a p.d.f. $f(z)$. Capital is valued at a common price p, so that project returns are pz. Firms take p as given.[32] Our setting incorporates project heterogeneity in returns rather than size, an assumption we later relax. Project returns only accrue upon project completion. The firm's discount rate is given by $r > 0$. We focus first on a partial equilibrium setting.

Let $i(t)$ be how much the firm spends on the project in period t. Normalize the start time of the project at $t = 0$. The project is then completed at the time T that satisfies

$$\int_0^T i(t)dt = \bar{x} \tag{6}$$

Equation (6) states projects are completed when investment flows add up to total size. Suppose p, z and r are constant (ie. as in a steady state). The firm chooses a sequence of investment flows $i(t)$ and a completion time T to maximize the net present value (NPV) of the project subject to (6)

$$\max_{i(t),T} e^{-rT}pz - \int_0^T e^{-rt}i(t)dt \quad \text{s.t.} \int_0^T i(t)dt = \bar{x}$$

In the absence of frictions, the solution is trivial: build positive NPV projects instantaneously, and discard negative NPV projects. To see why, note that for any fixed T, the firm minimizes the NPV of construction costs. Since $r > 0$, it is optimal to invest $\bar{x}$ immediately before T, and nothing at all other times. Discounted profits are then given by $e^{-rT}(pz - \bar{x})$, which are maximized at $T = 0$ so long as $pz > \bar{x}$. If $pz < \bar{x}$, the firm does not build the project.

But physical capital is not built in a day – our measurement exercises reject zero gestation lags. We introduce frictions into the construction process, modeled as additional costs $\kappa(i)$, assumed to be weakly increasing and weakly convex. The firm problem becomes

$$\max_{i(t),T} e^{-rT}pz - \int_0^T e^{-rt}\left[i(t) + \kappa(i(t))\right]dt \quad \text{s.t.} \int_0^T i(t)dt = \bar{x} \tag{7}$$

The frictional costs $\kappa(i)$ do not contribute to installed capacity, i.e. they are deadweight losses. Whenever necessary (i.e. in quantitative exercises), we assume that $\kappa(i)$ has the fol-

to study industry (Kalouptsidi, 2014) or macroeconomic dynamics.

[32]For now, firms will take terminal payoffs as fixed. When we do first-order dynamics below, p will not respond to changes in discount rates. This is consistent with the benchmark neoclassical framework with frictionless rental markets or secondary markets for capital resales. We can also think of p as a contractually stipulated payoff that cannot be renegotiated. We discuss deviations from this assumption in Section 7.4.

lowing functional form (none of our analytical results rely on the functional form):

$$\kappa(i) = \begin{cases} 0 & , i \leq \bar{i} \\ \frac{\kappa}{2}(i - \bar{i})^2 & , i > \bar{i} \end{cases}$$

The firm can invest without incurring any additional costs up to some upper threshold $\bar{i}$. We interpret $\bar{i}$ as a natural rate of project completion. Beyond this point, accelerating incurs a quadratic cost parametrized by κ. Our model nests one with exogenous time to build as $\kappa \to \infty$, in which case all projects are completed in $T = \bar{x}/\bar{i}$ periods.[33] The exogenous case illustrates how $\kappa(\cdot)$ can stand for a technological constraints on the rate at which the project can be completed.

The convex element captures many other complexities of construction. For instance, κ can capture congestion costs: building a factory is faster if more workers are putting up walls and laying down wires, but workers eventually start to bump heads. κ could also stand for moral hazard concerns. Suppose it is costly for contractors to exert effort; and that effort is not contractible (or, if so, not verifiable). In a second-best world, costs unrelated to capacity installation arise, since workers must be incentivized to exert effort. Additional stories for κ are plenty: costs associated with renegotiating contracts, with speeding up deliveries, with facing a convex labor supply curve for construction workers, or with many other frictions that characterize the construction industry. Bajari and Tadelis (2001) provide a comprehensive account of such issues.

Optimal Investment Towards a recursive representation, let $x(t)$ denote the state of the project

$$x(t) = \int_0^t i(s) ds \tag{8}$$

Projects are heterogeneous with respect to project stage x and project-specific returns z. We write the firm problem as a Hamilton-Jacobi-Bellman (HJB) equation with state x

$$rv(x,z) = \max_{i \geq 0} -i - \kappa(i) + i\frac{\partial v(x,z)}{\partial x} \tag{9}$$

and terminal condition

$$v(\bar{x},z) = pz \tag{10}$$

The HJB characterizes optimal firm behavior. The flow value of a project to the firm is

[33] The exogenous TTB formulation is inconsistent with our empirical findings under the assumption of optimizing firms. If firms speed up, they increase i following the shock. This would require that they were investing $i < \bar{i}$ at the time of the shock, increasing their investment to $i = \bar{i}$ at the time of the shock. But then, firms would not have been optimizing. This model is consistent with the data only if the credit supply shock increases the technological primitive $\bar{i}$. We find this explanation unlikely.

$rv(x,z)$. When the firm behaves optimally, it trades off the costs and benefits of investment. Costs are born today: placing i units towards the project costs $i+\kappa(i)$. Benefits accrue only when projects are completed. But projects close to completion are more valuable, since the firm is closer to reaping returns. More investment today thus raises the future value of the project. When time intervals are small, the marginal benefit of investing i is given by $\partial v(x,z)/\partial x$, so that the increment to the project's continuation value is $i\partial v(x,z)/\partial x$. The terminal condition restates that a completed project returns pz.

Since $\kappa(i)$ is differentiable, the necessary condition for an optimum is

$$1+\kappa'(i)=\frac{\partial v(x,z)}{\partial x} \tag{11}$$

so the firm invests up to the point that the marginal cost of investment inclusive of acceleration costs equals the marginal benefit of moving the project forward. With slight abuse of notation, we denote by $i(x,z)$ the resulting policy in the state space. We solve the HJB (9) numerically using finite-difference methods (Achdou et al., 2021). Appendix H collects all numerical routines.

Optimal Entry Firms optimally start new projects. For any drawn z, the firm can compute the NPV $v(0,z)$. If the NPV is positive, the firm will start the project. This pins down a rule for optimal entry:[34] start all projects with high enough return $z>\bar{z}$, where the break-even point $\bar{z}$ satisfies

$$v(0,\bar{z})=0 \tag{12}$$

Distribution of Projects Firms manage multiple projects. As firms invest, project stages change. The distribution of projects over stages thus evolves over time. Let $g(x,z)$ denote the steady state distribution. It satisfies a Kolmogorov Forward Equation (KFE)

$$\forall(x,z),\ \frac{\partial}{\partial x}\left[i(x,z)g(x,z)\right]=0 \tag{13}$$

The intuition for the KFE is simple.[35] The investment policy $i(x,z)$ is the speed according

[34]Soon, we will linearize the model and study first-order dynamics around a steady state. To do so while preserving an extensive margin, we introduce stochastic fixed costs of entry. The firm pays an entry cost distributed uniformly on $[0,\epsilon]$. The distribution allows us to convexify the problem and proceed with differentiable methods. We set ϵ to be small, so the fixed cost has insignificant effects on the steady state.

[35]A heuristic derivation of the KFE is as follows. Suppose we are in a steady state. A project (x,z) will receive an investment $i(x,z)$ in the next infinitesimal span of time ∂t. The "mass" of these projects is $g(x,z)$. Total investment is thus $i(x,z)g(x,z)$. Consider now two stages $x_0=x-\epsilon$ and $x_1=x+\epsilon$ in the vicinity of x. During the ∂t time span, projects in stage x_0 will move to stage x – a flow $i(x_0,z)g(x_0,z)$ into stage x. In turn, projects in stage x will move to stage x_1 – a flow $i(x_1,z)g(x_1,z)$ out of stage x. If we difference inflows from outflows, we would obtain by how much the measure of projects in stage x changes during a time span ∂t. As we bring x_0 and x_1 closer to x by letting $\epsilon\to 0$, the difference $i(x_0,z)g(x_0,z)-i(x_1,z)g(x_1,z)$ becomes $-\frac{\partial}{\partial x}\left[i(x,z)g(x,z)\right]$.

to which firms move along the x dimension. The flow of firms through a given point x is, therefore, $i(x,z)g(x,z)$: the product of the density of firms at this point and the velocity with which they travel. In steady state, the flow of firms through any given point must be constant, which is equivalent to equation (13).

Now recall that a a constant flow $f(z) \cdot 1\{z \geq \bar{z}\}$ of projects with productivity z enters the economy at each instant.[36] This implies that, in steady state, we must have $g(x,z)i(x,z) = f(z) \cdot 1\{z \geq \bar{z}\}$, which gives us

$$\forall(x,z),\ g(x,z) = \frac{f(z) \cdot 1\{z \geq \bar{z}\}}{i(x,z)} \tag{14}$$

Completion Times Optimal investment $i(x,z)$ tells us how quickly we are installing capacity. We also know how much more capacity we must install to complete the project. Therefore, we can compute how much time remains until project completion, which we denote $T(x,z)$. Solving the steady state Kolmogorov Backward Equation (see Appendix F) with an appropriate terminal condition $T(\bar{x},z) = 0$ yields

$$T(x,z) = \int_x^{\bar{x}} \frac{1}{i(y,z)} dy \tag{15}$$

An analogy to classical mechanics helps here. In a small time interval dt, a firm increases its project state by $dx = i(x,z)dt$. The time necessary to move dx units in the x dimension is, then, $dx/i(x,z)$. The total time until completion adds up these small time intervals until the projects reaches $\bar{x}$. (15) states this formally.

Steady State Solution Figure 4 shows steady state objects as a function of project stage ($100 \times x/\bar{x}$) for $z = 2$ and $z = 8$, and illustrative parameter values. Panel (a) shows the value function. It is monotonically increasing and reaches the terminal value pz at $x = \bar{x}$. Panel (b) shows optimal investment relative to $\bar{i}$. We can see that optimal investment is endogenously higher for more productive projects, reflecting the higher payoff at completion. Panel (c) shows the stationary distribution normalized by the flow of entrants, $g(x,z)/f(z)$, which is simply $1/i(x,z)$. There is a larger mass of active projects, given by the area below the curve, for lower values of z (relative to the flow of entrants $f(z)$), reflecting the fact that projects with lower z take longer to be completed. Finally, panel (d) shows the time until completion $T(x,z)$, which is higher for low-productivity projects until reaching zero at $x = \bar{x}$.

But this infinitesimal change in the mass of projects at stage x in time span ∂t has to be zero, because we are in a steady state. This is what the KFE (13) states.

[36]Note that the total mass of ongoing projects is not required to be 1 and is later going to vary over time following shocks.

Figure 4: Steady state objects as functions of project stage

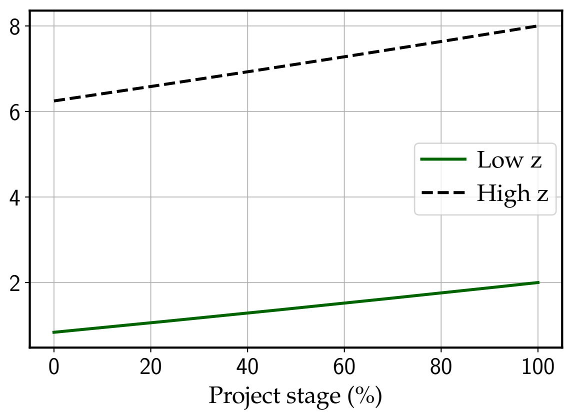

(a) Value function $v(x,z)$

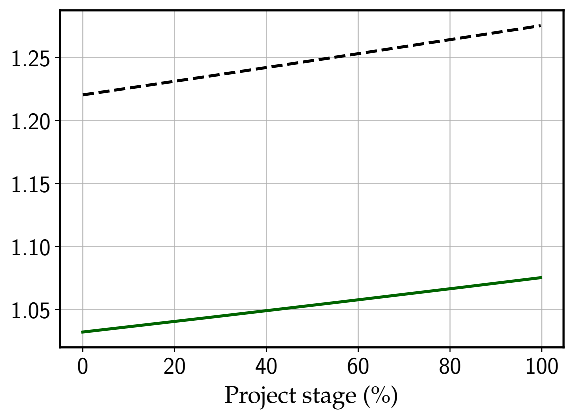

(b) Optimal investment $i(x,z)/\bar{\imath}$

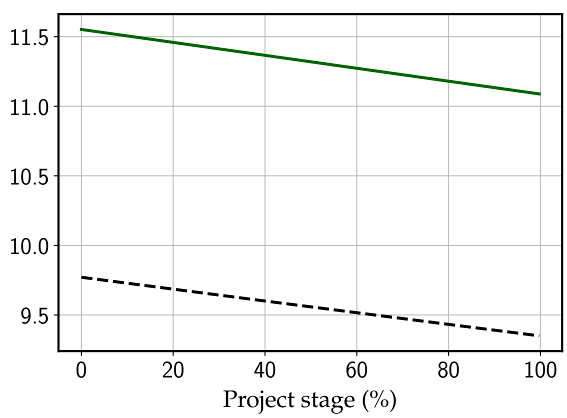

(c) Stationary distribution $g(x,z)/f(z)$

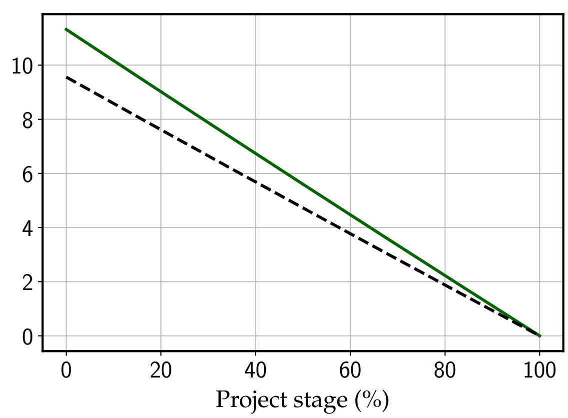

(d) Time until completion $T(x,z)$

Note: the figures plot the steady state solution of the model. Panel (a) plots the value function $v(x,z)$, panel (b) plots the investment policy $i(x,z)/\bar{\imath}$, panel (c) plots the stationary distribution $g(x,z)/f(z)$, and panel (d) plots the time to completion distribution $T(x,z)$. The x-axis is the project stage, defined as $100 \times x/\bar{x}$. Solid green lines represent low output projects ($z = 2$); black dashed lines represent high output projects ($z = 8$). Parameters are set at the median values obtained in our calibration in section 6.1: $r = 0.01$, $p = 1$, $\bar{x} = 1$, $\bar{\imath} = 0.083$, $\kappa = 30.1$.

5.1 First-Order Dynamics

To map our model to the empirical results in sections 4.2 and 4.3, we characterize the response of optimal investment policies to a credit shock. We model the banking shock from our empirical exercises as an increase in the discount rate r. Intuitively, as firms see themselves in a harder position to borrow, they raise their internal discount rates.

We model this shock as a perfect-foresight path for discount rates, i.e. $r(t) = r + se^{-\rho t}$. Let $\mathcal{V}(x,z,s,t)$ and $\mathcal{I}(x,z,s,t)$ be the value function and investment policy following the shock, respectively, which depend on the shock size s and calendar time t. We have the following result:

Lemma 1. *The first-order responses of the value function and investment policy to the shock are given by*

$$\mathcal{V}(x,z,s,t) = v(x,z) + se^{-\rho t}\hat{v}(x,z) + o(s)$$

$$\mathcal{I}(x,z,s,t) = i(x,z) + se^{-\rho t}\hat{i}(x,z) + o(s),$$

where

$$\hat{v}(x,z) = \left.\frac{\partial \mathcal{V}}{\partial s}\right|_{t,s=0} \quad \hat{i}(x,z) = \left.\frac{\partial \mathcal{I}}{\partial s}\right|_{t,s=0},$$

and $v(x,z)$ and $i(x,z)$ are the steady state value function and investment policy, respectively.

Proof. See appendix F for all proofs. □

Lemma 1 allows us to characterize value functions and investment policies analytically following the shock up to first-order as deviations from steady state. Analytical tractability stems from separating the time-dependent terms ($e^{-\rho t}$) from terms related to project stage x and productivity z. We start by characterizing the response of the project's value.

Proposition 2. *The first order response of project value is given by*

$$\hat{v}(x,z) = -\left(\frac{1-e^{-\rho T(x)}}{\rho}\right)e^{-rT(x)}pz + \int_0^{T(x)} \left(\frac{1-e^{-\rho\tau}}{\rho}\right)e^{-r\tau}\left[i(y(\tau),z) + \kappa\left(i(y(\tau),z)\right)\right]d\tau, \tag{16}$$

and, as $\rho \to 0$,

$$\lim_{\rho\to 0}\hat{v}(x,z) = -T(x)e^{-rT(x)}pz + \int_0^{T(x)} \tau e^{-r\tau}\left[i(y(\tau),z) + \kappa\left(i(y(\tau),z)\right)\right]d\tau,$$

where $y(\tau)$ is the stage of a project that starts in x after τ periods, i.e.

$$\dot{y}(\tau) = i(y(\tau),z), \text{ subject to } y(0) = x.$$

The limiting case of a one-time shock to discount-rates ($\rho \to 0$) is particularly illustrative. Project value changes by its **duration**. This is a familiar object to first-year MBAs, capital budgeting specialists, fixed income investors, and readers of asset pricing textbooks. To compute the change in project value, we take a weighted sum of time periods over the remaining life of the project. We weigh each time period by the NPV of cash flows accruing in the time period.[37] We use the steady state investment flows and remaining time to build to compute duration.

Intuitively, duration gives us how long it will take to reap back the cash flows of the project. Different from standard fixed income securities, a project requires outlays over its lifetime in order to yield a large terminal payoff. Similar to many fixed-income instruments, the cash flow paid at maturity is large relative to near-term coupons. As interest rates rise, project values fall. But they fall ***more*** for projects whose cash flows will only be received further out in the future. Just like longer duration bonds suffer more from interest rate increases, ***longer time to build projects suffer more from increases in discount rates***.

Duration is a crucial input for understanding the response of investment to discount rate shocks, as the following proposition illustrates.

Proposition 3. *The first order response of the optimal investment policy is given by*

$$\hat{i}(x,z) = \frac{1}{\kappa''(i(x,z))} \cdot \frac{\partial}{\partial x}\hat{v}(x,z), \tag{17}$$

where

$$\frac{\partial}{\partial x}\hat{v}(x,z) = \frac{1}{i(x,z)}\left[v(x,z) + (r+\rho)\hat{v}(x,z)\right] \tag{18}$$

Proposition 3 characterizes how to re-optimize investment plans following a change in discount rates. Equation (17) simply equates the marginal cost to the marginal value of reducing the project's duration, which in turn captures the marginal value of bringing the project closer to completion. How does this change alter the discount rate shock? Equation (18) shows two opposing effects. First, investors become more impatient, and therefore value receiving cash flows earlier. This effect is positive, and captured by $v(x,z)$. Second, there is a negative effect arising from revaluing future cash flows at the higher discount rate. This

[37] An analogy with classical mechanics once again helps with the notation. Define $\tau = T(x) - T(y)$, so that $d\tau = {}^{dy}/_{i(y,z)}$. From our discussion on expected completion times, we can interpret $d\tau$ as how long it takes to install dy units of capital when we are optimally installing capital at a rate $i(y,z)$. Note that $\tau = 0$ corresponds to stage x. And given optimal investment, can define the stage y where the project will be in τ periods as $\dot{y}(\tau) = i(y(\tau),z)$, where $y(0) = x$. Without defying any laws of physics, we have just renormalized space and time. Suppose we start counting time when the project is at stage x, and normalize that instant to $\tau = 0$. The project will be in stage $y(\tau)$ after τ periods. How quickly does it get there? At an optimum, the velocity of completion is given by $i(y(\tau),z)$ at each point $y(\tau)$.

is standard asset pricing: as discount rates rise, project values fall.[38] The following corollary shows that the latter effect becomes negligible as projects approach completion.

Corollary 4. *For x sufficiently close to $\bar{x}$, $\hat{i}(x,z) > 0$.*

Proof. Note that $\hat{v}(x,z)$ becomes arbitrarily small as x approaches $\bar{x}$, as it is the integral of a bounded function over a vanishingly small interval (i.e. $T(x) \to 0$ as $x \to \bar{x}$). On the other hand, $v(x,z)$ approaches $pz > 0$. Therefore, it follows from (18) that $\hat{i}(x,z)$ must eventually be positive. □

Corollary (4) rationalizes our intensive margin results. For projects close to completion, the firm increases investment following a discount rate shock. The intuition is simple: remaining time to build is small, and a big payoff awaits the firm if it speeds up the project. Equation (18) then shows it is optimal to invest more in the project. For projects with long time to build, the negative effect from large duration dominates, and it might become optimal to slow down the project.

Corollary (4) also qualifies the standard intuition that increases in discount rates necessarily cause investment to fall. At the *project-level*, investment may very well *rise*. How long it takes for the project to yield cash flows matters. This depends on how quickly the firm is building the project; and on how costly it is to accelerate – hence the denominator in (17) and (18). Where the project currently is in its lifecycle (x) summarizes these considerations.

Our results borrow intuition from concepts in fixed income pricing. In Appendix F.2, we demonstrate that this is not a loose mapping. The underlying economics between accelerating project investments and accelerating bond coupons is quite similar, and reflects a more general insight: changes in the time value of money incentivize agents to shift cash flows (or, more broadly, actions) over time.

Figure 5 illustrates our propositions by showing the impact responses of the value function $\hat{v}(x,z)$ and optimal investment $\hat{i}(x,z)$ to the shock. Panel (a) shows $\hat{v}(x,z)$. Note that it is strictly negative for $x < \bar{x}$, as the payoff upon project completion is more strongly discounted after the shock. In short, all projects fall in value. But project value falls by more for

[38] Equation (18) additionally provides a simple capital-budgeting rule: is the marginal dollar worth more ***inside*** or ***outside*** the project? Note that the proportional change in project value following the shock is $\hat{v}(x,z)/v(x,z)$. It is the return an equityholder would obtain by investing the marginal dollar into the project. The firm is not constrained to invest this marginal dollar into the project. Instead, it could invest in some corporate savings account, or simply return the marginal dollar to its equityholder. The return on this outside option $-(r+\rho)^{-1}$. If this exceeds the return on the project, the firm should invest less in the project and reallocate towards the outside option. In our setup, the outside option is a fixed-income perpetuity. We can interpret it in multiple ways. It could stand for a corporate cash account earning the riskless rate; it could also stand for the safe returns an equityholder would have earned on a marginal dividend payment; it could stand for the return on the market in a setting with aggregate risk, where we then interpret r as a riskless return under the risk-neutral measure. In any case, what matters is that this outside option provides the firm with a benchmark, risk-adjusted opportunity cost of investing the marginal dollar on the project.

Figure 5: Responses of the value function and optimal investment to discount rate shock, as functions of project stage

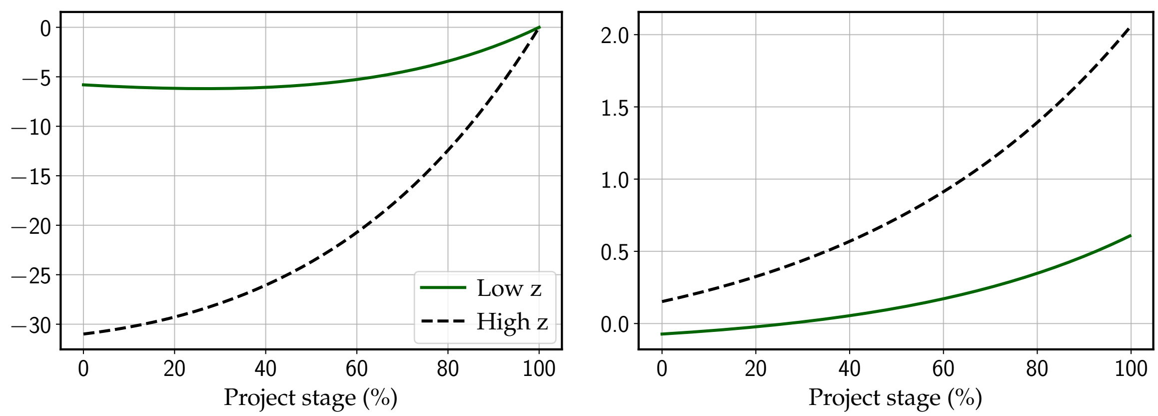

(a) Response of the value function $\hat{v}(x,z)$

(b) Response of optimal investment $\hat{i}(x,z)$

Note: the figure plots the first-order responses of the value function $\hat{v}(x,z)$ (panel a, left) and of the investment policy $\hat{i}(x,z)$ (panel b, right) following a shock to the discount rate. By definition, the y-axis is read as percent-deviations from steady state values. The x-axis is the project stage, defined as $100 \times x/\bar{x}$. Solid green lines represent low output projects ($z = 2$); black dashed lines represent high output projects ($z = 8$). Shock persistence is set at $\rho = 0.17$ and shock size is set at $s = 0.02$, as in appendix table A16. Remaining parameters are set at the same values as in figure 4.

projects that are further away from completion, illustrating the duration intuition of Proposition 2. Panel (b) shows $\hat{i}(x,z)$. For low levels of z, projects are accelerated if they are beyond a threshold stage (around 30-35% of total construction). For high levels of z, it may be optimal to increase investment not only for x close to $\bar{x}$, but even for newly started projects.

To understand changes in *firm-level* investment, we have to further account for the distribution of projects over their completion stages $g(x,z)$; as well as potential changes in this distribution over time following the shock, including entry effects. Proposition (3) only speaks to ongoing projects. It qualitatively rationalizes our intensive margin results, but it is silent on extensive margin effects. Our next result addresses this gap.

Proposition 5. *Let*

$$\bar{\mathcal{Z}}(t,s) = \bar{z} + se^{-\rho t}\hat{z} + o(s),$$

where

$$\hat{z} = \left.\frac{\partial \bar{\mathcal{Z}}}{\partial s}\right|_{t,s=0}.$$

The first-order response of the entry threshold $\bar{z}$ to the shock is given by

$$e^{-rT(0,\bar{z})} \times p\hat{z} = -\hat{v}(0,\bar{z}) \tag{19}$$

Proposition (5) states that the change in the threshold follows a standard NPV logic.

Given the rise in discount rates, the value of unbegun projects falls by $\hat{v}(0, \bar{z})$. Projects at $\bar{z}$ are no longer positive NPV. To select projects with positive NPV, we need to raise our hurdle by $\hat{z}$. We raise it such that the cash flows $p\hat{z}$ on the marginal projects $(\bar{z}, \bar{z} + \hat{z}]$, once discounted properly, makes up for the lost value $\hat{v}(0, \bar{z})$. To discount cash flows, we use the original discount rate r; and take the time to build of these marginal projects as the relevant time horizon for receiving the cash flows. Importantly, note that $\hat{z} > 0$, and so the flow of newly started projects decreases after the increase in discount rates, in line with our empirical results (and with standard investment theory).

5.2 Aggregation

Towards quantifying the implications of endogenous time to build to firm-level investment, we describe how we aggregate across projects.

The law of motion for the capital stock K is

$$\dot{K}(t) = I^K(t) - \delta K(t)$$

where instead of accounting investment, the gross increment in productive capacity is given by fixed capital formation $FCF(t)$. We have

$$I^K(t) = \int \underbrace{z}_{\text{capital output}} \cdot \underbrace{g(\overline{x}, z, t) \cdot i(\overline{x}, z, t)}_{\text{flow of finished projects yielding } z \text{ units of capital}} dz$$

that is, productive capital grows because more projects are completed. Since projects can differ in their output of capital, we aggregate across the distribution of output.[39] Accounting investment, in turn, is given by aggregating investment flows across all ongoing projects

$$I^{acc}(t) = \int \int g(x, z, t) \cdot i(x, z, t) \; dzdx$$

[39]In steady state, equation (14) implies that accounting investment is independent of time to build. In turn, fixed capital formation depends on time to build via its effect on the entry threshold for projects. Formally, letting $M = \int f(z)dz$ be the total mass of entrants, we have $I^{acc}_{ss} = M \cdot \bar{x}$. In turn, since only projects with z above $\bar{z}$ enter, we have

$$I^K_{ss} = \mathbb{E}\left[z | z \geq \bar{z}\right] P(z \geq \bar{z})$$

6 Taking the Model to the Data

6.1 Quantifying Endogenous Time to Build

Our reduced form exercises are inconsistent exogenous gestation lags, but are silent beyond local causal estimands. However, many pressing questions require quantifying counterfactual. What if contractual arrangements became so stiff that time to build indeed became an exogenous constraint on capital formation? Is investment more or less volatile once firms have the option to accelerate projects? What are the implications of endogenous gestation lags for stabilization policy?

Answering these questions require a quantitative model. Fortunately, the structure we built up above is up to the task, once sufficiently enriched. We calibrate the model to match our measurements from Section (3) and the estimated elasticities from Section (4). First, we confront the sizable heterogeneity in project costs and gestation lags. To do so, we let projects be heterogeneous in their sizes $\bar{x}$. We assume project sizes are drawn according to $\log \bar{x} \sim \mathcal{N}\left(\mu_{\bar{x}}, \sigma_{\bar{x}}^2\right)$. We set the parameters of the lognormal distribution to match the 10th, 25th, 50th, 75th, and 90th percentiles of the empirical distribution of log costs.

To match the relationship between project size and time to build, we postulate that each project's costless rate of completion $\bar{i}$ is drawn according to $\log \bar{i} \sim \mathcal{N}\left(\beta_0^{\bar{i}} + \beta_1^{\bar{i}} \log \bar{x}, \sigma_{\bar{i}}^2\right)$. This parametrization assumes that mean investment rates depend log-linearly in project sizes, but the standard deviation of investment rates is homoscedastic. Higher draws of $\bar{i}$ for any fixed $\bar{x}$ imply a quicker time to completion. For instance, if time to build was purely exogenous, a project's (random) gestation lag would be uniquely determined as $\bar{T} = \bar{x}/\bar{i}$. We set the parameters of the $\bar{i}$ distribution to match the average log time to build within each log size quintile. Figure A15a shows our results.

For analytical convenience, we parametrize the distribution of project returns $f(z)$ as a Pareto distribution with minimum draw z_{min} and tail exponent α. We use accounting data on firm revenue to proxy for project gross returns. We trim 1% of the revenue distribution to remove outliers, and run log-rank regressions on the n largest firms following Gabaix and Ibragimov (2011) to obtain the Pareto exponent α. Figure A15b shows our results.

We use the intensive and extensive margin semi-elasticities $(\beta_{int}, \beta_{ext})$ estimated in sections 4.2 and 4.3 to calibrate how endogenous time to build is.[40] The chief difficulty is how to map the reduced form regression exercise to the theoretical framework. Introducing banks, bank relationships, and banking shocks would significantly increase model complexity, with questionable gains. Instead, following our earlier theoretical results, we map the banking shock to an increase in firm's discount rates. But how can we map the quantitative magnitudes of the shock size across data and model?

[40]To be precise, we use the point estimates in columns (5) of Tables 4 and 5.

Our solution is to interpret the size of the model shock as a parameter to be calibrated. For any shock size s and persistence ρ, we can obtain the model-based semi-elasticities for project completion hazards and project starts. We estimate ρ by running panel autoregressions of Exposure_{ft} on its first lag, including firm and year fixed effects. To circumvent Nickell (1981) bias, we run a battery of different dynamic panel data estimators: OLS without fixed effects, the Han and Phillips (2010) GMM estimator, and several variations of the system GMM estimators of Arellano and Bover (1995) and Blundell and Bond (1998). Reassuringly, we find similar estimates across all models; and convert them into quarterly figures to arrive at our estimate of ρ. Figure A15c shows our results.

We are left with two parameters: the shock size s and the time-to-build endogeneity parameter κ. For any pair of values (s, κ), we can obtain values for the semi-elasticities in the model, so we can calibrate these parameters to match our empirical results using a standard minimum distance routine. Though all moments are matched jointly, the intuition for identification is as follows. β_{ext} is uninformative about κ, but informative about s. In turn, β_{int} is informative about both parameters; but β_{ext} and β_{int} move in opposite directions as we change s. The opposing responses to the shock and the differential informativeness of each regression coefficient to κ thus drive identification. Appendix Table A16 collects our calibrated parameters. We estimate a shock size $s = 2\%$, and a acceleration cost $\kappa = 3$. Appendix Table A17 reports that the model calibrated model fits well the targeted moments.

6.2 A Sharper Test: Heterogeneous Effects

The model lends itself to an additional, sharper test: the effect of discount rate shocks on investment is heterogeneous, depending on the stage of the project at the time that the shock hits.

Consider Panel (b) of Figure 5. The plot shows that investment rises when discount rates rise for either (i) sufficiently high output projects, or (ii) sufficiently late-stage projects. But the same plot reveals more. Fixing a project, the effect of the shock would have been greater had the project been at a later stage relative to an earlier stage. This holds for any z, and thus provides a sharper test of the model's mechanism.

Testing for heterogeneous treatment effects of this sort runs into two challenges. First, we do not observe a project's value put in place (the empirical counterpart of the model stage x) in the data. To add insult to injury, we further run into the fundamental problem of causal inference. At the time the shock hits, the project cannot simultaneously be in two distinct stages. We cannot observe counterfactual stages beyond the measured stage at the time of the shock.

Despite the econometric challenges, we attempt to provide the sharper test. We address the measurement challenge by constructing a proxy measure of project stage. We define the proxy stage as the project's age when entering the hazard regression sample[41], relative to an ex ante, predicted measure of its gestation lag

$$\text{Proxy Stage}_p = \frac{\text{Initial Sample Age}_p}{\text{Predicted Time to Build}_p}$$

The intuition is that, fixing a project's expected time to build, older projects are expected to be at a more advanced stage; conversely, fixing a project's age, projects with shorter average gestation lags are expected to be at a more advanced stage.

Since we do not observe expected time to build, we predict it in a preliminary step. We isolate all completed projects in CapexDx that are not included in the regression sample. For these projects, we regress time to build on exogenous predictors. We estimate different regressions, and compare them by assessing their R^2 in our regression sample. We pick the model with highest out-of-sample R^2, which turns out to be a model including industry dummies. Appendix Table A11 and Appendix Figure A11 expand on the prediction exercise.

Having estimated the proxy stage variable, we standardize it and include it along with its interaction with the firm exposure variable in model 3. Table 6 collects results. Across specifications, the results are supportive of the model's predictions. For the project at the mean proxy stage, a one-standard deviation increase in exposure raises the hazard by roughly 10%. If the project stage had been one standard deviation above (below) the mean, the effect would have been increased (reduced) by roughly 6.5%, so the total effect of the should would have been 16.5% (3.5%). In Appendix Table A12, we find similar results if we don't scale initial age by predicted time to build, suggesting our results are not driven by the prediction step.[42]

How does the model fare relative to the data? Figure 6 plots the results. The red dashed line plots the cumulative distribution of the proxy stage variable in the model.[43] The blue solid line plots the impact change in completion probability by the proxy stage variable. The results are quite encouraging: the effect is a linear function of the proxy stage, consistent with our regression specification. Quantitatively, a unit-increase in proxied stage increases the hazard by 8%. Qualitatively, we interpret the result as suggestive evidence that data and

[41] We prefer to use age when entering the hazard regression over age at the time of shock. Some projects may enter and leave the sample entirely before or after the start of the asset quality review audits; therefore we cannot define age at time of the shock for these projects, but we can always define the age at the time it enters the sample. Moreover, since we include initial age and year fixed effects in the regressions below, one can always partly "reconstruct" what would have been age at the time of the shock.

[42] The effects are significant at conventional levels regardless of the proxy we use for project stage, but we must be cautious interpreting the inferential step when scaling by predicted time to build given we do not account for sampling uncertainty in the prediction step. The fact that the estimates remain significant in Appendix Table A12 suggests that ignoring the noise introduced by the prediction step is not substantially compromising inference.

[43] We use the unconditional expected time to build in the model. The reason is that the model is minimalist relative to the data. In the model, conditioning on $\bar{i}$, $\bar{x}$ or z would strip out most of the variation. In the data, however, the predictive step is much noisier. Even predicting with industry bins, there remains sizable unexplained variation. The unconditional distribution thus provides a better match to the noisy reality of the data.

Table 6: Effects of Exposure by Proxied Stage

	(1)	(2)	(3)	(4)
Exposure	0.099^{***}	0.095^{***}	0.106^{***}	0.104^{***}
	(0.026)	(0.026)	(0.025)	(0.025)
(Age/ETTB) x Exposure	0.063^{***}	0.065^{***}	0.070^{***}	0.071^{***}
	(0.022)	(0.022)	(0.023)	(0.023)
Initial Age FE	✓	✓	✓	✓
Firm FE	✓	✓	✓	✓
State x Year FE	✓	✓	✓	✓
Sector x Year FE	✓	✓		
Industry x Year FE			✓	✓
Firm Controls		✓		✓

Notes: the table reports PPML estimates of the semi-elasticity of project-level hazards to firm exposure to the Asset Quality Review. The sample consists of projects observed quarterly between 2013-2019. Firm exposure is defined as the weighted-average of lender exposure, where weights correspond to outstanding loans in the previous year. The proxy for stage is "Age/ETTB", constructed as age of the project when first observed in the sample divided by its predicted time to build. This measure is winsorized at the 1st and 99th percentile, and then standardized before being included in the regression model. See the main text and Appendix Table A11 for details on the prediction step. The coefficient on exposure is interpreted as the percent effect on the outcome for a one standard-deviation increase in firm exposure for the project with mean proxy stage. The coefficient on the interaction is interpreted as the change in the main effect resulting from a one-standard deviation level change in proxy stage. For example, the first row in column (1) reports that moving from a firm with average exposure to a firm with one-standard deviation higher exposure increases the hazard by 9.9% for the project at the mean proxy stage. The second row in column (1) reports that if the project had been one standard deviation above (below) the mean stage at the time of the shock, the effect would have incresed (declined) by 6.3%, for a total effect of 16.2% (3.6%). In order to avoid strong collinearity concerns between the age spline, initial age, and the time fixed effects, we omit the age spline, control non-parameterically for the initial age, and further include dummies for quartiles of the proxy stage. Firm controls include sales, sales growth, profitability, debt over equity, book leverage, equity issuances, cash over liquid assets, and debt over EBITDA. All firm controls are winsorized at the 1st and 99th percentile, and then standardized. All firm controls are lagged by one-period. The reduction in sample size from columns (1) and (2) to columns (3) and (4) results from separation of additional observations by the more granular industry-by-year fixed effects relative to the sector-by-year fixed effects. Standard errors are clustered at the firm level and reported in parentheses. $*$, $**$, and $***$ denote statistical significance at the 10%, 5% and 1% levels, respectively.

Figure 6: Heterogeneous Effects in the Model

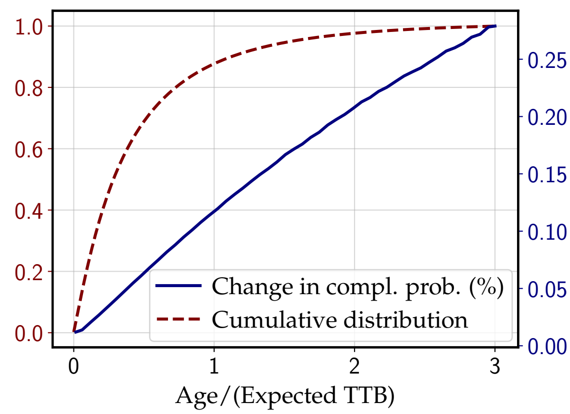

Notes: the figure plots the cumulative distribution of the proxy stage variable in the model (red dashed line) and the heterogeneous effects of a discount rate shock on hazards by the proxy stage variable (blue solid line). Expected TTB is the cross-sectional average gestation lag across all model projects, i.e. averaging over the joint distribution of $(\bar{x}, \bar{\imath}, z)$. Age is measured in quarters. All parameter values, including shock size and persistence, are as in appendix table A16.

model speak to each other well when it comes to heterogeneous effects by project stage.

7 Endogenous Time to Build and Investment Fluctuations

Equipped with our calibrated model, we study the implications of endogenous time-to-build for investment dynamics.

7.1 Partial Equilibrium

We first examine the response of firm-level investment to a shock in discount rates (as in our empirical setup) and to a shock in future cash flows. We do so first for our calibrated model. Then, we consider an extreme setting where gestation lags become completely exogenous. Formally, we let $\kappa \to \infty$, a world in which contractual arrangements are completely stiff or technological constraints on construction are severe.

Figure 7: Impulse response functions of investment to discount rate shock.

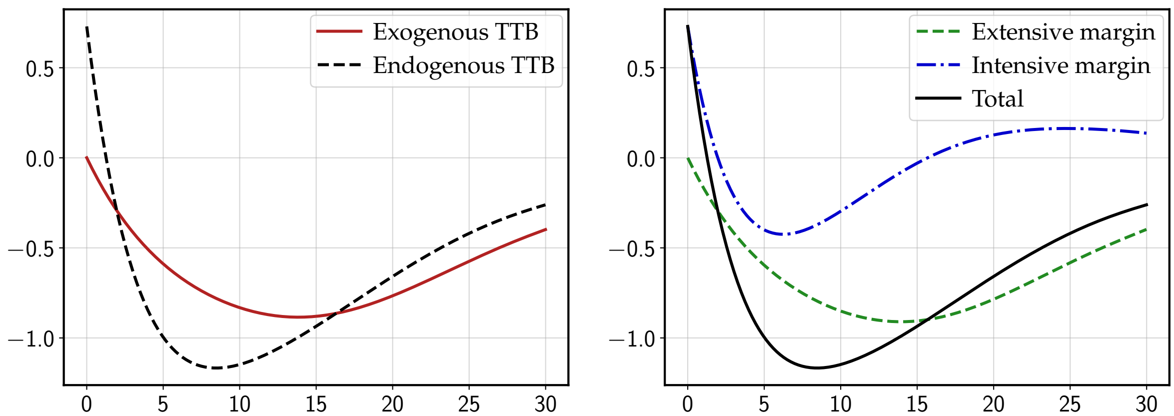

(a) Total IRF for endogenous versus exogenous time to build model

(b) Decomposition of endogenous time to build IRF between intensive and extensive margins

Note: the figure plots (quarterly) impulse response functions (IRFs) of aggregate investment to a discount rate shock in partial equilibrium (i.e. fixing p to its steady state level following the shock). The y-axis measures percent-deviations from steady state values. Panel (a) plots the IRFs for the exogenous time to build model (red solid line) and for the endogenous time to build model (black dashed line). Panel (b) decomposes the total IRF (black solid line) from the endogenous time to build model into the extensive margin (green dashed line) and the intensive margin (blue dot-dashed line). Shock persistence, shock size, and remaining parameters as in appendix table A16.

Figure 7 illustrates our results for a discount rate shock, holding fixed p. In panel (a), we plot the response of investment when time to build is endogenous (black dashed line) versus exogenous (red solid line). In a world with exogenous time to build, investment falls, reaches a trough 12-16 quarters later, and then returns to steady state, generating a hump-shaped impulse response function (IRF). The dynamics are driven exclusively by entry. Since investment flows are constant conditional on project entry, there is no impact (i.e. time 0) effect on investment. Higher discount rates lower the value of all projects and raise the threshold for a project to be positive NPV. Firms thus start fewer projects today. But fewer projects started today mean lower future investment on ongoing projects. This drives the downward-sloping section of the IRF. Eventually, the discount rate shock starts to revert, and entry starts to return to its normal level. This drives the upward-sloping section of the IRF. Previous theoretical work corroborates this result: to first-order, exogenous TTB is first-order equivalent to investment adjustment costs (Lucca, 2007; Bloesch and Weber, 2023).

In a world with endogenous time to build, investment initially *rises*. The intuition follows from Proposition 3 and Corollary 4. There is a mass of projects close to completion when the shock hits. They are sped up, which leads to an upsurge in investment flows. The extensive margin component is also present, but is zero at impact. Hence, for the first few quarters, we only have the acceleration effects taking place. The net effect is the short-lived investment spike after the rise in discount rates.

Investment collapses shortly after impact. Two reinforcing mechanisms drive the decline. The first is the extensive margin response. The mechanism is the same described in the exogenous time to build case. The second is the intensive margin effect, which becomes negative. To decompose effects, panel (b) splits the endogenous time to build IRF into an intensive margin IRF and an extensive margin IRF. The former fixes the composition of projects and accounts only for changes in investment rates of ongoing projects; the latter only accounts for entering projects and their subsequent investment flows.

To understand the reversal in the intensive margin effects, recall from proposition (3) that while some investments may be sped up, some might be slowed down. As projects close to maturity are completed, the distribution of remaining projects becomes skewed towards early-stage projects. These are precisely the projects for which lower investment is optimal, even in steady state – see Figure 4. The distributional change thus implies less capex on earlier-stage projects, driving the drop in investment. The intensive and extensive margins add to each other and generate a steeper fall in investment in the model with endogenous time to build.

Endogenous time to build also delivers faster convergence to the original steady state. The effect is partly driven by the intensive margin reverting back faster than the extensive margin; but partly driven by the intensive margin becoming positive later on. Investment in ongoing projects quickly booms, shortly collapses, and then overshoots relative to steady state. These boom-bust dynamics are reminiscent of the oscillatory behavior present in investment accelerator theories (Kalecki, 1935; Samuelson, 1939). While they are hardwired in these earlier models, they are microfounded by optimal investor behavior in our setting.

The rebound in the intensive margin results again from compositional changes over time. As the discount rate shock fades away, investment flows return to their steady state levels. Over time, the distribution of project shifts again towards later-stage projects – if there were no entry, all projects would eventually be completed. Even in steady state, it's optimal to invest more in these projects. Overall investment in ongoing projects eventually rises above the steady state level as the distribution shifts towards later-stage projects.

Figure 8 moves to an increase in terminal payoffs p, holding fixed discount rates. Once again, we consider an AR(1) perfect-foresight shock.[44] In the exogenous world, the IRF is analogous to the case of discount rate shocks, except the positive payoff shock raises entry. In the endogenous world, there is a short-run spike in investment. As projects are completed, the distribution skews to earlier-stage projects with smaller steady state investment. The intensive margin effect thus drops below zero, until it converges back to steady state. The resulting effect is a short-run boom followed by a long-lived mild slump.

Given investment is more sensitive to shocks in a world with endogenous time to build,

[44] One can interpret the price shock as a productivity shock under a competitive market structure and constant returns to scale.

Figure 8: Impulse response functions of investment to price shock.

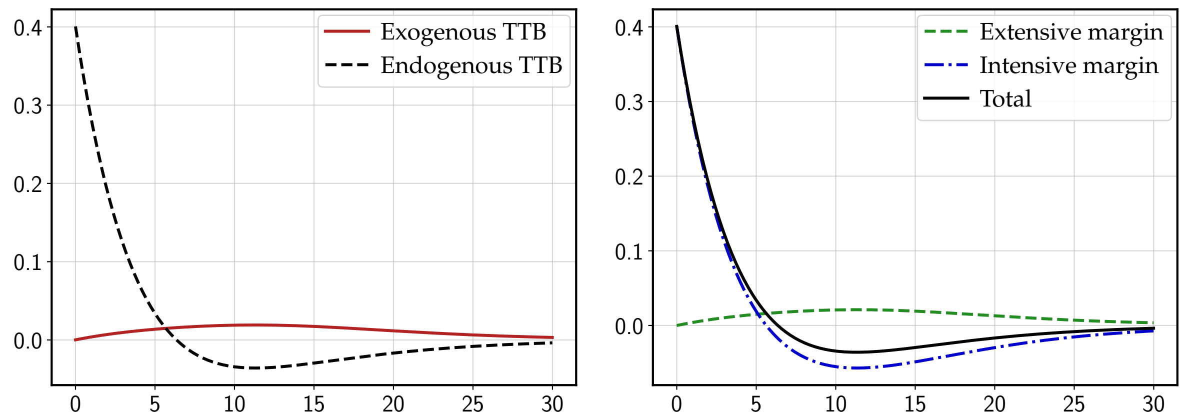

(a) Total IRF for endogenous versus exogenous time to build model

(b) Decomposition of endogenous time to build IRF between intensive and extensive margins

Note: the figure plots (quarterly) impulse response functions (IRFs) of aggregate investment to a price shock in partial equilibrium (i.e. fixing r to its steady state level following the shock). The y-axis measures percent-deviations from steady state values. Panel (a) plots the IRFs for the exogenous time to build model (red solid line) and for the endogenous time to build model (black dashed line). Panel (b) decomposes the total IRF (black solid line) from the endogenous time to build model into the extensive margin (green dashed line) and the intensive margin (blue dot-dashed line). Shock persistence, shock size, and remaining parameters as in appendix table A16.

casual intuition suggests it should also be more volatile. That is indeed the case: moving from exogenous to endogenous time to build increases investment volatility by 30%. Endogenous time to build allows models to account for investment volatility not by appealing to volatile or large shocks, but rather by generating endogenous amplification. The following proposition provides a lens to interpret the increased responsiveness:

Proposition 6. *Fix any $z > \bar{z}$. For a sufficiently small and temporary shock, the cumulative effect on the intensive margin of investment is zero.*

Proposition (6) shows that the intensive margin IRF has no cumulative effects. The statement is quite general: it holds for any shock, so long as the disturbance does not prevent the distribution from returning to its steady state. Endogenous time to build allows firms to reshuffle investment flows over time. It adds time series variation around a null cumulative effect. Therein lies the source of higher volatility.

Figures 7 and 8 highlight that our quasi-experimental setup is particularly well-suited to disentangle endogenous and exogenous time to build settings. Discount rate shocks generate positive effects on the intensive margin, but negative effects on the extensive margin. This underpins our an empirical test based on reduced-form exercises. Further, it facilitates calibrating the degree of endogenous time to build. If the extensive and intensive margin effects went in the same direction – as in a price shock – we would only be able to pin down the degree of endogenous time to build from the relative effect of these margins, rather than the direction of the effects given a shock size.

7.2 State-Dependent Monetary Policy

Recall from Figure 6 that the effects of a discount rate shock on project investment flows are heterogeneous. Effects are stronger for projects that are closer to completion – a central model prediction that is borne out in the data. Consequently, the distribution of projects at the time of the shock determines the aggregate effect. What if the economy had recently experienced an influx of new projects? Or what if the economy had been recovering from a period of sluggish starts, with the distribution being skewed towards older projects?

Figure 9 studies these settings. In Panel (a), we plot three distributions of projects over stages. The first is the steady state distribution (black solid line). The IRF in Figure 7 reflects the impulse starting from this distribution. We then consider two alternative scenarios. In the first, the economy has experienced a sharp influx of new projects, skewing the distribution towards less mature projects (green dashed line). In the second, the economy instead has experienced a prolonged period of no starts, with the distribution skewing towards old projects as a result (blue dot-dashed line).

Panel (b) plots the response of aggregate investment to the discount rate shock of Figure 7, but starting from each of the skewed distributions. When the economy starts filled with mature projects, the initial investment spike is larger. In turn, when the economy starts filled with fledgling projects, the impact response turns negative. The distributional dependence is so strong that the effect on impact can even flip signs depending on initial conditions.

Figure 9 carries lessons for monetary policy. The effects of interest rate changes on investment are state-dependent – an effect also documented in the context of mortgage refinancing (Berger et al., 2021; Eichenbaum et al., 2022) and durables purchases (McKay and Wieland, 2021). Interest rate hikes (cuts) will be less contractionary (expansionary) on impact when the distribution is skewed towards more mature projects. In short, the investment channel of monetary authority may face headwinds, which are absent when time to build is exogenous.

Figure 9: Endogenous Time to Build and State-Dependent Monetary Policy

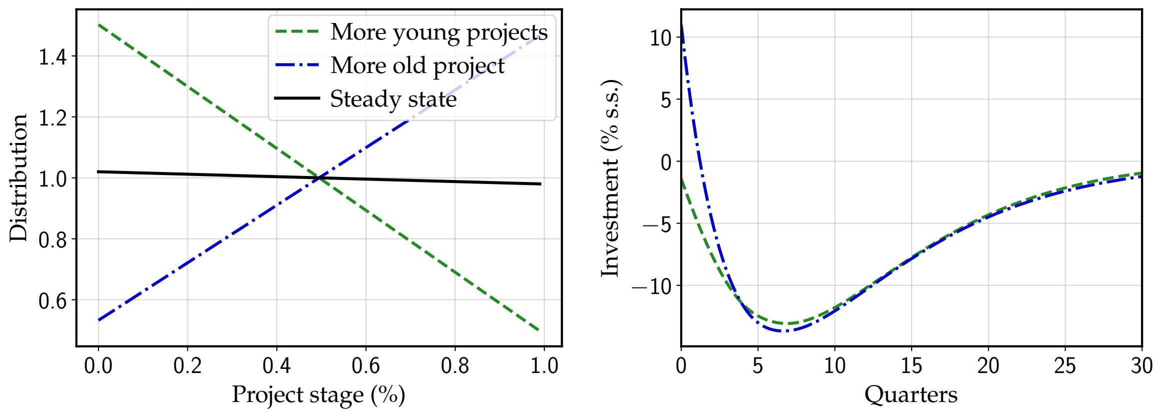

(a) Distributions of Projects Across Stages

(b) Investment IRF to r Shock Across Different Starting Distributions

Note: the figure illustrates the state-dependent effects of monetary policy in a world of endogenous time to build. Panel (a) plots the steady state distribution (black solid line), and two "counterfactual" distributions at the time of the shock: one skewed towards younger projects (green dashed line), and one skewed towards older projects (blue dot-dashed line). Panel (b) plots (quarterly) impulse response functions (IRFs) of aggregate investment to a discount rate shock in partial equilibrium (i.e. fixing p to its steady state level following the shock). The y-axis measures percent-deviations from steady state values. The green dashed line plots the IRF when the distribution is skewed towards younger projects at the time of the shock. The blue dot-dashed line plots the IRF when the distribution is skewed towards older projects at the time to the shock. The IRF for the steady state distribution is the black dashed line in Panel (a) of Figure 7. Shock persistence, shock size, and remaining parameters as in appendix table A16.

7.3 Fiscal Policy: Reshuffling Aggregate Demand over Time

If monetary policy can affect investment by changing discount rates, fiscal policy can pull its weight by changing cash flows. Indeed, investment tax credits are a widely used policy lever in recessions. They primarily spur investment by creating a tax shield, effectively generating a positive cash flow shock relative to the counterfactual world in which the tax credit had not been in place. But when gestation lags are endogenous, investment responses to cash flow shocks depend not only on whether projects are close or far from completion; but also on the time profile of the shocks. Motivated by these observations, we study the effects of introducing an unanticipated investment tax credit. Firms know once the policy is in place that it will last for 10 quarters. Figure 10 shows the resulting IRFs of accounting investment and productive capital.

In panel (a), the red dashed line plots the IRF of accounting investment in the exogenous time to build model. The positive cash flow shock induces project entry. Investment slowly rises as firms build nascent projects. As the tax credit expires, investment starts to returns to steady state. The long-lived effects of the shock result from distributional dynamics: since projects were started, the IRF accounts for cash flows that will be spent on these projects,

Figure 10: Endogenous Time to Build and Investment Tax Credits

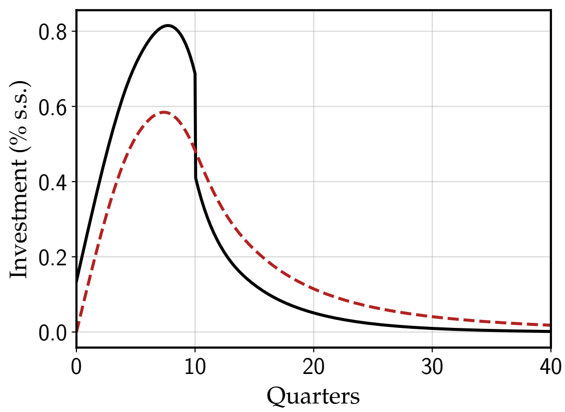

(a) I^{acc} IRF, Endogenous vs Exogenous TTB

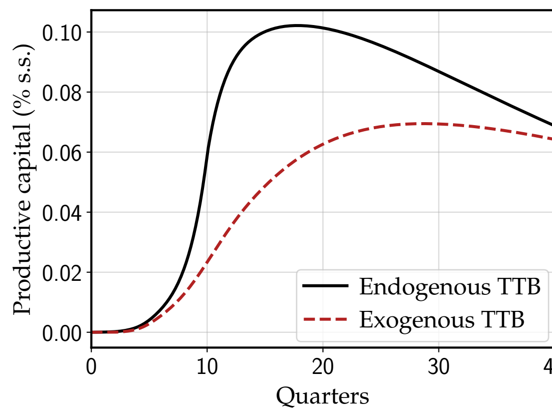

(b) K IRF, Endogenous vs Exogenous TTB

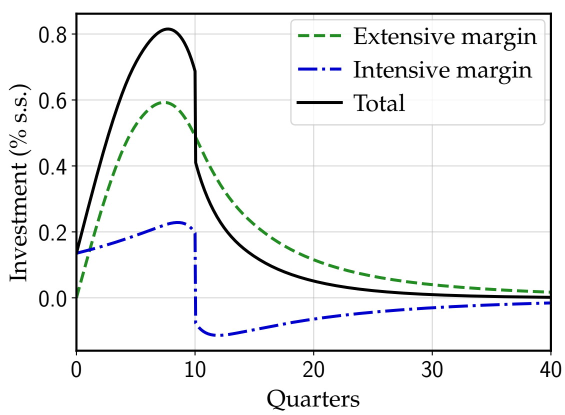

(c) IRF Decomposition: I^{acc}

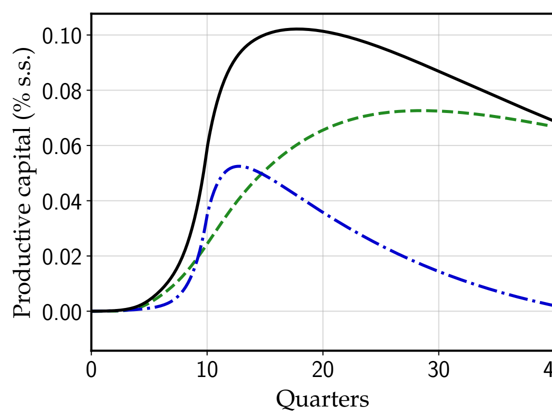

(d) IRF Decomposition: K

Note: the figure plots (quarterly) impulse response functions (IRFs) of aggregate investment (panels a and c) and aggregate capital stock (panels b and d) to the (unanticipated) introduction of an investment tax credit at $t = 0$ lasting for 10 quarters in general equilibrium (i.e. allowing $p(t)$ to vary following the change in policy). The y-axis measures percent-deviations from steady state values. In panels a and b, the solid black line plots the total IRF for the endogenous TTB model; and the red dashed line plots the total IRF for the exogenous TTB model. In panels c and d, the black dashed line reproduces the total IRF from panels a and b; the green dashed line plots the extensive margin component of the IRF (i.e. fixing investment policies to steady state levels, but allowing entry to respond to the shock); and the blue dot-dashed line plots the intensive margin component of the IRF (i.e. shutting down entry, but allowing investment rates to respond to the shock). Remaining parameters are as in appendix table A16.

even after the tax credit is long gone. Only as they are completed, investment returns to steady state.

The black dashed line turns to the endogenous time to build case. The striking feature of the IRF is its shape: investment rises on impact, keeps rising, but then collapses at the time that the credit expires. Investment continues to decline back to steady state, and eventually runs into negative territory. The IRF behave "as if" the tax credit had reallocated investment flows across time, front-loading them to periods when the credit was in place.

Panel (c) decomposes the total IRF into its extensive and intensive margins. The extensive margin accounts for the exogenous case. The intensive margin explains the reshuffling of investment flows over time. It is entirely responsible for the impact effect: investment flows rise for all ongoing projects. Investment keeps rising slightly as the tax credit is still in place, which follows from distributional dynamics: projects become more mature over time, and it's optimal – even in steady state – to invest more in these projects. When the credit expires, investment in ongoing projects collapses. Firms front-load investment to periods when it's cheaper to build.

The intensive margin slump reflects distributional dynamics. Though investment rises for projects across the stage distribution, more mature projects respond by more. As projects are completed, the distribution skews towards younger projects. It is optimal to invest less in these projects – even in steady state. When the credit expires, the distribution is skewed towards younger projects, resulting in the bust. As these projects age, investment flows pick up, and the aggregate impulse returns to steady state.

Panels (b) and (d) illustrate the implications for productive capital. It accumulates more quickly in the endogenous case, which reflects the front-loading of investment flows, particularly so for projects that are closer to completion.

The investment IRFs highlight our first policy takeaway. The cumulative response of investment to tax credits comes from project starts. To first order, temporary tax credits just imply an intertemporal reshuffling of expenditures. This is not to say that investment tax credits are useless. Quite on the contrary, our results suggest they provide incredible flexibility to policymakers. If aggregate demand is currently low, we can spur investment today by borrowing it from the future. The duration and the size of the tax credit provide two levers with which policy makers can control this intertemporal reshuffling of aggregate demand.

Our second policy takeaway is that investment tax credits are a particularly flexible tool to spur investment relative to monetary policy. While changes in discount rates affect the entire path of cash flows accruing to a project, changes in tax policy can more flexibly target the "short-end" of the cash flow curve relative to the "long-end".[45] Fiscal policy does share

[45]The results in Figure 8 suggest that the partial equilibrium effects of monetary policy must contend with general equilibrium headwinds – at least if capital prices are sufficiently sensitive to changes in the capital stock.

one commonality with monetary policy: its effects are state-dependent. If the distribution of projects had been skewed towards later-stage projects when the investment tax credit was introduced, we would have observed stronger intertemporal reshuffling of investment flows and therefore a stronger short-term investment boom. Conversely, had the distribution been skewed towards earlier stage projects, the aggregate effects would have been muted.

We have restricted our attention to investment tax credits as a tool for stabilization policy. Investment subsidies additionally speed up project completion. For example, targeted subsidies can accelerate the green transition or infrastructure development more broadly. Tax policy can be a powerful lever to speed up structural change. We leave exploration of these policies for future research.

7.4 General Equilibrium

Our estimated causal elasticities are informative about partial equilibrium responses: the inclusion of state-by-year and sector-by-year fixed effects absorb effects from changes in local-level prices for construction inputs, sector-level prices for finished or secondary structures, or aggregate-level movements in real rates or factor prices. On the one hand, the partial equilibrium analysis better reconstructs our empirical setting and cleanly spells out the theoretical mechanisms behind the reduced form results. On the other hand, they are not informative about general equilibrium responses to these shocks. We cannot escape the "missing intercept" problem.

We can, however, embed the partial equilibrium model into a general equilibrium setting, and therein address counterfactual questions. Our exercise is minimalist by design. Parsimony isolates the key effects arising from feedback effects from prices, and illustrates how the baseline setup can be embedded into richer models of industry or aggregate dynamics.

Given a capital stock $K(t)$, we assume that final goods are produced according to

$$Y(t) = A(t)K(t)^{\alpha}$$

where $\alpha \in (0, 1]$ indexes returns to scale, and $A(t)$ stands for stochastic productivity. The price of the final good is normalized to one. The final goods producer buys the aggregate capital bundle $K(t)$ in competitive markets.

Standard profit maximization implies that capital prices reflect future marginal products of capital

$$p(t) = \int_0^{\infty} e^{-(r+\delta)s} MPK(t+s)ds$$

In our model, these headwinds are much weaker for fiscal policy. Therein lies a potential third policy takeaway, though admittedly suggestive relative to the other two.

where $MPK(t) = \alpha A(t)K(t)^{\alpha-1}$. When $\alpha = 1$, production of final goods features constant returns to scale, so any partial equilibrium analysis taking as given a price sequence can be recast as a general equilibrium analysis with a particular path of productivity disturbances $A(t)$. When $\alpha < 1$, production of final goods incurs decreasing returns to scale. As more capital is built, the marginal product of capital falls. Capital prices drop, and so does investment. Under decreasing returns, general equilibrium forces dampen partial equilibrium responses.

The structure above suffices to endogenize capital prices and understand the joint behavior of returns on capital and investment fluctuations (e.g. as in industry equilibrium models such as Kalouptsidi 2014 or Greenwood and Hanson 2014). In many macroeconomic models, further assumptions would pin down discount rates $r(t)$. We opt for simplicity, and retain exogenous discount rates. This can be micro-founded either by retaining a neoclassical interpretation and assuming a representative household with linear preferences; or by assuming that interest rates are determined outside of the model, either by a monetary authority (e.g. as in a New Keynesian model) or by global capital markets (e.g. as in a small open economy).[46] Within this minimalist setup, we can entertain aggregate counterfactuals.

Proposition 7. *Let hatted variables denote linearized deviations from steady state. Suppose we announce paths $\hat{A}(t)$ and $\hat{r}(t)$ for aggregate productivity and real interest rates. Then, to first-order, the response of capital prices is given by*

$$\frac{\hat{p}(t)}{p} = (r+\delta)\int_0^\infty e^{-(r+\delta)s}\left[\frac{\hat{A}(t+s)}{A} - (1-\alpha)\frac{\hat{K}(t+s)}{K} - \frac{\hat{r}(t+s)}{r+\delta}\right] ds \tag{20}$$

Proposition 7 highlights three channels by which fluctuations in productivity and interest rates change capital prices. First, there is the direct effect of productivity shocks on the marginal product of capital. When productivity is higher, marginal products of capital are higher, holding fixed a future path $\hat{K}(t)$ for the capital stock. Since capital is more productive, final goods producers are willing to pay more for it, raising $p(t)$. Second, there is a direct effect of interest rate shocks on the discounting of future marginal products of capital, again holding fixed future paths for the capital stock.[47] As discount rates rise, future marginal

[46]It is now well-understood that standard specifications of household preferences tend to undo the partial equilibrium responses of investment to shocks in general equilibrium. The task of "taming" the interest-rate sensitivity of investment in fully-fledged DSGE models with concave preferences (as in Winberry, 2021) is important, but outside the scope of this paper.

[47]In our partial equilibrium analysis, terminal payoffs did not respond to changes in idiosyncratic discount rates. Firms could lower their valuation p in partial equilibrium if e.g. high capital specificity implied the firm must hold the project. This intermediate case would look like a combination of our partial and general equilibrium results. While aggregate investment in partial equilibrium might then fall on impact following the increase in discount rates, none of our other results related to intensive margin investment would be affected. That is, we would still find project acceleration among close to completion projects, state-dependent responses

Figure 11: Impulse response functions to discount rate shocks in GE

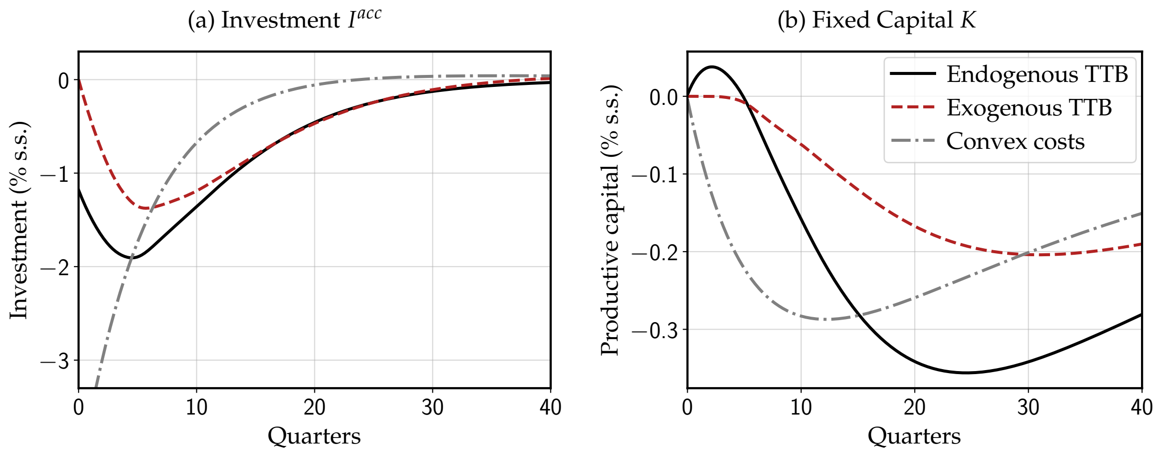

Note: the figure plots (quarterly) impulse response functions (IRFs) of aggregate investment (panel a) and aggregate capital stock (panel b) to a discount rate shock in general equilibrium (i.e. allowing $p(t)$ to vary following the shock as in (20), but shutting down productivity shocks). The y-axis measures percent-deviations from steady state values. The solid black line plots the IRF for the endogenous TTB model. The red dashed line plots the IRF for the exogenous TTB model. The gray dot-dashed line plots the response for a convex adjustment costs of installation model with no TTB. Shock persistence, shock size, and remaining parameters are as in appendix table A16.

products of capital are discounted more heavily. As future cash flows are discounted more heavily, $p(t)$ falls.

The key general equilibrium feedback effect between the stock of capital and capital prices is embedded in the ${}^{\hat{K}(t+s)}/_{K}$ terms. As productivity or discount rates change, firms revise their investment plans, thereby altering future capital stocks and marginal products of capital. In equilibrium, changes in capital prices are consistent with revisions in firms' investment plans and future capital stocks.

Figure 11 plots the impulse response function of accounting investment (panel a) and the fixed capital stock (panel b) to a shock in discount rates, holding fixed productivity. The magnitude of the shock is the same as in the partial equilibrium exercise (figure 7). The response of accounting investment in the exogenous time to build model is similar to the partial equilibrium response. The effects are more pronounced since the rise in discount rates lowers capital prices in general equilibrium, further discouraging project entry. As a result, productive capital falls. Notably, endogenous gestation lags generate persistent falls in productive capital, despite short-lived shocks.

The model with endogenous time to build now predicts a sharp contraction in investment on impact. Relative to the exogenous time to build IRF, firm's ability to alter project maturity now contributes to a bust rather than a boom in investment. The reason is two-fold. First,

to interest rates, front-loading of expenditures due to tax credits, and the ensuing consequence for shock- and state-dependence of adjustment costs.

there is the direct impact of the discount rate shock on capital prices, which lowers project returns and thereby discourages investment. Second, there is an effect on marginal product of capital. Projects that are sufficiently close to completion are accelerated, generating the upward bump in the IRF for productive capital and downward pressure in the marginal product of capital, further depressing investment incentives.

The bust in productive capital is stronger and more persistent in a world with endogenous time to build. Slower project completion drags down the capital stock relative to what we would have observed absent the shock. The result is a persistent slump. Note that this result holds despite no apparent differences in the speed of convergence of accounting investment back to steady state across the exogenous and endogenous time to build models. We conclude that the latter can generate sizable and persistent disruptions in capital supply even for modest contractions in investment.

Figure 12 decomposes the IRF for the endogenous time to build model into intensive and extensive margin IRFs. The first takeaway is that intensive margin effects are large and short-lived. They drive the short-run response of investment to discount rate shocks. Extensive margin effects instead generate sizable, but delayed investment responses; and drive the long-lived dynamics of the aggregate IRF.

The second takeaway is that the boom-bust investment dynamics generated by intensive margin effects survive in general equilibrium. The intuition is exactly as in the partial equilibrium dynamics. As the discount rate shock fades away, investment flows return to their steady state levels. Since the distribution of project shifts towards later-stage projects, investment in ongoing projects rises above the steady state level. Changes in capital prices do not undo these mechanisms.

A Comment on Counterfactual Policy Analysis. We return to Figure 11 and compare the results from time to build models with a comparable model featuring convex adjustment costs of installation (see Appendix G). Understanding departures from this standard formulation illustrates what time to build brings to the table. The model with convex installation costs predicts a sharp, short-lived contraction in investment, with no hump-shaped dynamics. Productive capital falls by more and more quickly, but converges faster to the original steady state. Time-to-build can generate hump-shaped investment dynamics and, in particular if endogenous, more persistent slumps in capital supply.

Adjustment costs at the level of the aggregate capital stock or investment can be seen as reduced-form devices to capture gestation lags (Taylor et al., 1982). They might perform well when matching particular evidence, e.g. the smoothness of investment impulses to interest rate shocks (Christiano et al., 2018). They might fail to match other features of investment dynamics. Should we be concerned?

In a world with endogenous gestation lags, discount rate and cash flow shocks generate

Figure 12: Decomposition of IRFs in Endogenous Time to Build Model

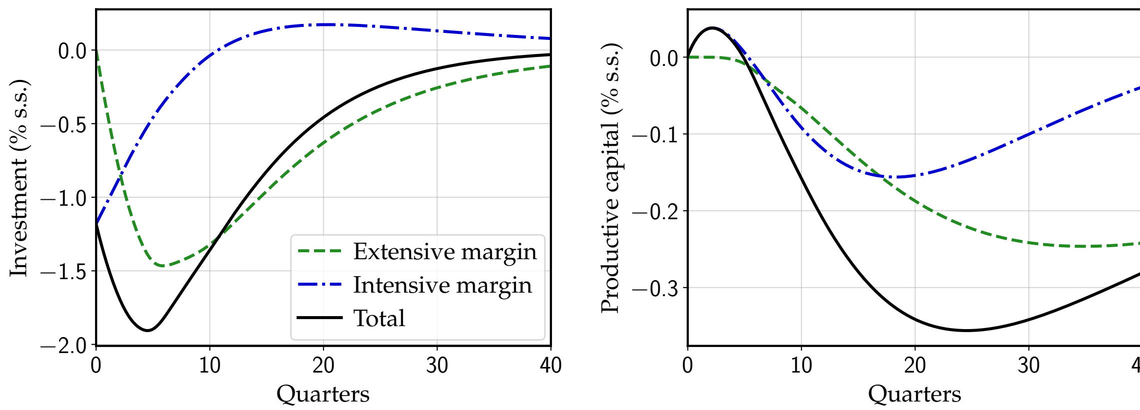

(a) Investment I^{acc}

(b) Fixed Capital K

Note: the figure plots (quarterly) impulse response functions (IRFs) of aggregate investment (panel a) and aggregate capital stock (panel b) to a discount rate shock in general equilibrium (i.e. allowing $p(t)$ to vary following the shock as in (20), but shutting down productivity shocks). The y-axis measures percent-deviations from steady state values. The solid black line plots the total IRF for the endogenous TTB model. The green dashed line plots the extensive margin component of the IRF (i.e. fixing investment policies to steady state levels, but allowing entry to respond to the shock). The blue dot-dashed line plots the intensive margin component of the IRF (i.e. shutting down entry, but allowing investment rates to respond to the shock). Shock persistence, shock size, and remaining parameters are as in appendix table A16.

impulses of different shapes, so that adjustment costs become *shock-dependent*. The effect of these shocks in turn vary with the underlying distribution of projects, which depends – even in steady state – on policy stances (e.g. the interest rate set by the monetary authority, or tax policy set by the fiscal authority). Adjustment costs thus become *state-dependent*. Counterfactual policy statements arising from models with ad hoc adjustment costs may therefore yield inappropriate conclusions.

We have provided two illustrations. In the context of monetary policy, the strength (and, in partial equilibrium, sign) of the impact investment response relied crucially on the distribution of projects along their completion stages. In the context of fiscal policy, the short-run effects of investment tax credits were stronger, and would also be state-dependent. Endogenous time to build provides an empirically plausible microfoundation for adjustment costs, and can be readily embedded in models of firm, industry, or aggregate dynamics.[48]

[48] That is not to say that our model is immune to a Lucas-type critique. We parametrized our acceleration costs and used data to discipline it. Our approach is econometrically no different from using IRFs of aggregate investment to discipline adjustment costs in DSGE models. But it provides a considerable methodological difference. We rely on microeconomic responses to identify the microeconomic primitives. Our work relies on the understanding that data at more disaggregated levels (e.g. project- or structure-level data) allows for more precise microeconomic measurement and thereby better microfounded models of capital formation.

8 Conclusion

Physical capital takes time to build. Leveraging detailed project histories of large capital expenditure projects in India, we moved beyond anecdotal evidence. Time to build is a pervasive microeconomic feature of fixed capital formation across industries, sponsor types, time periods, and project sizes. It is not exclusively an exogenous constraint that firms must accept. We provided causal evidence that firms have control over gestation lags; and rationalized our findings in a simple neoclassical investment model with a single new twist: endogenous time to build. We showed that this endogeneity has distinctive implications for the economics of capital investment and formation. Investment becomes more responsive to shocks, and thereby more volatile. More mature projects react more strongly relative to fledgling projects – a key model prediction that is borne out in the data. As a result, the distribution of projects along their completion lifecycle shapes aggregate dynamics. Monetary policy becomes state-dependent: the distribution of projects at the time of the interest rate change determines impact effects. The implications are not restricted to discount rate shocks. Investment in ongoing projects is extremely responsive to variations in cash flows, opening up space for time-varying investment credits to flexibly reshuffle aggregate demand across expansions and recessions; and suggesting novel roles for investment subsidies as structural policies to accelerate infrastructure development, such as the green transition.

Data at more disaggregated levels (e.g. project- or structure-level data) allows for further measurement and thereby better microfounded models of capital formation. We have provided a minimal setup that captures key economic forces, but is portable to other settings (e.g. managing a collection of workers or tasks) and invites extensions (e.g. time-varying project returns). But there is more to build. Questions are not in short supply. Can we structure project financing practices that internalize the resulting incentive effects in gestation lags? Do firms internalize gestation lags of their competitors when expanding capacity, and how does that shape industry dynamics? Do short-termist politicians speed up project and overrun budgets to claim credit for successful completion before their tenure ends, and can such incentives explain the political investment cycles? How much can targeted subsidies speed up infrastructure development in laggard economies or a global green transition? Can our insights from fixed capital formation be generalized to breakthrough innovations or other intangible forms of capital (e.g. job training as in Bloesch and Weber 2023)? Our future research will continue to tease out further implications of endogenous time to build.

References

Abadie, A., Athey, S., Imbens, G. W., and Wooldridge, J. M. (2022). When Should You Adjust Standard Errors for Clustering? *The Quarterly Journal of Economics*, 138(1):1–35.

Achdou, Y., Han, J., Lasry, J.-M., Lions, P.-L., and Moll, B. (2021). Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach. *The Review of Economic Studies*, 89(1):45–86.

Aftalion, A. (1927). The theory of economic cycles based on the capitalistic technique of production. *The Review of Economics and Statistics*, 9(4):165–170.

Alfaro, L. and Chari, A. (2014). Deregulation, Misallocation, and Size: Evidence from India. *Journal of Law and Economics*, 57(4):897–936.

Allison, P. D. (1982). Discrete-time methods for the analysis of event histories. *Sociological Methodology*, 13:61–98.

Antras, P. (2023). An 'Austrian' Model of Global Value Chains. NBER Working Papers 30901, National Bureau of Economic Research, Inc.

Arellano, M. and Bover, O. (1995). Another look at the instrumental variable estimation of error-components models. *Journal of Econometrics*, 68(1):29–51.

Asea, P. K. and Zak, P. (1999). Time-to-build and cycles. *Journal of Economic Dynamics and Control*, 23(8):1155–1175.

Ashcraft, A. B. (2005). Are banks really special? new evidence from the fdic-induced failure of healthy banks. *American Economic Review*, 95(5):1712–1730.

Bajari, P. and Tadelis, S. (2001). Incentives versus transaction costs: A theory of procurement contracts. *The RAND Journal of Economics*, 32(3):387–407.

Berger, D., Milbradt, K., Tourre, F., and Vavra, J. (2021). Mortgage prepayment and path-dependent effects of monetary policy. *American Economic Review*, 111(9):2829–78.

Bernanke, B. S. and Gertler, M. (1995). Inside the black box: The credit channel of monetary policy transmission. *Journal of Economic Perspectives*, 9(4):27–48.

Bloesch, J. and Weber, J. (2023). Congestion in onboarding workers and sticky r&d. techreport.

Blundell, R. and Bond, S. (1998). Initial conditions and moment restrictions in dynamic panel data models. *Journal of Econometrics*, 87(1):115–143.

Borusyak, K., Hull, P., and Jaravel, X. (2021). Quasi-Experimental Shift-Share Research Designs. *The Review of Economic Studies*, 89(1):181–213.

Brooks, P. K. (2000). The facts about time-to-build. IMF Working Papers 2000/138, International Monetary Fund.

BÃ¶hm-Bawerk, E. (1889). *The Positive Theory of Capital*. Macmillan.

Casares, M. (2006). Time-to-build, monetary shocks, and aggregate fluctuations. *Journal of Monetary Economics*, 53(6):1161–1176.

Chari, A., Jain, L., and Kulkarni, N. (2021). The unholy trinity: Regulatory forbearance, stressed banks and zombie firms. Working Paper 28435, National Bureau of Economic Research.

Chen, J. and Roth, J. (2022). Logs with zeros? Some problems and solutions. Papers 2212.06080, arXiv.org.

Chodorow-Reich, G. (2013). The Employment Effects of Credit Market Disruptions: Firm-level Evidence from the 2008-9 Financial Crisis *. *The Quarterly Journal of Economics*, 129(1):1–59.

Chodorow-Reich, G., Nenov, P. T., and Simsek, A. (2021). Stock market wealth and the real economy: A local labor market approach. *American Economic Review*, 111(5):1613–57.

Chopra, Y., Subramanian, K., and Tantri, P. L. (2020). Bank Cleanups, Capitalization, and Lending: Evidence from India. *The Review of Financial Studies*, 34(9):4132–4176.

Christiano, L., Eichenbaum, M., and Evans, C. (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy*, 113(1):1–45.

Christiano, L. and Todd, R. M. (1996). Time to plan and aggregate fluctuations. *Quarterly Review*, 20(Win):14–27.

Christiano, L. and Vigfusson, R. (2003). Maximum likelihood in the frequency domain: the importance of time-to-plan. *Journal of Monetary Economics*, 50(4):789–815.

Christiano, L. J., Eichenbaum, M. S., and Trabandt, M. (2018). On dsge models. *Journal of Economic Perspectives*, 32(3):113–40.

CMIE (2023a). Capexdx: Database of the life cycle of new investments. Last Accessed: 2023-09-01.

CMIE (2023b). Prowessdx: Database of financial performance of indian companies. Last Accessed: 2023-09-01.

Cohn, J. B., Liu, Z., and Wardlaw, M. I. (2022). Count (and count-like) data in finance. *Journal of Financial Economics*, 146(2):529–551.

Correia, S., Guimarães, P., and Zylkin, T. (2020). Fast poisson estimation with high-dimensional fixed effects. *The Stata Journal*, 20(1):95–115.

Correia, S., Guimarães, P., and Zylkin, T. (2021). Verifying the existence of maximum likelihood estimates for generalized linear models.

Edge, R. M. (2007). Time-to-build, time-to-plan, habit-persistence, and the liquidity effect. *Journal of Monetary Economics*, 54(6):1644–1669.

Eichenbaum, M., Rebelo, S., and Wong, A. (2022). State-dependent effects of monetary policy: The refinancing channel. *American Economic Review*, 112(3):721–61.

Flanagan, T. and Purnanandam, A. (2023). Why do banks hide losses? Working paper.

Freixas, X., Peydro, J.-L., and Laeven, L. (2015). *"Systemic Risk, Crises, and Macroprudential Regulation"*. The MIT Press.

Freixas, X. and Rochet, J.-C. (2008). *Microeconomics of Banking, 2nd Edition*, volume 1 of *MIT Press Books*. The MIT Press.

Gabaix, X. and Ibragimov, R. (2011). Rank1/2: A simple way to improve the ols estimation of tail exponents. *Journal of Business & Economic Statistics*, 29(1):24–39.

Gourieroux, C., Monfort, A., and Trognon, A. (1984). Pseudo maximum likelihood methods: Theory. *Econometrica*, 52(3):681–700.

Greenwood, R. and Hanson, S. G. (2014). Waves in Ship Prices and Investment *. *The Quarterly Journal of Economics*, 130(1):55–109.

Hall, R. (1977). Investment, interest rates, and the effects of stabilization policies. *Brookings Papers on Economic Activity*, 8(1):61–122.

Han, C. and Phillips, P. C. B. (2010). Gmm estimation for dynamic panels with fixed effects and strong instruments at unity. *Econometric Theory*, 26(1):119–151.

Kalecki, M. (1935). A macrodynamic theory of business cycles. *Econometrica*, 3(3):327–344.

Kalouptsidi, M. (2014). Time to build and fluctuations in bulk shipping. *American Economic Review*, 104(2):564–608.

Khwaja, A. I. and Mian, A. (2008). Tracing the impact of bank liquidity shocks: Evidence from an emerging market. *American Economic Review*, 98(4):1413–42.

Kulkarni, N., Ritadhi, S., Vij, S., and Waldock, K. (2021). Unearthing zombies. Working paper, Georgetown McDonough School of Business.

Kydland, F. E. and Prescott, E. C. (1982). Time to build and aggregate fluctuations. *Econometrica*, 50(6):1345–1370.

Leeper, E. M., Walker, T. B., and Yang, S.-C. S. (2010). Government investment and fiscal stimulus. *Journal of Monetary Economics*, 57(8):1000–1012.

Lewis, G. and Bajari, P. (2011). Procurement contracting with time incentives: Theory and evidence. *The Quarterly Journal of Economics*, 126(3):1173–1211.

Lewis, G. and Bajari, P. (2014). Moral Hazard, Incentive Contracts, and Risk: Evidence from Procurement. *The Review of Economic Studies*, 81(3):1201–1228.

Liu, E. and Tsyvinski, A. (2023). A dynamic model of input-output networks. *Review of Economic Studies*.

Lucca, D. O. (2007). Resuscitating Time-to-Build. 2007 Meeting Papers 909, Society for Economic Dynamics.

Mayer, T. (1960). Plant and equiptment lead times. *The Journal of Business*, 33(2):127–132.

McKay, A. and Wieland, J. F. (2021). Lumpy durable consumption demand and the limited ammunition of monetary policy. *Econometrica*, 89(6):2717–2749.

Meier, M. (2020). Supply Chain Disruptions, Time to Build, and the Business Cycle. Crc tr 224 discussion paper series.

Montgomery, M. R. (1995). Time-to-build completion patterns for nonresidential structures, 1961-1991. *Economics Letters*, 48(2):155–163.

Neyman, J. and Scott, E. L. (1948). Consistent estimates based on partially consistent observations. *Econometrica*, 16(1):1–32.

Nickell, S. (1981). Biases in dynamic models with fixed effects. *Econometrica*, 49(6):1417–1426.

Oh, H. and Yoon, C. (2020). Time to build and the real-options channel of residential investment. *Journal of Financial Economics*, 135(1):255–269.

Peek, J. and Rosengren, E. S. (2000). Collateral damage: Effects of the japanese bank crisis on real activity in the united states. *American Economic Review*, 90(1):30–45.

Pigou, A. (1920). *The Economics of Welfare*. Macmillan.

Pinardon-Touati, N. (2023). The crowding out effect of local government debt: Micro and macro-estimates. Technical report, Columbia University.

Ramey, V. (2020). The macroeconomic consequences of infrastructure investment. NBER Working Papers 27625, National Bureau of Economic Research, Inc.

Rouwenhorst, K. (1991). Time to build and aggregate fluctuations: A reconsideration. *Journal of Monetary Economics*, 27(2):241–254.

Samuelson, P. A. (1939). Interactions between the multiplier analysis and the principle of acceleration. *The Review of Economics and Statistics*, 21(2):75–78.

Silva, J. M. C. S. and Tenreyro, S. (2006). The Log of Gravity. *The Review of Economics and Statistics*, 88(4):641–658.

Singer, J. D. and Willett, J. B. (1993). It's about time: Using discrete-time survival analysis to study duration and the timing of events. *Journal of Educational Statistics*, 18(2):155–195.

Taylor, J. B., Baily, M. N., and Fischer, S. (1982). The swedish investment funds system as a stabilization policy rule. *Brookings Papers on Economic Activity*, 1982(1):57–106.

Van Den Berg, G. J. (2001). Chapter 55 - duration models: Specification, identification and multiple durations. volume 5 of *Handbook of Econometrics*, pages 3381–3460. Elsevier.

Winberry, T. (2021). Lumpy investment, business cycles, and stimulus policy. *American Economic Review*, 111(1):364–96.

Zhou, C. (2000). Time-to-Build and Investment. *The Review of Economics and Statistics*, 82(2):273–282.

A Data Appendix

A.1 CapexDx

We provide additional information on data collection, definitions, and limitations for our project-level data.

Defining Projects CMIE defines a project when it can assess the capacity that will be created. They illustrate the definition as follows. Suppose a sponsor decides to invest a billion dollars in India. This does not constitute a project: we don't know the capacity that is being created (i.e. whether it is a factory, a road, a school building, etc). Suppose further that the sponsor specified that it would invest in the glass industry. Again, it is not clear what capacity is being created. Is it a factory to produce glass? Or is it a warehouse to store glass? Perhaps it could be a new headquarter building for the sponsor.

Suppose then that the conglomerate specifies that it will invest a hundred million rupees in a glass manufacturing plant in Chennai. Now, there is clarity on the capacity that will be created. To quote the documentation, "Although the company has not categorically stated so, it does imply the creation of additional capacity, unless the context or circumstances suggest otherwise." CMIE therefore tracks the project. Importantly, CMIE would not track the project if the sponsor was just acquiring an already existing glass manufacturing plant, rather than developing the structure.

Some projects are multi-faceted and complex. They might combine several smaller projects. In that case, CMIE recognizes "the smaller logical components as independent projects if the cost for each of them is available separately and if the implementation is clearly separate". Consider the glass manufacturing example above. If the sponsor specified that there would be four plants in the complex, each costing 25 million rupees, with plants being simultaneously built and with equal capacity, then CMIE would track each plant individually. The exception would be if the sponsor explicitly states that it considers all sub-components of the larger project as a single, integrated structure. CMIE then follows the sponsor and defines a single project, despite detailed information on sub-components. In that case, or in the absence of such granular information on project sub-components, we would observe our example in the data as "Glass Manufacturing Project" with a total cost of 1 billion dollars located in Chennai.

Data Collection CMIE builds Capex by combing publicly available sources and by actively contacting sponsors via emails and phone calls. This detective work reflects the lack of systematic sources of project-level data. There is no universal standard of disclosure that sponsors must follow, either in company reports (i.e. disclosures of materially relevant information for investors in the case of public companies) or in government forms. CMIE therefore collects the data from any and all credible sources.

To quote CMIE, their "long experience in building and maintaining such a database and in applying it in its own research ensures the utility of the database." Indeed, many players in industry, government and academia rely on CMIE products to obtain reliable data on the Indian economy. Capex is no exception. Practitioners we spoke to suggest they actively used Capex to either track existing projects or make strategic decisions regarding future capacity creation.

Our conversations with the data provider further alleviates data quality concerns. First, monitoring is regular, as all projects "are touched for updates once in a quarter." Second, CMIE recognizes the need to validate sponsor information. To quote a conversation we had with the Capex team: "Media, government websites are also source of capturing of events that measure the progress of implementation of the project. If these are the sources of information then we do not rely for the company to disclose the same progress." Third, CMIE actively checks for outdated information in their tracking. Suppose CMIE learns after contacting promoters that implementation for a particular project happened one month before what CMIE had as the original implementation start date. In that case, CMIE checks the veracity of the sponsors' claims against other sources it has used previously, such as public disclosures by government sources or media outlets. If CMIE finds corroborating evidence in the public domain for the sponsors' claims, it then updates the outdated information to the most recent news.

Dating of events reflects data collection. In many cases, the date of an event is the date when the information was available to CMIE. For example, announcement dates correspond to the first time that the project is mentioned in the public domain; implementation start dates mark the first time that CMIE learns that an implementation-related event has taken place, either from public domain sources or from project sponsors.

Despite intense monitoring, CMIE may fail to follow a project's full history. While we observe announcement dates for all projects, we do not always know what eventually happened to a project. In short, some observations are right-censored. We deal with this feature of the data in two different ways. In some exercises, we restrict analyses to non-censored projects or exclusively to completed projects. In the parlor of survival analyses, we perform "complete case" analyses, aware we cannot circumvent sample selection concerns depending on our outcome of interest. In other exercises, we leverage econometric methods that take into account the right-censored nature of the data, thereby allowing us to maximize the number of observations we can learn from.

Reassuringly, we understand part of the data generating process for censoring — at least that under the discretion of CMIE. CapexDx discloses it stops tracking projects "if they are stalled or shelved or abandoned or information regarding them is not available for over 30 months in the case of manufacturing projects and for over 36 months in the case of infrastructure projects." Right-censoring is not a one-way street: CMIE resumes tracking previously censored projects if it manages to obtain reliable information on project status from the public domain. In short, the combination of intense monitoring and transparent censoring schemes bolsters data reliability.

B Additional Measurement Exercises

Table A1: Transitions Across Project Stages

	Announced	Completed	No Information	Abandoned	Stalled	Under Imp
Announced	0.00	6.52	20.49	2.47	2.45	68.07
Stalled	10.63	1.31	6.35	22.85	0.00	58.86
Under Imp	0.00	62.15	27.55	3.26	7.04	0.00

Notes: This table constructs a transition matrix across project stages for all projects available in the March 2022 vintage of CapexDx. Rows represent departing states. We only include non-terminal stages, i.e. we exclude completion, abandonment, and loss to follow-up. Columns denote arriving states, which include terminal stages. We only include transitions for which the departing and arriving stage differ. Each entry is the fraction of transitions from the row stage to the column stage (multiplied by 100). For example, 68.07% of all transitions starting from "announced" correspond to transitions into "under implementation".

Table A2: Quarterly Transition Probabilities

	Announced	Completed	No Information	Abandoned	Stalled	Under Imp
Announced	87.92	0.79	2.48	0.30	0.30	8.22
Abandoned	0.66	0.17	0.00	97.98	0.00	1.20
Stalled	0.63	0.08	0.38	1.36	94.05	3.50
Under Imp	0.00	4.02	1.78	0.21	0.46	93.53

Notes: This table constructs a transition matrix across project stages for all projects available in the March 2022 vintage of CapexDx. Rows represent departing states. We only include non-terminal stages, i.e. we exclude completion, abandonment, and loss to follow-up. In some cases, abandoned projects returns to activity. Since these are exceedingly rare events, we consider abandonment as terminal in all other results in the paper. Columns denote arriving states, which include terminal stages. We include all quarterly transitions. Each entry is the fraction of transitions from the row stage to the column stage (multiplied by 100). For example, 87.92% of all transitions starting from "announced" remain "announced", while 8.22% move into implementation.

Table A3: Cost-Time Schedules

	(1)	(2)	(3)	(4)	(5)
Log Cost	0.178	0.163	0.162	0.152	0.155
	(0.015)	(0.013)	(0.013)	(0.011)	(0.011)
Obs.	29988	29988	29983	29987	29982
Sector		✓	✓		
Industry				✓	✓
Start Year			✓		✓

Notes: This table reports correlations between log time to build and log costs estimated with linear regression from the microdata underlying the binned scatter plots of figure 3. The sample consists of completed projects for which cost data is available. Column (1) reports the bivariate relationship. Column (2) includes sector fixed effects, and column (3) adds year started fixed effects. Columns (4) and (5) replace sector fixed effects by industry fixed effects. Standard errors are clustered at the industry level and reported in parentheses.

Figure A3: Cost-Time Schedule by Subgroups

(a) By Ownership

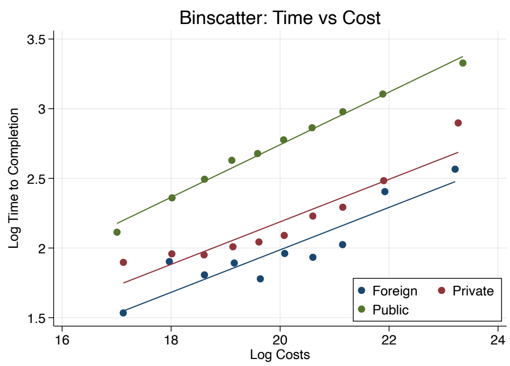

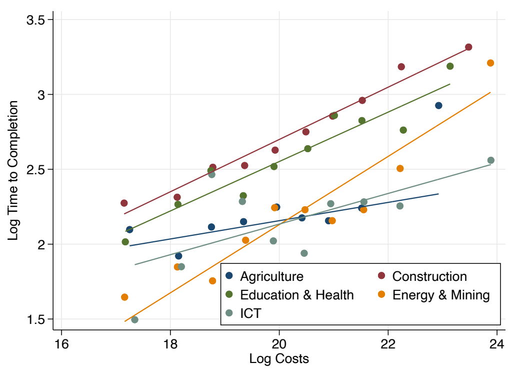

(b) By Sector

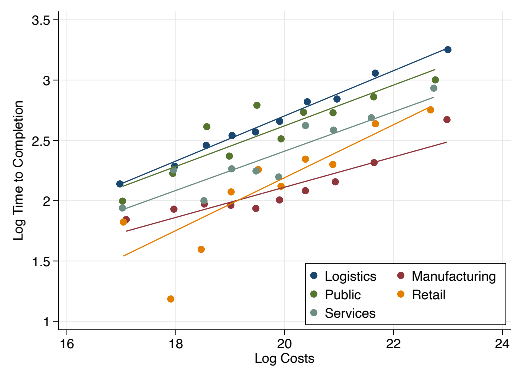

(c) By Sector

Notes: The figure disaggregates the binned scatter plots of log time to build on log costs (see panel a of figure 3) by sponsorship (panel a) and sectors (panels b and c, split for the sake of visualization).

C Additional Empirical Results

Table A4: Effects of Asset Quality Review on Credit Supply

	(1)	(2)	(3)	(4)	(5)	(6)
Exposure	-0.0021	-0.0069***	-0.0027*	-0.0046*	-0.0085***	-0.0051**
	(0.0015)	(0.0018)	(0.0015)	(0.0024)	(0.0028)	(0.0024)
Outc Mean	-0.08	-0.08	-0.08	-0.10	-0.10	-0.10
Obs	194072	194072	194072	78827	78827	78827
Full Sample	✓	✓	✓			
Hazard Sample				✓	✓	✓
Borrower + Year FE	✓		✓	✓		✓
Borrower x Year FE		✓			✓	
Firm Controls			✓			✓

Notes: the table reports estimates of β from the following model:

$$\Delta\text{Loans}_{bft} = \text{Firm-Year}_{ft} + \beta \times \text{Bank Shock}_{bt} + \varepsilon_{bft}$$

where ΔLoans_{bft} is the growth rate of loans between $t-1$ and t, constructed as in Davis and Haltiwanger (1992). The sample consists of firm-bank relationships for 2013-2019, conditional on an existing relationship (positive outstanding loans) in the 2013-2015 period. Exposure is defined as the difference between required provision after the audit and the original provisions, scaled by total bank assets. Coefficient estimates are read as the effect on the outcome for a one standard-deviation increase in lender exposure. For example, column (1) reports that moving from the lender with average exposure to a lender with one-standard deviation higher exposure decreases outstanding loans by -0.21 percentage points, relative to a mean growth rate of 8 percentage points. In columns (1)-(3), the sample comprises all borrower-lender pairs in our MCA-Prowess merged sample. In columns (4)-(6), the sample is restricted to firms for which we eventually observe projects in our hazard regressions (see model (3) and table 4). Columns (2) and (5) include borrower-by-year fixed effects. Columns (1) and (3) include borrower, and year fixed effects. Columns (3) and (6) add firm controls. Firm controls include sales, sales growth, profitability, debt over equity, book leverage, equity issuances, cash over liquid assets, and debt over EBITDA. All firm controls are winsorized at the 1st and 99th percentile, and then standardized. All firm controls are lagged by one-period. Standard errors are clustered at the lender level and reported in parentheses. $*$, $**$, and $***$ denote statistical significance at the 10%, 5% and 1% levels, respectively.

Table A5: Effects of Firm-Level Credit Supply Shocks on Project Completion Hazards

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Exposure (Pre)	0.1092***	0.0803***	0.0807***	0.0778***	0.0782***	0.0873***	0.0884***
	(0.0283)	(0.0283)	(0.0281)	(0.0276)	(0.0276)	(0.0274)	(0.0275)
Exposure (2015)	0.1102***	0.0728***	0.0734***	0.0705***	0.0713***	0.0842***	0.0858***
	(0.0281)	(0.0278)	(0.0277)	(0.0270)	(0.0270)	(0.0274)	(0.0275)
DepVar Mean	0.0392	0.0392	0.0392	0.0392	0.0392	0.0409	0.0409
Observations	123857	123857	123857	123857	123857	118314	118314
Age Splines	✓	✓	✓	✓	✓	✓	✓
Firm FE		✓	✓	✓	✓	✓	✓
State x Year FE		✓	✓	✓	✓	✓	✓
Sector x Year FE				✓	✓		
Industry x Year FE						✓	✓
Firm Controls			✓		✓		✓

Notes: the table reports PPML estimates of the semi-elasticity of project-level hazards to firm exposure to the Asset Quality Review (see model (3)). The sample consists of projects observed quarterly between 2013-2019. Each row corresponds to a measure of firm exposure. Firm exposure is defined as the weighted-average of lender exposure, where weights vary by row. "(Pre)" indicates weights are constructed as average outstanding loans between 2013-2015 for a borrower-lender pair relative to average outstanding loans between 2013-2015 for a given borrower. "(2015)" indicates weights are constructed as outstanding loans in 2015 between a borrower-lender pair relative to outstanding loans in 2015 for a given borrower. Coefficient estimates are interpreted as the percent effect on the outcome for a one standard-deviation increase in firm exposure. For example, column (1) reports that moving from a firm with average exposure to a firm with one-standard deviation higher exposure increases the hazard by 10.92%. Age splines refer to cubic splines on a project's age (measured in quarters). Firm controls include sales, sales growth, profitability, debt over equity, book leverage, equity issuances, cash over liquid assets, and debt over EBITDA. All firm controls are winsorized at the 1st and 99th percentile, and then standardized. All firm controls are lagged by one-period. The reduction in sample size from columns (5) to columns (6) and (7) results from separation of additional observations by the more granular industry-by-year fixed effects relative to the sector-by-year fixed effects. Standard errors are clustered at the firm level and reported in parentheses. $*$, $**$, and $***$ denote statistical significance at the 10%, 5% and 1% levels, respectively.

Table A6: Effects of Firm-Level Credit Supply Shocks on Project Completion Hazards

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Exposure	0.0061***	0.0033***	0.0033***	0.0031***	0.0031**	0.0033***	0.0033***
	(0.0018)	(0.0013)	(0.0013)	(0.0012)	(0.0012)	(0.0011)	(0.0011)
DepVar Mean	0.0392	0.0392	0.0392	0.0392	0.0392	0.0391	0.0391
Observation	123857	123857	123857	123857	123857	123835	123835
Age Splines	✓	✓	✓	✓	✓	✓	✓
Firm FE		✓	✓	✓	✓	✓	✓
State x Year FE		✓	✓	✓	✓	✓	✓
Sector x Year FE				✓	✓		
Industry x Year FE						✓	✓
Firm Controls			✓		✓		✓

Notes: the table reports linear regression estimates of the effect of firm exposure to the Asset Quality Review on project hazards (see model (3)). The sample consists of projects observed quarterly between 2013-2019. Firm exposure is defined as the weighted-average of lender exposure, where weights correspond to outstanding loans in the previous year. Coefficient estimates are interpreted as the level effect on the outcome for a one standard-deviation increase in firm exposure. For example, column (1) reports that moving from a firm with average exposure to a firm with one-standard deviation higher exposure increases the hazard by 0.61 percentage points, relative to an unconditional completion probability of 3.92 percentage points. Age splines refer to cubic splines on a project's age (measured in quarters). Firm controls include sales, sales growth, profitability, debt over equity, book leverage, equity issuances, cash over liquid assets, and debt over EBITDA. All firm controls are winsorized at the 1st and 99th percentile, and then standardized. All firm controls are lagged by one-period. The reduction in sample size from columns (5) to columns (6) and (7) results from separation of additional observations by the more granular industry-by-year fixed effects relative to the sector-by-year fixed effects. Standard errors are clustered at the firm level and reported in parentheses. $*$, $**$, and $***$ denote statistical significance at the 10%, 5% and 1% levels, respectively.

Table A7: Effects of Firm-Level Credit Supply Shocks on Project Starts

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Exposure (Pre)	-0.0740***	-0.0708**	-0.0673**	-0.0808**	-0.0753**	-0.0562*	-0.0556*
	(0.0278)	(0.0324)	(0.0321)	(0.0321)	(0.0319)	(0.0305)	(0.0303)
Exposure (2015)	-0.0755***	-0.0711**	-0.0676**	-0.0818***	-0.0765**	-0.0577*	-0.0572*
	(0.0280)	(0.0321)	(0.0319)	(0.0317)	(0.0315)	(0.0298)	(0.0296)
DepVar Mean	0.8362	0.8362	0.8362	0.8362	0.8362	0.8326	0.8326
Observation	8994	8994	8994	8994	8994	8869	8869
Firm FE		✓	✓	✓	✓	✓	✓
Year FE		✓	✓				
NIC Section x Year FE				✓	✓		
NIC Division x Year FE						✓	✓
Firm Controls			✓		✓		✓

Notes: the table reports PPML estimates of the semi-elasticity of firm-level project starts to firm exposure to the Asset Quality Review (see model (5)). The sample consists of the Capex-Prowess matched sample of firms observed yearly between 2013-2019. Each row corresponds to a measure of firm exposure. Firm exposure is defined as the weighted-average of lender exposure, where weights vary by row. "(Pre)" indicates weights are constructed as average outstanding loans between 2013-2015 for a borrower-lender pair relative to average outstanding loans between 2013-2015 for a given borrower. "(2015)" indicates weights are constructed as outstanding loans in 2015 between a borrower-lender pair relative to outstanding loans in 2015 for a given borrower. Coefficient estimates are interpreted as the percent effect on the outcome for a one standard-deviation increase in firm exposure. For example, column (1) reports that moving from a firm with average exposure to a firm with one-standard deviation higher exposure reduces project starts by 7.4%. NIC Section and Division classifications are constructed from NIC codes and categorizations. Firm controls include sales, sales growth, profitability, debt over equity, book leverage, equity issuances, cash over liquid assets, and debt over EBITDA. All firm controls are winsorized at the 1st and 99th percentile, and then standardized. All firm controls are lagged by one-period. The reduction in sample size from columns (5) to columns (6) and (7) results from separation of additional observations by the more granular division-by-year fixed effects relative to the section-by-year fixed effects. Standard errors are clustered at the firm level and reported in parentheses. $*$, $**$, and $***$ denote statistical significance at the 10%, 5% and 1% levels, respectively.

Table A8: Effects of Firm-Level Credit Supply Shocks on Project Starts

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Exposure	-0.0560***	-0.0534**	-0.0520**	-0.0574**	-0.0562**	-0.0471**	-0.0462**
	(0.0195)	(0.0226)	(0.0223)	(0.0226)	(0.0223)	(0.0218)	(0.0215)
DepVar Mean	0.8362	0.8362	0.8362	0.8362	0.8362	0.8281	0.8281
Observation	8994	8994	8994	8994	8994	8941	8941
Firm FE		✓	✓	✓	✓	✓	✓
Year FE		✓	✓				
NIC Section x Year FE				✓	✓		
NIC Division x Year FE						✓	✓
Firm Controls			✓		✓		✓

Notes: the table reports linear regression estimates of the effect of firm exposure to the Asset Quality Review on firm-level project starts (see model (5)). The sample consists of the Capex-Prowess matched sample of firms observed yearly between 2013-2019. Each row corresponds to a measure of firm exposure. Firm exposure is defined as the weighted-average of lender exposure, where weights correspond to outstanding loans in the previous year. Coefficient estimates are interpreted as the level effect on the outcome for a one standard-deviation increase in firm exposure. For example, column (1) reports that moving from a firm with average exposure to a firm with one-standard deviation higher exposure reduces project starts by 0.056 projects relative to an unconditional mean of 0.8362 projects. NIC Section and Division classifications are constructed from NIC codes and categorizations. Firm controls include sales, sales growth, profitability, debt over equity, book leverage, equity issuances, cash over liquid assets, and debt over EBITDA. All firm controls are winsorized at the 1st and 99th percentile, and then standardized. All firm controls are lagged by one-period. The reduction in sample size from columns (5) to columns (6) and (7) results from separation of additional observations by the more granular division-by-year fixed effects relative to the section-by-year fixed effects. Standard errors are clustered at the firm level and reported in parentheses. $*$, $**$, and $***$ denote statistical significance at the 10%, 5% and 1% levels, respectively.

Table A9: Balance Tests: Cross-Section

	(1)	(2)	(3)
Log Sales	3.644	3.967	-0.1998**
			(0.0826)
Sales Growth	0.040	0.043	-0.0002
			(0.0525)
Profitability (EBIT)	0.052	0.062	-0.0149
			(0.0325)
EBIT < 0	0.157	0.117	0.0419
			(0.0402)
Debt/Equity	57.100	42.253	-1.9474
			(4.5947)
Book Leverage	35.800	28.777	1.6844
			(3.1223)
Equity Issuance	-0.020	0.134	-0.2403**
			(0.1041)
Cash/Liquid Assets	0.156	0.111	0.0150
			(0.0116)
Debt/EBIT	0.103	2.566	-3.6173
			(2.8500)

Notes: the table reports linear regression estimates of covariate balance across sponsor exposure to the Asset Quality Review. In particular, we regress firm-level characteristics averaged over 2013-2015 X_f on average firm exposure over the 2016-2019 period $Expo_f$, where firm exposure is defined as the weighted-average of lender exposure and weights correspond to outstanding loans in the previous year. Each row corresponds to one of our firm controls. Columns (1) and (2) report raw means conditional on a binary exposure indicator. Column (3) reports the coefficient on the exposure variable in the cross-sectional linear regression model. All outcomes are winsorized at the 2.5th and 97.5th percentile, and then standardized. Standard errors are heteroscedasticity-robust and reported in parentheses. $*$, $**$, and $***$ denote statistical significance at the 10%, 5% and 1% levels, respectively.

Table A10: Balance Tests: Panel

	(1)	(2)	(3)
Log Sales	2.849	3.838	-0.3798
			(0.4663)
Sales Growth	0.085	0.045	-0.1046
			(0.0782)
Profitability (EBIT)	0.016	0.074	-0.2861**
			(0.1177)
EBIT < 0	0.179	0.119	0.0906
			(0.0644)
Debt/Equity	85.557	46.410	-23.0500
			(46.1865)
Book Leverage	51.422	30.764	-17.9970
			(24.3064)
Equity Issuance	0.108	-0.104	-0.1340
			(0.1350)
Cash/Liquid Assets	0.225	0.140	-0.0710*
			(0.0377)
Debt/EBIT	-6.070	2.478	0.2829
			(5.7189)

Notes: the table reports linear regression estimates of covariate balance across sponsor exposure to the Asset Quality Review. The sample consists of the Capex-Prowess matched sample of firms observed yearly between 2013-2019. In particular, we regress lagged firm-level characteristics $X_{f,t-1}$ on contemporaneous firm exposure $Expo_{f,t}$. Each row corresponds to one of our firm controls. Firm exposure is defined as the weighted-average of lender exposure, where weights correspond to outstanding loans in the previous year. Columns (1) and (2) report raw means conditional on a binary exposure indicator. Column (3) reports the coefficient on the exposure variable in the panel linear regression model, including firm and year fixed effects. All outcomes are winsorized at the 2.5th and 97.5th percentile, and then standardized. Standard errors are clustered at the firm level and reported in parentheses. $*$, $**$, and $***$ denote statistical significance at the 10%, 5% and 1% levels, respectively.

Table A11: Predicting TTB: Model Comparison

	Time to Build				Log Time to Build			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
				Panel A: RMSE				
Out-of-Sample	13.719	13.402	13.590	13.571	0.861	0.825	0.851	0.858
In-Sample	11.209	10.728	11.768	10.652	0.820	0.785	0.792	0.780
				Panel B: R-Squared				
Out-of-Sample	0.124	0.164	0.141	0.143	0.111	0.184	0.132	0.117
In-Sample	0.368	0.421	0.304	0.430	0.139	0.211	0.197	0.221
				Panel C: OLS				
Slope	1.069	1.083*	1.214***	0.946	1.316***	1.194***	0.997	0.954***
	(0.058)	(0.047)	(0.035)	(0.042)	(0.023)	(0.017)	(0.015)	(0.016)
Constant	1.559**	1.580***	-1.036**	3.033***	-0.558***	-0.303***	0.122***	0.237***
	(0.717)	(0.570)	(0.434)	(0.518)	(0.052)	(0.039)	(0.034)	(0.035)
Sector FE	✓		✓		✓		✓	
Industry FE		✓		✓		✓		✓
State FE			✓	✓			✓	✓
Ownership FE			✓	✓			✓	✓
Type of Unit FE			✓	✓			✓	✓

Notes: the table compares the predictive ability of different models for time to build. Each column correspond to a different (linear) predictive model including (non-parametrically) the covariates marked by ✓ in the bottom panel of the table. Panel A reports root-mean squared errors and Panel B reports R^2 of each model for both the estimating sample ("In-sample") and for the hazard regression sample ("Out-of-Sample"). Panel C reports OLS estimates a linear regression of observed gestation lags on predicted gestation lags for the hazard regression sample (i.e. the target sample of interest). Slope values closer to one and constant values closer to zero indicate more proximity to the 45 degree line. Standard errors are clustered at the firm level and reported in parentheses. $*$, $**$, and $***$ denote statistical significance at the 10%, 5% and 1% levels, respectively. In Panel C, the null hypothesis is that the slope is equal to 1, and that the constant is equal to zero. Tests for slope and constant are conducted separately rather than jointly. Our preferred model is that of column (2). It provides the second-best out-of-sample predictive power as measured by R^2, being outperformed only by the model in column (6); we pick it over the latter since it provides a closer fit to the 45 degree line.

Figure A11: Predictive Models of Time to Build

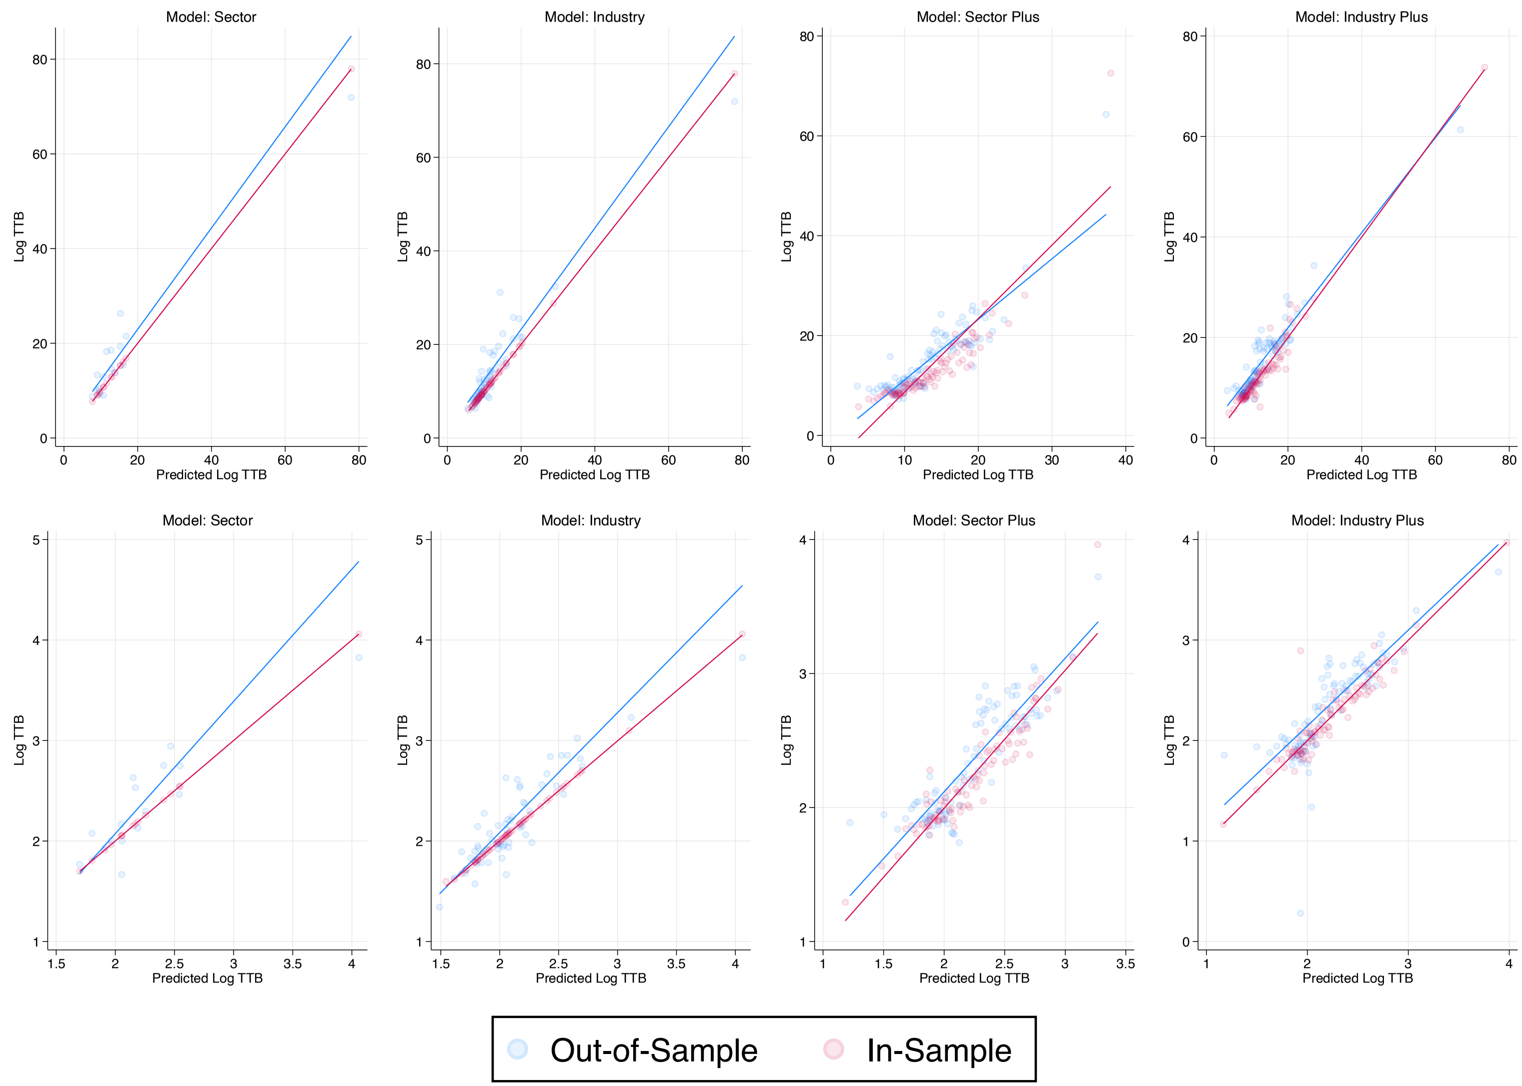

Notes: The figure plots binned scatter plots and lines of best fit for different predictive models of time to build. The y-axis plots either (quarterly) time to build in levels or in logs. The x-axis plots the predictions obtained from each model. If the model perfectly predicted the data, "in sample" dots would line up perfectly with the red line. The "in-sample" sample is all projects completed prior to 2015 that are not included in the hazard regression sample. The "out-of-sample" sample in blue is the hazard regression sample. The first row predicts gestation lags in levels. The second row predicts gestation lags in logs. Columns correspond to different models. "Sector" includes only sector dummies. "Industry" moves to industry dummies. "Plus" indicates that we add all other predictors in column 8 of Table 2 excluding year started.

Table A12: Heterogeneous Effects by Initial Age

	(1)	(2)	(3)	(4)
Exposure	0.101^{***}	0.098^{***}	0.110^{***}	0.109^{***}
	(0.026)	(0.027)	(0.026)	(0.026)
Age x Exposure	0.055^{**}	0.058^{**}	0.058^{**}	0.060^{**}
	(0.024)	(0.024)	(0.025)	(0.026)
Initial Age FE	✓	✓	✓	✓
Firm FE	✓	✓	✓	✓
State x Year FE	✓	✓	✓	✓
Sector x Year FE	✓	✓		
Industry x Year FE			✓	✓
Firm Controls		✓		✓

Notes: the table reports PPML estimates of the semi-elasticity of project-level hazards to sponsor exposure to the Asset Quality Review. The sample consists of projects observed quarterly between 2013-2019. Firm exposure is defined as the weighted-average of lender exposure, where weights correspond to outstanding loans in the previous year. The proxy for stage is Initial Age, constructed as age of the project when first observed in the sample. This measure is winsorized at the 1st and 99th percentile, and then standardized before being included in the regression model. The coefficient on exposure is interpreted as the percent effect on the outcome for a one standard-deviation increase in firm exposure for the project with mean proxy stage. The coefficient on the interaction is interpreted as the change in the main effect resulting from a one-standard deviation level change in proxy stage. For example, the first row in column (1) reports that moving from a sponsor with average exposure to a sponsor with one-standard deviation higher exposure increases the hazard by 10.1% for the project at the mean proxy stage. The second row in column (1) reports that if the project had been one standard deviation above (below) the mean stage at the time of the shock, the effect would have incresed (declined) by 5.5%, for a total effect of 15.6% (4.6%). In order to avoid strong collinearity concerns between the age spline, initial age, and the time fixed effects, we omit the age spline, and control non-parameterically for the initial age. Firm controls include sales, sales growth, profitability, debt over equity, book leverage, equity issuances, cash over liquid assets, and debt over EBITDA. All firm controls are winsorized at the 1st and 99th percentile, and then standardized. All firm controls are lagged by one-period. The reduction in sample size from columns (1) and (2) to columns (3) and (4) results from separation of additional observations by the more granular industry-by-year fixed effects relative to the sector-by-year fixed effects. Standard errors are clustered at the firm level and reported in parentheses. $*$, $**$, and $***$ denote statistical significance at the 10%, 5% and 1% levels, respectively.

D Backfilling, Censoring, and Selection

CapexDx strives to provide maximal coverage, despite the non-administrative nature of the data collection process. As a result, some project histories are backfilled. Once CapexDx obtains information on an already ongoing project, it reconstructs the project history by active research on the public domain and contact with project sponsors.

On the one hand, such retroactive filling improves coverage. On the other hand, it introduces a potential threat to our empirical strategy. The concern is that the RBI audits – and their resulting firm-level disruptions – alter the data collection process. For example, CapexDx could become less likely to find data on projects. This opens up space for a delicate form of selection bias: the sample of projects collected after the shock may differ from the sample of projects before the shock.[49]

If the resulting selection into sample is balanced across the treated and untreated group, it is *independent* of treatment. Our estimates will have a causal interpretation *within the selected sample* (Hernan and Robins, 2023). That is, we can still identify the effect of exposing units to the shock relative to a counterfactual of no exposure; but we can only speak to the effect on the units that eventually make it into our dataset. The causal effect on censored units is not identified from observational data. This concern is present in *any* study for which data collection might be affected by the treatment of interest (e.g. studies leveraging survey data).

We would not be able to make causal statements, even within the noncensored sample, if the treatment induced differential censoring. To be concrete, let's suppose then that CapexDx becomes less able to obtain data on treated relative to untreated projects. Perhaps, firms more exposed to the banking shock stop sharing information about their projects on the public domain or taking the time to respond to CapexDx inquiries. Let's further suppose there is no backfilling, so any new projects for which CapexDx obtains data are always recent starts.

The concern then is that the observed sample would be skewed towards ***fewer*** early projects among the treated group. Moreover, since older projects are more likely to be completed, comparing treated and untreated projects within the uncensored sample incurs a sample selection bias. That is, we may observe an association between completion and treatment, even if treatment does not influence completion. The observed association results simply from the missing early projects among the treated. If such projects had been included in the data, we might have been able to recover the true null causal effect of treatment on project completions. In a nutshell, there is no causal effect of treatment on completion, but only a causal effect of treatment on selection into the data. Matters are further complicated by the presence of backfilling. Projects entering in any given vintage of the data may have

[49]Such form of censoring can be pervasive in retrospective studies. Duration models allow us to learn from *right-censored* observations under the assumption of noninformative censoring (conditional on the model's covariates), but are not immune to the selective censoring concerns we highlight here.

been started in the more distant past. But the general concern remains: we worry that fewer young projects enter the sample for the treated group after the shock.

We provide two analyses to soothe concerns that our results are driven by sample selection bias due to differential changes in the data collection process for treated versus untreated projects. First, we restrict our completion analyses to projects that were in data vintages *prior* to the first audits. Even if the data collection process differentially changed following the shocks, projects in vintages prior to the shock would not have been affected by such changes. This sample is thus immune to the selection concerns. Second, we test a direct implication of differential censoring due to differential backfilling. If CapexDx had become less able to backfill data on treated relative to untreated projects, we should observe a downward trend in the number of treated projects relative to untreated projects before the shock. We thus test whether exposure predicts differential pre-trends in firm-level project stocks.

Figure A12: Total and incremental counts of projects by vintage

(a) Total (b) Incremental

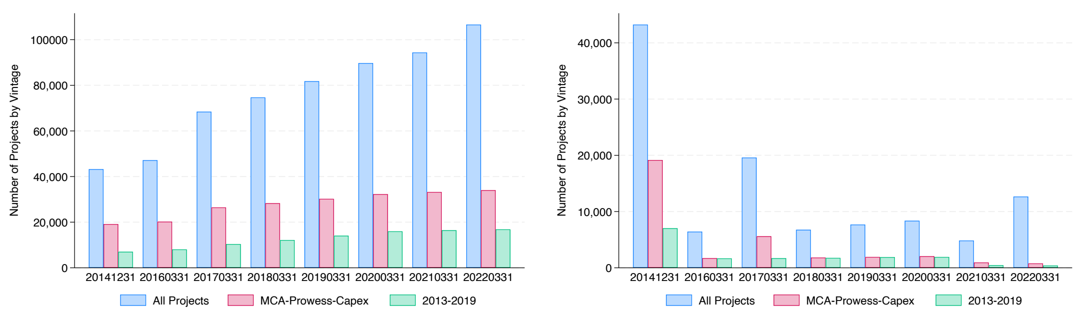

Notes: The figure plots the total (panel a) or incremental (panel b) number of projects for each available vintage of the CapexDx dataset. The x-axis indicates the vintage by its release date. Within a given vintage, we plot counts for three samples. The blue (left-most) bar corresponds to all projects. The red (middle) bar corresponds to projects in the MCA-Prowess-Capex matched dataset. The green (left-most) bar corresponds to projects in the 2013-2019 window around adoption of the Asset Quality Review.

For our first analysis, we collect previous vintages of CapexDx. The first vintage was released on December 31st, 2014. It thus predates the first Asset Quality Review audit. Further vintages were released yearly on March 31st starting in 2016, hence after the first audits had been carried out. Figure A12 plots the total number of projects in each vintage, as well as the number of projects added in each vintage. We split projects in three groups: all projects in the CapexDx data, projects with sponsors in the MCA-Prowess-Capex matched sample, and projects in the estimating sample for the completion regressions.

Most projects in our sample are already in the 2014 vintage, which allows us to reestimate the completion regressions in this subsample. Table A13 reports results from such exercise. In the first row, the firm-level exposure variable aggregates bank-level shocks using lagged, time-varying weights (LTVW). The point estimates are quantitatively similar to those ob-

Table A13: Effect on Completion Hazards: Subset of First Vintage Projects: Panel

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Exposure (LTV)	0.0235	0.0742*	0.0738*	0.0621	0.0624	0.0577	0.0560
	(0.0305)	(0.0409)	(0.0407)	(0.0416)	(0.0416)	(0.0454)	(0.0462)
Exposure (Pre)	0.0254	0.0979**	0.0997**	0.0898**	0.0932**	0.0976**	0.1016**
	(0.0312)	(0.0410)	(0.0410)	(0.0413)	(0.0414)	(0.0415)	(0.0416)
Exposure (2015)	0.0240	0.0951**	0.0984**	0.0880**	0.0922**	0.0936**	0.0985**
	(0.0304)	(0.0393)	(0.0392)	(0.0395)	(0.0394)	(0.0398)	(0.0399)
DepVar Mean	0.0379	0.0379	0.0379	0.0379	0.0379	0.0401	0.0401
Observation	54733	54733	54733	54733	54733	51245	51245
Age Splines	✓	✓	✓	✓	✓	✓	✓
Firm FE		✓	✓	✓	✓	✓	✓
State x Year FE		✓	✓	✓	✓	✓	✓
Sector x Year FE				✓	✓		
Industry x Year FE						✓	✓
Firm Controls			✓		✓		✓

Notes: the table reports PPML estimates of the semi-elasticity of project-level hazards to firm exposure to the Asset Quality Review (see model (3)). The sample consists of projects observed quarterly between 2013-2019; and is further ***restricted to projects that were present in the data as of the first vintage (Dec 31, 2014)***, thereby preceding the first Asset Quality Review. Firm exposure is defined as the weighted-average of lender exposure and weights vary by row. "(LTV)" indicates lagged, time-varying weights, as in our baseline results. "(Pre)" indicates weights are constructed as average outstanding loans between 2013-2015 for a borrower-lender pair relative to average outstanding loans between 2013-2015 for a given borrower. "(2015)" indicates weights are constructed as outstanding loans in 2015 between a borrower-lender pair relative to outstanding loans in 2015 for a given borrower. Coefficient estimates are interpreted as the percent effect on the outcome for a one standard-deviation increase in firm exposure. For example, row 1 of column (1) reports that moving from a firm with average exposure to a firm with one-standard deviation higher exposure increases the hazard by 2.35%. Age splines refer to cubic splines on a project's age (measured in quarters). Firm controls include sales, sales growth, profitability, debt over equity, book leverage, equity issuances, cash over liquid assets, and debt over EBITDA. All firm controls are winsorized at the 1st and 99th percentile, and then standardized. All firm controls are lagged by one-period. The reduction in sample size from columns (5) to columns (6) and (7) results from separation of additional observations by the more granular industry-by-year fixed effects relative to the sector-by-year fixed effects. Standard errors are clustered at the firm level and reported in parentheses. $*$, $**$, and $***$ denote statistical significance at the 10%, 5% and 1% levels, respectively.

Table A14: Cross-Sectional Balance Tests: Project Stocks

	Mean | Exp = 0 (1)	Mean | Exp > 0 (2)	Mean | Exp (3)
Active Projects	2.441	4.185	-3.3262***
			(1.1054)
Projects Started	0.504	0.683	-0.5060**
			(0.2276)
Projects Completed	0.233	0.342	-0.4336***
			(0.0998)
Projects Abandoned	0.043	0.075	-0.0377
			(0.0415)
Projects Censored	0.100	0.192	-0.0240
			(0.0779)
Projects Under Implementation	1.428	2.415	-1.9943***
			(0.6739)
Projects Announced	0.976	1.599	-1.3409***
			(0.4452)
Projects Stalled	0.176	0.355	0.0812
			(0.1501)
Exposure			✓

Notes: the table reports linear regression estimates of covariate balance across firm exposure to the Asset Quality Review. In particular, we regress firm-level outcomes averaged over 2013-2015 Y_f on average firm exposure over the 2016-2019 period $Expo_f$, where firm exposure is defined as the weighted-average of lender exposure, where weights correspond to outstanding loans in the previous year. Each row corresponds to either a stock (active, announced, under implementation, or stalled) or flow (completed, censored, abandoned, started) of projects. Columns (1) and (2) report raw means conditional on a binary exposure indicator. Column (3) reports the coefficient on the exposure variable in the cross-sectional linear regression model. Standard errors are heteroscedasticity-robust and reported in parentheses. $*$, $**$, and $***$ denote statistical significance at the 10%, 5% and 1% levels, respectively.

tained without restricting the sample. The standard errors are admittedly larger, which is unsurprising given the sample size is effectively halved by the vintage restriction. Despite the reduced power, we are content by the similar point estimates between the two specifications: if bias was a concern, we would have expected point estimates to change considerably between the two exercises.

In the second and third rows, the firm-level exposure variables weigh bank shocks by pre-audit exposures. The second row averages outstanding loans between 2013 and 2015 (PreW), whereas the third row simply takes the 2015 outstanding loan amounts (2015W). The standard errors are indistinguishable from those in the first row, but the point estimates are larger and statistically significant at conventional levels. This increase in point estimates might appear puzzling at first, but would be consistent with later-year weights becoming smaller if firms substituted away from more exposed banks.

The results in Table A13 corroborate that our main results are not driven by treatment inducing differential sample selection into later vintages. It does not, however, speak to the possibility of such differential censoring being present in the data. An implication of such selection would be differential pre-trends in the number of projects among treated and untreated firms. We test for this possibility as follows. First, we ask whether firms on average more exposed to the shock exhibit fewer projects or fewer starts before the audits ever began. The last column in table A14 shows that higher average exposure ex post predicts ex ante differences in firms' project stocks (each row corresponds to a different dependent variable). Prior to the shock, more exposed firms have on average fewer projects, start fewer projects, and complete fewer projects.

Table A14 performs a cross-sectional balance test. Such balance in levels would be reassuring, but is far more stringent than what's implied by the differential selection story (and than what's required by our identifying assumption). Level differences could be driven, for instance, by permanent differences in the set of firms that are eventually exposed or not. Our panel regressions control non-parametrically for such possibility with firm fixed effects; but firm-level regressions – as in table A14 cannot. We thus move to a panel regression, and test whether contemporaneous exposure predicts a decline in lagged number of projects.

Table A15 collects the results. The last two columns both regress lagged variables on contemporaneous treatments. The third column does not include any fixed effects. Consistent with table A14 , we cannot rule out differences in levels. But we can rule out differential pre-trends. Including firm and time effects in the last column, the predictability vanishes; point estimates become small in magnitude and economically insignificant. The results for number of projects is at odds with the predictions of differential selection.

Admittedly, our evidence directly ruling out differential backfilling by treatment status is suggestive. It reflects a fundamental data limitation. Nonetheless, we are reassured by the two robustness checks we started with. First, restricting the sample to projects already present in the 2014 vintage does not substantially alters the results from our completion regressions. Such robustness makes it unlikely that our main results are driven by sample selection concerns arising from differential changes in data collection after the shock for treated relative to untreated firms. Second, we fail to reject equal pre-trends on the number of started projects for treated versus untreated projects. Such evidence weighs against a direct implication of the differential censoring concern.

Table A15: Panel Balance Tests: Project Stocks

	Mean | Exp = 0 (1)	Mean | Exp > 0 (2)	Mean | Exp (3)	Mean | Exp (4)
L1 Active Projects	2.987	3.408	-2.0376***	-0.1223
			(0.5203)	(0.2349)
L1 Projects Started	0.473	0.649	0.0162	-0.1163
			(0.1167)	(0.1217)
L1 Projects Completed	0.255	0.332	-0.0630	0.0410
			(0.0730)	(0.0749)
L1 Projects Abandoned	0.042	0.032	-0.0510***	0.0292
			(0.0174)	(0.0242)
L1 Projects Censored	0.124	0.174	-0.0013	-0.0368
			(0.0414)	(0.0583)
L1 Projects Under Implementation	1.768	2.038	-1.3085***	-0.1103
			(0.3334)	(0.1944)
L1 Projects Announced	1.104	1.241	-0.6835***	-0.0655
			(0.2054)	(0.1568)
L1 Projects Stalled	0.228	0.259	-0.1360**	-0.0042
			(0.0572)	(0.0481)
Exposure			✓	✓
Firm + Year FE				✓

Notes: the table reports linear regression estimates of covariate balance across firm exposure to the Asset Quality Review. The sample consists of the Capex-Prowess matched sample of firms observed yearly between 2013-2019. In particular, we regress lagged firm outcomes $Y_{f,t-1}$ on contemporaneous firm exposure $Expo_{f,t}$. Each row corresponds to either a stock (active, announced, under implementation, or stalled) or flow (completed, censored, abandoned, started) of projects. Firm exposure is defined as the weighted-average of lender exposure, where weights correspond to outstanding loans in the previous year. Columns (1) and (2) report raw means conditional on a binary exposure indicator. Column (3) reports the coefficient on the exposure variable in the panel linear regression model (balance in levels). Column (4) adds firm and time fixed effects (balance in within-unit trends). Standard errors are clustered at the firm level and reported in parentheses. $*$, $**$, and $***$ denote statistical significance at the 10%, 5% and 1% levels, respectively.

E Additional Quantitative Results

Figure A15: Calibration Exercises

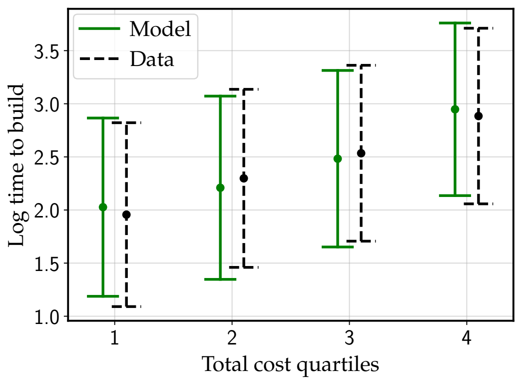

(a) Project Size and Time to Build

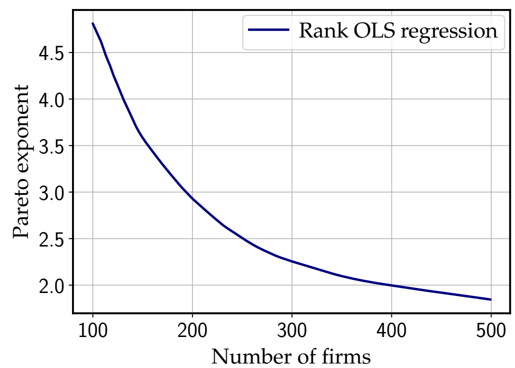

(b) Project Output

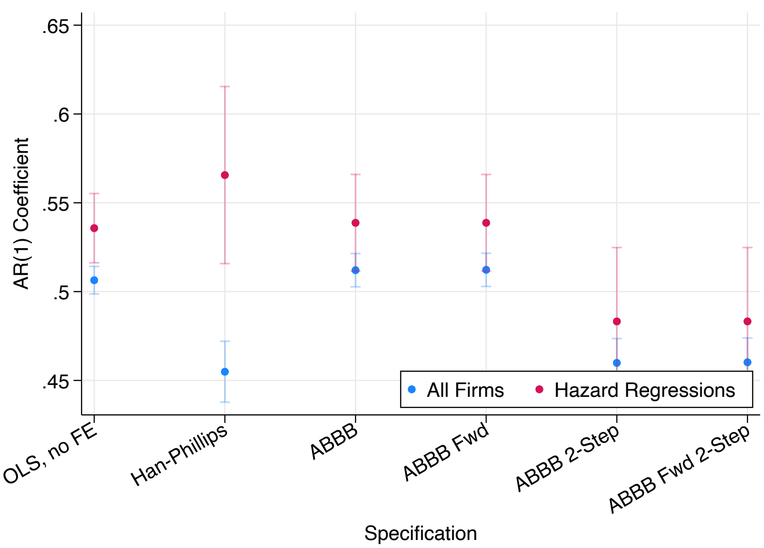

(c) Shock Persistence

Notes: The figure plots the results of our calibration exercises. Subfigure A15a plots the joint distribution of costs and time to build in the model (green, solid) and in the data (black, dashed). We leverage such moments to calibrate the parameters related to the distributions of $\bar{x}$ and $\bar{\iota}$. Subfigure A15b plots the results from estimating log-rank regressions. The x-axis plots the number of (top) firms included in the estimation. The y-axis plots the resulting coefficient of a regression of log sales on firm rank, which corresponds to an estimate of the Pareto exponent. We choose a midpoint of 3. Subfigure A15c plots coefficient estimates of multiple panel autoregressive models used to estimate the persistence of the discount rate shock ρ. We regress firm exposure on its first lag, clustering at the firm level. Each bin in the x-axis corresponds to a distinct estimator. From left to right, we report estimates from OLS without fixed effects, the Han-Phillips GMM estimator, and multiple variations of the Arellano-Bover-Blundell-Bond system GMM estimator ("Fwd" indicates taking forward deviations, and "2-Step" indicates use of two-step GMM with optimal covariance matrix). Blue dots plot estimates (and 95% confidence bands) in the sample of all firm-years in the MCA-Prowess matched data. Red dots restrict the sample to firms in our hazard regressions.

Table A16: Calibration Results: Parameter Estimates

Parameter	Value	Interpretation	Source
r	0.01	Real Rate	Post-1990 Time Series Average
p	1	Price of K	Normalization
$\mu_{\bar{x}}$	0	Avg Project Size	
$\sigma^2_{\bar{x}}$	2.12	SD Project Size	Cost-TTB Distribution
$\beta^{\bar{i}}_0$	-2.47	Mean $\bar{i}$: intercept	
$\beta^{\bar{i}}_1$	0.82	Mean $\bar{i}$: slope	
$\sigma^2_{\bar{i}}$	0.88	SD $\bar{i}$	
ρ	0.17	Shock Persistence	Panel Autoregressions
α	3	Pareto Tail	Log-Rank Regressions
s	0.2	Shock Size	Reduced Form Regressions
κ	3.09	Accel. Cost	

Notes: the table reports parameter estimates from our calibration exercise. See the notes to figure A15 and the main text for detailed explanations of the calibration strategy.

Table A17: Calibration Results: Targeted Moments

Moment	Model	Data
Avg Log Size (Q1)	-2.68	-2.58
Avg Log Size (Q2)	-1.42	-1.57
Avg Log Size (Q3)	0.01	-0.18
Avg Log Size (Q4)	1.45	1.42
Avg Log Size (Q5)	2.77	2.80
SD Log Size	0.91	0.90
Avg Log TTB (Q1)	1.96	2.02
Avg Log TTB (Q2)	2.30	2.21
Avg Log TTB (Q3)	2.53	2.48
Avg Log TTB (Q4)	2.88	2.95
Hazard Regression β	0.06	0.06
Starts Regression β	-0.08	-0.08

Notes: the table reports values for the targeted moments in our calibration exercise. See the notes to figure A15 and the main text for detailed explanations of the calibration strategy.

F Proofs and Additional Theoretical Results

F.1 Time to Completion and the Kolmogorov Backward Equation

Let $T(x,t)$ denote the time remaining until completion for a project in stage x at time t. In a time span Δt, to first-order, we have

$$T(x,t) \approx \Delta t + T(x + i(x)\Delta t, t + \Delta t)$$

Rearranging, we obtain

$$0 \approx 1 + \frac{T(x + i(x)\Delta t, t + \Delta t) - T(x,t)}{\Delta t}$$

For $i(x)\Delta t$ small,

$$T(x + i(x)\Delta t, t + \Delta t) \approx T(x, t + \Delta t) + i(x)\Delta t \frac{\partial}{\partial x} T(x,t)$$

Note further that

$$\frac{T(x, t + \Delta t) - T(x,t)}{\Delta t} \to \frac{\partial}{\partial t} T(x,t)$$

Hence, as $\Delta t \to 0$, we obtain a Kolmogorov Backward Equation

$$0 = 1 + i(x,t)\frac{\partial}{\partial x} T(x,t) + \frac{\partial}{\partial t} T(x,t) \tag{21}$$

with terminal condition

$$T(\overline{x}, t) = 0 \tag{22}$$

In steady state, we have $\frac{\partial}{\partial t} T(x,t) = 0$, which yields the $T'(x) = -1/i(x)$.

Proposition 8. *In steady state, completion times are uniformly distributed.*

Proof. The distribution of completion times is

$$F_T(t) = \text{Prob}(T(x) \leq t) = \int_{T^{-1}(t)}^{\bar{x}} g(x)dx$$

i.e. the mass of projects being completed up to time t is just the mass of projects about to be completed (close to $\bar{x}$) down to projects in stages such that after t periods, the project will be completed. Note that $g(x)$ does not integrate to 1, but this is not a serious concern.

Using Leibniz's rule, we differentiate the cumulative to obtain the density

$$f_T(t) = -g(T^{-1}(t)) \times (T^{-1})'(t)$$

Recall that the derivative of the inverse function is given by

$$(T^{-1})'(t) = \frac{1}{T'(T^{-1}(t))},$$

but $T'(x) = -1/i(x)$ and $g(x) = f/i(x)$, so for $x = T^{-1}(t)$, we obtain

$$\begin{aligned} f_T(t) &= \left(-\frac{f}{i(T^{-1}(t))}\right) \times \left(-\frac{1}{1/i(T^{-1}(t))}\right) \\ &= f \end{aligned}$$

$\square$

F.2 Parallel to Bond Pricing

The results in Section 5 capture a broad economic mechanism related to (changes in) the time value of money. We illustrate this point by returning to standard fixed income pricing.

Setup Consider a bond paying coupons $c(t)$ and with face value F. All cash flows are discounted at a rate r. The remaining time to maturity is T. The value of the bond is

$$V = \int_0^T e^{-rt} c(t) dt + e^{-rT} F$$

Recall that when interest rate changes, the change in the bond price is given by the bond's duration

$$\frac{\partial V}{\partial r} = -\int_0^T t e^{-rt} c(t) dt - e^{-rT} F \cdot T = -\text{Dur}(V)$$

This formulation maps to Section 5 as follows. The coupons are just the investment flows required to complete the project. The face value, in turn, is the return to completing the project, given by pz in the notation of Section 5. The bond price (or, value) V is akin to the project value $v(x,z)$, in which case the maturity of the project is instead the remaining time to project completion $T(x,z)$. The duration expression parallels Proposition 2, so that the $\frac{\partial V}{\partial r}$ is akin to $\hat{v}(x,z)$. The key difference is that we optimally choose investment flows for projects in Section 5. In turn, bond investors don't choose the path of coupons accruing to bonds (when they are not the issuers, at least).

"Accelerating" Bond Cash Flows Consider the following experiment. Suppose we "accelerate" the cash flows of the bond. To be precise, we shift all cash flows backwards in time by Δ. The last coupon is $c(T)$, which is now received at $T-\Delta$. The face value is also received at $T-\Delta$. At date 0, the coupon is now $c(\Delta)$ rather than $c(0)$. That is, the bond holder forfeits all coupon payments in the time interval $[0, \Delta]$. The value of this (new) bond is

$$V(\Delta) = \int_0^{T-\Delta} e^{-rt} c(t+\Delta) dt + e^{-r(T-\Delta)} F$$

where $\Delta = 0$ corresponds to the original bond. Consider a marginal change in the time shift Δ. We have,

$$\frac{\partial V(\Delta)}{\partial \Delta} = \int_0^{T-\Delta} e^{-rt} \frac{\partial c(t+\Delta)}{\partial \Delta} dt - e^{-r(T-\Delta)} c(T) + re^{-r(T-\Delta)} F \tag{23}$$

To build intuition about (23), consider first the last two terms. We have

$$e^{-r(T-\Delta)} \left(-c(T) + rF\right) = -\frac{\partial V(\Delta)}{\partial T}$$

so that the last two terms correspond to the case where we just exogenously reduce the maturity of the bond, but do not shift any of the other cash flows. In that case, we lose the terminal coupon $c(T)$, but we get the face value quicker, which we value at rF. We then discount the cash flows to today to obtain the marginal change in the bond value. For most bonds, coupon payments are small relative to face values, so that ${}^{\partial V(\Delta)}/_{\partial T}$ is positive. That is, bond prices should rise if maturities exogenously fell.

The first term of (23) reflects the fact that we are shifting all cash flows over time, not just the terminal cash flows. The cash flow that would have been received in period $t + \Delta$ is now received at time t. Rather than being discounted by $e^{-r(t+\Delta)}$, it is now discounted by e^{-rt}. Consequently, the value of the bond must change to reflect this new path of discount rates.

Note that

$$\frac{\partial c(t+\Delta)}{\partial \Delta} = \frac{\partial c(t+\Delta)}{\partial t},$$

which follows from the fact that we are uniformly shifting cash flows over time. Using this result, we can integrate the first term in (23) by parts

$$\begin{aligned}\int_0^{T-\Delta} e^{-rt} \frac{\partial c(t+\Delta)}{\partial \Delta} dt &= \left[e^{-rt} c(t)\right]_0^{T-\Delta} - \left[\int_0^{T-\Delta} -re^{-rt} c(t+\Delta) dt\right] \\ &= e^{-r(T-\Delta)} c(T-\Delta) - c(\Delta) + r\int_0^{T-\Delta} e^{-rt} c(t+\Delta) dt\end{aligned}$$

Bringing all together and rearranging, we find

$$\frac{\partial V(\Delta)}{\partial \Delta} = -c(\Delta) + rV(\Delta) \tag{24}$$

Expression (24) is a key object. When we shift all cash flows by Δ, we lose the marginal coupon today. What we gain is the flow value from holding a new bond without that coupon, but with a shorter maturity and the original cash flows. In bond markets, we see securities with the same coupons, but with different residual maturities being traded at the same time. For example, we might observe two sovereign bonds with equal coupons, but one has T days remaining until maturity while the other has $T - \Delta$ days remaining. For a bond investor with fixed discount rate r, how much would the investor value swapping the former for the latter? Since one bond is just an intertemporal shift of the cash flows of the other, exactly as in our experiment, expression (24) provides the answer.

Suppose now that the bond investor could pay $\kappa(\Delta)$ today to accelerate cash flows by Δ.[50] We assume κ is differentiable and increasing. The necessary condition for an optimal choice of "acceleration" is

$$\kappa'(\Delta^*) = \frac{\partial V(\Delta^*)}{\partial \Delta} \tag{25}$$

We assume that an interior solution exists. A priori Δ^* could be positive or negative. An investor could prefer to shift cash flows further into the future if it was expecting a coupon in the next Δ^* that was really large relative to the remaining cash flows accruing to the bond. For most bonds, however, we think that coupons are small relative to face values. That is, for must bonds, we tend to think that the capital gains accruing from receiving a large payoff earlier outweighs the lost coupons. Formally, (24) is positive and therefore $\Delta^* > 0$.

A notable feature of (25) is that it carries the same economic content as (11), the FOC for the investment problem with endogenous time to build. Both equate the marginal cost of accelerating cash flows to the marginal benefits of doing so. While expression (11) is written in stage space, expression (25) is written in time space. However, we know that $dx = idt$, so that (11) can be rewritten in time space and yield a parallel expression to (25). In short, the underlying economics carries across the bond and the project settings.

Suppose now that discount rates change, as in our empirical exercise. The value of "accelerating" cash flows changes, resulting in a change in Δ^*. From (25), it follows that

$$\frac{\partial \Delta^*}{\partial r} = \frac{1}{\kappa''(\Delta^*)} \cdot \frac{\partial^2 V(r, \Delta^*)}{\partial r \partial \Delta} = \frac{1}{\kappa''(\Delta^*)} \cdot \frac{\partial^2 V(r, \Delta^*)}{\partial \Delta \partial r} \tag{26}$$

which again bears a striking resemblance to (17), the expression characterizing the impact effect of a discount rate shock on investment flows. Assuming convexity of κ, bond (project) investors prefer to speed up cash (investment) flows if the increase in discount rate increases the marginal benefit of speeding up cash (investment) flows. This corresponds to the first equality in expression (26). By Fubini's theorem, an equivalent interpretation is that bond (project) investors speed up cash (investment) flows if doing so reduces the bond (project)

[50]One way to think about κ is to model the supply side of a bond market where broker-dealers must be compensated for supplying identical bonds of differential maturities to investors that would like to do such swaps. They must manage such inventory and charge a fee depending on the maturity swap.

duration. This is precisely the intuition we emphasized in the main text.

Differentiating (24) with respect to r yields

$$\frac{\partial V(r,\Delta)}{\partial r \partial \Delta} = \frac{\partial}{\partial r}\left(rV(r,\Delta)\right) = \underbrace{V(r,\Delta)}_{\text{Impatience}(+)} \underbrace{-r \cdot \mathrm{Dur}(V(r,\Delta))}_{\text{Discounting}(-)} \tag{27}$$

Expression (27) spells out the economic forces determining whether cash flows are sped up or slowed down. A first effect, which we call "impatience", is mechanical. Holding fixed the value of the asset, higher discount rates means investors are more impatient for cash flows. This effect is always positive. The second effect, which we call "discounting", comes from the fact that as we change discount rates, we change the NPV of future cash flows and therefore revalue the asset. To first-order, the changes in asset value are summarized by its duration. The total effect adds both terms, and is given by the NPV of all cash flows weighted by $(1 - rt)$.

Expression (27) immediately implies that the effect of discount rate shocks vary across similar bonds with different residual maturities. For bonds (projects) close to maturity (completion), duration is small, implying that it is optimal to speed up cash (investment) flows, consistent with Corollary (4). In turn, for bonds (projects) with large residual maturity (time to build), duration is large, implying it might even be optimal to slow down the pace of coupons (investments).

The parallel between bond pricing and project development is admittedly not one-to-one. Two key differences exist. First, we optimally choose the "coupons" (investment flows) of the project and therefore its residual "maturity" (time to build). Second, the project requires a fixed amount if investment to be paid in in order to reap the terminal payoff. In turn, the bond pays the face value at maturity regardless of the path of coupons beforehand. In that sense, an experiment where we shift all of the bonds' cash flows, but do not require the remaining coupons to "add up" in value to the value of the coupons before the shift is not a perfect parallel to the project investment case. But even without the additional complications of designing a variational path for coupons that preserved their sum, we obtained the same intuition for the decision to speed up cash flows. All in all, while the parallel between bond pricing and project development is not perfect, it illustrates how the key forces driving the responses in both cases relate to the time value of money, and how changes in the price of cash flows today versus tomorrow have differing implications for long-lived assets with differing residual maturities.

F.3 First-Order Dynamics: Terminal Payoff Shocks

In the main text, we restricted attention to the first-order dynamics of the value function and of investment following a AR(1) shock to discount rates. It is feasible (and easier) to derive similar dynamics in the case of a shock to the terminal payoff of the project, i.e. a shock of the form

$$p(t) = p(1 + se^{-\rho t})$$

We leverage the results from Lemma 1, which generalizes straightforwardly to p-shocks. With a slight abuse of notation, we again use "hatted", lower-case variables to denote the partial derivatives of a function with respect to the shock size s, evaluated at the steady state.

Proposition 9. *The first order response of project value to a AR(1) shock to terminal payoff is given by*

$$\hat{v}(x, z) = pze^{-(r+\rho)T(x,z)}, \tag{28}$$

and the first order response of investment flows to the same shock is given by

$$\hat{i}(x, z) = \frac{1}{\kappa''(i(x, z))} \frac{\partial \hat{v}(x, z)}{\partial x} \tag{29}$$

Proof. The time-dependent HJB is given by

$$r\mathcal{V}(x, z, s, t) = -\mathcal{I}(x, z, t, s) - \kappa(\mathcal{I}(x, z, t, s)) + \mathcal{I}(x, z, t, s)\frac{\partial}{\partial x}\mathcal{V}(x, z, t, s) + \frac{\partial}{\partial t}\mathcal{V}(x, z, t, s)$$

with terminal condition

$$\mathcal{V}(\bar{x}, z, s, t) = pz(1 + se^{-\rho t})$$

and where $\mathcal{I}(x, t, s)$ solves

$$1 + \kappa'(\mathcal{I}(x, z, t, s)) = \frac{\partial}{\partial x}\mathcal{V}(x, z, t, s)$$

Differentiate the HJB and the terminal condition with respect to s using the Envelope Theorem

$$r\frac{\partial}{\partial s}\mathcal{V}(x, z, s, t) = \mathcal{I}(x, z, t, s)\frac{\partial}{\partial s}\frac{\partial}{\partial x}\mathcal{V}(x, t, s) + \frac{\partial}{\partial s}\frac{\partial}{\partial t}\mathcal{V}(x, z, t, s)$$

$$\frac{\partial}{\partial s}\mathcal{V}(\bar{x}, z, s, t) = pze^{-\rho t}$$

Define the steady-state gradients

$$\hat{v}(x, z, t) := \left.\frac{\partial}{\partial s}\mathcal{V}(x, z, s, t)\right|_{s=0}$$

and evaluate the HJB at the steady state, after switching the order of differentiation

$$r\hat{v}(x, z, t) = i(x, z)\frac{\partial}{\partial x}\hat{v}(x, z, t) + \frac{\partial}{\partial t}\hat{v}(x, z, t)$$

Define $\hat{v}(x,z) := \frac{\partial}{\partial s}\mathcal{V}(x,z,s,t)\Big|_{s,t=0}$ and conjecture $\hat{v}(x,z,t) = e^{-\rho t}\hat{v}(x,z)$. Then, the HJB becomes

$$(r+\rho)\hat{v}(x,z) = i(x,z)\frac{\partial}{\partial x}\hat{v}(x,z)$$

which is an ODE on $\hat{v}(x,z)$ with terminal condition $\hat{v}(\overline{x}) = pz$. Recalling ${\partial T(x,z)}/{\partial x} = -1/{i(x,z)}$ in steady state, we can solve the ODE subject to the terminal condition and obtain (28).

Moving towards investment, differentiate the FOC with respect to s

$$\kappa''(\mathcal{I}(x,z,t,s))\frac{\partial}{\partial s}\mathcal{I}(x,z,t,s) = \frac{\partial}{\partial s}\frac{\partial}{\partial x}\mathcal{V}(x,z,t,s)$$

Switching the order of differentiation and evaluating at a steady state, we obtain

$$\kappa''(i(x,z))\hat{i}(x,z,t) = \frac{\partial}{\partial x}\hat{v}(x,z,t)$$

where we defined $\hat{i}(x,z,t) := \frac{\partial}{\partial s}\mathcal{I}(x,z,s,t)\Big|_{s=0}$. Once again, we conjecture $\hat{i}(x,z,t) = e^{-\rho t}\hat{i}(x,z)$ (and analogously for $\hat{v}(x,z,t)$), where as before $\hat{i}(x,z) := \frac{\partial}{\partial s}\mathcal{I}(x,z,s,t)\Big|_{s,t=0}$. Cancelling the exponential term, we readily obtain (29). □

Proposition 9 states that we can obtain a sense of the change in project value simply by discounting the terminal payoff. If the terminal payoff rises (falls), value rises (falls). The actual insight of the proposition is to tell us the correct discount factor and the correct time length for the discounting exercise. Since the shock in the terminal payoff is persistent, the discount rate rises from r to $r+\rho$. In the limit of a permanent change in terminal payoffs, the discount rate is just r. When $\rho < 1$, the shock is temporary, so we raise the discount factor to account for the fading dynamics.

Most crucially, we relevant time frame for discounting is just given by the steady state remaining time to build $T(x,z)$. Fledgling projects respond *less* to price shocks relative to mature projects. For any fixed discount rate, their terminal cash flows are far away. Discounting pushes their value downwards. Unless the shock is very persistent, it is unlikely that the firm will be able to capture the benefits of the shock. Granted, since our shock is a AR(1) shock, it only completely dies out in the limit of $t \to \infty$, so even projects that have just been started still have some reaction to the shock. But the benefits that will accrue to such young projects are smaller, for it will take longer to reach completion.

Turning to the second result of Proposition 9, we find that investment rises (falls) following a positive (negative) shock to terminal payoffs – the acceleration costs are assumed to be convex and $\frac{\partial}{\partial x}\hat{v}(x,z) > 0$. In the case where κ is quadratic, as in our quantitative exercises, the model makes an unconditional prediction that investment in more mature projects reacts more to shocks to terminal payoffs. The intuition falls off from the result that such mature projects have greater changes in value following the shock. When convexity depends on the level of investment, whether such monotonicity on project stage holds relies on quantitative exploration. Since steady state investment is increasing in project stage, the total effect depends on the relative magnitudes of valuation gains versus acceleration costs, both of which become increasing in project stage.

Last, we turn to the extensive margin. Proposition 10 shows that a positive terminal payoff shock lowers the entry threshold and therefore boosts project starts, consistent with standard investment theory.

Proposition 10. *The first-order response of the entry threshold $\bar{z}$ to an AR(1) terminal payoff shock is given by*

$$\hat{z} = -\bar{z}e^{-\rho T(0,\bar{z})} \tag{30}$$

Proof. We can follow the same steps as in the proof of Proposition 5. Using (28) and (37), we find

$$\begin{aligned} \hat{z} = \frac{\partial \bar{z}}{\partial s} &:= -\frac{\partial v(x,z)/\partial x}{\partial v(x,z)/\partial z}\bigg|_{(x,z)=(0,\bar{z})} \\ &= -\frac{p\bar{z}e^{-(r+\rho)T(0,\bar{z})}}{pe^{-rT(0,\bar{z})}} = -\bar{z}e^{-\rho T(0,\bar{z})} \end{aligned}$$

$\square$

F.4 Proofs from Main Text

F.4.1 Proof of Lemma 1 and Proposition 2

We omit z from the derivations to lighten notation. The results hold pointwise along the distribution of z.

Step 1. Linearizing the HJB (Proof of Lemma 1) The time-dependent HJB is given by

$$(r + se^{-\rho t})\mathcal{V}(x,s,t) = -\mathcal{I}(x,t,s) - \kappa(\mathcal{I}(x,t,s)) + \mathcal{I}(x,t,s)\frac{\partial}{\partial x}\mathcal{V}(x,t,s) + \frac{\partial}{\partial t}\mathcal{V}(x,t,s)$$

with terminal condition

$$\mathcal{V}(\bar{x},s,t) = pz$$

and where $\mathcal{I}(x,t,s)$ solves

$$1 + \kappa'(\mathcal{I}(x,t,s)) = \frac{\partial}{\partial x}\mathcal{V}(x,t,s)$$

Differentiate the time-dependent HJB with respect to the shock size s

$$
\begin{aligned}
r\frac{\partial}{\partial s}\mathcal{V}(x,s,t)+e^{-\rho t}\mathcal{V}(x,s,t) = &-\left(1+\kappa'(\mathcal{I}(x,t,s))\right)\frac{\partial}{\partial s}\mathcal{I}(x,s,t) \\
&+\frac{\partial}{\partial s}\mathcal{I}(x,s,t)\cdot\frac{\partial}{\partial x}\mathcal{V}(x,t,s) \\
&+\mathcal{I}(x,t,s)\frac{\partial}{\partial s}\frac{\partial}{\partial x}\mathcal{V}(x,t,s)+\frac{\partial}{\partial s}\frac{\partial}{\partial t}\mathcal{V}(x,t,s)
\end{aligned}
$$

and from the FOC, a standard Envelope condition arises, yielding

$$
r\frac{\partial}{\partial s}\mathcal{V}(x,s,t)+e^{-\rho t}\mathcal{V}(x,s,t)=\mathcal{I}(x,t,s)\frac{\partial}{\partial x}\frac{\partial}{\partial s}\mathcal{V}(x,t,s)+\frac{\partial}{\partial t}\frac{\partial}{\partial s}\mathcal{V}(x,t,s) \tag{31}
$$

where we switched the order of differentiation in the last terms.

We define the steady-state gradients

$$
\hat{v}(x,t) := \left.\frac{\partial}{\partial s}\mathcal{V}(x,s,t)\right|_{s=0}
$$

so evaluating (31) at the steady state,

$$
r\hat{v}(x,t)+e^{-\rho t}v(x)=i(x)\frac{\partial}{\partial x}\hat{v}(x,t)+\frac{\partial}{\partial t}\hat{v}(x,t)
$$

where $v(x)$ and $i(x)$ are the steady state value and policy functions, respectively.

We conjecture

$$
\begin{aligned}
\hat{v}(x,t) &= e^{-\rho t}\hat{v}(x) \\
\hat{v}(x) &:= \left.\frac{\partial}{\partial s}\mathcal{V}(x,s,t)\right|_{s=0,t=0}
\end{aligned}
$$

so that

$$
re^{-\rho t}\hat{v}(x)+e^{-\rho t}v(x)=i(x)e^{-\rho t}\frac{d}{dx}\hat{v}(x)-\rho e^{-\rho t}\hat{v}(x)
$$

Cancelling the $e^{-\rho t}$ terms and rearranging, we obtain

$$
(r+\rho)\hat{v}(x)+v(x)=i(x)\frac{d}{dx}\hat{v}(x)
$$

Differentiating the terminal condition with respect to s gives us

$$
\frac{\partial}{\partial s}\mathcal{V}(\bar{x},s,t)=0
$$

Evaluating at the steady state yields

$$\hat{v}(\bar{x}, t) = 0$$

and given our conjecture, we obtain

$$\hat{v}(\bar{x}) = 0$$

Hence, we obtain an ordinary differential equation in $\hat{v}(x)$ with associated terminal condition

$$\begin{aligned}(r+\rho)\hat{v}(x) + v(x) &= i(x)\frac{d}{dx}\hat{v}(x) \\ \hat{v}(\overline{x}) &= 0\end{aligned} \tag{32}$$

Step 2. Solving the ODE for $\hat{v}(x)$ We rearrange (32) to obtain

$$\frac{d}{dx}\hat{v}(x) - \frac{r+\rho}{i(x)}\hat{v}(x) = \frac{1}{i(x)}v(x)$$

Multiplying by the integrating factor $e^{(r+\rho)T(x)}$, we have

$$e^{(r+\rho)T(x)}\frac{d}{dx}\hat{v}(x) - \frac{r+\rho}{i(x)}e^{(r+\rho)T(x)}\hat{v}(x) = \frac{e^{(r+\rho)T(x)}}{i(x)}v(x)$$

And recall that $T'(x) = -1/i(x)$, so

$$\frac{d}{dx}\left[e^{(r+\rho)T(x)}\hat{v}(x)\right] = \frac{e^{(r+\rho)T(x)}}{i(x)}v(x)$$

Integrating from x to $\bar{x}$ and using the fundamental theorem of calculus, we have

$$e^{(r+\rho)T(\bar{x})}\hat{v}(\bar{x}) - e^{(r+\rho)T(x)}\hat{v}(x) = \int_x^{\overline{x}} \frac{e^{(r+\rho)T(y)}}{i(y)}v(y)dy$$

where we have changed the running variable for integration on the RHS from x to y. Now, from the terminal condition

$$\begin{aligned}-e^{(r+\rho)T(x)}\hat{v}(x) &= \int_x^{\overline{x}} \frac{e^{(r+\rho)T(y)}}{i(y)}v(y)dy \\ \Rightarrow \hat{v}(x) &= -\int_x^{\overline{x}} \frac{e^{-(r+\rho)(T(x)-T(y))}}{i(y)}v(y)dy\end{aligned} \tag{33}$$

Step 3. Changes of Variables The key step in simplifying expression (33) is to move the right-hand side integral from stage space to time space. Let

$$\tau = T(x) - T(y)$$
$$\Rightarrow d\tau = \frac{dy}{i(y)}$$

Note further that $y = x \Leftrightarrow \tau = 0$ and $y = \bar{x} \Leftrightarrow \tau = T(x)$, so changing the integration bounds, we find

$$\hat{v}(x) = -\int_0^{T(x)} e^{-(r+\rho)\tau} v(y(\tau)) d\tau \tag{34}$$

where $y(\tau)$ is defined as the position of the project that starts at x at $\tau = 0$ and evolves according to $dy = i(y)d\tau$.

Consider in turn the sequence formulation of the firm problem. Instead of summing cash flows over time, we can sum cash flows over the states that the project will go through until its eventual completion. Formally,

$$v(y) = -\int_y^{\bar{x}} \frac{e^{-r[T(y)-T(u)]}}{i(u)} \left[i(u) + \kappa\left(i(u)\right)\right] du + e^{-rT(y)} pz \tag{35}$$

The left-hand side is the value of a project starting in stage y. After $T(y)$ periods, it will be completed and yield pz. At each stage u until it reaches completion, the firm spends $i(u) + \kappa(i(u))$. Converting the sum from a time summation to a summation over positions follows from the same change of variables considered above: each time step becomes $du/i(u)$, and we renormalize time to run from the time at which we are in stage y onwards. The bounds of the integral follow directly from the fact that the project starts at y and is finished once it reaches $\bar{x}$.

Substituting (35) into (34), we obtain

$$\begin{aligned}
\hat{v}(x) &= -\int_x^{\overline{x}} \frac{e^{-(r+\rho)(T(x)-T(y))}}{i(y)} \left[-\int_y^{\bar{x}} \frac{e^{-r[T(y)-T(u)]}}{i(u)} \left[i(u) + \kappa\left(i(u)\right)\right] du + e^{-rT(y)} pz\right] dy \\
&= \underbrace{\int_x^{\overline{x}} \int_y^{\overline{x}} \frac{e^{-(r+\rho)(T(x)-T(y))}}{i(y)} \frac{e^{-r[T(y)-T(u)]}}{i(u)} \left[i(u) + \kappa\left(i(u)\right)\right] du dy}_{\hat{v}_1(x)} - \underbrace{\int_x^{\overline{x}} \frac{e^{-(r+\rho)(T(x)-T(y))}}{i(y)} e^{-rT(y)} pz dy}_{\hat{v}_2(x)}
\end{aligned}$$

Step 4. Simplifying Algebra We start with $\hat{v}_1(x)$. We change the order of integration, so

$$\begin{aligned}
\hat{v}_1(x) &= \int_x^{\overline{x}} \int_y^{\bar{x}} \frac{e^{-(r+\rho)(T(x)-T(y))}}{i(y)} \frac{e^{-r[T(y)-T(u)]}}{i(u)} \left[i(u) + \kappa\left(i(u)\right)\right] dudy \\
&= \int_x^{\overline{x}} \int_x^{u} \frac{e^{-(r+\rho)(T(x)-T(y))}}{i(y)} \frac{e^{-r[T(y)-T(u)]}}{i(u)} \left[i(u) + \kappa\left(i(u)\right)\right] dydu \\
&= \int_x^{\overline{x}} \left[\frac{i(u) + \kappa\left(i(u)\right)}{i(u)}\right] e^{-(r+\rho)T(x)} e^{rT(u)} \left(\int_x^{u} \frac{e^{\rho T(y)}}{i(y)} dy\right) du
\end{aligned}$$

Noting that

$$\begin{aligned}
\int_x^{u} \frac{e^{\rho T(y)}}{i(y)} dy &= \int_x^{u} \left(\frac{d}{dy} - \frac{e^{\rho T(y)}}{\rho}\right) dy \\
&= -\frac{1}{\rho}\left(e^{\rho T(u)} - e^{\rho T(x)}\right)
\end{aligned}$$

We can write

$$\hat{v}_1(x) = \frac{1}{\rho} \int_x^{\overline{x}} \left[\frac{i(u) + \kappa\left(i(u)\right)}{i(u)}\right] e^{-r(T(x)-T(u))} \left(1 - e^{-\rho(T(x)-T(u))}\right) du$$

Last, change variables from u to $\tau = T(x) - T(u)$. Then, ${}^{du}/_{i(u)} = d\tau$, the bounds now run from 0 to $T(x)$, and we define u as the stage $y(\tau)$ of the project starting in stage x, i.e. for which $y(0) = x$. Putting all together,

$$\hat{v}_1(x) = \int_0^{T(x)} \left[\frac{i(y(\tau)) + \kappa\left(i(y(\tau))\right)}{i(u)}\right] e^{-r\tau} \left(\frac{1 - e^{-\rho\tau}}{\rho}\right) d\tau$$

Moving to $\hat{v}_2(x)$, we intergrate

$$
\begin{aligned}
\hat{v}_2(x) &= \int_x^{\overline{x}} \frac{e^{-(r+\rho)(T(x)-T(y))}}{i(y)} e^{-rT(y)} pzdy \\
&= pze^{-(r+\rho)T(x)} \int_x^{\overline{x}} \frac{e^{\rho T(y)}}{i(y)} dy \\
&= e^{-rT(x)} pz \left(\frac{1-e^{-\rho T(x)}}{\rho} \right)
\end{aligned}
$$

So bringing all together yields (16) and completes the proof.

F.4.2 Proof of Proposition 3

We omit z from the derivations to lighten notation. The results hold pointwise along the distribution of z. Differentiate the FOC

$$1 + \kappa'(\mathcal{I}(x,t,s)) = \frac{\partial}{\partial x} \mathcal{V}(x,t,s)$$

with respect to s to obtain

$$\kappa''(\mathcal{I}(x,t,s)) \frac{\partial}{\partial s} \mathcal{I}(x,t,s) = \frac{\partial}{\partial s} \frac{\partial}{\partial x} \mathcal{V}(x,t,s)$$

Switching the order of differentiation and evaluating at a steady state, we obtain

$$\kappa''(i(x)) \hat{i}(x,t) = \frac{\partial}{\partial x} \hat{v}(x,t)$$

where we defined $\hat{i}(x,t) := \frac{\partial}{\partial s} \mathcal{I}(x,s,t)\Big|_{s=0}$. Once again, we conjecture $\hat{i}(x,t) = e^{-\rho t} \hat{i}(x)$ (and analogously for $\hat{v}(x,t)$), where as before $\hat{i}(x) := \frac{\partial}{\partial s} \mathcal{I}(x,s,t)\Big|_{s,t=0}$. Thus,

$$\hat{i}(x) = \frac{\partial \hat{v}(x)/\partial x}{\kappa''(i(x))}$$

Differentiating (33) with respect to x using Leibniz's rule, we obtain

$$
\begin{aligned}
\frac{\partial \hat{v}(x)}{\partial x} &= \frac{1}{i(x)} \left[v(x) - (r+\rho) \int_x^{\overline{x}} \frac{e^{-(r+\rho)[T(x)-T(y)]}}{i(y)} v(y) dy \right] \\
&= \frac{1}{i(x)} \left[v(x) - (r+\rho) \hat{v}(x) \right]
\end{aligned}
$$

which completes the proof.

To get the capital budgeting rule, note that

$$\frac{\partial \hat{v}(x)}{\partial x} = \frac{v(x)}{(r+\rho)i(x)}\left[\frac{1}{r+\rho} - \frac{\hat{v}(x)}{v(x)}\right]$$

It remains to show that $\hat{v}_{perp}/v_{perp} = -1/(r+\rho)$. To see this, write down the HJB equation for valuing a perpetual bond with unit cashflow

$$ru(t) = 1 + \frac{\partial u(t)}{\partial t}$$

(or, more simply, the ODE describing of the value of a perpetuity over time). Given a first-order shock to discount rates, we can write

$$\left(r + se^{-\rho t}\right) u(s,t) = 1 + \frac{\partial u(s,t)}{\partial t}$$

Defining $\hat{u}(t) := \partial u(s,t)/\partial s|_{s=0}$, differentiate with respect to s

$$r\hat{u}(t) + e^{-\rho t}u(t) = \frac{d\hat{u}}{dt}(t)$$

And then conjecture $\hat{u}(t) = e^{-\rho t}\hat{u}$, where $\hat{u} := \partial u(s,t)/\partial s|_{s,t=0}$, so

$$\begin{aligned} re^{-\rho t}\hat{u} + e^{-\rho t}u &= -\rho e^{-\rho t}\hat{u} \\ \Rightarrow \frac{\hat{u}}{u} &= -\frac{1}{r+\rho} \end{aligned}$$

for u the $t = 0$ value of the perpetual bond. Relabeling appropriately yields the result.

F.4.3 Proof of Proposition 5

Let $\bar{\mathcal{Z}}(t,s)$ denote the entry threshold for a project. It must satisfy

$$\mathcal{V}(0, \bar{\mathcal{Z}}(t,s), t, s) = 0$$

where the restriction $x = 0$ comes from the project being nascent. From the implicit function theorem, it is immediate that

$$\frac{\partial \bar{\mathcal{Z}}(t,s)}{\partial s} = -\left.\frac{\partial \mathcal{V}/\partial s}{\partial \mathcal{V}/\partial z}\right|_{(x,z)=(0,\bar{\mathcal{Z}}(t,s))} \tag{36}$$

We have already obtained the numerator, at least in steady state, when deriving the response of value to the shock. Towards the denominator, differentiate the time-dependent

HJB making use of the Envelope Theorem to obtain

$$r\frac{\partial \mathcal{V}}{\partial z}(x,z,s,t) = \mathcal{I}(x,z,s,t)\frac{\partial}{\partial x}\left(\frac{\partial \mathcal{V}}{\partial z}(x,z,s,t)\right)$$

$$\frac{\partial \mathcal{V}}{\partial z}(\bar{x},z,s,t) = p$$

and in the steady-state (i.e. $s = t = 0$), using lower-case letters to denote steady-state objects

$$r\frac{\partial v(x,z)}{\partial z} = i(x,z)\frac{\partial}{\partial x}\left(\frac{\partial v(x,z)}{\partial z}\right)$$

$$\left.\frac{\partial v(x,z)}{\partial z}\right|_{x=\bar{x}} = p$$

Once again, we have a first-order ODE in $\frac{\partial v(x,z)}{\partial z}$ with a given terminal condition. Here, recall that

$$\frac{\partial}{\partial x}e^{-rT(x,z)} = \frac{re^{rT(x,z)}}{i(x,z)}$$

so that

$$\frac{\partial v(x,z)}{\partial z} = pe^{-rT(x,z)} \tag{37}$$

satisfies the ODE and the terminal condition.

Returning to (36) and evaluating at a steady state, we have

$$\begin{aligned}
\frac{\partial \bar{z}}{\partial s} &:= \left.\frac{\partial \bar{\mathcal{Z}}(t,s)}{\partial s}\right|_{s,t=0} \\
&= -\left.\frac{\partial v(x,z)/\partial x}{\partial v(x,z)/\partial z}\right|_{(x,z)=(0,\bar{z})} = -\frac{\hat{v}(0,\bar{z})}{\partial v(0,\bar{z})/\partial z} \\
&= -\frac{e^{rT(0,\bar{z})}}{p}\hat{v}(0,\bar{z})
\end{aligned}$$

which is (19) after rearranging.

F.4.4 Proof of Proposition 6

We omit z from the derivations to lighten notation. The results hold pointwise along the distribution of z.

Fixed Capital Formation Fix any $z > \bar{z}$. After any shock, not necessarily first-order, the distribution evolves according to the time-dependent KFE. Integrating the KFE, we have

$$\frac{\partial}{\partial t} \int_0^{\bar{x}} g(x,t)dx = i(0,t)g(0,t) - i(\bar{x},t)g(\bar{x},t)$$

Now, integrate from $t = 0$ (the time the shock hits) to $t = \infty$, and so

$$\lim_{t\to\infty} \int_0^{\bar{x}} g(x,t)dx - \int_0^{\bar{x}} g(x,0)dx = \int_0^{\infty} \left[i(0,t)g(0,t) - i(\bar{x},t)g(\bar{x},t)\right] dt$$

where $g(x,0)$ is by definition the steady state distribution.

If the shock is temporary, the distribution reverts back to steady state, and so

$$\int_0^{\infty} \left[i(0,t)g(0,t) - i(\bar{x},t)g(\bar{x},t)\right] dt = 0$$

That is, the total outflow of projects must be equal to the total inflow of projects over time.

If the shock is small enough that z remains above $\bar{z}$ (or for any first order shock and $z > \bar{z}$), we have

$$\int_0^{\infty} \left[i(\bar{x},t)g(\bar{x},t) - f(z)\right] dt = 0$$

where we have used the initial boundary condition for the KFE for $z > \bar{z}$

$$i(0,t)g(0,t) = f(z)$$

But $f(z)$ is also the steady state outflow of firms, so this means that the area under the IRF of outflows (or capital formation if multiplied by z) is zero.

Accounting Investment The area "under" the curve is

$$\int_0^{\infty} \left[\int_0^{\bar{x}} i(x,t)g(x,t)dx - f(z)\bar{x}\right] dt$$

where $\int_0^{\bar{x}} i(x,t)g(x,t)dx$ is total accounting investment at time t and $f(z)\bar{x}$ is its steady state value. We want the area is zero.

Integrating the KFE, we have

$$\frac{\partial}{\partial t}\int_0^x g(y,t)dy = i(0,t)g(0,t) - i(x,t)g(x,t)$$

Now, integrate again from $x = 0$ to $\bar{x}$

$$\frac{\partial}{\partial t}\int_0^{\bar{x}}\int_0^x g(y,t)dydx = i(0,t)g(0,t)\bar{x} - \int_0^{\bar{x}} i(x,t)g(x,t)dx$$

Consider the left-hand side. Switching the order of integration, we have

$$\begin{aligned}\int_0^{\bar{x}}\int_0^x g(y,t)dydx &= \int_0^{\bar{x}}\int_y^{\bar{x}} g(y,t)dxdy \\ &= \int_0^{\bar{x}} (\bar{x}-y)g(y,t)dy\end{aligned}$$

And thus

$$\int_0^{\bar{x}} (\bar{x}-y)\frac{\partial}{\partial t}g(y,t)dy = i(0,t)g(0,t)\bar{x} - \int_0^{\bar{x}} i(x,t)g(x,t)dx$$

Integrating from $t = 0$ to $t \to \infty$, we obtain

$$\int_0^{\bar{x}} (\bar{x}-y)\left[\lim_{t\to\infty} g(y,t) - g(y,0)\right] dy = -\int_0^{\infty}\left[\int_0^{\bar{x}} i(x,t)g(x,t)dx - i(0,t)g(0,t)\bar{x}\right] dt$$

Last, if the shock is temporary and sufficiently small, we have $\lim_{t\to\infty} g(y,t) = g(y,0)$ and $i(0,t)g(0,t) = f(z)$, so

$$0 = \int_0^{\infty}\left[\int_0^{\bar{x}} i(x,t)g(x,t)dx - f(z)\bar{x}\right] dt$$

as desired.

G Convex Adjustment Cost Model

In this appendix, we describe the convex adjustment cost model we use to compare general equilibrium investment dynamics across models with and without time to build. See Figure 11 and the main text for an introductory discussion.

Setup The setup is largely the same as in Section 5. The firm still faces a distribution of projects $f(z)$, but now projects can be instantaneously built. This means that projects are started so long as $p(t)z(t) = \bar{x}$, and hence the (possibly time-varying) threshold for entry is $\bar{z}(t) = \bar{x}/p(t)$. As projects are completed, they are added to the aggregate capital stock $K(t)$. All projects are of the same size $\bar{x}$, and all heterogeneity is loaded on the output distribution $f(z)$.

Each period, a representative final goods producer uses the capital stock to produce and obtains revenues $A(t)K(t)$. It chooses investment in capital goods $I(t)$, and pays a quadratic adjustment cost if investment deviates from some level $\bar{I}$. The firm chooses $I(t)$ to maximize its discounted net present value of profits

$$\max \int_0^\infty e^{-r(t)t} \left[A(t)F(K(t)) - p(t)I(t) - p(t)\frac{\phi}{2}\left(I(t) - \bar{I}\right)^2 \right] dt$$

subject to the investment law of motion $I(t) = \dot{K}(t) + \delta K(t)$. The interest rate $r(t)$ is taken as given.

Solution The Hamiltonian is given by

$$\mathcal{H} = A(t)F(K(t)) - p(t)I(t) - p(t)\frac{\phi}{2}\left(I(t) - \bar{I}\right)^2 + \lambda(t)\left[I(t) - \delta K(t)\right] \tag{38}$$

where $\lambda(t)$ is the co-state (dynamic Lagrange multiplier) on the law of motion for the state $K(t)$.

Thet necessary conditions for optimality are given by

$$\begin{aligned} \frac{\partial \mathcal{H}}{\partial I(t)} &= 0 \\ r(t)\lambda(t) &= \frac{\partial \mathcal{H}}{\partial K(t)} + \dot{\lambda}(t) \end{aligned}$$

Differentiation of (38) yields

$$\lambda(t) = p(t) \left[1 + \phi \left(I(t) - \bar{I}\right)\right]$$
$$\dot{\lambda}(t) = (r(t) + \delta)\lambda(t) - MPK(t)$$

We can solve forward the co-state law of motion to obtain the co-state as

$$\lambda(t) = \int_0^\infty e^{-\int_0^s (r(t+u)+\delta)du} MPK(t+s)ds \tag{39}$$

which implies the following optimality condition for investment

$$p(t) \left[1 + \phi \left(I(t) - \bar{I}\right)\right] = \int_0^\infty e^{-\int_0^s (r(t+u)+\delta)du} MPK(t+s)ds \tag{40}$$

In the absence of adjustment costs ($\phi = 0$), the equation reflects the usual result that the price of capital in competitive markets reflects the discounted value of future marginal products, where depreciation is baked into the Jorgensonian user cost. This is effectively an equation capturing the demand for physical capital by the final goods producer.

By definition, fixed capital formation is given by the total amount of completed projects in any given instant

$$FCF(t) = \int_{\bar{z}(t)}^\infty z f(z)dz = \mathbb{E}\left[z | z \geq \bar{z}(t)\right] \tag{41}$$

where recall $\bar{z}(t) = \bar{x}/p(t)$. This equation represent the supply of physical capital in the economy.

Distributional Assumption Towards simplifying the model, we follow the assumption made in Section 6 and assume a Pareto distribution for project output

$$f(z) = \alpha \frac{z_{min}^\alpha}{z^{\alpha+1}} \tag{42}$$

Standard results for the truncated mean of a Pareto then allow us to express fixed capital formation, or capital supply, in closed form

$$\begin{aligned} FCF(t) &= \frac{\alpha}{\alpha - 1} \frac{z_{min}^\alpha}{\bar{z}(t)^{\alpha-1}} \\ &= \frac{\alpha}{\alpha - 1} \frac{z_{min}^\alpha}{\bar{x}^{\alpha-1}} p(t)^{\alpha-1} \end{aligned} \tag{43}$$

Equilibrium Equilibrium in the market for physical capital requires equating the supply of capital by capital producing firms and the demand for capital by the final goods producer

$$FCF(t) = I(t)$$

From (43), we have

$$p(t) = \left[\frac{\alpha-1}{\alpha}\frac{\overline{x}^{\alpha-1}}{z_{min}^{\alpha}}\right]^{\frac{1}{\alpha-1}} I(t)^{\frac{1}{\alpha-1}} \tag{44}$$

Finally, from (40), we can solve out the price of capital to obtain

$$I(t)^{\frac{1}{\alpha-1}}\left[1+\phi\left(I(t)-\bar{I}\right)\right] = \left[\frac{\alpha}{\alpha-1}\frac{z_{min}^{\alpha}}{\overline{x}^{\alpha-1}}\right]^{\frac{1}{\alpha-1}} \int_0^{\infty} e^{-\int_0^s (r(t+u)+\delta)du} MPK(t+s)ds \tag{45}$$

Steady State In our time to build models, the steady state capital stock and investment are independent of adjustment costs. To make the convex model comparable, we specify the $\bar{I} = \delta K$, i.e. the firm pays adjustment costs of investment if it deviates from the steady state level of investment. With a constant MPK in steady state, we have

$$I^{\frac{1}{\alpha-1}} = \left[\frac{\alpha}{\alpha-1}\frac{z_{min}^{\alpha}}{\overline{x}^{\alpha-1}}\right]^{\frac{1}{\alpha-1}} \frac{MPK}{r+\delta} \tag{46}$$

and the steady state price of capital is then just the annuity value of the steady state marginal product of capital

$$p = \frac{MPK}{r+\delta}$$

First-Order Dynamics In Figure 11, we plot the response of accounting investment and productive capital in this model to a AR(1) shock to discount rates, i.e. $r(t) = r + se^{-\rho t}$ for r the steady state interest rate. In the spirit of Proposition 7, we can derive the response of equilibrium objects to this first-order shock.

We first take first-order deviations of the co-state equation (39) relative to steady-state

$$\hat{\lambda}(t) = \int_0^{\infty} e^{-(r+\delta)s}\left(\widehat{MPK}(t+s) - MPK\frac{\hat{r}(t+u)}{r+\delta}\right)ds$$

We can write equation (45) as

$$I(t)^{\frac{1}{\alpha-1}}\left[1+\phi\left(I(t)-\delta K\right)\right] = \left[\frac{\alpha}{\alpha-1}\frac{z_{min}^{\alpha}}{\overline{x}^{\alpha-1}}\right]^{\frac{1}{\alpha-1}} \lambda(t)$$

so taking "hats" on both sides, we find

$$\left[\frac{1}{\alpha-1}\frac{1}{I}+\phi\right] I^{\frac{1}{\alpha-1}}\hat{I}(t)=\left[\frac{\alpha}{\alpha-1}\frac{z_{min}^{\alpha}}{\overline{x}^{\alpha-1}}\right]^{\frac{1}{\alpha-1}}\hat{\lambda}(t)$$

From steady state investment (46), we can write

$$\left[\frac{1}{\alpha-1}\frac{1}{I}+\phi\right]\hat{I}(t)=\frac{r+\delta}{MPK}\times\hat{\lambda}(t)$$

and hence

$$\frac{\hat{I}(t)}{I}=\frac{r+\delta}{\frac{1}{\alpha-1}+\phi I}\times\int_{0}^{\infty}e^{-(r+\delta)s}\left(\frac{\widehat{MPK}(t+s)}{MPK}-\frac{\hat{r}(t+s)}{r+\delta}\right)ds \tag{47}$$

Finally, capital price dynamics immediately follow from equation (44)

$$\frac{\hat{p}(t)}{p(t)}=\frac{\hat{I}(t)}{I(t)}$$

Panel (a) of Figure 11 plots "accounting" or balance sheet investment of the capital suppliers. We now derive the equation underlying the grey dot-dashed line in such figure. In this model, we can write accounting investment as

$$\begin{aligned}I^{acc}(t)&=\int_{\bar{z}(t)}^{\infty}\overline{x}f(z)dz\\&=\overline{x}P(z>\overline{z}(t))\end{aligned}$$

so under the Pareto assumption,

$$\begin{aligned}I^{acc}(t)&=\overline{x}\left(\frac{z_{min}}{\overline{z}(t)}\right)^{\alpha}=\frac{z_{min}^{\alpha}}{\overline{x}^{\alpha-1}}p(t)^{\alpha}\\&=\frac{z_{min}^{\alpha}}{\overline{x}^{\alpha-1}}p(t)^{\alpha}\end{aligned}$$

From equation (44), we obtain[51]

$$I^{acc}(t) = \overline{x}\left[\frac{\alpha-1}{\alpha}\frac{1}{z_{min}}\right]^{\frac{\alpha}{\alpha-1}} I(t)^{\frac{\alpha}{\alpha-1}}$$

Finally, taking "hats"

$$\frac{\widehat{I^{acc}}(t)}{I^{acc}} = \frac{\alpha}{\alpha-1}\frac{\hat{I}(t)}{I} \tag{48}$$

Towards the dynamics of the capital stock, recall that the capital stock reflects all underpreciated capital, which implies

$$\frac{\hat{K}(t)}{K} = \delta\int_0^t e^{\delta(s-t)}\frac{\hat{I}(s)}{I}ds \tag{49}$$

Equations (48) and (49) help us understand the gray dot-dashed lines in Figure 11. We assume a decreasing returns to scale production function with same curvature as that in the main text. The dynamics are as follows: higher discount rates push for lower capital prices as the future marginal products of capital (i.e. cash flows) are discounted more heavily. The lower price of capital acts to reduce capital supply, as capital-producing firms raise the threshold they use to start new projects. This drives the large impact drop in accounting investment in Panel (a). Indeed, up to first-order

$$\frac{\hat{\bar{z}}(t)}{\bar{z}(t)} = -\frac{\hat{p}(t)}{p(t)}$$

As the capital stock falls, marginal products of capital rise. At some point, the resulting increase in marginal products of capital relative to steady state – the $\widehat{MPK}/MPK$ terms – start to balance the effect of higher discount rates – the $\hat{r}/r$ terms decline over time and eventually reach zero given the AR(1) nature of the shock. As a result, the response of investment starts to become small. This drives the upward sloping part of the IRF in Panel (b). Accounting investment, fixed capital formation, and the productive capital stock eventually return to the steady state level.

[51]Moreover, steady state accounting investment is given by

$$\begin{aligned} I^{acc} &= \overline{x}\left(\frac{\alpha-1}{\alpha}\frac{I}{z_{min}}\right)^{\frac{\alpha}{\alpha-1}} \\ &= z_{min}\frac{MPK}{r+\delta} \end{aligned}$$

H Numerical Routines

This appendix collects all numerical routines used in solving and estimating the model. We borrow our methods mostly from Achdou et al (2020).

H.1 Solving the HJB

We start by solving the HJB in steady state, i.e. ignoring time-derivatives.

Fix some z drawn at inception. Define a grid for x: $x_0 = 0, \ldots, x_{J+1} = \overline{x}$, uniformly spaced with step Δx for simplicity. Discretize the value function as a vector of same dimension over this grid, i.e. $v_j = v(x_j)$. We will solve for the value function by iteration.

Let v^0 be an initial guess. Denote the discretized value function evaluated at point $x = x_j$ in the n^{th} iteration of the algorithm as $v_j^n = v^n(x_j)$. For each iteration n, perform the following steps.

1. Compute the discretized derivative of the value function for $j \leq J$ using forward-differences[52]

$$\left(v_j^n\right)' = \frac{v_{j+1}^n - v_j^n}{\Delta x}$$

   It is useful to express the forward-difference operator as a square matrix by stacking over j

$$(v^n)' = \frac{1}{\Delta x} \underbrace{\begin{bmatrix} -1 & 1 & & & & \\ & -1 & 1 & & & \\ & & -1 & 1 & & \\ & & & \ddots & \ddots & \\ & & & & -1 & 1 \\ & & & & & -1 \end{bmatrix}}_{\mathcal{D}_1} v^n + \begin{bmatrix} 0 \\ 0 \\ 0 \\ \vdots \\ 0 \\ pz/\Delta x \end{bmatrix}$$

   where we set $v_{J+1} = pz$. Note that for $j = J$, the firm jumps to the full completion value.

2. Solve the optimization problem for $j \leq J$

$$\max_{i \geq 0} i \left[(v_j^n)' - 1 \right] - \kappa(i)$$

   For convenient choices of $\kappa(i)$, this reduces to solving for the policy function along the grid point-wise. Obtain the policy functions $i_j^n = i^n(x_j)$.

[52] In many settings, other differencing schemes are recommended. In our case, we exploit the economics of the problem to simplify computations. Since firms cannot disinvest, $i \geq 0$, and $v(x)$ is increasing in x. The more sophisticated upwind scheme of Barles and Souganidis (1991) then reduces to a forward-differences operator.

3. Discretize $iv'(x)$ as

$$\underbrace{\begin{bmatrix} -i_0^n/\Delta x & i_0^n/\Delta x & & & & \\ & -i_1^n/\Delta x & i_1^n/\Delta x & & & \\ & & -i_2^n/\Delta x & i_2^n/\Delta x & & \\ & & & \ddots & \ddots & \\ & & & & -i_{J-1}^n/\Delta x & i_{J-1}^n/\Delta x \\ & & & & & -i_J^n/\Delta x \end{bmatrix}}_{A^n} v^n + \underbrace{\begin{bmatrix} 0 \\ 0 \\ 0 \\ \vdots \\ 0 \\ pzi_J^n/\Delta x \end{bmatrix}}_{a^n}$$

4. Given an iteration step Δ, update the value function using either an explicit or implicit scheme

   (a) Explicit method: start from

   $$rv_j^n + \frac{v_j^{n+1} - v_j^n}{\Delta} = -i_j^n - \kappa(i_j^n) + i_j^n (v_j^n)'$$

   Stacking over j,

   $$rv^n + \frac{1}{\Delta}(v^{n+1} - v^n) = -i^n - \kappa(i^n) + A^n v^n + a^n$$

   Rearranging,

   $$v^{n+1} = v^n + \Delta\left[-i^n - \kappa(i^n) + A^n v^n + a^n - rv^n\right]$$

   $$\text{or } v^{n+1} = \Delta\left[-i^n - \kappa(i^n) + a^n\right] + \left[(1 - r\Delta)I + \Delta A^n\right] v^n$$

   (b) Implicit method: start from

   $$rv_j^{n+1} + \frac{v_j^{n+1} - v_j^n}{\Delta} = -i_j^n - \kappa(i_j^n) + i_j^n (v_j^{n+1})'$$

   Stacking over j,

   $$rv^{n+1} + \frac{1}{\Delta}(v^{n+1} - v^n) = -i^n - \kappa(i^n) + A^n v^{n+1} + a^n$$

   $$\text{or } v^{n+1} = \left(\left(r + \frac{1}{\Delta}\right) I - A^n\right)^{-1} \left(\frac{1}{\Delta} v^n - i^n - \kappa(i^n) + a^n\right)$$

Terminate the algorithm when $\left\|v^n - v^{n-1}\right\| < \varepsilon$ for some tolerance ε.

The algorithm above yields a solution for the steady state value function $v(x)$. To study optimal policies following a shock, we must consider time-dependent terms. The HJB equation becomes

$$rv(x,t) = \max_{i\geq 0} -i - \kappa(i) + i\frac{\partial}{\partial x}v(x,t) + \frac{\partial}{\partial t}v(x,t)$$

We first create a grid for time $t_0 = 0, t_1, t_2, \ldots, t_K = T$. Then, we discretize the investment policy as $i_j^k = i(x_j, t_k)$, which is a $(J+1) \times K$ matrix representing the policy for stage $x = x_j$ and time period $t = t_k$. Discretization of the advection term $i\frac{\partial}{\partial x}v(x,t)$ is exactly as in the steady state case, but can be done point-wise for each grid point t_k.

For the time-dependent terms, we iterate backwards starting from the steady state. That is, for $t_K = T$, we set $v_j^K = v(x_j, t_K) = v(x_j, T) = v_{ss}(x_j)$. Since we are going back in time, we discretize the time-derivative using backward-differences

$$\frac{\partial}{\partial t}v(x_j, t_k) \approx \frac{v(x_j, t_k) - v(x_j, t_{k-1})}{(t_k - t_{k-1})}$$

For simplicity, we consider a equi-spaced grid for t with steps Δt.

H.2 Solving the KFE

Recall that in steady state, we have $g(x,z) = f(z)/i(x,z)$. From solving the steady state HJB, one thus immediately obtains the steady state distribution. We now turn to the time-dependent KFE

$$\frac{\partial g(x,t)}{\partial t} = -\frac{\partial}{\partial x}\left[i(x,t)g(x,t)\right]$$

First, we set up a grid for the distribution. We let $g^k = (g_1^k, \ldots, g_{J+1}^k)$ be the discretization of the distribution in the x-space (a column vector). The discretization of the time-dependent distribution is thus just $g(x_j, t_k)$, where we just collate the g^k along the column dimension.

To discretize the right-hand side term, we use the well-known result on the KFE operator being the adjoint of the differential operator of the HJB (Achdou et al 2021). That is, for any fixed t_k, we can discretize the RHS as

$$\left(A^k\right)^T g^k + \underbrace{\begin{bmatrix} f/\Delta x \\ 0 \\ 0 \\ \vdots \\ 0 \\ 0 \end{bmatrix}}_{\tilde{a}^k}$$

where the first row is equivalent to the boundary condition

$$g(0,t) = \frac{f}{i(0,t)}$$

which simply states that at $x = 0$, a flow f of firms enter the distribution.

To discretize the left-hand side term,we can use either an explicit scheme

$$\begin{aligned}\frac{g^{k+1} - g^k}{\Delta} &= \left(A^k\right)^T g^k + \tilde{a}^k \\ \Rightarrow g^{k+1} &= g^k + \Delta\left[\left(A^k\right)^T g^k + \tilde{a}^k\right]\end{aligned}$$

or an implicit scheme

$$\begin{aligned}\frac{g^{k+1} - g^k}{\Delta} &= \left(A^k\right)^T g^{k+1} + \tilde{a}^k \\ \Rightarrow \left[I - \Delta\left(A^k\right)^T\right] g^{k+1} &= g^k + \Delta\tilde{a}^k\end{aligned}$$

so that in both cases, we can write $g^{k+1} = Bg^k + b$, where B and b depend on the scheme used.