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

Adriano Fernandes
International Monetary Fund

Rodolfo Rigato *
European Central Bank

First Draft: November 2023
This Draft: October 2026

Abstract

We show that firms accelerate ongoing capital expenditure projects when credit tightens using project-level data from India and quasi-experimental variation in credit supply. We provide a theory of endogenous time to build that both rationalizes these findings and qualifies the standard capital budgeting principle that tighter financing necessarily lowers investment. Tighter credit incentivizes firms to accelerate near-completion projects for faster cash flows while delaying early-stage projects, a gradient borne out in the data. A quantitative model disciplined by our estimates shows that endogenous time to build generates amplification and state-dependence along the distribution of project stages. We derive implications for monetary policy and investment tax credits.

Keywords: Time to Build, Investment, Business Fluctuations
JEL classification codes: D25, E22, E32, G31, H25.

*Fernandes: afernandes@imf.org. Rigato: rodolfo.dinis_rigato@ecb.europa.eu.

We are indebted to our advisors Gabriel Chodorow-Reich, Samuel Hanson, Jeremy Stein, and Ludwig Straub for their invaluable support. We are grateful to Veronica DeFalco, Namrata Narain, Louise Paul-Delvaux and Adi Sunderam for continuous encouragement and suggestions. We would also like to thank our discussants, Rui Castro, Nico Crouzet, Bob McDonald, Matthias Meier, and Mindy Xiaolan; as well as Andy Abel, Laura Alfaro, Pol Antràs, Patrick Bajari, Robert Barro, Adrien Bilal, Michael Blank, Emily Breza, John Campbell, Damien Capelle, Pedro Degiovanni, Giovanni Dell'Ariccia, Jan Eberly, Martin Eichenbaum, Ester Faia, Xavier Gabaix, Shresth Garg, Edward Glaeser, Pierre-Olivier Gourinchas, Christian Hellwig, Johannes Hörner, Joe Kaboski, Myrto Kalouptsidi, Andrea Lanteri, Greg Lewis, Karen Lewis, Debbie Lucas, Greg Mankiw, Rui Mano, Ariel Pakes, Dev Patel, Sole Martinez Peria, Elie Tamer, Audrey Tiew, Tom Winberry, Christian Wolf, Lingxuan Wu, and numerous seminar participants at Boston College, Central Bank of Brazil, CREI, Fed Board, ECB, Harvard, IDB, IMF, JHU SAIS, Lubramacro, NBER Summer Institute, SED Annual Meetings, SF Fed, STEG-CEPR, Toulouse, UT Austin McCombs, Wharton, and WFA Conference. We thank Samuel Efraim, Harold Klapper, Irene Moreno, and Johnathan Sun for outstanding research assistance; and João Menegotto for stellar research assistance. Adriano Fernandes recognizes financial support during his graduate studies 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 International Monetary Fund or the European Central Bank, their Executive Boards, or their Management. Any errors are our own.

1 Introduction

Capital takes time to build. At least since Pigou (1920), Aftalion (1927), and Kalecki (1935), gestation lags have been central to how economists think about investment, its response to the cost of capital (Jorgenson, 1963; Hall, 1977), its role in propagating business cycles (Kydland and Prescott, 1982; Lucca, 2007), and its effect on the power of stabilization policies (Taylor, 1982; Ramey, 2020). Throughout, almost all investment models have treated the time it takes to build physical capital as a fixed technological constraint beyond the firm's control. Firms may choose how much to invest, but not how fast to complete what they have started. Shocks and policy can only reach investment through decisions to build new capital.

This paper challenges this assumption.[1] Using project-level data on large capital expenditure projects, we show that time to build is *endogenous*: when financing conditions tighten, firms start fewer new projects but complete ongoing ones faster, the more so the closer the projects are to completion. We rationalize these findings with an investment model in which firms choose how fast to build. In the model, a project yields cash flows once it is completed, so finishing faster brings those cash flows forward. When credit tightens, firms value cash today more relative to cash tomorrow. Projects that have barely started, or not started at all, become less attractive, because their payoffs lie far in the future. Bringing forward the payoff of a project close to completion, by contrast, becomes more valuable, so firms speed it up. The standard capital budgeting logic that tighter financing conditions lower investment therefore holds for new projects, but reverses for projects close to completion. Aggregating across projects, the response of investment to financial conditions depends on how mature the economy's pipeline of projects is. Endogenous time to build can (i) attenuate or even reverse standard transmission channels of interest rate policy; and (ii) reshuffle the intertemporal effects of fiscal policy.

Our research question requires us to track investment projects over time. This information is not readily available in accounting records or most data sources used in the extant investment literature. We overcome this measurement challenge by obtaining project-level data on large capital expenditure projects in India. By combining event diaries with extensive data cleaning, we construct a panel tracking projects over time. Another advantage of our data is its cross-sectional richness. In contrast with prior work, we observe public and private projects (e.g., private factories versus public roads) across a wide cross-section of end-user industries (e.g., pharmaceutical versus cement manufacturing plants). We leverage this richness to provide novel facts on time to build, offering new portable moments to investment models while providing context to our identification strategy. We find a median time to completion of 2.5 years, but the distribution has a long right tail, with some projects

[1] Curiously, Brainard (1977) had objected to Hall's (1977) treatment of investment lags as mechanical, while Baily (1982) had noted that firms could pay to "fast-track" construction in his discussion of Taylor (1982).

taking over 15 years. Consistent with a role for technological primitives, (i) a 1% increase in project size is associated with a 0.18% increase in time to build, and (ii) industry accounts for 30% of the cross-sectional variation in completion times.

We then provide evidence for an endogenous margin of time to build. Suppose a firm experiences a shock to its economic environment that is plausibly orthogonal to the daily affairs of project construction. Under a null of exogenous time to build, the firm-level shock should not affect project completion. We provide evidence against this view by studying a credit supply shock. For identification, we leverage large-scale bank audits by the Reserve Bank of India (RBI), which forced commercial banks to considerably increase loan loss provisions for non-performing assets. The audit led to a contraction in credit supply (Chopra et al., 2021): controlling for demand, banks that had to put up more provisions reduced lending relative to less exposed banks. We then leverage this audit shock as a quasi-experiment to study what firms do to project timelines when credit dries up.

To trace the shock from banks to capex projects, we construct a firm-level exposure variable: firms are more exposed to the bank shock if they borrowed more from exposed banks prior to the shock. Our main finding is that completion hazards increase for projects of more exposed relative to less exposed firms, consistent with *acceleration of ongoing projects*.[2] On the extensive margin, we also find that more exposed firms *start fewer new projects* following the shock, consistent with standard investment theory. In our preferred specifications, a one-standard-deviation increase in firm-level exposure raises the completion hazard by 6% and lowers project starts by 7.6%.

The acceleration result is at odds with exogenous time to build. Though surprising at first, the result is consistent with a model of investment with endogenous time to build. Our theory further provides a sharper test: following the shock, the acceleration should be concentrated among mature projects. This prediction is borne out in the data. Using survival-model estimates of each project's proximity to completion, we find that the effect of the shock on completion hazards is larger for projects close to completion.

Our empirical strategy relies on the bank-level shocks being exogenous relative to project and firm-level unobservables, conditional on controls. We cannot fully address threats related to unobserved project quality, a data limitation that qualifies our findings. That said, five lines of evidence support our interpretation. First, the audit was carried out by the central bank, simultaneously across banks, and more- and less-exposed banks and borrowers show no significant differences before the audit (Chopra et al., 2021). Second, we include a detailed set of fixed effects and firm-level controls to minimize confounding concerns. Third,

[2]Our interpretation mirrors a voluminous literature showing that less generous unemployment benefits accelerate re-employment through higher job finding hazards. These works widely interpret these dynamics as evidence of (endogenously) higher search effort – despite search effort being unobservable (Chetty, 2008; Rothstein, 2011; Schmieder and von Wachter, 2016). We apply the same logic: systematic changes in project completion hazards following a credit shock provide evidence of endogenous construction timelines.

our credit supply regressions absorb credit demand with borrower-by-year fixed effects, and estimates with only borrower and year effects are similar, ruling out correlated demand shocks and banks differentially cutting credit based on unobservable, time-varying firm characteristics. Fourth, panel and cross-sectional balance tests show that firm-level shock exposures are largely uncorrelated with firm-level observables. Fifth, bank shocks are not driven by lending to a small number of bad firms, reducing concerns of shock endogeneity. We further provide corroborating anecdotal evidence and a battery of tests to bolster our interpretation relative to competing stories, including differential censoring and ex post changes in scope.

To rationalize our empirical findings, we develop an investment model with endogenous time to build. Our setup is minimalist and portable, lending itself to further applications and extensions. In our basic setup, a firm invests in a collection of heterogeneous, irreversible projects. Each project takes time to build, and only generates cash flows once fully completed. Firms benefit from completing projects faster because they discount future cash flows. They 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 in ongoing projects and whether to start new projects. 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. Our framework nests exogenous time to build as a limiting case.

We first explain why our empirical results follow from endogenous time to build. To do so, we take an analytical approach. We map the banking shock (data) to an increase in firm discount rates (model). The mapping parsimoniously captures that firms value cash more today relative to tomorrow when credit dries up. It has deep roots in macro-finance, as discount rates are a central channel through which financial conditions transmit to real activity (Gilchrist and Zakrajšek, 2012) and are explicitly used by firms in project-level capital budgeting (Graham, 2022), with quantitatively large effects on firm investment (Gormsen and Huber, 2025). The mapping also follows models where a credit supply disruption increases the cost of bank credit (Crouzet, 2018).[3]

Formally, the change in a project's value following the discount rate shock is given by a duration formula, as for a bond whose maturity is the project's completion date. The project value falls by more for projects far from completion. The investment response balances this duration effect against an acceleration effect, the value of shifting cash flows forward by building faster. We prove that the duration effect vanishes as the project approaches completion, so acceleration dominates and mature projects are sped up; for early-stage projects, the

[3]The mapping is also consistent with extensive evidence that the cost of credit rises in episodes of credit supply disruptions; see Appendix H.7.

duration effect can dominate, and it is optimal to slow them down.

Our empirical and theoretical results qualify the standard intuition that capital expenditures necessarily decrease following adverse shocks to firm financial health. Capital formation may very well *rise*, as the optimal response to higher discount rates for sufficiently mature projects is to invest more, rather than less. Further, this has implications for interpreting extant evidence on the effect of shocks on capex: aggregation across projects averages the positive effects on projects close to completion and the negative effects on early-stage or unbegun projects, so the distribution of projects over their implementation stages matters for firm- or aggregate-level investment. Consistent with this channel, Selgrad and Siani (2025) find that contractionary monetary policy shocks lead firms to revise upward their investment plans in already ongoing projects, but to sharply revise downward plans in their unbegun projects. Our model provides a lens to interpret these results.

To quantify the implications of endogenous time to build for investment dynamics and policy, we enrich the model and take it to the micro-data. We introduce heterogeneity in project returns, sizes, and construction technologies; and estimate structural parameters to match our causal estimates by computing model counterparts to our intensive and extensive margin regression coefficients. We feed discount rate shocks into our quantitative model and compare investment fluctuations under endogenous versus exogenous time to build, both in partial and general equilibrium.

First, endogenous time to build makes investment dynamics more responsive to shocks, increasing investment volatility by 30% relative to the exogenous benchmark. 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 time to build. Hence, endogenous time to build amplifies investment responses relative to the exogenous benchmark, thereby requiring smaller primitive shocks to generate sizable investment fluctuations.

Second, endogenous time to build generates *state-dependence*: mature and early-stage projects differ in their response to discount rate shocks, so the initial distribution matters. With exogenous time to build, our model behaves to first order like the investment adjustment costs widely used in macro models (Christiano et al., 2005; Lucca, 2007). With endogenous time to build, the implied adjustment cost depends on the stage distribution of projects, which itself depends on past shocks and policy. We thus propose a novel source of state-dependence in the effects of monetary policy: interest rate hikes (cuts) will be less contractionary (expansionary) on impact when the distribution is skewed towards more mature projects. This state-dependence is quantitatively large. When the economy's project pipeline is skewed towards mature projects, the impact effect of an interest rate increase on aggregate investment can flip sign, from contractionary to expansionary.

Third, endogenous time to build provides lessons for investment stimulus policy. With

exogenous time to build, temporary tax credits increase investment only through entry. Endogenous time to build introduces an intensive-margin channel: investment in existing projects spikes, as firms rush to complete projects before subsidies expire. From a stabilization perspective, tax credits reallocate investment demand intertemporally, thereby providing flexibility to discretionary fiscal policy. Beyond stabilization, investment subsidies have been increasingly deployed to accelerate structural changes (e.g., the renewable tax credits from the Inflation Reduction Act). Our model clarifies how these policies trade off speed gains with acceleration costs.

Together, these results show that the investment channel of monetary and fiscal policy is richer than standard models imply. The effectiveness of interest rate changes depends not just on their magnitude, but on the composition of the economy's project pipeline at the time of the policy change – a source of state-dependence absent from models with convex adjustment costs. Endogenous time to build provides an empirically grounded microfoundation for this state-dependence, and can be readily embedded in models of firm, industry, or aggregate dynamics. The relevance of these mechanisms is visible in the current build-out of artificial intelligence infrastructure. Top policymakers and firm executives have been stressing time to build considerations.[4] Our results thus provide a lens to understand how fast new capacity can come online and how future monetary and fiscal policy will shape the AI buildout. We conclude by outlining avenues for new applications across a range of settings.

Relation to Literature. Our paper contributes to three literatures. First, a long literature asks which frictions govern how investment responds to shocks, from convex adjustment costs (Lucas, 1967; Hayashi, 1982; Eberly et al., 2012) to fixed costs and irreversibility (Abel and Eberly, 1994; Doms and Dunne, 1998; Cooper and Haltiwanger, 2006); see Caballero (1999) for a survey. In spirit, our paper parallels Lanteri (2018), who makes irreversibility, usually modeled as a constant technological friction, an endogenous outcome that microfounds state-dependent adjustment costs. We do the same for time to build: exogenous time to build acts as an investment adjustment cost (Lucca, 2007), and making it endogenous renders that cost state-dependent. Methodologically, we follow a tradition in macroeconomics and finance of bringing novel microdata to estimate micro facts, elasticities, and structural parameters and to discipline industry- or macro-models. We are closest to work that looks beyond accounting measures of investment expenditures to observe the adjustment of capital directly. Ramey and Shapiro (2001) use equipment-level data, and Bertolotti and Lanteri (2026) assemble data on aircraft models to study capital replacement and embodied innova-

[4] Fed Chairman Warsh has described the AI buildout as "a race between supply and demand" in which supply "is going to expand, but it'll take longer" (Warsh, 2026), while ECB President Lagarde has argued that "Europe needs to build at a different pace altogether" (Lagarde, 2026). Microsoft's CFO recently stressed the emphasis on speeding up delivery of projects, "landing this CapEx as quickly as we can, and converting it to revenue as quickly as we can" (Microsoft Corporation, 2026).

tion. Also close to us is Whited (2006), who estimates hazard models for investment spikes in Compustat and benchmarks the findings relative to a structural investment model. Like these papers, we provide moments and elasticities that can discipline models. We make three contributions: (i) new measurement and evidence on how firms adjust capex timelines with wide cross-sectional coverage using project-level data, (ii) a theoretical framework with sharp comparative statics that qualify conventional capital-budgeting rules, and (iii) a structural model with rich micro heterogeneity, general equilibrium dynamics, and implications for policy evaluation.

Second, our paper contributes to a literature that treats time to build as a driver of real dynamics.[5] In macroeconomics, Meier (2020) argues that time-series variation in delivery times due to supply chain disruptions yields GDP and TFP losses.[6] Kalouptsidi (2014), Greenwood and Hanson (2014), Oh and Yoon (2020), Gomtsyan (2024) and Oh et al. (2024) study how exogenous time to build drives supply elasticities, return predictability, and boom-bust dynamics in shipping or housing. Lewis and Bajari (2011, 2014) show that contractors accelerate road construction to avoid procurement penalties. Relative to such work, we (i) document quasi-experimental evidence of endogenous time to build in a wide cross-section of industries; and (ii) provide a portable, microfounded model of optimal time to build that extends the neoclassical theory of investment (Hall and Jorgenson, 1967; Abel and Eberly, 1994), providing new lessons for capital budgeting.[7]

Finally, our work relates to how investment frictions shape aggregate dynamics and the effects of stabilization policy. In models of monetary transmission, investment adjustment costs help match the gradual response of investment to policy shocks (Christiano et al., 2005; Casares, 2006; Edge, 2007). On the fiscal side, implementation lags weaken the short-run multipliers of public investment (Leeper et al., 2010; Ramey, 2020), and the response of private investment to tax incentives depends on adjustment costs and on the timing of the incentive (Summers, 1981; Cummins et al., 1994; House and Shapiro, 2008; Zwick and Ma-

[5] We also contribute to the measurement of time to build. Early literature relied on small samples or aggregated data from the US (Mayer, 1960; Montgomery, 1995; Brooks, 2000). Millar et al. (2016) measure time to plan for retail stores, offices, hotels and warehouses in the US. Glancy et al. (2025) uses the same data source to measure time to build and study the role of abandonment options for commercial real estate supply elasticities. We contribute by relying on microdata from multiple industries, sponsors, and project sizes in India, the world's third-largest construction market, with $1.7 trillion in anticipated investment outlays by 2030 (Invest India). Comprehensive project-level measurement like ours is non-existent for developing economies, to the best of our knowledge.

[6] Time as a factor of production dates back to Böhm-Bawerk (1891), and has been revisited by Antras (2023); Liu and Tsyvinski (2023) and Schaal and Taschereau-Dumouchel (2025) study dynamic input-output networks with exogenous delivery lags.

[7] We incorporate irreversibilities, but abstract from fixed costs. Majd and Pindyck (1987) introduce exogenous time to build in real options (McDonald and Siegel, 1986), showing it can strengthen "wait-and-see" effects (see also Aïd et al. (2015)). While our mechanism bears a resemblance to real options models, it differs in that it requires neither fixed costs nor risk. We further innovate by (i) analytically proving how changes in discount rates and in future cash flows shape optimal project value and investment; and (ii) working with heterogeneous projects, thus studying aggregation and distributional state-dependence.

hon, 2017). In these models, time to build and adjustment costs are fixed technological features. With fixed adjustment costs, aggregate investment responds more to shocks when more firms are close to adjusting (Caballero and Engel, 1999; Winberry, 2021).[8] There the state changes the size of the response; with endogenous time to build it can change its sign on impact. We thus identify a new source of state-dependence in aggregate investment, arising from the distribution of project stages, and introduce endogenous time to build as a channel with lessons for both monetary and fiscal policy.

The rest of the paper is organized as follows. Section 2 describes the data and measurement exercises. Section 3 presents the econometric evidence. Section 4 introduces our model and analytical results. Section 5 explains our calibration. Section 6 studies investment dynamics and policy implications. Section 7 concludes. The Online Appendix elaborates on data construction and provides further empirical and theoretical results.

2 Data and Measurement

We obtain data on large capital expenditure projects from CapexDx (CMIE, 2023). The data provider defines a project as a construction enterprise where "there is reasonable clarity on the capacity being created". The data includes project name, sponsor, cost, geographical location, and end-user industry from the mid-1990s onward. Examples of projects include steel plants, irrigation canals, airport terminals, office and residential buildings.[9] Crucially, CapexDx "is a monitoring machinery that records the various events that a project goes through in its life-cycle." This event history allows us to observe when projects were announced, implemented, finished (e.g., completed or abandoned), or lost-to-follow-up. We elaborate on project definition rules and data collection in Appendix A.1.

Table 1 describes our micro-data on project lifecycles. First, this fills a long-standing measurement gap. Second, it provides context to our quasi-experimental estimates. Panel A presents summary statistics on the cross-section of projects. We observe over 100,000 projects sponsored by over 30,000 entities across a broad range of end-user industries and geographies.[10] The median (average) project is sponsored by a firm with 11 (113) projects. Projects

[8] Existing work also locates the state-dependence of policy effects in firms' financial positions (Ottonello and Winberry, 2020; Cloyne et al., 2023; Jeenas, 2025), in the phase of the business cycle (Auerbach and Gorodnichenko, 2012), or in distributions shaped by past decisions, such as mortgage rate gaps (Berger et al., 2021; Eichenbaum et al., 2022) and durable stocks (McKay and Wieland, 2021). Ours arises from the stage composition of the project pipeline, which likewise depends on past shocks and policy.

[9] The System of National Accounts divides "fixed assets" (the K of macroeconomic models) into structures (what our data covers), equipment, and, more recently, intellectual property. We do not observe replacements or maintenance of existing equipment; nor do we observe capacity or capacity utilization.

[10] Sponsors include mostly Indian private firms and state-owned enterprises. Not all projects in our data are sponsored by firms, but we will use the terms sponsor and firm interchangeably. Note also that though public sponsors account for almost 40% of the projects, most such sponsors are (partially or fully) state-owned enterprises, not government agencies.

in our data are sizable enterprises: the median project costs 503 million rupees, while projects above the 90th percentile cost over 10 billion rupees.[11] Importantly, projects in our sample are either new units (81%), large expansions (15%) or modernizations (4%) of existing units. They are *not* acquisitions or other forms of secondary market transactions.

Panel B provides an overview of projects' lifecycles. To do so, we construct project histories from event-level data. Multiple difficulties arise in reconstructing a project lifecycle from event-level data, but we take several steps to ensure that our project histories are sensible, and we find no evidence that backfilling or censoring drives our results (Appendix E.1). We elaborate on our extensive data cleaning, including multiple examples, in Appendix A.2.

All projects start when they are announced. The median project spends 3 quarters in this phase. Most projects (70%) move to implementation. The median project spends 4.5 quarters under implementation. At any point during 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). Abandoned or stalled projects are rare in our data. As our data ends in March 2022, any projects announced or under implementation at that point are treated as right-censored observations.[12]

Thirty percent of projects are "lost to follow-up". CapexDx 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." Reassuringly, this 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 or the project sponsor. The combination of intense monitoring and transparent censoring rules bolsters data reliability. In any case, our econometric model explicitly accounts for censoring.[13]

Novel Facts on Time to Build We measure a project's time to build from announcement to completion. To do so, 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. Figure 1 plots the distribution of time to build. The median project takes 2.5 years to be completed.

[11] The median (average) project costs 7 million (86 million) USD, and the right tail features projects over 140 million USD. Given the large nominal depreciation of the rupee since the global financial crisis, these numbers likely provide a conservative estimate of project costs in USD.

[12] Appendix Table A2 tabulates the frequency of transitions across different stages, excluding non-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 A3 instead estimates quarterly transition probabilities between stages including transitions to itself, and shows that intermediate stages are persistent.

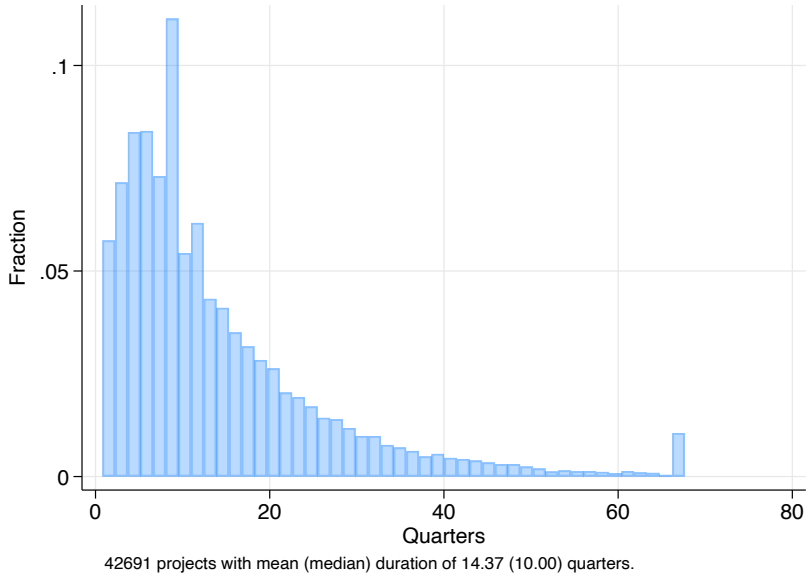
[13] That is, projects that are abandoned or lost to follow-up contribute to the completion likelihood through their survival up to the last observed period, without requiring assumptions about their eventual completion.

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					
Share: Section	0.19	0.11	0.02	0.07	0.23	0.31	0.31
Share: Industry	0.04	0.04	0.00	0.01	0.02	0.04	0.12
Share: State	0.06	0.04	0.02	0.03	0.05	0.09	0.14
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. Shares 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.

Figure 1: Distribution of Time to Build Across Completed Projects



Notes: the figure plots the histogram of time to build 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 (its announcement date) to its eventual completion. We winsorize the distribution at the 1st and 99th percentiles for visualization (but report the mean and median of the raw distribution).

The distribution is right-skewed, with a long right tail of projects taking over 15 years.

The aggregate distribution hides considerable cross-sectional heterogeneity. Figure A1 breaks down the distribution at the sector level. Transportation and construction feature the longest time to build, whereas retail and information technology feature the shortest. The figure also splits the distribution based on public or private ownership. Public sponsors build projects with longer time to build. This is partly driven by industrial composition: the public sector frequently sponsors heavy infrastructure projects. For example, Indian Railways is the largest public enterprise in India, and railways have notoriously long time to build. Last, Appendix Figure A2 plots the relationship between time to build and project size, as proxied by cost.[14] The elasticity between size and gestation lags appears robustly constant: a 1% increase in costs is associated with a 0.15–0.18% increase in completion times. Appendix Table A4 collects estimates from the regression specifications. Table A5 decomposes the cross-sectional variation with fixed effects. Industry dummies capture about 30% of the variation. Geography and ownership add little additional explanatory power, consistent with a strong role for technological primitives captured by the project's end-user industry.[15]

[14]Monetary cost provides a common metric to compare projects across different "size" definitions. That is, while it is feasible to compare two residential buildings based on how many floors are being built, it is not as straightforward to compare railways to irrigation canals.

[15]Industry may fail to proxy exclusively for technological factors. When we include sponsor and start-year fixed effects, we account for 58% of the variation in time to build. The sizable within-firm variation may re-

The rich cross-section of projects, industries, and sponsors documented above provides both the context and the variation needed to test whether time to build responds to economic conditions. We turn to the main empirical question: do firms have some control over how fast they build projects?

3 Testing for Endogenous Time to Build

Testing whether firms can change how fast they complete 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 were purely exogenous. At project onset, firms would choose an idealized builder, make payments upfront, and once the contract was signed, could not make changes to the construction process. In turn, the builder would consistently work on the project. In the absence of any unforeseen disturbances, time to build would be perfectly known. This thought experiment describes how time to build is treated in standard investment models. Whether the firm experienced distress — weak sales, cash shortages, or cuts in credit lines – would have no bearing on project completion. 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 completion hazards.

We contrast this null hypothesis against an alternative scenario where firms do have some control over how quickly to finish projects. If times were exceptionally good (bad), firms could speed up (slow down) completion of an ongoing manufacturing plant. In the absence of project-specific disturbances that are correlated with the firm-level disturbance, such episodes would offer encouraging evidence. Firms would have 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

flect the unpredictability of construction projects (Bajari and Tadelis, 2001) or active management by sponsors of project timelines. A stricter approach would obtain characteristics from engineering blueprints or construction records, which we do not observe. Further, increasing the dimensionality of our predictor set improves explanatory power, but at the cost of interpretability.

thus be uninformative about whether the firm could alter time to build.

We circumvent this identification challenge by considering a setting where firms face a contraction in credit supply originated in the banking sector. We trace the shock from banks to firms and examine what happens to firm projects. We leverage variation from firms being differentially exposed to troubled banks; and we will provide evidence that more and less exposed firms were observationally similar before banks got in trouble. If remaining reasons for potential project disruptions are just noise, comparing what happens to projects of more versus less exposed firms tests whether a firm-level shock unrelated to the natural progression of a project led to changes in time to build.

To implement our empirical strategy, we link projects in CapexDx to their sponsoring firms on ProwessDx. The latter data is sourced from regulatory filings and is *not* restricted to publicly traded companies (Alfaro and Chari, 2014). We collect firm- and bank-level financial statements from ProwessDx. We then hand-collect loan-level data from the Indian Ministry of Corporate Affairs (MCA) website. The data approximates a credit registry for secured loans[16], 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 names. We do not observe interest rates, contractual maturity, covenants, or any other contractual details on the loan.

Last, we link borrowers in the loan-level data to CMIE datasets by using the unique Corporate Identification Number (CIN) of a firm. These identifiers are assigned by the Indian government and present in both datasets, 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 ProwessDx company, thereby identifying the majority of entities lending to project sponsors.

3.1 Setting

Our quasi-experimental setting is the Indian Asset Quality Review, a widespread audit of Indian banks. We summarize its main features, closely following Chopra et al. (2021).

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 mar-

[16]Most bank lending in India is secured. Failing to register a secured loan effectively nullifies the benefits of secured lending, thereby incentivizing creditors to register their charges. 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.

kets. 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.[17] 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.[18] The audits uncovered sizable under-provisioning across systemic banks and thus had the potential to reduce bank lending. Importantly, the RBI did not structure capital infusions or backstops.

We first confirm that the audits led to a supply-led reduction in credit with borrower (Khwaja and Mian, 2008) or borrower-by-time fixed effects. The fixed effects estimator identifies the effect of the shock on equilibrium loan amounts by comparing more and less exposed banks, holding credit demand fixed. We include lender fixed effects to ensure our estimator is not capturing time-invariant heterogeneity across banks. Rather, the estimator asks whether banks more exposed to the shock differentially cut loans to a particular borrower in a given year, relative to their usual proclivity to extend loans.

We construct a lender-borrower-year panel from the MCA data. First, we isolate all outstanding loans between 2013 and 2015 (pre-audit period) 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 amount. If the pair had no outstanding loans in a given year, we assign it 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 audit, scaled by total bank assets (Chopra et al., 2021)

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

The measure captures that banks that must post more provisions experience larger reduc-

[17] Loan loss provisions reduce bank net income, thereby generating incentives to under-report. 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. (2021) document sizable increases in restructurings, but no increases in non-performing assets during the forbearance period.

[18] The reporting requirement was binding either if the additions to provisions exceeded 15% of the 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. (2021) document 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.

tions in lending capital, which should translate into fewer loans. We normalize by bank assets to avoid spuriously comparing large and small banks. We then regress total loans outstanding between lender b and borrower f in year t on the bank shock using a Poisson pseudo-maximum likelihood (PPML) estimator

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

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} collects 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.[19]

Appendix Table A6 presents the results. In our preferred specification with lender and borrower-by-year fixed effects, we find that a one-standard-deviation increase in bank exposure reduces outstanding credit in the average pair by 5.1%. The firm-year fixed effects control for time-varying demand confounders, so the coefficient represents a relative reduction in credit supply among more versus less affected banks. We obtain the same results when we restrict the sample to borrowers for whom we observe projects in the CapexDx dataset, and whose projects are included in our project-level regressions.[20]

3.2 From Banks to Projects: Estimation and Identification

Having established that the audit shock has bite, we map the banking shock to firms and their projects. If a firm borrows more from troubled banks, it potentially faces starker credit disruptions. We thus model exposure for firm f by aggregating bank-level exposures to the audits, weighting each bank b by its lagged share of total borrowing

$$\text{Firm Exposure}_{ft} = \sum_{b \in B_{f,t-1}} \left(\frac{\text{Loans}_{b,f,t-1}}{\sum_{b \in B_{f,t-1}} \text{Loans}_{b,f,t-1}} \right) \text{Bank Shock}_{bt}$$

[19] 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. As for the choice of estimator, we use PPML since the dependent variable is non-negative, heavily skewed, but contains a large mass of zeros (Silva and Tenreyro, 2006; Cohn et al., 2022). Chen and Roth (2024) show this is more appropriate than $\log(1+Y)$ specifications since one can recover arbitrarily sized effects in those regressions by appropriately rescaling the outcome Y.

[20] Appendix Table A7 reports results from a linear model in first-differences. The semi-elasticities obtained by dividing the coefficients by the mean of the outcome are similar in magnitude to those estimated in Table A6. This specification suggests a greater role for firm-time effects, but the control set included in the regressions with only borrower and year effects approximates well the estimates from the more saturated specifications.

To quantify the effect of shock exposure on project completion hazards, we estimate a discrete-time duration model (Van Den Berg, 2001)

$$Completed_{pft} = \exp\left(\beta \times \text{Firm Exposure}_{ft} + X_{pft}\Gamma\right) v_{pft} \tag{1}$$

where p refers to a project, f refers to the project's sponsor (or firm), and t refers to calendar time. Intuitively, conditional on surviving up to period t, completion arrives at a rate that depends on shock exposure, controls X_{pft} and unobservables v_{pft}.[21] The duration model explicitly accounts for right-censoring of ongoing and lost-to-follow-up projects. We cluster standard errors at the firm level (Abadie et al., 2022) and estimate the model with PPML (Gourieroux et al., 1984; Correia et al., 2020).

Proper identification of β requires that bank shocks are (conditionally) uncorrelated with (loan-weighted) project-level unobservables (Chodorow-Reich, 2014).[22] Concerns would arise 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. Several features of the shock support this assumption. RBI officials, not the banks, assessed each bank's non-performing assets, simultaneously across banks. Banks and borrowers with high and low exposure show no statistically significant differences in pre-audit characteristics, and future exposure does not significantly predict pre-audit lending (Chopra et al., 2021). We address potential confounders in turn, controlling for them where possible.

First, some firms may systematically perform poorly, build projects slowly and build relationships with unhealthy banks that are more likely to underreport non-performing assets. Firm fixed effects address this time-invariant sorting concern. Second, projects in our sample are immobile, which makes them susceptible to regional shocks (e.g., political instability, changes in construction regulations, variation in economic slack or construction labor supply). These factors threaten our design if banks cut credit by internalizing them. We include state-by-year fixed effects to absorb these regional confounders. Third, aggregate or industry-level shocks (e.g., recessions, supply chain disruptions) may delay project completion regardless of firm actions; and may lead bank loans to perform poorly and induce banks to hide losses. We include sector-by-year fixed effects to mitigate these concerns. Fourth, we include a cubic spline on project age, which removes baseline hazards and allows 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.,

[21] We organize the data into a quarterly panel of projects. 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 it or because the project is ongoing, $Completed_{pft}$ takes value zero. Restricting our attention to fully completed projects would run into selection. Duration models explicitly account for right-censoring and therefore do not suffer from such issues.

[22] Identification leverages exogenous variation in the shocks. Correlation of the shares $w_{b,f,t-1}$ with unobservables or the outcome does not invalidate the design (Borusyak et al., 2021).

Table 2: Effects of Firm-Level Credit Supply Shocks on Project Completion Hazards

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Exposure	0.1094***	0.0667**	0.0636**	0.0605**	0.0573**	0.0734***	0.0716***
	(0.0270)	(0.0272)	(0.0273)	(0.0264)	(0.0266)	(0.0264)	(0.0266)
DepVar Mean	0.0397	0.0397	0.0397	0.0397	0.0397	0.0397	0.0397
Observation	122123	122123	122123	122123	122123	122100	122100
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 (1)). The sample consists of projects observed quarterly between 2013 and 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 10.94%. 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 percentiles, and then standardized. All firm controls are lagged by one period. The reduction in sample size from column (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.

sales growth) and firm financial solvency (e.g., leverage), which attempt to span the space of time-varying firm-level factors that could influence both project completion likelihood and banks' propensity to hide losses.[23]

Table 2 presents estimates of model (1). Column (1) starts by estimating a 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%. Column (2) tightens identification by including firm fixed effects and project state-by-year fixed effects. A one-standard-deviation increase in project-level exposure now raises hazards by 6.6%, a statistically significant effect. Column (3) shows the results are unaffected by the inclusion of firm-year controls. Column (4) adds National Industrial Classification (NIC) sector-by-year fixed effects. The effect size drops to 6%, but remains significant. Column (5) is our preferred specification. The additional firm-level controls barely move the point esti-

[23]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 us to include other fixed effects that control for project-level confounders. 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 is measured at this frequency. Even many official macroeconomic series are released based on financial year.

mate: hazards increase by 5.7%, a significant effect. Last, columns (6) and (7) replace sector-by-year with industry-by-year fixed effects. If anything, the tighter specifications yield larger completion effects. Appendix Figure A4 shows the gap in completion probabilities cumulates over the year following the shock.

We interpret our hazard estimates as evidence of project acceleration. This parallels a voluminous literature on unemployment dynamics in labor and public economics. Spikes in re-employment hazards around unemployment insurance benefit expirations have provided an empirical foundation for search-based theories, despite search effort often being unobservable.[24] Our empirical design is closest to Rothstein (2011), who quantified the effects of unexpected UI duration extensions on unemployment exit hazards and explored heterogeneous effects by spell duration. Our empirical evidence follows the same spirit: though we do not observe firms' investment rates (search effort) throughout the life of the project (unemployment spell), hazard estimates provide an empirical ground for the theory of endogenous time to build (unemployment duration). We later describe robustness tests and narrative checks that corroborate this interpretation.

3.3 Extensive Margin Results

Firms discount future cash flows, so accelerating completion increases the project's net present value. If it were 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 behavior in times of low credit supply. 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

$$\text{New Project Starts}_{ft} = \exp\left(\beta \times \text{Firm Exposure}_{ft} + X_{ft}\Gamma\right)\nu_{ft} \tag{2}$$

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 3 collects results. In a bivariate specification (column 1), a one-standard-deviation increase in firm-level exposure reduces new project starts by 7.8%. 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 with NIC sector-by-year fixed effects to control for sector-specific shocks. In our

[24] Dating back to Moffitt (1985) and Katz and Meyer (1990), exit rates from unemployment have been shown to spike close to benefit expiration. These and many follow-up estimates have been interpreted as evidence of endogenous search effort by job seekers and used to quantify the moral hazard costs of unemployment insurance (Chetty, 2008; Kroft and Notowidigdo, 2016).

Table 3: 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 (2)). The sample consists of the Capex-Prowess matched sample of firms observed yearly between 2013 and 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 percentiles, and then standardized. All firm controls are lagged by one period. The reduction in sample size from column (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.

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 reduces the point estimates to 5.8%. All effects are significant at conventional levels. The magnitudes remain in a similar range across specifications.

3.4 Heterogeneous Effects

A key implication of endogenous time to build is that the acceleration effect should vary by project maturity. Near-completion projects have imminent payoffs, with strong incentives to accelerate. Just-started projects instead resemble unbegun projects: with most of the construction ahead, slowing down is the natural response to tighter credit. Our model in Section 4 formalizes this prediction. Here, we take it as given and provide supportive evidence.

Testing for heterogeneous treatment effects requires ranking projects by their proximity to completion at the time of the shock, which we do not observe directly. A project's age alone is a problematic proxy. Heterogeneity in project returns can generate an age distribution where old age may reflect low quality rather than proximity to completion. Indeed, out of sample, age classifies which projects complete within the next two years barely better than chance (Appendix D.1).

We therefore estimate proximity to completion directly with a conditional survival model.

For a project of age a with characteristics X_p, the probability of completing within the next H quarters given survival to a is

$$P_H(a, X_p) = 1 - \frac{S(a + H|X_p)}{S(a|X_p)} \tag{3}$$

where $S(\cdot|X_p)$ is the survival function of construction spells. We estimate the survival function with a discrete-time hazard model on pre-AQR data using flexible age controls and a rich battery of predetermined covariates. The maturity scores are cross-fitted (i.e., no firm's own outcomes enter its prediction) and frozen at the shock date.[25] To ensure our measure indeed captures project maturity, we provide several out-of-sample tests. First, the estimated probabilities discriminate well between projects that do and do not complete within the next two years. Second, they are well calibrated: among projects of the same age, predicted completion probabilities map almost one-to-one with realized completion rates. Appendix D.1 details the construction and validation.

From each project's estimated survival curve we take two summaries of maturity at the time of the shock: expected residual time to build (restricted mean survival time, RMST, truncated at ten years) and the probability of completing within two years P_8. The two summarize the same survival curve at different points and are highly correlated.[26] We then estimate model (1), including each maturity measure along with its interaction with the firm exposure variable. Table 4 collects the OLS estimates.[27]

Across both measures, the effect of exposure on completion hazards is increasing in project maturity. For the project at mean maturity, a one-standard-deviation increase in exposure raises the quarterly completion probability by roughly 0.2 percentage points. Had the project been one standard deviation more (less) mature, the effect would have been larger (smaller) by roughly 0.4 percentage points. These are large effects relative to the mean quarterly completion probability of 1.9% (0.6% before 2016). Appendix D.3 provides a battery of robustness exercises. The same pattern holds (i) across maturity quartiles, (ii) controlling for interactions of age with exposure, (iii) expanding the risk set to earlier or later projects,[28]

[25] Cross-fitting removes overfitting bias. The scores are also well calibrated (slope of 1.08, Table A16), so their prediction error is nearly uncorrelated with the scores themselves. Unlike classical measurement error, such error does not attenuate the estimated maturity gradient.

[26] Appendix D.3 reports results for the probability of completing within $\bar{T} = 17$ quarters. This threshold value is derived analytically from the model and calibrated to the parameter values from our quantitative evaluation. The theory-motivated statistic delivers nearly identical results.

[27] Table 4 reports linear probability estimates for the quarterly completion indicator (Angrist and Pischke, 2009; Wooldridge, 2010). The theory's prediction is about the level of the completion probability across projects of different maturities, which a linear model estimates directly. Baseline hazards are close to zero for the least mature projects, so proportional effects are not defined over part of the maturity distribution; Appendix D.3 reports the PPML counterparts and reconciles the two margins.

[28] Because maturity at the shock is defined only for projects already under way in 2015Q2, the sample conditions on that risk set. Survival to the shock date is predetermined, so this conditioning defines the estimand (het-

and (iv) under alternative constructions of the maturity proxies. Hence, the cross-sectional patterns we document are consistent with the predictions of endogenous time to build.

3.5 Robustness Exercises

Our empirical design is valid if the control set ensures that the identification condition holds. Confounding would result if there existed factors that (i) influence project completion and correlate with the credit supply shock; *and* (ii) that we fail to control for.[29]

A first concern is the absence of firm-time effects in model (1). Suppose there were unobservables 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. The question then is how well our control set spans the space of time-varying firm-level confounders. We follow Khwaja and Mian (2008) and Chodorow-Reich (2014) and use the results from our credit regressions as suggestive evidence on the absence of unobservable firm-year confounders. If borrower-year unobservables were the chief sources of confounding, moving from borrower-year to borrower and year fixed effects in Table A6 would change point estimates considerably. Instead, estimates barely move: firm and year effects approximate 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. To further alleviate this identification concern, we present balance checks in Appendix Table A14. If our exposure measure predicted pre-audit firm characteristics, we would be concerned that banks were differentially cutting credit on these observables. We regress one-period lagged firm characteristics on a contemporaneous measure of exposure, including firm and year effects 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] Regardless, we include all these covariates as controls.

A second concern is concentrated bank lending: if a few firms dominate a bank's loan

erogeneity among ongoing projects) without affecting identification (van Houwelingen, 2007; Anderson et al., 1983).

[29] We focus our discussion on robustness to identification concerns. We provide a battery of econometric robustness exercises in Appendix C. Tables A8 and A11 show that our choice of weights for aggregating the bank shocks is conservative: using pre-period weights instead strengthens the results. Tables A9 and A12 reassure our results are not driven by the choice of estimator: semi-elasticities obtained by dividing point estimates of linear models by dependent variable means are similar to the PPML estimates. In unreported results, we also find robustness to other combinations of fixed effects and analysis windows; and find that our specification captures well effect dynamics, as effects on hazards are largest on impact and decline over time. The short-lived effect on hazards, however, translates into longer-lived effects on survival functions that persist one to two years out, so that our effects on capital formation are meaningful beyond impact horizons.

[30] We also run cross-sectional balance tests. In column (3) of Table A13, 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 than what is required by our identification assumption, but is overall encouraging.

Table 4: Effects of Exposure by Project Maturity

	(1)	(2)	(3)	(4)
Panel A: Residual Time to Build (RMST)				
Exposure	0.0021**	0.0021**	0.0022**	0.0023**
	(0.0010)	(0.0010)	(0.0009)	(0.0009)
RMST x Exposure	-0.0044***	-0.0043***	-0.0041***	-0.0040***
	(0.0009)	(0.0009)	(0.0009)	(0.0009)
Panel B: Two-Year Completion Probability (P8)				
Exposure	0.0020**	0.0020**	0.0021**	0.0022**
	(0.0010)	(0.0010)	(0.0009)	(0.0009)
P8 x Exposure	0.0036***	0.0036***	0.0033***	0.0033***
	(0.0009)	(0.0009)	(0.0010)	(0.0010)
Age at Shock FE	✓	✓	✓	✓
Firm FE	✓	✓	✓	✓
State x Year FE	✓	✓	✓	✓
Sector x Year FE	✓	✓		
Industry x Year FE			✓	✓
Firm Controls		✓		✓
N	101,632	101,632	101,607	101,607

Notes: the table reports OLS estimates of the effect of firm exposure to the Asset Quality Review on project-level completion hazards. The sample consists of projects observed quarterly between 2013 and 2019 that were active in 2015Q2. Firm exposure is defined as the weighted average of lender exposure, where weights correspond to outstanding loans in the previous year. The maturity measures are constructed from a conditional survival model estimated on pre-treatment data and frozen at the shock date: RMST (Panel A) is the expected residual time to build in quarters; P_8 (Panel B) is the probability of completing within the next 8 quarters. Each measure is winsorized at the 1st and 99th percentiles, and then standardized before being included in the regression model. See the main text and Appendix D.1 for details on the construction and validation of the maturity measures. The coefficient on exposure is interpreted as the percentage point effect on the quarterly completion probability for a one-standard-deviation increase in firm exposure for the project with mean maturity. The coefficient on the interaction is interpreted as the change in the main effect resulting from a one-standard-deviation change in the maturity measure; a higher RMST corresponds to a less mature project, hence the negative signs in Panel A. For example, the first row in Panel B column (2) reports that moving from a firm with average exposure to a firm with one-standard-deviation higher exposure increases the quarterly completion probability by 0.20 percentage points for the project at the mean maturity. The second row reports that if the project had been one standard deviation above (below) the mean maturity at the time of the shock, the effect would have increased (declined) by 0.36 percentage points. We control non-parametrically for the age of the project at the time of the shock. 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 percentiles, 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 singleton observations dropped 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.

book, shocks to those firms could endogenize the bank shock. Appendix Figure A5 shows this is not the case – most firms comprise a small share of total charges, and banks most exposed to the audit had *less* concentrated lending portfolios. A third concern is that the shock could affect the censoring process. Appendix Table A10 re-estimates our specifications replacing the outcome with an indicator for a project becoming lost to follow-up. Across all specifications with fixed effects, exposure does not predict the censoring hazard.[31] A fourth concern is India's demonetization of November 2016. Chodorow-Reich et al. (2020) identify the real (but short-lived) effects of demonetization from the geographic distribution of new currency notes, which varied across states according to the RBI's regional allocation of notes. Our state-by-year fixed effects control non-parametrically for this variation.[32]

A final concern is that some projects could have had their scopes reduced after the shock, in which case a positive effect on the completion hazard could reflect a reduction in time to build even if it was purely exogenous. In Appendix E.2, we alleviate this concern surrounding the interpretation of our estimands. First, we provide qualitative evidence on the construction industry, arguing that commonly used contracts and technical considerations hinder changes in scope of ongoing projects. Such frictions are more pronounced for large capex projects undertaken by sizable sponsors, as in our sample. Second, we repeat our analyses discarding projects that are proxied to be most susceptible to potential reductions in scope. We classify such projects based on how the data provider keeps track of changes in project scopes when collecting the raw data; or based on technological or regulatory aspects of construction that preclude large variations while the project is ongoing. Across these quantitative exercises, our results are similar to our baseline, supporting our preferred interpretation that firms are speeding up project completion.

One concern we cannot speak to is changes to project quality. This is due to data limitations (e.g., we do not observe the quality of materials used). If sponsors responded by reducing input quality, and if low-quality inputs have lower exogenous lead times, projects could be completed faster. This story could threaten identification if banks differentially cut credit to firms based on their ability to accelerate projects via this quality downgrading channel. This is unlikely: quality is usually fixed in the contract specifications and bill of quantities (Appendix E.2). Further, it could be interpreted as a form of endogenous time to build (sub-

[31] In Appendix E.1, we alleviate residual censoring 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 due to exposure, projects in vintages prior to the shock would not have been affected and are thus immune to the selection concern. Second, we test whether exposure predicts differential pre-trends in firm-level project stocks. If CapexDx had become less able to backfill data on exposed projects, we should observe a downward pre-trend in the relative number of exposed projects. Across the two tests, we find no evidence that our results are driven by differential censoring.

[32] State-by-time fixed effects also mitigate concerns about equilibrium effects from local input markets. If the credit shock leads some companies to cut back on projects, localized inputs are freed up presumably at lower cost, which could spur other companies to speed up their ongoing projects. In any case, this mechanism would still be consistent with endogenous time to build.

stituting inputs to complete faster), though not one that accords with the theoretical model we will use to rationalize the data.

It is worth highlighting that media coverage at the time provides textual corroboration for our econometric findings. The Economic Times, a leading news source in India, reported in May 2016 how "(...) builders too are feeling the heat with cash flow drying up and banks no longer ready to play ball. So, the only option they now have is to complete projects," with builders "trying to rope in private equity funders (...) ask[ing] for astronomically high returns – in the range of over 25%"; and in May 2015, how "banks were not healthy and their NPAs were rising", with builders "focusing on completing projects which will then unlock cash flows that are locked in." Companies mentioned "high cost funds which we are using to finish projects. These do lower our margins but the focus is on trying to complete projects"; and "not launching projects and phases but rather concentrating on construction and delivery of existing sold projects".[33]

Last, we touch on external validity. Endogenous time to build is not specific to our setting. History offers well-known examples. During World War II, the US accelerated the Pentagon's construction with workers on site around the clock, and the building was delivered at nearly twice its budget (Goldberg, 1992; Vogel, 2007; Rollo, 2024). Recent episodes point the same way. In 2016, spurred by demand and tax incentives, Tesla "doubled the number of people constructing the 'gigafactory,' (...) 1,000 workers build seven days a week on two shifts in an effort to start churning out lithium-ion cells" (Ramsey, 2016). The current boom in data center construction further provides timely external validity. The CEO of Crusoe (the firm building one of OpenAI's largest data centers) described its aim as delivering on "the fastest schedule that a 100-megawatt-or-greater data center has ever been built", adding: "In this moment, speed matters a lot" (Ghaffary et al., 2025). And speed is costly. In the words of a construction lender at Citi: "Optimizing scope while maintaining an aggressive schedule results in an increase in project cost. However, the end users in this industry are less concerned about the cost component" (Gallo, 2026). These episodes admittedly reflect demand more than credit conditions, but they illustrate how firms choose how fast to build and how changes in firm demand or credit conditions can drive time to build.

All in all, our robustness checks and corroborating narratives reinforce our main empirical findings. We now formalize the economic theory that rationalizes our econometric evidence.

[33]The article titles aptly summarize the gist of the narratives: "Low on stock markets: Cash-strapped realty firms run from pillar to post to raise funds" (ET 2015) and "Speeding up delivery only way out for cash-strapped builders" (ET 2016). Though the articles focus on real estate developers, they corroborate the mechanism behind our estimates. We also read through selected earnings calls from listed firms that appear in our regression sample. Appendix E.3 presents some quotes. We found examples among firms in manufacturing, energy, etc., some of which mentioned unlocking cash flows faster as a reason for expediting projects. Though systematic narrative evidence is beyond the scope of this paper – our main econometric evidence comes from our quasi-experimental hazard estimates – the quotes exemplify how firms discuss project timelines and expedited completion.

4 A Theory of Investment with Endogenous Time to Build

We rationalize our empirical findings by developing a theory of investment with endogenous time to build. By endogenous, we mean that the firm *chooses* time to build, rather than treating it as an exogenous primitive. Our model is minimalist.[34] 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, which we defer to Section 5.

Setup Time is continuous and indexed by t. A firm manages a collection of capital investment projects.[35] The firm's discount rate is given by $r > 0$. Each period, a continuum of projects become available. Completing a project requires investing an amount $\bar{x}$. After the project is finished, say at time T, the firm obtains z units of capital, drawn at project inception from a p.d.f. $f(z)$. Capital is valued at a common price p, which firms take as given. We relax this assumption and prove that our analytical results hold in Appendix H.3. Project returns are pz and only accrue upon project completion. Our setting incorporates project heterogeneity in returns rather than size, an assumption we relax in Section 5. Last, there is no risk. In Appendix H.4, we incorporate risk and prove that Propositions 2–4 continue to hold. 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

$$\int_0^T i(t)dt = \bar{x} \tag{4}$$

Equation (4) states projects are completed when investment flows add up to total size. Suppose p, z and r are constant (i.e., 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 (4)

$$\max_{i(t),T} e^{-rT}pz - \int_0^T e^{-rt}\left[i(t) + \kappa(i(t))\right]dt \quad \text{s.t.} \int_0^T i(t)dt = \bar{x} \tag{5}$$

We introduce frictions into the construction process, modeled as additional costs $\kappa(i)$,

[34] Standard alternatives do not naturally generate the heterogeneous acceleration pattern documented in Table 4. Convex adjustment costs generate smooth responses uniform across project stages. Capital utilization operates on existing productive capital, not capital under construction. Exogenous time to build predicts constant investment conditional on entry (see also footnote 37). Spending constraints can generate a stage gradient, but not the positive average effect on completion (Appendix H.8).

[35] 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 to study industry dynamics.

assumed to be weakly increasing and weakly convex.[36] The frictional costs $\kappa(i)$ do not contribute to installed capacity, i.e., they are deadweight losses. 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 get in each other's way. 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 overtime pay, renegotiating contracts, speeding up deliveries, facing a convex labor supply curve for construction workers, etc. (Bajari and Tadelis, 2001).

Whenever necessary (i.e., in quantitative exercises), we assume that $\kappa(i)$ has the following 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.[37] The exogenous case illustrates how $\kappa(\cdot)$ can stand for a technological constraint on the rate at which the project can be completed.

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

$$x(t) = \int_0^t i(s)ds \tag{6}$$

[36] 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. 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. This is counterfactual. Physical capital is not built in a day, as the data shows $T > 0$ for most projects.

[37] 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 $i < \bar{i}$ before the shock, increasing their investment to $i = \bar{i}$ when it hit. 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. $\bar{i}$ captures the technological pace of construction, and a credit shock could move it only through the construction side, for instance through contractors that borrow from the same banks. But a contractor under financial stress would, if anything, slow work down, lowering $\bar{i}$; and a contractor that speeds up to collect payments sooner is itself choosing time to build in response to financing conditions.

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

$$rv(x,z) = \max_{i \geq 0} -i - \kappa(i) + i\frac{\partial v(x,z)}{\partial x} \tag{7}$$

and terminal condition

$$v(\bar{x},z) = pz \tag{8}$$

The HJB characterizes optimal firm behavior. The flow value of a project to the firm is $rv(x,z)$. When the firm behaves optimally, it trades off the costs and benefits of investment. Costs are borne 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

$$1 + \kappa'(i) = \frac{\partial v(x,z)}{\partial x} \tag{9}$$

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 – a project-level marginal q (Abel, 1981; Hayashi, 1982). With slight abuse of notation, we denote by $i(x,z)$ the resulting policy in the state space. We solve the HJB (7) numerically using finite-difference methods (Achdou et al., 2021).

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:[38] start all projects with high enough return $z > \bar{z}$, where the break-even point $\bar{z}$ satisfies

$$v(0,\bar{z}) = 0 \tag{10}$$

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

[38]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 negligible effects on the steady state.

the steady state distribution. It satisfies a Kolmogorov Forward Equation (KFE)

$$\forall (x,z),\ \frac{\partial}{\partial x}\left[i(x,z)g(x,z)\right] = 0 \tag{11}$$

The intuition for the KFE is simple. The investment policy $i(x,z)$ is the speed at which projects move along the x dimension. The flow of projects through a given point x is, therefore, $i(x,z)g(x,z)$: the product of the density of projects at this point and the velocity with which they travel. In steady state, the flow of projects through any given point must be constant, which is equivalent to equation (11).

Now recall that a constant flow $f(z) \cdot 1\{z \geq \bar{z}\}$ of projects with productivity z enters the economy at each instant.[39] 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

$$\forall (x,z),\ g(x,z) = \frac{f(z) \cdot 1\{z \geq \bar{z}\}}{i(x,z)} \tag{12}$$

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 H) with an appropriate terminal condition $T(\bar{x},z) = 0$ yields

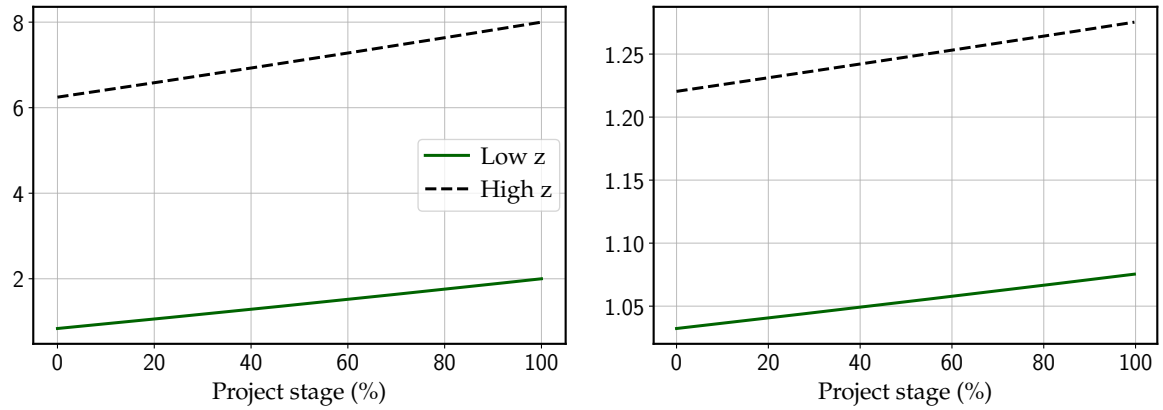
$$T(x,z) = \int_x^{\bar{x}} \frac{1}{i(y,z)} dy \tag{13}$$

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 project reaches $\bar{x}$. (13) states this formally.

Steady State Solution Figure 2 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. Appendix Figure A11 reports the distribution and time to completion as a function of stage.

[39] 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 2: Steady state objects as functions of project stage



(a) Value function $v(x, z)$ (b) Optimal investment $i(x,z)/\bar{i}$

Note: the figures plot the steady state solution of the model. Panel (a) plots the value function $v(x, z)$ and panel (b) plots the investment policy $i(x, z)/\bar{i}$. 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 5: $r = 0.01$, $p = 1$, $\bar{x} = 1$, $\bar{i} = 0.083$, $\kappa = 3.09$.

4.1 Dynamics

To map our model to the empirical results in Sections 3.2 and 3.3, we characterize the response of optimal investment policies to a credit shock. We model the credit shock as an increase in the firm's discount rate r. Appendix H.7 discusses theoretical channels and empirical evidence supportive of this mapping.[40] Intuitively, as credit dries up, firms value cash today more highly, which we capture as higher discount rates. This can be driven by banks raising credit costs today, firms engaging in precautionary behavior over potential future credit rationing or higher credit costs, or other factors that raise discount rates independently of credit costs (Graham, 2022). Our observed empirical patterns – and therefore the basic theory of endogenous time to build – can arise from a neoclassical framework under a more comprehensive view of discount rates.[41]

Mathematically, we specify a perfect-foresight AR(1) path for discount rates, i.e., $r(t) =$

[40] There, we also briefly sketch a microfoundation linking provisioning shocks to discount rates. Abstracting from the banking microfoundation in the main model allows us to focus on endogenous time to build for real investment, where our main contribution lies. Further, the change in discount rates may arise through other channels, which may be quantitatively important. Our estimation strategy will allow us to be agnostic over which channels matter most.

[41] The neoclassical chassis further clarifies that our empirical and theoretical results are not predicated on the particular form that financing frictions may take (for example, asset- vs earnings-based constraints). In Appendix H.8, we assess whether a model with borrowing constraints can rationalize our empirical findings, and elaborate on why we prefer the discount rate shock formulation. Additionally, and perhaps most pragmatically, this computationally tractable choice allows us to carry out meaningful estimation and counterfactual exercises focusing on the two main novelties of our paper: the theoretical modeling of endogenous time to build, and the rich disaggregated project-level data that can discipline it.

$r + se^{-\rho t}$. 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 H for all proofs. □

Lemma 1 allows us to characterize value functions and investment policies analytically following the shock 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*

$$\hat{v}(x, z) = -\left(\frac{1 - e^{-\rho T(x)}}{\rho}\right) e^{-rT(x)} pz + \int_0^{T(x)} \left(\frac{1 - e^{-\rho\tau}}{\rho}\right) e^{-r\tau} \left[i(y(\tau), z) + \kappa\left(i(y(\tau), z)\right)\right] d\tau, \tag{14}$$

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 permanent shock to discount rates ($\rho \to 0$) is particularly illustrative. Project value changes by its **duration**. To compute the change in project value, we take a weighted sum of time periods over the remaining life of the project. We weight each time period by the NPV of cash flows accruing in the time period.[42] We use the steady state investment flows and remaining time to build to compute duration.

[42] An analogy with classical mechanics once again helps with the notation. Define $\tau = T(x) - T(y)$, so that

Intuitively, duration gives us how long it will take to reap back the cash flows of the project. Unlike 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:

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

$$\hat{i}(x,z) = \frac{1}{\kappa''(i(x,z))} \cdot \frac{\partial}{\partial x}\hat{v}(x,z), \tag{15}$$

where

$$\frac{\partial}{\partial x}\hat{v}(x,z) = \frac{1}{i(x,z)} \left[\underbrace{v(x,z)}_{Acceleration} + (r+\rho) \underbrace{\hat{v}(x,z)}_{Duration} \right] \tag{16}$$

Proposition 3 characterizes how to re-optimize investment plans following a change in discount rates. Equation (15) 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 the discount rate shock alter this condition? Proposition (3) is the paper's central theoretical result. It 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 is standard asset pricing: as discount rates rise, project values fall. 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., $\lim_{x\to\bar{x}} T(x) = 0$). In turn, $v(x,z)$ approaches $pz > 0$. Therefore, it follows from (16) that $\hat{i}(x,z)$ must eventually be positive. □

$d\tau = {}^{dy}/_{i(y,z)}$. 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 $y(\tau)$.

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 (16) 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* – a result that qualifies conventional capital budgeting theory. How long it takes for the project to yield cash flows matters crucially. This depends on how quickly the firm is building the project; and on how costly it is to accelerate – hence the denominator in (15) and (16). Where the project currently is in its lifecycle (x) summarizes these considerations.

Our results borrow intuition from concepts in fixed income pricing. In Appendix G, 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 interest rates incentivize agents to shift cash flows (or, more broadly, actions) over time. Differently from bond pricing, firms may operate the project after its completion. In Appendix H.3, we extend this result to the case that the firm must operate the project following its completion. p then corresponds to the future value of operating cash flows, which falls as discount rates rise. Even then, we show that so long as cash flows depreciate at a positive rate, it is still optimal to accelerate completion for sufficiently mature projects.

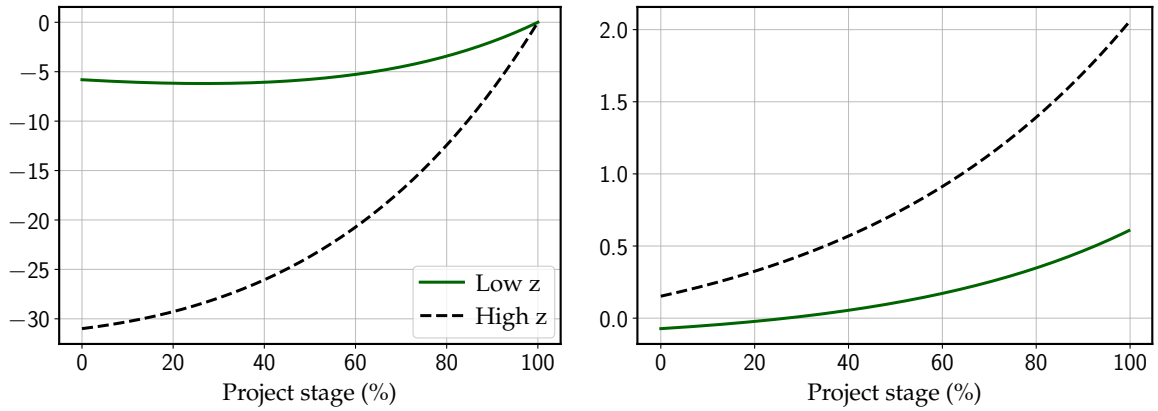
Figure 3 illustrates our propositions by showing the impact responses of the value function $\hat{v}(x,z)$ and optimal investment $\hat{i}(x,z)$ to the shock. Panel (a) shows $\hat{v}(x,z)$. Note that it is strictly negative for $x < \bar{x}$, as the payoff upon project completion is more strongly discounted after the shock. In short, all projects fall in value. But project value falls by more for 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. Appendix H shows that this threshold is unique for every project type.

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*

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

Figure 3: 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 A24. Remaining parameters are set at the same values as in Figure 2.

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

$$e^{-rT(0,\bar{z})} \times p\hat{z} = -\hat{v}(0,\bar{z}) \tag{17}$$

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, make 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 Quantitative Model

So far, we have presented evidence inconsistent with exogenous time to build and shown the theoretical implications of relaxing this assumption for project-level investment. We now enrich the model and bring it to the data. In short, we calibrate the model to the measurements from Section 2 and Appendix B; and match the estimated elasticities from Section 3.

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 $\log \bar{x} \sim \mathcal{N}\left(\mu_{\bar{x}}, \sigma^2_{\bar{x}}\right)$. 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 $\bar{i}$ depends log-linearly on project size: $\log \bar{i} \sim \mathcal{N}\left(\beta_0^{\bar{i}} + \beta_1^{\bar{i}} \log \bar{x}, \sigma^2_{\bar{i}}\right)$. Higher draws of $\bar{i}$ for any fixed $\bar{x}$ imply a quicker time to completion (e.g., if time to build were exogenous, a project's (random) time to build would be $\bar{T} = \bar{x}/\bar{i}$). We set the parameters of the $\bar{i}$ distribution to match the average log time to build within each log size quartile. Figure A9a 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 A9b shows our results.

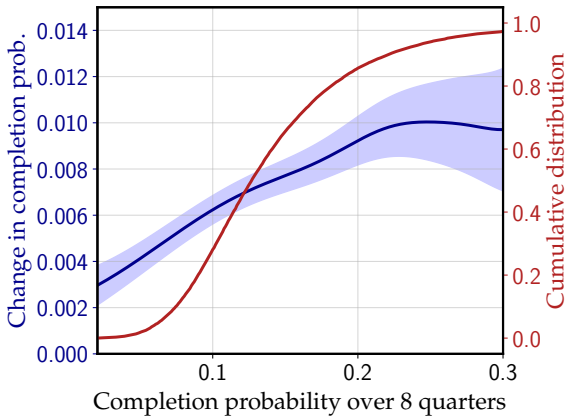
To calibrate the remaining parameters, we use the intensive and extensive margin semi-elasticities $(\beta_{int}, \beta_{ext})$ estimated in column (5) of Tables 2 and 3. Introducing banks, bank relationships, and banking shocks would add considerable complexity. Instead, following our earlier theoretical results, we map the banking shock to an increase in firms' discount rates. We first estimate ρ by running panel autoregressions of Exposure_{ft} on its first lag, firm, and year fixed effects.[43] Figure A9c 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 set these parameters to match our empirical results using a standard minimum distance routine.[44] 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}

[43] 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.

[44] The model-implied semi-elasticities are always computed in partial equilibrium. This is because our regression estimates are informative about partial equilibrium responses: the inclusion of state-by-year and sector-by-year fixed effects absorbs 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.

Figure 4: Heterogeneous Effects in the Model by Estimated Completion Probability



Notes: the figure plots the impact change in the quarterly completion probability after the discount rate shock for ongoing model projects versus their estimated probability of completing within the next eight quarters, P_8. We simulate a sample of 100,000 projects and compute for each one the predicted completion probabilities conditional on age a and total cost $\bar{x}$, and treating $(\bar{\imath}, z)$ as unobservables, as in Section 3.4. We then fit a nonparametric regression using a Gaussian kernel with bandwidth 0.03 to obtain the fitted line and the 95% confidence intervals, reported in blue. The red lines show the cumulative distribution of simulated projects. Age is measured in quarters. All parameter values, including shock size and persistence, are as in Appendix Table A24.

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 A24 collects our calibrated parameters. We estimate a shock size $s = 2\%$, and an acceleration cost $\kappa = 3.09$. Appendix Table A25 reports that the calibrated model fits the targeted moments well.

Heterogeneous effects. As an untargeted moment, we ask whether the model reproduces the maturity gradient documented in Section 3.4. For every project that is active when the shock hits, we compute the same maturity variable used in the data (P_8, the probability of completing within the next eight quarters given survival to the project's current age).[45] Figure 4 shows in blue a nonparametric kernel regression of the impact change in the quarterly completion probability against P_8 in a simulated sample with 100,000 projects, as well as the corresponding 95% confidence interval. The response rises with maturity, as in Section 3.4. Between the 10th and 90th percentiles, the marginal effect of a shock on the completion probability increases from approximately 0.5 p.p. to 1 p.p. This gradient is untargeted, as we match only the average intensive and extensive margin responses, leaving no free parameter in the model to govern how the response varies with maturity.

[45] In the model, a project of type $(\bar{x}, \bar{\imath}, z)$ and age a has a deterministic time to completion $T(0, z, \bar{x}, \bar{\imath}) - a$, where $T(0, z, \bar{x}, \bar{\imath})$ is its time to build. We mimic the econometrician's information set: $\bar{x}$ and a are observed, while $\bar{\imath}$ and z are not. The model score is therefore $P_8(\bar{x}, a) = 1 - S(a + 8 \mid \bar{x})/S(a \mid \bar{x})$, where S is the survival function of construction spells in the model, integrated over the unobserved $(\bar{\imath}, z)$. This is the formula of Section 3.4 applied to the model's own completion distribution; no regression is involved.

6 Endogenous Time to Build and Investment Fluctuations

Equipped with our calibrated model, we study how endogenous time to build shapes investment dynamics and policy. Endogenous time to build generates qualitatively distinct dynamics: an initial investment spike following adverse discount rate shocks (as near-completion projects are accelerated), oscillatory boom-bust patterns driven by the intensive margin, and state-dependent effects of monetary policy that depend on the maturity composition of the project pipeline. These features arise from the interaction of heterogeneous projects at different completion stages – a dimension absent from models with convex or fixed adjustment costs. We also show that endogenous time to build creates a measurement wedge between investment expenditures and productive capital formation, and provides insights for the design of investment tax credits.

6.1 Aggregating Across Projects

To examine the response of investment across projects (i.e., aggregate) to a shock in discount rates (as in our empirical setup), we must first 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 $I^K(t)$. We have

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

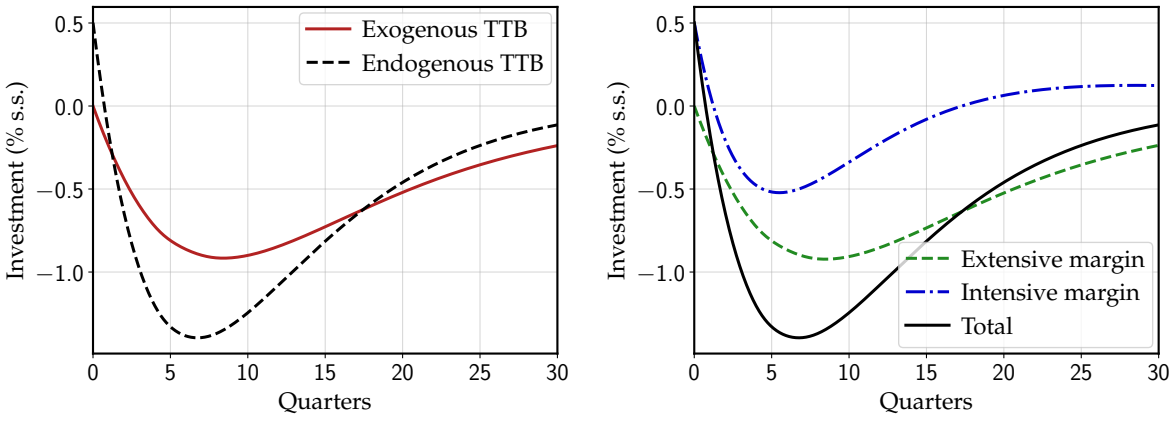
that is, productive capital grows because more projects are completed. Since projects can differ in their output of capital, we aggregate across the distribution of output.[46] Accounting investment, in turn, is given by aggregating investment flows across all ongoing projects

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

[46] In steady state, equation (12) 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 $M_e \equiv \int_{\bar{z}}^{\infty} f(z)dz$ be the flow of entrants in steady state, we have $I^{acc}_{ss} = M_e \cdot \bar{x}$. In turn, since only projects with z above $\bar{z}$ enter, we have

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

Figure 5: Impulse response functions of investment to discount rate shock in partial equilibrium



(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 A24.

6.2 Amplification and Propagation with Endogenous Time to Build

Figure 5 illustrates our results for a discount rate shock, holding fixed p. We start with partial equilibrium since this assumption is more aligned with our empirical setup. In panel (a), we plot the response of investment when time to build is endogenous (black dashed line) versus exogenous (red solid line, $\kappa \to \infty$). 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 equivalent to investment adjustment costs (Lucca, 2007). The exogenous benchmark thus stands in for the adjustment costs common in macro models; the endogenous version shows how that friction responds to the state of the project pipeline (Section 6.5). Lanteri (2018) makes the

analogous point for irreversibility, which becomes endogenously larger in downturns.

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, with intensive margin effects becoming negative. 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 toward early-stage projects, for which lower investment is optimal even in steady state – see Figure 2. The distributional change thus drives 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 oscillatory boom-bust dynamics as 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 effect is driven by the intensive margin: investment in ongoing projects quickly booms, shortly collapses, and then oscillates back to steady state.

Given investment is more sensitive to shocks in a world with endogenous time to build, 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. Importantly, this additional volatility has a qualitatively different character than what standard models produce: it features short-lived spikes driven by the intensive margin and oscillatory dynamics absent from models with convex adjustment costs. Two results provide a lens to interpret the increased responsiveness. The first characterizes the impact response of investment; the second shows that its cumulative response is pinned down by an accounting identity.

Recall from Proposition 3 that the response of an ongoing project trades off an acceleration motive, $v(x,z)$, against a duration effect, $(r+\rho)\hat{v}(x,z)$. Define the project's proportional

duration $D(x,z)$ and its *acceleration responsiveness* $\Lambda(x,z)$ as

$$D(x,z) \equiv -\frac{\hat{v}(x,z)}{v(x,z)}, \qquad \Lambda(x,z) \equiv \frac{v(x,z)}{\kappa''(i(x,z))\, i(x,z)^2}.$$

Proportional duration is the semi-elasticity of project value to the shock, and Λ collects the terms that scale the size of the response: the value at stake, the local ease of adjusting the pace of construction, and the conversion into a proportional change in investment. Equations (15)–(16) can then be written as

$$\frac{\hat{i}(x,z)}{i(x,z)} = \Lambda(x,z)\left[1-(r+\rho)D(x,z)\right].$$

Let $M_e \equiv \int_{\bar{z}}^{\infty} f(z)dz$ denote the flow of entrants in steady state, so that $I_{ss}^{acc} = M_e \bar{x}$ by equation (12), and define the average acceleration responsiveness and the responsiveness-weighted duration as

$$\bar{\Lambda} \equiv \frac{1}{M_e \bar{x}} \int_{\bar{z}}^{\infty} f(z) \int_0^{\bar{x}} \Lambda(x,z)\, dx\, dz, \qquad D^{\Lambda} \equiv \frac{\int_{\bar{z}}^{\infty} f(z) \int_0^{\bar{x}} \Lambda(x,z) D(x,z)\, dx\, dz}{\int_{\bar{z}}^{\infty} f(z) \int_0^{\bar{x}} \Lambda(x,z)\, dx\, dz}.$$

Proposition 6. *Suppose the economy is in steady state when the discount rate shock hits. To first order, the impact response of accounting investment is*

$$\frac{\hat{I}^{acc}(0)}{I_{ss}^{acc}} = \bar{\Lambda}\left[1-(r+\rho)D^{\Lambda}\right], \tag{18}$$

and average investment per ongoing project responds by the same proportion. Accounting investment rises on impact if and only if $D^{\Lambda} < 1/(r+\rho)$.

Proposition 6 is the aggregate counterpart of Proposition 3. The steady-state distribution places more mass on stages where projects move slowly, since $g \propto 1/i$, but those projects also spend less per unit of time. The two effects cancel in accounting investment, so that, conditional on z, each unit of physical progress carries the same weight. What governs the sign of the aggregate response is therefore not the average duration of ongoing projects, but duration weighted towards projects whose investment is more responsive.

While the impact response depends on the distribution of duration across projects, the cumulative response is pinned down by an accounting identity. Let $b(z,t) \equiv i(0,z,t)g(0,z,t)$ denote the flow of type-z projects started at time t, $B(t) \equiv \int b(z,t)dz$ the total flow of project starts, and define the stock of *work-in-progress* as

$$W(t) \equiv \int \int_0^{\bar{x}} (\bar{x}-x)\, g(x,z,t)\, dx\, dz. \tag{19}$$

Proposition 7. *Along any path of the economy,*

$$\dot{W}(t) = \bar{x}B(t) - I^{acc}(t). \tag{20}$$

Consequently, for any temporary shock after which the distribution of projects returns to its steady state: (i) the cumulative response of accounting investment equals the cumulative response of project starts,

$$\int_0^\infty [I^{acc}(t) - I^{acc}_{ss}]\, dt = \bar{x} \int_0^\infty [B(t) - B_{ss}]\, dt; \tag{21}$$

(ii) the projects ongoing at the time of the shock invest a total of $W(0)$ over their remaining lives, regardless of the shock, so that the cumulative effect on the intensive margin of investment is zero.

Proposition (7) 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. The statement does not rely on a first-order approximation and holds in general equilibrium, as it uses only the law of motion of the distribution. 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. The intuition behind the result is as follows. Every project that starts adds $\bar{x}$ units of work-in-progress, and every unit of accounting investment removes one. Endogenous time to build can therefore reshape *when* ongoing projects absorb investment, but not *how much* investment they absorb in total. Cumulative investment responds to a shock only to the extent that project starts do. Bringing the two propositions together, we can formalize how endogenous time to build generates boom-bust cycles:

Corollary 8. *If $D^\Lambda < 1/(r+\rho)$, the response of accounting investment to the discount rate shock is positive on impact and negative at some later date.*

Corollary 8 formalizes the pattern in Figure 5. The initial spike is an intertemporal reallocation of construction work that would have taken place anyway. Because that work is brought forward, and because fewer projects are started while discount rates are high, investment must subsequently fall below its steady-state level. The identity does not determine the timing or depth of the reversal, which depend on the full distributional dynamics, but shows how boom-bust patterns can be generic features in neoclassical investment theory.

6.3 Amplification and Propagation in General Equilibrium

We now show that general equilibrium forces do not undo our partial equilibrium results.

GE Setup. 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 from the capital goods producers, which we have studied throughout the paper so far. Assuming that costs of construction are rebated to consumers, the aggregate resource constraint is

$$Y(t) = C(t) + I^{acc}(t)$$

where $C(t)$ is household consumption. We assume a representative household has isoelastic preferences over consumption $C(t)$, with (inverse) elasticity of intertemporal substitution σ and discount rate β.

Standard profit maximization by final goods producers determines capital prices, which reflect future marginal products of capital

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

where $MPK(t) = \alpha A(t)K(t)^{\alpha-1}$ and r^f is the discount rate of final goods producers. 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. Hence, general equilibrium forces dampen partial equilibrium responses under decreasing returns.

The structure above suffices to endogenize capital prices and understand the joint behavior of returns on capital and investment in industry models (Kalouptsidi 2014, Greenwood and Hanson 2014). In macroeconomic models, further assumptions are needed to pin down discount rates $r(t)$. Standard specifications of household preferences tend to undo the partial equilibrium responses of investment to shocks in general equilibrium.[47] Winberry (2021) adds consumption habits to ensure interest rates are not too sensitive to consumption. We pursue a route that better aligns with our empirical setup by differentiating discount rates of

[47] See House (2014) and Koby and Wolf (2020). The intuition is that changes in investment shift consumption through the resource constraint, which then shifts interest rates through the consumption smoothing motive of households. The increase in interest rates due to a steeper consumption path mutes the semi-elasticity of investment to a range of shocks.

capital producers and final goods producers.

To be precise, after any MIT shock, final goods producers discount future profits according to the stochastic discount factor

$$M^f(t) = e^{-\beta t}\left(\frac{C(t)}{C}\right)^{-\sigma}$$

where C is steady state consumption. Capital goods producers in turn are subject to a discount rate shock $\varepsilon(t)$, so that their SDF is

$$M^c(t) = e^{-\int_0^t(\beta+\varepsilon(u))du}\left(\frac{C(t)}{C}\right)^{-\sigma}$$

To keep the same notation as before, let $r^f(t)$ and $r^c(t)$ be the discount rate of final and capital goods producers, respectively. We then have[48]

$$r^f(t) = \beta + \sigma\frac{\frac{d}{dt}C(t)}{C(t)} \tag{22}$$

$$r^c(t) = r^f(t) + \varepsilon(t) \tag{23}$$

The key insight is that we can simultaneously shock discount rates of capital producers while allowing GE forces to operate via the usual term $\sigma\frac{\frac{d}{dt}C(t)}{C(t)}$. We interpret $\varepsilon(t)$ as banks deciding to charge higher rates when lending to capital producers. This "spread" shock could be microfounded by a range of macro-financial models where a banking block is explicitly spelled out, e.g., Gertler and Karadi (2011). The key idea is that some form of market incompleteness (e.g., market segmentation or overlapping generations of bankers without perfect annuity markets) does not allow required rates of return by banks and households to be equalized, leading to a "credit spread" in the economy. This spread is exogenous, but riskless rates of return are endogenously determined by household consumption. With this setup, we can spell out how capital prices and interest rates respond in GE:

Proposition 9. *Let hatted variables denote linearized deviations from steady state. Suppose we announce paths $\hat{A}(t)$ and $\hat{\varepsilon}(t)$ for aggregate productivity and discount rate shocks. Then, to first order, the response of capital prices is given by*

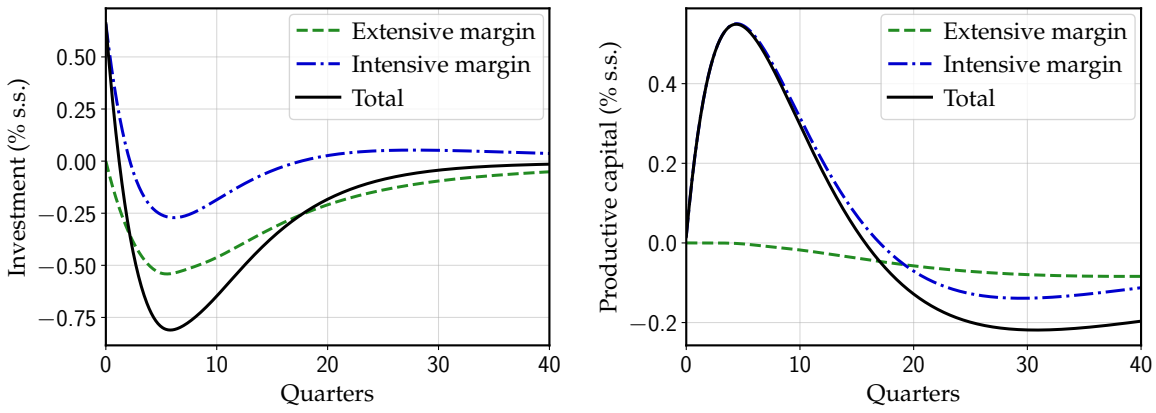
$$\frac{\hat{p}(t)}{p} = (r^f+\delta)\int_0^\infty e^{-(r^f+\delta)s}\left[\frac{\hat{A}(t+s)}{A} - (1-\alpha)\frac{\hat{K}(t+s)}{K} - \frac{\hat{r}^f(t+s)}{r^f+\delta}\right]ds \tag{24}$$

[48]This follows by defining the discount rate that matches the SDF $e^{-\beta t}\left(\frac{C(t)}{C}\right)^{-\sigma} = e^{-\int_0^t r^f(u)du}$, consistent with standard asset pricing.

The proof follows from linearizing the expression for $p(t)$ above around the steady state. Proposition 9 highlights three channels by which fluctuations in productivity and discount rates change capital prices. First, there is the direct effect of productivity shocks on the marginal product of capital. When capital is more productive, final goods producers are willing to pay more for it, raising $p(t)$. Second, as productivity or discount rates change, firms revise their investment, thereby altering future capital stocks $\hat{K}(t+s)/K$. Third, there is a direct effect of interest rate changes on the discounting of future marginal products of capital, again holding fixed future paths for the capital stock. As interest rates rise, future marginal products of capital are discounted more heavily, lowering $p(t)$.

Figure 6 shows the response to a discount rate shock $\hat{\varepsilon}(t)$ to capital producers, holding productivity fixed (i.e., $\hat{A} = 0$). First, there is an effect on capital stocks. Capital producers increase investment as they rush to complete projects close to completion. This pushes down the price of capital, lowering incentives to invest. As expected, the IRF magnitudes are smaller in GE. Regardless, the key results from Figure 5 – the short-run investment spike, the oscillatory dynamics, the endogenous amplification – all survive in general equilibrium.

Figure 6: IRFs of investment and capital to discount rate shock in general equilibrium



(a) Decomposition of investment IRF between intensive and extensive margins

(b) Decomposition of capital IRF between intensive and extensive margins

Note: the figure plots (quarterly) impulse response functions (IRFs) of aggregate investment and the capital stock to a shock in discount rates in general equilibrium. The y-axis measures percent-deviations from steady state values. Panels (a) and (b) plot the IRFs for investment and productive capital, respectively. The total IRF (black solid line) from the endogenous time to build model is decomposed 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 A24.

6.4 Tax Policy: Reshuffling Investment over Time

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.[49] But when time to build is endogenous, investment responses to cash flow shocks depend not only on whether projects are close to 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 7 shows the resulting IRFs of accounting investment and productive capital.

In panel (a), the green dashed line plots the IRF of accounting investment in the exogenous time to build model (i.e., the extensive margin of the endogenous 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 return 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, even after the tax credit is long gone. Only as they are completed does investment return to steady state.

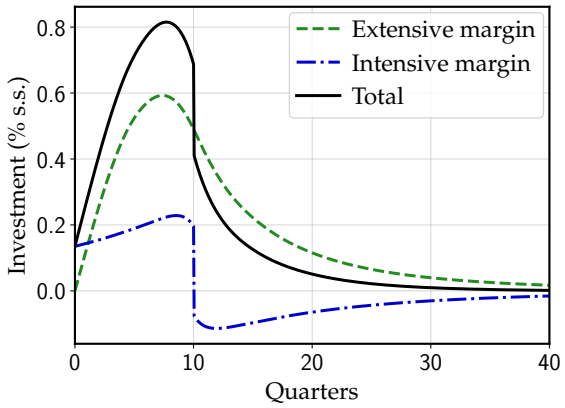
The blue dot-dashed line turns to the endogenous time to build effect. The striking feature of the IRF is its shape: investment rises on impact, keeps rising, but then collapses at the time that the credit expires, and then reverts to steady state. The intensive margin effectively allows firms to reshuffle investment over time. It is entirely responsible for the impact effect: investment flows rise for all ongoing projects as firms front-load investment to periods when it is cheaper to build. The intensive margin slump in turn 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.

Panel (b) shows the productive capital IRF. 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. Last, panels (c) and (d) repeat the exercise in general equilibrium. Consistent with our findings on discount rate shocks, we find that general equilibrium forces do not undo the partial equilibrium effects, but dampen the responses.

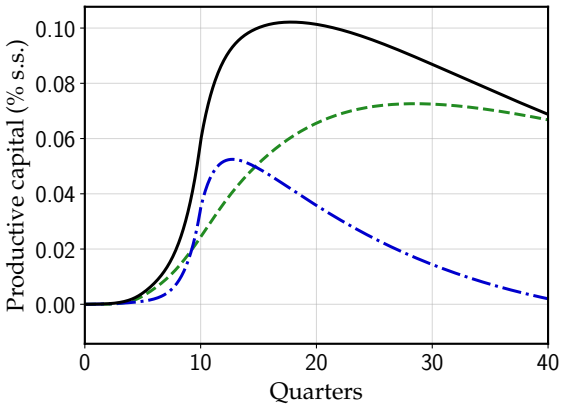
The investment IRFs highlight our first policy takeaway. The cumulative response of investment to tax credits comes from project starts (Proposition 7). For ongoing projects, temporary tax credits only reshuffle expenditures over time. This is not to say that investment tax credits are useless. Quite the contrary, our results suggest they give policymakers an additional lever over the timing of investment demand. If aggregate demand is currently low,

[49] For illustration, Appendix Figure A12 shows the effect of a (perfect foresight) AR(1) shock to p (i.e., a "cash flow" shock). When terminal payoffs rise, firms rush to complete projects, spurring an investment spike driven by both ongoing and newly started projects. Investment tax credits usually operate by lowering costs rather than increasing terminal payoffs, but the exercise clearly illustrates how higher cash flows induce acceleration and thus higher investment.

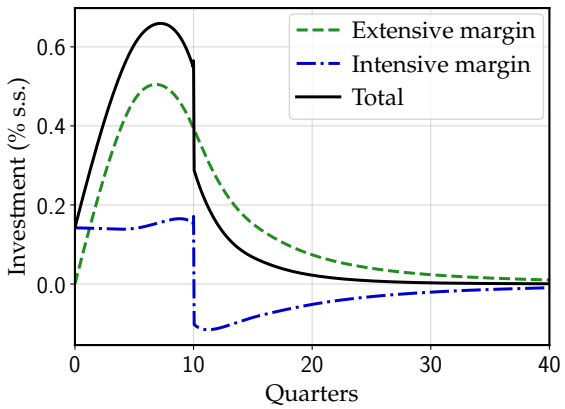
Figure 7: Endogenous Time to Build and Investment Tax Credits



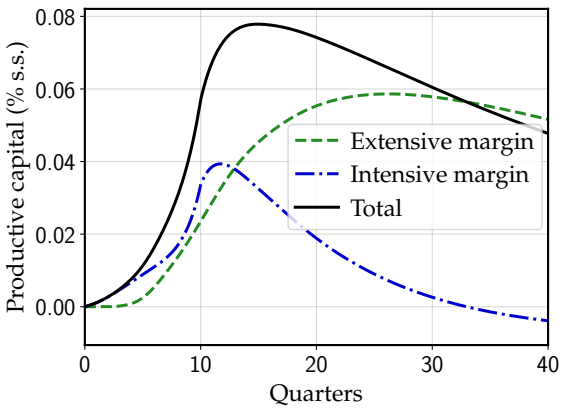
(a) IRF Decomposition: I^{acc} (Partial Equilibrium)



(b) IRF Decomposition: K (Partial Equilibrium)



(c) IRF Decomposition: I^{acc} (General Equilibrium)



(d) IRF Decomposition: K (General Equilibrium)

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. The y-axis measures percent-deviations from steady state values. Panels a and b show partial equilibrium responses. The black solid line captures the total IRF; 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). Panels c and d repeat the exercises in panels a and b, but allowing for general equilibrium responses. Remaining parameters are as in Appendix Table A24.

we can spur investment today by borrowing it from the future. The duration and the size of the tax credit provide two levers for policymakers to intertemporally reshuffle 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".[50] Fiscal policy does share 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 mostly focused on investment tax credits as a tool for stabilization policy. Investment subsidies can speed up structural change. The Inflation Reduction Act and the CHIPS Act of 2022 had explicit rationales linked to accelerating the green transition and the reshoring of strategic industries, respectively. Our model clarifies how these policies trade off speed gains with acceleration costs; and provides a portable framework to quantify such costs and benefits. This is a promising application for future research.

6.5 State-Dependence

Recall from Figure 4 that the effects of a discount rate shock on investment flows are heterogeneous. Projects that are closer to (farther from) completion are sped up (slowed down) – a model prediction that the data bear out in Section 3.4. 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 8 studies these settings. We operate in partial equilibrium to isolate the theoretical channel. In Panel (a), we plot three distributions of projects over stages. The first is the steady state distribution (black solid line, Figure 5). 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, so the distribution skews towards old projects (blue dot-dashed line).

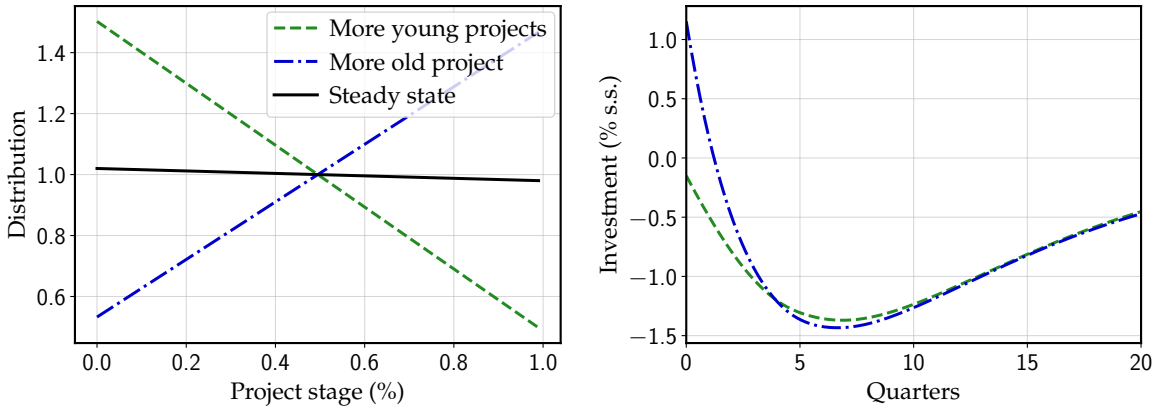
Panel (b) plots the response of aggregate investment to the discount rate shock of Figure 5, but starting from each of the skewed distributions. When the economy starts filled with

[50] While q-theory models with convex adjustment costs (Hayashi, 1982) can also generate our first policy takeaway, they do not generate this second implication, because in those models the response does not depend on how far along ongoing projects are.

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 sign depending on initial conditions.[51]

Figure 8 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). Here the state is the pipeline of projects under construction, which depends on past starts and hence on past shocks and policy. 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 policy may face headwinds that are absent when time to build is exogenous.

Figure 8: 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 of the shock. The IRF for the steady state distribution is the black dashed line in Panel (a) of Figure 5. Shock persistence, shock size, and remaining parameters as in Appendix Table A24.

[51] In CapexDx, the maturity of the project pipeline varies over time. Recalculating our maturity measures for every live project from 2006 to 2019 with rolling five-year versions of the survival model of Appendix D.1, we find that the average probability of completing within two years ranges from 15 to 23 percent (15 to 21.5 percent when completion rates are estimated only on data before each date).

7 Conclusion

We show that time to build is endogenous and that this endogeneity is consequential for investment. When firms can choose how fast to complete projects, the standard capital budgeting intuition that tighter financing must lower investment no longer holds. An increase in discount rates can raise project-level investment, the effect of monetary policy on aggregate investment becomes state-dependent, and investment subsidies gain a new intertemporal reshuffling channel. We provided quasi-experimental evidence for this mechanism using detailed project histories from India, and rationalized our findings in a neoclassical investment model. We showed that this endogeneity has distinctive implications for the economics of investment and capital formation, both in partial and general equilibrium. 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. In our model, investment in ongoing projects responds strongly to the timing of cash flows, opening up space for time-varying investment tax 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., execution risk or asymmetric information). But there is more to build, and questions are not in short supply. Take for instance the artificial intelligence buildup. Do firms internalize gestation lags of their competitors when expanding capacity, and how does that shape industry dynamics and competition? Can we structure financing practices that internalize endogenous duration risk? How much can targeted policies speed up infrastructure development? Can our insights from fixed capital formation be generalized to breakthrough innovations or other intangible forms of capital? Future research could further tease out 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.

Abel, A. B. (1981). A dynamic model of investment and capacity utilization*. *The Quarterly Journal of Economics*, 96(3):379–403.

Abel, A. B. and Eberly, J. C. (1994). A unified model of investment under uncertainty. *The American Economic Review*, 84(5):1369–1384.

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.

Aïd, R., Federico, S., Pham, H., and Villeneuve, B. (2015). Explicit investment rules with time-to-build and uncertainty. *Journal of Economic Dynamics and Control*, 51:240–256.

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

Anderson, J. R., Cain, K. C., and Gelber, R. D. (1983). Analysis of survival by tumor response. *Journal of Clinical Oncology*, 1(11):710–719.

Angrist, J. D. and Pischke, J.-S. (2009). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press, Princeton, NJ.

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.

Auerbach, A. J. and Gorodnichenko, Y. (2012). Measuring the output responses to fiscal policy. *American Economic Journal: Economic Policy*, 4(2):1–27.

Baily, M. N. (1982). Comment on "The Swedish Investment Funds System as a Stabilization Policy Rule" by John B. Taylor. *Brookings Papers on Economic Activity*, 1982(1):100–102.

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.

Bertolotti, F. and Lanteri, A. (2026). Capital replacement and innovation dynamics. *Working Paper*.

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

Böhm-Bawerk, E. (1891). *The Positive Theory of Capital*. Macmillan.

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

Brainard, W. C. (1977). Comment on "Investment, Interest Rates, and the Effects of Stabilization Policies" by Robert E. Hall. *Brookings Papers on Economic Activity*, 1977(1):112–117.

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

Caballero, R. J. (1999). Aggregate investment. In Taylor, J. B. and Woodford, M., editors, *Handbook of Macroeconomics*, volume 1, chapter 12, pages 813–862. Elsevier, Amsterdam.

Caballero, R. J. and Engel, E. M. R. A. (1999). Explaining investment dynamics in u.s. manufacturing: A generalized (s,s) approach. *Econometrica*, 67(4):783–826.

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

Chen, J. and Roth, J. (2024). Logs with zeros? some problems and solutions. *The Quarterly Journal of Economics*, 139(2):891–936.

Chetty, R. (2008). Moral hazard versus liquidity and optimal unemployment insurance. *Journal of Political Economy*, 116(2):173–234.

Chodorow-Reich, G. (2014). 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., Gopinath, G., Mishra, P., and Narayanan, A. (2020). Cash and the Economy: Evidence from India's Demonetization. *The Quarterly Journal of Economics*, 135(1):57–103.

Chopra, Y., Subramanian, K., and Tantri, P. L. (2021). 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.

Cloyne, J., Ferreira, C., Froemel, M., and Surico, P. (2023). Monetary policy, corporate finance, and investment. *Journal of the European Economic Association*, 21(6):2586–2634.

CMIE (2023). Capexdx: Database of the life cycle of new investments. 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.

Cooper, R. W. and Haltiwanger, J. C. (2006). On the nature of capital adjustment costs. *The Review of Economic Studies*, 73(3):611–633.

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

Crouzet, N. (2018). Aggregate implications of corporate debt choices. *The Review of Economic Studies*, 85(3):1635–1682.

Cummins, J. G., Hassett, K. A., and Hubbard, R. G. (1994). A reconsideration of investment behavior using tax reforms as natural experiments. *Brookings Papers on Economic Activity*, 1994(2):1–59.

Doms, M. and Dunne, T. (1998). Capital adjustment patterns in manufacturing plants. *Review of Economic Dynamics*, 1(2):409–429.

Eberly, J., Rebelo, S., and Vincent, N. (2012). What explains the lagged-investment effect? *Journal of Monetary Economics*, 59(4):370–380.

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.

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

Gallo, G. (2026). Data centers: Speed-to-market and the risks behind the race. Technical report, Construction Lender Risk Management Roundtable, https://www.construction-lender-risk-management.org/data-centers-speed-to-market-and-the-risks-behind-the-race/. Last accessed Oct 2026.

Gertler, M. and Karadi, P. (2011). A model of unconventional monetary policy. *Journal of Monetary Economics*, 58(1):17–34. Carnegie-Rochester Conference Series on Public Policy: The Future of Central Banking April 16-17, 2010.

Ghaffary, S., Ford, B., and Chang, E. (2025). Inside the first stargate ai data center. Technical report, Bloomberg Businessweek, https://www.bloomberg.com/news/features/2025-05-20/inside-stargate-ai-data-center-from-openai-and-softbank. Last accessed Oct 2026.

Gilchrist, S. and Zakrajšek, E. (2012). Credit spreads and business cycle fluctuations. *American Economic Review*, 102(4):1692–1720.

Glancy, D., Kurtzman, R. J., and Loewenstein, L. (2025). Shovel ready projects and commercial construction activity's long and variable lags. *Finance and Economics Discussion Series (FEDS)*.

Goldberg, A. (1992). *The Pentagon: The First Fifty Years*. Government Printing Office.

Gomtsyan, D. (2024). Building the city under financial frictions. *Review of Economic Dynamics*, 52:70–83.

Gormsen, N. J. and Huber, K. (2025). Corporate discount rates. *American Economic Review*, 115(6):2001–2049.

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

Graham, J. R. (2022). Presidential address: Corporate finance and reality. *Journal of Finance*, 77(4):1975–2049.

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.

Hall, R. E. and Jorgenson, D. W. (1967). Tax policy and investment behavior. *The American Economic Review*, 57(3):391–414.

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.

Hayashi, F. (1982). Tobin's marginal q and average q: A neoclassical interpretation. *Econometrica*, 50(1):213–224.

House, C. L. (2014). Fixed costs and long-lived investments. *Journal of Monetary Economics*, 68(C):86–100.

House, C. L. and Shapiro, M. D. (2008). Temporary investment tax incentives: Theory with evidence from bonus depreciation. *American Economic Review*, 98(3):737–68.

Jeenas, P. (2025). Firm balance sheet liquidity, monetary policy shocks, and investment dynamics. *Journal of Political Economy*. Forthcoming.

Jorgenson, D. W. (1963). Capital theory and investment behavior. *The American Economic Review*, 53(2):247–259.

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.

Katz, L. F. and Meyer, B. D. (1990). Unemployment insurance, recall expectations, and unemployment outcomes. *The Quarterly Journal of Economics*, 105(4):973–1002.

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.

Koby, Y. and Wolf, C. (2020). Aggregation in heterogeneous-firm models: Theory and measurement. *Working Paper*.

Kroft, K. and Notowidigdo, M. J. (2016). Should unemployment insurance vary with the unemployment rate? theory and evidence. *The Review of Economic Studies*, 83(3):1092–1124.

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

Lagarde, C. (2026). A new age of capital: Growth, sovereignty and AI. Technical report, European Central Bank, `https://www.ecb.europa.eu/press/key/date/2026/html/ecb.sp260914_2~a3f0efbee4.en.html`. Speech at Hofburg im Dialog, Vienna, 14 September 2026. Last accessed Oct 2026.

Lanteri, A. (2018). The market for used capital: Endogenous irreversibility and reallocation over the business cycle. *American Economic Review*, 108(9):2383–2419.

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*.

Lucas, Jr., R. E. (1967). Adjustment costs and the theory of supply. *Journal of Political Economy*, 75(4):321–334.

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

Majd, S. and Pindyck, R. S. (1987). Time to build, option value, and investment decisions. *Journal of Financial Economics*, 18(1):7–27.

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

McDonald, R. and Siegel, D. (1986). The value of waiting to invest. *The Quarterly Journal of Economics*, 101(4):707–727.

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.

Microsoft Corporation (2026). Earnings conference call, fiscal year 2026 third quarter. Technical report, Microsoft Investor Relations, https://www.microsoft.com/en-us/investor/events/fy-2026/earnings-fy-2026-q3. Remarks by Amy Hood, Executive Vice President and Chief Financial Officer, 29 April 2026. Last accessed Oct 2026.

Millar, J. N., Oliner, S. D., and Sichel, D. E. (2016). Time-to-plan lags for commercial construction projects. *Regional Science and Urban Economics*, 59(C):75–89.

Moffitt, R. (1985). Unemployment insurance and the distribution of unemployment spells. *Journal of Econometrics*, 28(1):85–101.

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

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

Oh, H., Yang, C., and Yoon, C. (2024). Land development and frictions to housing supply over the business cycle. *Finance and Economics Discussion Series*.

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.

Ottonello, P. and Winberry, T. (2020). Financial heterogeneity and the investment channel of monetary policy. *Econometrica*, 88(6):2473–2502.

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

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

Ramey, V. and Shapiro, M. (2001). Displaced capital: A study of aerospace plant closings. *Journal of Political Economy*, 109(5):958–992.

Ramsey, M. (2016). Tesla races to finish 'gigafactory' in time for model 3 rollout. Technical report, WSJ, https://www.wsj.com/articles/tesla-races-to-finish-gigafactory-in-time-for-model-3-rollout-1469404142.

Rollo, C. (2024). 10 things you probably didn't know about the pentagon. https://www.defense.gov/News/Feature-Stories/story/Article/1650913/10-things-you-probably-didnt-know-about-the-pentagon/.

Rothstein, J. (2011). Unemployment insurance and job search in the great recession. *Brookings Papers on Economic Activity*.

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

Schaal, E. and Taschereau-Dumouchel, M. (2025). Echoes and delays: Time-to-build in production networks. *Working Paper*.

Schmieder, J. F. and von Wachter, T. (2016). The effects of unemployment insurance benefits: New evidence and interpretation. *Annual Review of Economics*, 8(Volume 8, 2016):547–581.

Selgrad, J. and Siani, K. (2025). Monetary policy and investment plans. *Working Paper*.

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

Summers, L. H. (1981). Taxation and corporate investment: A q-theory approach. *Brookings Papers on Economic Activity*, 1981(1):67–127.

Taylor, J. B. (1982). The swedish investment funds system as a stabilization policy rule. *Brookings Papers on Economic Activity*, 1982(1):57–99.

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.

van Houwelingen, H. C. (2007). Dynamic prediction by landmarking in event history analysis. *Scandinavian Journal of Statistics*, 34(1):70–85.

Vogel, S. (2007). *The Pentagon – A History: The Untold Story of the Wartime Race to Build the Pentagon and to Restore it Sixty Years Later*. Random House.

Warsh, K. (2026). Transcript of chairman warsh's press conference, june 17, 2026. Technical report, Board of Governors of the Federal Reserve System, https://www.federalreserve.gov/mediacenter/files/FOMCpresconf20260617.pdf. Last accessed Oct 2026.

Whited, T. M. (2006). External finance constraints and the intertemporal pattern of intermittent investment. *Journal of Financial Economics*, 81(3):467–502.

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

Wooldridge, J. M. (2010). *Econometric Analysis of Cross Section and Panel Data*. MIT Press, Cambridge, MA, 2nd edition.

Zwick, E. and Mahon, J. (2017). Tax policy and heterogeneous investment behavior. *American Economic Review*, 107(1):217–48.

Online Appendix (Not for Publication)

A Data Appendix

A.1 CapexDx

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

A.1.1 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 hundred million rupees in India. This does not constitute a project: we do not 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 headquarters 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 100 million rupees located in Chennai.

A.1.2 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.

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 suggested they actively use Capex to either track existing projects or make strategic decisions regarding future capacity creation.

Our conversations with the data provider further alleviate 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 parlance of survival analysis, 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.

A.1.3 Raw Data Structure

Before we start describing how we construct event histories for our econometric exercises, it is convenient to describe the individual files we start from. All such files can be obtained with a subscription to CapexDx. We use the March 31st 2022 vintage for our analyses.

- Project Detail file: cross-sectional data on project characteristics, such as project name and id, sponsor name and id, state and district of project, most recent project status, etc.
- Project Costs file: data at the project-cost estimate level, including date of cost estimate. Projects may have multiple cost estimates over their lifetime, which we refer to as "restatements".
- Project Completion file: cross-sectional data for completed projects, specifying dates of completion and previous estimates of completion dates.
- Project Events file: data at the project-event-event date level, tracking the project lifecycle.

A.2 Constructing Event Histories

The ability to track events at the project level is a noteworthy feature of CapexDx. However, pre-processing is required before any meaningful econometric analysis can be carried out. This is mostly because of the structure of the raw data: the event histories contain an "event" for each row. This means that multiple events can happen on the same date, which is often the case.[52] For running our hazard models, however, we must construct a panel at

[52]For instance, there might be three entries on 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.

the project-calendar time level.

We therefore reduce the dimensionality of the event-level data. Our strategy is to use the event-level data to construct a project-level panel that tracks whether the project is in one of the following stages: announced, shelved, stalled, under implementation, completed, abandoned, or lost to follow-up. These correspond to the labels used for project statuses in the Project Detail file.[53] Completion, abandonment/shelving, and lost to follow-up are terminal (or, absorbing) stages. All remaining stages correspond to "intermediate stages", from which a project may transition in or out. For example, though all projects start as "announced", a project might temporarily move into "stalled", then back into "announced", and then become "under implementation".[54]

Below we describe the criteria we use to characterize a project over time in each one of these stages. We implement each step sequentially, carefully browsing and visualizing the data as we implement each step, and recurrently checking individual project histories to ensure accuracy. We start with the terminal stages:

- *Completed*: any event date that occurs at or after we observe one of the events "Completed" or "Commercial Production". In cases when "Completed" or "Commercial Production" both appear in the event history but do not share the same event date, we consider the last event as marking completion. The project is marked as "Completed" in the Detail file.

- *Abandoned (Shelved):* any event date that occurs at or after we observe the event "Abandoned on" ("Shelved on").[55] The project is marked as "Abandoned" in the Detail file.

- *Lost to Follow-Up:*

[53]Sadly, CapexDx does not have a variable keeping track of these stages over time in the event file, which would have made our wrangling much easier. The stage variable in the detail file corresponds only to the most recent project stage as of the vintage release.

[54]We first do some pre-processing of the events data to simplify our wrangling. First, we drop some events that are not informative for our purposes. Many events in the data are "Company contacted successfully" or "Company contacted but no response". These events can be used to measure the intensity at which CapexDx tracks a project, but do not convey any information about the project by themselves. We first remove these rows from the events dataset and manipulate them in a separate dataset to obtain monitoring intensity measures. To be precise, we drop such events if either (1) there is another event on the same date that is not a monitoring marker, or (2) they occur by themselves, but are neither the first nor last event in the project's history; or if (3) they are the last observations in the project history, but the preceding event in the project history is either completion, abandonment, or shelving. Second, some projects feature multiple subprojects or phases. Some events relate to particular phases (often marked by containing the substring "for pr/s/u/ph" in the event description). We remove these events as CapexDx also creates event markers that apply to the entire project.

[55]In a minuscule number of cases, the project flows in and out of abandonment or shelving. We contacted CMIE to inquire about these cases. Indeed, these are projects where, ex post, CMIE acquired information that the project sponsor reversed the decision and continued implementation of the project. In these cases, we treat abandonment as an intermediate rather than terminal stage. Our results are not sensitive to whether such projects are included.

- If there is only one event “No information as on” throughout the event history: any event date that occurs at or after we observe such event. The project must be marked as “No Information” in the Detail file.
- If there is more than one event “No information as on” throughout the event history: periods of temporary loss to follow-up can be identified by the initial events “No Information as on” and the terminal events “No information as up to”. For the purposes of understanding completion dynamics, we are not interested in "No information" as a time-varying state. If this "no information" will be succeeded by some other event, we take the last-occurring “No Information as on” marker as when the project becomes censored; and assume that whatever state the project was in before this no information period is the state in which the project remained during this period.[56]

For the remaining stages, a project may transition in and out of a particular stage. We thus have to keep track not only of markers indicative of entry, but also of markers indicative of exit.

- *Stalled:* a project is marked as stalled during the following time periods:
  - Event dates between the entry event “Implementation stalled on” and the exit event “Stalled up to” or any exit events marking transition into a terminal stage (abandonment, completion, loss to follow-up, or shelving).[57]
  - Event dates between the entry event “Announced and stalled on” and the exit event “Stalled up to” or any exit events marking transition into a terminal stage (abandonment, loss to follow-up, or shelving).[58]
  - Event dates starting at the entry event “Implementation stalled on” or “Announced and stalled on”, but without a transition into a terminal stage. For such cases, we manually check the histories of each project. These histories normally feature many monitoring markers such as “Company contacted successfully” or “Company contacted but no response”, or are the last entries in a project’s event history. These projects are more recent and still actively stalled.

[56] Granted, this does not deal with potential cases where the "no information" period is a period where the project went through a different state (e.g., if it was temporarily stalled during this non-information spell, but the data provider could not obtain such information). However, we cannot circumvent this non-observability; and, except for analyzing transition matrices, this data decision will not influence our results.

[57] There are 16 projects in the entire dataset that move from stalled to completed. We cannot disambiguate whether this means the project was completed to begin with, but its operations were put on hold, or whether the data provider was unable to ascertain at which date implementation resumed.

[58] There are no projects in the entire dataset that move to completion after this initial entry event. The marker “Announced and stalled on” is used by the data provider to indicate a project where the first contact with the project occurs at a point where the project is already stalled.

- *Under Implementation:* a project is marked as under implementation during the following time periods:
  - Event dates between the entry event "Implementation started" and the exit event "Implementation stalled on" or any exit events marking transition into a terminal stage (abandonment, completion, loss to follow-up, or shelving).[59]
  - Event dates starting at the entry event "Implementation stalled on" or "Stalled up to", but without a transition into a terminal stage. For such cases, we manually check the histories of each project. These histories normally feature many monitoring markers such as "Company contacted successfully" or "Company contacted but no response", or are the last entries in a project's event history. These projects are more recent and still in active implementation.
- *Announced:* a project is marked as announced during the following time periods:
  - Projects that were never started: event dates between the entry event "Date of Announcement" and the exit events "Announced and stalled on", "Shelved on".
  - Projects that are eventually started: event dates between the entry event "Date of Announcement" and the exit events "Announced and stalled on", "Implementation started", or any exit events marking transition into a terminal stage (abandonment, completion, or loss to follow-up).[60]
  - Very recent projects: event dates between the entry event "Date of Announcement", but without a transition into another stage. For such cases, we manually check the histories of each project. These histories normally feature many monitoring markers such as "Company contacted successfully" or "Company contacted but no response", or are the last entries in a project's event history. These projects are more recent and we simply do not know what will happen to them.

Once we construct these time-varying stages, we ensure the resulting histories are sensible. We implement the following assertion tests:

[59] A project might have more than one "Implementation started" marker. These cases correspond to multiple spells of implementation within a single history, often marked by temporary stalls. In these cases, we mark each spell separately so as to adequately capture the multiple spells in the history. In rare cases of multiple implementation spells, the entry event is "Stalled up to", unaccompanied by the "Implementation started" event but with a preceding "Implementation started" event earlier in the project history. For such cases, we manually check the histories of each project.

[60] A project might go through multiple spells of having been announced, but either not having started to be implemented yet (e.g., licenses are being procured), or being temporarily stalled. In these cases, we mark each spell separately so as to adequately capture the multiple spells in the history. In rare cases of multiple "announce and hold" spells, the entry event is "Stalled up to", unaccompanied by the "Date of announcement" event but with a preceding "Date of announcement" event earlier in the project history. For such cases, we manually check the histories of each project.

1. At any point in time, a project is assigned to one of the above stages.

2. At any point in time, a project is assigned to a single stage (i.e., a project cannot be simultaneously under implementation and stalled, or completed and lost to follow-up).[61]

At this point, a project's event history is effectively constructed. We now only need to convert the event date-event-project level dataset into a project-time level dataset that can be used for conventional analyses. In a first step, we convert the data to a project-day level data. Then, we aggregate to the quarterly frequency.[62] Last, to estimate transition probabilities and hazard models, it is important for the data to feature a particular structure: the last row should capture the transition into the terminal stage. After doing so, we obtain our final quarterly panel on project histories, to which we can merge any time-varying or cross-sectional information on projects, their sponsors, their geographies, etc.

Note that we could have simply characterized the project as "completed" or "not completed" for the purposes of our main regression exercises. However, we wanted to ensure that the project history was sensible, i.e., indeed transitions across different stages painted a coherent story of the project's lifecycle. This was a key motivation for engaging in the full-scale wrangling exercise and producing Appendix Table A3.

A.2.1 Examples

To illustrate the cleaning procedure, we consider a fictitious project with multiple events. Note that to preserve the confidentiality of the data provider's raw data, this is a fictitious project, merely for illustration of how we can go from raw event-level data to project-level data.

The first example (project ID 1) illustrates a straightforward case: the project is announced, starts implementation, and is eventually completed. In this example, we suppress events that are not essential for our classification of time-varying stages (e.g., monitoring events, markers for initial estimates of completion dates, etc) to avoid clutter. For the same reason, we omit columns for stages that are all zeros over the project's lifetime.

[61] In a minuscule number of cases, we find that our procedure produced overlapping stages. We manually disambiguated the event histories in such cases by looking at the entire project history and determining which stage made most sense and whether the entire project history looked sensible.

[62] We compute the number of events per month and per quarter before aggregating, and do not find considerable discrepancies, which alleviates concerns that we are losing key information about the project's history by aggregating. Whenever there is more than one stage in a particular quarter, we follow standard procedure in survival analysis and take the last stage as the stage for the entire quarter. For example, if the project moves from announcement to implementation in the middle of 2011:2, the project is considered "under implementation" for the entirety of 2011:2.

Event-Time Panel

ID	Event	Event Date
1	Date of announcement	05/20/2006
⋮	⋮	⋮
1	Implementation Started	09/02/2007
⋮	⋮	⋮
1	Completed	07/13/2012

→

Project-Time Panel

ID	Calendar	Announced	Under Imp	Completed
1	2006:2	1	0	0
⋮	⋮	⋮	⋮	⋮
1	2007:2	1	0	0
1	2007:3	0	1	0
⋮	⋮	⋮	⋮	⋮
1	2012:2	0	1	0
1	2012:3	0	0	1

The second example illustrates a more nuanced case: the project becomes temporarily stalled, and later on continues to be implemented. However, the project is then stalled again. The data provider fails to obtain more information on the project from the company, and eventually the project is marked as lost to follow-up. This example illustrates how projects may go in and out of particular stages of implementation.

In this example, we include the monitoring markers provided by CapexDx. This serves two purposes. First, it illustrates the intense monitoring efforts of the data provider; second, it illustrates how loss to follow-up occurs when the data provider cannot ascertain what is happening to the project after successive attempts to obtain information on the project status for a given period of time.

Event-Time Panel

ID	Event	Event Date
2	Date of announcement	11/05/2005
2	Contacted Successfully	08/01/2006
⋮	⋮	⋮
2	Land acquired	11/11/2008
2	Implementation started	11/11/2008
2	Contacted Successfully	11/11/2008
⋮	⋮	⋮
2	Implementation stalled on	04/20/2009
2	Contacted Successfully	04/20/2009
⋮	⋮	⋮
2	Stalled up to	04/01/2010
2	Construction commenced	04/01/2010
2	Contacted Successfully	04/01/2010
⋮	⋮	⋮
2	Implementation stalled on	08/21/2012
2	Contacted Successfully	08/21/2012
⋮	⋮	⋮
2	Stalled up to	10/09/2012
2	Contacted Successfully	10/09/2012
⋮	⋮	⋮
2	No response	02/09/2013
2	No response	07/10/2013
2	No response	12/01/2013
2	No response	04/13/2014
2	No information as on	06/27/2014

→

Project-Time Panel

ID	Calendar	Announced	Under Imp	Stalled	No Info
2	2005:4	1	0	0	0
⋮	⋮	⋮	⋮	⋮	⋮
2	2008:3	1	0	0	0
2	2008:4	0	1	0	0
⋮	⋮	⋮	⋮	⋮	⋮
2	2009:1	0	1	0	0
2	2009:2	0	0	1	0
⋮	⋮	⋮	⋮	⋮	⋮
2	2010:1	0	0	1	0
2	2010:2	0	1	0	0
⋮	⋮	⋮	⋮	⋮	⋮
2	2012:2	0	1	0	0
2	2012:3	0	0	1	0
2	2012:4	0	1	0	0
⋮	⋮	⋮	⋮	⋮	⋮
2	2014:1	0	1	0	0
2	2014:2	0	0	0	1

A.3 ProwessDx and Charges Data

We obtain annual and quarterly-level accounting data from ProwessDx, and merge any relevant variables with our project-level dataset at the firm-time level. We obtain charge-level data from the Ministry of Corporate Affairs. We collect data on entities whose CIN codes are

present in ProwessDx. As described in the main text, we undertake a combination of fuzzy matching algorithms and extensive manual checks to arrive at a crosswalk between the name of the lending entities in the MCA data and the names of companies or banks in ProwessDx. Once we have the names of the lending entities, we can obtain their Prowess company codes to ensure a smooth matching with other data. For the borrowing entities, matching with Prowess data is straightforward since we have their CIN codes in both datasets.

In order to arrive at a bank-firm-time panel, we first expand the data to get a panel at the loan-time level; and then compute any relevant variables (e.g., total amount outstanding in a year) at the appropriate bank-firm-time level. To move from a bank-firm-time to a firm-time or a bank-time level dataset, we again compute any relevant variables before collapsing the data. Relative to the project-level dataset, wrangling of these data is more straightforward.

Table A1: Concentration Statistics

	N	Mean	SD	10th	25th	50th	75th	90th
Projects per Company	32456	3.29	18.99	1.00	1.00	1.00	2.00	4.00
Share: Section	16	0.06	0.09	0.01	0.01	0.02	0.05	0.23
HHI: Section	16	0.19						
Share: Industry	134	0.01	0.02	0.00	0.00	0.00	0.01	0.02
HHI: Industry	134	0.04						
Share: State	42	0.02	0.03	0.00	0.00	0.01	0.04	0.05
HHI: State	42	0.06						

Notes: This table reports concentration measures for all projects available in the March 2022 vintage of CapexDx. The number of states includes Union Territories, and exceeds the current number of administrative units due to changes in administrative boundaries over time.

B Additional Measurement Exercises

Table A2: 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

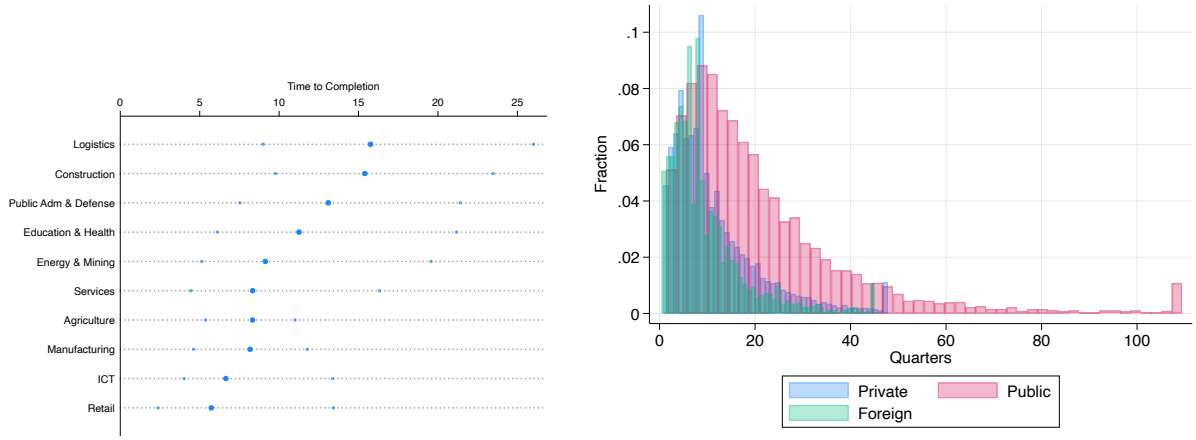
Notes: This table constructs a transition matrix across project stages for all projects available in the March 2022 vintage of CapexDx. Rows represent departing states. We only include non-terminal stages, i.e., we exclude completion, abandonment, and loss to follow-up. Columns denote arriving states, 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 A3: 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 return 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.

Figure A1: Disaggregating the Time to Build Distribution



(a) Medians and IQR of Gestation Lags by Sector (b) Distribution of Gestation Lags by Ownership

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 exclude 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.

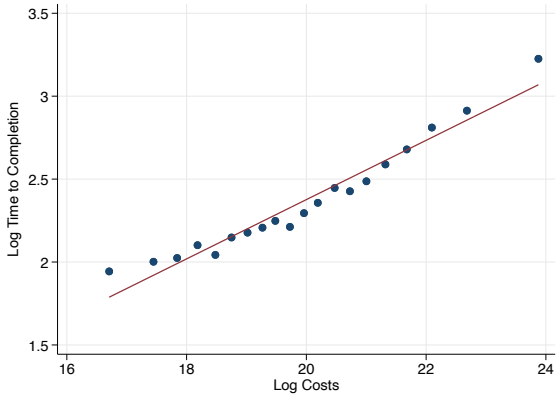
Table A4: Cost-Time Schedules

	(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			✓		✓

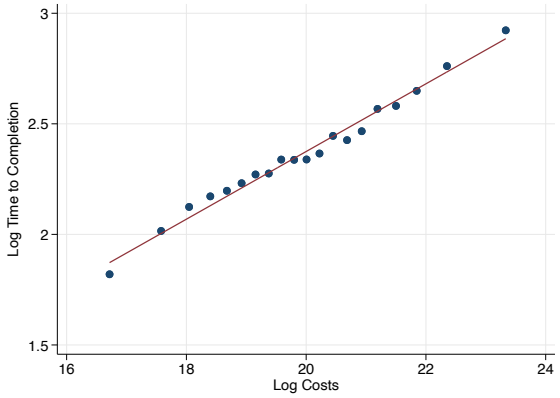
Notes: This table reports correlations between log time to build and log costs estimated with linear regression from the microdata underlying the binned scatter plots of Figure A2. The sample consists of completed projects for which cost data is available. Column (1) reports the bivariate relationship. Column (2) includes sector fixed effects, and column (3) adds year started fixed effects. Columns (4) and (5) replace sector fixed effects with industry fixed effects. Standard errors are clustered at the industry level and reported in parentheses.

Figure A2: Cost-Time Schedules

(a) Bivariate Correlation



(b) Within-Industry Correlation



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.

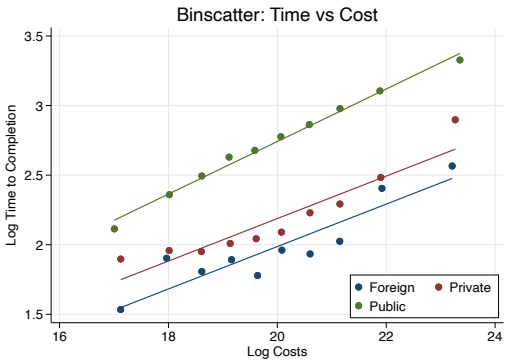
Table A5: Cross-Sectional Variance Decompositions: Time to Completion

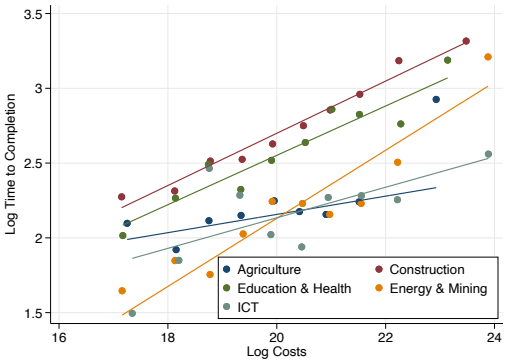
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
R^2	0.239	0.302	0.285	0.395	0.409	0.410	0.396	0.441	0.579
Unc. Var					0.014	0.015	0.001	0.046	
Partial R^2					0.023	0.024	0.002	0.077	
Ownership FE						✓		✓	
State FE					✓			✓	
Sector FE	✓			✓	☒	☒	☒	☒	
Industry FE		✓							
Year Started FE			✓	✓	☒	☒	☒	☒	✓
Type 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 power 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.

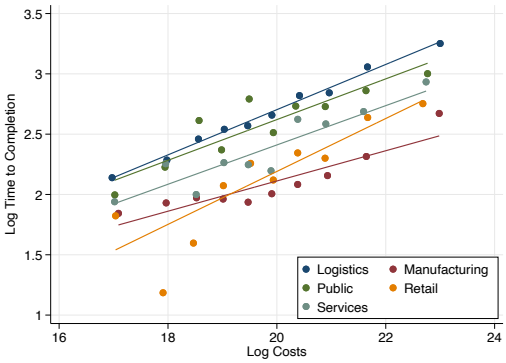
Figure A3: Cost-Time Schedule by Subgroups

(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 A2) by sponsorship (panel a) and sectors (panels b and c, split for the sake of visualization).

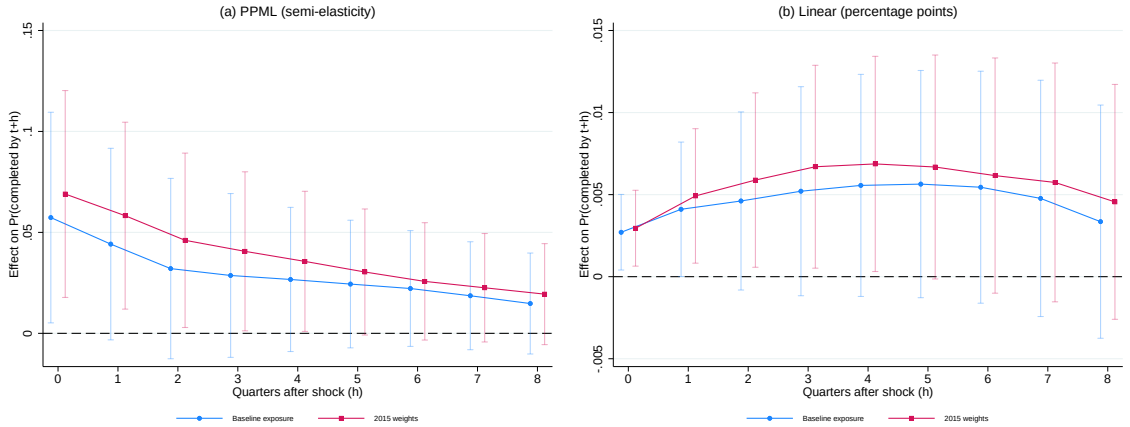
Table A6: Effects of Asset Quality Review on Credit Supply

	(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 (1) and table 2). 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 percentiles, 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.

C Additional Empirical Results

Figure A4: Effects on Cumulative Project Completion.



Notes: the figure plots local-projection estimates of the effect of firm-level exposure on an indicator equal to one if the project has been completed by quarter $t+h$, for horizons $h = 0, \ldots, 8$. Panel (a) estimates the cumulative completion probability by PPML, as in Table 2. Coefficients are semi-elasticities relative to a baseline completion probability that grows mechanically with the horizon, which is why they decline as h increases. Panel (b) estimates a linear probability model, with effects in percentage points. Both panels use the specification of column (5) of Table 2: firm fixed effects, sector-by-year and state-by-year fixed effects, and firm-level controls. Coefficients at negative horizons are not defined because completion is an absorbing state; only forward horizons are reported. 95% confidence intervals based on standard errors clustered at the firm level.

Table A7: 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			✓			✓

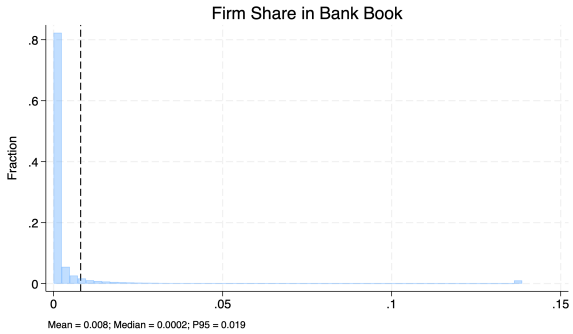
Notes: the table reports estimates of β from the following model:

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

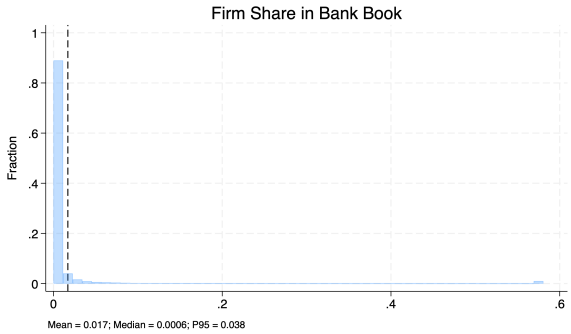
where ΔLoans_{bft} is the growth rate of loans between $t-1$ and t, constructed as in Davis and Haltiwanger (1992). The sample consists of firm-bank relationships for 2013-2019, conditional on an existing relationship (positive outstanding loans) in the 2013-2015 period. Exposure is defined as the difference 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 (1) and table 2). Columns (2) and (5) include borrower-by-year fixed effects. Columns (1) and (4) 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 percentiles, 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.

Figure A5: Distribution of Firm Shares of Bank Charges

(a) Universe of Charges



(b) Firms in Hazard Regression



Notes: the figures plot the distribution of firm shares in bank charges in the MCA registry. We compute $s_{bft} = {}^{Loans_{bft}}/_{\sum_{f'} Loans_{bf't}}$ for both the universe of charges (left panel) and for firms that eventually make it into our regression sample (right panel). The statistics at the bottom represent the mean, median and 95th percentile of the distribution. Histograms are winsorized at the 99th percentile for visibility.

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

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Exposure (Pre)	0.1057***	0.0772***	0.0766***	0.0754***	0.0749***	0.0844***	0.0845***
	(0.0277)	(0.0273)	(0.0273)	(0.0267)	(0.0267)	(0.0265)	(0.0266)
Exposure (2015)	0.1067***	0.0701***	0.0698***	0.0691***	0.0690***	0.0823***	0.0830***
	(0.0275)	(0.0270)	(0.0269)	(0.0261)	(0.0261)	(0.0265)	(0.0266)
DepVar Mean	0.0397	0.0397	0.0397	0.0397	0.0397	0.0414	0.0414
Observation	122123	122123	122123	122123	122123	116788	116788
Age Splines	✓	✓	✓	✓	✓	✓	✓
Firm FE		✓	✓	✓	✓	✓	✓
State x Year FE		✓	✓	✓	✓	✓	✓
Sector x Year FE				✓	✓		
Industry x Year FE						✓	✓
Firm Controls			✓		✓		✓

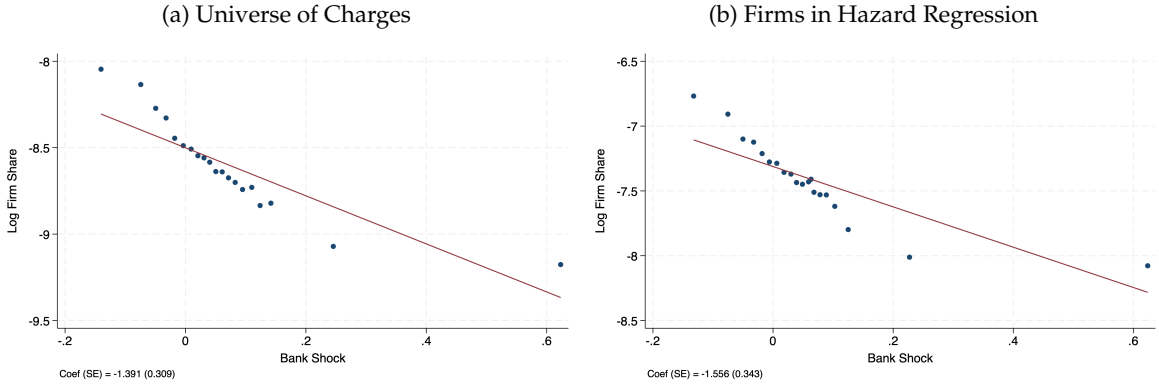
Notes: the table reports PPML estimates of the semi-elasticity of project-level hazards to firm exposure to the Asset Quality Review (see model (1)). The sample consists of projects observed quarterly between 2013 and 2019. Each row corresponds to a measure of firm exposure. Firm exposure is defined 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 10.57%. 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 percentiles, and then standardized. All firm controls are lagged by one period. The reduction in sample size from column (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 A9: Effects of Firm-Level Credit Supply Shocks on Project Completion Hazards

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Exposure	0.0057***	0.0030**	0.0029**	0.0028**	0.0027**	0.0029***	0.0028**
	(0.0017)	(0.0012)	(0.0012)	(0.0012)	(0.0012)	(0.0011)	(0.0011)
DepVar Mean	0.0397	0.0397	0.0397	0.0397	0.0397	0.0397	0.0397
Observation	122123	122123	122123	122123	122123	122100	122100
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 (1)). The sample consists of projects observed quarterly between 2013 and 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 0.57 percentage points, relative to an unconditional completion probability of 3.97 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 percentiles, and then standardized. All firm controls are lagged by one period. The reduction in sample size from column (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.

Figure A6: Correlations of Firm Shares and Future Bank Exposure



Notes: the figures plot binscatters of (log) firm shares $s_{bft} = {Loans_{bft}}/{\sum_{f'} Loans_{bf't}}$ on one-period ahead bank exposure Bank Shock$_{b,t+1}$ for both the universe of charges (left panel) and for firms that eventually make it into our regression sample (right panel). Consistent with the timing of our regressions, we restrict to FY 2013-2019 and control for firm and time fixed effects. The notes on the bottom of the charts reproduce the coefficient on the forward bank shocks obtained by regressing the log share log s_{bft} on future exposure Bank Shock$_{b,t+1}$, controlling for time and firm fixed effects and clustering standard errors at the lender level.

Table A10: Effects of Firm-Level Credit Supply Shocks on Censoring Hazards

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: PPML							
Exposure	0.0909***	0.0015	0.0000	0.0042	0.0022	0.0041	0.0022
	(0.0236)	(0.0352)	(0.0349)	(0.0355)	(0.0350)	(0.0414)	(0.0410)
DepVar Mean	0.0261	0.0261	0.0261	0.0261	0.0261	0.0280	0.0280
Observation	103458	103458	103458	103458	103458	96075	96075
Panel B: Linear Probability							
Exposure	0.0031***	-0.0002	-0.0003	-0.0003	-0.0005	-0.0006	-0.0008
	(0.0010)	(0.0011)	(0.0011)	(0.0011)	(0.0011)	(0.0012)	(0.0012)
DepVar Mean	0.0261	0.0261	0.0261	0.0261	0.0261	0.0260	0.0260
Observation	103458	103458	103458	103458	103458	103425	103425
Age Splines	✓	✓	✓	✓	✓	✓	✓
Firm FE		✓	✓	✓	✓	✓	✓
State x Year FE		✓	✓	✓	✓	✓	✓
Sector x Year FE				✓	✓		
Industry x Year FE						✓	✓
Firm Controls			✓		✓		✓

Notes: the table reports estimates of the effect of firm exposure to the Asset Quality Review on project-level censoring hazards. The outcome is an indicator for a project becoming permanently lost to follow-up, i.e., exiting the sample with a terminal "no information" status. Specifications mirror the columns of Table 2. Panel A reports PPML estimates, interpreted as the percent effect on the censoring hazard of a one-standard-deviation increase in firm exposure. Panel B reports estimates from a linear probability model, interpreted in percentage points. The estimation sample is smaller than in Table 2 because projects of firms that never experience a censoring event are dropped due to separation under firm fixed effects. The further reduction in columns (6) and (7) of Panel A results from separation of additional observations by the more granular industry-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 A11: 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 (2)). The sample consists of the Capex-Prowess matched sample of firms observed yearly between 2013 and 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 percentiles, and then standardized. All firm controls are lagged by one period. The reduction in sample size from column (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 A12: 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 (2)). The sample consists of the Capex-Prowess matched sample of firms observed yearly between 2013 and 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 0.836 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 percentiles, and then standardized. All firm controls are lagged by one period. The reduction in sample size from column (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 A13: 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 percentiles. Standard errors are heteroscedasticity-robust and reported in parentheses. $*$, $**$, and $***$ denote statistical significance at the 10%, 5% and 1% levels, respectively.

Table A14: 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 and 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 percentiles. 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 Maturity Effects

This appendix elaborates on the construction of the maturity measures and related robustness exercises.

D.1 Measuring Project Maturity

We do not observe project quality or stage directly in our data. Since we allow for unobserved project heterogeneity (the z in our model), the cross-sectional distribution of age observed in the data does not reflect a one-to-one mapping between observed age and residual time to build. Simply put, if we observe a project that is old, we cannot differentiate whether the advanced age reflects a low-productivity project that is far away from completion or a high-productivity project that is about to be completed.[63]Figure A7 (left panel) illustrates the issue by plotting the quarterly completion probability of as a function of age, estimated non-parametrically on pre-AQR data. The baseline hazard is non-monotonic and hump-shaped. The declining arm reflects dynamic selection: conditional on still running at high ages, projects are disproportionately slow-moving. Simply put, two projects of the same age can sit at very different distances from completion.

To move beyond these concerns, we rely on survival models again. We compute maturity scores that ascertain how close a project is to completion based on a rich set of project characteristics and the project progress. By using moments from the estimated survival functions rather than observed age alone, we can map the observed age to a posterior likelihood of completion in the coming quarters. Operationally, we first estimate a discrete-time hazard model for completion on pre-AQR data with flexible age-bin dummies and predetermined covariates only (industry, project cost, ownership and unit type, state, the project's own announced completion schedule, and pre-period stall history). Missing covariate values enter as explicit categories, so every at-risk project receives a maturity score.[64] For estimation, we split firms into five folds and predict each project's hazards over time from a model estimated on the other four folds. This alleviates concerns of over-fitting, as in standard machine learning prediction tasks.

Once we have the fitted hazards, we integrate over the project histories to obtain the conditional completion probability P_H of equation (3) and the restricted mean survival time

[63]This issue is also familiar to epidemiologists and labor economists. For example, in clinical settings, patients of similar ages may react differently to the same drug, even conditioning on pre-existing conditions, due to unobserved differences in metabolic factors, routine habits, etc. In labor economics, workers can remain unemployed for the same time, but we do not know if this reflects low search effort or low unobserved skill.

[64]We include project-quarters from 2001Q1 through 2015Q2, covering 858,000 project-quarters and 20,200 completions. The window ends at the last full quarter before the AQR announcement window opens, so neither outcomes nor covariates dated after 2015Q2 enter the model anywhere. Spells still running in 2015Q2 are right-censored at the shock, as estimating on completed spells only would oversample fast projects and bias survival downward. We bin some categories to ensure at least 100 observations within each cell.

Table A15: Frozen maturity scores at the shock date: distributions across the at-risk cross-section.

Score	N	Mean	SD	p25	p50	p75
P8	22,042	0.176	0.157	0.039	0.144	0.264
P17	22,042	0.338	0.234	0.138	0.317	0.506
RMST(40) (quarters)	22,042	26.025	8.091	20.031	26.376	32.219

Notes: the table reports the cross-sectional distribution of the frozen maturity scores over all projects at risk in 2015Q2. The scores are computed from the cross-fitted discrete-time hazard model described in the text, with each project's age and covariates evaluated at their last pre-shock values: P_8 and P_{17} are the predicted probabilities of completing within the next 8 and 17 quarters, respectively, and RMST(40) is the expected residual construction time in quarters, truncated at 40 quarters. Scores are reported in raw units, before the winsorization and standardization applied in the regression tables. The estimation samples of Table 4 are smaller than the at-risk cross-section because the regressions additionally require the firm exposure measure and firm controls.

$\text{RMST}(L) = \sum_{h=1}^{L} [1 - P_h]$, i.e., the expected residual construction time. The main text uses $\text{RMST}(40)$ and P_8. The truncation at $L = 40$ ensures RMST is not distorted by the far tail of the survival curve, which is poorly identified. Table A15 reports the distributions. All maturity measures are evaluated at each project's age in 2015Q2, with covariates at their last pre-shock values.

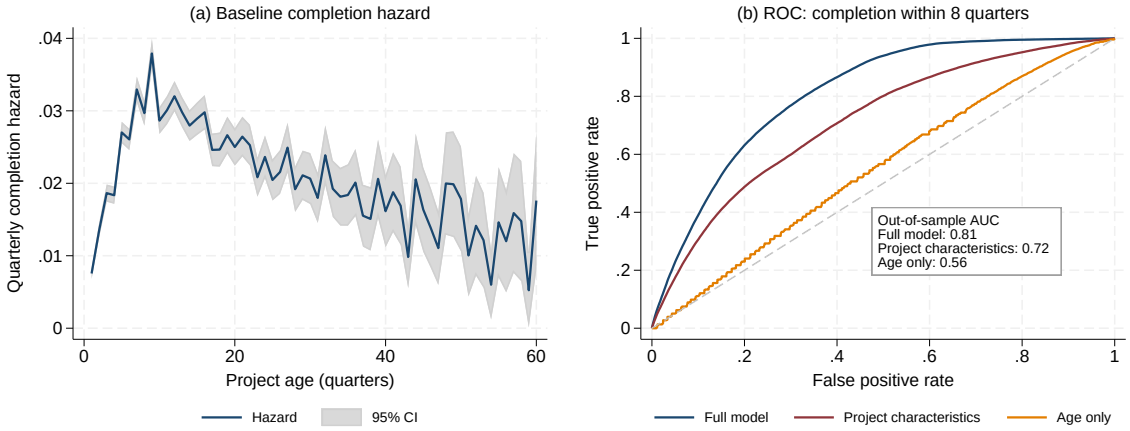
D.2 Validation

Before doing any regressions, we provide three validation exercises that indicate our measures indeed capture project maturity. First, from an out-of-sample prediction standpoint, the survival-based proxies discriminate well between projects that complete within two years and those that do not. Figure A7 (right panel) plots receiver operating characteristic (ROC) curves and computes the associated areas under the curve (AUC). Higher values indicate that the prediction model is better at discriminating projects that do from those that do not complete. Our preferred model has an AUC of 0.81. For comparison, a model with age only performs not much better than a coin flip (i.e., the AUC is about 0.5).

Second, our maturity measures are nearly unbiased and well calibrated predictions of the realized completion probabilities. Regressing the realized within-8-quarter completion indicator on the P8 measure with age bins absorbed yields a coefficient of 1.08 (s.e. 0.013). In other words, among projects of the same age, a 10 pp higher predicted probability means a roughly 11 pp higher realized completion rate (Table A16). To ensure this is not an artifact of the two-year horizon, we repeat the exercise for one-year and three-year horizons, obtaining similar results.

Third, our maturity measures capture information that is materially distinct from age alone. Most of the variation in maturity measures lives *within* age-by-industry cells (91% for P_8), and the score is stable across estimation subperiods. More detailed results are available upon request.

Figure A7: Baseline Completion Hazard and Out-of-Sample ROC Curves



Notes: the figure characterizes the information content of project age and of the survival-model maturity scores. Panel (a) plots the quarterly completion hazard by project age, estimated non-parametrically on pre-AQR project-quarters (2001Q1–2015Q2), with a 95% confidence band. The hazard is hump-shaped: the declining arm reflects dynamic selection, since projects still under way at high ages are disproportionately slow-moving. Panel (b) plots receiver operating characteristic (ROC) curves for out-of-sample classification of whether a project completes within the next 8 quarters, for three discrete-time hazard models: age only; age plus project characteristics (industry, project cost, ownership and unit type, and state); and the full model, which adds the project's pre-period stall history. Predictions are cross-fitted across five firm folds, so no firm's own outcomes enter its prediction. The inset box reports the out-of-sample area under the curve (AUC) of each model; the 45-degree line corresponds to a classifier no better than chance.

Table A16: Predicted maturity and realized completion rates

	(1)	(2)	(3)
Horizon	$H = 4$	$H = 8$	$H = 12$
Predicted probability (P_H)	1.128***	1.079***	1.056***
	(0.0154)	(0.0129)	(0.0123)
Age bin FE	✓	✓	✓
N	697,842	588,805	509,085
R^2	0.132	0.205	0.249

Notes: each column reports an OLS regression of an indicator equal to one if the project completes within the next H quarters on the out-of-sample predicted completion probability P_H, absorbing project-age-bin fixed effects. The sample pools all pre-AQR project-quarters for which the H-quarter outcome is observable. A coefficient of one indicates that, among projects of the same age, predicted probabilities map one-to-one into realized completion rates. 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.3 Robustness Exercises

The results for the heterogeneity analysis are presented in the main text. We compile here additional robustness exercises.

Model-implied thresholds. The model delivers an analytical bound on the acceleration region that depends only on the discount rate and the persistence of the shock. Write $T \equiv T(x,z)$ for remaining time to build, $h_\rho(u) \equiv (1-e^{-\rho u})/\rho$, and $c(\tau) \equiv i(y(\tau),z) + \kappa(i(y(\tau),z))$ for the construction outlay τ periods ahead along the steady-state path, as in Proposition 2. Along that path the value of the project is the discounted terminal payoff net of discounted outlays,

$$v(x,z) = P(x,z) - C(x,z), \qquad P(x,z) \equiv e^{-rT} pz, \qquad C(x,z) \equiv \int_0^T e^{-r\tau} c(\tau)\, d\tau,$$

and equation (14) reads $-\hat{v}(x,z) = h_\rho(T) P(x,z) - \int_0^T h_\rho(\tau) e^{-r\tau} c(\tau)\, d\tau$. Let

$$H_C(x,z) \equiv \frac{\int_0^T h_\rho(\tau) e^{-r\tau} c(\tau)\, d\tau}{C(x,z)}$$

denote the outlay-weighted average of $h_\rho(\tau)$ over the remaining construction period.

Proposition 10. *For any ongoing project with $v(x,z) > 0$,*

$$D(x,z) = h_\rho\big(T(x,z)\big) + \frac{C(x,z)}{v(x,z)} \Big[h_\rho\big(T(x,z)\big) - H_C(x,z) \Big] \;\geq\; h_\rho\big(T(x,z)\big), \tag{25}$$

with strict inequality whenever $x < \bar{x}$. As $\rho \to 0$, the decomposition becomes $D = T + (C/v)\,[T - T_C]$, where $T_C \equiv \int_0^T \tau e^{-r\tau} c(\tau) d\tau / C$ is the duration of the remaining outlays.

Proof. Substitute $P = v + C$ into the expression for $-\hat{v}$ and divide by v:

$$D = -\frac{\hat{v}}{v} = \frac{h_\rho(T)(v+C) - C\, H_C}{v} = h_\rho(T) + \frac{C}{v}\big[h_\rho(T) - H_C\big].$$

Since h_ρ is increasing and every outlay occurs at some $\tau \leq T$, $H_C \leq h_\rho(T)$, with strict inequality when $C > 0$, i.e., when construction is not yet complete. The limit follows from $h_\rho(u) \to u$. □

The first term is the duration of a claim that pays only at the physical completion date. The second term is nonnegative because the project pays out construction costs before it receives the terminal payoff: those early outflows make net project value more sensitive to discounting than the completion date alone would suggest. Physical remaining time to build

is therefore a lower bound on economic duration, and the gap is the financing component of the project's exposure to the shock.

Combining Proposition 10 with the duration condition (Appendix H.6, Proposition 18) gives a bound on the acceleration threshold $T^*(z)$ of Proposition 19 that holds for every project type.

Corollary 11. *Define* $\bar{T}(r,\rho) \equiv \rho^{-1}\log(1+\rho/r)$. *For every type* z,

$$T^*(z) \le \bar{T}(r,\rho), \qquad \textit{so that} \qquad T(x,z) > \bar{T}(r,\rho) \implies \hat{\imath}(x,z) < 0.$$

The horizon depends only on r *and* ρ, *not on* z, p, *the size of the project, or the acceleration cost* κ. *As* $\rho \to 0$, $\bar{T} \to 1/r$.

Proof. By Proposition 10, $D \ge h_\rho(T)$. Acceleration requires $D < 1/(r+\rho)$, hence $h_\rho(T) < 1/(r+\rho)$, i.e., $e^{-\rho T} > r/(r+\rho)$, i.e., $T < \bar{T}$. Thus every accelerating stage has $T < \bar{T}$, which gives $T^*(z) \le \bar{T}$ and the contrapositive. □

At the calibrated quarterly discount rate $r = 0.01$ and the quarterly decay rate of the shock $\rho = 0.17$ (Appendix Table A24), $\bar{T} = \log(18)/0.17 \approx 17$ quarters. A project with more than 17 quarters of construction remaining is slowed down by the shock whatever its type; a project within 17 quarters of completion may or may not be accelerated, since its own threshold $T^*(z)$ lies below $\bar{T}$ by an amount governed by the financing component in equation (25). The probability of completing within 17 quarters, P_{17}, is therefore the probability that a project lies inside the universal horizon: an upper bound on the probability that it is in the acceleration region, and one minus the probability that it is certainly in the slowdown region. Table A17 repeats Table 4's specifications with P_{17} as the maturity proxy, finding the same results.

PPML vs OLS. The estimates in Table 4 are level effects, i.e., percentage points of quarterly completion probability. Meanwhile, in many other analyses, we reported PPML estimates. Here, we argue that PPML estimates are not appropriate to test the maturity mechanism, because the two estimators deliver different answers for mechanical rather than economic reasons.

PPML estimates proportional hazard effects. Hence, how the baseline hazard changes as the project matures crucially determines its estimate. Baseline completion rates rise steeply with maturity in our data. The top maturity quartile completes at 5.5% per quarter against 0.04% in the bottom, a 150x difference. Consequently, mildly increasing *level* gains across the maturity distribution can translate into decreasing *proportional* effects in the maturity distribution as baseline hazards increase faster than the level effects. Note that this is *not* an

issue for average estimates (all other PPML estimates in the paper), since the average baseline hazard is stable. However, to study specifically maturity heterogeneity, the mechanical dependence of proportional estimators on the behavior of baseline hazards artificially obfuscates the economic mechanism.

More concretely, we find *level* gains across the maturity distribution. However, PPML weights observations by their baseline hazard. This is not an economically motivated weighting scheme, but rather imposed by quasi-maximum likelihood. In our data, we find that PPML concentrates roughly half its weight on the top three maturity deciles. As baseline hazards rise steeply with maturity, PPML places more weight on the negative proportional effects, pushing the interaction coefficient downwards.

To illustrate how the PPML weighting can mask economically meaningful level effects, consider the ladder presented in Table A18. For projects that are not separated by the PPML estimator, the PPML interaction coefficient is negative and marginally significant. Converting the PPML interaction into level units rather than proportional units yields a noisy zero. Meanwhile, estimating OLS in the same sample yields a positive and significant level effect. Dividing the OLS interaction by the pre-2016 hazard at mean maturity gives a proportional analog of $+0.19$ (s.e. 0.11), against the log link's -0.07. Hence, the equal-weighted proportional estimator and the hazard-weighted one disagree precisely because the baseline gradient is so steep.[65]

Other robustness tests Table A19 collects many other robustness exercises. First, we include age at the shock and its exposure interaction. This further alleviates any concern, despite our validation exercises, that our maturity proxies could be simply capturing age. The maturity interactions are essentially unchanged from Table 4.

The full-panel row of Table A19 removes the conditioning and reproduces the full estimation panel of Table 2. Maturity at the shock is defined only for projects on the 2015Q2 risk set.[66]

The score rows of Table A19 re-estimate Table 4's interaction with scores from four progressively leaner survival models: the full covariate set; excluding stall history; excluding schedule-position indicators; and excluding the announced schedule entirely. The maturity interaction is stable and significant throughout.

Redefining exposure with fixed pre-period loan weights (or weights frozen in 2015) barely moves the interaction cells. Replacing age-at-shock bins with the panel-age splines of Table 2 changes nothing. Replacing the linear interaction with maturity-quartile bins yields a mono-

[65] What the proportional margin does detect is concentrated among overdue projects. Once an overdue dummy is included, the maturity interaction survives additively (OLS) while it disappears proportionally (PPML), loading instead on the overdue dummy (results available upon request).

[66] To include other projects, we include dummies for whether we can construct the maturity variables and their interaction with exposure.

Table A17: Effects of exposure by P_{17}, the model-implied acceleration probability (linear probability model, landmark sample; columns as in Table 4).

	(1)	(2)	(3)	(4)
Exposure	0.0021**	0.0022**	0.0022**	0.0023**
	(0.0010)	(0.0010)	(0.0009)	(0.0009)
P17 x Exposure	0.0045***	0.0045***	0.0042***	0.0042***
	(0.0009)	(0.0009)	(0.0009)	(0.0009)
Age at Shock FE	✓	✓	✓	✓
Firm FE	✓	✓	✓	✓
State x Year FE	✓	✓	✓	✓
Sector x Year FE	✓	✓		
Industry x Year FE			✓	✓
Firm Controls		✓		✓
N	101,632	101,632	101,607	101,607

Notes: the table repeats the specifications of Table 4 with P_{17} — the model-implied probability of completing within $\bar{T} = 17$ quarters — as the maturity measure. The sample, fixed effects (reported in the table body), firm controls, winsorization and standardization, and the interpretation of the coefficients are identical to Table 4. Standard errors are clustered at the firm level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5% and 1% levels, respectively.

tonically increasing and statistically significant exposure-effect profile, matching the linear gradient (figure available upon request).

Last, we dispel concerns that the maturity score could simply be picking up the *type* of project rather than progress toward completion. We measure "maturity at birth" for every project at its announcement date, using announcement-time characteristics only, and re-run Table 4 with the birth maturity in place of residual maturity at the shock. The interactions are precise zeros (Table A19, placebo row), while the same sample produces the full gradient when maturity proxies are measured at the shock date. The heterogeneity therefore indeed captures realized progress. Projects that have advanced toward completion respond more, as an acceleration mechanism implies.

Table A18: Level and proportional estimates of the maturity interaction.

	coef.	SE
Panel A: log-link sample (N = 45,257)		
PPML interaction (proportional)	-0.0670*	(0.0369)
PPML implied average cross-partial (level)	-0.0002	(0.0027)[a]
OLS interaction (level)	0.0047**	(0.0023)
OLS ÷ pre-2016 hazard at mean maturity (proportional)	0.194*	(0.107)[b]
Pre-2016 hazard at mean maturity	0.0243	(0.0017)
Panel B: full sample, incl. separation-dropped projects (N = 101,632)		
OLS interaction (level)	0.0036***	(0.0009)
OLS ÷ pre-2016 hazard at mean maturity (proportional)	0.650***	(0.186)[b]
Pre-2016 hazard at mean maturity	0.0056	(0.0004)

Notes: all regressions include firm exposure (per standard deviation), the standardized P_8 score, and their product as regressors, with firm, sector-by-year, state-by-year, and project-age-at-shock-bin fixed effects and the firm controls of Table 4; standard errors are clustered at the firm level unless noted. Panel A fits both estimators on identical rows — the project-quarters on which the log-link (PPML) estimator is defined; Panel B fits OLS on the full harmonized sample, which retains the projects the log link drops by separation. The PPML interaction is a semi-elasticity and is directly proportional; the implied average cross-partial converts it into percentage-point units by averaging the fitted model's second cross-derivative over the estimation sample. The OLS-÷-hazard rows divide the OLS interaction by the pre-2016 completion hazard of the project at mean maturity in the same sample — the intercept of a firm-clustered regression of the completion indicator on the standardized maturity score over pre-2016 rows — giving the OLS a proportional reading, at the mean-maturity project, comparable to the PPML row; the hazard rows report that denominator with its firm-clustered standard error. *, **, and *** denote statistical significance at the 10%, 5% and 1% levels, respectively (the descriptive mean rows carry no stars). [a] Standard error from 100 firm-cluster bootstrap replications. [b] Standard error from 100 firm-cluster bootstrap replications resampling the interaction coefficient and the pre-2016 mean jointly, so it accounts for the estimated denominator.

Table A19: Maturity-interaction point estimates across designs (linear probability model).

Design	RMST(40) x Exp	P_{17} x Exp	P_8 x Exp
Headline (landmark sample)	-0.0043***	0.0045***	0.0036***
Age race: maturity and age jointly	-0.0041***	0.0042***	0.0034***
Full panel (no conditioning)	-0.0047***	0.0048***	0.0037***
Score: no stall history	-0.0035***	0.0036***	0.0024***
Score: no overdue indicators	-0.0031***	0.0031***	0.0029***
Score: no announced schedule	-0.0036***	0.0036***	0.0033***
Exposure: pre-period weights	-0.0043***	0.0046***	0.0038***
Exposure: 2015 weights	-0.0044***	0.0047***	0.0039***
Log-link estimation rows	-0.0053***	0.0055***	0.0047**
Placebo: announcement-date scores	-0.0009	0.0006	0.0005

Notes: each cell is the maturity-interaction coefficient from a separate regression, collected from the designs of this appendix; each row is labeled by its single deviation from the headline design. All rows share the headline architecture: firm, sector-by-year, state-by-year, and project-age-at-shock-bin fixed effects, the firm controls of Table 4, and firm-level clustering. The age-race row adds the interaction of exposure with age at the shock to the headline specification, so maturity and age compete directly; the maturity measure and age are both winsorized and standardized. "Log-link estimation rows" re-estimates the linear probability model on the subsample where the log-link estimator is defined (Table A18); the placebo row replaces shock-date maturity with announcement-date scores constructed from announcement-time characteristics only, on the full estimation panel of Table 2. *, **, and *** denote statistical significance at the 10%, 5% and 1% levels, respectively, based on firm-clustered standard errors.

E Further Robustness Checks

E.1 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.

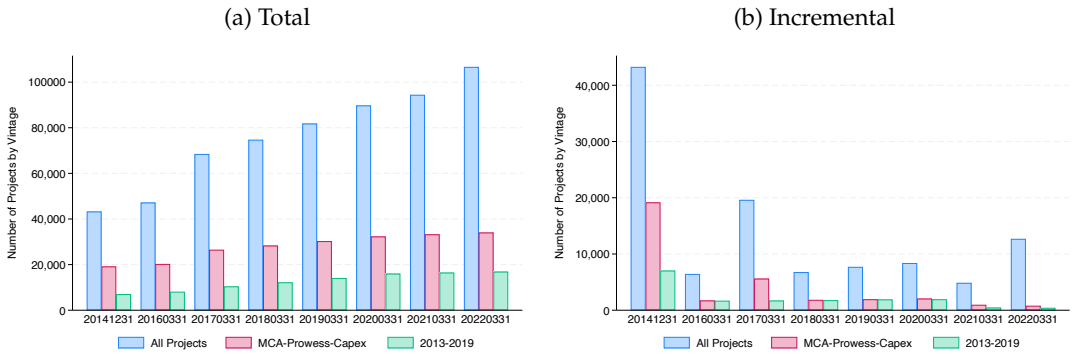
On the one hand, such retroactive filling improves coverage. On the other hand, it introduces a potential threat to our empirical strategy. The concern is that the RBI audits – and their resulting firm-level disruptions – alter the data collection process. For example, CapexDx could become less likely 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.[67]

If selection into the sample is independent of treatment, our estimates may retain a causal interpretation within the selected sample (Hernán and Robins, 2020). The threat is differential censoring. Suppose CapexDx became less able to obtain data on treated relative to untreated projects, for example because more exposed firms stop sharing information in the public domain or stop responding to inquiries. The treated sample would then be skewed toward fewer young projects. Because older projects are more likely to be completed, treatment would then predict completion even under a true null effect. A direct implication is that fewer young projects should 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.

[67] This 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.

Figure A8: Total and incremental counts of projects by vintage



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 (right-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 A8 plots the total number of projects in each vintage, as well as the number of projects added in each vintage. We split projects into 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.

We reestimate the completion regressions in the first vintage subsample. Table A20 reports results. In the first row, the firm-level exposure variable aggregates bank-level shocks using lagged, time-varying weights (LTV). The point estimates are quantitatively similar to those obtained without restricting the sample. The standard errors are larger, but this is unsurprising given the sample size is effectively halved. Despite the reduced power, the similar point estimates between the two specifications are highly encouraging: if bias were a concern, point estimates would have changed considerably between the two exercises.

In the second and third rows, the firm-level exposure variables weight 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. This increase in point estimates might appear puzzling at first, but would be consistent with later-year weights becoming smaller if firms slowly substitute away from more exposed banks.

The results in Table A20 corroborate that our main results are not driven by treatment inducing differential sample selection into later vintages. They do not, however, speak to

Table A20: Effect on Completion Hazards: Subset of First Vintage Projects: Panel

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Exposure (LTV)	0.0213	0.0637	0.0627	0.0580	0.0578	0.0469	0.0449
	(0.0302)	(0.0422)	(0.0422)	(0.0425)	(0.0425)	(0.0453)	(0.0457)
Exposure (PreW)	0.0260	0.0802**	0.0817**	0.0782*	0.0804*	0.0717*	0.0743*
	(0.0306)	(0.0409)	(0.0410)	(0.0418)	(0.0420)	(0.0434)	(0.0437)
Exposure (2015W)	0.0241	0.0769*	0.0789**	0.0756*	0.0783*	0.0697*	0.0726*
	(0.0299)	(0.0397)	(0.0397)	(0.0405)	(0.0405)	(0.0417)	(0.0419)
DepVar Mean	0.0385	0.0385	0.0385	0.0385	0.0385	0.0408	0.0408
Observation	53750	53741	53741	53741	53741	50277	50277
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 (1)). The sample consists of projects observed quarterly between 2013 and 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. "(PreW)" 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. "(2015W)" 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 2.13%. 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 percentiles, and then standardized. All firm controls are lagged by one period. The reduction in sample size from column (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 A21: Panel Balance Tests: Project Stocks

	Mean | Exp = 0	Mean | Exp > 0	Mean | Exp	Mean | Exp
	(1)	(2)	(3)	(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 and 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.

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. Balance in the levels of project stocks is far more stringent than what the differential-selection story implies and than what our identifying assumption requires: levels could differ simply because of permanent differences in the set of firms that are eventually exposed, which our panel regressions absorb with firm fixed effects. We thus test for differential pre-trends directly in a panel regression, asking whether contemporaneous exposure predicts a decline in the lagged number of projects.

Table A21 collects the results. The last two columns both regress lagged variables on contemporaneous treatments. The third column does not include any fixed effects. We cannot rule out differences in levels.[68] But we find no evidence of differential pre-trends. Once we

[68]The binary and continuous comparisons answer different questions. Columns (1) and (2) compare firms with any exposure to firms with none, whereas column (3) measures how project stocks vary with the level of

include firm and time effects in the last column, the predictability vanishes. Point estimates become small in magnitude and economically insignificant. The results for the number of projects are 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 alter 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.

E.2 Changes in Scope

In the main text, we argued that our results were not driven by firms simply changing the scope of their projects ex post, reducing their size and thereby completing them faster. This appendix supplements such discussion. First, we provide some institutional details about the construction industry. Second, we elaborate on discussions we had with our data provider on these concerns. Third, we implement robustness checks to our main specifications. Despite data limitations, qualitative and quantitative evidence alleviates concerns that our results are driven by the credit supply shock leading to ex-post changes in project scope.

Construction in India: Institutional Details Construction is a risky endeavor, and established legal frameworks allocate risks between the parties involved in a construction project. Although many such contracts are incomplete in the sense of Hart and Moore (1988), the sheer monetary value at stake has led to the emergence of standard contractual forms, which are then adapted to specific situations.[69] In this appendix, we describe the contractual features most relevant to our interpretation of the empirical results.

Construction contracts involve two key parties: the sponsor and the contractor. The former wants to build a particular structure, whereas the latter executes the construction. As the project grows in scope, the number of parties can increase substantially, with the addition of architects, engineers, subcontractors, subcontractors to subcontractors, and so on.

exposure. Their signs therefore need not agree.

[69] While the construction industry in India does not formally subscribe to a particular standard contract, it is in practice common to adhere to the suite of contracts published by the Fédération Internationale des Ingénieurs-Conseils (FIDIC, the International Federation of Consulting Engineers) or by the Indian Institute of Architects. In the US, most contracts are similarly based on forms published by the American Institute of Architects. Among contracts with a private sponsor, the most frequently used FIDIC form is the design-build contract.

Construction contracts normally combine multiple documents.[70] The drawings and the "specs" describe the full scope and quality of the work, and the associated bill of quantities lists the materials required.[71] In the language of our model, these engineering primitives are captured by $(\bar{x}, \bar{i})$. Under the design-build contractual forms most commonly used by private sponsors, these documents are produced at the design stage, before contractors bid for the right to implement the work.

At the time of signature, the contract will specify the scope of the project. If the sponsor moves to change the scope of the project ex post – a "variation in the works" – the contractor must approve the variation and is entitled to compensation for the additional work required. In fact, any term or clause that allowed the sponsor to unilaterally and substantially change the contract by omitting obligations to the contractor would be considered unenforceable in court. Crucially, the variations must be of limited nature. If they substantially alter the contract or if they are not within the contractor's ability to execute, such variations may be denied or result in contractual breaches.

Particular to our setting, variations map into either the adjustment costs $\kappa(\cdot)$ or into the project scope $\bar{x}$. Requesting that the work be accelerated with no changes in scope is a valid variation so long as the contractor is (i) duly compensated for the resulting costs (e.g., increased wage payments for overtime, as captured by additional marginal costs $\kappa'(\cdot)$ in the model) and (ii) the acceleration did not entail impossible deadlines on the contractor. But changing the project ex post to reduce its scope – and therefore reduce the originally contractually stipulated pay to the contractor – can result in legal action against the sponsor, unless explicitly permitted by the contract. It is not uncommon for any major negative variation to induce a partial termination of contract, which then likely entails indemnifications from the sponsor to the contractor. When disputes do arise, arbitration is the usual resolution mechanism; yet congestion in the legal system and other judicial delays make dispute resolution costly, so good-faith negotiations to amend the contract by mutual consent are typically pursued first.

Robustness based on Communications with Data Provider Besides gathering institutional know-how about variations in works, we communicated with the data provider about our

[70] The Indian Contract Act of 1872 – largely based on English Common Law – governs construction contracts. According to Accounting Standard (AS) 7 on Construction Contracts, a construction contract is "a contract specifically negotiated for the construction of an asset or a combination of assets that are closely interrelated or interdependent in terms of their design, technology and function or their ultimate purpose or use".

[71] Specs can be extremely complex from an engineering perspective (e.g., structural calculations for a suspension bridge) and include detailed instructions of the materials used in the enterprise (e.g., the particular steel beams for the structure of a commercial building). The bill of quantities is normally used to estimate in good faith the project cost. The contract's Conditions further allocate responsibilities and risks to each party, listing the sponsor's expectations on the scope and quality of the work, the desired performance of the final structure (e.g., the generating capacity of a power plant or the final square footage of a residential building), and the procedures for dispute adjudication.

concerns. The Capex team actively monitors projects for changes in scope. We further asked *"Are there instances where a project might get marked as 'Complete', but the company just reduced the scope of the original project and considered it complete?"*

First, Capex told us that they hardly found any drastic changes in scope across many years in the data. Second, were such drastic changes ever found, Capex followed a very structured logic to record project revisions.

If the change in scope is small, the change is recorded directly in a project's profile. Their example was as follows: suppose the initial announcement for a project stated that the sponsor would build a plant with production capacity of X at a cost C. Before implementation starts, the sponsor revises the plans and reduces capacity to 0.9X and costs fall to 0.95C. Capex will then revise the capacity of the project and add a new row in the cost dataset indicating when the revision in cost was made and what the new cost was.

If the change in scope is large, modifications to the project file would depend on the particular change. A hypothetical situation we were given was a change in the use of the structure. If the initial plan was to build a shopping mall, but then the type of the project changed to a residential complex, Capex would record a new residential project under the same sponsor, and stall the original shopping mall project.

Perhaps more relevant as a threat to our interpretation of the regression results, we probed specifically about a hypothetical drastic reduction in scope. We asked *"Suppose a company was going to build three factories in a single complex. The company had completed two out of the three plants, but then the company runs out of cash. It decides to kill the third plant, and considers the complex completed only with two factories."* In this case, Capex would proceed in two steps. First, they would revise downwards the capacity and the costs of the original project. Second, they would push the incomplete plant into a separate project and mark it as stalled. Their logic is that a project can be considered completed only if all the subunits comprising the project have been completed.

Our communications with the Capex team motivate a first battery of robustness checks to our results. First, we exclude any projects that have more than one phase or unit. Capex has less granular information on costs and production independence of such projects, as otherwise they would have been coded as independent projects within the same sponsor. Excluding multi-unit projects safeguards against the possibility that reductions in scope such as eliminating one unit out of the originally planned are driving the results on completion hazards. Second, we exclude from our regression all projects that experience downward cost restatements. Such revisions proxy for negative variations. Last, we exclude from our sample all firms that had stalled projects over the sample period. Following the example provided by Capex, a stalled project coinciding with completion of other projects within a firm could indicate a negative variation.

Appendix Table A22 presents the results. Panel A removes multi-unit projects. The re-

Table A22: Effects of Firm-Level Credit Supply Shocks on Project Completion Hazards

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Remove Multi-Unit Projects							
Exposure	0.1112***	0.0901***	0.0852***	0.0822***	0.0775***	0.0947***	0.0928***
	(0.0285)	(0.0290)	(0.0291)	(0.0277)	(0.0279)	(0.0297)	(0.0298)
N	90289	90289	90289	90289	90289	85812	85812
Panel B.1: Remove Downward Cost Revisions							
Exposure	0.0872***	0.1075***	0.1030**	0.1023**	0.0981**	0.1068**	0.1065**
	(0.0255)	(0.0410)	(0.0415)	(0.0400)	(0.0403)	(0.0462)	(0.0462)
N	51589	51589	51589	51589	51589	47911	47911
Panel B.2: Remove Projects with Missing Costs							
Exposure	0.0970***	0.0997***	0.0941***	0.0941***	0.0885***	0.1247***	0.1198***
	(0.0261)	(0.0345)	(0.0349)	(0.0329)	(0.0332)	(0.0375)	(0.0376)
N	72958	72336	72336	72958	72958	67968	67968
Panel C: Remove Ever Stalled Projects							
Exposure	0.0982***	0.0570**	0.0542*	0.0516*	0.0487*	0.0651**	0.0639**
	(0.0267)	(0.0279)	(0.0280)	(0.0270)	(0.0272)	(0.0276)	(0.0278)
N	104186	104155	104155	104155	104155	99623	99623
Panel D: Remove Any Companies with Projects Stalled							
Exposure	0.0967***	0.0667*	0.0648*	0.0469	0.0451	0.0737**	0.0740**
	(0.0279)	(0.0370)	(0.0369)	(0.0341)	(0.0341)	(0.0375)	(0.0373)
N	52493	52493	52493	52493	52493	49117	49117
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 (1)). The sample consists of projects observed quarterly between 2013 and 2019. Each panel assesses the robustness of the main estimates to eliminating a particular subset of projects that could have experienced reductions in scope following the credit shock. These subsamples were defined based on communications with the data provider. 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 11.12%. 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 percentiles, and then standardized. All firm controls are lagged by one period. The reduction in sample size from column (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.

sults are similar to those reported in Table 2. In Panel B.1, we remove any projects that experience any downward revision in cost over their entire lifecycle. For comparison, Panel B.2 estimates the regressions in the full sample of projects with cost information. The results are similar across both panels, and remain statistically significant in both cases.

Panel C discards any project that was ever stalled during its lifecycle. The point estimates are smaller. Last, Panel D excludes all companies that had any projects stalled in our sample period (2013-2019). The point estimates are similar to those in Panel C, but the much smaller sample size increases standard errors and precludes us from claiming the results are significant in some specifications. In both panels, we obtain smaller p-values when we instead use exposure variables based on borrower shares from the pre-shock period, as in Appendix Table A8; the estimates we present are thus the most conservative.

Robustness based on Ex-post Malleability Classifications To further assess the robustness of our estimates, we draw on institutional knowledge gathered through conversations with industry practitioners to classify projects as "ex-post malleable". The idea is that changing the scope of the project ex post is more difficult for some projects. Some of these barriers are technical, while others are regulatory. We identify ex-post malleable projects; and reestimate our main specifications while removing such projects.

We consider three metrics of ex-post malleability. First, we consider an industry-based measure. Commercial or residential real estate projects are technically less complex than industrial real estate projects or works of heavy civil engineering like bridges, aqueducts, etc. We thus remove projects that are classified as serving any of the following end-user industries: housing construction, computer software, ITES, business services and consultancy, retail trading, wholesale trading, commercial complexes, hotels and restaurants, education, health services, and tourism. Panel A of Appendix Table A23 reports regression estimates excluding such projects. The results are similar to those obtained in Table 2.

Next, we consider a regulatory-based measure. Using the event-level dataset, we identify projects that ever had an event related to regulatory approvals, such as obtaining environmental clearances or other government licenses, locating in a special economic zone, requiring approval by planning commissions, etc. We then reestimate our regression in this subsample of projects. The idea is that any large changes in scope would trigger additional regulatory burden, potentially derailing the project altogether. As regulatory approvals are granted based on pre-determined plans and conditioned on good-faith implementation of such plans, pursuing large scope revisions with disregard for additional reviews could result in judicial action to halt the project. Overall, such regulatory burden would discourage any changes in project scopes. Panel B reports estimates. The point estimates are larger and remain significant, except for the results with industry-time fixed effects. We believe the latter arises as the sharp reduction in sample size coupled with the high granularity of the

industry-time fixed effects compromises our power. But the direction of the point estimates is consistent with our baseline interpretation.

Last, we consider an implementation-based measure. The key idea is that changing the project scope is easier before the project actually starts being implemented, i.e., in the design stage. This logic holds for at least three reasons. First, once construction has started, it often becomes impossible from an engineering standpoint to reduce or increase the scope of the project. For example, one cannot convert a four-lane suspension bridge into a five- or three-lane suspension bridge without threatening the structural soundness of the project. Second, if the project is under a design/build contractual arrangement, changes in scope might require approval of the design firm, which can substantially delay completion. Third, changes in scope midway may induce large variations in costs and schedules, which can trigger partial contractual terminations and follow-on costly arbitration if contractors oppose the changes. This is even more pronounced for complex projects, where networks of different subcontractors are normally interlocked in the project's execution through contractual chains, thereby increasing the potential for hold-up problems (Hart and Moore, 1988). Our implementation-based subsample only keeps projects that were marked as "under implementation" before the AQR shock. Panel C reports results. As in Panel B, the point estimates are larger and remain significant, except for the specifications with industry-time fixed effects, whose point estimates are close to the baseline but less precisely estimated.

Taking Stock. This appendix addressed the concern that our results are driven by firms prematurely completing their projects by reducing their scope following the shock. A perfect test would require observing changes in contractual arrangements or in specifications, which are not available through CapexDx. We instead combined institutional evidence that ex-post reductions in project scope are contractually costly with robustness checks that discard the projects most susceptible to such changes. Across these exercises, our point estimates remain similar to the baseline results, bolstering our interpretation of the findings.

E.3 Narrative Corroboration from Earnings Calls

This appendix elaborates on our reading of corporate transcripts. We use transcripts of corporate earnings calls of Indian listed firms from LSEG Refinitiv Transcripts (Hassan et al., 2025), linked to firms in Prowess by stock-exchange symbol and company name. Coverage is the main limitation. Of the 1,821 firms in our estimation sample, only 52 have an earnings call before 2016, and 45 during the audit window. This thin coverage rules out a more systematic analysis of the textual data. Our main evidence thus remains the quasi-experimental hazard regressions from Section 3. We use the transcripts as *corroborating* narrative evidence. We read the passages in which managers of firms in our sample discuss the timing of their

Table A23: Effects of Firm-Level Credit Supply Shocks on Project Completion Hazards

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
				Panel A: Remove Malleable			
Exposure	0.0851***	0.0774**	0.0768**	0.0757**	0.0748**	0.0976***	0.0981***
	(0.0269)	(0.0320)	(0.0324)	(0.0322)	(0.0324)	(0.0338)	(0.0340)
N	74483	74439	74439	74439	74439	69960	69960
				Panel B: Some Regulatory Event			
Exposure	0.0789***	0.2180***	0.2193***	0.2186***	0.2182***	0.1650	0.1634
	(0.0258)	(0.0805)	(0.0789)	(0.0799)	(0.0779)	(0.1009)	(0.1006)
N	29966	29966	29966	29966	29966	24765	24765
				Panel C: Under Implementation before Shock			
Exposure	0.1180***	0.1402***	0.1348***	0.1269***	0.1220***	0.0665*	0.0614
	(0.0351)	(0.0388)	(0.0390)	(0.0388)	(0.0391)	(0.0398)	(0.0400)
N	65046	65046	65046	65046	65046	61611	61611
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 (1)). The sample consists of projects observed quarterly between 2013 and 2019. Each panel assesses the robustness of the main estimates to eliminating a particular subset of projects that could have experienced reductions in scope following the credit shock. These subsamples were defined based on our conversations with practitioners and LLM parsing of construction manuals. 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 8.51%. 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 percentiles, and then standardized. All firm controls are lagged by one period. The reduction in sample size from column (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.

capital projects. These passages were located using keyword and AI-assisted search. Though we cannot reproduce full transcripts, the following examples illustrate narratives consistent with endogenous time to build. In these passages, managers describe completing projects that are close to completion while holding back less advanced ones, refer to the timing of cash flows and to interest costs during construction, and describe adding resources to speed up construction.

- Prestige Estates Projects Ltd, May 2014: "lot of our CapEx projects are nearing completion. We believe that we should finish them off as fast as possible, because there is a lot of demand in terms of leasing for them, (...) We started spending more so that these projects get completed faster and then rental income starts kicking in (...) the increase in the net debt is due to speeding up of all CapEx projects, because the longer we are going to take the higher would be interest during construction."

- Vedanta Ltd, April 2015: "The shaft sinking at Rampura Agucha mine to extend the underground ore production (...) is slower than the original plan and we are reinforcing contractor resources to speed up this work."

- Jindal Steel and Power Ltd, November 2015: "(...) we are not making any new investments at this moment. What we're trying to do right now is the projects which are very close to commissioning, we are trying to just commission those projects (...) So the Angul Phase II has been put back (...)"

- Sterlite Technologies Ltd, April 2018: "We are happy to inform that we are on an accelerated path to bring about 10 million of the expanded capacity at Phase 1 of this project by end of this current calendar year. This would allow some benefit of the expanded optical fiber capacity (...) which we are currently running over 95% utilization to flow into the current financial year."

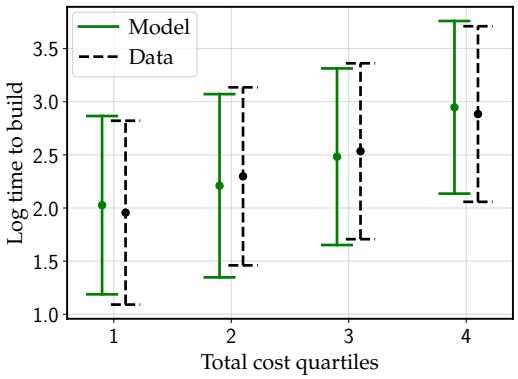
- Sobha Ltd, February 2019: "you are investing back money faster in completing the project to ensure that (...) you can get money faster (...) That kind of a philosophy is allowing us to ensure that the investment is made on completing the projects a little bit ahead of time."

- Tata Steel Ltd, May 2019: "(...) money is already being spent on the ground. So the focus is to get the revenues as fast as possible. (...) So we are basically prioritizing to see how we will reduce our cash to cash cycle for the project and get revenues as fast as possible (...) One has the option of further deleveraging if you stop that CapEx, but then the money you spent on the ground is not getting any returns."

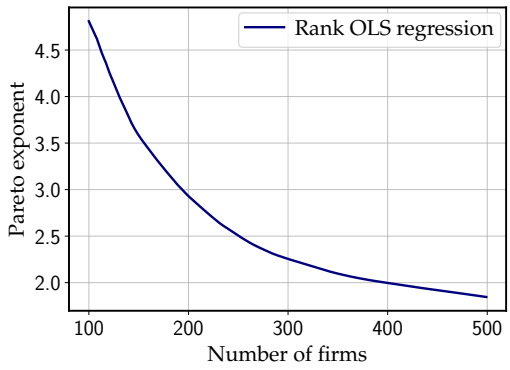
- Arvind Ltd, May 2019: "we are currently taking a pause for the first half of this year to complete what investments we have begun and not add new investments (...)"

F Additional Quantitative Results

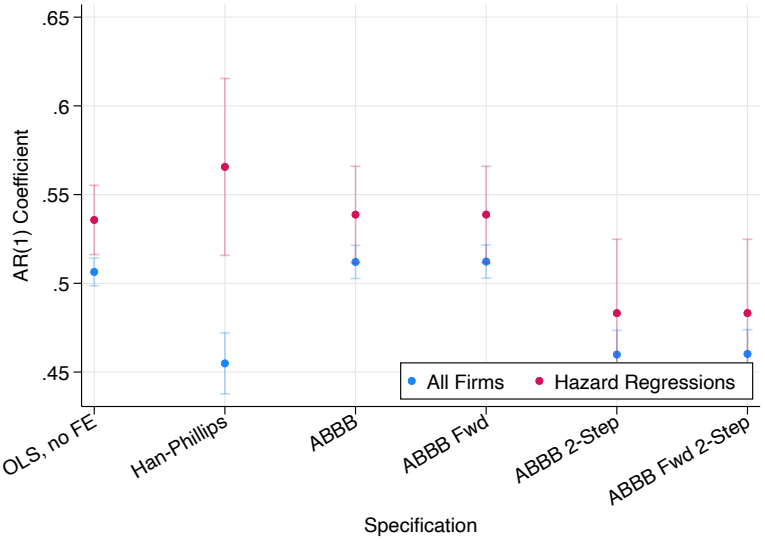
Figure A9: Calibration Exercises



(a) Project Size and Time to Build



(b) Project Output



(c) Shock Persistence

Notes: The figure plots the results of our calibration exercises. Subfigure A9a 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 $\bar{\imath}$. Subfigure A9b 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 A9c 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 A24: 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	
$\beta^{\bar{i}}_0$	-2.47	Mean $\bar{i}$: intercept	Cost-TTB Distribution
$\beta^{\bar{i}}_1$	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	0.02	Shock Size	Reduced Form Regressions
κ	3.09	Accel. Cost	

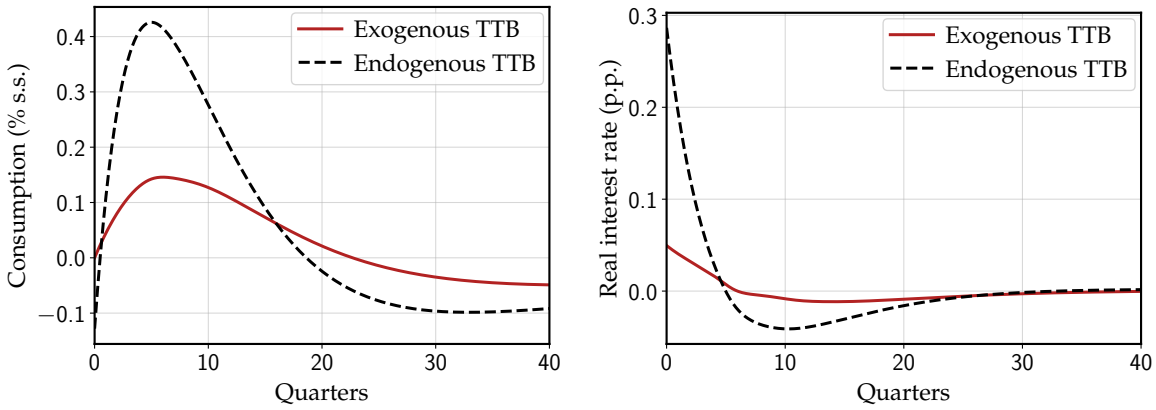
Notes: the table reports parameter estimates from our calibration exercise. See the notes to Figure A9 and the main text for detailed explanations of the calibration strategy.

Table A25: 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 A9 and the main text for detailed explanations of the calibration strategy.

Figure A10: IRFs of consumption and interest rates to discount rate shock in GE.



(a) Consumption

(b) Interest Rates

Note: the figure plots (quarterly) impulse response functions (IRFs) of aggregate consumption (panel a) and discount rates (panel b) to a shock in discount rates in general equilibrium. The y-axis measures percent-deviations from steady state values. Shock persistence, shock size, and remaining parameters as in Appendix Table A24.

G Endogenous Time to Build and Bond Pricing

The results in Section 4 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 the 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 4 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 4. 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 4. In turn, bond investors do not 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,

$$\frac{\partial V(\Delta)}{\partial \Delta} = \int_0^{T-\Delta} e^{-rt} \frac{\partial c(t + \Delta)}{\partial \Delta} dt - e^{-r(T-\Delta)} c(T) + r e^{-r(T-\Delta)} F \tag{26}$$

To build intuition about (26), 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 (26) 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 (26) 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

$$\frac{\partial V(\Delta)}{\partial \Delta}=-c(\Delta)+rV(\Delta) \tag{27}$$

Expression (27) 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 (27) provides the answer.

Suppose now that the bond investor could pay $\kappa(\Delta)$ today to accelerate cash flows by

Δ.[72] We assume κ is differentiable and increasing. The necessary condition for an optimal choice of "acceleration" is

$$\kappa'(\Delta^*) = \frac{\partial V(\Delta^*)}{\partial \Delta} \tag{28}$$

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 most bonds, we tend to think that the capital gains accruing from receiving a large payoff earlier outweigh the lost coupons. Formally, (27) is positive and therefore $\Delta^* > 0$.

A notable feature of (28) is that it carries the same economic content as (9), 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 (9) is written in stage space, expression (28) is written in time space. However, we know that $dx = idt$, so that (9) can be rewritten in time space and yield a parallel expression to (28). 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 (28), it follows that

$$\frac{\partial \Delta^*}{\partial r} = \frac{\frac{\partial^2 V(r,\Delta^*)}{\partial r \partial \Delta}}{\kappa''(\Delta^*) - \frac{\partial^2 V(r,\Delta^*)}{\partial \Delta^2}} \propto \frac{\partial^2 V(r,\Delta^*)}{\partial \Delta \partial r} \tag{29}$$

which again bears a striking resemblance to (15), the expression characterizing the impact effect of a discount rate shock on investment flows.[73] 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 (29). By Fubini's theorem, an equivalent interpretation is that bond (project) investors speed up cash (investment) flows if doing so reduces the bond (project) duration. This is precisely the intuition we emphasized in the main text.

[72] 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.

[73] The only difference is that instead of just having the second derivative of the acceleration cost in the denominator, we also have $\partial^2 V(r,\Delta^*)/\partial \Delta^2$. For a bond with constant coupons, it is easy to show that this second derivative is always positive for any bond of finite maturity. Since the denominator is the second-order condition for a maximum, we assume it holds for simplicity, anchoring on the results for constant coupon bonds. The reason why only the second derivative of the acceleration cost appears in our expressions for the full model follows from the fact that the value function is defined directly as a function of states only, and not of controls, while the value of the bond $V(r,\Delta)$ depends on both states and controls (very much like an arbitrary utility or profit function).

Differentiating (27) with respect to r yields

$$\frac{\partial V(r,\Delta)}{\partial r \partial \Delta} = \frac{\partial}{\partial r}\left(rV(r,\Delta)\right) = \underbrace{V(r,\Delta)}_{\text{Acceleration}(+)} \quad \underbrace{-r \cdot \text{Dur}(V(r,\Delta))}_{\text{Discounting}(-)} \tag{30}$$

Expression (30) spells out the economic forces determining whether cash flows are sped up or slowed down. A first effect, which we call "acceleration", is mechanical. Holding fixed the value of the asset, higher discount rates mean investors place more value on earlier cash inflows. 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 (30) immediately implies that the effect of discount rate shocks varies 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 of investment to be paid 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 bond's 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.

H Proofs and Additional Theoretical Results

H.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

$$0 = 1 + i(x,t)\frac{\partial}{\partial x} T(x,t) + \frac{\partial}{\partial t} T(x,t) \tag{31}$$

with terminal condition

$$T(\overline{x}, t) = 0 \tag{32}$$

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

Proposition 12. *In steady state, completion times are uniformly distributed.*

Proof. The distribution of completion times is

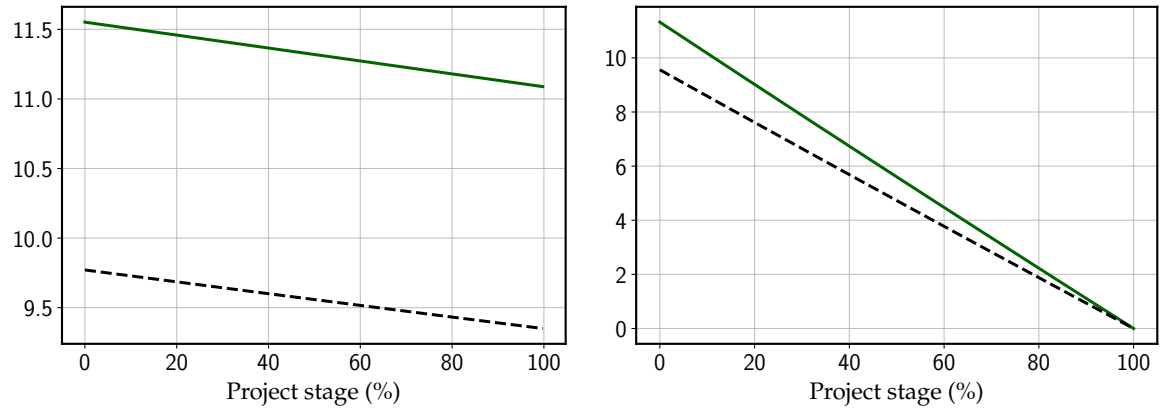
$$F_T(t) = \text{Prob}(T(x) \leq t) = \int_{T^{-1}(t)}^{\bar{x}} g(x)dx$$

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

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

Figure A11: Steady state objects as functions of project stage



(a) Stationary distribution $g(x,z)/f(z)$ (b) Time until completion $T(x,z)$

Note: the figures plot the steady state solution of the model. Panel (a) plots the stationary distribution $g(x,z)/f(z)$, and panel (b) plots the time to completion $T(x,z)$. 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 5: $r = 0.01$, $p = 1$, $\bar{x} = 1$, $\bar{\imath} = 0.083$, $\kappa = 3.09$.

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

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

Figure A11 shows the stationary distribution and the time to completion as a function of the project stage. Panel (a) 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. Panel (b) shows the time until completion $T(x,z)$, which is higher for low-productivity projects until reaching zero at $x = \bar{x}$.

□

H.2 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 an 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 13. *The first-order response of project value to an AR(1) shock to terminal payoff is given by*

$$\hat{v}(x,z) = pze^{-(r+\rho)T(x,z)}, \tag{33}$$

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

$$\hat{i}(x,z) = \frac{1}{\kappa''(i(x,z))}\frac{\partial \hat{v}(x,z)}{\partial x} \tag{34}$$

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

$$\hat{v}(x,z,t) \coloneqq \frac{\partial}{\partial s}\mathcal{V}(x,z,s,t)\Big|_{s=0}$$

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 $\hat{v}(x,z) \coloneqq \frac{\partial}{\partial s}\mathcal{V}(x,z,s,t)\Big|_{s,t=0}$ 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 (33).

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 $\hat{i}(x,z,t) \coloneqq \frac{\partial}{\partial s}\mathcal{I}(x,z,s,t)\Big|_{s=0}$. 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 $\hat{i}(x,z) \coloneqq \frac{\partial}{\partial s}\mathcal{I}(x,z,s,t)\Big|_{s,t=0}$. Cancelling the exponential term, we readily obtain (34). □

Proposition 13 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 to the terminal payoff decays at rate ρ, 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 $\rho > 0$, the shock is temporary, so we raise the discount rate to account for the fading dynamics.

Most crucially, the 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 an 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 13, we find that investment rises (falls) fol-

lowing 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 follows 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 14 shows that a positive terminal payoff shock lowers the entry threshold and therefore boosts project starts, consistent with standard investment theory.

Proposition 14. *The first-order response of the entry threshold $\bar{z}$ to an AR(1) terminal payoff shock is given by*

$$\hat{z} = -\bar{z}e^{-\rho T(0,\bar{z})} \tag{35}$$

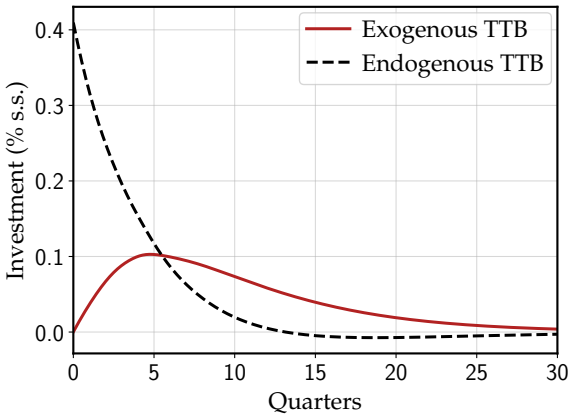
Proof. We can follow the same steps as in the proof of Proposition 5. Using (33) and (46), we find

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

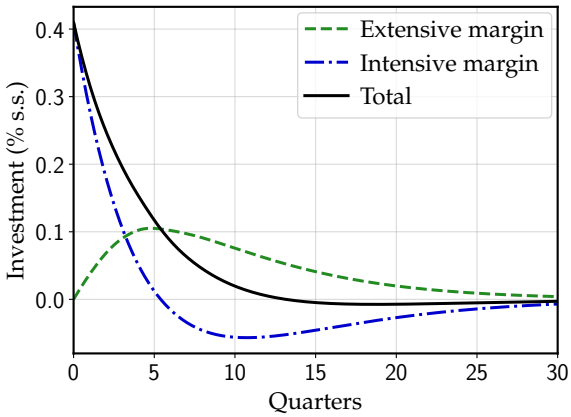
□

Figure A12 illustrates our analytical results using our quantitative model. We consider an AR(1) perfect-foresight increase in terminal payoffs p, holding fixed discount rates. One can interpret the price shock, for example, as a productivity shock under a competitive market structure and constant returns to scale. In the exogenous world, the effect is strictly positive, as the positive payoff shock raises entry. In the endogenous world, there is a short-run spike in investment: firms rush to complete projects while returns are higher. 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.

Figure A12: 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 A24.

H.3 Endogenous p

In our baseline model, we restricted attention to an exogenous terminal payoff pz. 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. This meant that p was independent of changes in r. We now relax this assumption and show that our main analytical results still hold.

If the plant must be operated after construction, the cash flows accruing to the project after its completion should also respond to changes in discount rates. Suppose that the completed project depreciates at a constant rate $\delta > 0$ and generates exogenous cash flows $c > 0$. The present value of operating the project is given by

$$p(t) = \int_t^\infty c e^{-\int_t^\tau (r(u)+\delta)du} d\tau$$

so that $p = {}^{c}/_{(r+\delta)}$ in steady state. We normalize cash flows c such that, for any price p and steady state discount rate r, we have $c = p(r+\delta)$.

Consider again a perfect foresight, AR(1) discount rate shock $r(t) = r + se^{-\rho t}$. The following lemma computes the change in $p(t)$ resulting from this shock.

Lemma 15. *The first-order response of the present value of operating the project is given by*

$$\frac{\hat{p}(t)}{p} = -\frac{e^{-\rho t}}{r+\delta+\rho} \tag{36}$$

Proof. After the interest rate shock,

$$\begin{aligned} p(t) &= c\int_t^\infty \exp\left[-\left((r+\delta)(\tau-t) + s\int_t^\tau e^{-\rho u}du\right)\right]d\tau \\ &= c\int_t^\infty \exp\left[-\left((r+\delta)(\tau-t) + s\frac{1}{\rho}(e^{-\rho t}-e^{-\rho\tau})\right)\right]d\tau \end{aligned}$$

Taking the derivative with respect to s and evaluating at $s=0$, we obtain

$$\begin{aligned} \hat{p}(t) &= -c\int_t^\infty \frac{1}{\rho}(e^{-\rho t}-e^{-\rho\tau})e^{-(r+\delta)(\tau-t)}d\tau \\ &= -\frac{c}{\rho}e^{-\rho t}\int_t^\infty \left(e^{-(r+\delta)(\tau-t)} - e^{-(r+\delta+\rho)(\tau-t)}\right)d\tau \\ &= -\frac{c}{\rho}e^{-\rho t}\left(\frac{1}{r+\delta} - \frac{1}{r+\delta+\rho}\right) \\ &= -\underbrace{\frac{c}{(r+\delta)}}_{=p}\frac{e^{-\rho t}}{(r+\delta+\rho)} \end{aligned}$$

□

With this result, we can extend Proposition 3 to the case where p depends directly on r.

Proposition 16. *The first-order response of the optimal investment policy is given by*

$$\hat{i}(x,z) = \frac{1}{i(x,z)\cdot\kappa''(i(x,z))}\cdot\left[\underbrace{v(x,z)+(r+\rho)\hat{v}(x,z)}_{\text{Fixed } p} - \underbrace{\frac{r+\rho}{r+\delta+\rho}e^{-(r+\rho)T(x)}pz}_{\text{Price Response}}\right], \tag{37}$$

Proof. As p now depends on r, we have, by the chain rule,

$$\hat{i}(x,z) = \hat{i}_r(x,z) + \hat{i}_p(x,z)\cdot\hat{p}$$

We have already computed all these derivatives. First, $\hat{i}_r(x,z)$ is the response of investment with respect to r keeping p fixed, which is given by (15) and (16) in Proposition 3. Second,

$\hat{i}_p(x,z)$ is the response of investment with respect to p keeping r fixed, which is given by (34) in Proposition 13. Last, $\hat{p}$ is given by (36) in Lemma 15. Substituting all expressions and simplifying yields the result. □

Expression 37 is composed of two terms. The first, "Fixed p", is the marginal value of accelerating the project, or reducing the project's duration, in the case where p is independent of r, i.e., 16. The second, "Price Response", captures the change in the value of operating the project due to changes in the discount rate. As discounting rises, future cash flows are less valuable, which reduces the incentives to complete the project and therefore pushes towards delaying project completion. Even with these headwinds, we can still prove that acceleration is optimal for projects sufficiently close to completion.

Corollary 17. *For x sufficiently close to $\bar{x}$, we have*

$$\lim_{x\to\bar{x}} \hat{i}(x,z) = \frac{1}{i(\bar{x},z)\cdot\kappa''(i(\bar{x},z))}\frac{\delta}{r+\delta+\rho}pz > 0 \tag{38}$$

Proof. Note that $T(x) \to 0$ as $x \to \bar{x}$. From the proof of Corollary 4, we know $\hat{v}(x,z) \to 0$ and $v(x,z) \to pz$. Simplifying yields the result. □

Therefore, when discount rates rise, it remains optimal to accelerate the pace of completion for projects sufficiently close to completion, *so long as the project depreciates at a positive rate*. The incentives to accelerate shrink, however, by a factor $\delta/r+\rho+\delta \in (0,1)$.

H.4 Risk

This section introduces risk into the endogenous time to build model from the main text, with the objective of understanding how this affects our main theoretical results. We find that introducing risk related to the value of a finished project does not affect the result that firms accelerate projects close to completion following an increase in their discount rate. More generally, we show that risk does not affect investment behavior close to completion.

Suppose that the project productivity z_t is now time-varying and follows a Poisson process. Let $z_1 < \cdots < z_J$ be the possible values for z_t and μ_{jk} be the rate at which a jump occurs from state z_j to state z_k, i.e.,

$$Prob(z_{t+\Delta} = z_k | z_t = z_j) = \mu_{jk}\Delta + o(\Delta).$$

The optimization problem of firms is now characterized by a set of value functions $V_j(x)$, corresponding to the value of having a project at stage x and productivity z_j, given by

$$r(t)V_j(x,t) = \max_{i\geq 0} -i - \kappa(i) + \frac{\partial V_j(x,t)}{\partial x}i + \frac{\partial V_j(x,t)}{\partial t} + \sum_{k\neq j}\mu_{jk}\left[V_k(x,t) - V_j(x,t)\right].$$

The terminal condition is now

$$V_j(\bar{x}, t) = pz_j.$$

Importantly, we assume that all firms immediately sell their projects upon completion. This rules out the possibility that a firm completes with very low productivity and then waits for its value to increase before selling it. Optimal investment is characterized by the first order condition:

$$1 + \kappa'(i_j(x,t)) = \frac{\partial V_j(x,t)}{\partial x}.$$

Now suppose, as in the main text, there is an exponentially decaying shock to firms' discount rate:

$$r(t) = r + se^{-\rho t}.$$

As in the case without risk, let variables with hats denote first-order perturbations relative to steady state. The derivation of results here follows the same steps as those in the model from the main text. We plug the time-varying discount rate into the HJB equation, take derivatives with respect to shock size s, and guess that $\hat{V}_j$ has the form $\hat{V}_j(x,t) = e^{-\rho t}\hat{V}_j(x)$ to obtain

$$(r+\rho)\,\hat{V}_j(x) + V_j(x) = i_j(x)\hat{V}_j'(x) + \sum_k \mu_{jk}\left(\hat{V}_k(x) - \hat{V}_j(x)\right),$$

$$\hat{V}_j(\bar{x}) = 0.$$

By defining $\mathcal{I}(x) = \text{diag}(i_1(x), \cdots, i_J(x))$, $\tilde{\mu}_j = \sum_{k \neq j} \mu_{jk}$ and

$$\mathcal{Q} = \begin{bmatrix} -\tilde{\mu_1} & \cdots & \mu_{1j} & \cdots & \mu_{1J} \\ \vdots & & \vdots & & \vdots \\ \mu_{j1} & \cdots & -\tilde{\mu}_j & \cdots & \mu_{jJ} \\ \vdots & & \vdots & & \vdots \\ \mu_{J1} & \cdots & \mu_{Jj} & \cdots & -\tilde{\mu_J} \end{bmatrix},$$

this equation can be written as

$$((r+\rho)I - \mathcal{Q})\,\hat{\mathbf{V}}(x) + \mathbf{V}(x) = \mathcal{I}(x)\hat{\mathbf{V}}'(x). \tag{39}$$

Above, $\mathbf{V} = (V_1(x), \ldots, V_J(x))'$, and $\hat{\mathbf{V}}$ and $\hat{\mathbf{V}}'$ are defined similarly.

To solve equation 39, consider first the matrix $\Psi(x)$ that solves the homogeneous system

$$((r+\rho)I - \mathcal{Q})\,\Psi(x) = \mathcal{I}(x)\Psi'(x).$$

$\Psi(x)$ is more precisely defined as a matrix whose columns consist of the J linearly indepen-

dent solutions of the homogeneous equations above. We can express $\hat{\mathbf{V}}(x)$ as

$$\hat{\mathbf{V}}(x) = -\Psi(x) \int_x^{\bar{x}} \Psi^{-1}(y)\mathcal{I}^{-1}(y)\mathbf{V}(y)dy,$$

which can be easily checked to satisfy the terminal condition $\hat{\mathbf{V}}(\bar{x}) = 0$.

Having solved for $\hat{\mathbf{V}}(x)$, define $\hat{\mathbf{I}}(x) = (\hat{i}_1(x), \ldots, \hat{i}_J(x))'$. Start from

$$\hat{i}_j(x) = \frac{1}{\kappa''(i_j(x))}\hat{V}_j'(x)$$

and use 39 to obtain

$$\hat{\mathbf{I}}(x) = \mathcal{I}_\kappa^{-1}\left[((r+\rho)I - \mathcal{Q})\,\hat{\mathbf{V}}(x) + \mathbf{V}(x)\right],$$

where $\mathcal{I}_\kappa = \text{diag}\left(i_1(x)\kappa_1''(i_1(x)), \ldots, i_J(x)\kappa_J''(i_J(x))\right)$. It is straightforward to see that, for x close enough to $\bar{x}$, each entry of $\hat{\mathbf{I}}(x)$ will be positive given that $\mathbf{V}(\bar{x})$ is component-wise positive and $\kappa'' > 0$. Hence, firms will optimally decide to speed up projects close enough to completion even in a setting with risk. More generally, since $\hat{\mathbf{V}}(\bar{x}) = 0$, risk will not influence optimal investment decisions for projects close to completion. The intuition is simple: when projects are sufficiently close to completion, the project behaves as if it is en route to deterministic completion, so that the role of risk vanishes.

H.5 Proofs from Main Text

H.5.1 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

$$r\frac{\partial}{\partial s}\mathcal{V}(x,s,t) + e^{-\rho t}\mathcal{V}(x,s,t) = \mathcal{I}(x,t,s)\frac{\partial}{\partial x}\frac{\partial}{\partial s}\mathcal{V}(x,t,s) + \frac{\partial}{\partial t}\frac{\partial}{\partial s}\mathcal{V}(x,t,s) \tag{40}$$

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

We define the steady-state gradients

$$\hat{v}(x,t) \coloneqq \frac{\partial}{\partial s}\mathcal{V}(x,s,t)\Big|_{s=0}$$

so evaluating (40) 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

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

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

$$\begin{aligned} (r+\rho)\hat{v}(x) + v(x) &= i(x)\frac{d}{dx}\hat{v}(x) \\ \hat{v}(\overline{x}) &= 0 \end{aligned} \tag{41}$$

Step 2. Solving the ODE for $\hat{v}(x)$ We rearrange (41) 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

$$\begin{aligned} -e^{(r+\rho)T(x)}\hat{v}(x) &= \int_x^{\overline{x}} \frac{e^{(r+\rho)T(y)}}{i(y)}v(y)dy \\ \Rightarrow \hat{v}(x) &= -\int_x^{\overline{x}} \frac{e^{-(r+\rho)(T(x)-T(y))}}{i(y)}v(y)dy \end{aligned} \tag{42}$$

Step 3. Changes of Variables The key step in simplifying expression (42) is to move the right-hand side integral from stage space to time space. Let

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

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

$$\hat{v}(x) = -\int_0^{T(x)} e^{-(r+\rho)\tau} v(y(\tau)) d\tau \tag{43}$$

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,

$$v(y) = -\int_y^{\bar{x}} \frac{e^{-r[T(y)-T(u)]}}{i(u)} \left[i(u) + \kappa\left(i(u)\right) \right] du + e^{-rT(y)} pz \tag{44}$$

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 (44) into (43), we obtain

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

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

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

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,

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

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

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

So bringing all together yields (14) and completes the proof.

H.5.2 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 $\hat{i}(x,t) \coloneqq \frac{\partial}{\partial s} \mathcal{I}(x,s,t) \Big|_{s=0}$. Once again, we conjecture $\hat{i}(x,t) = e^{-\rho t} \hat{i}(x)$ (and analogously for $\hat{v}(x,t)$), where as before $\hat{i}(x) \coloneqq \frac{\partial}{\partial s} \mathcal{I}(x,s,t) \Big|_{s,t=0}$. Thus,

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

Differentiating (42) 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

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

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 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 $\hat{u}(t) := \partial u(s,t)/\partial s\big|_{s=0}$, 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 $\hat{u} := \partial u(s,t)/\partial s\big|_{s,t=0}$, 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.

H.5.3 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

$$\frac{\partial \bar{\mathcal{Z}}(t,s)}{\partial s} = -\frac{\partial \mathcal{V}/\partial s}{\partial \mathcal{V}/\partial z}\bigg|_{(x,z)=(0,\bar{\mathcal{Z}}(t,s))} \tag{45}$$

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

$$\frac{\partial v(x,z)}{\partial z} = pe^{-rT(x,z)} \tag{46}$$

satisfies the ODE and the terminal condition.

Returning to (45) and evaluating at a steady state, we have

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

which is (17) after rearranging.

H.5.4 Proof of Proposition 6

The shock is unanticipated and the distribution of ongoing projects evolves continuously through the KFE, so $g(x,z,0) = g(x,z)$: neither entry nor completion changes the stock of ongoing projects at the instant of the shock. Differentiating $I^{acc}(0) = \int\int g(x,z)\,\mathcal{I}(x,z,s,0)\,dx\,dz$ with respect to s at $s = 0$ gives

$$\hat{I}^{acc}(0) = \int\int g(x,z)\hat{i}(x,z)\,dx\,dz = \int_{\bar{z}}^{\infty} f(z)\int_0^{\bar{x}}\frac{\hat{i}(x,z)}{i(x,z)}\,dx\,dz,$$

where the second equality substitutes the steady-state distribution (12). Dividing by $I_{ss}^{acc} = M_e \bar{x}$ and using the proportional project response $\hat{i}/i = \Lambda [1 - (r + \rho)D]$ together with the definitions of $\bar{\Lambda}$ and D^{Λ} yields equation (18). The mass of ongoing projects is predetermined at impact, so average investment per ongoing project responds by the same proportion. The sign statement follows from $\bar{\Lambda} > 0$.

H.5.5 Proof of Proposition 7

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 projects, 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 to show that 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.

Work-in-Progress The two paragraphs above establish the result type by type, with entry fixed at $f(z)$. The general statement keeps entry free. Differentiating equation (19) with respect to time and using the time-dependent KFE, $\partial_t g + \partial_x[i\,g] = 0$, integration by parts in x gives

$$\begin{aligned}
\dot{W}(t) &= -\int \int_0^{\bar{x}} (\bar{x} - x) \frac{\partial}{\partial x} \left[i(x,z,t)g(x,z,t) \right] dx\, dz \\
&= -\int \left\{ \Big[(\bar{x} - x)\, i\, g \Big]_{x=0}^{x=\bar{x}} + \int_0^{\bar{x}} i(x,z,t)g(x,z,t)\, dx \right\} dz \\
&= \bar{x} \int b(z,t)\, dz - I^{acc}(t),
\end{aligned}$$

which is equation (20). The boundary term vanishes at $x = \bar{x}$, because a project exits exactly when its remaining work reaches zero, and equals $-\bar{x}\, b(z,t)$ at $x = 0$.

For (i), subtract the steady-state version of equation (20), $0 = \bar{x}B_{ss} - I_{ss}^{acc}$, integrate from 0 to ∞, and use $W(0) = W_{ss}$ (the distribution is predetermined) and $\lim_{t\to\infty} W(t) = W_{ss}$ (the shock is temporary).

For (ii), consider the cohort of projects ongoing at $t = 0$ and let $g^0(x,z,t)$ denote its density. It obeys the same KFE with no entry, so its work-in-progress $W^0(t)$ satisfies $\dot{W}^0(t) = -I^0(t)$, where $I^0(t)$ is the cohort's accounting investment. Since $W^0(0) = W(0)$ and every project in the cohort eventually completes, $\lim_{t\to\infty} W^0(t) = 0$ and $\int_0^\infty I^0(t)dt = W(0)$. This holds for any path of the shock, including no shock, so the cumulative response of the cohort's investment is zero.

H.5.6 Proof of Corollary 8

By Proposition 6, $\hat{I}^{acc}(0) > 0$. By Proposition 5, $\hat{z} > 0$, since $\hat{v}(0,\bar{z}) < 0$. To first order, project starts are $B(t) = \int_{\bar{\mathcal{Z}}(t,s)}^\infty f(z)dz$, so their response is $-f(\bar{z})\hat{z}e^{-\rho t} < 0$ and, by equation (21),

$$\int_0^\infty \hat{I}^{acc}(t)\, dt = -\frac{\bar{x} f(\bar{z})\hat{z}}{\rho} < 0.$$

A response that is positive at $t = 0$, continuous, and integrates to a negative number must be negative on a set of dates of positive measure.

H.6 Duration and the Acceleration Threshold

Proposition 3 expresses the investment response as the difference between an acceleration motive, $v(x,z)$, and a duration effect, $(r+\rho)\hat{v}(x,z)$. The two can be combined into a single condition. Recall the project's proportional duration with respect to the shock and its acceleration responsiveness,

$$D(x,z) \equiv -\frac{\hat{v}(x,z)}{v(x,z)}, \qquad \Lambda(x,z) \equiv \frac{v(x,z)}{\kappa''(i(x,z))\, i(x,z)^2},$$

so that, by equations (15)–(16),

$$\frac{\hat{i}(x,z)}{i(x,z)} = \Lambda(x,z)\left[1-(r+\rho)D(x,z)\right].$$

Proposition 18. *For any ongoing project with $v(x,z) > 0$ at an interior optimum with $\kappa''(i(x,z)) > 0$,*

$$\hat{i}(x,z) > 0 \quad \Longleftrightarrow \quad D(x,z) < \frac{1}{r+\rho}.$$

Proof. Since $v > 0$, $\kappa'' > 0$, and $i > 0$, we have $\Lambda > 0$, and the sign of $\hat{i}$ in the display above is the sign of $1-(r+\rho)D$. □

The condition parallels the pricing of a bond: a rise in the discount rate lowers the value of any claim, but the loss is small when duration is short. A project accelerates when its duration is shorter than $1/(r+\rho)$, the duration of a perpetuity discounted at the shock-adjusted rate. Since $\hat{v}(\bar{x},z) = 0$, duration vanishes at completion, which is Corollary 4. The magnitude of the response depends on Λ (how much value is at stake, how cheaply the pace of construction can be changed, and on the level of investment being scaled), but its sign does not.

The duration condition signs the response point by point. We can further show that the acceleration region is an interval. Define

$$G(x,z) \equiv v(x,z) + (r+\rho)\hat{v}(x,z),$$

so that, by Proposition 3, $\hat{i}(x,z) = G(x,z)/\left[\kappa''(i(x,z))\, i(x,z)\right]$ and $\hat{i}$ has the sign of G.

Proposition 19. *Fix a type z with $v(x,z) > 0$ for all $x \in [0,\bar{x}]$, and suppose $\kappa''(i(x,z)) > 0$ along the project's path. Either $\hat{i}(x,z) > 0$ at every stage, or there exists a unique $x^*(z) \in (0,\bar{x})$ such that*

$$\hat{i}(x,z) < 0 \text{ for } x < x^*(z), \qquad \hat{i}(x,z) > 0 \text{ for } x > x^*(z).$$

Equivalently, in terms of remaining time to build, with $T^*(z) \equiv T(x^*(z),z)$ and the convention

$T^*(z) = T(0,z)$ *when every stage accelerates,*

$$\hat{i} > 0 \text{ for } T(x,z) < T^*(z), \qquad \hat{i} < 0 \text{ for } T(x,z) > T^*(z).$$

Proof. Differentiate the definition of G in x and substitute equation (16) for $\partial\hat{v}/\partial x$:

$$G_x(x,z) = v_x(x,z) + \frac{r+\rho}{i(x,z)}\, G(x,z).$$

At any stage where $G(x,z) = 0$, this gives $G_x(x,z) = v_x(x,z) = 1 + \kappa'(i(x,z)) > 0$ by the first-order condition. Hence G can only cross zero from below as x increases. A continuous function that crosses zero only upward has at most one zero, and once positive it stays positive. At completion, $\hat{v}(\bar{x},z) = 0$ and $G(\bar{x},z) = v(\bar{x},z) = pz > 0$, so G is positive on an interval below $\bar{x}$. Therefore either $G > 0$ on $[0,\bar{x}]$, or G has exactly one zero $x^*(z)$ with $G < 0$ below it and $G > 0$ above it. The statement in remaining time follows from $T_x = -1/i < 0$. □

Proposition 19 says that, for each type, the acceleration region is the set of projects within $T^*(z)$ of completion, and the slowdown region is its complement. The numerical threshold at about a third of total construction in Figure 3 is an instance of this result. It is not a feature of the quadratic cost function, since the argument uses only $v_x > 0$ at an interior optimum. Appendix D.3 bounds $T^*(z)$ from above by a horizon that depends only on r and ρ.

H.7 Discount Rates

We mapped the empirical shock into an increase in firms' discount rates. This appendix elaborates on our modeling choice. First, it reviews the several margins through which banking distress raises firms' discount rates, together with supporting evidence. Second, it sketches a brief microfoundation that would directly connect the provisioning shock to higher discount rates.

Our starting point is a comprehensive view of the discount rate: the required return that firms apply to project cash flows in capital budgeting. Graham (2022) documents that explicit hurdle rates are the near-universal tool with which firms evaluate projects. Gormsen and Huber (2025) measure these discount rates directly from corporate conference calls. They show that discount rates respond to financial conditions, carry time-varying wedges relative to the cost of capital, and have quantitatively large effects on investment. Discount rates thus aggregate many margins, including the cost of debt and equity, the shadow value of internal funds, and managerial practice. This is why we model the credit disruption as a discount rate shock. The shock summarizes how banking turmoil raises firms' valuation of cash today relative to cash tomorrow, without restricting the pass-through to any single

financing margin.

The most direct margin is the cost of debt financing, which banks influence via lending rates. Hence, turbulence in the banking sector would translate into more expensive debt financing and higher discount rates. A large empirical literature documents this pass-through across many contexts and traces its mechanisms. Troubled banks pass higher funding costs into loan prices, and sticky lending relationships and adverse selection make it costly for borrowers to switch away from affected lenders.[74]

Morais et al. (2020) provide direct causal evidence that shocks that increase bank loan loss provisions pass through to lower borrowing, shorter maturity, and higher lending rates. They study the introduction of expected-loss provisioning requirements in Colombia in the mid-2000s. The reform made changes to provisioning requirements based on arbitrary firm-size thresholds, thereby bolstering identification claims; and was implemented during a period of economic stability and growth, similar to the macroeconomic context of India at the time of the AQR. Following the reform, affected borrowers faced higher lending rates, lower loan volumes, and shorter maturities; these effects were economically sizable and materialized within a year of implementation.

As an illustration of the cost-of-debt margin, it is possible to microfound the pass-through from shocks to loan loss provisions into lending rates as follows.[75] Consider a setup with multiple banks and firms as in Herreño (2023), which delivers (i) a downward-sloping demand curve for credit and (ii) lending shares across banks for each firm. Lending is risky: a fraction of loans become non-performing and may default. Banks set lending rates not only based on loan demand and default risk, but also on provisioning rules (Bouvatier and Lepetit, 2012). The marginal loan increases required provisions and reduces bank earnings (Walter, 1991), so optimal loan pricing consists of a risk-adjusted mark-up over funding costs (i.e., deposit rates or interbank market loans) plus the provisioning opportunity costs. Regulators can only imperfectly detect underprovisioning. Banks thus hide NPAs to reduce provisioning costs, offer more competitive lending rates, and gain market share. The Indian Asset Quality Review is a monitoring shock that forces banks to put up provisions, which then increases lending rates. This sketch formalizes a single margin because it is the most direct to model; it is an illustration of how the provisioning shock raises discount rates, not a claim that lending rates are the only conduit.

Discount rates, however, also respond through margins that do not operate via lending

[74] See Hubbard et al. (2002) for early cross-sectional evidence in the US; Chava and Purnanandam (2011) for the US following the 1998 Russian crisis; Santos (2010) and Chodorow-Reich (2014) for the US following the collapse of Lehman Brothers and the GFC; Sette and Gobbi (2015) and Bonaccorsi di Patti and Sette (2016) for Europe; Abuka et al. (2019) for a developing economy context; and Hale et al. (2019) for cross-border evidence. On the mechanisms – funding-cost pass-through, costly lender switching, relational contracting, and adverse selection – see Stiglitz and Weiss (1981), Sharpe (1990), Sette and Gobbi (2015), Bolton et al. (2016), Banerjee et al. (2021), and Jiménez et al. (2017).

[75] For brevity, we describe only the model setup and the main result.

rates at all. The bank shock can raise the value of cash inside the firm by reducing credit supply. The discount rate shock partly captures the higher value of cash, as more recent cash flows become more valuable relative to delayed cash flows when discount rates rise. This margin operates even for firms that never reprice or renew a bank loan during the episode: scarcer credit raises the precautionary value of liquidity, and with it the required return on sinking funds into long-gestation projects. Another modeling device used to capture this channel is firm-level borrowing constraints. We elaborate on our choice of discount rates versus borrowing constraints in Appendix H.8. Last, higher discount rates may reflect behavioral or organizational frictions. Managers may become more risk-averse in times of scarce credit; or CFOs might use the shock as a disciplining device for internal capital markets, forcing managers to adopt more stringent discount rates. These channels are not merely theoretical possibilities: Gormsen and Huber (2024) show that firms' perceived cost of capital deviates substantially from measured financing costs; and Graham (2022) documents that hurdle rates embed strategic and organizational considerations well beyond the cost of funds. We therefore treat the discount rate as a summary statistic for how the credit disruption changes the relative price of present versus future cash flows inside the firm – encompassing, rather than adjudicating among, these complementary channels. This aggregation also matters for the quantification: the model is disciplined by matching our estimated investment responses, so its conclusions do not hinge on which margin delivers the rise in discount rates.

H.8 Spending Constraints

In the main text, we modeled the credit shock as an increase in the discount rate. An alternative approach maps the credit shock to a tightening of the firm's ability to borrow. In this appendix, we sketch such a model. We then argue in favor of the discount rate shock as the preferred channel to rationalize our empirical findings.

Consider a firm building $N \geq 1$ projects. The discreteness simplifies exposition and is without loss of generality. The building technology is the same as in the main text. There is now a spending constraint. If a firm has N projects, it can invest up to E_N across all projects. To fix ideas, consider a firm with a single project ($N = 1$). Optimal investment $i \geq 0$ solves

$$\begin{aligned} rv(x) &= \max_i -[i + \kappa(i)] + iv'(x) \\ \text{s.t.} & i + \kappa(i) \leq E_1 \\ & v(\bar{x}) = p \end{aligned}$$

Unconstrained investment, as before, equates the marginal cost $1 + \kappa'(i^*(x))$ to the marginal

value $v'(x)$. If investment is constrained, the firm sets

$$\bar{i}(x) + \kappa\left(\bar{i}(x)\right) = E_1$$

so that optimal investment is given by $i(x) = \max\left\{0, \min\left\{i^*(x), \bar{i}(x)\right\}\right\}$.

Suppose instead the firm were building two projects. Denote by x_n the stage of project $n = 1, 2$. The firm now solves

$$\begin{aligned} rV(x_1, x_2) = \max_{i_1, i_2} & -\left[i_1 + \kappa(i_1) + i_2 + \kappa(i_2)\right] + i_1 \frac{\partial V}{\partial x_1} + i_2 \frac{\partial V}{\partial x_2} \\ \text{s.t. } & i_1 + \kappa(i_1) + i_2 + \kappa(i_2) \leq E_2 \end{aligned}$$

Note that E_2 can be different from E_1. For example, once a project is finished, available capital inside the firm increases and could relax the spending constraint. That would be the case if (i) the windfall from project completion became cash that could be immediately spent; or if (ii) completion of the project relaxed borrowing constraints through a collateral-based or earnings-based channel. In such cases, $E_2 < E_1$. Alternatively, ex ante budget allocations could increase with the number of projects, making $E_2 > E_1$.

The terminal conditions for each project are

$$V(x_1, \bar{x}) = p + v(x_1)$$

$$V(\bar{x}, x_2) = p + v(x_2)$$

The first condition states that the value of having project 2 completed is the return to project 2 plus the continuation value of still having project 1 under implementation. The second condition states the same when project 1 is completed first. For expositional simplicity, we assume equal returns across projects. The necessary conditions for optimal investment are now

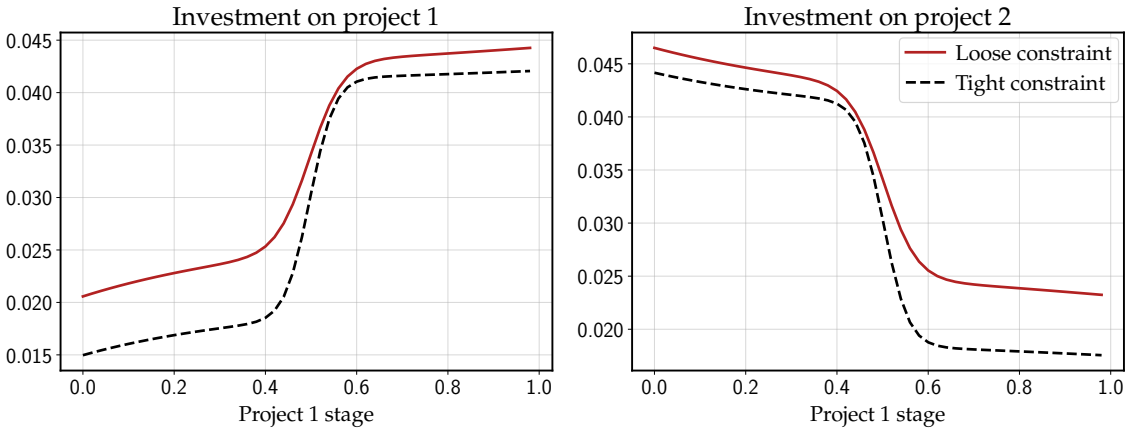
$$1 + \kappa'(i_j) = \frac{1}{1 + \lambda_j} \frac{\partial V}{\partial x_j}$$

where λ_j is the Lagrange multiplier on the spending constraint. We index λ by j to make explicit the dependence on the total budget E_j. Optimal capital budgeting thus implies equalizing the ratio of marginal values to marginal costs across projects.

To illustrate the economics, we resort to a numerical example. We assume $E_1 > E_2$, so completing one project *relaxes* the spending constraint. We fix project 2 at 50% completion. We then ask how investment in each project varies as a function of project 1's stage under two scenarios: tight vs loose spending constraints. Appendix Figure A13 presents the results.

As in the model without spending constraints, investment in project 1 is an increasing function of its own stage, independent of the constraint. With the aggregate spending con-

Figure A13: Optimal Investment with Spending Constraints



Notes: the figure plots the optimal investment rate on each project as a function of project 1's stage, fixing the stage of project 2 at 50%. Parameter values for the numerical example are $\bar{x}_1 = \bar{x}_2 = z_1 = z_2 = 1$ and $\kappa(i) = \kappa i^2/2$. The solid (dashed) line corresponds to a loose (tight) spending constraint.

straint, investment in project 2 becomes a decreasing function of project 1 when the constraint binds. The intuition is simple: the more advanced project 1 is, the higher is the marginal value of investing in project 1 relative to project 2. The firm prefers to prioritize the project that is closer to completion in order to receive the terminal cash flows. When financial constraints are tighter, the investment schedule shifts down across both projects, reflecting the shared factor $(1 + \lambda_j)$ increasing the marginal cost of investment.

Suppose we start from the loose scenario. A credit shock then lowers the spending limit, moving us to the tight scenario. For illustration, suppose this is a permanent, unanticipated shock. Aggregate investment must mechanically drop. In the cross-section, both projects cut investment. As such, the "average" project within the firm is scaled back. The investment cut for the average project is the empirical counterpart of our regression estimate in this model. As such, a model with spending constraints *cannot* rationalize the robustly positive hazard effects on the average investment project. What spending constraints *can* rationalize are the heterogeneous hazard effects across projects of different stages.

We refrained from writing down a firm simultaneously optimizing (i) its investment across multiple projects and (ii) its aggregate capital structure subject to (iii) occasionally binding borrowing constraints. Such an approach would require three features. First, we would need to keep track of the entire distribution of projects, as we currently do. Second, we would need to solve for policies under an aggregate, occasionally binding constraint on total investment.[76] Third, we would need to solve for the resulting transitional dynamics following a temporary shock to the constraint. Individually, each task is manageable. But

[76] An easier, though economically questionable approach would be to specify borrowing constraints at the project level, as if such projects could be individually collateralized. Project finance, however, is still nascent in India. Most of the corporate loans are not secured against specific assets.

the interaction of these three elements tremendously increases the model's computational complexity.[77]

Taking stock, we see four reasons why we prefer to map the credit disturbance to a discount rate shock. First, it illustrates that the empirical patterns we observe – and therefore the basic theory of endogenous time to build – could arise from a neoclassical benchmark and a more nuanced understanding of the time value of money. Our empirical and theoretical results are not predicated on the particular form that financing frictions might take (e.g., asset- vs earnings-based).

Second, and perhaps most pragmatically, we make a computationally tractable choice that allows us to carry out meaningful estimation and counterfactual exercises while taking full advantage of the rich underlying heterogeneity in the project-level data. Third, we have sketched how the pass-through of the credit shock to the firm's discount rate could be microfounded without appealing to borrowing constraints at the firm level – though that need not be the only channel through which credit shocks would increase firms' discount rates. Last, from a more epistemological viewpoint, we appeal to Occam's razor.

This is not to say that interactions between endogenous time to build and firm-level financing frictions are not important. We believe they are an exciting area for further work: capital budgeting often requires managers to screen across multiple projects precisely due to organizational constraints on divisional spending. We hope our work ignites further progress in that direction.

Online Appendix References

Abuka, C., Alinda, R. K., Minoiu, C., Peydró, J.-L., and Presbitero, A. F. (2019). Monetary policy and bank lending in developing countries: Loan applications, rates, and real effects. *Journal of Development Economics*, 139:185–202.

Banerjee, R. N., Gambacorta, L., and Sette, E. (2021). The real effects of relationship lending. *Journal of Financial Intermediation*, 48.

Bolton, P., Freixas, X., Gambacorta, L., and Mistrulli, P. E. (2016). Relationship and transaction lending in a crisis. *Review of Financial Studies*, 29(10):2643–2676.

Bolton, P., Wang, N., and Yang, J. (2019). Investment under uncertainty with financial constraints. *Journal of Economic Theory*, 184:104912.

[77] Indeed, models with liquidity or borrowing frictions like Bolton et al. (2019) study joint optimization of investment and capital structure, but abstract from time to build and do not model the dynamic evolution of an entire distribution of projects.

Bonaccorsi di Patti, E. and Sette, E. (2016). Did the securitization market freeze affect bank lending during the financial crisis? evidence from a credit register. *Journal of Financial Intermediation*, 25:54–76.

Bouvatier, V. and Lepetit, L. (2012). Provisioning rules and bank lending: A theoretical model. *Journal of Financial Stability*, 8(1):25–31.

Chava, S. and Purnanandam, A. (2011). The effect of banking crisis on bank-dependent borrowers. *Journal of Financial Economics*, 99(1):116–135.

Chodorow-Reich, G. (2014). 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.

Davis, S. J. and Haltiwanger, J. (1992). Gross job creation, gross job destruction, and employment reallocation. *The Quarterly Journal of Economics*, 107(3):819–863.

Gormsen, N. J. and Huber, K. (2024). Firms' perceived cost of capital. *Working Paper*.

Gormsen, N. J. and Huber, K. (2025). Corporate discount rates. *American Economic Review*, 115(6):2001–2049.

Graham, J. R. (2022). Presidential address: Corporate finance and reality. *Journal of Finance*, 77(4):1975–2049.

Hale, G., Kapan, T., and Minoiu, C. (2019). Shock transmission through cross-border bank lending: Credit and real effects. *The Review of Financial Studies*, 33(10):4839–4882.

Hart, O. and Moore, J. (1988). Incomplete contracts and renegotiation. *Econometrica*, 56(4):755–785.

Hassan, T. A., Hollander, S., Kalyani, A., van Lent, L., Schwedeler, M., and Tahoun, A. (2025). Text as data in economic analysis. *Journal of Economic Perspectives*, 39(3):193–220.

Hernán, M. A. and Robins, J. M. (2020). *Causal Inference: What If*. Chapman & Hall/CRC, Boca Raton.

Herreño, J. (2023). Aggregating the effect of bank credit supply shocks on firms. Working paper, UC San Diego.

Hubbard, R., Kuttner, K., and Palia, D. N. (2002). Are there bank effects in borrowers' costs of funds? evidence from a matched sample of borrowers and banks. *The Journal of Business*, 75(4):559–81.

Jiménez, G., Ongena, S., Peydró, J.-L., and Saurina, J. (2017). Macroprudential policy, countercyclical bank capital buffers, and credit supply: Evidence from the Spanish dynamic provisioning experiments. *Journal of Political Economy*, 125(6):2126–2177.

Morais, B., Ormazabal, G., Peydró, J.-L., Roa, M., and Sarmiento, M. (2020). Forward looking loan provisions: Credit supply and risk-taking. CEPR Discussion Paper 15278, Centre for Economic Policy Research.

Santos, J. A. C. (2010). Bank corporate loan pricing following the subprime crisis. *The Review of Financial Studies*, 24(6):1916–1943.

Sette, E. and Gobbi, G. (2015). Relationship lending during a financial crisis. *Journal of the European Economic Association*, 13(3):453–481.

Sharpe, S. A. (1990). Asymmetric information, bank lending, and implicit contracts: A stylized model of customer relationships. *Journal of Finance*, 45(4):1069–1087.

Stiglitz, J. E. and Weiss, A. (1981). Credit rationing in markets with imperfect information. *American Economic Review*, 71(3):393–410.

Walter, J. R. (1991). Loan loss reserves. *Federal Reserve Bank of Richmond Economic Review*, 77(4):20–30.