Investment Under Upstream and Downstream Uncertainty

FOTIS GRIGORIS and GILL SEGAL[⋆]

ABSTRACT

The impact of uncertainty shocks on firm-level economic activity depends on their origin in supply chains. Upstream (downstream) uncertainty from suppliers (customers) is associated with variability over future input (output) prices. Consequently, a real-option production model with time-to-build suggests that only upstream uncertainty suppresses investment, since upstream (downstream) uncertainty affects the shorter (longer) run. Production network data show that upstream uncertainty negatively affects firm-level outcomes. Conversely, downstream uncertainty affects firm-level outcomes more weakly but positively. At the macro level, these two uncertainties oppositely predict aggregate growth and asset prices. Overall, downstream uncertainty has an expansionary effect, in contrast to other facets of uncertainty.

JEL classification: E32, D81, D50, L14, G11, G12

Keywords: Production Network, Macro-Finance, Uncertainty, Investment, Real Option, Time to Build

---

[⋆]Fotis Grigoris is at Tippie College of Business, University of Iowa. Gill Segal (corresponding author, Gill_Segal@kenan-flagler.unc.edu) is at Kenan-Flagler Business School, University of North Carolina at Chapel Hill. We are indebted to Editor Urban Jermann, an anonymous Associate Editor, and two anonymous referees for invaluable comments. For helpful comments and suggestions, we also thank Vikas Agrawal, Hengjie Ai (discussant), Ravi Bansal, Matteo Binfarè, Nick Bloom, Daniel Carvalho, Bhagwan Chowdhry, Riccardo Colacito, Andrea Eisfeldt, Jack Favilukis, Michael Gallmeyer (discussant), Jon Garfinkel, Lorenzo Garlappi, Naveen Ghondi (discussant), Christian Heyerdahl-Larsen, Lixin Huang (discussant), Philipp Illeditsch, Rustom Irani (discussant), Ravi Jagannathan, Gustavo Joaquim (discussant), Shane Johnson, Preetesh Kantak, Lars-Alexander Kuehn (discussant), Erik Lie, Sydney Ludvigson, Michael O'Doherty, Richard Peter, Pavel Savor (discussant), Petra Sinagl, Michael Woeppel, Wei Wu (discussant), Keer Yang (discussant), Nancy Xu, and Michael Young, and seminar and conference participants at Indiana University, The University of Iowa, The University of Missouri, The University of Western Australia, Korean Business School, North Carolina State University, University of Virginia, Penn State University, the 2024 econometric society meeting, the 2022 Midwest Finance Association Annual Meeting, the 2022 Young Scholar Finance Consortium, the 2022 Conference on Financial Economics and Accounting, the 2022 Financial Management Association Annual Meeting, the 2022 Eastern Finance Association Annual Meeting, the 2021 Indian Business School Summer Research Conference, the 2021 Northern Finance Association Annual Meeting, the 2021 European Economic Association Annual Meeting, the 2021 Wabash River Finance Conference, the 2021 New Zealand Finance Meeting, the 34th AFBC, the 2021 Financial Markets and Corporate Governance Conference, the 2023 SITE - Asset Pricing conference, and the 2023 SITE - Uncertainty and Volatility conference. We have read *The Journal of Finance* disclosure policy and have no conflicts of interest to disclose. All errors are our own.

The global economy experienced unprecedented uncertainty during the Great Recession and the COVID-19 pandemic. Accordingly, a growing body of literature in macroeconomics and finance documents that increased uncertainty is typically associated with lower investment rates, economic output, and asset prices.

By and large, the literature has taken a "top-down" approach to study the effects of uncertainty, either by constructing measures of aggregate uncertainty or by assuming common uncertainty shocks in equilibrium models. In contrast, this paper takes a "bottom-up" approach to study firm-level investment decisions under micro-level uncertainties that originate from different locations in each firm's supply chain. In particular, questions that we address include: what is the relation between a firm's upstream uncertainty (i.e., the uncertainty originating from suppliers) and its investment rate, does a similar relation hold with a firm's downstream uncertainty (i.e., the uncertainty originating from customers), and moving from the micro to the macro level, do these uncertainties suppress economic growth and financial valuations, and do they rise in recessions?

By combining two comprehensive data sets on dynamic supplier-customer relationships, we document a great deal of asymmetry between the impacts of each firm's upstream and downstream uncertainty. We empirically show that higher upstream uncertainty tends to decrease firm-level investment, inventory, and hiring, whereas higher downstream uncertainty has a positive effect on these variables. The fact that downstream uncertainty can spur firms' economic activity is novel and stands in contrast to the commonly documented negative interaction between real outcomes and other facets of uncertainty.

We also show that this asymmetry between upstream and downstream uncertainty holds at the aggregate level. Using a measure of upstreamness (i.e., distance to final consumers) that is based on the Bureau of Economic Analysis's (BEA's) Input-Output tables, we classify industries as upstream or downstream and compute the macro-level uncertainty of each group. We then show that these macro-level upstream and downstream uncertainties induce polar opposite and significant effects on aggregate variables such as consumption, output, investment, and the market's valuation, in line with the micro-level evidence. Moreover, we show that the marginal utility of investors rises (falls) with more upstream (downstream) uncertainty. This stark difference between the two uncertainties, and the positive effects of downstream uncertainty in particular, may help reconcile why the causality or the magnitude of the relation between uncertainty and the economy is sometimes called into question.

To motivate the empirical analysis, we consider several a priori differences between upstream and downstream uncertainty. One prominent difference relates to the horizon at which each facet of uncertainty impacts firms via relative prices. To see this, consider a typical investment project that comprises three steps: purchasing inputs from suppliers, using these inputs to assemble a product, and then selling this product to customers. In reality, the middle stage — often referred to as the firm's "time-to-build" — can be lengthy due to the time required to process the inputs, perform research and development (R&D), clear regulatory bars, or construct any necessary equipment. Consequently, upstream uncertainty, which is associated with the input prices paid to suppliers, typically affects firms in the short run, while downstream uncertainty, which is associated with the sale price received from customers, typically affects firms in the longer run.[1] We check whether this intrinsic horizon-based asymmetry between the two uncertainties yields different implications for firm investment.[2]

To fix ideas about investment under these supply-chain uncertainties, we construct a production-based framework in which the firm is endowed with a growth option. Exercising this option involves the firm buying inputs from its supplier, building extra capacity, and then selling a stream of output generated by the additional productive capacity to its customers. The project involves time-to-build since any inputs purchased from the supplier take time

[1] Economic shocks to a firm's suppliers (e.g., technology or markup shocks) affect the firm's input price by shifting the supply of inputs (assuming supply and demand curves are not perfectly elastic). Thus, upstream economic uncertainty is related to input price uncertainty. Likewise, economic shocks to a firm's customers affect the firm's selling price, so downstream economic uncertainty is related to output price uncertainty.

[2] Of course, there could be exceptions to this horizon-based asymmetry. In some cases, investments involve pre-paid orders. In other instances, futures contracts that hedge output price variation (e.g., for certain commodity producers) may exist. Nonetheless, to the extent that actual production entails an exchange of goods with suppliers *before* an exchange of goods with customers, this horizon-based asymmetry is largely descriptive of the typical investment project. In our empirical analysis, we find evidence that supports this assumption of horizon asymmetry, through the lens of our model.

to become operational. The project is also partially reversible, as the firm can shed excess capacity through disinvestment (i.e., abandoning the project). Every period, the firm chooses between exercising the option immediately or waiting. In the model, upstream (downstream) uncertainty, broadly defined as the return volatility of suppliers (customers), is endogenously linked to the future input (output) price uncertainty the firm faces.

As is standard in real-option models, there is value in waiting for new information to arrive before committing to the project. By waiting, firms can learn good news about the input price dropping while avoiding bad news about the output price dropping. Since higher upstream and downstream uncertainty are associated with an increased likelihood of extreme input and output prices, both types of uncertainty increase the option value of waiting.

However, there is also an opportunity cost of waiting to invest. By not investing today, the firm forgoes any profits it could have earned during the waiting period. These forgone revenues become stochastic with time-to-build, as the missed revenues depend on the *future* output price that is realized only when the time-to-build period ends. The option to abandon the project limits potential losses and makes this opportunity cost of waiting a convex function of the future output price. Consequently, this cost of waiting increases with downstream (or output price) uncertainty but is unaffected by upstream (or input price) uncertainty.

The fact that the opportunity cost of waiting rises with downstream uncertainty but is unrelated to upstream uncertainty creates an asymmetry between these two types of uncertainty. If the time-to-build period is sufficiently long, then downstream uncertainty can cause the opportunity cost of waiting to dominate the benefit and hasten investment. This leads us to hypothesize that (i) the relation between upstream uncertainty and investment is unambiguously negative, while (ii) the relation between downstream uncertainty and investment is weaker but can be positive.

Motivated by these hypotheses, we examine the relations between corporate investment and these two types of uncertainty by collecting granular data on supplier-customer links from Compustat Segments and Factset Revere. We combine these sources to obtain the most complete dynamic data for each firm's customers and suppliers given data availability. Notably, while input and output prices are not directly observed in the data, the model predicts that a firm's input (output) price uncertainty is positively related to uncertainty over the valuations of its suppliers (customers) in equilibrium. We verify this relation using the NBER-CES manufacturing database and consequently measure each firm's upstream (downstream) uncertainty using its suppliers' (customers') realized stock return volatility. Robustness tests show similar results using a forward-looking uncertainty proxy based on option-implied volatility.

We next estimate panel regressions in which the dependent variable is a firm's investment rate, working capital growth, intangibles growth, hiring growth, or cost of goods sold (COGS) growth, and the independent variables include the firm's upstream or downstream uncertainty, the firm's own uncertainty, and a host of controls. We show that upstream uncertainty has a strong downstream effect and is associated with a significant reduction in all variables of interest, consistent with our first hypothesis. In contrast, higher downstream uncertainty impacts focal firms more weakly in absolute terms and often with the opposite sign: it has either no effect or a *positive* impact on these outcomes, in line with our second hypothesis.

Given the challenge of finding exogenous shocks to these supply-chain uncertainties, we interpret these empirical results through the lens of our motivating model and the broader literature on the established effects of supply-side shocks. Nonetheless, we also provide suggestive causal evidence using an instrumental variables (IV) approach.[3]

We test two additional model-implied hypotheses to tighten the link between the former results and our theoretical framework. Our first test classifies firms as having long or short time-to-build based on one of three proxies: the firm's sector, R&D intensity, or depreciation rate. The first proxy is motivated by the fact that the construction of nondurable consump-

[3] See Section II for a discussion on the paths through which uncertainty shocks operate. We note two additional caveats. First, the validity of this empirical evidence does not depend on the particular assumptions of our economic model, and we do not rule out that the findings may also arise for reasons outside of our framework. Second, even if the relation between downstream uncertainty and investment is noncausal, its sign stands fundamentally in contrast to the common negative interaction between uncertainty and growth.

tion goods and services is generally quicker than the construction of durable consumption or investment goods. The second proxy is motivated by the fact that R&D-intensive projects often have long incubation periods that precede the sale of output (e.g., the time it takes for a biotechnology firm to develop and sell a drug can exceed 10 years). The third proxy reflects the notion that building more persistent capital typically requires more planning and construction time. Using each proxy, we find that the relation between downstream uncertainty and investment is significantly higher (more positive) among firms with longer time-to-build, consistent with the theory. Our second test shows that the positive (negative) effect of downstream (upstream) uncertainty is amplified (attenuated) for firms whose abandonment option is less costly, as proxied by the degree of capital redeployability.

While our motivating framework pertains to firm-level outcomes, we also examine whether the asymmetry between upstream and downstream uncertainty holds at the aggregate level. To this end, we use the BEA's Input-Output tables to construct "upstreamness" scores that capture each industry's weighted distance from the production of final consumption goods (see Antràs and Chor (2018) and Gofman, Segal, and Wu (2020)). We then dynamically classify each industry as upstream (downstream) if its upstreamness score is at the top (bottom) of the cross-sectional distribution of these scores. To parallel our firm-level analyses, we proxy for these macro-level upstream and downstream uncertainties by computing the average realized stock return volatility of each of these two groups of firms.

We then compute impulse responses from each type of macro-level uncertainty to output, consumption, and investment growth, the market's price-to-dividend ratio, and the risk-free rate using smooth local projections (Barnichon and Brownlees, 2019). Controlling for both uncertainties, the impulse responses from macro-level upstream (downstream) uncertainty shocks to these variables are negative (positive), in line with the micro-level evidence. In absolute terms, the negative impact of macro-level upstream uncertainty is stronger than the positive impact of macro-level downstream uncertainty. Lastly, we estimate the market prices of risk of the predictable components of the macro-level uncertainties using a Generelized Method of Moments (GMM) procedure. In line with the impulse responses to consumption, we find that macro-level upstream (downstream) uncertainty has a negative (positive) price of risk. That is, macro-level upstream (downstream) uncertainty increases (decreases) the marginal utility of investors.

Our results have potential implications for policymaking. First, understanding whether increased macro-level uncertainty stems from upstream or downstream firms can help policymakers anticipate whether a contractionary shock will be amplified by uncertainty fluctuations. For instance, we find that the rise in uncertainty at the onset of the COVID-19 crisis was driven largely by downstream firms, potentially accelerating recovery. This stands in contrast to other recessions, where upstream uncertainty dominated from the outset. Second, we show that the downstream uncertainty can positively influence investment in longer-duration projects. Thus, regulations that extend the incubation period of projects could attenuate the adverse effects of uncertainty on the economy.

In sum, our evidence highlights that not all types of uncertainty cause or are associated with recessions: downstream uncertainty may induce the opposite — expansionary — effect.

The remainder of the paper is organized as follows. Section I discusses related literature. Section II presents our theoretical framework. Section III provides micro-level evidence consistent with the framework, and Section IV extends the analysis to the macro level. Finally, Section V concludes.

I. Related Literature

The empirical literature in macroeconomics and finance typically documents a negative relation between uncertainty and economic activity. On the macro side, the negative association between growth and uncertainty is shown in Ramey and Ramey (1995), Martin and Rogers (2000), Engel and Rangel (2008), Bloom (2009), and Baker and Bloom (2013), among others. Nakamura, Sergeyev, and Steinsson (2017) report a negative association between growth and volatility shocks for a panel of countries.[4] Likewise, most asset pricing studies suggest that higher aggregate uncertainty lowers valuations. Bansal, Khatchatrian,

[4] See Bloom (2014) for a comprehensive survey of uncertainty and macroeconomic growth.

and Yaron (2005) show that consumption volatility is negatively related to price-dividend ratios. Boguth and Kuehn (2013) and Bansal et al. (2014) show that consumption uncertainty has a negative price of risk, whereas Drechsler and Yaron (2011) report that higher macroeconomic uncertainty positively predicts risk premia. The negative effect of uncertainty on asset prices also holds for other types of uncertainty, such as fiscal or monetary policy uncertainty (Pastor and Veronesi, 2012; Croce, Nguyen, and Schmid, 2012; Baker, Bloom, and Davis, 2016; Bretscher, Hsu, and Tamoni, 2020; Husted, Rogers, and Sun, 2020) and the common component of idiosyncratic volatility (Herskovic et al., 2016).[5]

A small set of papers find evidence that uncertainty's impact on growth and asset prices can be positive, at least for certain outcome variables or particular types of uncertainty. For example, Stein and Stone (2013) find that uncertainty depresses investment and hiring but encourages R&D. In contrast, we show that the positive effects of downstream uncertainty are not confined to R&D, but also extend to a wide range of outcomes, including capital investment, consumption, intangibles, and stock prices. Other papers find ambivalent links between different facets of uncertainty and asset prices. Segal, Shaliastovich, and Yaron (2015) distinguish between the negative and positive semivariances of industrial production and show that the latter semivariance can increase stock prices.[6] We focus on total variation and show that total variance has different implications depending on where it originates in the supply chain. While Segal (2019) considers cross-sectoral heterogeneity in uncertainty and finds that the total factor productivity (TFP) volatility of final consumption (investment)

[5]Some studies also argue that in contrast to prior papers, the negative relation between uncertainty and growth or asset prices is either weak or noncausal. Dew-Becker et al. (2017) find that only short-run (realized) volatility is priced by investors. Ludvigson, Ma, and Ng (2021) use structural vector autoregressions to show that macroeconomic uncertainty shocks do not lead to a decline in industrial production, but the converse is true – a drop in output predicts higher macroeconomic uncertainty. Our paper can perhaps reconcile the mixed evidence regarding the strength or the causal impact of aggregate uncertainty on the real and the financial economy. We show that aggregate uncertainty has a component (namely, the downstream part) that often exhibits a positive correlation with growth and asset prices. Due to the offsetting effects of upstream and downstream uncertainty, the association between total macro-level uncertainty and economic growth may become less pronounced or ambiguous.

[6]Other papers that depart from Gaussianity or feature an ambivalent relation between uncertainty and asset prices include Okou et al. (2018), Ai and Kiku (2016), Dou (2017), Bianchi, Kung, and Tirskikh (2023), Chang, d'Avernas, and Eisfeldt (2024), and Bekaert, Engstrom, and Xu (2022).

good producers negatively (positively) impacts asset prices, our baseline decomposition of uncertainty into its upstream and downstream components is firm-specific and not based on firms' sectors. In most cases, the firm, its customers, and its suppliers operate *within* the same sector.

In line with extant studies, we demonstrate that not all types of uncertainty are negatively correlated with the real and financial economy. However, our evidence is granular: we start with micro-level analysis and show that the same conclusions hold at the macro level. The findings also demonstrate that the effect of uncertainty shocks can have both upstream and downstream effects, albeit with different implications and magnitudes.[7]

The theoretical literature that treats uncertainty shocks as exogenous typically shows that uncertainty leads to lower investment, consumption, and output. Higher uncertainty can lead to higher markups (e.g., Fernández-Villaverde et al., 2015; Basu and Bundick, 2017), lower utilization (Segal and Shaliastovich, 2025), or higher costs of capital (e.g., Christiano, Motto, and Rostagno, 2010; Gilchrist, Sim, and Zakrajšek, 2014; Arellano, Bai, and Kehoe, 2019), which reduces investment. Real-option studies (see, for example, Dixit (1992), Pindyck (1993), Bloom (2009), and Alfaro, Bloom, and Lin (2024)) show that uncertainty can inhibit investment due to a "bad news principle" (Bernanke, 1983), whereby uncertainty implies that future bad states are more severe. With investment irreversibility, there is a greater benefit of waiting for new information to avoid the potential loss incurred by investing just before a bad state is realized, as the opportunity cost of waiting (namely, the forgone profits during the waiting period) is typically deterministic.[8]

Nonetheless, related studies suggest that real-option models can be enriched to introduce a countervailing force. Under certain assumptions, uncertainty can raise the opportunity cost of waiting and/or expected profits (i.e., the "growth option" value). In these theories,

[7]While our paper considers the effect of second-moment shocks, existing studies consider the effect of first-moment shocks along the supply chain (see, for example, Atalay, Hortaçsu, and Syverson, 2014; Acemoglu, Akcigit, and Kerr, 2016; Baqaee, 2018; Carvalho et al., 2021).

[8]See, for example, Abel and Eberly (1994) for a detailed exposition of investment under uncertainty in a model featuring adjustment costs and irreversibility, and Gulen and Ion (2016) for recent evidence. Related, Grullon, Lyandres, and Zhdanov (2012) show that the risk-return trade-off is driven by real options, while Chaney, Sraer, and Thesmar (2012) consider the effect of financing frictions on investment.

a mean-preserving spread leads to an unbounded upside but a capped downside, suggesting an increase in expected profits by convexity (e.g., Oi, 1961; Hartman, 1972; Abel, 1983; Lyandres and Zhdanov, 2010). One assumption that yields an increase in the growth option value, and the opportunity cost of waiting, is time-to-build (e.g., Majd and Pindyck (1987) and Bar-Ilan and Strange (1996)).

Our contribution relative to these studies is twofold. First, we cast the notion of time-to-build into a *quantitative production* environment and study a model that features two *stochastic* types of uncertainty. We show that the time-to-build mechanism applies to downstream (output price) uncertainty but not to upstream (input price) uncertainty, and we derive further theoretical predictions on when this outcome is more pronounced. Second, our empirical findings provide strong support for the prediction of these "growth option" theories, particularly the time-to-build channel. Empirical evidence in favor of this channel hinges on our decoupling of uncertainty into its upstream and downstream parts.

II. Motivating Theoretical Framework

Ex-ante, several theoretical arguments suggest that upstream uncertainty should have a stronger impact on economic activity than downstream uncertainty at both the micro and the macro levels. We discuss several such arguments based on extant studies before proposing an economically distinct channel through which the effect of the two types of uncertainty is asymmetric in not only its magnitude but also its sign.

At the micro level, several studies consider the impact of *first-moment* shocks in production networks (e.g., Boehm, Flaaen, and Pandalai-Nayar (2019), Barrot and Sauvagnat (2016), and Carvalho et al. (2021)). In particular, Acemoglu, Akcigit, and Kerr (2016) employ minimal assumptions to show how shocks to the level of productivity, governing the supply side of the economy, exert a stronger impact downstream than upstream. These supply shocks to the level of productivity could change the input prices faced by customers and lead to sizable downstream effects. In our context, upstream and downstream uncertainty are both *second-moment* supply shocks. Thus, while Acemoglu, Akcigit, and Kerr (2016) do not explicitly consider uncertainty shocks either theoretically or empirically, the comparative statics from their model imply that firms should be impacted more strongly by upstream uncertainty than downstream uncertainty.

At the macro level, Hulten (1978) shows that the impact of industry-specific productivity shocks on aggregate productivity reflects the industry's sales share of GDP, that is, its Domar weight. As most links in the production network point downstream, the output of upstream industries is directly and indirectly used by more firms (i.e., all downstream firms), leading to a larger Domar weight for upstream firms. First-moment shocks to upstream firms may therefore contribute to a larger fraction of TFP. Consequently, and all else equal, upstream uncertainty should drive a larger share of total uncertainty and induce a more sizable impact on the macroeconomy than downstream uncertainty.

The former explanations help reconcile some of the empirical results that follow, in particular, why downstream uncertainty typically has a *quantitatively* smaller effect on investment than upstream uncertainty. However, these theories cannot explain why downstream uncertainty has the *qualitatively* opposite effect on investment than upstream uncertainty. That is, downstream uncertainty often spurs, rather than suppresses, investment.

We therefore highlight a separate channel that motivates our empirical analysis: there is an economically important difference in the horizon at which downstream uncertainty and upstream uncertainty impact a firm's cash flows. To see why, consider Figure 1 that represents the typical "textbook" investment project.

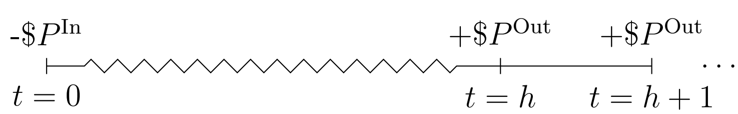

Figure 1. Investment project cash flow diagram. This figure illustrates a stylized cash flow diagram for a representative investment project. At time 0, the firm purchases an input good from its supplier at a price of P^{In}. The firm then uses the next $h \geq 1$ periods to convert this input good into its own output good. Here, h represents the firm's "time-to-build." Finally, the firm sells its output good to customers at a price of P^{Out}.

The firm spends P^{In} to purchase inputs from its *supplier* at time 0, converts these inputs into its own output product over the following $h \geq 1$ periods, and sells this output to its customers for a cash flow of P^{Out} from time h onwards. The second step of this process is often lengthy due to "time-to-build." Thus, uncertainty from suppliers impacts the cost of the input good (P^{In}) and affects firms in the short run, whereas uncertainty from customers associated with the price of the output good (P^{Out}) affects firms in the longer run.

While the former example is highly stylized, in what follows we construct a production-based model that formally illustrates how this conceptual horizon-based asymmetry between upstream and downstream uncertainty leads to different implications for investment.

A. Setup

A.1. Focal Firm

The focal firm has assets-in-place in the form of an installed capacity of inputs k_t that depreciate at rate δ. This installed capacity yields an operating cash flow of π_t per period:

$$\pi_t = P_t^{\text{Out}} k_t^{\alpha} - \omega k_t - P_t^{\text{In}} \delta k_t. \tag{1}$$

The first term in equation (1) captures the revenue from sales. Specifically, each unit of installed input is converted into an output via decreasing returns to scale technology, where $\alpha \in (0, 1)$. Each unit of the firm's output is then sold to a customer for a stochastic output price of P_t^{Out}. The second term in the equation captures operating costs, which are proportional to the stock of inputs via a nonnegative scalar ω. The last term captures the cost of maintaining the existing stock of inputs. Maintenance requires the firm to purchase the depreciated amount of inputs from its supplier, each unit of which is sold to the focal firm for a stochastic input price of P_t^{In}.

At any time t, the focal firm has a growth option to expand its capacity to a level $k' > k_t$, where k' is endogenously chosen by the firm. Exercising the growth option at time t entails paying a fixed cost of $f \cdot k_t$ upfront, thereby making expansion a real option. After exercising the option, the firm must wait for H periods of time-to-build to elapse before the additional capacity becomes productive and starts generating extra output. The standard assumption in most neoclassical models is that H is equal to one period, and our setup nests that case. During the build-up stage, the focal firm (gradually) acquires the $(k' - k_t)$ excess units of inputs from the supplier. Specifically, at each time $\tau \in [t, t+H-1]$, the firm purchases a fraction $w_{\tau-t+1}$ of the total excess capacity from the supplier at the input price prevailing at time τ, with $\sum_{j=1}^{H} w_j = 1$. While building this excess capacity, the focal firm still generates a periodic cash flow of π_τ from its existing assets-in-place. The growth option can also be abandoned in the future, as the firm can choose to shed capacity via disinvestment and receive the revenue from selling part of its operational capacity at the prevailing input price.

Both the input price, paid to suppliers, and the output price, paid by customers, are time-varying. To ensure that no asymmetric effects arise exogenously, we deliberately assume that the laws of motion for these prices are symmetric. The logarithms of both P_{t+1}^{In} and P_{t+1}^{Out} evolve as mean-reverting processes with time-varying conditional volatilities of σ_t^{In} and σ_t^{Out}, respectively. Specifically, for $j \in \{\text{In, Out}\}$, the law of motion is given by

$$p_{t+1}^j = \rho_p p_t^j + \sigma_p e^{\sigma_t^j} \varepsilon_{t+1}^j, \tag{2}$$

$$\sigma_{t+1}^j = \rho_\sigma \sigma_t^j + \sigma_w \eta_{t+1}^j, \tag{3}$$

where $p_t^j = \log(P_t^j)$ and the innovations ε_{t+1}^j and η_{t+1}^j are standard normal and independent and idendically distributed (i.i.d). The parameter ρ_p (ρ_σ) governs the persistence of the price (stochastic volatility) process, while σ_p (σ_w) governs the unconditional volatility of the process innovations.

Let $\Gamma_t = \left[p_t^{\text{In}}, \sigma_t^{\text{In}}, p_t^{\text{Out}}, \sigma_t^{\text{Out}}\right]$ denote the vector of exogenous state variables. The focal firm's problem can be recursively written as

$$V(k,\Gamma) = \max_{k'} \Bigg\{ \pi(k,\Gamma) + \Phi(k,k') +$$

$$\max \Bigg[P^{\text{In}}(k-k') + \beta \mathbb{E}[V(k',\Gamma')], \qquad \text{(if } k' \leq k\text{, Contraction)}$$

$$- f \cdot k - P^{\text{In}} w_1 (k' - k) + \beta \mathbb{E}[V^{\text{Build}}(k, k', \Gamma', H-1)]\Big]\Bigg\}, \quad (\text{if } k' > k, \text{ Expansion}) \tag{4}$$

where

$$V^{\text{Build}}(k, k', \Gamma, \tau) = \begin{cases} V(k', \Gamma), & \text{if } \tau = 0 \\ \pi(k, \Gamma) - P^{\text{In}} w_{H-\tau+1}(k' - k) + \beta \mathbb{E}[V^{\text{Build}}(k, k', \Gamma', \tau - 1)], & \text{if } \tau > 0. \end{cases} \tag{5}$$

Bellman equation (4) suggests that the firm selects its future stock of inputs k' to maximize firm value, $V(k, \Gamma)$, which is the cum-dividend value of a firm with an unexercised growth option. This value includes three parts: (i) the operating cash flows $\pi(k, \Gamma)$ from equation (1), (ii) a twice-differentiable adjustment cost for expanding or contracting capacity, $\Phi(k, k')$, such that $\Phi(k', k) = 0$ if $k' = k$ and $\Phi''(\cdot) > 0$, and (iii) a continuation value that depends nonsmoothly on whether the firm chooses to exercise its option to expand. If the firm opts to contract or keep its capacity unchanged (i.e., if $k' \leq k$), then it receives both the proceeds from selling inputs, $P^{\text{In}}(k - k')$, and the expected discounted continuation value.

If the firm chooses to exercise the growth option and expand its operational capacity (i.e., if $k' > k$), then the firm pays the fixed cost of expansion and enters the time-to-build phase by purchasing a fraction w_1 of the excess capacity $(k' - k)$. The firm's continuation value then depends on V^{Build}, a separate valuation function that reflects the firm's value during this build-up stage (see equation (5)). This function is determined by (i) installed capacity (k) and the exogenous states (Γ), which determine the level of profits from existing assets-in-place, (ii) future capacity (k'), which determines the level of operational capacity when the building is complete, and (iii) the elapsed time-to-build (τ), which determines the fraction of additional inputs to be acquired in the current period. When the remaining time-to-build is zero, V^{Build} equals V and all previously purchased excess capacity becomes operational.[9]

A.2. Supplier and Customer Firms

The supplier firm (s) produces the same good that the focal firm uses as its input. Likewise, the input good of the customer firm (c) is the same good that the focal firm produces and sells as its output. The customer- and supplier-firm problems are symmetric and reflect the standard neoclassical firm programs. Specifically, for $j \in \{s, c\}$,

$$V^j(k^j, \Gamma) = \max_{k^{j\prime}, I^j} \quad \underbrace{P^{j,\text{Out}}\left(k^j\right)^{\alpha} - \omega k^j - P^{j,\text{In}} I^j - \Phi(k^j, k^{j\prime})}_{d^j} + \beta \mathbb{E}[V^j(k^{j\prime}, \Gamma')], \tag{6}$$

$$\text{subject to} \quad k^{j\prime} = (1-\delta)k^j + I^j, \tag{7}$$

where k^j (I^j) is the stock of inputs (investment) of firm j. Each firm's dividend (d^j) depends on revenues net of operating costs, investment expenditures, and capital adjustment costs.

Notably, the prices faced by the focal firm depend on its trade partner's prices: the input price of the focal firm is equal to the output price of the supplier, s, and the output price of the focal firm is equal to the input price of the customer, c. This implies that

$$P_t^{\text{In}} = P_t^{s,\text{Out}} \quad \text{and} \quad P_t^{\text{Out}} = P_t^{c,\text{In}}. \tag{8}$$

Upstream and Downstream Uncertainty. Based on the former definitions, the realized stock returns of the supplier, R_t^s, and the customer, R_t^c, are given by each firm's cum-dividend value divided by the firm's lagged ex-dividend value:

$$R_t^s = \frac{V_t^s}{V_{t-1}^s - d_{t-1}^s} \quad \text{and} \quad R_t^c = \frac{V_t^c}{V_{t-1}^c - d_{t-1}^c}.$$

[9]Equations (4) and (5) jointly embed a law of motion for capital that depends on the firm's chosen action. The investment quantity is inferred by collecting all terms that are pre-multiplied by P^{In}: Assume that the current installed capacity is k_t. If the firm chooses to exercise the growth option and increase its capacity to a higher level k', then $k_{t+H} = k'$, while $k_\tau = k_t$ and $I_\tau = \delta k_\tau + w_{\tau-t+1}(k_{t+H} - k_t)$ for $\tau \in [t, t+H-1]$. Note that if $H = 1$, then this reduces to $k_{t+1} = (1-\delta)k_t + I_t$. If the firm chooses to keep its capacity unchanged or shrink to a lower capacity k', then $k_{t+1} = k'$ and $I_t = \delta k_t + (k_{t+1} - k_t) = k_{t+1} - (1-\delta)k_t$.

Consequently, we define upstream and downstream uncertainty at each time t as the realized stock return volatility of the supplier and the customer, respectively, over a rolling window that covers the previous W periods:[10]

$$\sigma_t^{\text{Upstream}} = \text{Std}(R_{t-W}^s, ..., R_t^s) \quad \text{and} \quad \sigma_t^{\text{Downstream}} = \text{Std}(R_{t-W}^c, ..., R_t^c). \tag{9}$$

In the model, common shocks to the supplier (customer) and the focal firm pass through the input (output) price. Thus, while input and output price uncertainty, σ^{In} and σ^{Out} from equation (3), are not directly observed, we propose and test their observable proxies, $\sigma_t^{\text{Upstream}}$ and $\sigma_t^{\text{Downstream}}$, respectively, through the lens of the model.[11]

B. *Model Solution and Quantification*

We solve the model numerically using a value function iteration method. Gaussian autoregressive processes are discretized using a variant of the Tauchen (1986) method that allows for time-varying conditional volatility, similar to Alfaro, Bloom, and Lin (2024). To ensure the accuracy of the implied policy functions, we use a refined and endogenous grid for capital that is centered around the model's stochastic steady state. The grids for prices are made dense around the endogenous free boundaries where the growth option is exercised. Details on the model's solution and simulations are relegated to Section VI of the Internet Appendix.

The model is calibrated at the quarterly frequency. Table I lists the model parameters.

[10] Based on equations (6) and (8), only one type of price uncertainty is considered when constructing the return volatility of supplier and customer firms. The supplier (customer) firm's return volatility, σ^{Upstream} ($\sigma^{\text{Downstream}}$), depends only on the input (output) price uncertainty of the focal firm, σ^{In} (σ^{Out}).

[11] A largely equivalent modeling assumption would be to explicitly assume that the customer and the supplier firms have economic fundamentals (e.g., technology or markup shocks), denoted by $z_{c,t}$ and $z_{s,t}$, respectively, that fluctuate over time, where an increase in $z_{c,t}$ would increase P_t^{Out} (more demand for the firm's output) whereas an increase in $z_{s,t}$ would decrease P_t^{In} (more supply of the firm's input).

While first-moment shocks to $z_{c,t}$ and $z_{s,t}$ can each affect relative prices with opposite signs, second-moment shocks to these fundamentals always increase the variability of relative prices regardless of the underlying nature of the fluctuations in $z_{c,t}$ or $z_{s,t}$. Consequently, fundamental upstream uncertainty, or $\sigma(z_{s,t})$, is positively related to input price uncertainty, whereas fundamental downstream uncertainty, or $\sigma(z_{c,t})$, is positively related to output price uncertainty. Moreover, in equilibrium, these fundamental uncertainties would also be related to $\sigma_t^{\text{Upstream}}$ and $\sigma_t^{\text{Downstream}}$. Our current modeling approach is effectively isomorphic but does not take a strong stance on what the fundamental shocks $z_{c,t}$ or $z_{s,t}$ are.

Table I

Calibration

The table shows the focal firm's model parameters. The model is calibrated to the quarterly frequency.

Symbol	Value	Description
α	0.400	Return-to-scale
β	0.997	Discount factor
δ	0.025	Depreciation rate
f	0.020	Fixed cost of investment
b	0.010	Quadratic adjustment cost
ρ_p	0.950	Persistence of input and output prices
σ_p	0.200	Volatility of input and output prices
ρ_v	0.950	Autocorrelation of price volatility
σ_w	0.100	Volatility of price volatility shocks

We set α, the physical input (capital) share of output to 0.4, consistent with empirical evidence. The quarterly depreciation rate is 2.5%, suggesting an annual depreciation of 10%, in line with capital's annual depreciation rate from the BEA. The time discount rate β is set such that the annualized real risk-free rate is approximately 1%, close to its empirical counterpart. We consider two values for H. In the baseline, we set H to one period, consistent with the standard one-period time-to-build in neoclassical models. We also obtain model-implied results when the parameter H is set to eight quarters, consistent with average estimates reported by Bar-Ilan and Strange (1996).

The benchmark adjustment cost function reflects the standard quadratic specification $\Phi(k, k') = \frac{b}{2}(\frac{k'}{k} - 1)^2 k$. We assume that all investment costs for capital expansions are made upfront, suggesting that $w_1 = 1$. However, we relax this assumption in Section II.C.3. The six remaining parameters $\{\rho_p, \sigma_p, \rho_\sigma, \sigma_w, b, f\}$ govern the driving forces of prices and capital frictions.[12] We select these six parameters to target eight annual moments related to firm-level sales, profits, and investment. These annual moments are computed by simulating a population path of the model and then aggregating quarterly observations to the annual

[12]In the benchmark analysis, we shut down the operating leverage ($\omega = 0$) in order to conserve on the number of free parameters. We examine the case of $\omega > 0$ in Section VII of the Internet Appendix.

frequency. The model-implied moments and their empirical counterpart are listed in Table II.

Table II

Model-Implied Moments Versus Data

The table reports model-implied moments against their empirical counterparts under the baseline calibration. The moments pertain to firm-level investment rates (I/K), sales growth (ΔY), and return-on-assets (Profit/K). In this table, σ represents the standard deviation of a given variable, whereas ρ_i denotes the i^{th}-order autocorrelation. The model-implied moments are based on a simulation of a population path from the model that allows the focal firm's capital, price, and associated uncertainty to stochastically vary over time. Quarterly observations are aggregated to form annual observations. Brackets correspond to the 95% confidence interval for each moment.

Statistic	Model	Data	
$\sigma(I/K)$	0.165	0.165	[0.049, 0.370]
Skew(I/K)	0.626	0.701	[-0.551, 2.186]
$\rho_1(I/K)$	0.430	0.341	[-0.213, 0.764]
$\rho_2(I/K)$	0.196	0.065	[-0.426, 0.603]
$\sigma(\Delta Y)$	0.362	0.210	[0.073, 0.602]
$\rho_1(\Delta Y)$	0.162	0.176	[-0.341, 0.710]
$\rho_2(\Delta Y)$	-0.080	-0.012	[-0.494, 0.563]
$\sigma(\text{Profit}/K)$	0.160	0.144	[0.045, 0.389]

The persistence parameters ρ_p and ρ_σ jointly target the autocorrelations of firm-level investment and sales growth rates. The model-implied first-order autocorrelations of these two variables are 0.43 and 0.16, respectively, close to their values of 0.34 and 0.18 in the data. While not directly targeted, the model also generates a mildly negative second-order autocorrelation of sales growth, in line with the data. The volatility parameters σ_p and σ_w jointly target the volatility of sales growth and profit margins, measured by return on assets (ROA). These moments take on model-implied values of 0.362 and 0.160, respectively. In model simulations, we set the smooth adjustment cost parameter b to target the volatility of investment rates to the data. Under the baseline, we obtain an almost perfect fit, with an annual volatility of 0.165. The fixed cost of expansion f governs the lumpiness of investment, and hence its skewness, which amounts to 0.626 (0.701) in the model (data). Importantly, all model-implied moments fall within the empirical confidence intervals provided in brackets.

C. *Model Results*

C.1. Policies

Figure 2 shows the focal firm's steady-state investment policy around the growth option's free boundaries as a function of the log input price (in the left column) and log output price (in the right column). Panel A reports these policies when the time-to-build process is short and lasts one period, which corresponds to the standard neoclassical assumption, while Panel B reports the policies for a longer time-to-build process that lasts eight periods. In all plots, capital is held at the stochastic steady-state level, as are the remaining exogenous state variables that do not vary along the horizontal axis. Each plot displays the relevant policy for three values of input and output price uncertainty. We focus on the "medium" value of each uncertainty, which corresponds to the midpoints of the σ^{In} and σ^{Out} grids, as well as the "high" and "low" values of each uncertainty, which correspond to the cases in which each type of uncertainty is one point above or below the midpoint, respectively.

Because investment is a real option, each policy is characterized by a flat inaction region in which the investment rate is equal to the depreciation rate (i.e., $k' = k$).[13] This is followed by a region in which the firm expands capacity (i.e., sets $k' > k$), provided that the output price is sufficiently high (right column) or the input price is sufficiently low (left column). When the time-to-build period is short (Panel A), the inaction region increases with both higher input and output price uncertainty. In broad terms, this is consistent with the traditional "wait and see" effect of uncertainty.

However, when the time-to-build period is long (Panel B), then the inaction region *shrinks* as output price uncertainty increases from low to medium, and further decreases as output price uncertainty rises from medium to high (i.e., the bottom right figure). This pattern stands in sharp contrast to the impact of output price uncertainty on inaction when juxtaposed with Panel A. We explore the implications of this finding, along with its mechanism,

[13] Put differently, inaction in Figure 2 is equivalent to the investment rate being equal to δ, or 2.5%, whereas in an action region investment rate is higher. To provide a complementary illustration of how the inaction region changes with uncertainty as a step function, refer to Section VII.F in the Internet Appendix.

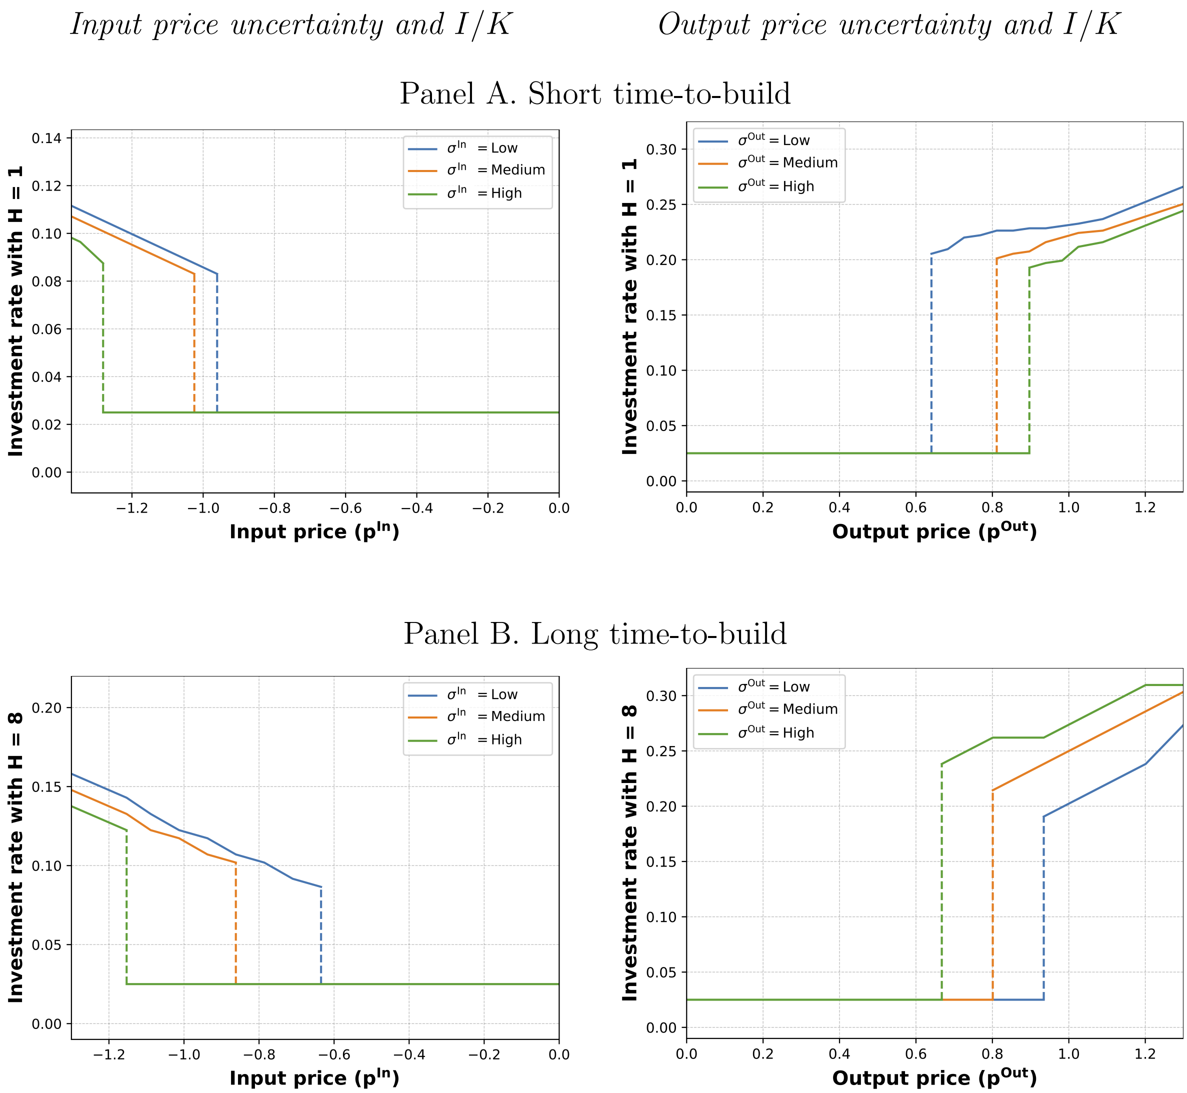

Figure 2. Investment policies. This figure shows the focal firm's investment policy around the free boundary of expansion at the model's steady state. These policies are displayed as a function of the firm's log input price (left column) and log output price (right column). In all figures, capital is held at its stochastic steady-state value. Moreover, the plots in the left (right) column hold P^{Out} and σ^{Out} (P^{In} and σ^{In}) fixed at their unconditional means. Each plot displays the policy for three levels of uncertainty. The orange line in the left (right) column plots the investment policy when σ^{In} (σ^{Out}) is fixed at its unconditional mean (i.e., "Medium") value, which corresponds to the midpoint of its grid. The blue and green lines then show each policy when the relevant uncertainty is one grid point below or above the midpoint, which corresponds to states of "Low" and "High" uncertainty, respectively. Finally, in the top row (i.e., in Panel A), the time-to-build period is set to $H = 1$ period, while the bottom row (i.e., Panel B), sets the time-to-build period to $H = 8$ quarters.

in the next subsections.

C.2. Upstream and Downstream Uncertainty and Investment

The policies in Figure 2 show that at the model's steady state, and for $j \in \{\text{In}, \text{Out}\}$, there exists an investment threshold $p^*(\sigma^j)$ such that if $|p^j| > |p^*(\sigma^j)|$, then in optimum $k' > k$. Put differently, depending on the type of uncertainty, there is a sufficiently large log input or output price (in absolute magnitude) that the firm is in inclined to exercise its growth option.

In practice, firm-specific input and output prices may not be observed or available at a sufficiently high frequency to recover the underlying input and output price uncertainties (i.e., σ_t^{In} and σ_t^{Out} in equation (3)). In contrast, the realized stock returns of a firm's customers and suppliers are more readily available. Consequently, we derive testable implications from the model by examining how the investment thresholds vary with upstream and downstream uncertainty, the observable measures defined in equation (9), which correlate with the fundamental input and output price uncertainty in equilibrium.

Admittedly, as upstream and downstream uncertainty are backward-looking proxies and based on return data, there is no one-to-one correspondence between these return-based measures and the underlying price-based uncertainties that can be expressed in closed form. We therefore proceed with the following simulations around the model's stochastic steady state. When analyzing the effects of upstream uncertainty on investment, we simulate finite-sample paths of the model, allowing the focal firm's log input price (p^{In}) and price volatility (σ^{In}), and the supplier's capital (k^s) and valuations (V^s), to vary over time. We save these finite-sample paths in the matrix $\mathcal{P} = [p_t^*, \sigma_t^{\text{Upstream}}]$,[14] which contains the investment thresholds associated with each realized value of upstream uncertainty. For each unique investment threshold p^* in the matrix $\mathcal{P}$, we compile a set $\text{RV}(p^*) = \{\sigma^{\text{Upstream}} | (p^*, \sigma^{\text{Upstream}}) \in \mathcal{P}\}$ that contains all observed values of upstream uncertainty that are associated with a given investment threshold.

[14] For brevity, we omit uncertainty from the threshold price and write p_t^* rather than $p^*(\sigma_t^{\text{In}})$.

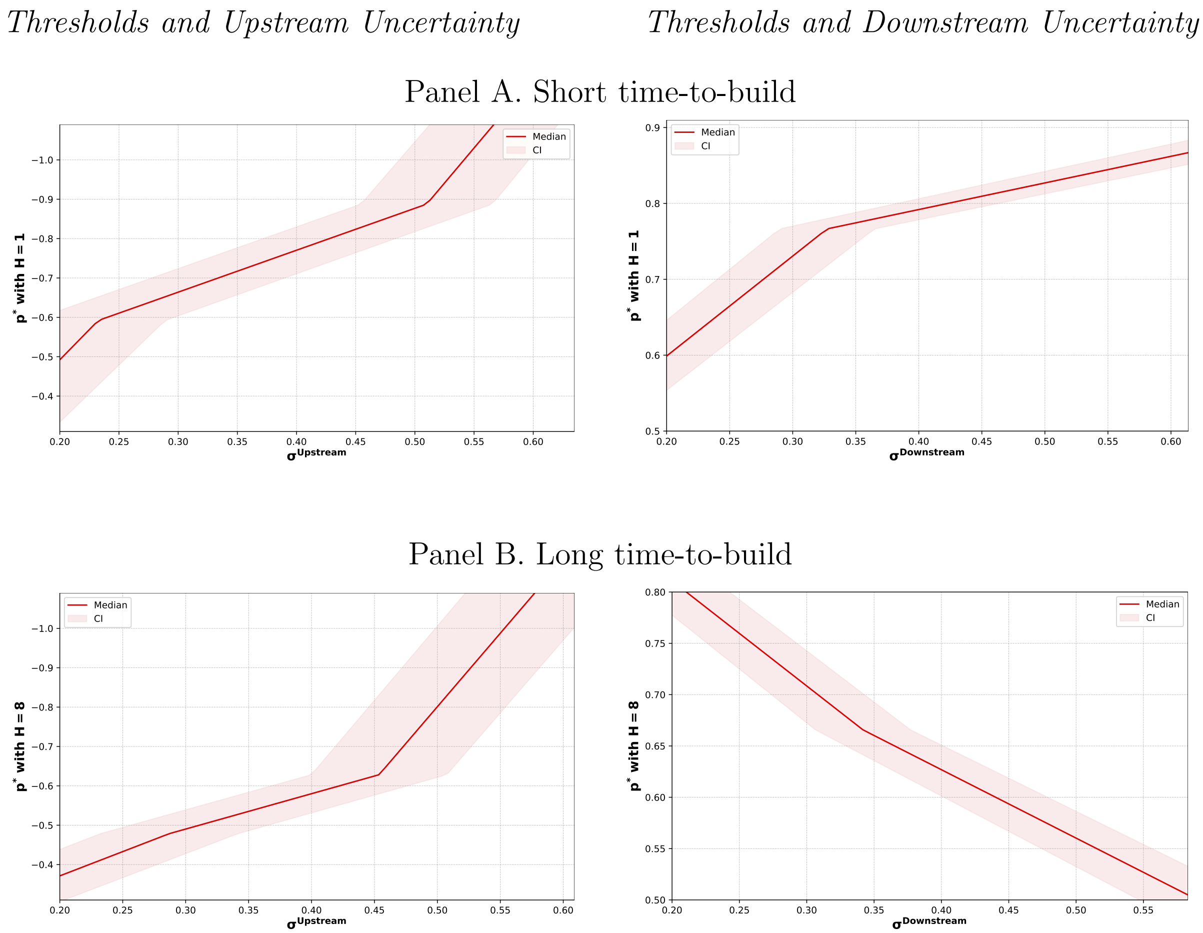

Figure 3. Investment thresholds as a function of return uncertainty. This figure shows the investment thresholds, p^*, in order of increasing (absolute) magnitude as a function of upstream uncertainty σ^{Upstream} (left column) and downstream uncertainty $\sigma^{\text{Downstream}}$ (right column) at the model's stochastic steady state. Accordingly, in the left (right) column, the thresholds p^* reflect free boundaries based on critical values of the log input (output) price, p^{In} (p^{Out}). The horizontal shaded region shows the model-implied 90% confidence interval of the effects of upstream or downstream uncertainty on the investment threshold, displayed on the vertical axis. The red line shows the median value of each confidence interval (i.e., median upstream or downstream uncertainty). Upstream and downstream uncertainty are defined in equation (9) and reflect either the supplier's or the customer's realized stock return volatility over the previous year. Panel A (B) shows the results when time-to-build period H is set to one (eight) quarters.

The left column in Figure 3 plots the results. Each value on the vertical axis corresponds to an investment threshold p^*, displayed in order of increasing absolute magnitude, and the red-shaded region shows the confidence interval of the upstream uncertainties observed at each investment threshold (i.e., $\text{RV}(p^*)$). The red line represents the median value of $\text{RV}(p^*)$

across all possible p^*. We repeat the former analysis for downstream uncertainty by replacing p^{In}, σ^{In}, k^s, V^s, and σ^{Upstream} with p^{Out}, σ^{Out}, k^c, V^c, and $\sigma^{\text{Downstream}}$, respectively. The left column in Figure 3 shows the corresponding results.

If the time-to-build period is short, as in Panel A, both red curves are upward-sloping. This suggests that even after accounting for any noise embedded in the return-based measures of uncertainty from equation (9), the investment thresholds are increasing (in absolute value) with upstream and downstream uncertainty. Thus, with short time-to-build, both types of uncertainty delay the focal firm's investment decision.

Switching to a long time-to-build period in Panel B shows that the investment thresholds still rise with upstream uncertainty (bottom left figure), suggesting that higher supplier-level return volatility is still associated with suppressed investment as in the case of short time-to-build. The relation between the investment thresholds and customer-level return volatility, however, turns downward-sloping with a longer time-to-build (bottom right figure). This downward-sloping curve implies that higher downstream uncertainty could actually spur or hasten investment, in stark contrast to the effects of higher upstream uncertainty, which consistently works to delay investment.[15]

Model Intuition. When downstream uncertainty rises, which in equilibrium is associated with higher values of σ^{Out}, the benefit of waiting to invest increases. This is the "bad news principle" (Bernanke, 1983). Waiting to invest allows the firm to avoid the net loss of investing in an ex-post uneconomical project if the output price falls. As such, higher downstream uncertainty increases the potential loss from investing in the project immediately and increases the benefit of waiting to invest.

Similarly, greater upstream uncertainty — as reflected by higher equilibrium values of σ^{In} — also creates a benefit of waiting for new information to arrive before investing. Waiting allows the firm to learn that the input price has fallen, which allows the firm to exercise its

[15]The impacts of upstream and downstream uncertainty on the investment thresholds differ not only in direction but also in magnitude. The slope of each curve represents how changes in each type of uncertainty affect the respective thresholds. In Panel B of Figure 3, the slope of the curve associated with upstream (downstream) uncertainty is 1.61 (-0.89). Consequently, the positive effect of downstream uncertainty is weaker in absolute value than the negative effect of upstream uncertainty.

growth option at a lower cost. This benefit of waiting increases in σ^{In} as the potential gain from a lower input price becomes more extreme with higher upstream uncertainty.

A key point of asymmetry between the effects of the two uncertainties relates to the opportunity cost of waiting. This cost captures the forgone revenues during the inaction period. In the extreme case with no time-to-build (e.g., McDonald and Siegel (1986)), the firm's excess capacity becomes operational immediately after investment. In this "standard" case, the forgone revenues from waiting are not stochastic. Rather, they are deterministic in both input and output prices. This makes the opportunity cost of waiting independent of both σ^{In} and σ^{Out}, the drivers of upstream and downstream uncertainty, respectively.

With positive time-to-build (i.e., when $H > 0$), the forgone revenues from waiting depend on the output price that is realized *after* the time-to-build period ends. As such, these forgone profits become a stochastic function of the output price but do not depend on the input price when all investment is made upfront, as in the benchmark. Moreover, these forgone profits are capped below by the firm's ability to disinvest capacity (i.e., the abandonment option). This makes the forgone profits at time t — the cost of waiting to invest — a convex function of $P^{\text{Out}}_{\tau \geq t+H}$ that increases with σ^{Out}. We refer to this as the "good news principle" associated with output price uncertainty, and hence with observed downstream uncertainty. In sum, whereas upstream uncertainty only raises the benefit of waiting, downstream uncertainty increases both the benefit *and* the cost of waiting to invest.

As Panels B of Figures 2 and 3 show, the net effect of downstream uncertainty on investment depends critically on the length of the time-to-build period. Intuitively, if the time-to-build period is short, then the opportunity cost of waiting to invest is also small, as the profits forgone during the build-up stage cannot become too extreme. This allows the benefit of waiting to invest, that is, of delaying investment to avoid bad news about the future output price dropping, to dominate the firm's investment decision. In contrast, if the time-to-build period is sufficiently long but remains empirically justified, then the opportunity cost of waiting to invest, that is, the forgone future profits, can become sizable and dominate the benefit of waiting, leading the firm to invest sooner than it otherwise would have.

To conclude, both types of uncertainty suppress investment when time-to-build is short. In contrast, longer time-to-build reduces the net benefit of waiting to resolve *output* price uncertainty, linked to the volatility of customer firms. Therefore, there is an intrinsic asymmetry between the effects of the two types of uncertainty on investment. We emphasize that the forces above, and their effects on the decision to invest, do not depend on the specific calibration of the model. To show that these results hold without loss of generality, in Section I of the Internet Appendix we express these forces in closed form under a simplified (and qualitative) version of the model.

C.3. Model Extensions

Sequential Investment. The benchmark model assumes that all investment in excess capacity is made upfront. We explore an alternative case in which the investment is equally spaced across the time-to-build period, such that $w_j = 1/H$ of the investment is made during each period $j \in [1, H]$ of the build-up stage. Figure 4 shows the results.

With sequential investment, the input price is not fully known at the time the growth option is exercised. Consequently, the forgone profits from waiting to invest now depend qualitatively not only on downstream uncertainty, as in the baseline, but also on the prevailing input prices over the future time-to-build period and consequently on upstream or input price uncertainty. Quantitatively, however, we find that the results still mimic the benchmark case — downstream uncertainty spurs investment by lowering the investment threshold while upstream uncertainty delays investment — although the latter effect is slightly attenuated (i.e., the curve's slope is mildly lower relative to the benchmark).

Intuitively, the "good news principle" of uncertainty is larger when the horizon over which it impacts the firm is longer. This is due to the effects of volatility compounding over time and the convexity induced by the ability to disinvest. Upstream uncertainty affects the short run only, with its effects capped at the horizon of the time-to-build period (i.e., eight quarters). In contrast, downstream uncertainty affects the longer run, from the conclusion of the build-up stage onward. This leads to an intrinsically larger impact on the opportunity cost of waiting and consequently maintains the asymmetry between the two types of uncertainty.

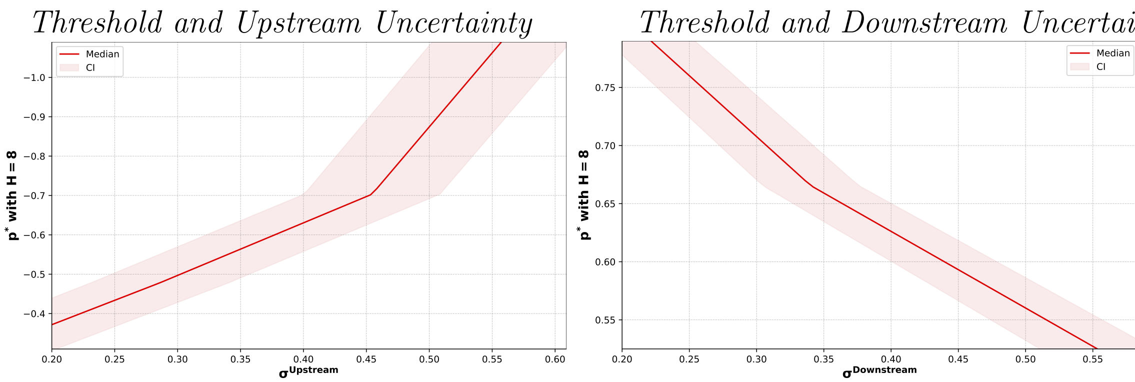

Figure 4. Investment thresholds as a function of return uncertainty: Sequential investment. This figure shows the investment thresholds p^* in order of increasing (absolute) value as a function of upstream uncertainty σ^{Upstream} (left) and downstream uncertainty $\sigma^{\text{Downstream}}$ (right), using an identical procedure to that outlined in Figure 3, but for the case in which the investment is made sequentially over the time-to-build period so that $w_\tau = 1/8$ for $\tau \in [1, 8]$. In the left (right) column, the thresholds p^* reflect free boundaries based on critical values of the log input (output) price, p^{In} (p^{Out}).

Disinvestment Real Option. The benchmark model assumes only smooth adjustment costs for shrinking the capital stock. We explore the implications of an alternative calibration of the model in which investment is made less reversible by introducing a fixed cost for disinvestment. This makes the decision to shrink capacity a real option. Specifically, we assume that if the focal firm chooses an investment policy such that $k' < k$, then it must pay a cost $f^- \cdot k$, where $f^- > 0$ is the fixed cost of disinvestment.

The results are shown in Figure 5. Relative to the benchmark case (displayed in red), all investment thresholds in the alternative calibration rise in absolute value (displayed in blue). This decreases the firm's overall propensity to invest.[16] Put differently, when investment is

[16]Moreover, while not immediately visible in the plots, the disinvestment option increases the slope of the curve for σ^{Upstream} by 25%(=2.016/1.611-1) relative to the benchmark. This higher slope means that as input price uncertainty rises, the *marginal* propensity to invest is lower (i.e., upstream uncertainty delays investment more). Likewise, the disinvestment option reduces the slope of the curve associated with $\sigma^{\text{Downstream}}$ (in absolute value) by 36%(=0.890/0.655-1) compared to the benchmark. The shallower slope means that as output price uncertainty rises, the investment threshold still falls, but not by as much as in the benchmark. This attenuates the hastening effect of downstream uncertainty on investment.

less reversible, the impact of upstream or input price uncertainty becomes more negative, while the impact of downstream or output price uncertainty becomes less positive.

As f^- increases, the growth opportunity is less reversible, and higher upstream uncertainty magnifies the "bad news principle," as waiting to invest becomes the only way for the firm to avoid (opportunity) losses if the investment is ex-post uneconomical. These losses become larger with more upstream uncertainty. Conversely, when investment is more reversible (as in the benchmark case), the "good news principle" associated with downstream uncertainty is magnified. By not investing immediately, the firm forgoes future cash flows, but can also more easily truncate potential losses from the project should bad news arrive. As such, the cost of waiting to invest following a downstream uncertainty shock is larger when the abandonment cost is lower.

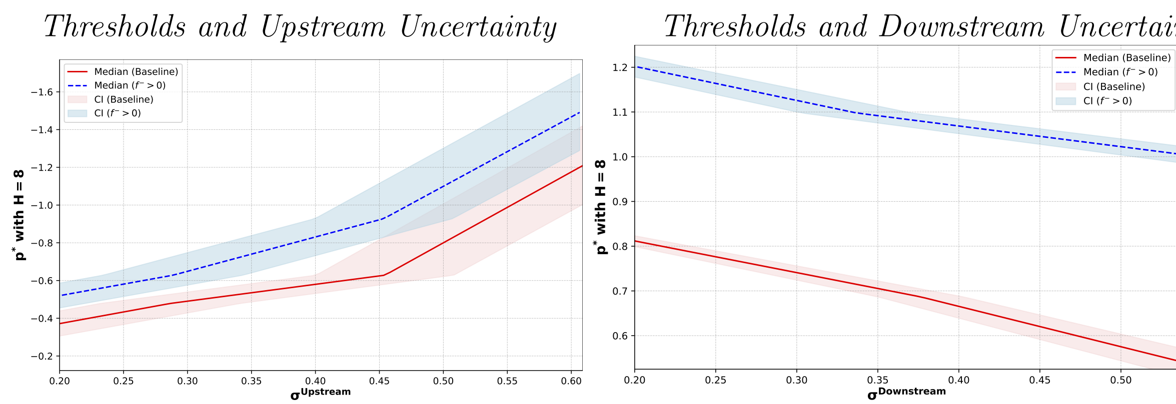

Figure 5. Investment thresholds as a function of return uncertainty: Disinvestment option. This figure shows the investment thresholds p^* in order of increasing (absolute) value as a function of upstream uncertainty σ^{Upstream} (left) and downstream uncertainty $\sigma^{\text{Downstream}}$ (right), using an identical procedure as that outlined in Figure 3, but for the case in which disinvestment is also a real option. Specifically, we assume that the firm pays a cost $f^- \cdot k$ if and only if the firm chooses a policy such that $k' < k$. We set f^- to 0.025. The solid red line in each plot represents the median investment thresholds in the baseline model with long time-to-build, while the dashed blue line represents the median thresholds when $f^- > 0$. Similarly, the red and blue shaded regions represent the 90% confidence interval associated with each curve. In the left (right) column, the thresholds p^* reflect free boundaries based on critical values of the log input (output) price, p^{In} (p^{Out}). Here, we set the time-to-build period H to correspond to eight quarters.

Other Extensions. Section VII in the Internet Appendix considers several other model extensions. First, we show that if investment is concentrated at the end of the time-to-build period, there is a qualitative attenuation effect on the negative impact of upstream uncertainty on investment, but this attenuation is quantitatively negligible. Second, we consider the case in which prices and their associated uncertainties are negatively correlated. Under an empirically realistic correlation structure, this negative correlation decreases the effects of the supply-chain uncertainties on investment, but the results maintain the asymmetry whereby downstream (upstream) uncertainty has a positive (negative) effect on investment. Third, we recalibrate the model to allow upstream uncertainty to exhibit moderately higher volatility than downstream uncertainty (as later confirmed in the data) and report a qualitatively similar effect. Lastly, we recalibrate the model to match the aggregate volatility of dividends and show that the effects of the supply-chain uncertainties are largely invariant to the magnitude of the unconditional levels of volatility.

C.4. Hypotheses

Given the theoretical results, we posit the following testable hypotheses.

HYPOTHESIS 1: The association between upstream (supplier-level) uncertainty and investment is unambiguously negative.

HYPOTHESIS 2: The association between downstream (customer-level) uncertainty and investment is weaker in absolute value, but can be positive.

HYPOTHESIS 3: The association between downstream (customer-level) uncertainty and investment is more positively pronounced for firms with longer time-to-build.

HYPOTHESIS 4: A negative (positive) interaction between upstream (downstream) uncertainty and investment is magnified among firms with harder- (easier-) to abandon projects.

Below, we confirm these conjectures using micro- and macro-level data.

III. Micro-Level Evidence

We use granular and dynamic data on supplier-customer relationships to provide novel empirical evidence that is consistent with the motivating model from Section II. We find that (i) the association between upstream (i.e., supplier-level) uncertainty and investment, hiring, and working capital is negative and statistically significant, whereas (ii) the association between downstream (i.e., customer-level) uncertainty and economic activity is either statistically insignificant or positive and significant.

We describe the construction of the firm-level upstream and downstream uncertainty measures in Section III.A. Using these measures, Sections III.B and III.C document the asymmetry between the two facets of uncertainty for both firm-level investment rates and other firm-level outcomes. We also show that, in line with the model, the positive effects of downstream uncertainty are concentrated among firms that have longer time-to-build periods (Section III.D) and greater investment reversibility (Section III.E). We further establish the asymmetric impact of the two uncertainties on firm-level investment using a battery of robustness checks and an IV approach in Section III.F.

A. *Data*

Our sample includes all firms in the Center for Research in Security Prices (CRSP)/Compustat universe listed on the NYSE, AMEX, and NASDAQ exchanges except for financial firms (SIC 6000 to 6999) and public utilities (SIC 4900 to 4999). We exclude these firms as their supply-chain environments differ greatly from those underlying our model. The firm-level analyses range from 1976 to 2019 due to the availability of granular data on interfirm relationships, described below.

Identifying Customers and Suppliers. We construct the upstream and downstream uncertainties by dynamically identifying the sets of firms that supply to (i.e., are upstream of) and buy from (i.e., are downstream of) each firm in our sample. To this end, we employ two data sets on supplier-customer relationships: Compustat Segments and the FactSet Revere Relationship database. By combining these two data sets we overcome the limitations associated with each specific data set and produce a comprehensive network of interfirm links that spans the longest time period possible.[17]

[17]On the one hand, Compustat Segments data contain considerably fewer links between firms than FactSet. This is because firms are only required to disclose relationships with customers that account for at least 10% of total sales at the annual frequency. However, the Compustat data are available for an extended time

We combine both data sets to construct a panel of inter-firm links between 1976 and 2019 as follows. First, for firm-year observations between June 2003 and June 2019, we start by looking for each firm's suppliers and customers in the FactSet data. By using FactSet data in the first step, we obtain the most comprehensive coverage of each firm's supplier-customer relationships for the most recent part of our sample period. Next, for the years prior to 2003 (when FactSet is unavailable), we obtain each firm's suppliers and customers from the Compustat Segments database. Importantly, using FactSet data for the period in which the two data sources overlap ensures that we capture the union of both data sets as the links reported in the Compustat Segments data are a *subset* of those reported in the FactSet data.

Measuring Uncertainty. While input and output prices, and their associated uncertainties, are not observed for each firm, the model offers an observable proxy by predicting that the return volatility of suppliers (customers) inherits the same implications for investment as input (output) price uncertainty, since return volatility and price uncertainty are linked in equilibrium (recall Section II.A.2 in the model).

As a preliminary step, we empirically verify the link between price uncertainty and return volatility using the NBER-CES database in Section IV of the Internet Appendix. There, we show that the average correlation between industries' input price uncertainty and the return volatility of suppliers is positive and significant (about 0.3). We also obtain a similar result for the correlation between output price uncertainty and the return volatility of customers.

We construct our baseline measures of uncertainty in the spirit of Leahy and Whited (1996) as follows. First, for each firm-year observation between June 1976 and June 2019, we identify the customers and suppliers associated with firm i in year t as described above. Second, for each supplier (customer) firm linked to firm i at time t, we use CRSP daily data to compute the volatility of the supplier's (customer's) daily stock returns in the year preceding time t — the same window used by the model.[18] We then compute the average stock return

period, starting in 1976. On the other hand, FactSet contains around 10 times as many links as Segments, as interfirm relationships in FactSet are obtained using comprehensive data from accounting statements, press releases, interviews, and firms' websites, among other sources. However, the FactSet data are available only from 2003 to 2019.

[18] We calculate these stock return volatilities by (i) adjusting returns for delisting events, and (ii) requiring

volatility across the suppliers (customers) that trade with firm i at time t and refer to this average as the upstream (downstream) uncertainty of firm i. Finally, we control for firm i's own inherent uncertainty, which we also measure using the realized stock return volatility of the firm over the past year, to account for aspects of uncertainty that are exogenous to the model and potentially unrelated to the firm's trading partners.

Notably, while the above measures of uncertainty are backward-looking by construction, the fact that we define uncertainty using variation in *realized* rather than *expected* stock returns does not drive our results (as is also the case in the model). First, since the time series of uncertainty measures are typically highly persistent, ex-post volatility is highly correlated with ex-ante uncertainty but does not rely on parametric assumptions. Second, and for the purpose of robustness, we show that our baseline results hold when we use a forward-looking measure of uncertainty that is extracted from option prices for a smaller cross-section and shorter time series for which option data are available. We report and discuss these results in Section III.F. We provide details on the construction of each variable in Section III.A of the Internet Appendix. In Section VIII.A of the Internet Appendix we also report summary statistics associated with the uncertainty measures and other outcome variables used in the analysis that follows.

B. *Firm-Level Investment Under Supply-Chain Uncertainty*

We establish new facts that are consistent with the main predictions of the model: upstream uncertainty and downstream uncertainty have (i) *independent* interactions with firms' decisions to invest — beyond a firm's own uncertainty — and (ii) *asymmetric* effects on firms' investments. Upstream uncertainty suppresses investment, whereas downstream uncertainty never depresses and potentially spurs investment. We document these effects of the supply-chain uncertainties via the panel regression

$$y_{i,t} = \alpha_i + \delta_t + \beta_1 \sigma(\text{Own})_{i,t} + \beta_2 \sigma(\text{SupplyChain})_{i,t} + \boldsymbol{\gamma}' \boldsymbol{Z}_{i,t} + \varepsilon_{i,t},$$

that each firm have at least 200 nonmissing daily stock returns over the previous year.

$$\text{where } \sigma(\text{SupplyChain}) \in \{\sigma(\text{Upstream}), \sigma(\text{Downstream})\}. \tag{10}$$

Here, $y_{i,t}$ is the investment rate of firm i at time t, defined in line with Belo, Lin, and Bazdresch (2014), α_i denotes firm fixed effects, and δ_t denotes time fixed effects that subsume common shocks to all firms in a given time period (e.g., the Great Recession). The key variable of interest in equation (10) is $\sigma(\text{SupplyChain})_{i,t}$, which represents either $\sigma(\text{Upstream})_{i,t}$, the uncertainty associated with firm i's suppliers, or $\sigma(\text{Downstream})_{i,t}$, the uncertainty associated with firm i's customers, at time t. The term $\sigma(\text{Own})_{i,t}$ denotes firm i's own uncertainty at time t and is an economically important control. The baseline model in Section II assumes that $\sigma^{\text{In}}_{i,t}$ — the conditional volatility of input prices that governs upstream uncertainty — and $\sigma^{\text{Out}}_{i,t}$ — the conditional volatility of output prices that governs downstream uncertainty — are orthogonal. In reality, however, these two uncertainties may become correlated, in part through fundamental shocks that are specific to firm i, which are excluded from the model for computational tractability and parsimony. These shocks to firm i can simultaneously affect both upstream and downstream firms and jointly impact the firm's supplier- and customer-level uncertainties. Consequently, including $\sigma(\text{Own})_{i,t}$ in equation (10) accounts for common shocks to $\sigma(\text{Downstream})_{i,t}$ and $\sigma(\text{Upstream})_{i,t}$ that originate from their common trading partner — firm i.[19,20]

Finally, $\boldsymbol{Z}_{i,t}$ is a vector of time-varying and firm-level controls that capture relevant aspects of firm i's economic environment other than uncertainty that can influence its decision to invest. This set of controls is motivated by those in Leary and Roberts (2014) and includes firm size, leverage, asset tangibility, profitability, Tobin's q, and past returns. Section III.A in the Internet Appendix provides details on the construction of these variables, while summary statistics for these variables are found in Section VIII.A of the Internet Appendix.

Table III reports results of estimating equation (10). In Table III and all subsequent firm-

[19]Indeed, the correlation between $\sigma(\text{Downstream})_{i,t}$ and $\sigma(\text{Upstream})_{i,t}$ is about 0.60. However, the correlation between $\sigma(\text{Downstream})_{i,t}$ and the component of $\sigma(\text{Upstream})_{i,t}$ that is orthogonal to $\sigma(\text{Own})_{i,t}$ is only 0.20 and indistinguishable from zero.

[20]Section VIII.A in the Internet Appendix shows that the negative (positive) association between upstream (downstream) and investment is unchanged if we exclude $\sigma(\text{Own})_{i,t}$ from projection equation (10).

level results, each independent variable is scaled by its unconditional standard deviation to facilitate interpretation of each marginal effect. In all projections, we report robust standard errors that are clustered at the firm level. Moreover, observations are measured in June (i.e., time t refers to June of the given year) to ensure that accounting data are publicly available.

Columns (1) and (2) of the table show that, before accounting for supply-chain uncertainties, higher firm-level uncertainty is associated with a statistically significant decrease in investment. Specifically, a one-standard-deviation increase in a firm's own uncertainty is associated with the firm decreasing its investment rate by 0.12 (0.05) when we (do not) include $\boldsymbol{Z}_{i,t}$ in equation (10). This result is consistent with the wide-held notion that uncertainty is negatively related to investment (e.g., Bloom (2009)).

Columns (3) and (4) extend the previous analysis and consider the incremental effect of upstream uncertainty on firms' investment rates. Upstream uncertainty is negatively related to investment, beyond the negative impact of the firm's own uncertainty. For example, column (4) shows that a one-standard-deviation increase in upstream uncertainty decreases investment rates by an additional 3% when we also condition on the controls $\boldsymbol{Z}_{i,t}$. This effect is significant at the 1% level and consistent with Hypothesis 1 of the model in Section II.

In contrast, and more surprisingly, columns (5) and (6) show that higher downstream uncertainty has a positive relation to firm-level investment. A one-standard-deviation increase in downstream uncertainty *increases* the investment rate by 3%, controlling for the firm's own uncertainty and $\boldsymbol{Z}_{i,t}$. The positive association between downstream uncertainty and investment is in line with Hypothesis 2 of Section II. We report several robustness checks related to these micro-level results in Section III.F.

C. Other Firm-Level Outcomes Under Supply-Chain Uncertainty

While a firm's investment rate is the only policy examined within the motivating model, the effects of the supply-chain uncertainties should not be confined to investment only. This is because a firm's decision to expand or shrink its productive capacity is typically accompanied by changes in marginal costs, such as those related to labor and working capital.

Table III
Firm-Level Investment Under Supply-Chain Uncertainties

The table reports the relation between a firm's investment rate and the contemporaneous level of the firm's upstream uncertainty (i.e., the uncertainty of the firm's suppliers), the firm's downstream uncertainty (i.e., the uncertainty of the firm's customers), and the firm's own uncertainty. The results are based on estimating equation (10), where $y_{i,t}$ is firm i's investment rate obtained from the most recent annual report as measured at time t. The benchmark upstream and downstream uncertainty measures are constructed at time t following the procedure outlined in Section III.A. All specifications include firm and year fixed effects, and even-numbered columns also include the following set of additional control variables: firm size, financial leverage, tangibility, Tobin's q, profitability, and past stock returns. Section III.A of the Internet Appendix provides definitions of all variables. All regressions are estimated using a panel of firm-year observations over the period 1976 to 2019, and t-statistics, reported in parentheses, are based on robust standard errors clustered at the firm level.

	(1)	(2)	(3)	(4)	(5)	(6)
σ(Own)	-0.11	-0.05	-0.08	-0.04	-0.12	-0.05
	(-12.60)	(-5.25)	(-4.72)	(-2.53)	(-11.81)	(-4.78)
σ(Upstream)			-0.03	-0.03		
			(-2.88)	(-2.80)		
σ(Downstream)					0.03	0.03
					(3.78)	(3.15)
Controls	No	Yes	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Adj.-R^2	0.38	0.41	0.50	0.52	0.37	0.40
Obs.	50,786	50,319	17,563	17,456	34,675	34344

Consistent with the investment-related results in Section III.B and the aforementioned intuition, Table IV shows that the asymmetry between upstream and downstream uncertainty also applies to other firm-level policies. We establish this by estimating equation (10) after replacing $y_{i,t}$ with either firm i's working capital growth (Panel A), employment growth (Panel B), COGS growth (Panel C), or intangibles growth (Panel D) at time t. Section III.A of the Internet Appendix provides details on the construction of each of these variables.

All panels in Table IV show that higher upstream uncertainty is associated with a statistically significant decline in each measure of interest. In contrast, higher downstream uncertainty has either a strictly positive impact (e.g., working capital) or an insignificant

Table IV
Firm-Level Economic Activity Under Supply-Chain Uncertainties

The table reports the relation between a firm's working capital growth rate (Panel A), employment growth rate (Panel B), COGS growth rate (Panel C), and intangible capital growth rate (Panel D) and the contemporaneous level of the firm's upstream uncertainty (i.e., the uncertainty of the firm's suppliers), the firm's downstream uncertainty (i.e., the uncertainty of the firm's customers), and the firm's own uncertainty. The results are based on estimating equation (10), where $y_{i,t}$ is a measure of firm i's economic activity or valuation ratio obtained from the most recent annual report as measured at time t. The benchmark upstream and downstream uncertainty measures are constructed at time t following the procedure outlined in Section III.A. All specifications include firm and year fixed effects, and even-numbered columns also include the following set of additional control variables: firm size, financial leverage, tangibility, Tobin's q, profitability, and past stock returns. Section III.A of the Internet Appendix provides definitions for all variables. All regressions are estimated using a panel of firm-year observations over the period 1976 to 2019, and t-statistics, reported in parentheses, are based on robust standard errors clustered at the firm level.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Panel A: Working capital				Panel B: Employment			
σ(Own)	-0.11	-0.06	-0.14	-0.08	-0.13	-0.12	-0.14	-0.11
	(-5.82)	(-2.68)	(-12.04)	(-6.17)	(-6.91)	(-5.41)	(-12.54)	(-8.41)
σ(Upstream)	-0.03	-0.03			-0.03	-0.03		
	(-2.48)	(-2.32)			(-2.17)	(-2.14)		
σ(Downstream)			0.03	0.02			0.02	0.01
			(3.32)	(2.18)			(2.13)	(0.72)
Adj.-R^2	0.13	0.17	0.14	0.18	0.13	0.16	0.13	0.15
Obs.	17,551	17,118	36,064	34,217	17,912	17,478	36,115	34,280
	Panel C: COGS				Panel D: Intangibles			
σ(Own)	-0.07	-0.04	-0.10	-0.05	-0.16	-0.09	-0.17	-0.09
	(-3.72)	(-1.70)	(-8.34)	(-3.33)	(-7.34)	(-3.86)	(-13.78)	(-7.06)
σ(Upstream)	-0.04	-0.04			-0.02	-0.03		
	(-2.97)	(-3.01)			(-1.83)	(-2.55)		
σ(Downstream)			0.03	0.02			0.02	0.01
			(2.86)	(1.73)			(2.31)	(0.80)
Adj.-R^2	0.15	0.16	0.13	0.12	0.18	0.22	0.14	0.19
Obs.	17,863	17,433	36,330	34,463	17,025	16,919	33,925	33,600
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

but nonnegative effect (e.g., employment). When accounting for firm-level controls, in all cases the negative marginal effect of upstream uncertainty is weakly larger (in absolute value) than the marginal effect of downstream uncertainty.

D. Downstream Uncertainty and Time-to-Build

Hypothesis 3 predicts that the asymmetry between upstream and downstream uncertainty is more pronounced for firms facing longer time-to-build periods. Notably, downstream uncertainty can hasten investment if a firm's time-to-build is sufficiently long. We verify this prediction empirically. We use several proxies for the length of each firm's time-to-build to establish that the positive interaction between downstream uncertainty and investment is much stronger and more significant among firms with longer time-to-build. To test the interaction between downstream uncertainty and time-to-build, we use the regression

$$\begin{aligned} y_{i,t} = \alpha_i + \delta_t + \beta_1 \sigma(\text{Own})_{i,t} + \beta_2 \sigma(\text{Downstream})_{i,t} \times \mathbb{I}\left[\text{Long}\right]_{i,t} \\ + \beta_3 \sigma(\text{Downstream})_{i,t} \times \mathbb{I}\left[\text{Short}\right]_{i,t} + \boldsymbol{\gamma}' \boldsymbol{Z}_{i,t} + \varepsilon_{i,t}, \end{aligned} \tag{11}$$

where $\mathbb{I}\left[\text{Long}\right]_{i,t}$ ($\mathbb{I}\left[\text{Short}\right]_{i,t}$) is an indicator variable that takes a value of one if firm i is defined as having a long (short) time-to-build at time t, and zero otherwise. All other variables in equation (11) follow the same definitions as in equation (10).

As time-to-build is unobservable at the firm level, we split our sample into two groups of firms based on three distinct proxies for time-to-build. Our first proxy uses each firm's depreciation rate.[21] As firms with lower depreciation rates typically have assets that are longer-lasting and require lengthy periods of planning, construction, and commissioning, we proxy for the length of a firm's time-to-build using the inverse of its depreciation rate. Our second proxy uses the sector classification of Gomes, Kogan, and Yogo (2009). This proxy is motivated by the fact that perishable goods typically require a shorter construction time than durable goods. Thus, we assign firms in industries that produce nondurable consumption goods or services to the short time-to-build group, while the long time-to-build group includes firms operating in other sectors (specifically, investment goods or durable consumption goods). Our third proxy is based on R&D intensity since firms with more R&D expenditures engage in longer-duration projects that require considerable time to develop

[21] We thank an anonymous associate editor for suggesting this proxy.

Table V

Investment, Downstream Uncertainty, and Time-to-Build

The table reports the relation between a firm's investment rate and the contemporaneous level of the firm's downstream uncertainty (i.e., the uncertainty of the firm's customers), conditional on one of three proxies for the length of the firm's time-to-build period. The results are based on estimating regression (11), where $y_{i,t}$ is firm i's investment rate obtained from the most recent annual report as measured at time t and downstream uncertainty is constructed at time t following the procedure outlined in Section III.A. In Panel A, the "Short" ("Long") dummy takes the value of one for firms that have an inverse depreciation rate that is below (above) the cross-sectional median value of inverse depreciation. In Panel B, the "Short" dummy takes the value of one for firms that produce nondurable consumption goods or services, while the "Long" dummy takes the value of one for firms operating in other sectors (namely, investment good producers and durable consumption goods producers). Firms are assigned to the two groups on the basis of the sector classification scheme of Gomes, Kogan, and Yogo (2009). In Panel C, the "Long" ("Short") dummy takes the value of one for firms that have a level of industry-adjusted R&D intensity that is greater (less) than the cross-sectional median value. Year and firm fixed effects are included in all specifications and additional control variables are included in all even-numbered columns. This set of additional control variables includes firm size, financial leverage, tangibility, Tobin's q, profitability, and past stock returns. Section III.A of the Internet Appendix provides definitions for all variables. All regressions are estimated using a panel of firm-year observations over the period 1976 to 2019, and t-statistics, reported in parentheses, are based on robust standard errors clustered at the firm level. Finally, the table also reports the p-value from a Wald test on the null hypothesis that the relation between downstream uncertainty and investment is the same for both short and long time-to-build firms ($H_0 : \beta_2 = \beta_3$ in equation (11)).

	Panel A: Depreciation		Panel B: Sector		Panel C: R&D	
	(1)	(2)	(3)	(4)	(5)	(6)
σ(Own)	-0.11	-0.05	-0.12	-0.05	-0.12	-0.05
	(-10.94)	(-4.58)	(-11.81)	(-4.78)	(-9.64)	(-3.85)
σ(Downstream) × Short	0.03	0.02	-0.01	-0.02	0.03	0.02
	(2.13)	(1.63)	(-2.08)	(-2.27)	(1.95)	(1.49)
σ(Downstream) × Long	0.06	0.05	0.03	0.03	0.06	0.05
	(4.60)	(4.04)	(3.79)	(3.15)	(3.91)	(3.24)
p(Long=Short)	0.06	0.07	0.00	0.00	0.04	0.07
Controls	No	Yes	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Adj.-R^2	0.38	0.41	0.37	0.40	0.37	0.39
Obs.	34,647	34,317	34,675	34,344	23,717	23,516

and produce.[22]

[22] Another possibility is that a firm's age can also proxy for time-to-build. Younger firms are typically endowed with more growth opportunities, and being young often entails undertaking longer-duration projects

Using each of the three aforementioned proxies to distinguish between short and long time-to-build firms, we conjecture and verify that β_2 in equation (11) is qualitatively and quantitatively larger than β_3. Table V reports the estimated values of these slope coefficients and shows that the impact of downstream uncertainty is substantially smaller, and at times insignificant, for firms with short time-to-build periods. In contrast, when using controls, the marginal effect of downstream uncertainty is strictly positive and often two times larger among firms with long time-to-build periods. Moreover, a Wald test on the null hypothesis that $\beta_2 = \beta_3$ is rejected at the 10% level of significance or better in all specifications, indicating that the positive effect of downstream uncertainty is indeed concentrated among firms with long time-to-build periods.[23]

E. *Supply-Chain Uncertainty and Investment Reversibility*

A key ingredient of the model in Section II is that the firm holds an option to abandon growth opportunities via disinvestment. As stated in Hypothesis 4, reversibility has a direct impact on the net benefit of waiting to invest if either upstream or downstream uncertainty rises. Greater reversibility should weaken the negative impact of upstream uncertainty but strengthen the positive impact of downstream uncertainty.

We check these predictions of the model by splitting the effects of upstream and downstream uncertainty among firms with high and low investment reversibility. Specifically, we estimate projections that are similar to equation (11) but interact each type of supply-chain uncertainty with indicator variables that classify firms with high or low reversibility. Here, we use the asset redeployability measure of Kim and Kung (2017) as our proxy for reversibility. We assume that firms with below-median asset redeployability have less reversible

(i.e., the investment projects of young firms are not "matured"). Consistent with this view, Section VIII.B in the Internet Appendix shows that the positive effect of downstream uncertainty is concentrated among younger firms.

[23]By contrast, the model predicts that the effect of upstream uncertainty is not materially changed with either the time-to-build horizon (see Figure 3) or with sequential investment (see Figure 4). In Section VIII.B of the Internet Appendix we confirm this in the data. There, we split the effect of upstream uncertainty between the short and long time-to-build groups and consistently fail to reject the null hypothesis that the marginal negative effect of upstream uncertainty is equal between the short and long time-to-build groups.

Table VI
Investment, Supply-Chain Uncertainties, and Reversibility

The table reports the relation between a firm's investment rate and the contemporaneous level of the firm's upstream uncertainty (i.e., the uncertainty of the firm's suppliers), downstream uncertainty (i.e., the uncertainty of the firm's customers), and the firm's own uncertainty, conditional on a proxy for the firm's degree of investment reversibility. The results are based on estimating a regression similar to equation (11), where $y_{i,t}$ is firm i's investment rate obtained from the most recent annual report as measured at time t, and upstream and downstream uncertainty are constructed at time t following the procedure outlined in Section III.A. The proxy for the degree of investment reversibility is based on the capital redeployability measure of Kim and Kung (2017). $HighReverse$ ($LowReverse$) is a dummy variable that takes the value of one for firms that have above- (below-) median capital redeployability. All specifications include firm and year fixed effects, and even-numbered columns also include the following set of additional control variables: firm size, financial leverage, tangibility, Tobin's q, profitability, and past stock returns. Section III.A of the Internet Appendix provides definitions for all variables. All regressions are estimated using a panel of firm-year observations from 1985, when the asset redeployability data become available, to 2019, and t-statistics, reported in parentheses, are based on standard errors clustered at the firm level.

	(1)	(2)	(3)	(4)
σ(Own)	-0.09	-0.05	-0.12	-0.05
	(-4.84)	(-2.74)	(-10.87)	(-4.24)
σ(Upstream) $\times$ HighReverse	-0.03	-0.03		
	(-1.37)	(-1.42)		
σ(Upstream) $\times$ LowReverse	-0.05	-0.05		
	(-3.39)	(-3.31)		
σ(Downstream) $\times$ HighReverse			0.06	0.05
			(4.06)	(3.60)
σ(Downstream) $\times$ LowReverse			0.03	0.02
			(2.18)	(1.32)
Controls	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Adj.-R^2	0.49	0.51	0.38	0.41
Obs.	16,611	16,512	31,753	31,499

investments, corresponding to higher abandonment costs.

Table VI confirms the predictions outlined above. First, upstream uncertainty is more negatively associated with investment for firms with harder-to-abandon projects. For instance, column (2) shows that a one-standard-deviation increase in upstream uncertainty decreases investment by 0.05 among firms with harder-to-abandon projects, an effect that is significant at the 1% level, but by only 0.03 among firms with easier-to-abandon projects, a statistically insignificant change. Second, downstream uncertainty has a more positive effect on the investment rates of firms with easier-to-reverse investments. While a one-standard-deviation increase in downstream uncertainty increases investment rates by 0.05 among firms with reversible projects, it has no significant effect among firms with hard-to-abandon projects.

F. *Robustness of the Firm-Level Results*

Section VIII.B of the Internet Appendix demonstrates that the main results regarding investment under supply-chain uncertainty are robust along several dimensions. In particular, we obtain results that are qualitatively and quantitatively similar to those reported in Table III when we (i) measure uncertainty using either option-implied volatility or idiosyncratic stock return volatility, (ii) construct each supply-chain uncertainty measure by weighting the uncertainty of each supplier by its sales and the uncertainty of each customer by its COGS, (iii) focus on the recent period 2009 to 2019, in which the number of observations for suppliers and customers is more balanced, (iv) cluster standard errors over firm and time, (v) include industry-by-time fixed effect, and (vi) exclude one sector at a time. In all cases, upstream and downstream uncertainty have an asymmetric impact on a firm's decision to invest beyond the firm's own inherent uncertainty. We also show that the results are pervasive across both small and large firms, but are stronger for younger firms, which typically have longer time-to-build projects.

Section VIII.C in the Internet Appendix discusses the connection between our results and the degree of competition among suppliers or customers. For instance, we show that the effects of supply-chain uncertainty are more pronounced when the customer (supplier) base of the suppliers (customers), who sell (buy) their goods to (from) the focal firm, is more diverse. Section VIII.D of the Internet Appendix also uses an IV approach in the spirit of Alfaro, Bloom, and Lin (2024) to establish a positive (negative) effect of upstream (downstream) uncertainty on investment.

IV. Macro-Level Evidence

This section shows that the micro-level results documented in Section III also hold at the macro level. That is, the asymmetric effect of changes in upstream and downstream uncertainty on firm-level investment and other outcomes carries over to the macroeconomy. Intuitively, since supplier-specific (customer-specific) uncertainty is associated with decreased (increased) economic activity, common uncertainty shocks to firms that tend to operate relatively upstream (downstream) should induce a negative (positive) effect on other firms. Related, the effects of uncertainty shocks often display a "fractal" property that manifests at different levels of aggregation (see, for example, Bloom (2014)).[24]

We construct *macro-level* measures of the supply-chain uncertainties in Section IV.A. Section IV.B then shows that shocks to macro-level upstream (downstream) uncertainty lead to a deterioration (improvement) in aggregate economic growth and asset prices. Finally, Section IV.C considers the cyclicality of each type of macro-level uncertainty.

A. Data

Our macro-level analysis uses the same set of firms and the same time frame as in Section III.A above.

Constructing Vertical Position. The macro-level analysis of the effects of the supply-chain uncertainties requires moving beyond the measures of "upstream" and "downstream" uncertainty used in Section III. This aggregate-level analysis requires that we consider the *absolute* vertical position of each firm in the production network, rather than the *relative* position of each firm in a given supply chain.[25] To see why, consider a firm i that has a high vertical position (i.e., is further from final consumers). In the firm-level analysis, it was sufficient to define any supplier firm s that sells to firm i as an upstream firm *relative* to firm i. However, firm s is not necessarily more upstream than firm i in an *absolute* sense.

[24] As a broad motivation, we set the model parameters from Section II to match aggregate volatility and confirm that the key results hold using an aggregate calibration (see Section VII.E of the Internet Appendix.)

[25] We use the term "vertical position" or "upstreamness" to denote the distance of a firm from the producers of final consumption goods.

For example, while firm s may sell some of its output to firm i, firm s may sell most of its goods to final consumers. Therefore, while firm s is *relatively* upstream from the perspective of firm i in a specific supply chain, firm s may still have a lower *absolute* vertical position when considering the production network as a whole.

We measure firms' vertical position in the production network using the Input-Output (I-O) tables constructed by the Bureau of Economic Analysis (BEA) following Gofman, Segal, and Wu (2020). These tables, which are also referred to as the BEA's Make and Use tables, record the dollar flows of commodities between industries and their use for final consumption. There are two primary benefits of using the BEA I-O tables rather than Compustat Segments or FactSet data to measure *absolute* vertical position. First, the I-O tables span a significantly longer time period than either alternative data set, allowing us to measure the vertical positions of firms as far back as the 1970s. Second, the I-O tables account for both the *existence* and the *importance* of inter-industry links, as they record the dollar flows of goods between different parts of the economy.[26] Nonetheless, the I-O tables are published only once every five years.

The intuition behind the absolute vertical position measure based on the BEA's I-O tables is that industries that produce a higher (lower) dollar value of commodities that flow to final consumers are less (more) upstream in the production network. With this intuition in mind, we follow Antràs and Chor (2018) to combine the normalized I-O tables, denoted by $\tilde{M}$ and $\tilde{U}$, respectively, to compute a vertical position (upstreamness) score for each industry. Specifically, the vertical position score of industry k at time t is the k^{th} element of

$$VP_t = \left(I_{N_t \times N_t} - \tilde{M}_{N_t \times C_t} \times \tilde{U}_{C_t \times N_t}\right)^{-1} \iota. \tag{12}$$

Here, VP_t is a $N_t \times 1$ vector of vertical position scores, I is an identity matrix, ι is a vector of

[26] While Compustat Segments contains a small number of observations for the 1970s, there is an insufficient number of interfirm links to compute absolute measures of vertical position. For instance, there are only 1087 links in Compustat Segments for the entire 1976 to 1979 period. Moreover, although Compustat Segments data also contain information on the flow of sales between a firm and its most critical customers, these data are limited to the small set of firms that appear in the Segments data. Similar data on flows are unavailable for the more comprehensive FactSet data set.

ones, and the subscripts N_t and C_t represent the total number of industries and commodities in the BEA tables at time t, respectively. The Leontief inverse in equation (12) captures the importance of each industry as a *direct* and *indirect* supplier to all other industries (Carvalho and Tahbaz-Salehi, 2019).[27]

We apply the procedure underlying equation (12) to the BEA's I-O tables reported for years 2012, 2007, 2002, 1997, 1992, 1987, 1982, and 1977. For observations preceding (following) 1977 (2012), we assume that an industry's vertical position is identical to that in 1977 (2012). For observations between successive releases of the I-O tables, we assume that an industry's vertical position in year t is the same as that based on the previously released I-O tables. For the purpose of brevity, Section V of the Internet Appendix provides a detailed description of how we clean, filter, and normalize the BEA I-O tables, and then match the resulting industry-level vertical position scores to the CRSP/Compustat universe.

Constructing Macro-Level Uncertainties. At the end of each month from January 1976 to December 2019, we sort firms into two groups based on the cross-sectional distribution upstreamness scores from equation (12). Firms in industries with a vertical position score above (below) the 90^{th} (10^{th}) percentile of the distribution of these scores are considered upstream (downstream).[28] These cutoffs provide a clear distinction between upstream and downstream firms while producing portfolios that contain hundreds of individual firms. The results are robust to instead using the 20^{th} and 80^{th} percentiles as coarser cutoffs.

To mimic our micro-level measures of upstream and downstream uncertainty from Section III, we construct our baseline measures of macro-level upstream and downstream uncertainty using the realized volatility of firm-level stock returns. Specifically, we use the daily returns

[27]The properties of this matrix satisfy the conditions listed in Carvalho and Tahbaz-Salehi (2019) for this matrix's inverse to exist. Therefore, defining $T_t \equiv \tilde{M}_{N_t \times C_t} \times \tilde{U}_{C_t \times N_t}$, and noting that $(I_{N_t} - T_t)^{-1} = \sum_{z=0}^{\infty} T_t^k$, we can write element (i,j) of $(I_{N_t} - T_t)^{-1}$ as $t_{i,j} + \sum_{k=1}^{N_t} t_{i,k} t_{k,j} + \ldots$. Here, the first term reflects the importance of industry j as a supplier to industry i, the second term accounts for links between industry j and each industry r that, in turn, supplies to industry i, and the ellipses account for all other indirect links between industry j and industry i.

[28]In Section X of the Internet Appendix, we report those industries that fall most frequently in the macro-level upstream and downstream categories. The upstream industries generally produce investment goods, raw materials, defense, and vehicles, while the downstream industries generally produce durable consumer goods, such as footwear, or services, such as social services.

of each firm i in each month t to compute each firm's realized stock return volatility. We then compute the value-weighted average stock return volatility across all firms assigned to the upstream and downstream portfolio, respectively.[29] This procedure yields a monthly time series of aggregate uncertainty for both upstream and downstream firms. Finally, since most macroeconomic variables of interest are recorded at a lower frequency than stock returns, we aggregate the monthly time series of the macro-level upstream and downstream uncertainties into quarterly time series by computing the time-series mean of upstream (downstream) uncertainty over the three months preceding the end of each quarter t.[30]

B. *Upstream and Downstream Uncertainty: Aggregate Impulse Response Functions*

We show that macro-level upstream (downstream) uncertainty negatively (positively) predicts aggregate economic growth and asset prices by estimating impulse response functions (IRFs) that determine how macro-level uncertainty shocks impact key variables of interest.

The IRFs are estimated using the Smooth Local Projections (SLPs) from Barnichon and Brownlees (2019). Unlike IRFs from the local projections (LPs) of Jordà (2005), which require the estimation of *separate* predictive regressions for each forecast horizon, SLPs assume that IRFs are a *smooth* function of the forecast horizon. This allows SLPs to strike a balance between the benefits of IRFs from vector autoregressions, which are efficient for correctly specified models, and IRFs from LPs which are more robust to model misspecification but are potentially noisy.[31] Specifically, our IRFs are based on the following h-step-ahead predictive

[29] To compute these measures of firm-level volatility, we (i) adjust stock returns in CRSP daily for delisting returns, (ii) require that each firm has at least 15 valid stock returns in a given month, and (iii) construct realized variation of stock returns within each month.

[30] Rather than average uncertainty over the three months in a given quarter, we could alternatively define quarterly uncertainty as the value of uncertainty in the final month of the quarter. The results are robust to using this alternative method of aggregation.

[31] Barnichon and Brownlees (2019) strike a balance between VARs and LPs by estimating LPs and then using penalized B-splines to shrink the resulting IRF towards a smooth polynomial function.

regressions for horizons of $h \in \{1, ..., H\}$ quarters ahead:

$$y_{t+h} = \beta_{0(h)} + \beta_{1(h)} y_t + \beta_{2(h)} \sigma_{U,t} + \beta_{3(h)} \sigma_{D,t} + \sum_{p=1}^{P} \boldsymbol{\gamma}'_{p(h)} \boldsymbol{\Gamma}_{t-p} + \varepsilon_{t+h}. \tag{13}$$

Here, y_{t+h} denotes one of the six following aggregate-level variables at time $t + h$: the quarterly real growth rates of industrial production, consumption, private investment, and gross domestic product, the level of the market's price-dividend ratio, and the risk-free rate. The term $\sigma_{U,t}$ ($\sigma_{D,t}$) denotes macro-level upstream (downstream) uncertainty at time t, constructed following the procedure in Section IV.A, and $\boldsymbol{\Gamma}_{t-p}$ is a vector of controls in quarter $t - p$. These controls include the dependent variable of interest, the two macro-level supply-chain uncertainties, the excess market return, the term spread, the default spread, and the inflation rate. These last four variables are included to account for the negative correlation between the level of macroeconomic activity and macroeconomic uncertainty. We set P equal to four, but results are robust to setting P equal to one. Moreover, our results are robust to controlling for the contemporaneous values of the various controls in $\boldsymbol{\Gamma}_t$. Finally, we standardize all variables in equation (13) for ease of interpretation and comparability between figures.

Figure 6 (Figure 7) shows the IRFs for the six outcome variables with respect to a one-standard-deviation increase in macro-level upstream (downstream) uncertainty. Each figure displays the mean response of each variable to the given shock (solid lines) alongside the 90% confidence interval (dashed lines).

Figure 6 shows that the IRFs from macro-level upstream uncertainty to all variables are negative and significant. That is, higher upstream uncertainty leads to an economic contraction. One quarter after an upstream uncertainty shock, industrial production and GDP decline by about 0.15 standard deviations whereas consumption and investment decline by about 0.10 standard deviations. These effects are statistically significant and persist for about one year. Higher upstream uncertainty also drops the market's valuation ratio and the risk-free rate for about 12 and six quarters, respectively. These results together echo the

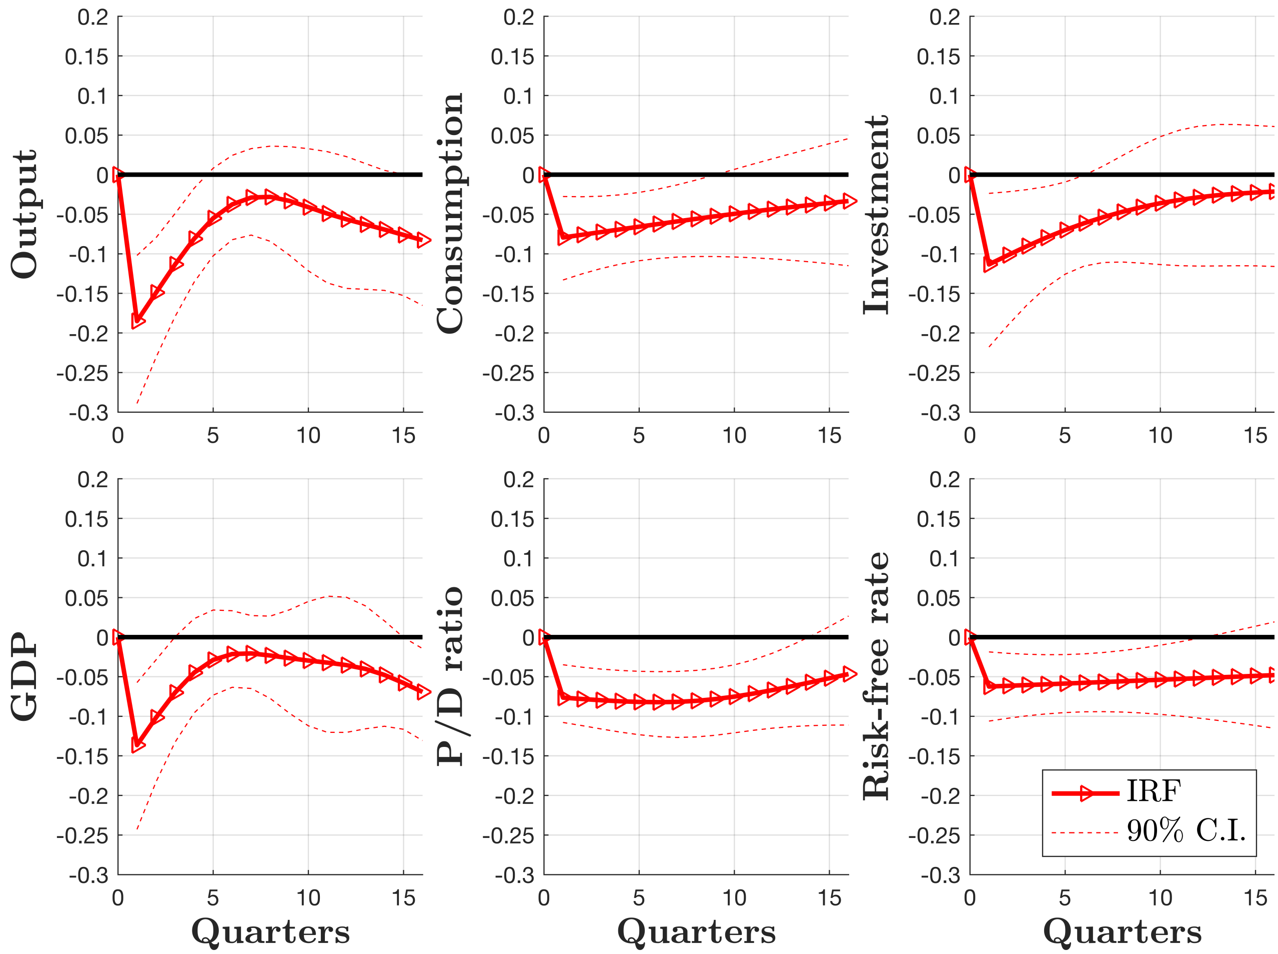

Figure 6. Impulse responses from macro-level upstream uncertainty shocks. This figure shows impulse response functions (IRFs) that display the effects of a one-standard-deviation shock to macro-level upstream uncertainty ($\sigma_{U,t}$) on the quarterly growth rates of industrial production, real consumption, real nonresidential investment, and real GDP, and the levels of the aggregate price-dividend ratio and the risk-free rate. We estimate the IRFs by applying the smooth local projection method of Barnichon and Brownlees (2019) to equation (13) for forecast horizons that range from one to 16 quarters ahead. We measure macro-level upstream uncertainty by following the procedure described in Section IV.A. Section III.B of the Internet Appendix provides detailed descriptions of the variables included in equation (13). The estimated IRFs are denoted by solid lines, while 90% confidence intervals are represented by dashed lines. The sample period is 1976Q1 to 2019Q4.

firm-level results for upstream uncertainty and are consistent with the traditional negative association between uncertainty and economic growth (e.g., Ramey and Ramey (1995)).

In contrast to upstream uncertainty, Figure 7 shows that macro-level downstream un-

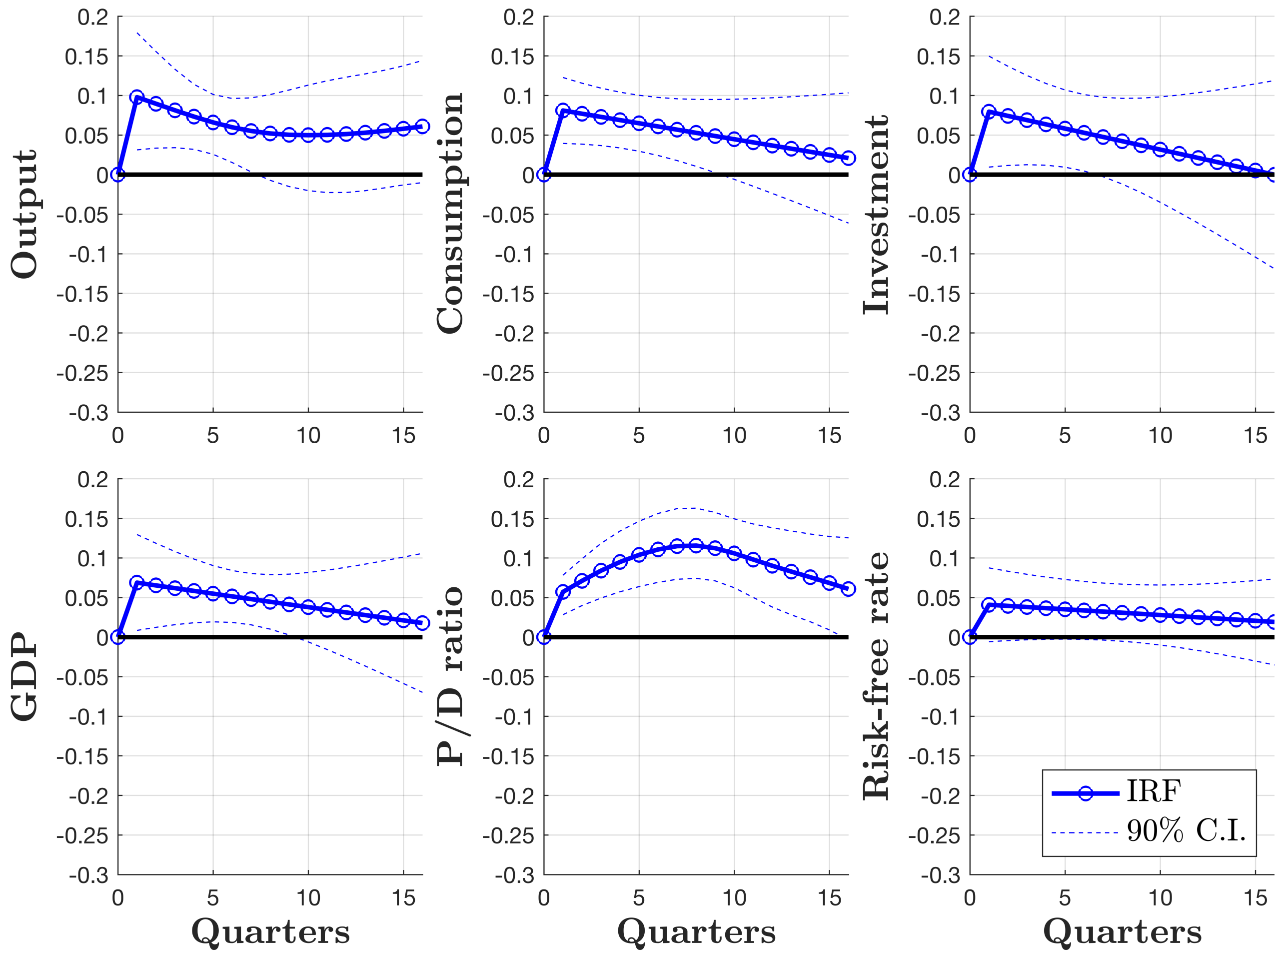

Figure 7. Impulse responses from macro-level downstream uncertainty shocks. This figure shows impulse response functions (IRFs) that display the effects of a one-standard-deviation shock to macro-level downstream uncertainty ($\sigma_{D,t}$) on the quarterly growth rates of industrial production, real consumption, real nonresidential investment, and real GDP, and the levels of the aggregate price-dividend ratio and the risk-free rate. We estimate the IRFs by applying the smooth local projection method of Barnichon and Brownlees (2019) to equation (13) for forecast horizons that range from one to 16 quarters ahead. We measure macro-level downstream uncertainty by following the procedure described in Section IV.A. Section III.B of the Internet Appendix provides detailed descriptions of the variables included in equation (13). The estimated IRFs are denoted by solid lines, while 90% confidence intervals are represented by dashed lines. The sample period is 1976Q1 to 2019Q4.

certainty shocks have an expansionary impact and lead to higher economic growth. A one-standard-deviation shock to macro-level downstream uncertainty induces industrial production, consumption, investment, and GDP growth to each rise by about 0.1 standard deviations for at least four quarters, a statistically significant change. Likewise, the market's price-dividend ratio rises by a similar and statistically significant amount, remaining elevated for almost 12 quarters. However, the risk-free rate is largely unaffected. Overall, the positive impact of downstream uncertainty is even more pronounced at the macro level than the firm level.

Comparing the IRFs in Figures 6 and 7 shows that macro-level upstream and downstream uncertainty have an *asymmetric* effect on aggregate variables in terms of signs and magnitudes. The positive impacts of macro-level downstream uncertainty shocks are quantitatively more muted than the negative impacts of macro-level upstream uncertainty shocks (in absolute value). In general, a one-standard-deviation shock to upstream uncertainty leads to responses that are up to 50% larger in magnitude than the responses to downstream uncertainty shocks. These differences in the sign and magnitude of the supply-chain uncertainty shocks are broadly consistent with the motivating model.

Marginal Utility of Investors. Section IX of the Internet Appendix demonstrates that higher macro-level upstream (downstream) uncertainty is associated with bad (good) economic states, in which the marginal utility of investors is higher (lower). This finding is consistent with the fact that macro-level upstream (downstream) uncertainty predicts future consumption growth negatively (positively). As investors are concerned primarily about changes in *expected* rather than *realized* outcomes, we first construct forward-looking measures for each type of uncertainty. We then posit that the pricing kernel (potentially) depends on these forward-looking uncertainties, and elicit the market price of risk associated with each type of uncertainty using a GMM estimation of an Euler equation for asset prices. We find that macro-level upstream uncertainty shocks do indeed carry a negative market price of risk (raise marginal utility), whereas macro-level downstream uncertainty shocks carry a positive market price of risk (lower marginal utility).

Robustness and Extensions. Section X of the Internet Appendix shows IRFs of macro-level upstream and downstream uncertainty on both (i) the subcomponents of private fixed investment from the BEA and (ii) aggregate outcomes based on public firms only. The results are generally stronger for types of investment with longer time-to-build (e.g., structures) and when restricting the sample to public firms, the returns of which are used to construct the macro-level uncertainty measures. Section X of the Internet Appendix also shows the results of controlling for investment-specific technology (IST) shocks in the local projection (13) and finds virtually identical results as the benchmark.

We conduct a host of other robustness checks and verify that we obtain similar IRFs when we (i) use either option-implied volatility or idiosyncratic stock return volatility to measure uncertainty; (ii) estimate the IRFs using the local projection method of Jordà (2005); (iii) use the 20^{th} and 80^{th} percentiles of the cross-sectional distribution of vertical position scores to define the sets of upstream and downstream firms; and (iv) convert the monthly time series of uncertainty into a quarterly time series by retaining the last monthly observation in a given quarter rather than averaging uncertainty across observations in each quarter.

C. Supply-Chain Uncertainties and the COVID-19 Crisis

Cyclicality. By jointly controlling for both upstream and downstream uncertainty in equation (13), we isolate the component of downstream uncertainty that is orthogonal to upstream uncertainty. This orthogonal component of downstream uncertainty is procyclical. To illustrate this point, Figure 8 plots the time series of both macro-level upstream uncertainty (Panel A) and the orthogonal component of downstream uncertainty (Panel B). This orthogonal component is obtained by projecting downstream uncertainty on contemporaneous upstream uncertainty and computing the residuals.

Figure 8 shows that macro-level upstream uncertainty is countercyclical, typically peaking around the NBER recession, including the Great Recession, the dot-com crash, and the 1987 stock market crash. In contrast, the orthogonal component of downstream uncertainty typically declines in recessions and rises to its highest level during the technological boom of the late 1990s. The countercyclicality (procyclicality) of upstream (orthogonal downstream) macro-level uncertainty is in line with the IRFs in Section IV.B.

COVID-19. The COVID-19 pandemic brought an almost unprecedented level of un-

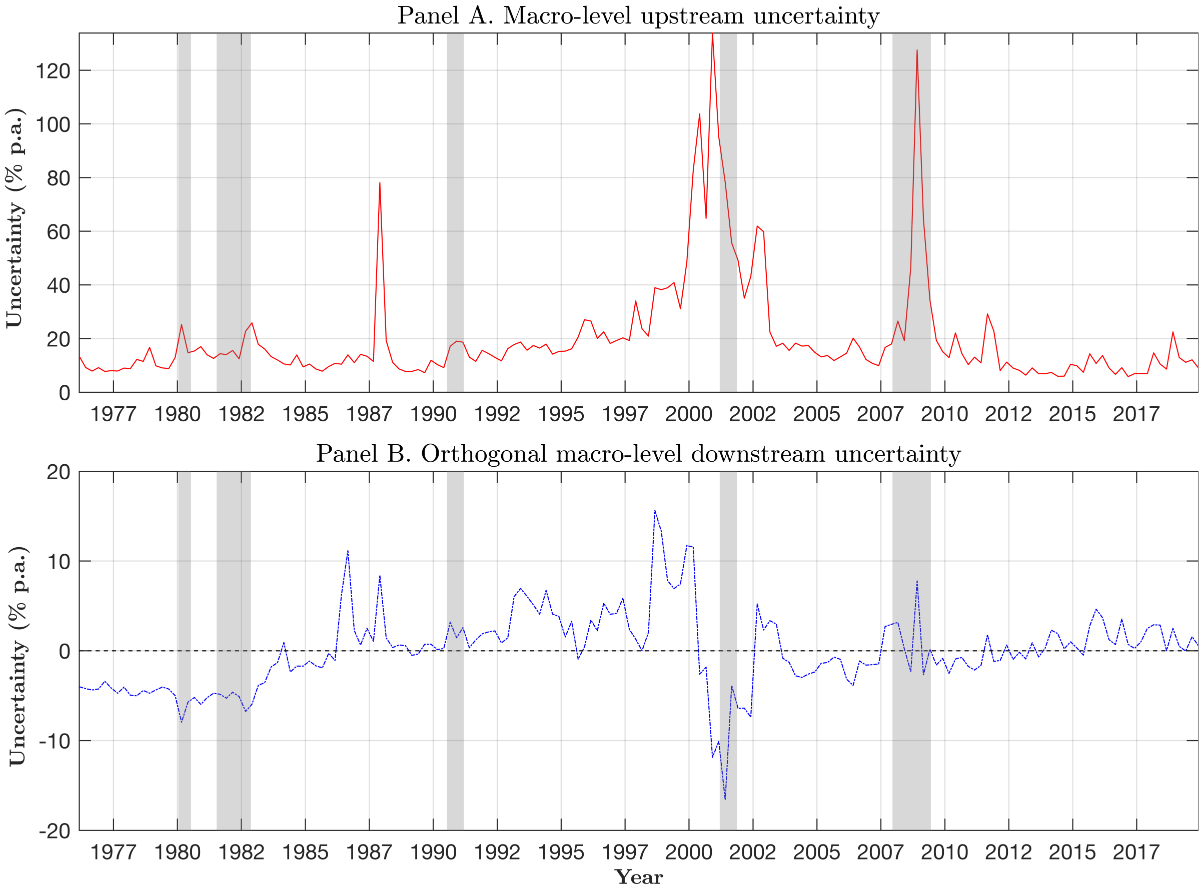

Figure 8. Macro-level upstream uncertainty and orthogonal downstream uncertainty. This figure shows the quarterly time series of macro-level upstream uncertainty (Panel A) and the orthogonal component of macro-level downstream uncertainty (Panel B). We compute the macro-level uncertainties following the procedure outlined in Section IV.A. The orthogonal component of downstream uncertainty represents the residuals from a projection of macro-level downstream uncertainty ($\sigma_{D,t}$) on upstream uncertainty ($\sigma_{U,t}$). The data span the period 1976Q1 to 2019Q4 and shaded regions represent NBER recessions.

certainty to financial markets. This was particularly manifested by the VIX index, which increased by nearly 200% in the first quarter of 2020, as shown in Panel A of Figure 9. While part of this increase likely reflects higher risk aversion, there is little doubt that the global pandemic was both a negative first-moment shock (as seen by the steep decline in the industrial production index in Panel B) and a positive second-moment shock. At the onset of the pandemic, uncertainty spiked, for example, due to the unknown nature of the disease, duration of imposed shutdowns, and timing of a potential vaccine.

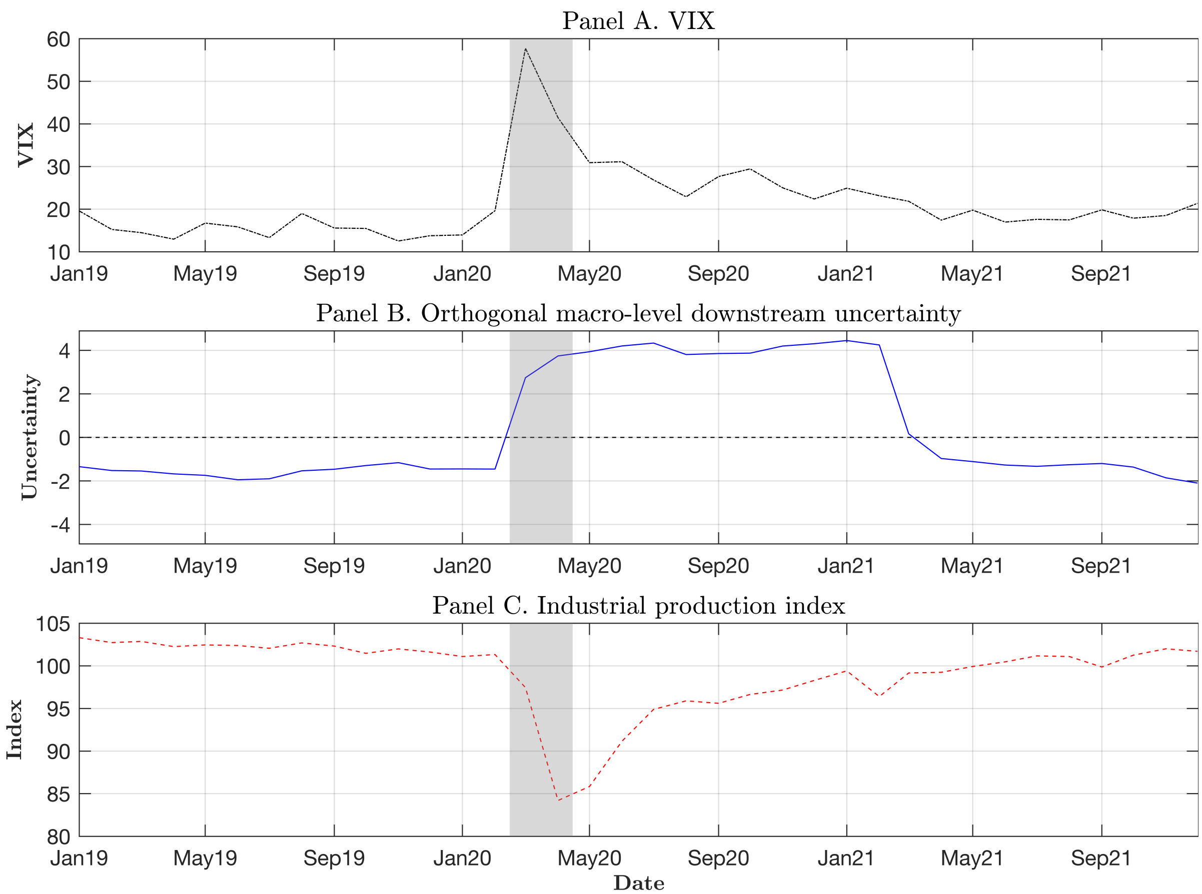

Figure 9. COVID-19 crisis: VIX, downstream uncertainty, and industrial production. This figure shows the monthly time series of (i) the VIX index (Panel A), (ii) the orthogonal component of macro-level downstream uncertainty (Panel B), and (iii) the level of the industrial production index (Panel C). Data on the VIX and industrial production index are from FRED. The component of downstream uncertainty that is orthogonal to upstream uncertainty is obtained from the residuals of a contemporaneous projection of downstream uncertainty onto upstream uncertainty, annualized. Each time series spans January 2019 to December 2021.

Was the increased uncertainty associated with the COVID-19 crisis driven by macro-level upstream or downstream uncertainty? We address this question by extending the time series of each uncertainty to December 2021, in an out-of-sample fashion. We assume that each industry's vertical position is identical to that implied by our sample's latest I-O tables, and extend the vertical position time series by following the same procedure described in Section IV.A.

Panel B of Figure 9 shows the orthogonal component of downstream uncertainty at the onset of the COVID-19 crisis. By construction, the mean of the orthogonal downstream uncertainty is zero and, as discussed above, this orthogonal component has turned negative in most past recessions. Notably, we find that orthogonal downstream uncertainty spikes in March 2020. Thus, although downstream uncertainty is dominated by upstream uncertainty in most contractions, the opposite was true at the beginning of the COVID-19 crisis. Nonetheless, the orthogonal downstream uncertainty turned negative in early 2021, in line with the generally more dominant role of upstream uncertainty.

These results shed light on the potential origins of the economic recession induced by COVID-19 and its progression. First, and a priori, it is unclear whether the initial COVID-19-induced contraction represents a supply-side or demand-side shock. Our results suggest that the latter is more likely. A possible narrative is that the sharp decline in economic activity during the first half of 2020 was driven largely by lockdown restrictions and precautionary consumer behavior. This negative demand shock induced an impact upstream, but primarily affected downstream firms. Second, our findings in Section IV.B suggest that higher downstream uncertainty positively predicts future economic growth. Consistent with these findings, along with the fact that downstream uncertainty was more dominant at the onset of the COVID-19 episode, the recession was short-lived.[32]

V. Conclusion

We examine the implications of uncertainties that originate in different locations of a firm's supply-chain environment on the firm's investment and real economic activity. Higher

[32] In Section X of the Internet Appendix we present the projected path of industrial production from September 2020 onward based on estimates of the smooth local projections from equation (13) over 1976 to 2019 and the state of the macro-level upstream and downstream uncertainties in September 2020. Notably, this out-of-sample analysis demonstrates that when downstream uncertainty's impact is shut down, the recovery from the COVID-19 recession would have been much more muted compared to the data.

upstream (downstream) uncertainty that stems from suppliers (customers) is negatively (positively) related to firm-level investment. This asymmetry not only holds at the firm level but is even stronger at the aggregate level. Specifically, macro-level upstream (downstream) uncertainty leads to a decline (increase) in key macroeconomic and financial market variables, such as output, consumption, investment, and the market's price-to-dividend ratio, and increases (decreases) the marginal utility of investors.

Our motivating production-based model predicts the aforementioned asymmetry. The key feature of the model is the realistic delay between the time when a firm initiates an investment project and the time when the firm receives its first revenues from the project. This "time-to-build" period implies that upstream uncertainty is related to uncertainty about the input price paid to suppliers and impacts the firm in the short run whereas downstream uncertainty is related to uncertainty about the output price received from customers and impacts the firm in the longer run. Upstream uncertainty depresses investment due to the traditional "bad news principle," while downstream uncertainty can sharply increase the opportunity cost of waiting via growth option channels and hasten investment.

Empirically, our study takes a bottom-up approach to examine the relationship between the two facets of supply-chain uncertainty and firm-level outcomes. Higher upstream uncertainty affects focal firms by suppressing firm-level investment. Higher downstream uncertainty impacts focal firms more weakly and never decreases, but rather often spurs, firm-level economic activity. Moreover, the positive relation between downstream uncertainty and investment *strengthens* for firms with longer time-to-build periods and greater investment reversibility— consistent with our production-based real-option model. Moreover, we construct *aggregate* measures of the supply-chain uncertainties and show that the results scale up to the macro level.

Overall, our findings suggest that while higher uncertainty is often associated with lower investment and asset prices, the effects of uncertainty are more nuanced and depend on where the uncertainty arises in the supply chain. Although upstream uncertainty is associated with contractions, downstream uncertainty may have an *expansionary* impact. This finding has implications for policymakers who may choose to react to the two supply-chain uncertainties differently. For example, policies that increase the time required to enter the product market (e.g., more trials and testing) may *promote* investment if downstream uncertainty is dominant. We leave the theoretical exploration of such policies for future research.

REFERENCES

Abel, Andrew B., 1983, Optimal investment under uncertainty, *American Economic Review* 73, 228–233.

Abel, Andrew B., and Janice C. Eberly, 1994, A unified model of investment under uncertainty, *American Economic Review* 84, 1369–1384.

Acemoglu, Daron, Ufuk Akcigit, and William Kerr, 2016, Networks and the macroeconomy: An empirical exploration, *NBER Macroeconomics Annual* 30, 273–335.

Ai, Hengjie, and Dana Kiku, 2016, Volatility risks and growth options, *Management Science* 62, 741–763.

Alfaro, Ivan, Nicholas Bloom, and Xiaoji Lin, 2024, The finance uncertainty multiplier, *Journal of Political Economy* 132, 577–615.

Antràs, Pol, and Davin Chor, 2018, On the measurement of upstreamness and downstreamness in global value chains, *World trade evolution: Growth, productivity and employment* 5, 126–194.

Arellano, Cristina, Yan Bai, and Patrick J. Kehoe, 2019, Financial frictions and fluctuations in volatility, *Journal of Political Economy* 127, 2049–2103.

Atalay, Enghin, Ali Hortaçsu, and Chad Syverson, 2014, Vertical integration and input flows, *American Economic Review* 104, 1120–1148.

Baker, Scott R., and Nicholas Bloom, 2013, Does uncertainty reduce growth? Using disasters as natural experiments, NBER Working paper 19475.

Baker, Scott R., Nicholas Bloom, and Steven J. Davis, 2016, Measuring economic policy uncertainty, *Quarterly Journal of Economics* 131, 1593–1636.

Bansal, Ravi, Varoujan Khatchatrian, and Amir Yaron, 2005, Interpretable asset markets?, *European Economic Review* 49, 531–560.

Bansal, Ravi, Dana Kiku, Ivan Shaliastovich, and Amir Yaron, 2014, Volatility, the macroeconomy, and asset prices, *Journal of Finance* 69, 2471–2511.

Baqaee, David Rezza, 2018, Cascading failures in production networks, *Econometrica* 86, 1819–1838.

Bar-Ilan, Avner, and William C. Strange, 1996, Investment lags, *American Economic Review* 86, 610–622.

Barnichon, Regis, and Christian Brownlees, 2019, Impulse response estimation by smooth local projections, *Review of Economics and Statistics* 101, 522–530.

Barrot, Jean-Noël, and Julien Sauvagnat, 2016, Input specificity and the propagation of idiosyncratic shocks in production networks, *Quarterly Journal of Economics* 131, 1543–1592.

Basu, Susanto, and Brent Bundick, 2017, Uncertainty shocks in a model of effective demand, *Econometrica* 85, 937–958.

Bekaert, Geert, Eric C. Engstrom, and Nancy R. Xu, 2022, The time variation in risk appetite and uncertainty, *Management Science* 68, 3975–4004.

Belo, Frederico, Xiaoji Lin, and Santiago Bazdresch, 2014, Labor hiring, investment, and stock return predictability in the cross section, *Journal of Political Economy* 122, 129–177.

Bernanke, Ben S., 1983, The determinants of investment: Another look, *American Economic Review* 73, 71–75.

Bianchi, Francesco, Howard Kung, and Mikhail Tirskikh, 2023, The origins and effects of macroeconomic uncertainty, *Quantitative Economics* 14, 855–896.

Bloom, Nicholas, 2009, The impact of uncertainty shocks, *Econometrica* 77, 623–685.

Bloom, Nicholas, 2014, Fluctuations in uncertainty, *Journal of Economic Perspectives* 28, 153–176.

Boehm, Christoph E., Aaron Flaaen, and Nitya Pandalai-Nayar, 2019, Input linkages and the transmission of shocks: Firm-level evidence from the 2011 Tōhoku earthquake, *Review of Economics and Statistics* 101, 60–75.

Boguth, Oliver, and Lars-Alexander Kuehn, 2013, Consumption volatility risk, *Journal of Finance* 68, 2589–2615.

Bretscher, Lorenzo, Alex Hsu, and Andrea Tamoni, 2020, Fiscal policy driven bond risk premia, *Journal of Financial Economics* 138, 53–73.

Carvalho, Vasco M., Makoto Nirei, Yukiko U. Saito, and Alireza Tahbaz-Salehi, 2021, Supply chain disruptions: Evidence from the great East Japan earthquake, *Quarterly Journal of Economics* 136, 1255–1321.

Carvalho, Vasco M., and Alireza Tahbaz-Salehi, 2019, Production networks: A primer, *Annual Review of Economics* 11, 635–663.

Chaney, Thomas, David Sraer, and David Thesmar, 2012, The collateral channel: How real estate shocks affect corporate investment, *American Economic Review* 102, 2381–2409.

Chang, Huifeng, Adrien d'Avernas, and Andrea L. Eisfeldt, 2024, Bonds vs. equities: Information for investment, *Journal of Finance* 79, 3893–3941.

Christiano, Lawrence J, Roberto Motto, and Massimo Rostagno, 2010, Financial factors in economic fluctuations, Working paper, ECB.

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

Dew-Becker, Ian, Stefano Giglio, Anh Le, and Marius Rodriguez, 2017, The price of variance risk, *Journal of Financial Economics* 123, 225–250.

Dixit, Avinash, 1992, Investment and hysteresis, *Journal of Economic Perspectives* 6, 107–132.

Dou, Winston Wei, 2017, Embrace or fear uncertainty: Growth options, limited risk sharing, and asset prices, Working paper, University of Pennsylvania.

Drechsler, Itamar, and Amir Yaron, 2011, What's vol got to do with it, *Review of Financial Studies* 24, 1–45.

Engel, Robert F., and Jose G. Rangel, 2008, The spline GARCH model for low-frequency volatility and its macroeconomics causes, *Review of Financial Studies* 21, 1187–1222.

Fernández-Villaverde, Jesús, Pablo A. Guerrón-Quintana, Keith Kuester, and Juan Rubio-Ramírez, 2015, Fiscal volatility shocks and economic activity, *American Economic Review* 105, 3352–3384.

Gilchrist, Simon, Jae W. Sim, and Egon Zakrajšek, 2014, Uncertainty, financial frictions, and investment dynamics, NBER Working paper 20038.

Gofman, Michael, Gill Segal, and Youchang Wu, 2020, Production networks and stock returns: The role of vertical creative destruction, *Review of Financial Studies* 33, 5856–5905.

Gomes, João F., Leonid Kogan, and Motohiro Yogo, 2009, Durability of output and expected stock returns, *Journal of Political Economy* 117, 941–986.

Grullon, Gustavo, Evgeny Lyandres, and Alexei Zhdanov, 2012, Real options, volatility, and stock returns, *Journal of Finance* 67, 1499–1537.

Gulen, Huseyin, and Mihai Ion, 2016, Policy uncertainty and corporate investment, *Review of Financial Studies* 29, 523–564.

Hartman, Richard, 1972, The effects of price and cost uncertainty on investment, *Journal of Economic Theory* 5, 258–266.

Herskovic, Bernard, Bryan Kelly, Hanno Lustig, and Stijn Van Nieuwerburgh, 2016, The common factor in idiosyncratic volatility: Quantitative asset pricing implications, *Journal of Financial Economics* 119, 249–283.

Hulten, Charles R., 1978, Growth accounting with intermediate inputs, *Review of Economic Studies* 45, 511–518.

Husted, Lucas, John Rogers, and Bo Sun, 2020, Monetary policy uncertainty, *Journal of Monetary Economics* 115, 20–36.

Jordà, Òscar, 2005, Estimation and inference of impulse responses by local projections, *American Economic Review* 95, 161–182.

Kim, Hyunseob, and Howard Kung, 2017, The asset redeployability channel: How uncertainty affects corporate investment, *Review of Financial Studies* 30, 245–280.

Leahy, John V., and Toni M. Whited, 1996, The effect of uncertainty on investment: Some stylized facts, *Journal of Money, Credit and Banking* 28, 64–83.

Leary, Mark T., and Michael R. Roberts, 2014, Do peer firms affect corporate financial policy?, *Journal of Finance* 69, 139–178.

Ludvigson, Sydney C., Sai Ma, and Serena Ng, 2021, Uncertainty and business cycles: exogenous impulse or endogenous response?, *American Economic Journal: Macroeconomics* 13, 369–410.

Lyandres, Evgeny, and Alexei Zhdanov, 2010, Accelerated investment effect of risky debt, *Journal of Banking & Finance* 34, 2587–2599.

Majd, Saman, and Robert S. Pindyck, 1987, Time to build, option value, and investment decisions, *Journal of Financial Economics* 18, 7–27.

Martin, Philippe, and Carol Ann Rogers, 2000, Long-term growth and short-term economic instability, *European Economic Review* 44, 359–381.

McDonald, Robert, and Daniel Siegel, 1986, The value of waiting to invest, *Quarterly Journal of Economics* 101, 707–727.

Nakamura, Emi, Dmitriy Sergeyev, and Jon Steinsson, 2017, Growth-rate and uncertainty shocks in consumption: Cross-country evidence, *American Economic Journal: Macroeconomics* 9, 1–39.

Oi, Walter Y, 1961, The desirability of price instability under perfect competition, *Econometrica* 29, 58–64.

Okou, Cedric, Mohammad Jahan-Parvar, and Bruno Feunou, 2018, Downside variance risk premium, *Journal of Financial Econometrics* 16, 341–383.

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

Pindyck, Robert S., 1993, A note on competitive investment under uncertainty, *American Economic Review* 83, 273–277.

Ramey, Gary, and Valery Ramey, 1995, Cross-country evidence on the link between volatility and growth, *American Economic Review* 85, 1138–1151.

Segal, Gill, 2019, A tale of two volatilities: Sectoral uncertainty, growth, and asset prices, *Journal of Financial Economics* 134, 110–140.

Segal, Gill, and Ivan Shaliastovich, 2025, Uncertainty, risk, and capital growth, *Review of Financial Studies* Forthcoming.

Segal, Gill, Ivan Shaliastovich, and Amir Yaron, 2015, Good and bad uncertainty: Macroeconomic and financial market implications, *Journal of Financial Economics* 117, 369–397.

Stein, Luke CD, and Elizabeth Stone, 2013, The effect of uncertainty on investment: Evidence from equity options, Working paper, Babson College.

Tauchen, George, 1986, Finite state Markov-chain approximations to univariate and vector autoregressions, *Economics Letters* 20, 177–181.