The Origin of Risk*

Alexandr Kopytov
University of Rochester

Mathieu Taschereau-Dumouchel
Cornell University

Zebang Xu
Cornell University

July 28, 2025

Abstract

We propose a model in which risk, at both the micro and macro levels, is endogenous and driven by incentives. In the model, each firm chooses the mean and variance of its productivity process, as well as how it covaries with the productivity of other firms. Aggregate risk arises when firms select productivity processes that are correlated with one another. The theory predicts that firms with larger sales and lower markups are less volatile and less correlated with aggregate productivity. We find support for these predictions in the data. Through their impact on risk-taking decisions, distortions such as taxes and markups can make GDP more volatile in equilibrium. In a calibrated version of the model, removing distortions significantly reduces GDP volatility.

JEL Classifications: E32, D81, C67, D57

*We thank Robert Barro, Stefanie Stantcheva, and three anonymous referees for helpful suggestions. We are also grateful to Julieta Caunedo, Foti Grigoris, Philipp Kircher, Narayana Kocherlakota, Christian vom Lehn, Jeremy Lise, Ezra Oberfield, Giulio Trigilia, Chen Yeh, and participants at various seminars and conferences for useful comments.

1 Introduction

Risk influences a wide range of economic activities. At the micro level, it shapes household and firm decisions, while at the macro level it drives aggregate fluctuations and influences decisions that are important for growth. Policymakers often worry about large-scale risks such as wars and climate change. Given its importance, understanding the determinants of risk is crucial. Yet most macroeconomic models treat risk as exogenous, assuming that key objects like productivity follow predefined stochastic processes. In contrast, this paper explores the idea that risk is endogenous, and that its quantity and properties are driven, like most things in economics, by incentives. Under this view, risk depends on economic conditions, and policymakers can exert some control over it.

That economic agents can influence the risk that they face seems natural. Firms, for instance, affect how much risk they are exposed to by deciding whom to hire, how to organize production, what projects to undertake, where to locate a plant, which markets to enter, etc. For instance, a US firm might relocate part of its production to China to reduce costs, but doing so increases its exposure to geopolitical risks and trade wars. Growing crops closer to the shore might provide reliable irrigation, but it also increases vulnerability to flooding. When aggregated, these individual exposure decisions shape the risk profile of the entire economy. If many firms choose to locate by the shore, for instance, crop yields become correlated, making flooding an aggregate risk factor.

Some stylized features of the data support the idea that risk is, at least in part, endogenous. At the aggregate level, the volatility of a country's GDP is related to its income level, its political system, the quality of its institutions, etc. At the firm level, the volatility of productivity and its correlation with aggregate GDP vary systematically with firm size and markup levels. These patterns suggest that productivity risk is to some degree shaped by the broader economic environment.[1]

To explore the origin of risk and its implications, we develop a parsimonious general equilibrium endogenous risk model. We focus on the decisions made by firms and, given its importance for macroeconomic outcomes, on the productivity risk that they face. Instead of modeling every single decision affecting productivity risk, we adopt a holistic approach and assume that firms can select their productivity process directly. Specifically, we let them choose the mean and the variance of their productivity, as well as how it correlates with that of other firms. We make this choice operational in the model by assuming that there are underlying sources of risk, and that firms can adjust their exposure to those fundamental risk factors.

In the model, managing a firm's risk exposure can reduce its productivity. For example, a firm might adopt rules and procedures to guard against accidents, cyberattacks, or data loss, but the

[1] See Ramey and Ramey (1995), Alesina et al. (1996), Acemoglu and Zilibotti (1997), and Koren and Tenreyro (2007; 2013) for work on the link between GDP volatility and country characteristics. In Appendix B.5, we show cross-country evidence that aggregate TFP volatility decreases with per-capita income, while it increases with taxes and markups. In Section 8, we provide firm-level evidence that the productivity of larger firms and those with lower markups is less volatile and less correlated with aggregate TFP.

constraints that these measures impose on normal workplace operations can lower productivity. Risk management might also require resources: for instance, avoiding droughts by planting crops near the shore may involve renting an expensive piece of land, while navigating a rapidly changing market may require hiring skilled managers. In the model, these resources are supplied by a representative, risk-averse household at a cost in terms of utility. When making risk-exposure decisions, firms balance productivity and resource costs against the benefit of risk exposure. Consequently, the presence of risk in this economy is a choice. Eliminating aggregate risk entirely is possible, but given the high cost that this would entail, agents generally prefer to tolerate some amount of risk.

We assume that markets are complete, and firms therefore use the household's stochastic discount factor when comparing cash flows in different states of the world. As a result, firms are concerned with how their risk exposure correlates with the aggregate economy. During downturns, the household's marginal utility of consumption is high, making cash flows in those states particularly valuable. Firms therefore aim to lower their exposure to pro-cyclical risk factors. Since consumption risk is influenced by the aggregate of firms' exposure decisions, firms diversify by seeking exposure to risk factors that are less correlated with those affecting other firms.

One of our objectives is to evaluate the impact of firm-level risk-exposure decisions on aggregate volatility. To properly capture how micro shocks translate into macro fluctuations, we embed the firms into an otherwise standard production network economy. The input-output structure implies that the risk decisions of one firm affect its neighbors through supply chain linkages. The model also features exogenous wedges, potentially from taxes or markups, that create a gap between the price at which goods are sold and their production cost. These wedges are the only source of inefficiency in our economy. By varying those wedges, we can examine how taxes and markups affect risk taking decisions and, through that channel, aggregate risk.

We show that there exists a unique equilibrium in this economy, and that this equilibrium can be characterized as the solution to a distorted planning problem. This problem implies that equilibrium risk-exposure decisions seek to increase expected GDP, reduce the variance of GDP and lower the cost of managing risk. Because of the wedges, risk-exposure decisions are generally inefficient, and the economy tends to be overexposed to harmful risk factors. We also establish conditions under which this overexposure decreases expected GDP and increases aggregate volatility.

Our theory predicts how risk-taking behavior varies with firm characteristics. In a firm's production function, productivity multiplies the input bundle. As a result, the marginal benefit of managing productivity risk naturally increases with the size of the firm. In contrast, the cost of risk management, which includes the productivity and resource components, might not scale as fast with firm size. Larger firms therefore have stronger incentives to manage risk, and they tend to be less exposed to harmful risks in equilibrium. As wedges reduce firm sizes, more distorted firms also tend to be more exposed to harmful risks. We use detailed firm-level data that covers the near-universe of Spanish firms to evaluate these predictions. Consistent with the theory, we find that the variance

of a firm's TFP and its covariance with GDP decrease with its size and increase with its markup.

Well-established facts about stock returns are also consistent with our model. As documented by Fama and French (1992), larger firms tend to comove less with the aggregate stock market. The mechanisms of the model offer an explanation for this stylized fact, provided that increases in productivity benefit a firm's stock price. Additionally, we find that firms with higher markups exhibit greater covariance with the stock market, in line with our theoretical predictions.

To evaluate the quantitative implications of endogenous risk for the macroeconomy, we provide a basic calibration of the model to the Spanish economy. Our data covers most revenue-generating firms in Spain, and we replicate each of them in the calibrated model. We pick key parameters to precisely match data features relevant to the model's mechanisms, including each firm's estimated markup, the volatility of its productivity, and the covariance of its productivity with GDP.

One key prediction of our model is that wedges tend to make GDP more volatile through their impact on risk-taking decisions. To evaluate the importance of this mechanism, we conduct an experiment in which we remove all wedges from the calibrated model. Without wedges, firms find it valuable to manage risk more aggressively, and the standard deviation of GDP falls from 2.3% to 1.8%. This suggests that the quantitative impact of distortions such as taxes and markups on aggregate volatility can be substantial. Without endogenous risk, the removal of wedges would have no impact on the volatility of GDP.

We also conduct a second experiment that mimics an uncertainty shock by doubling the volatility of the economy's underlying risk factor. If firms cannot adjust their risk-taking decisions, this increases GDP volatility by 53 basis points. In contrast, when firms are free to manage their risk, they sharply reduce their exposure in response to the higher fundamental risk, and the volatility of GDP increases by only 12 basis points. This last finding suggests that frictions and policies that impede firms' ability to manage risk might have sizable adverse effects on aggregate volatility.

We test the robustness of these results to alternative wedge measures and in an augmented model with disaster risk (Rietz, 1988; Barro, 2006). The latter model is able to match key asset pricing moments, such as the equity premium. We find that removing wedges significantly reduces the equity premium, as firms find it more valuable to manage risks more aggressively.

Literature review

The main contribution of this paper is to propose a theory of aggregate fluctuations driven by endogenous risk exposure. In contrast, standard representative agent models of the business cycles such as Kydland and Prescott (1982) assume that aggregate TFP follows an exogenous process. In models with individual firms, firm risk is generally exogenous but aggregate TFP risk can be endogenously driven by the decisions of the firms (Khan and Thomas, 2008; Clementi and Palazzo, 2016; Bloom et al., 2018, among many others). In contrast, our paper builds on the idea that firms

can choose their own productivity process.

The idea that firms have some control over their productivity has a long tradition in the growth literature. In endogenous (Romer, 1990; Grossman and Helpman, 1991; Aghion and Howitt, 1992) and semi-endogenous (Jones, 1995) growth models, firms invest in R&D to increase their productivity. A similar process is at work in models with firm dynamics in the spirit of Klette and Kortum (2004). In our model, each firm can also select the mean of its productivity process but, in contrast to those works, it can also pick its variance and its covariance structure. Aggregate fluctuations arise as a consequence.

Closer to our work are papers in the development and growth literature such as Greenwood and Jovanovic (1990), Acemoglu and Zilibotti (1997) and Cole, Greenwood, and Sanchez (2016) in which agents can choose between low-risk, low-return projects and high-risk, high return projects. When risk is undiversifiable, agents favor safer projects. As markets become more complete, however, riskier projects are preferred, leading to faster economic growth. In contrast, markets are always complete in our model, and we focus on business-cycle risk rather than long-run growth. We also investigate how distortions such as markups and taxes influence aggregate risk-taking behavior.

In the corporate finance literature, managers often select between projects with different risk levels, but the focus of these papers is usually on the agency problem between firm owners and managers (Jensen and Meckling, 1976; Ross, 1977). One distinguishing feature of our work is that firms also select the correlation of their productivity with that of other firms.

Our model shares features with network economies such as Long and Plosser (1983) and Acemoglu et al. (2012). Some of the propagation mechanisms at work in these models are also active in our paper. We also relate to the literature on wedges in production network economies, which includes Jones (2011), Baqaee and Farhi (2019), Liu (2019) and Bigio and La'O (2020). In standard Cobb-Douglas network economies with wedges, aggregate volatility depends only on the variance of shocks and on the vector of cost-based Domar weights. Both of these objects are independent of wedges. In contrast, wedges in our setup influence aggregate risk through their impact on exposure decisions. We are also related to Pellet and Tahbaz-Salehi (2023) and Kopytov et al. (2024) who study (exogenous) uncertainty in network economies.

The rest of the paper is organized as follows. We introduce our model of endogenous risk in Section 2 and derive some of its aggregate properties in Section 3. We characterize the firm's equilibrium risk-taking decisions in Section 4 and discuss the existence, uniqueness and efficiency of the equilibrium in Section 5. We then explore how primitives of the model shape aggregate risk (Section 6) and GDP (Section 7). In Section 8, we provide firm-level evidence that supports the predictions of the model. In Section 9, we calibrate the model and evaluate its quantitative implications. Section 10 concludes.

2 A model of endogenous risk

We study the origin of risk in an otherwise standard production network economy. The economy is populated by a set of firms, each producing a differentiated good that can be used as intermediate input or for final consumption. Firms can choose the distribution of their productivity, and economic conditions influence those choices. Collectively, these micro-level decisions shape the aggregate risk profile of the economy.

2.1 Firms and production functions

There are N firms, indexed by $i \in \{1, \ldots, N\}$, each producing a differentiated good.[2] Firm i has access to the constant returns to scale Cobb-Douglas production function

$$F\left(\delta_i, L_i, X_i\right) = e^{a_i(\varepsilon,\delta_i)} \zeta_i L_i^{1-\sum_{j=1}^{N} \alpha_{ij}} \prod_{j=1}^{N} X_{ij}^{\alpha_{ij}}, \tag{1}$$

where L_i is labor, $X_i = (X_{i1}, \ldots, X_{iN})$ is a vector of intermediate inputs, and ζ_i is a normalization term.[3] The matrix α contains the input shares $0 < \alpha_{ij} < 1$ and therefore describes the production network in this economy.[4]

2.2 Firm productivity

Firms can choose the probability distribution of their total factor productivity $a_i\left(\varepsilon, \delta_i\right)$. To make this idea operational, we assume that $a_i\left(\varepsilon, \delta_i\right)$ depends on several sources of fundamental risk, collected in the $M \times 1$ vector ε, and on firm i's exposure to those risks, captured by the $M \times 1$ vector δ_i. Specifically, we let

$$a_i\left(\varepsilon, \delta_i\right) = \delta_i^{\top} \varepsilon - b_i\left(\delta_i\right), \tag{2}$$

such that δ_{im} determines i's exposure to risk factor m. We do not restrict δ_i to be positive, so that firms can be negatively exposed to a factor. The vector ε follows a normal distribution, $\varepsilon \sim \mathcal{N}\left(\mu, \Sigma\right)$, where the vector μ represents the expected level of the risk factors, and the positive definite matrix Σ determines both the uncertainty about individual elements of ε and their correlations. In addition to the random component $\delta_i^{\top} \varepsilon$, firm i's productivity $a_i\left(\varepsilon, \delta_i\right)$ depends on a productivity cost of risk exposure $b_i\left(\delta_i\right)$, which we discuss in more detail below.

By choosing δ_i, firm i implicitly selects the mean, the variance, and the covariance of its pro-

[2] Equivalently, we can think about a continuum of identical firms producing good i.

[3] We set $\zeta_i^{-1} = \left(1 - \sum_{j=1}^{N} \alpha_{ij}\right)^{1-\sum_{j=1}^{N} \alpha_{ij}} \prod_{j=1}^{N} \alpha_{ij}^{\alpha_{ij}}$ to simplify the unit cost expression (10) below.

[4] It is natural to think that the risk-taking decisions of a firm affect its (perhaps removed) suppliers and customers and, through that channel, aggregate quantities. We include input-output linkages in the model to study that interaction, but they are not essential for some of the mechanisms of the model.

ductivity,

$$\mathrm{E}\left[a_i\right] = \delta_i^\top \mu - b_i\left(\delta_i\right), \qquad \mathrm{V}\left[a_i\right] = \delta_i^\top \Sigma \delta_i, \qquad \text{and} \qquad \mathrm{Cov}\left[a_i, a_j\right] = \delta_i^\top \Sigma \delta_j. \tag{3}$$

It follows that exposure to more volatile or correlated factors leads to higher TFP volatility. The choice of δ_i also affects how i's productivity correlates with that of other firms. As we will see, these correlations create aggregate risk, so that firm decisions matter for the risk profile of the overall economy. For example, if many firms choose to operate in the same geographic region, a localized earthquake could lead to a large drop in GDP.

Our specification of the fundamental sources of risk ε is purposefully abstract, and we do not take a stance on what exactly a specific risk factor ε_m is. In reality, firms make a large number of decisions that affect their risk profile: where to locate a plant, whom to hire, what project to develop, where to get financing, and many others. Each of these decisions involves particular tradeoffs and including them all in the model would make it intractable. Instead, we adopt a holistic approach and focus on *how much* risk firms take on, and on how *correlated* that risk is across firms. We introduce the risk factor structure as a way to tractably model these choices.[5,6]

2.3 Costs of risk exposure

Aggressively managing risk might lead to a loss in productivity. For instance, a firm might reduce its risk by putting in place rules to avoid mistakes and accidents. But those rules might also constrain workers, which would lower their productivity. We capture these costs through the term $b_i\left(\delta_i\right)$ in (2), and assume that it takes the quadratic form

$$b_i\left(\delta_i\right) = \frac{1}{2}\left(\delta_i - \delta_i^\circ\right)^\top B_i \left(\delta_i - \delta_i^\circ\right), \tag{4}$$

where B_i is an $M \times M$ positive definite matrix. As firm i can achieve risk exposure $\delta_i = \delta_i^\circ$ without any costs, we refer to δ_i° as the *natural risk exposure* of firm i.

In addition to this productivity cost, managing risk might also require the use of resources. For instance, renting a piece of land away from the shore might be needed to reduce flood risk or vulnerability to climate change. Hiring an experienced lobbyist might alleviate political risk. Employing research scientists might make the firm better equipped to face a disruptive new technology, etc. In the model, these resources are provided by the representative household at a cost in terms of utility,

[5] Ideally, we would let firms pick their productivity distributions in an arbitrary way, but modeling correlations across firms creates some issues. For instance, should firm i or firm j decide how their two productivity processes correlate? We use the underlying risk factors as a tractable modeling device to sidestep these problems while allowing firms to influence their correlation structure.

[6] While we keep the risk factors abstract for most of the paper, we provide an empirical exercise in Appendix D.1 in which ε corresponds to climate change risk. We show there that the response of Spanish firms to high temperature shocks is in line with our theoretical predictions.

and they are available to the firm at a price W_R.[7] To achieve risk exposure δ_i, we assume that firm i requires $g_i(\delta_i)$ units of risk management resources. We further let $g_i(\delta_i)$ take the form

$$g_i(\delta_i) = \frac{1}{2}(\delta_i - \delta_i^\circ)^\top G_i(\delta_i - \delta_i^\circ), \tag{5}$$

where G_i is an $M \times M$ positive-definite matrix. As for the productivity cost b_i, setting $\delta_i = \delta_i^\circ$ can be done without using any resources.[8]

The cost structure implied by b_i and g_i is sufficiently rich to allow for both aggregate and firm-specific sources of risk.[9] Likewise, the model can accommodate risks that cannot be mitigated by assuming that b or g are sufficiently steep. A constant firm-specific TFP shifter can also be introduced by assuming that there is a degenerate risk factor m with $\Sigma_{mm} = 0$ and that deviating from a fixed exposure to m is arbitrarily costly. Finally, note that firms can choose $\delta_i = 0$ and eliminate risk completely, although this choice might be costly. As such, the presence of risk in this economy is a choice made by economic agents responding to incentives.

Since productivity a_i multiplies the input bundle, the benefit of risk management increases with firm size. On the cost side, the two functions b_i and g_i affect firms of different sizes differently. Since b_i is a productivity cost, its influence also scales up with the size of the input bundle and, as we will see, its impact on risk decisions is size invariant. In contrast, g_i does not scale with the size of the firm, and it therefore provides larger firms with a risk management advantage. There is plenty of evidence for such an advantage in the data. For instance, larger firms are more likely to 1) have Chief Risk Officers (Liebenberg and Hoyt, 2003), 2) implement holistic Enterprise-wide Risk Management (ERM) systems (Beasley, Clune, and Hermanson, 2005; Hoyt and Liebenberg, 2011), 3) and hedge risks using derivatives (Géczy, Minton, and Schrand, 1997; Bodnar, Hayt, and Marston, 1998). We will also see in Section 8 that the size advantage provided by g_i can explain how risk varies with firm size in the data. When combined, b_i and g_i allow us to flexibly capture the link between firm size and risk-taking decisions.

Events in this economy unfold in two stages. Before ε is realized, firms choose how much risk management resources to purchase and pick δ. The household also decides how many risk management resources to supply at that time. After ε is drawn, consumption, labor and intermediate inputs

[7] We show in Appendix F.2 that the model can be written in terms of different resources, each with their own supply elasticity. We show how these different resources can be aggregated into a single composite resource. For simplicity, we therefore write the model directly in terms of that unique composite.

[8] We do not take a stance on the underlying adjustment margins that give rise to the cost functions b_i and g_i. In Appendix F.4, we propose different possible microfoundations that involve 1) researching a new technology, 2) a location problem, 3) workplace rules, and 4) hiring managers with different skills. In reality, many such microfoundations probably interact to give rise to b_i and g_i. Assuming quadratic costs allows us to explore how endogenous risk interacts with the macroeconomy in a tractable manner. As we show in Appendix F.3, most of our results also hold in some form without assuming that b_i and g_i are quadratic, but the exploration of the model is much more complicated in that case and does not yield important additional insights.

[9] Indeed, there can be a risk factor ε_m for which deviating from $\delta_{jm}^\circ = 0$ is extremely costly for any firm $j \neq i$. In this case, ε_m is effectively specific to firm i.

are chosen and their respective markets clear. This timing reflects the fact that risk management decisions (where to locate a plant, which technology to adopt, what workplace rules to impose, etc.) are long-lived and cannot be changed rapidly once shocks are realized.

2.4 Household preferences

A risk-averse representative household owns the firms and supplies one unit of labor inelastically. It also supplies a variable amount R of risk management resources.[10] The household values consumption of a bundle $Y = \prod_{i=1}^{N} \left(\beta_i^{-1} C_i\right)^{\beta_i}$ of the different goods. We impose $\sum_{i=1}^{N} \beta_i = 1$, so that $\beta_i \geq 0$ corresponds to good i's share of consumption expenditure in equilibrium. Since Y corresponds to aggregate value added in this economy, we refer to it as GDP.

The household has King, Plosser, and Rebelo (1988) preferences that are given by[11]

$$\mathcal{U}(Y)\,\mathcal{V}(R), \tag{6}$$

where $\mathcal{U}(Y) = Y^{1-\rho}/(1-\rho)$ is CRRA with relative risk aversion parameter $\rho \geq 1$, and $\mathcal{V}(R) = V(R)^{1-\rho}$ with $V(R) = \exp(-\eta R)$.[12] The elasticity $\eta > 0$ controls how costly it is for the household to supply risk management resources. The exponential form implies that the price of these resources is a constant fraction of nominal GDP. We adopt it for tractability but show in Appendix F.5 that several of our results extend to general convex V functions at the cost of extra complications.

The household decides on how much risk management resources to supply before uncertainty is realized. In contrast, it makes consumption decisions after ε is drawn. It follows that in each state of the world the household's budget constraint is

$$\sum_{i=1}^{N} P_i C_i \leq W_L + W_R R + \Pi, \tag{7}$$

where P_i is the price of good i, Π is the profit of the firms, W_R is the price of risk management resources, and W_L is the wage. We take the wage as numeraire so $W_L = 1$.[13] We further define $\overline{P} = \prod_{i=1}^{N} P_i^{\beta_i}$ as the consumption price index.

[10] We keep the supply of workers fixed to simplify the exposition. It is straightforward to make it elastic instead.

[11] The balanced-growth utility function of King, Plosser, and Rebelo (1988) implies that the substitution and the income effects associated with the price of risk management resources cancel each other and allows us to derive closed-form expressions for aggregate quantities. We show in Appendix F.5 how to extend our analysis to more general preferences.

[12] When $\log Y$ is normally distributed, maximizing $\mathrm{E}\left[(1-\rho)^{-1} Y^{1-\rho}\right]$ amounts to maximizing $\mathrm{E}[\log Y] - \frac{1}{2}(\rho - 1)\,\mathrm{V}[\log Y]$ such that $\rho \lessgtr 1$ indicates whether the household likes uncertainty in log consumption or not. This is a consequence of the usual increase in the mean of a log-normal variable from an increase in the variance of the underlying normal variable. The presence of ρ in $\mathcal{V}$ is an innocuous normalization to simplify some expressions.

[13] Since the labor market clears in each state of the world, there should technically be a different wage in each state ε. But given the specific structure of the economy, we show in Appendix F.6 that we can normalize $W_L = 1$ in each state of the world.

Importantly, markets are complete in this economy, meaning that there is a complete set of state-ε contingent claims. Since the household can trade these claims, their prices reflect the marginal utility of consumption in each state of the world. As usual, this implies that firms use the household's stochastic discount factor,

$$\Lambda = \frac{1}{\overline{P}} \frac{d\left[\mathcal{U}\left(Y\right)\mathcal{V}\left(R\right)\right]}{dY}, \tag{8}$$

to discount profits (see Appendix A for a derivation).

The problem of the household also provides a relation between the utility of supplying risk management resources R and their price W_R, given by

$$-\bar{P}Y\frac{V'\left(R\right)}{V\left(R\right)} = W_R. \tag{9}$$

Since $V\left(R\right)$ is exponential, this equation simplifies to $\eta\bar{P}Y = W_R$, so that the price of risk management resources is a fraction η of nominal GDP. A higher η therefore implies that risk management is more expensive.

2.5 Firm problem

We solve the problem of a firm i in two stages: 1) before ε is realized, the firm chooses its risk exposure δ_i, and 2) after ε is realized, it decides on how much labor to hire and on how many intermediate inputs to purchase, given its δ_i. In that second stage, ε is known and δ_i is fixed, and firm i's cost minimization problem implies a standard Cobb-Douglas unit cost of production

$$K_i\left(\delta_i, P\right) = \frac{1}{e^{a_i(\varepsilon,\delta_i)}}\prod_{j=1}^{N} P_j^{\alpha_{ij}}, \tag{10}$$

where $P = (P_1, \ldots, P_n)$ is the vector of prices. The cost of producing one unit of good i is equal to the geometric average of the individual input prices, weighted by their respective shares, and adjusted for total factor productivity. Total factor productivity, in turn, depends on the firm's risk-exposure decision δ_i.

Given (10), the first stage of the firm's problem involves choosing risk exposure δ_i to maximize expected discounted profits

$$\delta_i^* \in \arg\max_{\delta_i} \mathrm{E}\left[\Lambda\left[P_iQ_i - K_i\left(\delta_i, P\right)Q_i - g_i\left(\delta_i\right)W_R\right]\right], \tag{11}$$

where Q_i is the equilibrium demand for good i. The terms in the inner square bracket are sales, cost of goods sold, and risk management expenditure, respectively. They are multiplied by the stochastic discount factor Λ, which reflects the value of cash flows in different states of the world from the perspective of the representative household. As a result, firms are influenced by the representative

household's attitude toward risk.

2.6 Equilibrium conditions

Firm i sets its price at an exogenous wedge $\tau_i \geq 0$ above its unit cost K_i, such that

$$P_i = (1 + \tau_i) K_i (\delta_i, P) \text{ for all } i \in \{1, \ldots, N\}. \tag{12}$$

These wedges can be interpreted as markups, taxes, or other distortions. When $\tau_i = 0$ for all i, prices are fully competitive. For a set of risk-exposure decisions, equations (10) and (12) allow us to fully characterize the price system as a function of the productivity shocks ε.[14]

An equilibrium is defined by the optimality conditions of the household and the firms holding simultaneously, and all markets clearing.

Definition 1. An *equilibrium* is a choice of risk-exposure decisions $\delta^* = (\delta_1^*, \ldots, \delta_n^*)$ and a stochastic tuple $(P^*, W_R^*, C^*, L^*, R^*, X^*, Q^*)$ such that

1. (Optimal risk-exposure choice) For each $i \in \{1, \ldots, N\}$, the risk-exposure decision δ_i^* solves (11) given prices P^*, the price for risk management resources W_R^*, demand Q_i^*, and the stochastic discount factor Λ^* given by (8).

2. (Optimal input choice) For each $i \in \{1, \ldots, N\}$, factor demands per unit of output L_i^*/Q_i^* and X_i^*/Q_i^* are a solution to the cost minimization problem, $K_i (\delta_i, P) = \min_{L_i, X_i} \left(L_i + \sum_{j=1}^N P_j X_{ij} \right)$ subject to $F (\delta_i, L_i, X_i) \geq 1$, given prices P^* and W_R^*, and the chosen risk exposure δ_i^*.

3. (Consumer maximization) The consumption vector C^* and the supply of risk management resources R^* maximize expected utility (6) subject to (7) given prices P^* and W_R^*.

4. (Unit cost pricing) For each $i \in \{1, \ldots, N\}$, P_i^* solves (12) where $K_i (\delta_i^*, P^*)$ is given by (10).

5. (Market clearing) For each $i \in \{1, \ldots, N\}$,

$$C_i^* + \sum_{j=1}^N X_{ji}^* = Q_i^* = F_i (\delta_i^*, L_i^*, X_i^*), \quad \sum_{i=1}^N L_i^* = 1, \text{ and } \sum_{i=1}^N g_i (\delta_i^*) = R^*. \tag{13}$$

Conditions 2 to 5 are standard and imply that firms and the household optimize and that all markets clear at equilibrium prices. Condition 1 emphasizes that risk-taking decisions are equilibrium objects that depend on the primitives of the economy. Together, these decisions shape the aggregate risk structure of the economy.

[14] Due to constant returns to scale, risk management expenses lead to negative profits for τ near zero. To keep firms in operation, lump-sum transfers from the government can be introduced without affecting the forces of the model.

We have kept the model parsimonious to keep the exposition transparent, but it is straightforward to extend it in several ways. For instance, many of our results still hold when we relax the assumption that b_i and g_i are quadratic. We can similarly relax the functional form assumption on the disutility of supplying risk management resources $\mathcal{V}$. We work out these extensions in Appendix F. At the same time, some assumptions are required to keep the model tractable. Without the Cobb-Douglas structure, for instance, we cannot write a closed-form *distribution* for the price vector P under a given δ, which makes solving the model analytically infeasible. Departing from marginal product pricing would also introduce important complications.

Some features of our setup resemble the classical CCAPM model with normally distributed shocks (Lucas Jr, 1978). While it provides great tractability, it is well known that this model struggles to match standard asset pricing moments, such as the equity premium, implying that agents in our economy may not exhibit a realistic attitude toward risk (Mehra and Prescott, 1985). To address this limitation, Appendix F.7 extends the model to incorporate disaster risk (Rietz, 1988; Barro, 2006). We also explore the implications of disasters more thoroughly when we calibrate the model in Section 9.

3 Domar weights, equilibrium prices and GDP

When making risk decisions, firms take into account the equilibrium distributions of prices, demand and real GDP, the latter being relevant for the stochastic discount factor. In this section, we characterize these objects.

3.1 Domar weights and prices

Following Baqaee and Farhi (2019), we define the revenue-based input-output matrix Ω and the cost-based input-output matrix $\tilde{\Omega}$ as

$$\Omega_{ij} := \frac{P_j X_{ij}}{P_i Q_i} \qquad \text{and} \qquad \tilde{\Omega}_{ij} := \frac{P_j X_{ij}}{W_L L_i + \sum_{k=1}^{N} P_k X_{ik}}.$$

The elements Ω_{ij} and $\tilde{\Omega}_{ij}$ correspond to firm i's expenditure on good j as a share of its total sales and as a share of its input cost gross of risk management expenses. Since firms have Cobb-Douglas production functions, these shares are constant, with $\tilde{\Omega}_{ij} = \alpha_{ij}$ and $\Omega_{ij} = \alpha_{ij}/(1+\tau_i)$. This last equation follows since τ_i captures the gap between a firm's revenue and its cost.

Using Ω and $\tilde{\Omega}$, we also define the revenue-based and cost-based Leontief inverse matrices as

$$\mathcal{L} := (I-\Omega)^{-1} = I + \Omega + \Omega^2 + \dots \qquad \text{and} \qquad \tilde{\mathcal{L}} := \left(I - \tilde{\Omega}\right)^{-1} = I + \tilde{\Omega} + \tilde{\Omega}^2 + \dots. \tag{14}$$

While Ω_{ij} and $\tilde{\Omega}_{ij}$ capture the direct exposure of firm i to firm j as a share of revenues and costs

respectively, $\mathcal{L}_{ij}$ and $\tilde{\mathcal{L}}_{ij}$ capture both direct and indirect exposures through network linkages.

Finally, we define firm i's revenue-based Domar weight ω_i as the ratio of its sales to nominal GDP, such that $\omega_i := P_i Q_i / \left(\bar{P} Y\right)$. The market clearing condition (13) and the first-order conditions of the household imply that $\omega^\top = (\omega_1, \ldots, \omega_N)^\top = \beta^\top \mathcal{L}$. We define the cost-based Domar weight vector $\tilde{\omega}$ in a similar fashion: $\tilde{\omega}^\top := \beta^\top \tilde{\mathcal{L}} = \beta^\top (I - \alpha)^{-1}$. Both ω_i and $\tilde{\omega}_i$ reflect the importance of firm i as a supplier, either in terms of the revenues or of the costs of its customers. Intuitively, the Domar weight of a firm i is high if the household values its good (β_i large), or if another firm that the household values relies on i (large α_{ji} and β_j). Higher order connections also matter as indicated by the geometric expansions in (14). Without input-output linkages ($\alpha = 0$), both Domar weight vectors are equal to β.

Using the definition of $\mathcal{L}$ and the fact that $\Omega_{ij} = \alpha_{ij}/(1 + \tau_i)$, we can write the revenue-based Domar weights as

$$\omega^\top = \beta^\top \left(I - \left[\operatorname{diag} (1 + \tau) \right]^{-1} \alpha \right)^{-1}, \tag{15}$$

where $\operatorname{diag}(1 + \tau)$ is the diagonal matrix whose ith diagonal element is $1 + \tau_i$. It is easy to see that ω is weakly decreasing in τ_i. Indeed, an increase in τ_i implies a bigger gap between i's expenditure on inputs and its sales (Ω_{ij} goes down for all j). Since revenue-based Domar weights measure the importance of all suppliers in terms of the revenue of other firms, ω goes down as well. Cost-based Domar weights, in contrast, are unaffected by wedges. Without wedges ($\tau = 0$), $\omega = \tilde{\omega}$. Notably, neither cost-based nor revenue-based Domar weights depend on risk-exposure decisions.

Combining these definitions with (10) and (12), we can write the vector of log prices as

$$p = -\tilde{\mathcal{L}} \left(\delta \varepsilon - b(\delta) - \log(1 + \tau) \right), \tag{16}$$

where $p = (\log P_1, \ldots, \log P_N)$, $\log(1 + \tau) = (\log(1 + \tau_1), \ldots, \log(1 + \tau_N))$, $b(\delta) = (b_1(\delta_1), \ldots, b_N(\delta_N))$, and δ is the $N \times M$ matrix whose typical element δ_{im} is firm i's exposure to risk factor m. It follows that the price p_i of good i is high if firm i's productivity $a_i(\varepsilon_i, \delta_i) = \delta_i^\top \varepsilon - b_i(\delta_i)$ is low, or if the productivity of one of its important suppliers is low, and so on. Equation (16) also makes clear that wedges affect prices in the same way as a decline in productivity. Finally, since ε is normally distributed, so is p.

3.2 Moments of GDP

While δ describes firms' individual risk exposure, it will be convenient to work with an *aggregate* risk-exposure measure, which we define as $\Delta := \delta^\top \tilde{\omega}$. An element of that vector, $\Delta_m = \sum_{i=1}^N \tilde{\omega}_i \delta_{im}$, is simply the (cost-based) Domar-weighted sum of each firm's exposure to factor m.

With that definition in hand, we can derive an expression for real GDP.

Lemma 1. *Log (real) GDP* $y = \log Y$ *is given by*

$$y = \Delta^\top \varepsilon - \tilde{\omega}^\top b(\delta) - \tilde{\omega}^\top \log(1+\tau) - \log \Gamma_L, \tag{17}$$

where the labor share of income Γ_L *is given by*

$$\Gamma_L := \frac{W_L L}{\bar{P} Y} = 1 - \tau^\top (diag(1+\tau))^{-1} \omega. \tag{18}$$

The first term in (17) shows that the contribution of factor ε_m to GDP is proportional to the aggregate exposure to this factor Δ_m. A factor, then, is an important driver of GDP if many firms, or those with high Domar weights, are heavily exposed to it. The additional terms in (17) are deterministic. The second term reflects productivity losses due to managing risk exposures. The last two terms capture the role of wedges. When $\tau = 0$, cost-based and revenue-based Domar weights coincide, such that (17) collapses to $y(\delta) = \omega^\top a(\varepsilon, \delta)$, and Hulten's (1978) theorem applies. As wedges grow, GDP becomes distorted from its efficient level.

Equation (17) implies that log GDP is normally distributed. Its moments are

$$\mathrm{E}[y] = \Delta^\top \mu - \tilde{\omega}^\top b(\delta) - \tilde{\omega}^\top \log(1+\tau) - \log \Gamma_L \qquad \text{and} \qquad \mathrm{V}[y] = \Delta^\top \Sigma \Delta. \tag{19}$$

The first equation shows how risk exposure affects expected log GDP. On the one hand, a marginal increase in aggregate exposure Δ_m raises $\mathrm{E}[y]$ by the expected productivity μ_m of that factor. It follows that, unsurprisingly, gaining exposure to a high-mean risk factor is typically beneficial for $\mathrm{E}[y]$, although it might come at a large TFP cost b. That equation also implies that an increase in μ_m has a beneficial impact on $\mathrm{E}[y]$ if the economy is positively exposed to ε_m ($\Delta_m > 0$), but reduces expected log GDP otherwise.

The second equation in (19) describes the determinants of aggregate risk and is central to our analysis. It shows that an increase in the variance Σ_{mm} of a factor always leads to an increase in the volatility of GDP. The magnitude of that increase is proportional to Δ_m^2, such that its impact is stronger when Δ_m is very positive or very negative. In both cases, the economy is particularly vulnerable to ε_m. Correlations also matter. The impact of Σ_{mn} on $\mathrm{V}[y]$ depends on the sign of $\Delta_m \Delta_n$. If the economy is positively or negatively exposed to *both* risk factors, a higher covariance increases aggregate risk. In contrast, with positive exposure to one factor and negative exposure to the other, an increase in covariance stabilizes the economy and $\mathrm{V}[y]$ declines with Σ_{mn}.

Exposure decisions also affect the variance of GDP. One can show that

$$\frac{d\,\mathrm{V}[y]}{d\Delta_m} = 2\,\mathrm{Cov}[y, \varepsilon_m] = 2 \sum_{n=1}^{M} \Delta_n \,\mathrm{Cov}[\varepsilon_n, \varepsilon_m]. \tag{20}$$

The first equality implies that if ε_m is positively correlated with GDP, an increase in Δ_m adds risk

to the economy and raises $\mathrm{V}[y]$. The second equality follows since the stochastic part of GDP is $\Delta^\top \varepsilon$. It implies that the response of $\mathrm{V}[y]$ to Δ_m depends on how correlated ε_m is with other risk factors, and on the economy's exposure to those factors. If the economy is positively exposed to a factor ε_n ($\Delta_n > 0$), and if ε_m and ε_n are positively correlated, an increase in Δ_m adds on to the risk generated by ε_n, which contributes to a higher $\mathrm{V}[y]$. In contrast, if ε_m and ε_n are negatively correlated, increasing Δ_m offsets some of the fluctuations generated by ε_n.

Equation (19) also shows that wedges have no *direct* impact on $\mathrm{V}[y]$, and so their only possible influence on the variance of log GDP must operate through their impact on risk-taking decisions.

4 Firm decisions

As in any competitive model with constant returns to scale, firms take prices and the demand for their good as given.[15] In our setup, this implies that they seek to minimize a risk-adjusted version of their production costs. We can rewrite the problem of the firm (11) as

$$\delta_i^* \in \arg\min_{\delta_i} \mathrm{E}\left[K_i\left(\delta_i, P\right) Q_i\right] + \mathrm{Cov}\left(K_i\left(\delta_i, P\right) Q_i, \Lambda\right)/\mathrm{E}\left[\Lambda\right] + g_i\left(\delta_i\right) W_R. \tag{21}$$

Naturally, when firm i chooses its risk exposure δ_i, it seeks to minimize the expected cost of goods sold $\mathrm{E}[K_i Q_i]$ and risk management expenses $g_i(\delta_i) W_R$. Risk also plays a role, and the second term in (21) indicates that the firm seeks to minimize the covariance of its production cost $K_i Q_i$ with the stochastic discount factor Λ. The firm prefers to have low costs when the marginal utility of consumption is high or, equivalently, when GDP is low. Extra production is particularly valuable to the household in those states.

It is helpful to expand this covariance term as

$$\mathrm{Cov}\left(K_i Q_i, \Lambda\right) = \mathrm{Corr}\left(K_i Q_i, \Lambda\right) \times \sqrt{\mathrm{V}\left[K_i Q_i\right]} \times \sqrt{\mathrm{V}\left[\Lambda\right]}.$$

The first term in this product is positive for the typical firm. Indeed, as GDP and TFP tend to move hand-in-hand, firm costs move opposite to GDP—and therefore in the same direction as the marginal utility of consumption. Since the last term is also positive, it follows that the optimal risk exposure δ_i typically seeks to minimize the variance of the cost of goods sold, $\mathrm{V}[K_i Q_i]$. This can be done, for instance, by relying less on volatile risk factors, or by diversifying exposure by relying on risk factors that offset each other.[16]

[15] Firms behave atomistically, with no control over prices or the demand that they face. As in Kopytov et al. (2024), this assumption can be microfounded by assuming that each index i corresponds to a sector with many firms, and that buyers and sellers meet through a search process.

[16] The equilibrium productivity of some firms might be countercyclical, such that $\mathrm{Corr}\left(K_i Q_i, \Lambda\right) < 0$. For those firms, maximizing $\mathrm{V}[K_i Q_i]$ would be optimal.

4.1 The marginal value of risk exposure

The first-order condition from (21) provides more insights into the determinants of risk exposure.

Lemma 2. *In equilibrium, the risk exposure decision of firm i solves*[17]

$$\underbrace{\mathcal{E} K_i Q_i}_{\substack{\text{Marginal benefit} \\ \text{of risk exposure}}} = \underbrace{\nabla b_i(\delta_i) K_i Q_i + \nabla g_i(\delta_i) W_R}_{\substack{\text{Marginal cost} \\ \text{of risk exposure}}}, \tag{22}$$

where the marginal value of risk exposure $\mathcal{E}$ is defined as

$$\mathcal{E} := \mathrm{E}[\varepsilon] + \mathrm{Cov}[\lambda, \varepsilon], \tag{23}$$

and where $\lambda := \log \Lambda$ is the log of the stochastic discount factor.

Firm optimization implies that the marginal cost of risk exposure, given by the right-hand side of (22), must be equal to its marginal benefit, $\mathcal{E} K_i Q_i$. This benefit, in turn, depends on the vector $\mathcal{E}$ that captures the marginal benefit of exposure per unit of firm size, and on the size of the firm itself, measured by its cost of goods sold $K_i Q_i$. Intuitively, since productivity a_i multiplies the input bundle, changing its distribution is more consequential for large firms.

The marginal value of risk exposure $\mathcal{E}$ plays an important role in our analysis. Equation (23) shows that, unsurprisingly, risk factors with higher expected productivity $\mathrm{E}[\varepsilon] = \mu$ provide a higher marginal benefit of exposure $\mathcal{E}$. This benefit is also higher for risk factors that deliver high productivity in states of the world with high λ. These factors perform well when the household is poor, and they thus provide insurance in high marginal utility states. We will later see that $\mathcal{E}$ also corresponds to the marginal *welfare* benefit of aggregate exposure to ε. As a result, we say that a risk factor ε_m is *good* if $\mathcal{E}_m > 0$, and that it is *bad* if $\mathcal{E}_m < 0$.

We see from the right-hand side of (22) that the marginal cost of exposure depends on the curvature of the productivity and resource cost functions b_i and g_i. Importantly, these two costs interact with firm size differently. Since b_i is a productivity cost, its importance rises as the firm grows, and so its impact in (22) scales with the size of the firm $K_i Q_i$. In contrast, the resource cost g_i does not, and it therefore matters less for larger firms. When combined, b_i and g_i allow us to flexibly capture the link between firm size and risk-taking decisions.

Because the influence of g_i on risk decisions declines with firm size, larger firms in the model enjoy a risk-management advantage and manage risk more aggressively. Plenty of empirical evidence is consistent with such an advantage: larger firms are more likely to 1) have Chief Risk Officers (Liebenberg and Hoyt, 2003), 2) implement holistic Enterprise-wide Risk Management (ERM) systems (Beasley, Clune, and Hermanson, 2005; Hoyt and Liebenberg, 2011), and 3) hedge risks using

[17] As we discuss below, the ratio $K_i Q_i / W_R$, and therefore the whole expression (22), is deterministic *in equilibrium*.

derivatives (Géczy, Minton, and Schrand, 1997; Bodnar, Hayt, and Marston, 1998). Furthermore, we will see in Section 8 that the size advantage provided by g_i can explain various relationships between risk and firm characteristics in the data.[18] When $G_i = 0$ this size advantage disappears, and firm size $K_i Q_i$ no longer affects risk decisions.

4.2 Individual risk-taking decisions

The right-hand side of (22) shows the separate influence of b_i and g_i on risk-exposure decisions, but it will often be convenient to combine their impact into the *effective* cost of exposure function[19]

$$h_i(\delta_i) := \frac{1}{2}(\delta_i - \delta_i^\circ)^\top H_i(\delta_i - \delta_i^\circ), \text{ where } H_i := B_i + \frac{W_R}{K_i Q_i} G_i. \tag{24}$$

Using this notation, we can rewrite the first-order condition (22) of the firm as

$$\mathcal{E} = \nabla h_i(\delta_i), \tag{25}$$

and we can solve it for the optimal risk-taking decision δ_i as a function of $\mathcal{E}$.

Lemma 3. *The equilibrium risk-taking decision of the firm is*

$$\delta_i = \delta_i^\circ + H_i^{-1}\mathcal{E}. \tag{26}$$

Proof. The result follows immediately from solving (25) for δ_i. □

Equation (26) characterizes individual risk-taking decisions. Whenever risk factors are not neutral ($\mathcal{E} \neq 0$), firms actively manage their risk exposure, which leads to deviations from the natural level δ_i°. The magnitude of these deviations depends on H_i^{-1}, which captures the elasticities of the productivity and resource cost functions b_i and g_i.

Since H_i is positive definite, H_i^{-1} has a positive diagonal so (26) implies that an increase in $\mathcal{E}_m$ always leads to an increase in δ_{im}. Firms are therefore more exposed to attractive risk factors. In contrast, the response of δ_{in} for $n \neq m$ can be positive or negative depending on the sign of $\left[H_i^{-1}\right]_{mn}$. If $\left[H_i^{-1}\right]_{mn} > 0$, we say that risk factors m and n are *local complements*. In this case, an increase in $\mathcal{E}_m$ also leads to an increase in δ_{in}. If, instead, $\left[H_i^{-1}\right]_{mn} < 0$, we say that m and n are *local substitutes*, and a beneficial change to $\mathcal{E}_m$ makes the firm move away from risk factor n.

[18] To our knowledge, there is no comprehensive accounting data on firms' risk management expenditures, but many of these costs would feature in general operating expenses. We show in Appendix B.6 that those expenses do shrink with firm size as a proportion of sales, consistent with the resource cost g_i. In Appendix B.7, we also provide evidence for the productivity cost b_i: firms with lower TFP volatilitiesconsistent with intensive risk management—tend to have lower average TFPs.

[19] Expression (27) below, together with the household optimality condition $\eta \bar{P} Y = W_R$ implies that $W_R / (K_i Q_i) = \eta(1 + \tau_i)/\omega_i$, so that in equilibrium H_i is only a function of parameters.

These substitution patterns can be used to describe physical constraints a firm might face when managing different risks. For instance, suppose that technological constraints push a firm to locate its plants in a single location: either by the shore, where floods can happen, or in a plain, where droughts sometimes occur. The Hessian H_i can be parametrized to capture that situation by imposing that exposure to either flood or drought risk is affordable, but that exposure to both risks is prohibitively costly. In this case, flood and drought risks would be local substitutes.

4.3 Sales, wedges, and risk management

Since firm size K_iQ_i matters for risk-exposure decisions, it is important to understand its determinants. Wedges play an important role here, with an increase in τ_i leading to a decline in K_iQ_i. The intuition is straightforward. Since a firm's unit cost K_i is driven by the price of its inputs, it does not depend on τ_i in partial equilibrium. An increase in τ_i therefore leads directly to an increase in the price $P_i = (1+\tau_i)\,K_i$, and that higher price pushes for a decline in demand Q_i. Since fewer units now have to be produced, the total cost of goods sold K_iQ_i also declines. A firm's revenue-based Domar weight also plays a role. In equilibrium, firm i's cost of goods sold can be written as

$$K_iQ_i = \frac{P_iQ_i}{1+\tau_i} = \frac{\omega_i}{1+\tau_i}\bar{P}Y, \tag{27}$$

where the second equality follows from the definition of revenue-based Domar weights, $\omega_i := P_iQ_i/\left(\bar{P}Y\right)$. Equation (27) makes clear that higher sales ω_i lead to higher production costs K_iQ_i. Demand from the household (β) and the firms (α) influence K_iQ_i through their impact on ω_i.

Equation (27) implies that firms with higher ω_i or lower τ_i have higher costs of goods sold K_iQ_i. Since, by (22), the benefit of managing risk grows one-for-one with K_iQ_i, while its cost grows at a slower pace, those firms are willing to spend more on risk management. Their risk exposure is therefore improved, as the following result shows.

Corollary 1. *Firms with higher revenue-based Domar weights ω_i and lower wedges τ_i manage risk more aggressively, in the sense that*[20]

$$\frac{\partial\left[\mathcal{E}^\top\delta_i\right]}{\partial\omega_i} > 0 \qquad \textit{and} \qquad \frac{\partial\left[\mathcal{E}^\top\delta_i\right]}{\partial\tau_i} < 0.$$

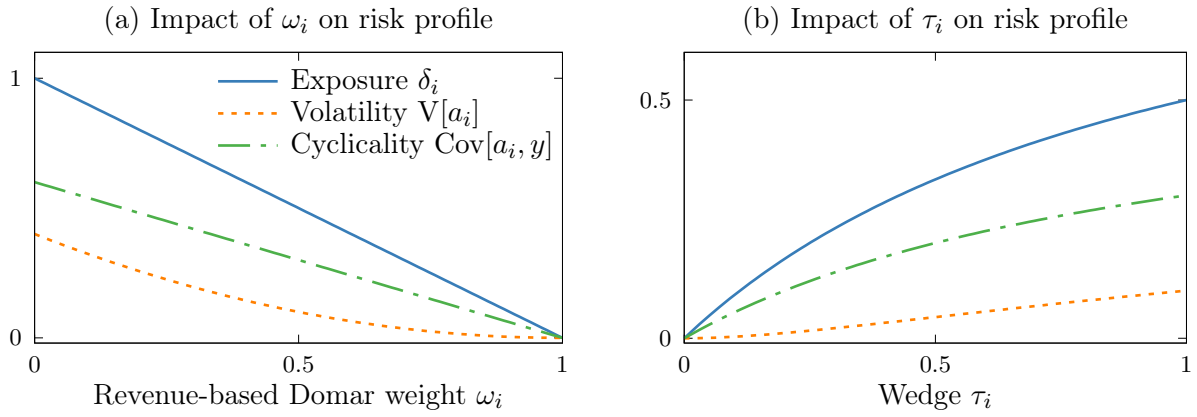
This result implies that firms with larger sales ω_i and smaller wedges τ_i tend to be more exposed to good risk factors ($\mathcal{E} > 0$) and less exposed to bad ($\mathcal{E} < 0$) ones. Since bad risk factors are more correlated with GDP, firms with less exposure to those factors tend to be less procyclical.

Example. Figure 1 illustrates these forces in a simple economy with a single risk factor ε. To fix ideas, we assume that ε is bad ($\mathcal{E} < 0$), and that the aggregate economy is positively exposed to

[20] The partial derivatives imply a partial equilibrium analysis in which $\mathcal{E}$ is kept constant.

it ($\Delta > 0$). A risk factor that drives the business cycle could fit that description. Panels (a) and (b) show the impact of increasing ω_i and τ_i on firm i's risk profile. As ω_i increases or τ_i decreases, the production cost $K_i Q_i$ of the firm rises per (27). The benefit of adjusting the productivity process a_i becomes more important, while the resource cost $g_i(\delta_i)$ associated with that adjustment stays constant. Managing risk more aggressively by reducing δ_i is therefore optimal. Figure 1 also shows the impact of this change on the variance of the firm's productivity a_i, given by (3), and its covariance with GDP, given by $\text{Cov}[a_i, y] = \delta_i^\top \Sigma \Delta$. Both quantities increase with δ_i. It follows that high-ω_i and low-τ_i firms have lower δ_i and hence are less volatile and less correlated with GDP. We will see in Section 8 that these patterns are also visible in firm-level Spanish data.

Figure 1: Impact of the sales share ω_i and wedge τ_i on firm i's risk profile



Notes: Parameters are $\mathcal{E} = -1$, $G_i = W_R = \bar{P}Y = \delta_i^\circ = 1$, $\Sigma = 0.4$, $B_i = 0$ and $\Delta = 1.5$. Panel (a): Responses of δ_i, $\text{V}[a_i]$ and $\text{Cov}[a_i, y]$ to a change in ω_i from 0 to 1. Panel (b): Responses of δ_i, $\text{V}[a_i]$ and $\text{Cov}[a_i, y]$ to a change in τ_i from 0 to 1.

5 Equilibrium existence, uniqueness and efficiency

In the previous section, we characterized firm risk-exposure decisions δ as a function of the marginal value of exposure $\mathcal{E}$. In equilibrium, however, $\mathcal{E}$ itself depends on exposure decisions. In this section, we consider the full equilibrium mapping in which $\mathcal{E}$ is endogenous. We start by introducing the problem of a social planner and by describing the efficient allocation. We then characterize the equilibrium of the model as the outcome of an optimization problem. We also show that there exists a unique equilibrium and that it is in general inefficient.

5.1 A planning problem

We consider the problem of a social planner that maximizes the expected utility of the household subject to the resource constraints (13). To describe that problem, it is convenient to introduce the

aggregate cost function $\bar{h}_{SP}$, defined as the value function

$$\bar{h}_{SP}(\Delta) := \min_{\delta} \tilde{\omega}^{\top} b(\delta) - \log V\left(\sum_{i=1}^{N} g_i(\delta_i)\right), \qquad \text{subject to } \Delta = \delta^{\top}\tilde{\omega}. \tag{28}$$

For a given aggregate risk-exposure vector Δ, $\bar{h}_{SP}(\Delta)$ corresponds to the smallest possible (log) utility cost required to achieve Δ. That cost is the sum of a TFP loss (first term) and a utility cost of supplying risk-management resources (second term). The minimization problem (28) also implicitly defines a function $\delta_{SP}(\Delta)$ that provides the individual risk-exposure matrix δ that minimizes the welfare cost of achieving Δ. We can further decompose $\bar{h}_{SP}$ into its components

$$\bar{h}_{SP}(\Delta) = \underbrace{\tilde{\omega}^{\top} b(\delta_{SP}(\Delta))}_{\bar{b}_{SP}(\Delta)} \underbrace{- \log V\left(\sum_{i=1}^{N} g_i(\delta_{SP,i}(\Delta))\right)}_{\bar{g}_{SP}(\Delta)}, \tag{29}$$

where $\bar{b}_{SP}(\Delta)$ and $\bar{g}_{SP}(\Delta)$ are the optimal aggregate productivity and resource costs of achieving aggregate risk exposure Δ.

With those definitions in hand, the problem of the social planner can be written as

$$\mathcal{W}_{SP} := \max_{\Delta} \underbrace{\Delta^{\top}\mu - \bar{b}_{SP}(\Delta)}_{\mathrm{E}[y_{SP}]} - \frac{1}{2}(\rho - 1)\underbrace{\Delta^{\top}\Sigma\Delta}_{\mathrm{V}[y_{SP}]} - \bar{g}_{SP}(\Delta), \tag{30}$$

where $\mathcal{W}_{SP}$ provides a measure of (log) welfare in the efficient allocation. This problem shows that the efficient aggregate risk exposure seeks to maximize the expected value of log GDP $\mathrm{E}[y_{SP}]$, while minimizing its volatility $\mathrm{V}[y_{SP}] = \Delta^{\top}\Sigma\Delta$ and the risk management utility cost $\bar{g}_{SP}$. Unsurprisingly, the importance of the variance term increases with the risk aversion ρ of the household.[21]

5.2 A distorted planning problem

In the presence of wedges, the equilibrium allocation generally differs from the efficient one. We can, however, write down a *distorted* version of the planning problem (30) whose solution coincides with the equilibrium. This alternative problem relies on a distorted version of $\bar{h}_{SP}$, defined as

$$\bar{h}(\Delta) := \min_{\delta} \tilde{\omega}^{\top} b(\delta) - \log V\left(\sum_{i=1}^{N} \kappa_i g_i(\delta_i)\right), \qquad \text{subject to } \Delta = \delta^{\top}\tilde{\omega}. \tag{31}$$

[21] Since the objective function (30) is strictly concave, there is a unique efficient exposure vector Δ_{SP} that solves the planner's problem (see proof of Proposition 1). There is also a unique efficient matrix $\delta_{SP}(\Delta_{SP})$ that minimizes the utility cost (28) associated with this risk exposure.

We also define $\delta^*(\Delta)$ as the solution to (31), and the functions $\bar{b}(\Delta)$ and $\bar{g}(\Delta)$ as the analogues to their efficient counterparts from (29).[22]

The cost function (31) differs from its efficient counterpart $\bar{h}_{SP}$ solely because of the *efficiency gaps* $\kappa_i := (1+\tau_i)\,\tilde{\omega}_i/\omega_i \geq 1$. As it multiplies g_i in (31), a higher κ_i effectively increases the cost of allocating risk management resources to firm i. When $\tau = 0$, revenue-based and cost-based Domar weights coincide, and all the efficiency gaps are equal to one. In this case, the cost functions of the distorted planner and of the social planner are identical: $\bar{h}(\Delta) = \bar{h}_{SP}(\Delta)$ for all Δ. Away from $\tau = 0$, efficiency gaps are increasing functions of τ, so that higher wedges lead to bigger distortions in $\bar{h}$.[23] Without input-output linkages ($\alpha = 0$), the efficiency gaps take a particularly simple form, $\kappa_i = 1 + \tau_i$, and a change in wedge τ_i only affects the efficiency gap of firm i.

Why do we need to distort $\bar{h}_{SP}(\Delta)$ to replicate the equilibrium allocation? It helps to compare the social and the private benefits of managing risk. Recall from (17) that the contribution of δ_i to GDP—a measure of its social value—is proportional to the cost-based Domar weight $\tilde{\omega}_i$. In contrast, the first-order condition (22) of the firm shows that the cost $K_i Q_i \propto \omega_i/(1+\tau_i)$ is a key driver of the private incentives to manage risk. Without wedges ($\tau = 0$), $\tilde{\omega}_i = \omega_i/(1+\tau_i)$ and the social and private incentives for risk management are aligned. When τ increases, however, $\tilde{\omega}_i$ stays the same but, as we have discussed, $K_i Q_i$ shrinks. It follows that the gap between private and social incentives to manage risk increases. Adjusting $\bar{h}_{SP}$ by the ratio $\kappa_i := (1+\tau_i)\,\tilde{\omega}_i/\omega_i$ therefore distorts the incentives of the fictitious planner to reach the equilibrium.

We will use $\bar{h}$ to characterize the equilibrium, but it is useful to first derive some of its properties.

Lemma 4. *The aggregate cost function $\bar{h}$ is given by*

$$\bar{h}(\Delta) = \frac{1}{2}\left(\Delta - \Delta^\circ\right)^\top \bar{H}\left(\Delta - \Delta^\circ\right), \tag{32}$$

where $\Delta^\circ := (\delta^\circ)^\top \tilde{\omega}$. Furthermore, the Hessian matrix of $\bar{h}$ is given by

$$\bar{H} = \left[\sum_{i=1}^{N} \tilde{\omega}_i H_i^{-1}\right]^{-1}, \tag{33}$$

where H_i is given by (24).

This result shows that the distorted utility cost $\bar{h}(\Delta)$ of reaching aggregate exposure Δ is a quadratic function of Δ. When all firms adopt their natural risk-exposure levels, $\delta = \delta^\circ$, the aggregate exposure is also at its natural level $\Delta^\circ = (\delta^\circ)^\top \tilde{\omega}$, and the cost of reaching that exposure is $\bar{h}(\Delta^\circ) = 0$. Whenever Δ departs from Δ°, the utility cost increases by an amount that depends

[22] To be precise, $\bar{b}(\Delta) := \tilde{\omega}^\top b(\delta^*(\Delta))$ and $\bar{g}(\Delta) := -\log V\left(\sum_{i=1}^N \kappa_i g_i(\delta_i^*(\Delta))\right)$.

[23] As we have seen, increasing τ_j lowers the revenue-based Domar weights of all firms. From the definition of κ_i, it follows that $\frac{d\kappa_i}{d\tau_j} \geq 0$. See (86) in the appendix for the exact formula.

on the curvature $\bar{H}$ of $\bar{h}$. As (33) shows, this curvature is a weighted harmonic average of the underlying individual effective cost matrices H_i. Intuitively, if all firms must pay a large cost—either in terms of productivity or in terms of resources—to gain exposure to some risk factor m (large $[H_i]_{mm}$ for all i), then the aggregate cost of increasing Δ_m is also large.

We can finally use the definition of $\bar{h}$ to characterize the equilibrium and its basic properties.

Proposition 1. *There exists a unique equilibrium, and its aggregate risk exposure Δ^* solves*

$$\mathcal{W}_{dist} := \max_{\Delta} \underbrace{\Delta^\top \mu - \bar{b}(\Delta) - \tilde{\omega}^\top \log(1+\tau) - \log \Gamma_L}_{\mathrm{E}[y]} - \frac{1}{2}(\rho - 1) \underbrace{\Delta^\top \Sigma \Delta}_{\mathrm{V}[y]} - \bar{g}(\Delta), \tag{34}$$

where $\bar{b}(\Delta) + \bar{g}(\Delta) = \bar{h}(\Delta)$ is given by (32). Furthermore, without wedges ($\tau = 0$), the equilibrium is efficient.

Proposition 1 shows that the equilibrium allocation is the solution to a *fictitious planning problem* distorted by wedges. Wedges enter (34) by distorting expected log GDP $\mathrm{E}[y]$ and by increasing the cost of risk exposure embedded in $\bar{h} = \bar{b} + \bar{g}$. However, since Δ does not interact with the wedges in $\mathrm{E}[y]$, only the distortions in $\bar{h}$ affect the equilibrium exposure decisions. The last part of the proposition shows that τ is the only source of distortions in this environment.

6 Forces shaping risk-exposure decisions

Primitives of the economy, such as the moments of the risk factors and the wedges, affect aggregate risk exposure Δ through their influence on the risk-management incentives of the firms. In this section, we solve the model and provide closed-form expressions for Δ in terms of model primitives. Our analysis relies on the optimality condition of the fictitious planner's problem (34).

Lemma 5. *The equilibrium aggregate risk exposure Δ solves*

$$\mathcal{E}(\Delta) = \nabla \bar{h}(\Delta), \tag{35}$$

where the marginal value of aggregate risk exposure $\mathcal{E}$ can be written as

$$\mathcal{E} = \underbrace{\mu}_{\mathrm{E}[\varepsilon]} \underbrace{- (\rho - 1)\Sigma\Delta}_{\mathrm{Cov}[\lambda, \varepsilon]}. \tag{36}$$

The first part of the lemma is the aggregate analogue of the firm's optimality condition (25). It states that in equilibrium, the marginal cost $\nabla \bar{h}$ of aggregate exposure is equal to its marginal benefit to the household's expected utility, which is given by the vector $\mathcal{E}$. Recall that $\mathcal{E}$ also captures the benefit of marginal exposure from the firm's perspective.

The second part of the lemma provides an expression for $\mathcal{E}$ and its two components, $\mathrm{E}[\varepsilon]$ and $\mathrm{Cov}[\lambda, \varepsilon]$. To better understand that expression, consider first an economy with a single risk factor. In this case, we can represent the relationship between $\mathcal{E}$ and Δ on a simple graph, as in the left panel of Figure 2. That figure also shows the quadratic relationship between Δ and $\mathrm{V}[y]$ implied by (19). It is clear from this figure that if the economy is positively exposed to the risk factor ($\Delta > 0$), any further increase in Δ would add to aggregate volatility. In contrast, under negative exposure ($\Delta < 0$), any marginal increase in Δ lowers aggregate volatility. Since the household dislikes uncertainty, a positive-exposure risk factor tends to have a negative marginal value of exposure $\mathcal{E}$, while a negative-exposure factor has a positive $\mathcal{E}$. In fact, $\mathcal{E}$ is always linearly decreasing in Δ, as (36) shows. When there is more than one risk factor, correlations matter as well. If the economy is heavily exposed to factors that positively covary with ε_m, any increase in Δ_m would raise aggregate volatility. In that case, the marginal value of exposure to ε_m would be low.

The moments of the risk factors also affect $\mathcal{E}$. Let $\mathbf{1}_m$ be the mth standard basis vector. Then[24]

$$\frac{\partial \mathcal{E}}{\partial \mu_m} = \mathbf{1}_m \qquad \text{and} \qquad \frac{\partial \mathcal{E}}{\partial \Sigma_{mn}} = -\frac{1}{2}(\rho - 1)(\Delta_m \mathbf{1}_n + \Delta_n \mathbf{1}_m). \tag{37}$$

An increase in μ_m always makes factor m more attractive but leaves the marginal value of exposure to all other factors unchanged. To understand the intuition behind the second equation in (37), consider the right panel of Figure 2, which shows what happens when uncertainty Σ increases. In this case, we see that Δ has a stronger impact on the variance of GDP. It follows that marginal increases in Δ are now more harmful for $\Delta > 0$ and more beneficial for $\Delta < 0$. As a result, the relationship between $\mathcal{E}$ and Δ becomes steeper, as Figure 2 and equation (37) show. When there is more than one factor, a similar reasoning applies to changes in covariances.[25]

6.1 Aggregate risk exposure

From (35), we can solve for the equilibrium exposure decisions.

Proposition 2. *The equilibrium aggregate risk-exposure decisions are given by*

$$\Delta = \Delta^\circ + \mathcal{H}^{-1}\mathcal{E}^\circ, \tag{38}$$

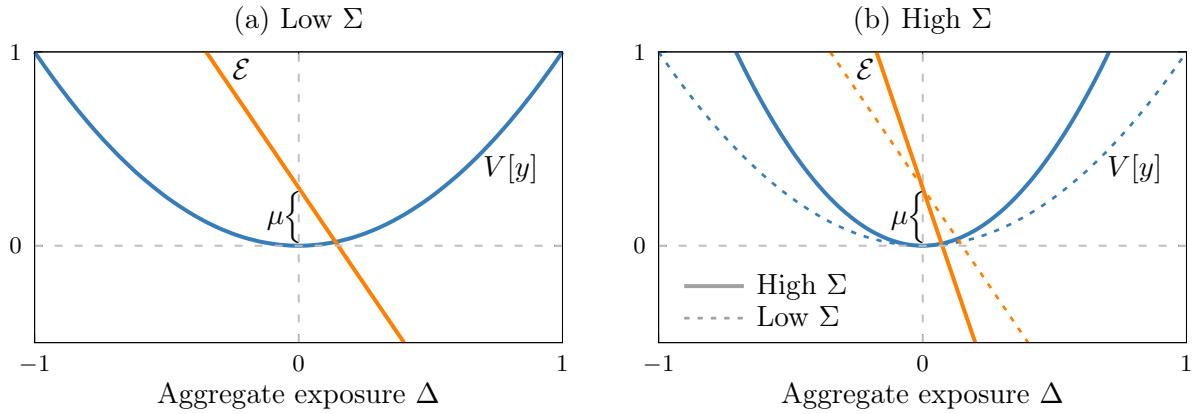
where $\mathcal{E}^\circ := \mu - (\rho - 1)\Sigma\Delta^\circ$ *and where the* $M \times M$ *positive definite matrix* $\mathcal{H}^{-1}$ *is*

$$\mathcal{H}^{-1} := \left(\bar{H} + (\rho - 1)\Sigma\right)^{-1}. \tag{39}$$

[24] Whenever we take derivatives with respect to off-diagonal elements of Σ, we simultaneously change Σ_{mn} and Σ_{nm} to keep Σ symmetric, and divide the result by two.

[25] Specifically, an increase in Σ_{mn} decreases the marginal value of exposure to both factors m and n if $\Delta_m > 0$ and $\Delta_n > 0$. In this case, the larger covariance implies that a marginal increase in Δ_m or Δ_n would translate into a larger increase in aggregate risk. The opposite happens when $\Delta_m < 0$ and $\Delta_n < 0$. Equation (37) also shows that a higher risk aversion ρ magnifies the impact of Σ on $\mathcal{E}$.

Figure 2: Impact of Δ on $\mathrm{V}[y]$ and $\mathcal{E}$ in an economy with a single risk factor



Equation (38) provides a closed-form expression for aggregate risk exposure Δ, and is the aggregate analogue of (26) which describes firm-level risk decisions.[26] We see from (38) that Δ depends on a vector $\mathcal{E}^\circ$, which corresponds to the marginal benefit vector $\mathcal{E}$ when aggregate exposure Δ is at its natural level Δ°. The impact of $\mathcal{E}^\circ$ on Δ is mediated by a matrix $\mathcal{H}^{-1}$, given by (39). Since $\mathcal{H}^{-1}$ is positive definite, if a risk factor m is naturally attractive ($\mathcal{E}^\circ_m > 0$), firms and the aggregate economy tend to be positively exposed to it. Whether a naturally attractive risk factor m contributes to exposure to another factor $n \neq m$ depends, however, on global substitution patterns that are encoded in $\mathcal{H}^{-1}$. If $\left[\mathcal{H}^{-1}\right]_{mn} > 0$, we say that risk factors m and n are *global complements*, in which case an increase in $\mathcal{E}^\circ_m$ leads, all else equal, to an increase in Δ_n. If, instead, $\left[\mathcal{H}^{-1}\right]_{mn} < 0$, m and n are *global substitutes*, and an increase in $\mathcal{E}^\circ_m$ reduces Δ_n.

These global substitution patterns depend on the Hessian matrix $\bar{H}$ of the aggregate cost function $\bar{h}$. An element $\bar{H}_{mn}$ captures how an increase in exposure to factor m changes the *marginal* cost of gaining exposure to factor n. Recall from (33) that $\bar{H}$ is a Domar-weighted average of the Hessians $(H_1, \ldots, H_N)$ of the firm-level effective cost functions. It follows that $\bar{H}$ reflects the global impact of the local substitution patterns embedded in those functions. The global substitution patterns also depend on the covariance matrix Σ. Intuitively, a positive correlation between two risk factors m and n contributes to these two factors being global substitutes. Indeed, if Δ_m increases, the fictitious planner would favor a decline in Δ_n to avoid too much aggregate risk. Unsurprisingly, the intensity of that channel depends on the household's risk aversion ρ.

6.2 Beliefs and aggregate risk exposure

Changes in the moments (μ, Σ) affect how attractive the risk factors are, which in turn affects risk exposure decisions.

[26] Equation (38) describes Δ directly in terms of parameters. From it, we can compute $\mathcal{E}$ using (36) and δ using (26). Other objects can be computed directly using these expressions. Proposition 2 therefore leads to a closed-form characterization of the whole equilibrium.

Proposition 3. *Let γ denote either the mean μ_m or an element Σ_{mn} of the covariance matrix. The response of the equilibrium aggregate risk exposure Δ to a change in γ is given by*

$$\frac{d\Delta}{d\gamma} = \mathcal{H}^{-1}\frac{\partial \mathcal{E}}{\partial \gamma}, \tag{40}$$

where $\partial\mathcal{E}/\partial\gamma$ is given by (37).

Proposition 3 shows that the impact of a change in a moment γ on Δ operates through its *direct* impact on exposure value $\mathcal{E}$, as captured by the partial derivative $\partial\mathcal{E}/\partial\gamma$. That is, a change in γ makes some risk factors more or less attractive than before, and this triggers an adjustment of risk-taking decisions through the matrix $\mathcal{H}^{-1}$. Since $\mathcal{H}^{-1}$ is positive definite, it is straightforward to characterize the outcome of that adjustment process.

Corollary 2. *An increase in the expected value μ_m of risk factor m leads to an increase in aggregate risk exposure Δ_m. An increase in the variance Σ_{mm} of risk factor m leads to a decrease in Δ_m if $\Delta_m > 0$ and to an increase in Δ_m if $\Delta_m < 0$.*

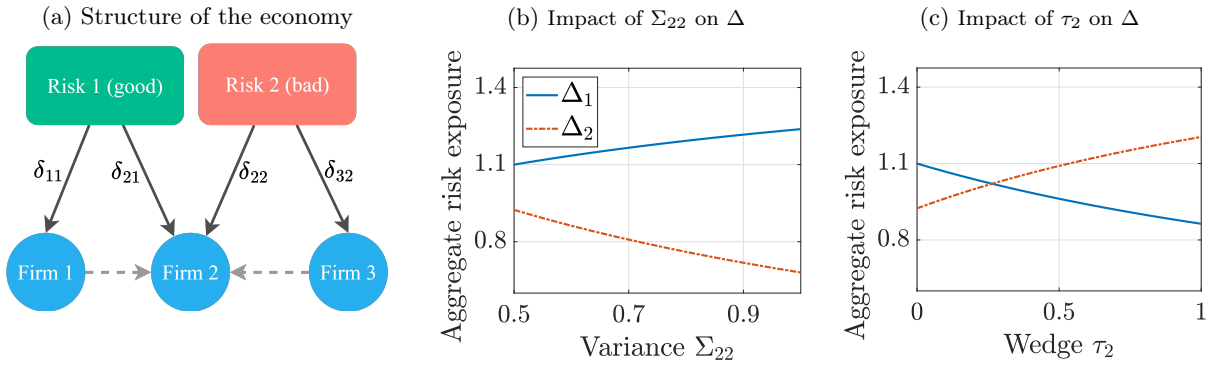
Increasing μ_m makes exposure to ε_m more attractive without affecting the cost $\bar{h}$ of achieving that exposure. Consequently, the fictitious planner chooses a higher Δ_m. The impact of Σ is more subtle. Recall from (19) that the sensitivity of $\mathrm{V}[y]$ to Σ_{mm} depends on the absolute value of Δ_m. Indeed, an increase in Σ_{mm} makes GDP more volatile even if $\Delta_m < 0$. It follows that a higher Σ_{mm} makes the fictitious planner reduce $|\Delta_m|$, which explains the second part of the corollary.

Example. As in Proposition 2, the global substitution matrix $\mathcal{H}^{-1}$ controls how changes in the attractiveness of a risk factor ε_m affect exposure to another factor, say ε_n. To better understand how these substitution patterns work, consider a simple economy in which three firms can be exposed to two sources of risk, as in panel (a) of Figure 3. Risk 1 is good with $\mathcal{E}_1 > 0$, while risk 2 is bad with $\mathcal{E}_2 < 0$. The cost functions g are parametrized so that firm 1 is only exposed to risk 1 and cannot adjust its exposure. Similarly, firm 3 is only exposed to risk 2. Firm 2, in contrast, is free to change its exposure to the two risk factors, and its cost function h_2 is such that δ_{21} and δ_{22} are substitutes. All firms use labor to produce but, in addition, firm 2 uses goods from firms 1 and 3 (see the caption of Figure 3 for exact parametrization).

One can think of many real-world examples that roughly fit this simple framework. For instance, firm 2 might only have a limited amount of resources (management attention or specialized scientists) that can be allocated to promising R&D research (good risk) or to mitigating the impact of a potential epidemic on operations (bad risk). Another example may involve two regions, such that region 1 is more productive than region 2 (for instance, due to better climate). Firm 2 must decide where to locate its plants. In this context, δ_{21} and δ_{22} are substitutes if it is challenging for firm 2 to manage geographically dispersed plants.

Panel (b) in Figure 3 shows what happens to the equilibrium aggregate risk exposure when the bad factor becomes more risky. When Σ_{22} increases, $\mathcal{E}_2$ declines and risk factor 2 becomes less attractive. As a result, firm 2 reduces its exposure to it, and since the other firms' exposures are fixed, Δ_2 declines. As the two risk factors are substitutes, increasing δ_{21} then becomes cost-effective, and the overall economy's exposure to the good factor increases.

Figure 3: Impact of parameters in a simple economy



Notes: Panel (a): The structure of the economy; there is an arrow from firm j to firm i if $\alpha_{ij} > 0$, and from risk factor m to firm i if $\delta_{im} \neq 0$. Panels (b) and (c): effects of changes in parameters. Initial parametrization is as follows. Household: $\rho = 2$ and $\beta_2 = 0.8$, $\beta_1 = \beta_3 = 0.1$. Network: $\alpha_{21} = \alpha_{23} = 0.25$, all other entries of α are zero. Beliefs: $\mu = (0.75, 0)$, Σ is diagonal with $\text{diag}(\Sigma) = (0.5, 0.5)$. Risk exposures: $\delta^\circ_{11} = \delta^\circ_{32} = 1$, $\delta^\circ_{22} = 1.9$, $\delta^\circ_{12} = \delta^\circ_{21} = \delta^\circ_{31} = 0$, $G_1 = G_3$ are diagonal with very large entries on the main diagonals; $G_{2,11} = G_{2,22} = 1$, $G_{2,12} = G_{2,21} = 0.75$. We set $B_i = 0$ for all i. Wedges: $\tau = 0$. In panel (b), Σ_{22} changes from 0.5 to 1. In panel (c), τ_2 changes from 0 to 1.

6.3 Wedges and aggregate risk exposure

Wedges affect firms' risk-taking incentives and, through that channel, aggregate risk exposure.

Proposition 4. *The response of the equilibrium aggregate risk exposure Δ to a change in wedge τ_i is given by*

$$\frac{d\Delta}{d\tau_i} = \frac{d\mathcal{H}^{-1}}{d\tau_i}\mathcal{E}^\circ, \tag{41}$$

where

$$\frac{d\mathcal{H}^{-1}}{d\tau_i} = -\mathcal{H}^{-1}\left(\sum_{j=1}^{N}\frac{\partial \bar{H}}{\partial \kappa_j}\frac{d\kappa_j}{d\tau_i}\right)\mathcal{H}^{-1},$$

and where the response $d\kappa_j/d\tau_i$ of the efficiency gap is nonnegative, and $\partial \bar{H}/\partial \kappa_j$ is a positive definite matrix.[27]

Wedges affect the equilibrium risk exposure Δ through the curvature of the aggregate cost function $\bar{h}$. Specifically, an increase in τ_i leads to a higher efficiency gap κ_j for all j. This, in

[27] We give expressions for $d\kappa_j/d\tau_i$ and $\partial \bar{H}/\partial \kappa_j$ in the proof of this proposition.

turn, means that firms find risk management less appealing, and the curvature of the aggregate cost function increases, in the sense that $\partial \bar{H}/\partial \kappa_j$ is a positive definite matrix. These changes in curvature then affect $\mathcal{H}^{-1}$ and, hence, Δ.

Proposition 4 provides a general characterization of the impact of τ, but we can get sharper results under a particular parametrization of the model.

Definition 2. An economy is *diagonal* if the risk factors are uncorrelated (diagonal Σ), and the individual cost functions $(b_1, \ldots, b_N)$ and $(g_1, \ldots, g_N)$ feature neither local complementarities nor local substitutabilities (diagonal B_i and G_i for all i).

Correlations between risk factors and complex substitution patterns can give rise to interesting mechanisms, but they often obscure some simpler forces that are at work in the economy. To highlight those forces, we will sometimes simplify the analysis by focusing on diagonal economies. The following result describes the impact of τ in such a setting.

Corollary 3. *In a diagonal economy, a higher wedge τ_i increases Δ_m for all m such that $\mathcal{E}_m < 0$ (bad risks) and decreases Δ_m for all m such that $\mathcal{E}_m > 0$ (good risks).*

This result shows that wedges push firms to increase their exposure to bad risks and to decrease their exposure to good risks. Intuitively, an increase in τ_i makes firms shrink in terms of their cost of goods sold. This implies that risk management becomes less cost-effective: since the firm is smaller, managing its TFP is now less rewarding. As a result, firms spend less to increase exposure to good risk factors and to reduce exposure to bad risk factors. As we will see in the next section, this has important implications for welfare.[28]

Example. The last panel of Figure 3 shows what happens to the example economy of panel (a) when firm 2's wedge increases. As a result of its higher sales price, firm 2 shrinks, which makes risk management relatively more costly. It follows that aggregate exposure to the good risk factor declines and exposure to the bad risk factor increases.

6.3.1 Equilibrium and efficient risk exposure

Proposition 4 implies that distortions matter for risk-taking decisions. As a result, the equilibrium Δ departs from its efficient level.

Lemma 6. *Suppose that $\tau_j > 0$ for at least one firm j. Then $(\Delta - \Delta_{SP})^\top \mathcal{E}^\circ < 0$, where Δ and Δ_{SP} are the aggregate risk exposure vectors in the equilibrium and the efficient allocation, respectively. Furthermore, in a diagonal economy the sign of $\Delta_i - \Delta_{SP,i}$ is the opposite of that of $\mathcal{E}_i^\circ$.*

[28] In addition to μ, Σ and τ, it is straightforward to describe how other parameters affect Δ. For instance, in Appendix F.1 we explore the role of the production network α. Our main finding is that an increase in α_{ij} typically leads to an increase in revenue-based Domar weights, which in turn implies that firms manage risk more aggressively in response.

When there is a single risk factor, the first part of Lemma 6 implies that the equilibrium is overexposed to bad natural risks ($\mathcal{E}^\circ < 0$) and underexposed to good natural risks ($\mathcal{E}^\circ > 0$), compared to the efficient allocation. Intuitively, wedges distort firms' incentives to manage their risk exposures and, when aggregated, those decisions lead to a departure from Δ_{SP}. When there are multiple risk factors, this result shows that these forces operate on average. For instance, if there are two risk factors, one good and one bad, the equilibrium can be overexposed to the good natural risk only if it is severely overexposed to the bad risk as well. The second part of the lemma shows that this over/underexposure result applies factor by factor in a diagonal economy.

7 GDP and welfare

Since changes in the environment affect the economy's risk exposure, they also affect GDP and welfare. In this section, we study the impact of endogenous risk on these quantities.

7.1 Moments of GDP

An increase in a parameter χ can influence the moments of GDP through two channels. First, it might trigger an adjustment in Δ, and that adjustment would, in turn, affect GDP. This is the main channel we are interested in in this paper. Second, an increase in χ may also have a direct impact on GDP that operates independently of χ's influence on Δ. For instance, an increase in wedges τ, on its own, can reduce GDP. This second channel is not specific to our model and, in many cases, has already been studied in the literature. In what follows, we therefore focus on the role played by changes in risk-exposure decisions, and filter out the impact of the second channel (denoted using partial derivatives in the expressions).

The following result describes how a change in a parameter affects GDP through Δ.[29]

Proposition 5. *Let χ denote either μ_m, Σ_{mn}, or τ_i. Then the impact of a change in χ on the moments of log GDP is given by*

$$\frac{d\,\mathrm{E}\,[y]}{d\chi} - \frac{\partial\,\mathrm{E}\,[y]}{\partial\chi} = \left(\mu - \nabla\bar{b}\right)^\top \frac{d\Delta}{d\chi} \qquad \textit{and} \qquad \frac{d\,\mathrm{V}\,[y]}{d\chi} - \frac{\partial\,\mathrm{V}\,[y]}{\partial\chi} = 2\Delta^\top\Sigma\frac{d\Delta}{d\chi}, \tag{42}$$

where the partial derivatives indicate that Δ is kept fixed, and where $d\Delta/d\chi$ is given by (40) for $\chi = \mu_i$ or $\chi = \Sigma_{mn}$, and by (41) for $\chi = \tau_i$.

Proof. The result follows directly from (19). □

Increasing χ triggers an adjustment in aggregate exposure Δ which, in turn, affects expected log GDP in two ways. First, the response of $\mathrm{E}\,[y]$ to χ depends on the mean μ of the risk factors

[29] Proposition 5 also applies when χ denotes a network link α_{ij}. In that case, the term $d\Delta/d\chi$ is given by (93) in Appendix F.1.

whose exposure responds to χ. Second, the adjustment in Δ also implies a different productivity cost of exposure $\bar{b}$. Recall, for instance, that an increase in μ_m raises Δ_m (Corollary 2). If $\mu_m > 0$, this additional exposure has a beneficial impact on $\mathrm{E}[y]$. However, an increase in Δ_m can also lead to a higher productivity cost $\bar{b}$. If this cost increase is sufficiently large, increasing μ_m can lower $\mathrm{E}[y]$. In general, a change in χ affects the economy's exposure to many risk factors, and so the first equation in (42) sums over all of them to get the overall effect on $\mathrm{E}[y]$.

The impact of χ on the variance of log GDP $\mathrm{V}[y]$ works in a similar way, but in this case, the correlations between the affected risk factors must be taken into account. Indeed, we can rewrite the second expression in (42) as $d\,\mathrm{V}[y]/d\chi - \partial\,\mathrm{V}[y]/d\chi = 2\,\mathrm{Cov}\left[y, (d\Delta/d\chi)^\top \varepsilon\right]$. It follows that the impact of χ depends on whether the risk factors whose exposure changes are positively or negatively correlated with log GDP. For example, since an increase in μ_m raises Δ_m, that change will lead to an increase in $\mathrm{V}[y]$ if ε_m is positively correlated with y.

We can further characterize the impact of (μ, Σ) on GDP in a diagonal economy. To keep our discussion concise, we focus on the case with resource costs only ($B_i = 0$ for all i). Empirical evidence presented in the next section suggests that resource costs are important, which motivates this choice. We tackle the general case with $B_i \neq 0$ in the proofs of Corollaries 4 and 5.

Corollary 4. *Let $B_i = 0$ for all i. In a diagonal economy, the following statements hold.*

1. *The impact of an increase in μ_m on GDP satisfies*

$$sign\left(\frac{d\,\mathrm{E}[y]}{d\mu_m} - \frac{\partial\,\mathrm{E}[y]}{\partial\mu_m}\right) = sign(\mu_m) \quad and \quad sign\left(\frac{d\,\mathrm{V}[y]}{d\mu_m} - \frac{\partial\,\mathrm{V}[y]}{\partial\mu_m}\right) = sign(\Delta_m). \tag{43}$$

2. *The impact of an increase in Σ_{mm} on GDP satisfies*

$$sign\left(\frac{d\,\mathrm{E}[y]}{d\Sigma_{mm}} - \frac{\partial\,\mathrm{E}[y]}{\partial\Sigma_{mm}}\right) = -sign(\mu_m\Delta_m) \quad and \quad \frac{d\,\mathrm{V}[y]}{d\Sigma_{mm}} - \frac{\partial\,\mathrm{V}[y]}{\partial\Sigma_{mm}} < 0. \tag{44}$$

The first result describes how an increase in μ_m affects GDP through Δ. Without productivity costs, $\mathrm{E}[y]$ rises in response if $\mu_m > 0$ and decreases otherwise, as explained above. The response of $\mathrm{V}[y]$, in contrast, depends on whether the economy is positively or negatively exposed to m. If $\Delta_m > 0$, the increase in μ_m makes Δ_m even more positive which makes the economy exposed to more risk, and $\mathrm{V}[y]$ rises as a result. If $\Delta_m < 0$ instead, the increase in μ_m makes Δ_m less negative, and so the economy is less vulnerable to shock m.

The second part of Corollary 4 offers similar expressions for the impact of an increase in variance. From Corollary 2, we know that an increase in Σ_{mm} leads to a decline in $|\Delta_m|$. This raises $\mathrm{E}[y]$ if $\mu_m\Delta_m < 0$, but lowers it otherwise. In contrast, the impact of Σ_{mm} on the variance of GDP via Δ is unambiguous. Since $|\Delta_m|$ shrinks, the economy becomes less sensitive to ε_m, and $\mathrm{V}[y]$ declines as well.

We can similarly characterize how wedges affect GDP. When there is a unique risk factor, this characterization is straightforward.

Corollary 5. *Let $B_i = 0$ for all i, and suppose that there is a single risk factor. Then*

$$sign\left(\frac{d\,\mathrm{E}\,[y]}{d\tau_i} - \frac{\partial\,\mathrm{E}\,[y]}{\partial\tau_i}\right) = -sign\left(\mu\mathcal{E}\right) \qquad and \qquad sign\left(\frac{d\,\mathrm{V}\,[y]}{d\tau_i} - \frac{\partial\,\mathrm{V}\,[y]}{\partial\tau_i}\right) = -sign\left(\Delta\mathcal{E}\right). \quad (45)$$

To get some intuition for this result, suppose there is a unique risk factor that is bad ($\mathcal{E} < 0$), and that firms are on average positively exposed to it ($\Delta > 0$). That risk factor could be, for example, an underlying productivity shock that drives the business cycle. Since wedges raise exposure to bad risk factors (Corollary 3), a higher τ_i increases Δ. Whether this has a positive impact on $\mathrm{E}\,[y]$ depends on the sign of μ. If $\mu > 0$, then $\mathrm{E}\,[y]$ also increases, while it declines if $\mu < 0$. Wedges also affect GDP volatility. In this case, since $\Delta > 0$, increasing τ_i makes the economy even more exposed to the risk factor, and $\mathrm{V}\,[y]$ increases.

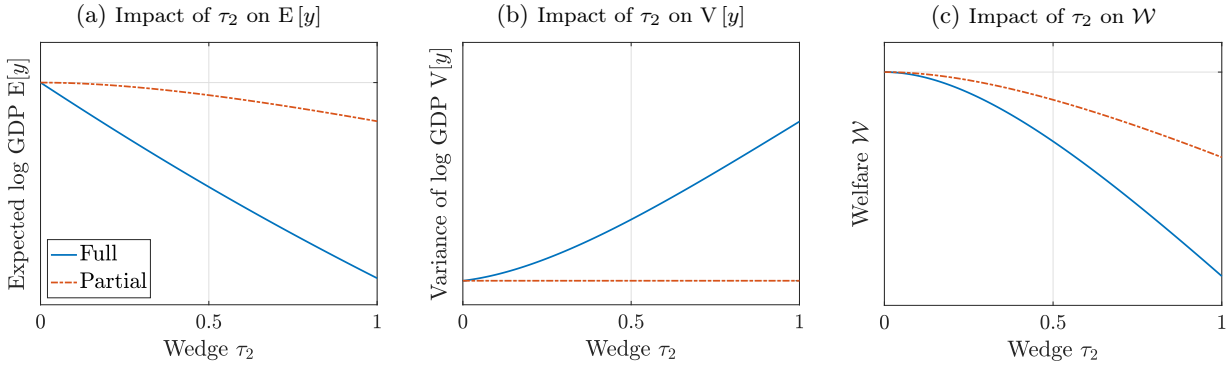
Example. When there are more than one risk factor, the response of the economy to a change in wedges depends on how correlated the risk factors are and on the substitution patterns in the firms' cost functions. To get a better sense of the forces at work, we can go back to the example economy of Figure 3. The first panel of Figure 4 shows the impact of raising τ_2 on $\mathrm{E}\,[y]$. Recall from the second panel of Figure 3 that increasing τ_2 makes the economy more exposed to risk 2 (the bad risk) and less exposed to risk 1 (the good risk). Since under our parametrization $\mu_1 > 0$ and $\mu_2 = 0$, this triggers a decline in $\mathrm{E}\,[y]$. Notice that this decline is more pronounced than if Δ remained fixed (red line). In that case, the increase in τ_2 would lower GDP only through its usual distortionary effect. We see that the endogenous risk management decisions of the firms make the response to an increase in wedges more severe.

The second panel of Figure 4 shows the response of $\mathrm{V}\,[y]$ to the same increase in τ_2. Both risk factors have the same variance, but given our parametrization, the increase in Δ_2 is stronger than the decline in Δ_1. It follows that the increase in τ leads to a rise in the volatility of log GDP. Notably, since $\mathrm{V}\,[y]$ does not directly depend on τ (see (19)), $\mathrm{V}\,[y]$ would be unaffected by wedges if Δ remained constant. This is illustrated by the red line in the panel. This example highlights that endogenous uncertainty is essential for wedges to affect aggregate volatility in our model.

7.2 Welfare

Model primitives affect GDP through their impact on risk-exposure decisions, but what the household ultimately cares about is welfare, which combines the moments of GDP with the disutility of managing risk. In this section, we evaluate the consequences of endogenous risk for welfare. We first characterize how welfare behaves in an environment without distortions.

Figure 4: The moments of GDP and welfare react to changes in the environment



Notes: The structure of the economy is given in panel (a) of Figure 3. Solid blue curve: full model. Red dashed curve: model with fixed Δ. Initial parametrization is as follows. Household: $\rho = 2$ and $\beta_2 = 0.8$, $\beta_1 = \beta_3 = 0.1$. Network: $\alpha_{21} = \alpha_{23} = 0.25$, all other entries of α are zero. Beliefs: $\mu = (0.75, 0)$, Σ is diagonal with diag $(\Sigma) = (0.5, 0.5)$. Risk exposures: $\delta^{\circ}_{11} = \delta^{\circ}_{32} = 1$, $\delta^{\circ}_{22} = 1.9$, $\delta^{\circ}_{12} = \delta^{\circ}_{21} = \delta^{\circ}_{31} = 0$, $G_1 = G_3$ are diagonal with very large entries on the main diagonals; $G_{2,11} = G_{2,22} = 1$, $G_{2,12} = G_{2,21} = 0.75$. We set $B_i = 0$ for all i. Wedges: $\tau = 0$. In panel (b), Σ_{22} changes from 0.5 to 1. In panel (c), τ_2 changes from 0 to 1.

Proposition 6. *Without wedges* ($\tau = 0$)*, the impact of the moments* (μ, Σ) *on welfare is given by*

$$\frac{d\mathcal{W}}{d\mu_m} = \frac{\partial \mathcal{W}}{\partial \mu_m} = \Delta_m \qquad \textit{and} \qquad \frac{d\mathcal{W}}{d\Sigma_{mn}} = \frac{\partial \mathcal{W}}{\partial \Sigma_{mn}} = -\frac{1}{2}(\rho - 1)\Delta_m \Delta_n.$$

Furthermore, the impact of the wedge τ_i *on welfare is given by*

$$\frac{d\mathcal{W}}{d\tau_i} = \frac{\partial \mathcal{W}}{\partial \tau_i} = 0.$$

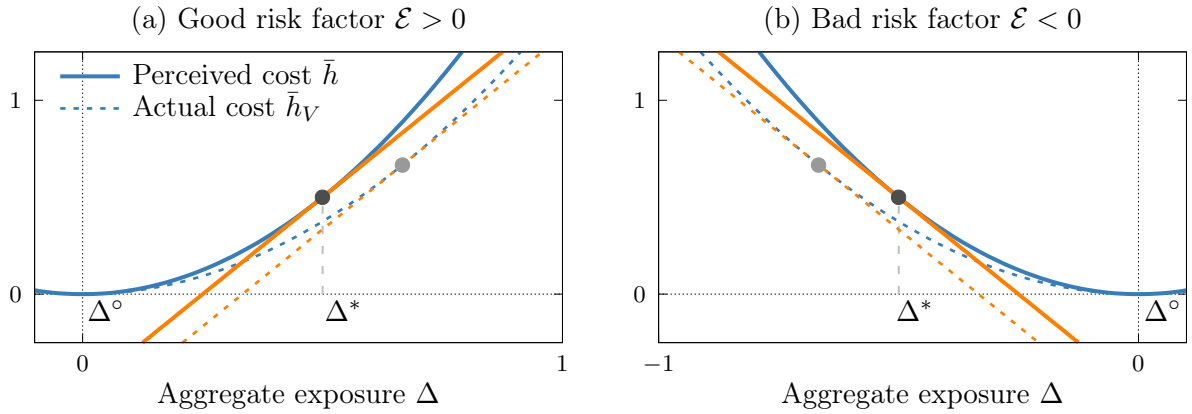
Proof. When $\tau = 0$, the objective function $\mathcal{W}_{dist}$ of the fictitious planner coincides with welfare. The result then follows from the envelope theorem. □

This result shows that making ε_m more productive on average (higher μ_m) benefits welfare if the economy is positively exposed to ε_m. In contrast, if $\Delta_m < 0$, increasing μ_m leads to a welfare loss. The proposition also shows that making ε_m more risky (higher Σ_{mm}) is always detrimental to welfare. Finally, an increase in the covariance Σ_{mn} hurts welfare if the economy is positively or negatively exposed to both risk factors. If, instead, Δ_m and Δ_n have opposite signs, the increase in Σ_{mn} is beneficial since the shocks are more likely to offset each other in this case.

Proposition 6 also shows that welfare responds to a marginal change in a parameter *as if* the risk-exposure decisions were kept fixed. This is a consequence of the envelope theorem. In the absence of wedges, the equilibrium is efficient and so Δ maximizes welfare. Any marginal movement around Δ must therefore have no impact on welfare.[30] The situation is, however, different in the presence

[30] Similarly, since $\tau = 0$ maximizes $\mathcal{W}$, a marginal increase in wedges from that point has no impact on welfare, as Proposition 6 shows (Baqaee and Farhi, 2019; Bigio and La'O, 2020).

Figure 5: Actual and perceived costs of adjusting risk exposure



of wedges, in which case the response of risk exposure can have a first-order effect.

The next result characterizes how welfare responds to changes in beliefs in a distorted equilibrium. To keep the analysis concise, we focus on diagonal economies here and derive the general results in Appendix E.15.

Lemma 7. *Suppose that the economy is diagonal and that $\tau_j > 0$ for at least one firm j. Then,*

$$sign\left(\frac{d\mathcal{W}}{d\mu_m} - \frac{\partial \mathcal{W}}{\partial \mu_m}\right) = sign\left(\mathcal{E}_m\right) \qquad \text{and} \qquad sign\left(\frac{d\mathcal{W}}{d\Sigma_{mm}} - \frac{\partial \mathcal{W}}{\partial \Sigma_{mm}}\right) = -sign\left(\Delta_m \mathcal{E}_m\right).$$

The first part of the lemma shows that increasing the expected value μ_m of a good risk factor ($\mathcal{E}_m > 0$) leads to a bigger welfare gain when risk is endogenous. To understand why, it helps to illustrate how the marginal benefit and marginal cost of exposure interact, as in the left panel of Figure 5. The blue solid curve shows the function $\bar{h}$, which reflects the cost of exposure as it is *perceived* by the fictitious planner, while the slope of the orange line is equal to $\mathcal{E}$. In equilibrium, we know from (35) that $\mathcal{E} = \nabla \bar{h}$ (black point in the figure). But recall from (31) that because of the distortions, the fictitious planner perceives the cost of exposure to be higher than it really is. The dashed blue curve in the figure represents the true utility cost of exposure, which we denote by $\bar{h}_V(\Delta)$.[31] It follows that at the equilibrium Δ^*, the true marginal cost $\nabla \bar{h}_V$ is smaller than the marginal benefit $\mathcal{E}$, and an increase in Δ would be welfare increasing. This is precisely what happens when μ_m increases: Δ_m goes up, which leads to an additional increase in welfare, as Lemma 7 shows. The right panel of Figure 5 shows the same quantities when the risk factor is bad ($\mathcal{E}_m < 0$). In this case, the perceived marginal cost $\nabla \bar{h}$ is *smaller* than the true marginal cost $\nabla \bar{h}_V$. In this case, increasing Δ would lead to a decline in welfare.

The second part of Lemma 7 shows that a similar result holds for an increase in Σ_{mm}. If ε_m is good, a higher exposure Δ_m would improve welfare. After an increase in Σ_{mm}, the fictitious

[31] We define $\bar{h}_V$ formally in (89) in Appendix E.15

planner reduces $|\Delta_m|$ to limit the economy's vulnerability to ε_m. It follows that if $\Delta_m > 0$, there is a decline in exposure, which lowers welfare. If instead $\Delta_m < 0$, the increase in risk Σ_{mm} leads to a higher exposure Δ_m, which improves welfare.

A similar result holds for the impact of wedges.

Lemma 8. *In a diagonal economy, an increase in wedges is more detrimental to welfare when risk-exposure decisions can adjust, that is,* $d\mathcal{W}/\partial\tau_i \leq \partial\mathcal{W}/\partial\tau_i$.

This result shows that endogenous risk increases the welfare cost of wedges. By Corollary 3, a higher τ_i increases exposure to bad risks and reduces exposure to good risks. Then Lemma 8 follows since additional exposure to bad risks and diminished exposure to good risks are both detrimental to welfare.

The right panel of Figure 4 illustrates how these forces operate in the model economy of Figure 3. If Δ is kept fixed, an increase in the wedge τ_2 of the central firm makes the distortions more important, which is detrimental for welfare (red line). But if the risk exposure Δ of the economy is free to adjust, the adverse impact of τ_2 on welfare is exacerbated (blue line) due to the higher exposure to the bad risk factor 2 and to the lower exposure to the good risk factor 1 (see second panel of Figure 3).

These results describe a new channel through which distortions like taxes and markups might reduce welfare. Beyond their conventional distortionary effects, these wedges also interfere with firm risk-taking incentives, altering the economy's aggregate risk profile in potentially harmful ways. This suggests that incorporating endogenous risk responses into the evaluation of tax policy and other interventions might be needed to accurately assess their welfare consequences.

8 Reduced-form evidence

Our theory predicts that the characteristics of a firm affect its risk-exposure choice and, as a result, the variance of its productivity and how it covaries with GDP. In this section, we verify that these predictions are visible in the data. We focus on business cycle risk and keep the analysis simple by assuming that there is a unique risk factor $\varepsilon_t \sim \text{iid}\,\mathcal{N}(0,\Sigma)$. We assume that firm i's TFP is given by

$$\log TFP_{it} = \delta_{it}\varepsilon_t - b_i(\delta_{it}) + \gamma_i t + v_{it}, \tag{46}$$

where γ_i is a firm-specific deterministic trend, and $v_{it} \sim \text{iid}\,\mathcal{N}(\mu_i^v, \Sigma_i^v)$ is a firm-specific shock to capture variation that is uncorrelated with GDP.[32]

[32] For simplicity, we assume that firms cannot adjust their exposure to v_{it} and γ_i. The analysis is similar if they can, but requires additional assumptions on the cost functions b_i and g_i. We can also allow for autocorrelation in ε_t and v_{it}, and for correlations between the v_{it} of different firms at the cost of extra complications.

In this setup, risk-taking decisions are time invariant ($\delta_{it} = \delta_i$) and Lemma 3 implies that[33]

$$\delta_i = \delta_i^\circ + \left(B_i + \eta \frac{1+\tau_i}{\omega_i} G_i \right)^{-1} \mathcal{E}. \tag{47}$$

It is natural to think that the aggregate economy is positively exposed to business cycle risk such that $\Delta > 0$.[34] In this case, aggregate risk is bad ($\mathcal{E} < 0$), and (47) implies that firms with larger revenue-based Domar weights ω_i or smaller wedges τ_i choose a lower risk exposure δ_i.

Using (17), we can compute the moments

$$\mathrm{V}\left[\Delta \log TFP_{it}\right] = 2\delta_i^2 \Sigma + 2\Sigma_i^v \quad \text{and} \quad \mathrm{Cov}\left[\Delta \log TFP_{it}, \Delta y_t\right] = 2\Delta\Sigma\delta_i + 2\tilde{\omega}_i \Sigma_i^v, \tag{48}$$

where $\Delta \log TFP_{it} = \log TFP_{it} - \log TFP_{it-1}$ and $\Delta y_t = y_t - y_{t-1}$. Combining these expressions with (47) implies that if $\delta_i > 0$, the TFP growth of firms with larger Domar weights and smaller wedges is less volatile and covaries less with GDP growth.[35]

To see if these predictions of the model are at work in reality, we use detailed firm-level data from Orbis that contain the near-universe of Spanish firms between 1995 and 2019. Our sample contains 7,465,325 firm-year observations. We use that data to construct firm-level measures of TFP, markups (our measure of wedges) and Domar weights. We briefly describe how we construct our sample below, and include more detail in Appendix B.2.1.

We compute each firm's revenue-based Domar weight as the ratio of its nominal sales to Spain's nominal GDP. Markups are estimated using the control function approach of De Loecker and Warzynski (2012). Specifically, the estimated markup is given by $1+\tau_{it} = \hat{\alpha}_{Li} / \left(\text{Wage Bill}_{it}/\text{Sales}_{it}\right)$, where $\text{Wage Bill}_{it}/\text{Sales}_{it}$ is the share of labor expenditure in sales and $\hat{\alpha}_{Li}$ is the estimated labor elasticity in production (specific to each 2-digit sector) using the Levinsohn and Petrin (2003) estimator with the Ackerberg, Caves, and Frazer (2015) correction.[36] The median markup is 1.25 and the median Domar weight is 4.96×10^{-7} across all firm-year observations. Finally, we compute each firm's TFP as a markup-corrected deflated Solow residual.

With these data, we first explore how firm characteristics affect the volatility of a firm's TFP growth. To do so, we compute the standard deviation of TFP growth for each firm, $\sigma_i\left(\Delta \log TFP_{it}\right)$,

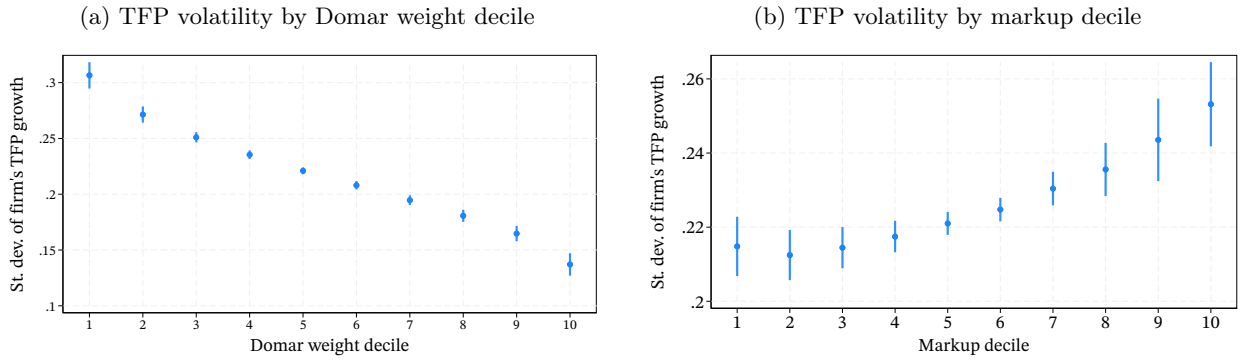
[33] See Appendix B.1 for a proof of time invariance and for the derivation of the equations of this section.

[34] Since $\mathrm{E}\left[\varepsilon_t\right] = 0$, the choice of $\Delta > 0$ is essentially a normalization that implies that firms have, on average, positive exposure $\delta_{it} > 0$. Alternatively, we can write an equivalent $\hat{\cdot}$ economy in which $\hat{\delta}^\circ = -\delta^\circ$, $\hat{\delta}_{it} = -\delta_{it}$, $\hat{\Delta} = -\Delta < 0$ and $\hat{\mathcal{E}} = -\mathcal{E} > 0$. That economy is indistinguishable from the original in terms of the variance of firm-level TFP growth and its covariance with GDP, as (48) shows below. Indeed, our results in Section 7 on the effect of Σ on the moments of GDP and welfare depend on the sign of the product of Δ and $\mathcal{E}$, not on their individual signs.

[35] With $\delta_i < 0$, a higher ω_i or a lower τ_i implies more volatile TFP growth. In our calibration, $\delta_i > 0$ for most firms.

[36] We discuss the robustness of our results to other markup measures at the end of this section. Rovigatti and Mollisi (2018) show that the output elasticity estimates with the Ackerberg, Caves, and Frazer (2015) correction are sensitive to the initial conditions used for optimization. To ensure that our estimates are less affected by this issue, we use 3,000 starting points. See Appendix B.3.5 for details.

Figure 6: TFP volatility, Domar weights and markups



Notes: Estimation results of (49) using a sample of Spanish firms from Orbis. Sample construction is described in Appendix B.2. The sample is winsorized at the top and bottom 0.5% of observations of TFP growth, Domar weight and markup in the panel in each sector before the computing the firm-level moments and only includes firms that have 1) at least 5 non-missing observations of TFP growth, and 2) positive average log markup. Panel (a): $\beta_j^{Domar} + \alpha + \beta_5^{Markup}$ by Domar weight decile j ($j = 1$ is lowest Domar weight, $j = 10$ is highest Domar weight). Panel (b): $\beta_j^{Markup} + \alpha + \beta_5^{Domar}$ by markup decile ($j = 1$ is lowest markup, $j = 10$ is highest markup). 90% confidence intervals are constructed using standard errors that are clustered at the industry level.

and the time-series averages of its markups and Domar weights. We then construct within-sector deciles based on these averages to create a set of dummy variables, D_{ji}^{Domar} and D_{ji}^{Markup}, such that $D_{ji}^{Domar} = 1$ if firm i's Domar weight is in decile j, and analogously for markups.[37] We then run the cross-sectional regression

$$\sigma_i \left(\Delta \log TFP_{it} \right) = \alpha + \sum_{j=1}^{10} \beta_j^{Domar} D_{ji}^{Domar} + \sum_{j=1}^{10} \beta_j^{Markup} D_{ji}^{Markup} + \theta_s + \varepsilon_i, \tag{49}$$

where θ_s are sector-specific fixed effects, and plot the estimated coefficients β_j^{Domar} in panel (a) and β_j^{Markup} in panel (b) of Figure 6. Consistent with the theory, firms with lower Domar weights and higher markups tend to have more volatile TFP growth. These relationships are statistically and economically significant. A firm in the top decile of the Domar weight distribution is about 17 p.p. less volatile than a firm in the bottom decile. In contrast, firms in the top decile of the markup distribution are about 4 p.p. more volatile than firms in the bottom decile.

In Appendix B.3, we run a simple cross-sectional regression of the standard deviation of firm-level TFP growth on Domar weights and markups. We also conduct an analogous analysis with US firms from Compustat. In both cases, we find similar results.

Next, we turn to our model's prediction that the TFP of firms with higher Domar weights and lower markups covaries less with GDP. Again, we construct a set of dummy variables, D_{jit}^{Domar} and

[37] We consider variation within 2-digit sectors in our reduced form analysis because firms in the same sector are more likely to have similar natural risk exposure δ_i° and risk management cost functions g_i and b_i. As shown by De Ridder, Grassi, Morzenti, et al. (2024), within-sector markup dispersion is more accurately measured than cross-sector dispersion when using revenue data. Our results hold even if cross-sector variation is taken into account.

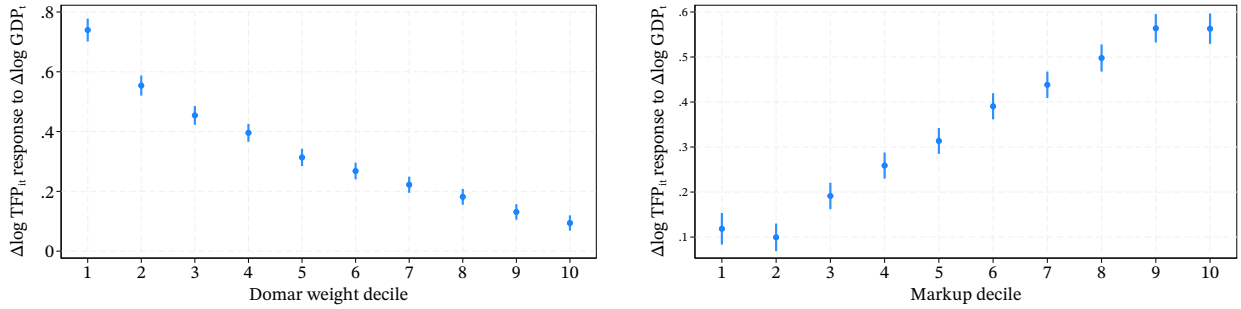
D_{jit}^{Markup}, such that $D_{jit}^{Domar} = 1$ if firm i's Domar weight is in within-sector decile j in year t, and analogously for markups. We then run the following panel regression,

$$\Delta \log TFP_{it} = \sum_{j=1}^{10} \beta_j^{Domar} \left(D_{jit}^{Domar} \times \Delta \log GDP_t \right) + \sum_{j=1}^{10} \beta_j^{Markup} \left(D_{jit}^{Markup} \times \Delta \log GDP_t \right) + \alpha + \beta_0 \Delta \log GDP_t + \sum_{j=1}^{10} D_{jit}^{Domar} + \sum_{j=1}^{10} D_{jit}^{Markup} + \theta_s + \varepsilon_{it}, \tag{50}$$

where θ_s are sector-specific fixed effects. The coefficients of interest, β_j^{Domar} and β_j^{Markup}, are reported in Figure 7. Consistent with the model, firms with higher Domar weights and lower markups covary less with GDP.[38] In Appendix B.3.2, we confirm these results through panel regressions with various sets of fixed effects.

Figure 7: Sensitivity of firm-level TFP to GDP

(a) Sensitivity of firm TFP to GDP by Domar weight decile (b) Sensitivity of firm TFP to GDP by markup decile



Notes: Estimation results of (50) using a sample of Spanish firms from Orbis. Sample construction is described in Appendix B.2. The estimation sample is winsorized at the top and bottom 0.5% of observations of TFP growth, Domar weight and markup in each sector. Panel (a): $\beta_j^{Domar} + \beta_0 + \beta_5^{Markup}$ by Domar weight decile j ($j = 1$ is lowest Domar weight, $j = 10$ is highest Domar weight). Panel (b): $\beta_j^{Markup} + \beta_0 + \beta_5^{Domar}$ by markup decile ($j = 1$ is lowest markup, $j = 10$ is highest markup). 90% confidence intervals are constructed using standard errors that are clustered at the firm level.

In Appendix B, we provide several additional exercises to ensure the robustness of our results. First, we estimate markups using materials instead of labor as the flexible input (Raval, 2023). Second, since markups estimated using the production function approach with revenue data could be biased (Bond et al., 2021), we consider profit accounting markups based on the Lerner index (Baqaee and Farhi, 2019; Burstein, Carvalho, and Grassi, 2024).[39] Third, since the control function

[38] Sales of large firms also also covary less with GDP than those of small firms in US data (Crouzet and Mehrotra, 2020).

[39] In response to the Bond et al. (2021) critique, De Ridder, Grassi, Morzenti, et al. (2024) use firm-level price and quantity data from France to estimate markups directly and find that the production approach that we use performs well when looking at the *cross-section* of markups across firms, which is the variation that we rely on in all of our reduced-form exercises.

approach of production function estimation requires specific assumptions on the timing of input choices, we use an IV-GMM production function estimator instead (Blundell and Bond, 2000). In all cases, we find that firms with higher Domar weights and lower markups are less volatile, and their productivities covary less with GDP.

Well-established facts about stock returns are also consistent with our model. As documented by Fama and French (1992), larger firms tend to comove less with the aggregate stock market. The mechanisms of the model provide an explanation for this stylized fact, as long as productivity increases generally raise a firm's stock price. Additionally, we show in Appendix B.4 that firms with higher markups exhibit greater covariance with the aggregate stock market, consistent with our theoretical predictions. There are also indications that the model mechanisms are visible in aggregate data. In Appendix B.5, we show that richer countries have smaller TFP volatility. If those countries have access to better risk-management technologies (smaller H_i or η), they should be better at reducing their risk exposure, consistently with this pattern. We also find that countries with higher markups feature more volatile TFP, in line with the forces of the model.[40]

Overall, the findings of this section suggest that key predictions of our theory are in line with firm-level productivity data, stock market return data, and stylized country-level evidence. Of course, other mechanisms could also be at work (Yeh, 2023). For instance, large firms might be able to average out many plant-level shocks, which would reduce their volatility. We view our mechanism as complementing these other stories. We are also reassured by the fact that our theory can explain patterns related to the variance and the covariance of firms jointly.

9 Calibration to the Spanish economy

To evaluate the quantitative importance of endogenous risk-taking decisions for the macroeconomy, we provide a basic calibration of the model to the Spanish economy. We rely on the detailed firm-level data introduced in the previous section to identify the parameters of the model. With the calibrated model in hand, we investigate how the economy handles an increase in the variance of a risk factor, and how wedges affect aggregate volatility. We also explore the implications of endogenous uncertainty in a calibrated economy with rare disasters. We present an overview of our calibration strategy below and include more detail in Appendix C.

9.1 Model with sectors

We specialize the general model of Section 2 to better map the moments of the data to model quantities. Specifically, we assume that some firms act as sectoral aggregators, and that individual

[40] In Appendix D.2, we estimate the full firm-level productivity factor structure of the model (Stock and Watson, 2002; Kelly, Pruitt, and Su, 2020). Our estimated factor loadings vary with Domar weights and markups as predicted by the theory.

firms purchase intermediate inputs from these aggregators directly. In that setup, we can use the Spanish sectoral input-output data to discipline the matrix α of network connections.

There are S sectors. In each sector s, there is an aggregator that converts the output of the N_s firms in that sector into a sector-specific good according to the production function $Q_s = e^{z_s}\prod_{i=1}^{N_s}\left(\theta_{si}^{-1}Q_{si}\right)^{\theta_{si}}$, where $\sum_{i=1}^{N_s}\theta_{si}=1$, $z_s \sim \text{iid}\,\mathcal{N}\left(\mu_s^z,\Sigma_s^z\right)$ are sectoral productivity shocks, and where Q_{si} denotes the output of firm i in industry s. This firm, in turn, operates the production function

$$Q_{si} = \exp\left(\delta_{sit}\varepsilon_t - b_i\left(\delta_{sit}\right) + \gamma_{si}t + v_{sit}\right)\zeta_{si}L_{si}^{1-\sum_{s'=1}^{S}\hat{\alpha}_{ss'}}\prod_{s'=1}^{S}X_{si,s'}^{\hat{\alpha}_{ss'}}, \tag{51}$$

where $X_{si,s'}$ denotes the use of the composite good of sector s' by firm i in sector s. Notice that the input elasticities $\hat{\alpha}$ are sector-specific, in line with the available data. As before, firms price their goods at a markup τ_{si} above marginal cost.[41] Finally, the household only consumes goods produced by the sectoral aggregators so that GDP is given by $Y = \prod_{s=1}^{S}\left(\beta_s^{-1}C_s\right)^{\beta_s}$.

As in Section 8, the TFP of firm i is the sum of a trend $\gamma_{si}t$, an idiosyncratic shock $v_{sit} \sim \text{iid}\,\mathcal{N}\left(\mu_{si}^v,\Sigma_{si}^v\right)$, a risk exposure cost $b_i\left(\delta_{sit}\right)$, and an aggregate shock $\varepsilon_t \sim \text{iid}\,\mathcal{N}\left(0,\Sigma\right)$ to which firms can adjust their exposure.[42] The stationary structure of the model implies that firms make the same risk-exposure decisions every period, and so we drop the subscript t and write δ_{si}.

Below, we will look at how changes in the environment affect the variance of GDP growth. In the context of this model, this quantity is given by

$$\text{V}\left[y_t - y_{t-1}\right] = 2\Sigma\Delta^2 + 2\tilde{\omega}_f^\top\Sigma^v\tilde{\omega}_f + 2\tilde{\omega}_s^\top\Sigma^z\tilde{\omega}_s, \tag{52}$$

where $\tilde{\omega}_f$ and $\tilde{\omega}_s$ are the vectors of firm-level and sectoral cost-based Domar weights. The last two terms in that equation reflect exogenous sources of risk. They are the aggregates of the firm idiosyncratic and sectoral productivity shocks, and correspond to the granular contributions of the firms and sectors to aggregate risk (Gabaix, 2011). The first term, which depends on the aggregate risk exposure Δ, is where the endogenous risk mechanism is at work.

9.2 Calibration strategy

Our calibration strategy aims at replicating the firm-level Orbis dataset in our model economy. Our calibrated economy therefore features 62 sectors and 531,944 individual firms. The biggest sectors are "real estate (including imputed rents)" with a consumption share of 13% and "accommodation and food services" with a consumption share of 12%.

[41] We think of the sectoral producers as an aggregation device and so we assume that they have no markups and that they make no risk-exposure decisions.

[42] We assume that there is a unique aggregate risk factor for simplicity and tractability. In Appendix D.2, we perform a factor model estimation (Kelly, Pruitt, and Su, 2020) in the spirit of Stock and Watson (2002) and find that the predominant factor driving firm-level productivity resemble aggregate TFP, in line with our calibration assumption.

Some model quantities can be identified directly from the data. For instance, we set $(\beta_1, \ldots, \beta_S)$ to match the sectoral consumption shares in the 2010 Spanish National Accounts. The same data provide the matrix of sectoral input shares $\hat{\alpha}$ directly. Consistent with the optimization problem of the sectoral aggregators, we set θ_{si} to match the share of firm i in sector s's sales.

Our calibrated model is consistent with the reduced-form framework of Section 8. We therefore rely on our markup estimates from that section to pin down the firm-level markups τ_{si}. Our reduced-form analysis also provided estimates for the firm-level variance of TFP growth and for the covariance of each firm's TFP growth with GDP growth. We use that information together with (48) to identify the risk exposure δ_{si} and the variance Σ^v_{si} of each firm in the economy.

Next, we can use the information gathered so far to pin down the value of exposure $\mathcal{E}$. Doing so requires setting a value for the risk-aversion ρ and the variance Σ of the underlying risk factor. It turns out that given our calibration procedure, both numbers do not matter for the counterfactual exercises that we conduct. Changing ρ and Σ only leads to a rescaling of some other objects in the model, and so we do not need to take a stance on their values.[43]

Finally, we can estimate the parameters of the firm-level risk-management cost functions. To do so, we rely on Lemma 3, which states that

$$\delta_{si} = \delta^{\circ}_{si} + \left(B_s + \eta \frac{1+\tau_{si}}{\omega_{si}} G_s \right)^{-1} \mathcal{E}, \tag{53}$$

where we have imposed that B_s and G_s are sector-specific. We proceed through a curve-fitting exercise and find the values of B_s and ηG_s that best fit (53) for the firms in sector s. Firm-level variation in sales and markups allows us to identify the respective roles of B_s and G_s in influencing risk decisions.[44]

Overall, our calibration procedure exactly matches six key data features: 1) sectoral shares of consumption, 2) the firm-level shares of sectoral sales, 3) the sectoral input-output cost shares, 4) the variance of each firm's TFP growth, 5) the covariance of each firm's TFP growth with GDP growth, and finally 6) the variance of GDP growth. Appendix C.5 provides more details about the calibrated economy.

9.3 Increases in aggregate risk

When risk-exposure decisions are endogenous, the economy has more flexibility to handle an increase in risk. To evaluate how important this mechanism is, we conduct an experiment in which the variance Σ of the risk factor doubles. We compare the response of the economy to that change

[43] To fix the scales in some of the figures and tables we set $\rho = 5$ and $\Sigma = 1$. With that normalization, the units of Δ are such that if Δ increases by 1%, the standard deviation of GDP growth also increases by 1%.

[44] Our calibration procedure cannot distinguish between η and G_s, and only provides an estimate for ηG_s. But since these quantities enter the model equations only through that product, there is no need for us to determine η and G_s separately. We therefore remain agnostic about the supply elasticity of risk management resources η.

when firms must keep their previous risk exposure δ fixed, and when they can adjust δ.

When Σ doubles and δ remains fixed, (52) implies that the endogenous component of GDP volatility doubles as well. As Table 1 shows, this leads to a large increase in aggregate risk, and the standard deviation of GDP moves from 2.35% in the calibrated economy to 2.88%. Now that aggregate risk is more important, the value of exposure $\mathcal{E}$ becomes even more negative than in the calibrated economy, and exposure to the aggregate risk factor is particularly unwanted. In the economy in which firms are allowed to adjust their risk exposure, they decide to do so, and Δ falls from 0.012 to 0.009 in response. This decline in Δ leads to a decline in aggregate volatility, and the standard deviation of GDP is reduced to 2.47%. Overall, while the doubling of Σ makes GDP more volatile compared to the calibrated economy, we see that the endogenous response of the firms' risk-taking decisions makes the increase in volatility substantially less severe.

Table 1: Exposure and GDP volatility in different environments

	Calibration	Double Σ		No wedges	
		Fixed δ	Flexible δ	Fixed δ	Flexible δ
Std. Dev. of GDP growth	2.35%	2.88%	2.47%	2.35%	1.80%
Aggregate risk exposure Δ	0.012	0.012	0.009	0.012	0.005
Exposure value $\mathcal{E}$	−0.05	−0.09	−0.07	−0.05	−0.02
Share of endogenous vol.	50%	67%	55%	50%	15%

Notes: "Fixed δ" refers an to an economy in which δ is fixed at its calibrated value. "Share of endogenous volatility" is $2\Sigma\Delta^2/(\mathrm{V}\left[\log \mathrm{GDP}_t - \log \mathrm{GDP}_{t-1}\right])$.

9.4 Wedges and inefficient risk exposure

The presence of wedges pushes firms to make inefficient risk-exposure decisions. To evaluate the quantitative importance of these inefficiencies, we compute the equilibrium in a version of the calibrated economy without wedges. When wedges disappear, firms become larger and their revenue-based Domar weights grow to reach their cost-based counterparts. Keeping δ fixed, this has no impact on GDP volatility, as we can see from the next to last column in Table 1. Indeed, recall from (52) that the volatility of GDP growth only depends on cost-based Domar weights (which are independent of τ), risk exposure δ (which is kept fixed), and the properties of the random variables (which do not depend on τ).

The situation is different when firms can adjust their risk exposure. Since the removal of the wedges makes firms larger, they manage their risk exposure more aggressively and, from (53), this leads to a decline in δ. As a result, aggregate exposure Δ falls from 0.012 to 0.005 (last column of Table 1), and aggregate volatility declines by 70 basis points compared to the calibrated economy. This last exercise shows that while wedges have no impact on volatility in standard economic models

without a risk-taking margin, they can have a large impact on fluctuations when risk is endogenous.

Overall, our findings suggest that disregarding the risk-taking behavior of economic agents can lead to an overestimation of the negative impact of fundamental risk on the economy and to an underestimation of the adverse effect of taxes and markups on aggregate volatility. We interpret these findings with caution. On the one hand, our estimation procedure is parsimonious, and does not impose any parameter values on the risk aversion (ρ), the elasticity of risk-management resources (η) and the underlying variance (Σ). This is reassuring given that these parameters could be hard to estimate. At the same time, the model is stylized in some dimensions, most notably in terms of functional forms and the nature of competition. We therefore do not view the exercises of this section as providing precise quantitative answers about how changes in the environment affect aggregate volatility. We believe, however, that these experiment capture some of the key mechanisms at work in reality, and that they provide a rough estimate of their importance.

We provide several additional results related to the calibrated economy in Appendix C.7. There, we reproduce our counterfactual experiments using different wedge measures. When wedges are large, the equilibrium is far from the efficient allocation and is less able to counteract an increase in Σ. It follows that aggregate volatility increases more when Σ doubles. As expected, reducing wedges to zero also has a stronger impact on volatility when τ is large. We also explore the cross-sectional implications of a change in Σ and of removing wedges in Appendix C.6.

9.5 Disaster risk

Our benchmark model has the advantage of being simple and tractable, but as is well-known, it cannot replicate traditional measures of the importance of aggregate risk, such as the equity premium (Mehra and Prescott, 1985). In this section, we consider an alternative model with disaster risk (Rietz, 1988; Barro, 2006). Namely, we assume that the unique risk factor ε consists of two independent components, $\varepsilon = \varepsilon^n + \varepsilon^d$, where $\varepsilon^n \sim \mathcal{N}(0, \Sigma)$ and the disaster component ε^d takes two values: $\varepsilon^d = -d$ with probability π_d and $\varepsilon^d = 0$ otherwise.[45] We calibrate this economy through the same procedure used above and set π_d and d in line with international evidence (Barro, 2006).[46,47]

Table 2 describes the outcome of various experiments in this economy. It reports the standard deviation of GDP growth associated with normal business cycles ("Conditional on no disasters" columns) and including potential disasters ("Unconditional" columns). We see that doubling Σ has a larger impact in the economy with disasters than in the baseline model. When disasters are possible, the household cares mostly about ε^d, and it sets its exposure Δ accordingly. Since an

[45] Appendix F.7 considers a more general model with disasters and extends our theoretical results to that setting.

[46] Barro (2006) reports that "the probability of entering into a 15 percent or greater event was 1.7 percent per year." Table I in Barro (2006) shows that on average GDP falls by 29% in such events. We therefore set $\pi_d = 0.017$ and $\Delta \times d = 0.29$ in our calibration.

[47] We calibrate the disaster risk model in the same way as the benchmark model, under the assumption that there are no disasters in the observed data. Annual real GDP growth never reached -15% in our sample.

increase in Σ does not affect ε^d, Δ does not respond much, and the economy behaves almost as if Δ were fixed.[48] It follows that regular business cycle volatility increases substantially. Since ε^n only accounts for a small portion of the overall volatility of ε, unconditional volatility responds relatively less. Removing wedges has a large impact in the economy with disasters. Given the destructive nature of ε^d, the social planner is willing to spend massively to reduce its impact, but wedges impede that process. When they are removed, spending on risk mitigation increases, and aggregate volatility, conditional and unconditional, falls substantially.

We also consider the impact of doubling the likelihood of a disaster, so that $\pi_d = 3.4\%$. When firms are not able to change their risk exposure, this leads to a large increase of almost 2 percentage points in unconditional volatility. In contrast, when risk is endogenous, the increase in π_d pushes firms to spend more on risk mitigation, and Δ declines. Unconditional volatility increases by only 56 bps in that case. Since Δ controls exposure to the whole business cycle risk factor ε, this decline in Δ reduces volatility associated with normal business cycles. Through this mechanism, a higher likelihood of disasters can tame run-of-the-mill business cycle volatility.

Because of the presence of rare disasters, this model is able to generate a realistic equity premium of 4.7%, and since risk is endogenous, changes in the environment can influence that premium.[49] For instance, when all wedges are set to zero, the equity premium falls to 0.74%. We interpret this number cautiously. The model is stylized and our estimation of the cost of managing disaster risk is indirect given the available data. Nonetheless, it suggests that taxes and other distortions, through their influence on risk, might have a sizable impact on the return on equity.

Table 2: Standard deviation of GDP growth with and without disaster risk

	Benchmark			With disaster risk					
				Cond. on no disasters			Unconditional		
	Calib.	Fixed δ	Flex. δ	Calib.	Fixed δ	Flex. δ	Calib.	Fixed δ	Flex. δ
Double Σ	2.35%	2.88%	2.47%	2.35%	2.88%	2.86%	5.79%	6.02%	5.97%
No wedges	2.35%	2.35%	1.80%	2.35%	2.35%	1.74%	5.79%	5.79%	2.44%
Double π_d	–	–	–	2.35%	2.35%	2.12%	5.79%	7.78%	6.25%

Notes: "Fixed δ" refers an to an economy in which δ is fixed at its calibrated value. "Cond. on no disasters" refers to the moments in a sample without disasters. "Unconditional" referes to the moments in a sample with disasters.

[48] In this setting, Δ controls exposure to both the normal and disaster shock components. Alternatively, one could consider a model in which firms can separately choose their exposures to ε^n and ε^d. However, given the lack of disasters in our sample, it would be challenging to calibrate the cost of changing disaster exposure.

[49] See Appendix F.7.5 for details. We set the time discount rate to about 0.98 to match a risk free rate of 2% and assume no trend growth in GDP, roughly in line with the Spanish economy since 2008.

10 Conclusion

This paper proposes a parsimonious theory of endogenous risk in which firms are free to choose the properties of their productivity processes. The model is intentionally simple but can explain key features of the data related to how firm characteristics influence their risk profiles. In a basic calibration of the model, we find that changes in the environment can have a large impact on aggregate volatility through their influence on risk-taking decisions.

The model is stylized, and many extensions merit exploration. For instance, we have assumed that markets are complete such that firms value cash flows using the household's stochastic discount factor. However, in a model with entrepreneurs unable to diversify risks related to their businesses, risk management decisions would likely become more consequential, potentially amplifying the model's mechanisms (Greenwood and Jovanovic, 1990).

References

Acemoglu, Daron and Fabrizio Zilibotti (1997). "Was Prometheus unbound by chance? Risk, diversification, and growth". *Journal of Political Economy* 105.4, pp. 709–751.

Acemoglu, Daron et al. (2012). "The Network Origins of Aggregate Fluctuations". *Econometrica* 80.5, pp. 1977–2016.

Ackerberg, Daniel A, Kevin Caves, and Garth Frazer (2015). "Identification properties of recent production function estimators". *Econometrica* 83.6, pp. 2411–2451.

Aghion, Philippe and Peter Howitt (1992). "A Model of Growth Through Creative Destruction". *Econometrica*, pp. 323–351.

Alesina, Alberto et al. (1996). "Political instability and economic growth". *Journal of Economic Growth* 1, pp. 189–211.

Baqaee, David Rezza and Emmanuel Farhi (2019). "Productivity and Misallocation in General Equilibrium". *Quarterly Journal of Economics* 135.1, pp. 105–163.

Barro, Robert J (2006). "Rare disasters and asset markets in the twentieth century". *Quarterly Journal of Economics* 121.3, pp. 823–866.

Beasley, Mark S, Richard Clune, and Dana R Hermanson (2005). "Enterprise risk management: An empirical analysis of factors associated with the extent of implementation". *Journal of accounting and public policy* 24.6, pp. 521–531.

Bigio, Saki and Jennifer La'O (2020). "Distortions in Production Networks". *Quarterly Journal of Economics* 135.4, pp. 2187–2253.

Bloom, Nicholas et al. (2018). "Really Uncertain Business Cycles". *Econometrica* 86.3, pp. 1031–1065.

Blundell, Richard and Stephen Bond (2000). "GMM estimation with persistent panel data: an application to production functions". *Econometric reviews* 19.3, pp. 321–340.

Bodnar, Gordon M, Gregory S Hayt, and Richard C Marston (1998). "1998 Wharton survey of financial risk management by US non-financial firms". *Financial management*, pp. 70–91.

Bond, Steve et al. (2021). "Some unpleasant markup arithmetic: Production function elasticities and their estimation from production data". *Journal of Monetary Economics* 121, pp. 1–14.

Burstein, Ariel, Vasco M Carvalho, and Basile Grassi (2024). "Bottom-up Markup Fluctuations".

Clementi, Gian Luca and Berardino Palazzo (2016). "Entry, exit, firm dynamics, and aggregate fluctuations". *American Economic Journal: Macroeconomics* 8.3, pp. 1–41.

Cole, Harold L, Jeremy Greenwood, and Juan M Sanchez (2016). "Why doesn't technology flow from rich to poor countries?" *Econometrica* 84.4, pp. 1477–1521.

Crouzet, Nicolas and Neil R Mehrotra (2020). "Small and large firms over the business cycle". *American Economic Review* 110.11, pp. 3549–3601.

De Loecker, Jan and Frederic Warzynski (2012). "Markups and firm-level export status". *American Economic Review* 102.6, pp. 2437–2471.

De Ridder, Maarten, Basile Grassi, Giovanni Morzenti, et al. (2024). "The Hitchhiker's Guide to Markup Estimation".

Fama, Eugene F and Kenneth R French (1992). "The cross-section of expected stock returns". *Journal of Finance* 47.2, pp. 427–465.

Gabaix, Xavier (2011). "The Granular Origins of Aggregate Fluctuations". *Econometrica* 79.3, pp. 733–772.

Géczy, Christopher, Bernadette A Minton, and Catherine Schrand (1997). "Why firms use currency derivatives". *the Journal of Finance* 52.4, pp. 1323–1354.

Greenwood, Jeremy and Boyan Jovanovic (1990). "Financial development, growth, and the distribution of income". *Journal of political Economy* 98.5, Part 1, pp. 1076–1107.

Grossman, Gene M and Elhanan Helpman (1991). *Innovation and growth in the global economy*. MIT press.

Hoyt, Robert E and Andre P Liebenberg (2011). "The value of enterprise risk management". *Journal of risk and insurance* 78.4, pp. 795–822.

Hulten, C. R. (1978). "Growth Accounting with Intermediate Inputs". *Review of Economic Studies* 45.3, pp. 511–518.

Jensen, Michael C and William H Meckling (1976). "Theory of the firm: Managerial behavior, agency costs and ownership structure". *Journal of Financial Economics* 3.4, pp. 305–360.

Jones, Charles I (1995). "R&D-based models of economic growth". *Journal of Political Economy* 103.4, pp. 759–784.

— (2011). "Intermediate Goods and Weak Links in the Theory of Economic Development". *American Economic Journal: Macroeconomics* 3.2, pp. 1–28.

Kelly, Bryan T, Seth Pruitt, and Yinan Su (2020). "Instrumented principal component analysis". *Available at SSRN 2983919*.

Khan, Aubhik and Julia K Thomas (2008). "Idiosyncratic shocks and the role of nonconvexities in plant and aggregate investment dynamics". *Econometrica* 76.2, pp. 395–436.

King, Robert G, Charles I Plosser, and Sergio T Rebelo (1988). "Production, growth and business cycles: I. The basic neoclassical model". *Journal of Monetary Economics* 21.2-3, pp. 195–232.

Klette, Tor Jakob and Samuel Kortum (2004). "Innovating firms and aggregate innovation". *Journal of Political Economy* 112.5, pp. 986–1018.

Kopytov, Alexandr et al. (2024). "Endogenous production networks under supply chain uncertainty". *Econometrica* 92.5, pp. 1621–1659.

Koren, Miklós and Silvana Tenreyro (2007). "Volatility and development". *Quarterly Journal of Economics* 122.1, pp. 243–287.

— (2013). "Technological diversification". *American Economic Review* 103.1, pp. 378–414.

Kydland, Finn E. and Edward C. Prescott (1982). "Time to Build and Aggregate Fluctuations". *Econometrica* 50.6, pp. 1345–1370.

Levinsohn, James and Amil Petrin (2003). "Estimating production functions using inputs to control for unobservables". *Review of Economic Studies* 70.2, pp. 317–341.

Liebenberg, Andre P and Robert E Hoyt (2003). "The determinants of enterprise risk management: Evidence from the appointment of chief risk officers". *Risk management and insurance review* 6.1, pp. 37–52.

Liu, Ernest (2019). "Industrial Policies in Production Networks". *Quarterly Journal of Economics* 134.4, pp. 1883–1948.

Long, John B. and Charles I. Plosser (1983). "Real Business Cycles". *Journal of Political Economy* 91.1, pp. 39–69.

Lucas Jr, Robert E (1978). "Asset prices in an exchange economy". *Econometrica*, pp. 1429–1445.

Mehra, Rajnish and Edward C Prescott (1985). "The equity premium: A puzzle". *Journal of monetary Economics* 15.2, pp. 145–161.

Pellet, Thomas and Alireza Tahbaz-Salehi (2023). "Rigid production networks". *Journal of Monetary Economics* 137, pp. 86–102.

Ramey, Garey and Valerie A Ramey (1995). "Cross-Country Evidence on the Link Between Volatility and Growth". *American Economic Review* 85.5, pp. 1138–1151.

Raval, Devesh (2023). "Testing the production approach to markup estimation". *Review of Economic Studies* 90.5, pp. 2592–2611.

Rietz, Thomas A (1988). "The equity risk premium a solution". *Journal of monetary Economics* 22.1, pp. 117–131.

Romer, Paul M (1990). "Endogenous technological change". *Journal of Political Economy* 98.5, Part 2, S71–S102.

Ross, Stephen A (1977). "The determination of financial structure: the incentive-signalling approach". *The Bell Journal of Economics*, pp. 23–40.

Rovigatti, Gabriele and Vincenzo Mollisi (2018). "Theory and practice of total-factor productivity estimation: The control function approach using Stata". *The Stata Journal* 18.3, pp. 618–662.

Stock, James H and Mark W Watson (2002). "Forecasting using principal components from a large number of predictors". *Journal of the American statistical association* 97.460, pp. 1167–1179.

Yeh, Chen (2023). "Revisiting the origins of business cycles with the size-variance relationship". *Review of Economics and Statistics*, pp. 1–28.

Online Appendix

A Derivation of the stochastic discount factor

In this appendix, we derive equations (8) and (9). The household chooses how much to consume after uncertainty about ε is realized but how much risk-management resource to supply before uncertainty is realized. We first consider the problem of the household once the state ε has been realized. Its Lagrangian is

$$\mathscr{L} = \frac{\left(\left(\frac{C_1}{\beta_1}\right)^{\beta_1} \times \cdots \times \left(\frac{C_N}{\beta_N}\right)^{\beta_N}\right)^{1-\rho} (V(R))^{1-\rho}}{1-\rho} - \Lambda\left(\sum_{i=1}^{N} P_i C_i - W_L L - W_R R - \Pi\right),$$

where we assume that the household supplies $L = 1$ units of production labor inelastically. The first-order condition with respect to C_i is

$$\beta_i \mathcal{U}'(Y)\, Y\, (V(R))^{1-\rho} = \Lambda P_i C_i,$$

where $\mathcal{U}(Y) = \frac{Y^{1-\rho}}{1-\rho}$. Summing over i on both sides of this equation and using the binding budget constraint, we find

$$\mathcal{U}'(Y)\, Y\, (V(R))^{1-\rho} = \Lambda\left(W_L L + W_R R + \Pi\right). \tag{54}$$

Combining with the first-order condition implies

$$P_i C_i = \beta_i\left(W_L L + W_R R + \Pi\right). \tag{55}$$

Combining the first-order condition with $Y = \prod_{i=1}^{N}\left(\beta_i^{-1} C_i\right)^{\beta_i}$ yields

$$\begin{aligned} Y = \prod_{i=1}^{N}\left(\beta_i^{-1} C_i\right)^{\beta_i} = \prod_{i=1}^{N}\left(\beta_i^{-1}\frac{\beta_i \mathcal{U}'(Y)\, Y\, (V(R))^{1-\rho}}{\Lambda P_i}\right)^{\beta_i} \Leftrightarrow \\ \Lambda = \mathcal{U}'(Y)\,(V(R))^{1-\rho}\prod_{i=1}^{N} P_i^{-\beta_i}, \end{aligned} \tag{56}$$

which we can combine with (54) to find

$$Y = \left(W_L L + W_R R + \Pi\right)\prod_{i=1}^{N} P_i^{-\beta_i}. \tag{57}$$

This last equation implicitly defines a price index,

$$\bar{P} = \prod_{i=1}^{N} P_i^{\beta_i}, \tag{58}$$

such that $\bar{P}Y = W_L L + W_R R + \Pi$. Therefore, (56) is equivalent to (8).

Finally, we can also compute the household's first-order condition with respect to R:

$$\mathrm{E}\left[Y^{1-\rho}\right] (V(R))^{-\rho} V'(R) + \mathrm{E}[\Lambda] W_R = 0. \tag{59}$$

Using (56), we find

$$\mathrm{E}\left[Y^{1-\rho}\right] \frac{V'(R)}{V(R)} + \mathrm{E}\left[\frac{Y^{1-\rho}}{\bar{P}Y}\right] W_R = 0. \tag{60}$$

By Lemma 1, $\bar{P}Y = W_L L \Gamma_L^{-1}$. Given our choice of numeraire, $W_L L = 1$ and so $\bar{P}Y = \Gamma_L^{-1}$, which is non-stochastic by Lemma 1. Therefore, (60) simplifies to (9).

B Appendix for Section 8

This appendix contains details about the reduced-form exercises of Section 8.

B.1 Derivation of (48)

From (17), we can write log real GDP in the model of Section 8 as

$$y = \Delta\varepsilon + \tilde{\omega}^\top (-b(\delta) + v + \gamma t) - \tilde{\omega}^\top \log(1+\tau) - \log \Gamma_L,$$

where v is the column vector of the firm-level productivity shocks, and γ is the vector of the firm-level growth trends γ_i. The fictitious planner's problem is therefore

$$\begin{aligned}\mathcal{W}_{dist} := \max_{\Delta} & \underbrace{\Delta \times 0 - \bar{b}(\Delta) + \tilde{\omega}^\top (\mu^v + \gamma t) - \tilde{\omega}^\top \log(1+\tau) - \log \Gamma_L}_{\mathrm{E}[y]} \\ & - \frac{1}{2}(\rho - 1) \underbrace{\left(\Sigma \Delta^2 + \tilde{\omega}^\top \Sigma^v \tilde{\omega}\right)}_{\mathrm{V}[y]} - \bar{g}(\Delta),\end{aligned}$$

where $\mathrm{E}[\varepsilon] = 0$, μ^v is the expected value of v, and Σ^v is the covariance matrix of v. Notice that the only non-stationary term, the growth trend vector γt, does not interact with the choice of Δ, and so Δ is constant over time. Consequently, δ is also constant over time as it solves (31).

Next, the TFP process for firm i is given by $\log TFP_{it} = \delta_i \varepsilon_t - b_i(\delta_i) + \gamma_i t + v_{it}$, where $v_{it} \sim$

iid $\mathcal{N}(\mu_i^v, \Sigma_i^v)$ and $\varepsilon_t \sim$ iid $\mathcal{N}(0, \Sigma)$. It follows that

$$\begin{aligned}
\mathrm{V}\left[\log TFP_{it} - \log TFP_{it-1}\right] &= \mathrm{V}\left[\delta_i \varepsilon_t - b_i(\delta_i) + v_{it} + \gamma_i t - \delta_i \varepsilon_{t-1} - v_{it-1} - \gamma_i (t-1)\right] \\
&= \mathrm{V}\left[\delta_i \varepsilon_t + v_{it} - \delta_i \varepsilon_{t-1} - v_{it-1}\right] \\
&= \mathrm{V}\left[\delta_i (\varepsilon_t - \varepsilon_{t-1}) + v_{it} - v_{it-1}\right] \\
&= 2\delta_i^2 \Sigma + 2\Sigma_i^v.
\end{aligned}$$

Similarly, the covariance between aggregate output growth and firm-level TFP growth is

$$\begin{aligned}
\mathrm{Cov}\left[y_t - y_{t-1}, \log TFP_{it} - \log TFP_{it-1}\right] &= \mathrm{Cov}\left[\Delta(\varepsilon_t - \varepsilon_{t-1}) + \tilde{\omega}^\top (v_t - v_{t-1}), \delta_i (\varepsilon_t - \varepsilon_{t-1}) + v_{it} - v_{it-1}\right] \\
&= 2\Delta\Sigma\delta_i + 2\tilde{\omega}_i \Sigma_i^v.
\end{aligned}$$

These are the equations reported in Section 8.

B.2 Data sources and variable construction

B.2.1 Orbis data

Our data of Spanish firms comes from the Orbis Historical Disk Product. The Orbis data is known to be the largest cross-country firm-level database that covers public and private firms' financial and real activities. We choose Spain for our analysis due to its near-universal coverage of firms (covering more than 95% of total industry gross output after 2010) as detailed in Gopinath et al. (2017) and Kalemli-Özcan et al. (2024). We use the sample period 1995-2019 for our analysis.[50]

Sample cleaning Following the procedure of Kalemli-Özcan et al. (2024), we link each firm's descriptive information with its financial information via the unique BVD firm identifier (`BVDID`). We restrict our analysis to Spanish firms, defined as firms that satisfy two criteria: 1) their latest address is in Spain and 2) their `BVDID` starts with the ISO-2 code `ES`. Within the Orbis Spanish sample, we apply the following standard cleaning procedure:

1. We harmonize the calendar year of each firm-year observation using the variable `closing_date`: if the closing date is after or on July 1, the current year is assigned as the calendar year. Otherwise, the previous year is assigned.[51]

2. In a given year, a firm in the Orbis database might have multiple financial statements from different sources (local registry, annual report, or others), for consolidated or unconsolidated

[50] Orbis offers good coverage of the Spanish economy starting from 1995. Moreover, the most recent observations in the version of Orbis Historical Disk Product that we use are from 2021. We therefore use 2019 as the last year of the sample since there is usually a two-year reporting lag for some variables (see Kalemli-Özcan et al. (2024) for details).

[51] This adjustment matters little for the Spanish sample, as 99% of firms close their books on December 31.

accounts. When several source-consolidation combinations exist for a firm, we deduplicate by selecting the account that, in order of priority, 1) shows the most consistent reporting frequency (closest to regular annual reporting), 2) offers the longest non-missing time series for key financial variables (fixed assets and/or sales), and 3) is consolidated, if the first two criteria are tied.

3. We only keep firms with non-missing and positive sales (`operating_revenue_turnover`), fixed assets (`fixed_assets`), wage bills (`costs_of_employees`), and material costs (`material_costs`). We also harmonize the units of all monetary values to be in current EUROs.

4. To prevent outliers from affecting the production function and markup estimation, we exclude any firm-year observation whose average revenue product for any input (fixed assets, wage bills, or material costs) lies above the 99th percentile or below the 1st percentile of that year's distribution. Furthermore, over the full sample period, we sequentially remove the top 0.1% of observations in sales, fixed assets, wage bills, and material costs (in that order).

Variable definitions and estimation of firm-level markup and TFP We calculate a firm i's revenue-based Domar weight as $\omega_{it} = \frac{\text{sales}_{it}}{\text{GDP}_{\text{nom},t}}$, where $\text{GDP}_{\text{nom},t}$ is the Spanish nominal GDP (in euros) obtained from the Annual Spanish National Accounts produced by the National Statistics Institute (INE).

We use the production function estimation approach to obtain estimates of firm-level markup and productivity growth. To implement the estimation procedure, we assume a Cobb-Douglas production function of the following form:

$$\log Q_{it} = \alpha_{Li} \log L_{it} + \alpha_{Mi} \log M_{it} + \alpha_{Ki} \log K_{it} + \varepsilon_{it}, \tag{61}$$

where Q_{it}, L_{it}, M_{it}, and K_{it} are the deflated values of sales, the wage bill, material costs, and fixed assets for each firm i in calendar year t. Sales and material costs are deflated using the industry-specific gross output price indices. We deflate fixed assets using the capital price indices and the wage bill using the GDP deflator. The industry-specific price indices are from the EU-KLEMS dataset, and we use the most detailed sector-level price indices at the NACE 2-digit level whenever available. The output elasticities are estimated using the Levinsohn and Petrin (2003) methodology with the Ackerberg, Caves, and Frazer (2015) (hereafter ACF) correction. Following Levinsohn and Petrin (2003), we use capital as the "state" variable, labor as the "free" variable and materials as the "proxy" variable. We estimate the production function for each NACE 2-digit sector. As in De Loecker, Eeckhout, and Unger (2020), we control for markups using firms' sales shares (sales share at the NACE 3-digit and 4-digit industry level) in the production function estimation. Rovigatti

and Mollisi (2018) show that output-elasticity estimates from the ACF correction are sensitive to the optimizer's starting values. To mitigate this potential issue, we run 3,000 sector-specific second-stage GMM estimations with different initial values and retain the solution with the lowest objective function as our estimate. We provide a detailed discussion in Appendix B.3.5.

Following De Loecker and Warzynski (2012), we compute the markup as $1 + \tau_{it} = \hat{\alpha}_{Li} \Big/ \left(\frac{\text{Wage Bill}_{it}}{\text{Sales}_{it}} \right)$, where $\hat{\alpha}_{Li}$ is the estimated labor elasticity in production and $\frac{\text{Wage Bill}_{it}}{\text{Sales}_{it}}$ is the share of labor expenditure in firms' sales. Our baseline markup measure regards labor as a variable input in production.

To translate our production function estimates into productivity growth, we use an adjusted Solow residual of the following form,[52]

$$\begin{aligned} \Delta \log \text{TFP}_{it} = & \Delta \log Q_{it} - \alpha_{Li} \Delta \log L_{it} - \alpha_{Mi} \Delta \log M_{it} - \alpha_{Ki} \Delta \log K_{it} \\ & - \left(\Delta \log \left(1 + \tau_{it}\right) - \Delta \log \left(1 + \tau_{s(i)t}\right) \right). \end{aligned} \tag{62}$$

The first line is the standard Solow residual, where (deflated) output growth is adjusted by the contribution of (deflated) input growth. We further adjust our measure by $\Delta \log \left(1 + \tau_{it}\right) - \Delta \log \left(1 + \tau_{s(i)t}\right)$, which accounts for the firm-specific markup growth net of the sectoral markup growth, where $s(i)$ represents the NACE 2-digit sector s to which firm i belongs. This adjustment allows us to remove the changes in firm-specific nominal price that are not taken into account by the sector-level price deflator.[53] We thus move closer to a quantity-based interpretation of our TFP measure.

For all reduced-form exercises that involve panel regressions, we restrict our sample to firms with at least two non-missing observations of measured TFP growth ($\Delta \log \text{TFP}_{it}$). Additionally, for all cross-sectional regressions, we only keep firms with at least 5 non-missing observations of TFP growth (such that TFP volatility can be precisely estimated) and positive average log markup, meaning that they do not constantly operate at a loss.

B.2.2 CRSP/Compustat Merged and US stock market data

This subsection describes the firm-level financial and stock market variables used in Appendix B.4 below. We use two main sources of data: 1) the CRSP/Compustat Merged Fundamentals Annual data that allows us to compute firm-level Domar weights and markups, and 2) the WRDS Beta Suite database that provides stocks' loading on the aggregate return, i.e., stock market betas.

[52] To avoid outliers in TFP growth, we drop the top and bottom 0.5% firm-year observations for output growth, (any of the three) input growth or markup growth for each sector.

[53] The sectoral markup growth is constructed as the revenue-weighted firm-level markup growth, which is consistent with the Cobb-Douglas sectoral price aggregator in our quantitative model.

CRSP/Compustat Merged Our sample selection and cleaning procedure follows standard practices in the literature, particularly those carried out in De Loecker, Eeckhout, and Unger (2020).

1. We apply the following standard filters when processing the data: 1) consolidation level is `C`; 2) industry format is `INDL`; 3) data format is `STD`; 4) population source is `D`; 5) currency is `USD`; 6) we include both active and inactive companies; 7) we exclude financial, utilities and public sector firms (i.e. those with SIC code in the following range: $4900-4999$, $6000-6999$ and $9000-9999$). We use the data from 1980 to 2016.

2. We drop firm-year observations for which firms are not incorporated (`fic`) or headquartered (`loc`) in the USA.

3. We use PERMNO as firm identifier and drop the firm-year observations with duplicates.

4. We only keep firms with non-missing and positive sales (`sale`), gross book value of fixed assets (`ppegt`), employees (`emp`), and cost of goods sold (`cogs`).

5. We drop firm-year observations with COGS-to-sales ratio in the top and bottom 1% of the corresponding year's distributions.

WRDS Beta Suite We estimate each stock's equity beta in each year using the Beta Suite by WRDS. We use PERMNO as our firm identifier. For each firm-month pair, we estimate the CAPM market model using 5-year backwards-looking rolling window of monthly firm and aggregate stock market returns. Specifically, the equity beta β_{it} for each firm-month pair is obtained from the time series regression

$$r_{it} - rf_t = \alpha_i + \beta_{it} mktrf_t + \varepsilon_{it},$$

where r_{it} is the firm i's stock return, $mktrf_t$ is the Fama-French Excess Return on the Market, and rf_t is the risk-free rate during month t. The estimation procedure is internally executed inside the Beta Suite with the following options: 1) frequency is `monthly`; 2) both the estimation window and the minimum window are 60 months; 3) risk model is `Market Model`; 4) return type is `Regular Return`. After the estimation, we only keep the beta estimates in December to be the firm's equity beta for the calendar year.

Variable definitions and estimation of firm-level markup and TFP The variable construction procedure for the firm-level data from Compustat is similar to that for Spanish firms from Orbis. We define firm i's revenue-based Domar weight as $\omega_{it} = \frac{\text{sales}_{it}}{\text{GDP}_{\text{nom},t}}$, where $\text{GDP}_{\text{nom},t}$ is US nominal GDP (the `GDPA` series from the Bureau of Economic Analysis).

Instead of estimating the production functions ourselves, we use the time-varying production function estimates from De Loecker, Eeckhout, and Unger (2020), who consider a Cobb-Douglas production function of the following form:

$$\log Q_{it} = \alpha_{Vit} \log V_{it} + \alpha_{Kit} \log K_{it} + \varepsilon_{it}, \tag{63}$$

where Q_{it}, V_{it}, and K_{it} are deflated values of, respectively, sales, cost of goods sold, and fixed assets for each firm i in calendar year t.[54] To be consistent with the estimates in De Loecker, Eeckhout, and Unger (2020), all variables are deflated using the GDP deflator.

De Loecker, Eeckhout, and Unger (2020) estimate output elasticities using the Olley and Pakes (1996) methodology with the ACF correction for each 2-digit NAICS sector and each year. As panel data are required in the Olley and Pakes (1996) approach, they use 5-year rolling windows such that the elasticity estimates for year t are obtained using data from year $t-2$ to year $t+2$. In the estimation, they also control for markups using measures of market share at various levels of aggregation.

We merge the estimates of $\hat{\alpha}_{Vit}$ and $\hat{\alpha}_{Kit}$ from De Loecker, Eeckhout, and Unger (2020)'s replication package with the Compustat data. Treating cost of goods sold (COGS) as a variable input in production, we compute the markup as $1 + \tau_{it} = \hat{\alpha}_{Vit} \Big/ \left(\frac{\text{Costs of Goods Sold}_{it}}{\text{Sales}_{it}} \right)$, where $\hat{\alpha}_{Vit}$ is the estimated variable input elasticity in production in year t and $\frac{\text{Costs of Goods Sold}_{it}}{\text{Sales}_{it}}$ is the share of variable input expenditure in firms' sales. We compute firm-level TFP growth as

$$\begin{aligned} \Delta \log \text{TFP}_{it} = \Delta \log Q_{it} - \frac{1}{2} \left(\alpha_{Vit} + \alpha_{Vit-1} \right) \Delta \log V_{it} - \frac{1}{2} \left(\alpha_{Kit} + \alpha_{Kit-1} \right) \Delta \log K_{it} \\ - \left(\Delta \log \left(1 + \tau_{it} \right) - \Delta \log \left(1 + \tau_{s(i)t} \right) \right), \end{aligned}$$

where $\Delta \log(1 + \tau_{s(i)t})$ is the sector-level markup growth and $s(i)$ represents the NAICS 2-digit sector that firm i belongs to. For all reduced-form exercises that involve panel regressions and equity beta in Section B.4, we only consider firms with at least two observations of measured TFP growth. Moreover, for the cross-sectional analysis, we only keep the firms with at least 5 non-missing observations of TFP growth and positive average log markup, following the procedure used for the Orbis sample.

Constructing NYSE breakpoints for the deciles of Domar weights We now discuss how we construct the NYSE breakpoints for the deciles of Domar weights. Our approach resembles the procedure used in Fama and French (1992) to construct the NYSE breakpoints for market equity (ME), but instead is applied to Domar weights in our annual CRSP/Compustat Merged sample. We make the following restrictions to the sample: 1) we exclude all firms from the financial, utilities, and public sectors along with those with invalid annual sales data, 2) we drop all observations that do not have a matched 5-year CAPM beta at the PERMNO-year level, and (3) we only consider firms with a stock exchange code (`exchg`) of 11 such that we only keep the firms that are traded in

[54] De Loecker, Eeckhout, and Unger (2020) use "costs of goods sold" as a composite variable input because material costs and wage bills are not separately reported in Compustat.

the NYSE.

After restricting the sample, we compute the 10th, 20th, 30th, 40th, 50th, 60th, 70th, 80th, and 90th percentiles of Domar weights for each year. These annual percentile values are then used to form the Domar weight deciles in the main sample.

B.3 Additional results

This section contains additional results related to the reduced-form exercises of Section 8.

B.3.1 Volatility of TFP growth

To complement the results reported in Figure 6, we run a cross-sectional regression of the firm-level standard deviation of TFP growth on average log Domar weight and average log markup while controlling for sector fixed effects,

$$\sigma_i\left(\Delta \log TFP_{it}\right) = \alpha + \beta^{Domar}\overline{\log\left(\text{Domar Weight}\right)}_i + \beta^{Markup}\overline{\log\left(\text{Markup}\right)}_i + \theta_s + \varepsilon_i. \tag{64}$$

The results are given in Table 3. Consistent with our theory, we find $\beta^{Domar} < 0$ and $\beta^{Markup} > 0$ for both the Orbis Spain sample and the Compustat sample.

Table 3: Volatility of TFP growth, Domar weights and markups

	Dependent variable: Volatility of firm-level TFP growth	
	(1): Orbis Spain	(2): Compustat
$\overline{\log\left(\text{Domar Weight}\right)}_i$	-0.037^{***}	-0.004^{***}
	(0.002)	(0.000)
$\overline{\log\left(\text{Markup}\right)}_i$	0.026^{***}	0.002^{*}
	(0.009)	(0.001)
Sector FE	Yes	Yes
Observations	367,213	5,885
R^2	0.383	0.594

Notes: The table presents estimation results of (64). The sample includes all Orbis Spain (column 1) and Compustat (column 2) firms. See text for a detailed description of variable construction and sample selection. The Orbis Spain sample is winsorized at the top and bottom 0.5% of observations of TFP growth, Domar weight and markup in each sector before computing firm-level moments. The Compustat sample is winsorized at 1%. Standard errors (in parentheses) are clustered at the NACE 2-digit industry level for the Orbis Spain sample in column (1), and are heteroskedasticity-robust for the Compustat sample in column (2). *,** ,*** indicate significance at the 10%, 5%, and 1% levels, respectively.

B.3.2 Sensitivity of firm-level TFP to GDP

To complement the results reported in Figure 7, we run a panel regression to examine whether firms with different Domar weights and markups display different comovement patterns between

firm-level TFP growth and aggregate GDP growth:

$$\Delta \log TFP_{it} = \beta_1 \log\left(\text{Domar Weight}_{it}\right) \times \Delta \log GDP_t + \beta_2 \log\left(\text{Markup}_{it}\right) \times \Delta \log GDP_t + \alpha + \beta_0 \Delta \log GDP_t + \beta_3 \log\left(\text{Domar Weight}_{it}\right) + \beta_4 \log\left(\text{Markup}_{it}\right) + \text{FE} + \varepsilon_{it}. \quad (65)$$

The coefficients of interest are β_1 and β_2, and we expect to find $\beta_1 < 0$ and $\beta_2 > 0$. Column (1) of Table 4 shows the results of estimating (65) with sector fixed effects. In column (2), we add sector-by-year fixed effects; in column (3), we use firm and year fixed effects and find the same patterns even within firms; and in column (4), we use firm and sector-year fixed effects. In all specifications, we find statistically significant coefficients whose signs are consistent with our theory.

Table 4: Sensitivity of firm-level TFP to GDP

	Dependent variable: Firm-level TFP growth			
	(1)	(2)	(3)	(4)
$\log\left(\text{Domar Weight}_{it}\right) \times \Delta \log GDP_t$	-0.059^{***}	-0.122^{***}	-0.099^{***}	-0.143^{***}
	(0.003)	(0.003)	(0.003)	(0.004)
$\log\left(\text{Markup}_{it}\right) \times \Delta \log GDP_t$	0.135^{***}	0.225^{***}	0.050^{***}	0.174^{***}
	(0.006)	(0.008)	(0.007)	(0.009)
Sector FE	Yes	No	No	No
Firm FE	No	No	Yes	Yes
Sector-Year FE	No	Yes	No	Yes
Year FE	No	No	Yes	No
Observations	7,190,624	7,190,624	7,190,624	7,190,624
R^2	0.029	0.037	0.207	0.214

Notes: Table presents estimation results of (65) using a sample of Spanish firms from Orbis. Sample construction is described in Appendix B.2. The estimation sample is winsorized at the top and bottom 0.5% of observations of TFP growth, Domar weight and markup in each sector. Coefficients on levels of GDP growth, (log) Domar weight and (log) markup are not reported for brevity. Standard errors (in parentheses) are clustered at the firm level. *,** ,*** indicate significance at the 10%, 5%, and 1% levels, respectively.

B.3.3 Robustness with alternative markup measures

Some of our results rely on markups estimated using the production function approach with revenue data (De Loecker and Warzynski, 2012). Some authors have raised concerns about that approach. In this section, we first argue that these concerns do not apply to our reduced-form exercises. We also show that our results are robust to alternative markup measures and estimation procedures.

Bond et al. (2021) have argued that markups estimated using the production function approach with revenue data could be biased. In response to that critique, De Ridder, Grassi, Morzenti, et al. (2024) use firm-level price and quantity data from France to estimate markups

directly. They find that the production approach performs well when looking at the *cross-section* of markups across firms, which is the variation that we rely on in all of our reduced-form exercises. To further address the Bond et al. (2021) critique, we follow Baqaee and Farhi (2019) and Burstein, Carvalho, and Grassi (2024) and consider profit accounting markups based on the Lerner index (Lerner, 1934). Namely, we compute accounting markups as the ratio of a firm's sales to total cost (measured as sales net of operating profits) from accounting data, that is $1+\tau_{it}^{Lerner} = \frac{1}{1-Lerner_{it}} = \frac{\text{Sales}_{it}}{\text{Sales}_{it}-\text{Operating Profits}_{it}}$.[55] Results from estimating (64) and (65) with those accounting markups are reported in columns (3) and (6) of Table 5.

Raval (2023) also shows that when using the production function approach, the estimated markups could differ based on whether labor or materials are treated as the flexible input. We therefore repeat our analysis by treating materials as the flexible factor. We compute this materials-based markup as $1+\tau_{it}^{M} = \hat{\alpha}_{Mi} \Big/ \left(\frac{\text{Material Costs}_{it}}{\text{Sales}_{it}}\right)$, where $\hat{\alpha}_{Mi}$ are the estimated output elasticities of materials obtained from the same set of production function estimates. Results from estimating (64) and (65) with this alternative markup measure are reported in columns (2) and (5) of Table 5, alongside our benchmark estimates in columns (1) and (4).

The estimates reported in Table 5 confirm that our results are qualitatively similar and statistically significant across all specifications, regardless of the markup measure used. As in our benchmark regressions, smaller firms and those with higher markups exhibit greater TFP volatility and greater sensitivity of firm-level TFP to GDP, in line with our theoretical model.

B.3.4 Robustness with the Blundell and Bond (2000) production function estimator

While the control function approach to production function estimation can potentially address measurement errors in output, it hinges on specific assumptions about the timing of input choices. We therefore conduct robustness checks by employing an alternative IV-GMM (Blundell and Bond, 2000) production function estimator as in De Ridder, Grassi, Morzenti, et al. (2024). Specifically, we follow the specification of Table III, column (v) of Blundell and Bond (2000) by using the `xtabond2` command in Stata. We use lagged levels dated $t-2$ and earlier as instruments for the equations in first differences and the lagged first-differences dated $t-1$ as instruments for the equations in levels. Results from estimating (64) and (65) with the estimates of markup and productivity obtained with the Blundell and Bond (2000) estimator are reported in columns (2) and (4) (labeled as "BB") of Table 6, alongside those obtained with our benchmark control function estimates in columns (1) and (3) (labeled as "LP + ACF"). We find that the estimates from the two production function estimators are quantitatively very similar.

[55] We further compute a consistently defined measure of firm-level TFP growth using (62).

Table 5: Results with alternative markup measures

Dependent variable	Firm-level TFP growth			Vol. of firm-level TFP growth		
	(1)	(2)	(3)	(4)	(5)	(6)
Markup estimate	Labor	Materials	Lerner	Labor	Materials	Lerner
$\Delta \log GDP_t$	-0.330^{***}	0.004	0.153^{***}			
	(0.041)	(0.049)	(0.030)			
$\Delta \log GDP_t$	-0.059^{***}	-0.010^{***}	-0.012^{***}			
$\times \log(\text{Domar Weight}_{it})$	(0.003)	(0.004)	(0.002)			
$\Delta \log GDP_t$	0.135^{***}	0.053^{***}	0.547^{***}			
$\times \log(\text{Markup}_{it})$	(0.006)	(0.010)	(0.032)			
$\log(\text{Domar Weight}_i)$				-0.037^{***}	-0.031^{***}	-0.021^{***}
				(0.002)	(0.004)	(0.001)
$\log(\text{Markup}_i)$				0.026^{***}	0.061^{***}	0.288^{***}
				(0.009)	(0.007)	(0.059)
Sector FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	7,190,624	7,158,659	7,189,677	367,213	331,454	390,839
R^2	0.029	0.056	0.006	0.383	0.558	0.336

Notes: Table presents estimation results of (65) (columns (1), (2) and (3)) and (64) (columns (4), (5) and (6)) using a sample of Spanish firms from Orbis. See text for a detailed description of variable construction and sample selection. The estimation samples for columns (1), (2) and (3) are winsorized at the top and bottom 0.5% of observations of TFP growth, Domar weight and markup in each sector. The estimation samples for columns (4), (5) and (6) are winsorized at the top and bottom 0.5% of observations of TFP growth, Domar weight and markup in the panel in each sector before computing the firm-level moments; it only includes firms that have 1) at least 5 non-missing observations of TFP growth, and 2) positive average log markup. In columns (1) and (4), firms' markups and TFP growth are computed using wage bill-based markups. In columns (2) and (5), firms' markups and TFP growth are computed using material-based markups. In columns (3) and (6), firms' markups and TFP growth are computed using accounting markups based on the Lerner index. Coefficients on levels of (log) Domar weight and (log) markup are not reported for brevity in columns (1), (2) and (3). Standard errors (in parentheses) are clustered at the firm level for the panel regressions in columns (1), (2) and (3) and are clustered at the NACE 2-digit industry level for the cross-sectional regressions in columns (4), (5) and (6). *,** ,*** indicate significance at the 10%, 5%, and 1% levels, respectively.

B.3.5 Robustness related to numerical challenges from the Ackerberg, Caves, and Frazer (2015) correction

As discussed in Appendix B.2.1, our baseline results rely on output elasticities estimated using the Levinsohn and Petrin (2003) methodology with the ACF correction. While the ACF correction is essential for the correct identification of the labor coefficient, Rovigatti and Mollisi (2018) show that the output elasticity estimates it provides are quite sensitive to the starting points passed to the underlying optimization algorithm. Specifically, their Monte Carlo simulations show that if the starting points in the second-stage GMM deviate from the true parameters, many commonly used optimization algorithms often converge to a local minimum and thus produce biased estimates.

To mitigate this issue, we run 3,000 iterations of the second-stage GMM with varied initial values. Each set of starting points is generated by adding random noise with a standard deviation

Table 6: Results with alternative production function estimator

Dependent variable	Firm-level TFP growth		Vol. of firm-level TFP growth	
	(1)	(2)	(3)	(4)
Production function estimator	LP + ACF	BB	LP + ACF	BB
$\Delta \log GDP_t$	-0.330^{***}	-0.544^{***}		
	(0.041)	(0.041)		
$\Delta \log GDP_t$	-0.059^{***}	-0.060^{***}		
$\times \log(\text{Domar Weight}_{it})$	(0.003)	(0.003)		
$\Delta \log GDP_t$	0.135^{***}	0.189^{***}		
$\times \log(\text{Markup}_{it})$	(0.006)	(0.006)		
$\log(\text{Domar Weight}_i)$			-0.037^{***}	-0.035^{***}
			(0.002)	(0.002)
$\log(\text{Markup}_i)$			0.026^{***}	0.023^{**}
			(0.009)	(0.009)
Sector FE	Yes	Yes	Yes	Yes
Observations	7,190,624	7,222,136	367,213	507,534
R^2	0.029	0.034	0.383	0.323

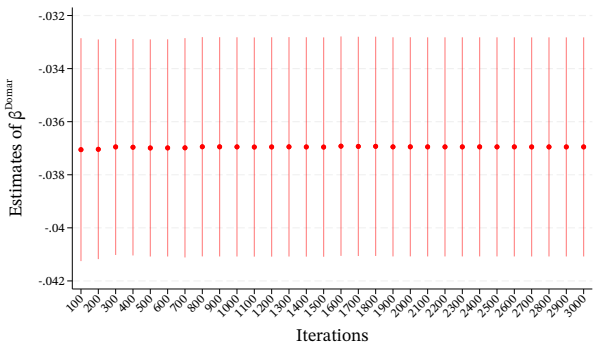
Notes: Table presents the estimation results of (65) (columns (1) and (2)) and (64) (columns (3) and (4)) using a sample of Spanish firms from Orbis. See text for a detailed description of variable construction and sample selection. In columns (1) and (3), a firm's markup and TFP growth are computed using wage bill-based markups along with production function estimates from the Levinsohn and Petrin (2003) estimator with ACF correction. In columns (2) and (4), a firm's markup and TFP growth are computed using wage bill-based markups along with production function estimates from the Blundell and Bond (2000) estimator. Standard errors (in parentheses) are clustered at the firm level for the panel regressions in columns (1) and (2) and are clustered at the NACE 2-digit industry level for the cross-sectional regressions in columns (3) and (4). *,** ,*** indicate significance at the 10%, 5%, and 1% levels, respectively.

of 0.2 to the OLS estimates.[56] We then use the Nelder-Mead algorithm and record the objective function value at numerical convergence for each iteration. Finally, we select the coefficient estimates corresponding to the global minimum among these 3,000 trials. Figure 8 illustrates how the estimates from column (1) of Table 3 and column (1) of Table 4 vary depending on the number of starting-point trials used when selecting the global minimum. This figure confirms the finding of Rovigatti and Mollisi (2018) that the estimates of production function coefficients can be biased when few iterations are used (most commonly used code packages in the literature default to one iteration). However, after around 500 trials of starting points, there is no further noticeable improvement in the optimization procedure and the estimates become stable, suggesting numerical convergence toward the true global minimum. Therefore, by taking the global minimum among the 3,000 trials, we largely resolve the numerical challenges raised by Rovigatti and Mollisi (2018).

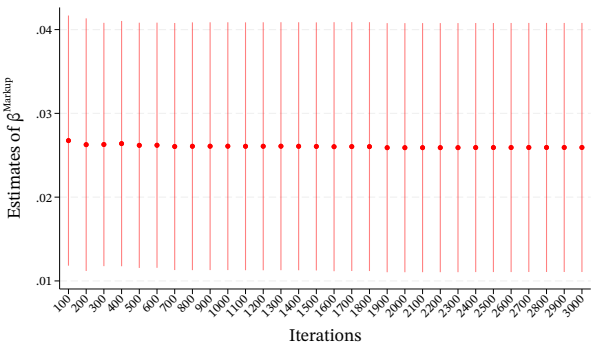
[56] We implement this by modifying the `prodest` package in Stata.

Figure 8: Key estimates of Table 3 and 4 across trials of starting points in the ACF correction

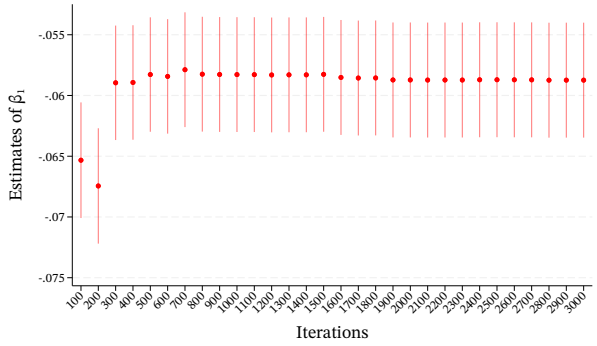
(a) Col. 1, row "$\overline{\log(\text{Domar Weight})_i}$" of Table 3



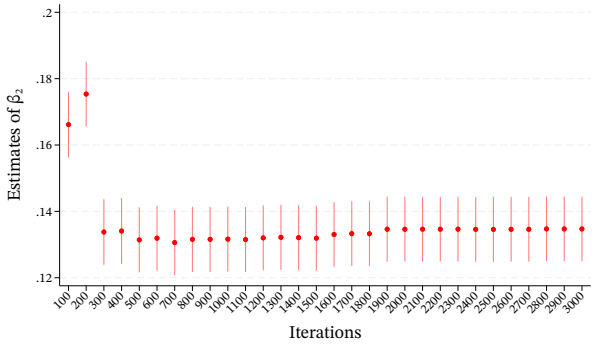
(b) Col. 1, row "$\overline{\log(\text{Markup})_i}$" of Table 3



(c) Col. 1, row "$\log(\text{Domar Weight}_{it}) \times \Delta \log GDP_t$" of Tab. 4



(d) Col. 1, row "$\log(\text{Markup}_{it}) \times \Delta \log GDP_t$" of Tab. 4



Notes: Estimation results of (64) and (65) using a sample of Spanish firms from Orbis against the maximum number of second-stage GMM iterations (shown in multiples of 100) allowed in the LP-ACF production-function estimation. For each iteration cap, the algorithm is run from multiple starting values and halted once the limit is reached; the best objective value achieved within that limit is then used to construct TFP growth and markups for the regressions. Panel (a): estimates of β^{Domar} in (64) (Col. 1, row "$\overline{\log(\text{Domar Weight})_i}$" of Table 3). Panel (b): estimates of β^{Markup} in (64) (Col. 1, row "$\overline{\log(\text{Markup})_i}$" of Table 3). Panel (c): estimates of β_1 in (65) (Col. 1, row "$\log(\text{Domar Weight}_{it}) \times \Delta \log GDP_t$" of Table 4). Panel (d): estimates of β_2 in (65) (Col. 1, row "$\log(\text{Markup}_{it}) \times \Delta \log GDP_t$" of Table 4. See text for a detailed description of variable construction and sample selection. 90% confidence intervals are constructed using standard errors that are clustered at the firm level.

B.3.6 Robustness on the scaling of TFP volatility

In this appendix, we show that our reduced-form evidence that smaller firms and those with higher markups exhibit greater TFP volatility in the cross-section is robust to an alternative scaling of TFP volatility. In particular, we perform a robustness check to ensure our results still hold when using the natural log of firm-level standard deviation of TFP growth, $\log(\sigma_i(\Delta \log TFP_{it}))$, as the dependent variable and re-estimate equation (64). Column (2) of Table 7 presents the estimates under this alternative scaling and column (1) reproduces our baseline results with $\sigma_i(\Delta \log TFP_{it})$ as the dependent variable. Results under this alternative scaling remain statistically significant and qualitatively similar to the baseline.

Table 7: Cross-sectional results with alternative scaling of TFP volatility

	(1): $\sigma_i(\Delta \log TFP_{it})$	(2): $\log(\sigma_i(\Delta \log TFP_{it}))$
$\overline{\log(\text{Domar Weight})_i}$	-0.037^{***}	-0.181^{***}
	(0.002)	(0.008)
$\overline{\log(\text{Markup})_i}$	0.026^{***}	0.119^{***}
	(0.009)	(0.029)
Sector FE	Yes	Yes
Observations	367,213	367,213
R^2	0.383	0.445

Notes: Table presents estimation results of (64) with different scaling of TFP volatility using a sample of Spanish firms from Orbis. See text for a detailed description of variable construction and sample selection. The estimation sample is winsorized the top and bottom 0.5% of observations of TFP growth, domar weight and markup in the panel in each sector before the computing the firm-level moments and only includes firms that have 1) at least 5 non-missing observations of TFP growth, and 2) positive average log markup. Standard errors (in parentheses) are clustered at the NACE 2-digit industry level. *,** ,*** indicate significance at the 10%, 5%, and 1% levels, respectively.

B.4 US stock market evidence

Our theory predicts that, all else equal, larger firms (as measured by Domar weights) and those with lower markups tend to exhibit weaker comovement with aggregate shocks. One common way to investigate this covariance is to use stock market equity betas that can be estimated using high-frequency financial data. As is well-known in the finance literature (e.g., Fama and French, 1992), firms' equity betas are strongly negatively correlated with their stock market capitalizations. Since firms with high market capitalizations tend to have high Domar weights, we expect to find a similar relationship between betas and Domar weights, as our theory predicts. We also explore how equity betas vary with firms' wedges.

We use a sample of US public firms from the CRSP/Compustat Merged database (see Appendix B.2.2 for a detailed description of the data). For each firm, we estimate its 5-year rolling-window CAPM beta. As in the previous section, we use markups to proxy for wedges, where markups are constructed as in De Loecker, Eeckhout, and Unger (2020). To explore how firms' equity betas correlate with their Domar weights and markups, we use the same approach as in the previous section. Specifically, for each year we construct firms' within-sector markup deciles and NYSE-based Domar-weight deciles. [57] We then construct a set of dummy variables, D_{jit}^{Domar} and D_{jit}^{Markup}, such that $D_{jit}^{Domar} = 1$ if firm i's Domar weight is in decile j in year t, and analogously for markups. We

[57] Following Fama and French (1992), we use NYSE breakpoints when defining the Domar weight deciles. The size distribution of firms in the other large market, NASDAQ, has changed substantially over time (historically, Nasdaq was dominated by small tech companies). In contrast, the size distribution of the NYSE firms is more stable.

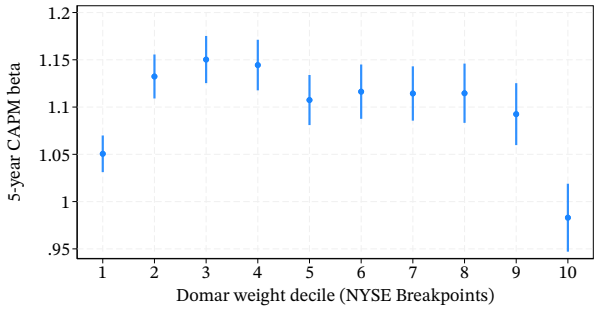
then run the following regression with sector fixed effects,

$$\beta_{it}^{CAPM} = b_0 + \sum_{j=1}^{10} b_j^{Domar} D_{jit}^{Domar} + \sum_{j=1}^{10} b_j^{Markup} D_{jit}^{Markup} + \theta_s + \varepsilon_{it}, \tag{66}$$

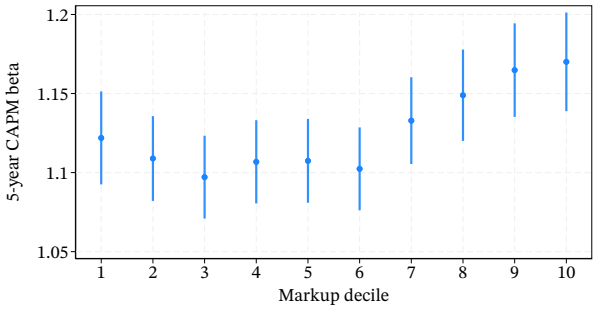
and plot the estimated regression coefficients b_j^{Domar} and b_j^{Markup} in Figure 6. Panel (a) shows that larger firms tend to have lower CAPM betas. That relationship holds quite monotonically except for firms in the first decile. For such small firms, the estimated betas can be noisy or even biased downwards due to liquidity issues (Ibbotson, Kaplan, and Peterson, 1997). Panel (b) shows that firms with higher markups tend to have higher CAPM betas. Both these results are in line with our theory as long as increases in productivity growth and stock market returns are correlated.

Figure 9: CAPM betas, Domar weights and markups

(a) CAPM beta by Domar weight decile



(b) CAPM beta by markup decile



Notes. Estimation results of (66) using a sample of US firms from Compustat. See text for a detailed description of variable construction and sample selection. The estimation sample is winsorized at the top and bottom 3% of observations of equity beta, Domar weight and markup in each sector. Panel (a): $b_j^{Domar} + b_5^{Markup} + b_0$ by Domar weight decile j ($j = 1$ is lowest Domar weight, $j = 10$ is highest Domar weight). Panel (b): $b_j^{Markup} + b_5^{Domar} + b_0$ by markup decile ($j = 1$ is lowest markup, $j = 10$ is highest markup). 90% confidence intervals are constructed using standard errors that are clustered at the firm level.

B.5 Cross-country evidence

In this appendix, we show that countries with higher GDP per capita and smaller distortions tend to have lower TFP growth volatility. We use three measures of country-level distortions: 1) the share of government consumption in GDP (as a proxy for tax burden), 2) country-level aggregate markups and 3) the inverse of credit-to-GDP ratio (as a proxy for financial frictions, see Arellano, Bai, and Zhang, 2012; Gorodnichenko and Schnitzer, 2013; Midrigan and Xu, 2014).[58] We combine data from the Penn World Table (version 10.01) with the country-level markup estimates from De Loecker and Eeckhout (2018) and credit-to-GDP ratios from the World Bank. For each country, we compute the standard deviation of TFP growth σ_{it}^{TFP} using 10-year rolling windows (we use TFP at constant national prices). The cross-country average of σ_{it}^{TFP} in Penn World Table is 2.69%.

[58] Bigio and La'O (2020) show that some financial frictions can show up as wedges in production network economies.

Following Acemoglu and Zilibotti (1997), we then run the following regression:

$$\log \sigma_{it}^{TFP} = \alpha + \beta^{GDP} \log(\text{GDP}_{\text{pc},it}) + \beta^{\tau} \log \tau_{it} + FE_i + FE_t + \varepsilon_{it}, \tag{67}$$

where $\log(\text{GDP}_{\text{pc},it})$ is log of GDP per capita of country i in year t (we use expenditure-side real GDP divided by total population), and $\tau_{it} \in \left\{ \frac{\text{Gov Cons}}{\text{GDP}}_{it}, \text{Markup}_{it}, \left(\frac{\text{Credit}}{\text{GDP}}\right)_{it}^{-1} \right\}$ represents various measures of distortions. The variables FE_i and FE_t are country and year fixed effects. The estimated coefficients are given in Table 8, and the associated binscatters are shown in Figure 10.

We find that countries with larger government consumption shares (heavier tax burdens), higher aggregate markups and worse financial development tend to havehigher TFP volatility while richer economies tend to experience lower TFP volatility. Our model can rationalize these findings. Insofar as some of those distortions affect the wedge between the price at which goods are sold and their production cost, our theory would suggest that firms in these countries are less aggressive with their risk management, and aggregate volatility should be higher as a result. Similarly, richer countries might have better risk management capabilities (lower B_i, G_i, or η) which would explain the lower aggregate volatility.

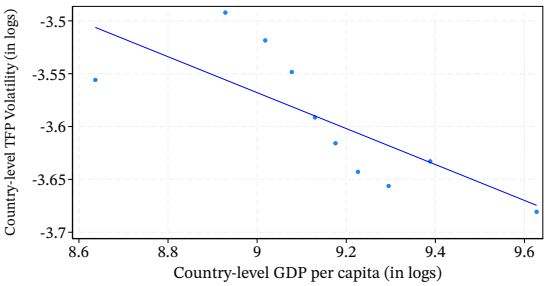
Table 8: Cross-country TFP volatility, GDP per capita and distortions

	Dependent variable: Country-specific TFP volatility				
	(1)	(2)	(3)	(4)	(5)
$\log(\text{GDP}_{\text{pc},it})$	-0.170^{**}			-0.178^{**}	-0.164
	(0.075)			(0.072)	(0.200)
$\log\left(\frac{\text{Gov Cons}}{\text{GDP}}_{it}\right)$		0.114^{*}		0.126^{**}	
		(0.062)		(0.057)	
$\log(\text{Markup}_{it})$			0.287^{*}		0.253^{*}
			(0.151)		(0.143)
Country FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	5,227	5,227	1,382	5,227	1,382
R^2	0.674	0.671	0.733	0.677	0.736

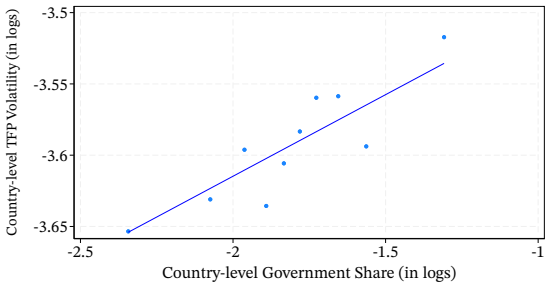
Notes: Results of estimating (67) using a cross-country sample from Penn World Tables. See text for a detailed description of variable construction and sample selection. Standard errors (in parentheses) are clustered at the country level. *,**,*** indicate significance at the 10%, 5%, and 1% levels, respectively.

Figure 10: Binscatter plots of cross-country TFP volatility, GDP per capita and government share

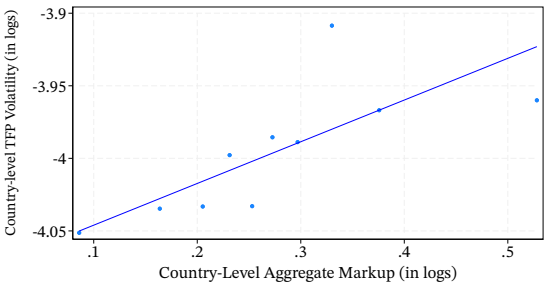
(a) TFP volatility and GDP per capita



(b) TFP volatility and government share



(c) TFP volatility and aggregate markup



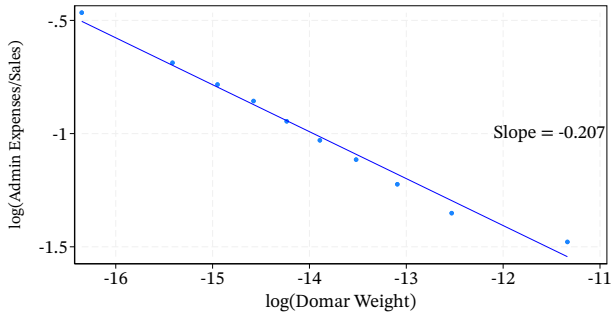
Notes: Binned scatter plots of the log of cross-country TFP volatility against the log of GDP per capita (panel a), the log of government share (panel b), the log of country-level aggregate markup (panel c), and the log of the inverse of private-credit share of GDP (panel d), using a cross-country sample from the Penn World Tables. See text for a detailed description of variable construction and sample selection. Country and year fixed effects are controlled for in all binscatter plots.

B.6 Evidence on the return-to-scale of risk management inputs

As we explored in Section 4, firm i benefits from a scale advantage in risk management whenever the resource cost function g_i is not identically equal to zero. Figure 11 provides some evidence for that advantage in the data. It shows that the share of management expenditures (measured as `other_operating_expenses` in Orbis) declines with a firm's Domar weight. Insofar as some of these expenditures involve managing risk, this finding is consistent with the model's implication that risk management expenditures as a share of sales decrease with firm size. We can reproduce Figure 11 in our calibrated economy. Doing so, we find that the slope of the slope of the regression line is -0.08, while the equivalent number in the data is -0.21.[59] It is reassuring that the two numbers are broadly in line. If anything, the scale advantage of large firms appears more important in the data.

[59] In the calibrated economy, the slope is estimated by regressing the log of the ratio of risk management resource costs to firm sales on log sales. We do not include the productivity cost in the regression as it would not show up in the accounting data.

Figure 11: Administrative expenditure share and firm size

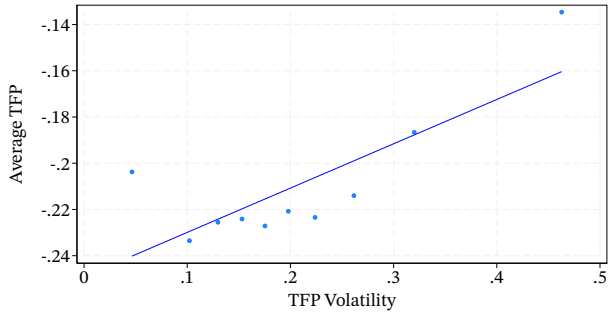


Notes: Binned scatter plots of the log of administrative expenditure share against the log of Domar weight using a sample of Spanish firms from Orbis. The administrative expenditure share is defined as the ratio between "other operating expenses" and sales. The estimation sample is winsorized at the top and bottom 0.5% of both variables. See text for a detailed description of variable construction and sample selection.

B.7 Evidence on the productivity cost of risk management inputs

Whenever $B_i \neq 0$, the model implies that firms managing their risks more aggressively have lower productivities. Specifically, a firm choosing a high exposure to bad risks would typically have high TFP volatility but, at the same time, high average TFP. Figure 12 shows that a firm's TFP volatility is positively correlated with its average TFP, confirming this prediction.

Figure 12: Average TFP vs. TFP Volatiltiy



Notes: Binned scatter plots of the average log TFP against TFP volatility using a sample of Spanish firms from Orbis. The estimation sample includes Spanish firms from Orbis that have 1) at least 5 non-missing observations of TFP growth, and 2) positive average log markup. The average log TFP is computed as the time-series average of a firm's log TFP net of sector-year fixed effects, and the TFP volatility is computed as the standard deviation of TFP growth. The estimation sample is winsorized at the top and bottom 0.5% of both variables.

C Appendix for Section 9

This appendix contains details about the calibration of Section 9.

C.1 Construction of the sample

This appendix describes the datasets used in the calibration and how the associated aggregate and firm-level moments are computed.

1. We measure real GDP using the chain-linked volume index from the Annual Spanish National Accounts (Gross domestic product at market prices and its components, Table 4). To be consistent with the Orbis data, we use the mean and variance of GDP growth, $y_t - y_{t-1}$, during the 1995-2019 period.

2. We calibrate the sectoral parameters using the 2010 input-output table from the Annual Spanish National Accounts. This table partitions the Spanish economy into 62 sectors which are usually defined at the 2-digit NACE industry level.[60] Conforming to the accounting conventions in the data, we calibrate the input elasticities of good s' in the production of sector s as $\hat{\alpha}_{ss'} = \frac{\text{Input from } s' \text{ at basic price}_s}{\text{Total input at basic price}_s} \times \frac{\text{Intermediate consumption at purchaser's price}_s}{\text{Output at basic price}_s}$, such that the residual labor share, $1 - \sum_{s'} \hat{\alpha}_{ss'}$, corresponds to the value added share of output at basic price in the data.[61] We calibrate the consumption share β_s to be the share of final consumption expenditure of good s in the sum of consumption expenditure spent on the 62 sectors.

3. We compute firm-level objects from the Orbis sample. After steps 1-4 in Section B.2.1, we perform a few additional steps.

   (a) We drop firm-year observations with average revenue product of any input (fixed asset, wage bills, and material costs) in the top and bottom 1% of the corresponding year-specific distributions.

   (b) We drop a firm if its sales, fixed assets, costs of employees, or material costs ever exceed the top or bottom 0.1% of the distribution of all firm-year observations.

   (c) We winsorize markup growth, real sales growth, real input (capital, labor, material) growth, and the level of markups at the top or bottom 0.5% of the distribution of all firm-year observations. We compute TFP growth using (62) and drop a firm if it does not have more than 5 valid observations. We then compute the correlation of each firm's TFP growth with GDP growth. We collapse the panel data into a cross-section of firms and compute the time series average of markups, which correspond to $1 + \tau_{si}$ in the model. We further winsorize firms' markups at 1% at the top and the bottom in the cross-section of firms to avoid outliers. Lastly, we compute θ_{si} as the share of the firm's Domar weight in the sum of all firms' Domar weights in the sector.

[60] Sector 63 (household-related production activities) and sector 64 (services by extraterritorial organizations and bodies) are also present in the 2010 input-output table, but their input-output data is missing.

[61] Because sector-to-sector data at purchasers' prices (i.e., adjusted for taxes less subsidies on products) are unavailable, we calibrate the intermediate-input expenditure share of inputs from sector s' used by sector s with the share computed in basic prices, $\frac{\text{Input from } s' \text{ at basic prices}_s}{\text{Total input at basic prices}_s}$.

4. With $\hat{\alpha}$, β, θ, and τ from the previous steps, we construct model-implied Domar weights using (15).

C.2 Stationarity of the risk-exposure decision

In the model of Section 9, log real GDP is given by

$$\tilde{y} = \Delta\varepsilon + \left(\tilde{\omega}^f\right)^\top \left(-b\left(\delta\right) + v + \gamma t\right) + \left(\tilde{\omega}^s\right)^\top z - \left(\tilde{\omega}^f\right)^\top \log\left(1+\tau\right) - \log\Gamma_L, \tag{68}$$

where $\tilde{\omega}^f$ and $\tilde{\omega}^s$ are the cost-based Domar-weight vectors of the firms and of the sectoral aggregators, v and z are the column vectors of the firm-level and sector-level productivity shocks, and γ is the vector of the firm-level growth trends γ_{si}. The fictitious planner's problem is therefore

$$\begin{aligned}\mathcal{W}_{dist} := \max_{\Delta} \underbrace{\Delta \times 0 - \bar{b}\left(\Delta\right) + \left(\tilde{\omega}^f\right)^\top \left(\mu^v + \gamma t\right) + \left(\tilde{\omega}^s\right)^\top \mu^z - \left(\tilde{\omega}^f\right)^\top \log\left(1+\tau\right) - \log\Gamma_L}_{\mathrm{E}[y]} \\ - \frac{1}{2}\left(\rho - 1\right) \underbrace{\left(\Sigma\Delta^2 + \left(\tilde{\omega}^f\right)^\top \Sigma^v \tilde{\omega}^f + \left(\tilde{\omega}^s\right)^\top \Sigma^s \tilde{\omega}^s\right)}_{\mathrm{V}[y]} - \bar{g}\left(\Delta\right),\end{aligned}$$

where $\mathbb{E}\left[\varepsilon\right] = 0$, μ^v and μ^z are the expected value vectors of v and z, and Σ^v and Σ^z are the covariance matrices of v and z. Notice that the only non-stationary term, the growth trend vector γt, does not interact with the choice of Δ, and so Δ is constant over time. Consequently, δ is also constant over time as it solves (31).

C.3 Identifying δ_{si} and Σ^v_{si}

Combining firm i's TFP from (51) with log real GDP from (68), we can write

$$\begin{aligned}A_{si} := \mathrm{V}\left[\log TFP_{si,t} - \log TFP_{si,t-1}\right] &= \mathrm{V}\left[\delta_{si}\varepsilon_t + v_{sit} + \gamma_{si}t - \delta_{si}\varepsilon_{t-1} - v_{sit-1} - \gamma_{si}\left(t-1\right)\right] \\ &= \mathrm{V}\left[\delta_{si}\left(\varepsilon_t - \varepsilon_{t-1}\right) + v_{si,t} - v_{si,t-1}\right] \\ &= 2\delta_{si}^2\Sigma + 2\Sigma^v_{si}.\end{aligned} \tag{69}$$

Similarly, for the covariance, we have

$$\begin{aligned}B_{si} := \mathrm{Cov}\left[y_t - y_{t-1}, \log TFP_{si,t} - \log TFP_{si,t-1}\right] = \mathrm{Cov}\Big[&\Delta_t\varepsilon_t - \Delta_{t-1}\varepsilon_{t-1} + \left(\omega^f\right)^\top \left(v_t - v_{t-1}\right) \\ &+ \left(\omega^s\right)^\top \left(z_t - z_{t-1}\right), \delta_{si}\varepsilon_t - \delta_{si}\varepsilon_{t-1} + v_{si,t} - v_{si,t-1}\Big] \\ &= 2\Delta\Sigma\delta_{si} + 2\tilde{\omega}^f_{si}\Sigma^v_{si}.\end{aligned} \tag{70}$$

In these equations, we define A_{si} and B_{si} to simplify the notation in what follows. Notice also that (69) and (70) are the same as in Section 8, and that we can therefore measure both A_{si} and B_{si} directly from the Spanish data, as explained in that section.

From the observed A_{si} and B_{si}, we can then identify key objects of the calibrated economy. From (69), we can write

$$\Sigma_{si}^{v} = \frac{A_{si}}{2} - \delta_{si}^{2}\Sigma. \tag{71}$$

Combining with (70), we find a quadratic equation in δ_{si}, whose solutions are given by

$$\delta_{si} = \frac{\Delta}{2\tilde{\omega}_{si}^{f}} \left(1 \pm \sqrt{1 - 2\tilde{\omega}_{si}^{f} \frac{B_{si} - \tilde{\omega}_{si}^{f} A_{si}}{\Delta^{2}\Sigma}} \right). \tag{72}$$

Since $\Delta = \sum_{s=1}^{S} \sum_{i=1}^{N_s} \tilde{\omega}_{si}^{f} \delta_{si}$, we find

$$2\Delta = \sum_{s=1}^{S} \sum_{i=1}^{N_s} \left(\Delta \pm \sqrt{\Delta^{2} - 2\tilde{\omega}_{si}^{f} \frac{B_{si} - \tilde{\omega}_{si}^{f} A_{si}}{\Sigma}} \right),$$

or

$$1 = \frac{1}{N-2} \sum_{s=1}^{S} \sum_{i=1}^{N_s} \sqrt{1 - 2\tilde{\omega}_{si}^{f} \frac{B_{si} - \tilde{\omega}_{si}^{f} A_{si}}{\Sigma\Delta^{2}}}.$$

Given a normalization for Σ, we can solve this equation numerically for Δ^2. Since aggregate exposure is positive, this gives us Δ. Combining with (72), we find δ_{si} (we pick the negative root because it corresponds to a positive Δ, in line with our normalization). Finally, combining with (71) gives us Σ_{si}^{v}.

We use this procedure to pin down the risk exposure δ_{si} and the idiosyncratic volatility Σ_{si}^{v} of about 99% of the firms in the sample, but there are some firms that are too extreme to fit our setup. Those are firms with covariance B_{si} that are extremely large or extremely small compared to their variance A_{si}.[62] Those firms are generally small and we suspect that measurement errors might contribute to their extreme properties. To include these firms in the model, we simply endow them with a new pair of measurements (A_{si}, B_{si}) from the distribution of firms. We have experimented with other ways to include these firms and have found that they only matter minimally for the counterfactual exercises.

[62] This happens when a firm covaries strongly with GDP while having a small variance. From (70), the strong covariance implies that δ_i must be large, but then a negative Σ_{si}^{v} might be required to match the variance through (69).

C.4 Identifying B_s and ηG_s

To estimate the parameter of the exposure cost functions, we rely on Lemma 3, which implies

$$\delta_{si} = \delta_s^\circ + u_{si} + \left(B_s + \eta \frac{1+\tau_{si}}{\omega_{si}^f} G_s \right)^{-1} \mathcal{E}, \tag{73}$$

where we have imposed that B_s and G_s are sector-specific, and that $\delta_{si}^\circ = \delta_s^\circ + u_{si}$ can be written as the sum of a sector-specific term δ_s° and a residual u_{si}. At this stage in the calibration, we have already identified δ_{si}, τ_{si}, ω_{si}^f, and $\mathcal{E}$, and so we can estimate B_s and G_s as the values that best fit (73). The inverse function in (73) introduces strong nonlinearities when B_s and G_s are small, which makes this curve-fitting exercise challenging. We therefore work instead with an approximated version of (73). Since Domar weights ω_{si}^f are small, we define $x_{si} := \omega_{si}^f / (1+\tau_{si})$ and approximate (73) with a quadratic function around $x_{si} = 0$ as

$$\delta_{si} \approx \delta_s^\circ + u_{si} + (\eta G_s)^{-1} \mathcal{E} x_{si} - (\eta G_s)^{-2} B_s \mathcal{E} x_{si}^2.$$

We then estimate ηG_s and B_s as those terms that minimize the sum of the squares of the residual $\sum_i^{N_s} u_{si}^2$ in sector s. Finally, with those values in hand, we set $\delta_s^\circ + u_{si} = \delta_{si} - \left(B_s + \eta \frac{1+\tau_{si}}{\omega_{si}^f} G_s \right)^{-1} \mathcal{E}$ so that (73) is satisfied.

C.5 Calibrated economy

This appendix provides more detail about the calibrated economy.

C.5.1 Distributions

The calibrated economy matches several features of the data perfectly. Figure 13 shows the distributions of four such features: 1) the sales share of each firm within its sector (panel a), 2) the firm-level markups (panel b), 3) the correlation between firm-level TFP growth and GDP growth (panel c),[63] and 4) the standard deviation of firm-level TFP growth (panel d).

In contrast, our calibration strategy implies that some other quantities are not fitted perfectly. This is the case for revenue-based Domar weights. We show the distribution of those weights in the data and in the model in Figure 14. The calibration matches well the right-tail of the distribution, where the biggest firms are located. This is reassuring as those firms are important drivers of GDP fluctuations. In contrast, the calibration features too many small firms and too few mid-size firms. The right panel of Figure 14 shows that the calibrated distribution of cost-based Domar weights, which we do not observe in the data, follows a similar shape.

[63] The jumps in the tails of the distribution are a consequence of the redrawing procedure for extreme firms described at the end of Appendix C.3.

Figure 13: Data distributions that the calibration matches exactly

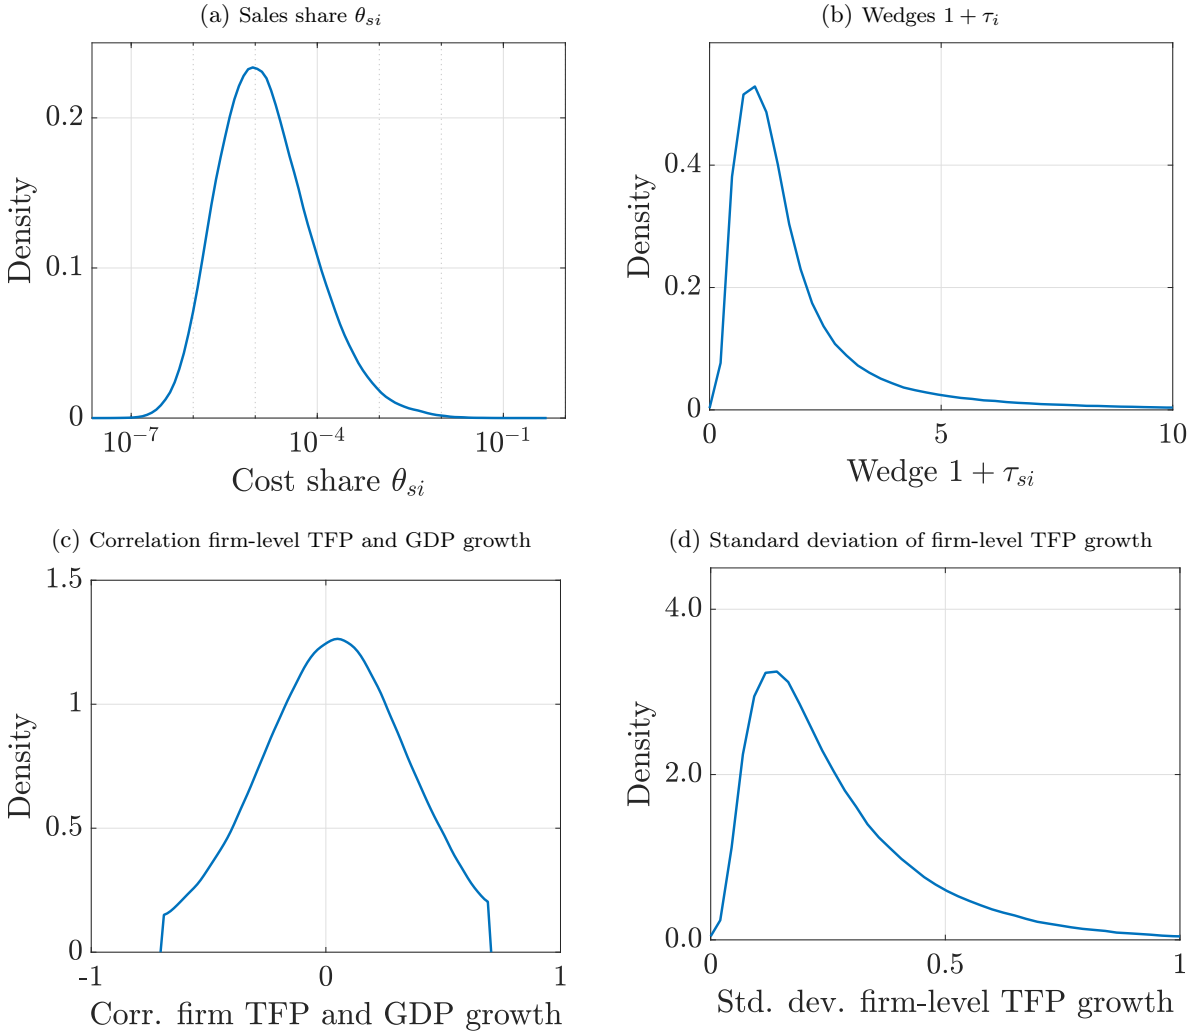
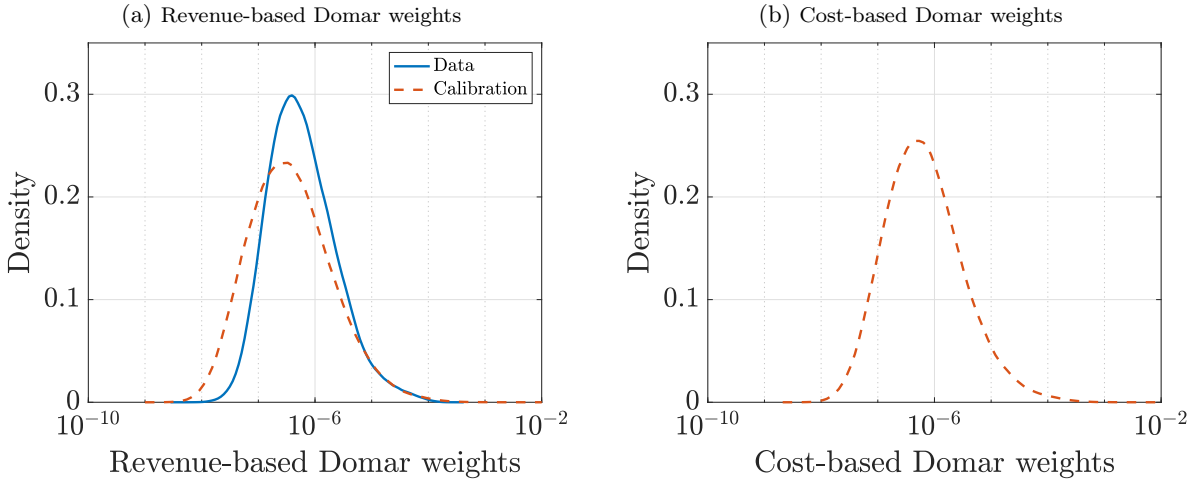


Figure 14: Domar weights of the firms in the data and in the model



C.5.2 Cost of risk exposure

Figure 15 shows the estimated value of H_s^{-1} for each sector. Several sectors have an estimated value of H_s^{-1} near zero, implying that their risk-taking behavior is rigid. The figure also shows the productivity cost and resource cost components of H_s. Figure 16 reports the distribution of the estimated firm-level natural risk exposure δ_i°.

C.5.3 Link between firm size, wedges and risk

To verify that our calibrated economy is able to replicate the empirical relationship between firm-level volatility, firm size and markups, we run the regressions in Appendix B.3.6 on the simulated data. The results are presented in Table 9. Reassuringly, the estimated coefficients are close to those of Table 7 which performs the same analysis on the data.

Table 9: Cross-sectional results with alternative scaling of TFP volatility

	(1): $\sigma_i\left(\Delta \log TFP_{it}\right)$	(2): $\log\left(\sigma_i\left(\Delta \log TFP_{it}\right)\right)$
$\overline{\log\left(\text{Domar Weight}\right)_i}$	-0.044^{***}	-0.178^{***}
	(0.002)	(0.006)
$\overline{\log\left(\text{Markup}\right)_i}$	0.045^{***}	0.163^{***}
	(0.009)	(0.013)
Sector FE	Yes	Yes
Observations	531,944	531,944
R^2	0.222	0.313

Notes: The table presents estimation results of (64) with different scaling of TFP volatility using simulated data from the calibrated economy. Standard errors are in parentheses. *,** ,*** indicate significance at the 10%, 5%, and 1% levels, respectively.

C.6 Cross-sectional effects of wedges and uncertainty

In this appendix, we investigate how doubling Σ and removing wedges affect the cross-section of firms. It is useful to rewrite (52) as

$$\mathrm{V}\left[y_t - y_{t-1}\right] = \sum_{s=1}^{S}\sum_{i=1}^{N_s} \tilde{\omega}_{si}\,\mathrm{Cov}\left[y_t - y_{t-1}, \log \mathrm{TFP}_{si,t} - \log \mathrm{TFP}_{si,t-1}\right] + 2\tilde{\omega}_s^\top \Sigma^z \tilde{\omega}_s. \tag{74}$$

It follows that the endogenous part of the variance of GDP is driven by the Domar-weighted average of the covariance of firm TFP growth with GDP growth. We will rely on that relation below to explore the impact of changes in the environment.

Consider the effect of doubling Σ first. The left panel in Figure 17 shows that, keeping δ fixed, the distribution of the covariance between firm-level TFP growth and GDP growth widens. Intuitively, as Σ increases, ε_t becomes a more important driver of firm-level TFP. As a result, firms that are positively exposed to ε_t now covary more with GDP while the opposite happens for negatively exposed firms.[64] Since on average firms are positively exposed to ε_t ($\Delta > 0$), this widening leads to an increase in the average covariance and, by (74), to an increase in aggregate volatility. When δ is flexible, firms respond to this large increase in aggregate volatility by managing their risk more aggressively, which undoes some of the widening, and mitigates the overall increase in aggregate risk.

Next, we look at the impact of the removal of the wedges. The right panel of Figure 17 shows that the distribution of the covariance of firm-level TFP growth with GDP growth becomes more peaked as a result. Indeed, the removal of the wedges leads to a decline in $\Delta\Sigma$ which, by (48), leads to a decline in the absolute value of each firm's covariance with GDP. Since the average firm is positively exposed to the aggregate risk ($\Delta > 0$), this compression in the covariance distribution leads to a decline in GDP volatility by (74).

C.7 Calibration using different wedge measures

In this appendix, we recalibrate the benchmark economy of Section 9.5 under alternative wedge measures: 1) no wedges (efficient allocation), 2) markups estimated using the Lerner index approach, 3) adjusted markups in line with De Ridder, Grassi, Morzenti, et al. (2024), and 4) wedges set to replicate the Spanish value-added tax. In each case, we recalibrate the model and repeat the experiments of Section 9.5. We report the results in Table 10. We can see from the table that without an endogenous response of risk exposure (fixed δ), doubling Σ or setting $\tau = 0$ would lead to the same outcome regardless of our choice of wedges. In the rest of this section, we describe what happens when δ is free to adjust to the change in the environment.

Efficient allocation

For this exercise, we recalibrate the model with $\tau = 1$, so that there are no wedges at all. When we double Σ, the standard deviation of GDP growth goes up by 9 bps under a flexible δ. Compared to the benchmark model, this suggests that the efficient allocation is better able to handle the increase in fundamental uncertainty.

[64] This can be seen more readily in (48). When Σ increases, the term $2\Delta\delta_i\Sigma$ grows for firms with $\delta_i > 0$ but shrinks for firms with $\delta_i < 0$, which explains the widening of the covariance distribution. In equilibrium, δ responds to mitigate that effect without offsetting it fully.

Table 10: Standard deviation of GDP growth under different wedge measures

	Calibration	Double Σ		No wedges	
		Fixed δ	Flex. δ	Fixed δ	Flex. δ
Benchmark	2.35%	2.88%	2.47%	2.35%	1.80%
Efficient allocation	2.35%	2.88%	2.44%	2.35%	2.35%
Lerner index markups	2.35%	2.88%	2.44%	2.35%	2.30%
Adj. markups	2.35%	2.88%	2.57%	2.35%	1.82%
Value-added tax	2.35%	2.88%	2.44%	2.35%	2.31%

Alternative markup measures

One worry about our baseline markups is that the estimates using the production function approach might be biased or imprecise. To alleviate these concerns, we recalibrate the model using markups based on the Lerner index, as described in Appendix B.3.3. These markups are very small—the average value of τ is 2.9%—and we see from Table 10 that when Σ doubles, the equilibrium behaves similarly to the efficient allocation. The last column of the table shows that removing wedges leads to a small decline in aggregate volatility.

In a second exercise, we adjust our baseline markups in line with findings from De Ridder, Grassi, Morzenti, et al. (2024). That paper relies on price and quantity data from France and finds that the production approach to markup estimation (using revenue data) captures well markup variation across firms, but not their overall level. In their estimates, the (log) quantity-based markups is about 2.9 times larger than their revenue-based counterparts. We therefore provide a version of our calibrated economy in which we adjust our baseline markups to be in line with those of De Ridder, Grassi, Morzenti, et al. (2024). That is, we set

$$\log\left(1+\tau_{adjusted}\right) = 2.9 \times \log\left(1+\tau_{baseline}\right).$$

We see from Table 10 that under these larger markups, the economy is less able to handle an increase in Σ. After Σ doubles, the standard deviation of GDP growth increases by almost 22 bps. Unsurprisingly, removing these large wedges leads to a large decline in volatility.

Value added tax

Spain has a value-added tax (VAT) that can be as high as 21% for certain sectors. We calibrate a version of our baseline model in which wedges are set to match the VAT instead of the estimated wedges. Since the VAT varies across sectors, we rely on data from the European Commission to set a sector-level wedge τ_s that is in line with statutory information.[65] Since wedges in the model

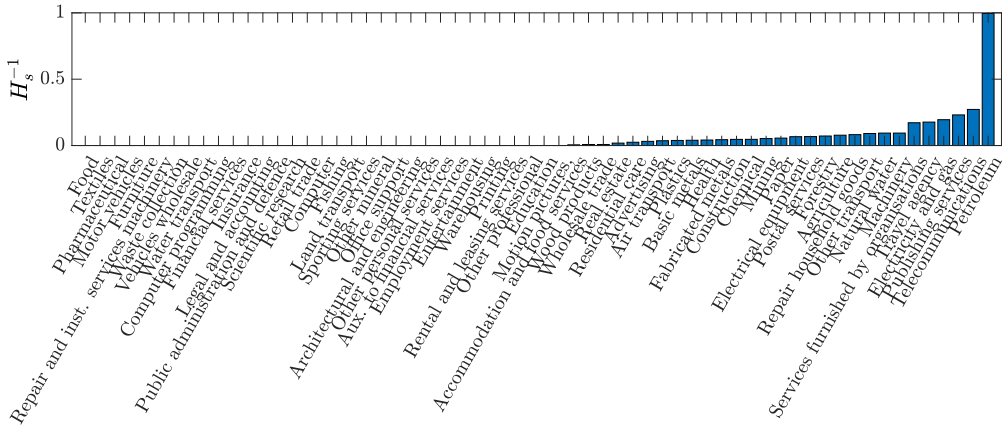
[65] That data is available at the website https://taxation-customs.ec.europa.eu/system/files/2021-06/vat_rates_en.pdf. We downloaded that data on February 26th 2025.

correspond to a sales tax instead of a value-added one, we adjust the statutory VAT rate with the value-added cost share of each firm. That is, we set

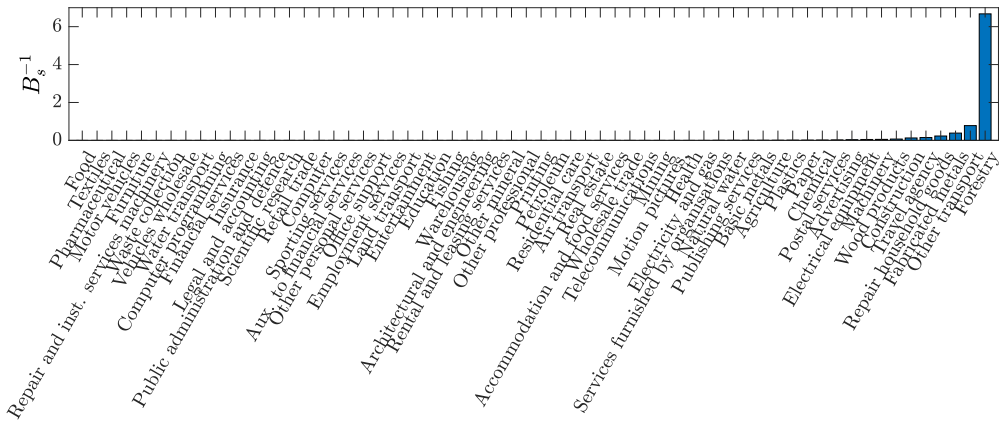
$$\tau_s = VAT_s \times \left(1 - \sum_{s'=1}^{S} \hat{\alpha}_{ss'}\right).$$

These wedges are fairly small, with a mean of 8%. It follows that the economy handles increases in Σ well, and that removing wedges does not lead to a large decline in aggregate volatility.

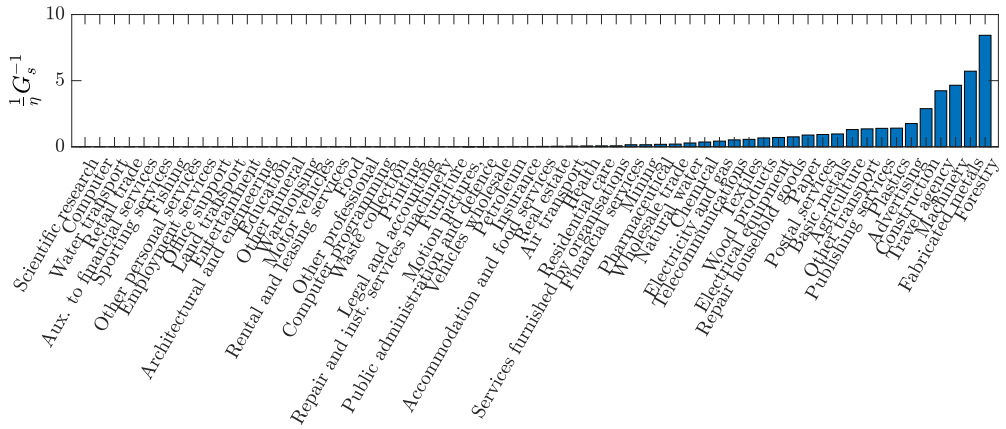
Figure 15: The costs of exposure in the calibrated economy



(a) Estimated value of H_s^{-1} for each sector.



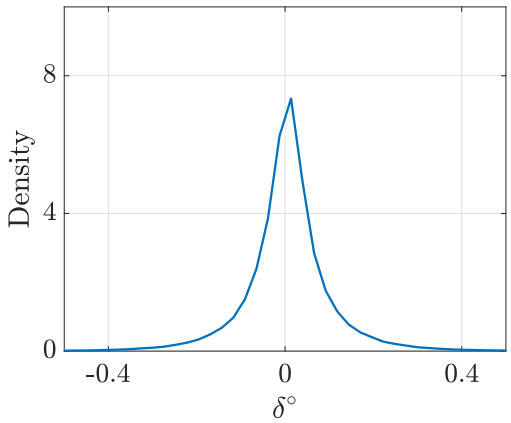
(b) Estimated value of B_s^{-1} for each sector.



(c) Estimated value of $\frac{1}{\eta} G_s^{-1}$ for each sector.

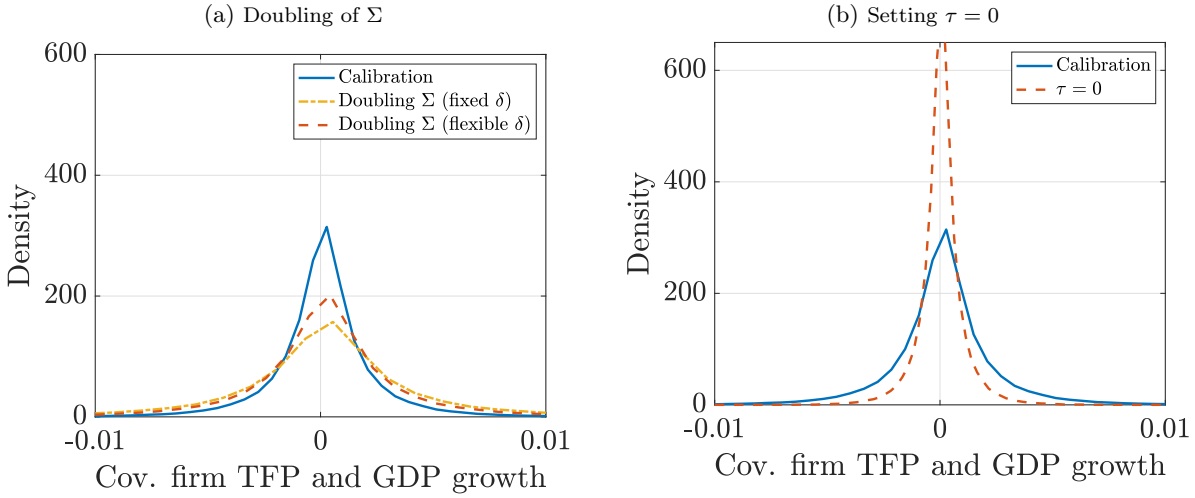
Notes. The scale of H_s^{-1} depends on our choice of ρ and Σ. We set $\rho = 5$ and $\Sigma = 1$ for this figure.

Figure 16: Distribution of the estimated firm-level natural risk exposure δ_i°.



Notes. The scale of δ_i° depends on our choice of ρ and Σ. We set $\rho = 5$ and $\Sigma = 1$ for this figure.

Figure 17: Changes in the environment and the covariance distribution



Notes: Both panels show the cross-sectional distribution of the covariance of firm-level TFP growth with GDP growth. Panel (a) shows how that distribution changes when Σ doubles. Panel (b) shows how that distribution changes under $\tau = 0$.

D Empirical evidence with other risk factors

In this appendix, we show that our model has broader quantitative applications to risk factors beyond business-cycle fluctuations. Specifically, we demonstrate that the model's predictions continue to hold when applied to climate risk and to risk factors estimated from a dynamic factor model.

D.1 Climate risk

We consider the implications of our endogenous risk model for climate change. It is well known that climate shocks (mostly, heat shocks) can cause significant productivity damage at the firm-level (Zhang et al., 2018) and at the macro level (Dell, Jones, and Olken, 2012). In the context of our model, we therefore think of climate risk as a bad shock ($\mathcal{E} < 0$) and consider different adaptation actions that firms can take to lower their climate risk exposure (e.g. install air conditioning units, or take measures against floods and storms). Corollary 1 implies that smaller firms and those with higher markups should be more exposed to climate change since those firms find it less cost-efficient to aggressively manage their climate risk exposure.

We test whether the model's key predictions are supported by the data. To do so, we merge our firm panel with detailed weather data for Spain and firm-location data, compiled by Liu and Xu (2024) (see that paper for more details about the data). The weather data are derived from ERA5 hourly data and provide daily temperatures for each of Spain's 50 NUTS-3 regions (province equivalents) from 1998 to 2018. The firm-location data were constructed from the postcode information in the Orbis dataset, using Eurostat's postcode–NUTS-3 crosswalk. Using the merged dataset, we can therefore explore how the link between firm-level productivity and temperature changes with firm characteristics.

We focus on a simple exercise to test the prediction of the model that smaller firms and those with larger markups are more exposed to climate change. The productivity of those firms should decline more during years with particularly warm weather. To test that prediction, we first define the temperature bin variables

$$\left\{\text{Tbin}_{r,t}^{<0^\circ\text{C}},\ \text{Tbin}_{r,t}^{0\text{–}20^\circ\text{C}},\ \text{Tbin}_{r,t}^{>20^\circ\text{C}}\right\},$$

where $\text{Tbin}_{r,t}^{b}$ is the number of days in geographical region r and year t with a mean daily temper-

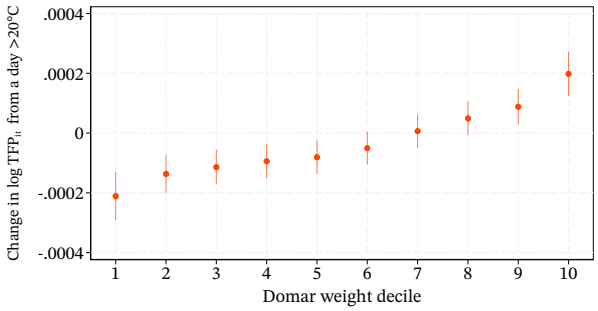
ature that falls in the range b. We then run the following regression:

$$\begin{aligned}\log\left(TFP_{i,t}\right) = & \sum_{b\in B/\{10-20°C\}} \delta_b \times \text{Tbin}^{b}_{r(i),t} \\ & + \sum_{b\in B/\{10-20°C\}} \sum_{j=2}^{10} \delta^{\text{Domar}}_{b,j} \times \text{Tbin}^{b}_{r(i),t} \times D^{\text{Domar}}_{ji} \\ & + \sum_{b\in B/\{10-20°C\}} \sum_{j=2}^{10} \delta^{\text{Markup}}_{b,j} \times \text{Tbin}^{b}_{r(i),t} \times D^{\text{Markup}}_{ji} \\ & + \delta_i + \alpha_{s(i),t} + \varepsilon_{i,t}\end{aligned}$$

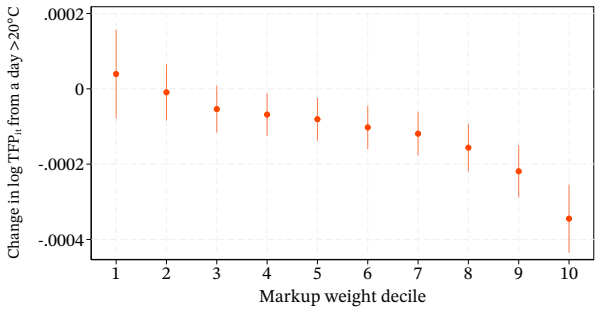
where $r\left(i\right)$ is the NUTS-3 region that firm i resides in. D^{Domar}_{ji} is a dummy variable that equals one if firm i is in the jth decile of the Domar weight distribution and zero otherwise, and $D^{\text{Markup}}_{ji}$is defined similarly.[66] Conditional on firm and sector-year fixed effects, $\text{Tbin}^{>20°C}_{r(i),t}$ can be thought of as a measure of region-level heat shock. The coefficients of interest are thus $\delta^{\text{Domar}}_{>20°\text{C},j}$ and $\delta^{\text{Markup}}_{>20°\text{C},j}$. The first one captures the differential impact on productivity of an extra day of high temperature on firms in Domar weight decile j, compared to those in the first decile. The coefficient $\delta^{\text{Markup}}_{>20°\text{C},j}$ provides similar information but for markups.

Figure 18: Impact of high temperature on productivity by Domar weights and markups

(a) Impact of high temperature on TFP by Domar weight decile



(b) Impact of high temperature on TFP by markup decile



The estimates are plotted in Figure 18. As we can see, the productivity of firms with higher sales Domar weights and lower markups responds less to high temperatures. This is consistent with our model. Indeed, in the model firms find it more beneficial to take measures to counteract the adverse effects of climate change on their operations. As a result, hot temperatures are less detrimental to their productivity.

We can also gauge the quantitative importance of our mechanisms using future climate-change projections. Under the SSP3-7.0 warming scenario, an average NUTS-3 region in Europe is projected

[66] We implicitly choose $10-20°\text{C}$ as the reference temperature bin to avoid perfect collinearity since all temperature bins sum up to 365 days.

to experience about 35 more days with temperatures above 20°C in 2100 relative to the 1951–1980 climate. Our estimates indicate that, due to climate change, firms in the bottom Domar-weight decile will suffer 1.4 percentage points more productivity loss than those in the top decile, while firms in the top markup decile will face 1.3 percentage points more damage than those in the bottom decile. These results suggest that our mechanisms could play a quantitatively important role when assessing the heterogeneous costs of climate change in the cross-section.

D.2 Risk factors estimated from a factor model

Our model assumes that the TFP of a firm follows $a_i(\varepsilon, \delta_i) = \delta_i^\top \varepsilon - b_i(\delta_i)$ which closely resembles a typical dynamic factor model à la Stock and Watson (2002). That literature studies factor models of the form

$$X_{it} = \lambda_i^\top f_t + u_{it},$$

where f_t is a vector of unobserved factors, λ_i is a vector of factor loadings, and u_{it} is an idiosyncratic disturbance. This precisely maps to the structure in our model: X_{it} is firm i's productivity, λ_i is the exposure vector δ_i, and f_t corresponds to the vector of fundamental risks ε. We therefore use technique from the factor model literature to 1) estimate the factors ε and the loadings δ_i, and 2) test whether larger and less distorted firms comove less with "bad" risk factors, as implied by our theory.

A frontier econometric method well suited to our unbalanced firm panel is the Instrumented Principal Component Analysis (IPCA), recently developed by Kelly, Pruitt, and Su (2019) and Kelly, Pruitt, and Su (2020), which allows us to achieve both goals in a single step. We estimate a particular IPCA model for demeaned firm-level TFP growth $\Delta \log \widetilde{TFP}_{it+1}$:

$$\Delta \log \widetilde{TFP}_{it+1} = \alpha_{i,t} + \beta_{i,t}^k \sum_{k=1}^{K} f_{t+1}^k + \epsilon_{i,t+1}$$

with

$$\alpha_{i,t} = \alpha_{\text{Domar}} \log(\text{Domar Weight}_{it}) + \alpha_{\text{Markup}} \log(\text{Markup}_{it}) + \alpha_c + v_{\alpha,i,t},$$

and

$$\beta_{i,t}^k = \beta_{\text{Domar}}^k \log(\text{Domar Weight}_{it}) + \beta_{\text{Markup}}^k \log(\text{Markup}_{it}) + \beta_c^k + v_{\beta,i,t}^k, \quad \text{for all } k.$$

We estimate this system using data on 531,944 firms observed over 25 years. The model features dynamic factor loadings, $\beta_{i,t}^k$, on K latent factors f_{t+1}^k. These loadings potentially depend on three observable instruments—a constant and the firm characteristics $\log(\text{Domar Weight}_{it})$ and

$\log(\text{Markup}_{it})$—as suggested by our theory. Mirroring the structure of $\beta_{i,t}^k$, the intercept $\alpha_{i,t}$ is also expressed as a linear combination of the same instruments. IPCA estimates the model by minimizing the sum of squared residuals across the full panel, yielding the intercept parameters $\alpha_{\text{Domar}}, \alpha_{\text{Markup}}, \alpha_c$; the loading parameters $\{\beta_{\text{Domar}}^k, \beta_{\text{Markup}}^k, \beta_c^k\}_{k=1}^K$; and the K series of latent factors $\{f_{t+1}^k\}_{t=2}^T$.

We estimate two specifications, with $K = 1$ and $K = 2$ factors. The estimated parameters are reported in Table 11, and the time series of the estimated factors are plotted in Figure 19. Figure 19 panel (a) reports the estimated time series of the single risk factor in the $K = 1$ case. Whether this factor is good or bad (positive or negative sign for $\mathcal{E}$), depends on its level μ_k and its covariance with the stochastic discount factor (Lemma 2). Our estimation procedure cannot identify the level of the factor, but it covaries positively with Spanish GDP growth which pushes for making it a bad factor. . Consistent with that interpretation, column $K = 1$ in Table 11 shows that larger and less-markup-distorted firms load less on this factor. When we allow for two risk factors ($K = 2$), the dominant factor still covaries positively with Spanish GDP growth, whereas the second factor covaries negatively. According to our theory, the $k = 2$ factor is a "better" risk than the $k = 1$ factor. As expected, column $k = 2$ in Table **??** shows that larger and less-markup-distorted firms load relatively more on the $k = 2$ factor than on the $k = 1$ factor.[67] Therefore, the cross-sectional patterns of firm-level loadings on the estimated risk factors are also qualitatively consistent with our model.

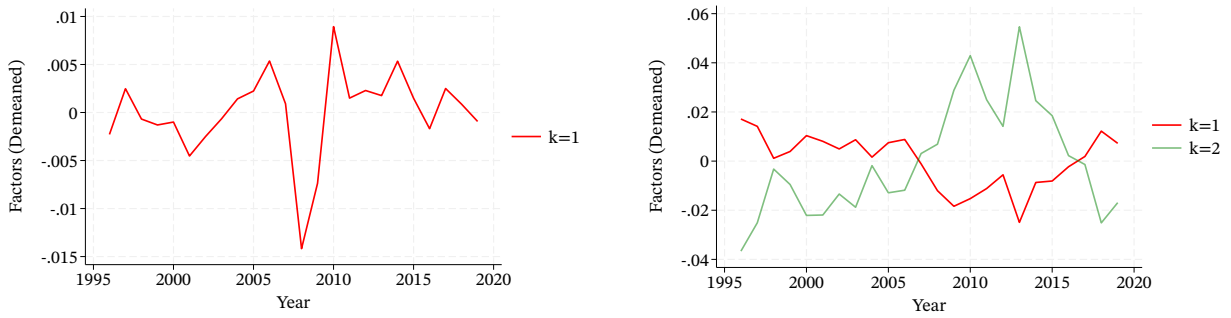
Table 11: Factor model estimation

Estimation cases	$K = 1$	$K = 2$	
Factor	$k = 1$	$k = 1$	$k = 2$
Correlation with GDP growth	0.092	0.710	-0.674
Covariance with GDP growth	9.7×10^{-6}	0.183×10^{-3}	-0.366×10^{-3}
β_{Domar}^k	-0.123	-0.341	-0.083
β_{Markup}^k	0.623	0.780	0.530
R^2	2.77%	2.79%	

[67] There are two ways to interpret $k = 2$ in the context of the model. First, its unobserved level might be negative enough to make it a bad risk factor despite the negative covariance with GDP growth. In this case, the signs of β_{Domar}^k and β_{Markup}^k in Table 11 are natural. Second, it could be that the level of $k = 2$ is not too negative and that its negative cyclicality is strong enough to make it a good risk factor. In this case, the signs of β_{Domar}^k and β_{Markup}^k can be explained by the model if the cost functions g_i and b_i imply that exposures to $k = 1$ and $k = 2$ are complements in the sense of Section 4.2. In that situation, exposures to both factors would move together, which would explain why a good factor ($k = 2$) behaves like a light version of a bad factor ($k = 1$). What is more informative in our view, is the difference between columns $k = 1$ and $k = 2$ in Table 11. Since $k = 2$ is "better" than $k = 1$, larger and less-markup-distorted firms load relatively more on that factor than on $k = 1$, in line with the model.

Figure 19: Estimated factors from the IPCA model

(a) $K = 1$ (b) $K = 2$



E Proofs

E.1 Proof of Lemma 1

Lemma 1. *Log (real) GDP $y = \log Y$ is given by*

$$y(\delta) = \Delta^{\top}\varepsilon - \tilde{\omega}^{\top} b(\delta) - \tilde{\omega}^{\top} \log(1+\tau) - \log \Gamma_L, \tag{17}$$

where the labor share of income Γ_L is given by

$$\Gamma_L := \frac{W_L L}{\bar{P}Y} = 1 - \tau^{\top} \left(diag(1+\tau)\right)^{-1} \omega. \tag{18}$$

Proof. Total profit in this economy is

$$\Pi = \sum_{i=1}^{N} \Pi_i = \sum_{i=1}^{N} \tau_i K_i Q_i - W_R \sum_{i=1}^{N} g_i(\delta_i) = \sum_{i=1}^{N} \frac{\tau_i}{1+\tau_i} \omega_i \bar{P}Y - W_R R,$$

where the last equality follows from the definition of revenue-based Domar weight. This equation implies that the share of profit, gross of risk management expenditure, in household income is

$$\Gamma_{\Pi} := \frac{\Pi + W_R R}{P^{\top} C} = \sum_{i=1}^{N} \frac{\tau_i}{1+\tau_i} \omega_i = \tau^{\top} \left(\text{diag}(1+\tau)\right)^{-1} \omega.$$

Since $\Gamma_L + \Gamma_{\Pi} = 1$, (18) follows. Next, from the definition of the labor share, we can write $\bar{P}Y = \Gamma_L^{-1} W_L L$. Taking the log and combining with (16) and (58), we find

$$y = \log\left(\bar{P}Y\right) - \log \bar{P} = \log\left(\Gamma_L^{-1} W_L L\right) + \tilde{\omega}^{\top}\left(\delta\varepsilon - b(\delta) - \log(1+\tau)\right).$$

Since $W_L L = 1$ given our choice of numeraire, we find (17). □

E.2 Proof of Lemma 2

Lemma 2. *In equilibrium, the risk exposure decision of firm i solves*

$$\underbrace{\mathcal{E} K_i Q_i}_{\substack{\text{Marginal benefit} \\ \text{of risk exposure}}} = \underbrace{\nabla b_i(\delta_i) K_i Q_i + \nabla g_i(\delta_i) W_R}_{\substack{\text{Marginal cost} \\ \text{of risk exposure}}}, \tag{22}$$

where the marginal value of risk exposure $\mathcal{E}$ is defined as

$$\mathcal{E} := \mathrm{E}[\varepsilon] + \mathrm{Cov}[\lambda, \varepsilon], \tag{23}$$

and where $\lambda := \log \Lambda$ is the log of the stochastic discount factor.

Proof. One can show that the firm's problem (21) is strictly convex, such that its first-order conditions are necessary and sufficient. They are given by

$$\mathrm{E}[(\varepsilon - \nabla b_i(\delta_i)) K_i Q_i] + \mathrm{Cov}\left((\varepsilon - \nabla b_i(\delta_i)) K_i Q_i, \frac{\Lambda}{\mathrm{E}[\Lambda]}\right) = \nabla g_i(\delta_i) W_R,$$

where we used the fact that $\partial K_i / \partial \delta_i = -(\varepsilon - \nabla b_i(\delta_i)) K_i$. Next, the definition of revenue-based Domar weights implies that $\omega_i = \frac{P_i Q_i}{PY} \Leftrightarrow Q_i = \frac{\omega_i \Gamma_L^{-1}}{P_i}$. Together with (12), we can therefore write $K_i Q_i = \omega_i \Gamma_L^{-1} / (1 + \tau_i)$, such that $K_i Q_i$ is deterministic. The first-order conditions become

$$\mathrm{E}[\varepsilon] - \nabla b_i(\delta_i) + \mathrm{Cov}(\varepsilon, \lambda) = \nabla g_i(\delta_i) \frac{W_R}{\omega_i \Gamma_L^{-1} / (1 + \tau_i)}, \tag{75}$$

where we have used Stein's Lemma to replace $\Lambda / \mathrm{E}[\Lambda]$ by $\lambda = \log \Lambda$. □

E.3 Proof of Corollary 1

Corollary 1. *Firms with higher revenue-based Domar weights ω_i and lower wedges τ_i manage risk more aggressively, in the sense that*[68]

$$\frac{\partial [\mathcal{E}^\top \delta_i]}{\partial \omega_i} > 0 \qquad \textit{and} \qquad \frac{\partial [\mathcal{E}^\top \delta_i]}{\partial \tau_i} < 0.$$

Proof. From (26), we can write

$$\mathcal{E}^\top \delta_i = \mathcal{E}^\top \delta_i^\circ + \mathcal{E}^\top H_i^{-1} \mathcal{E},$$

where $H_i := B_i + \eta \frac{1+\tau_i}{\omega_i} G_i$. Differentiating with respect to ω_i, we get

$$\frac{\partial [\mathcal{E}^\top \delta_i]}{\partial \omega_i} = -\mathcal{E}^\top H_i^{-1} \frac{\partial H_i}{\partial \omega_i} H_i^{-1} \mathcal{E} = \eta \frac{1 + \tau_i}{\omega_i^2} \mathcal{E}^\top H_i^{-1} G_i H_i^{-1} \mathcal{E} > 0,$$

[68] The partial derivatives imply a partial equilibrium analysis in which $\mathcal{E}$ is kept constant.

where the inequality follows since G_i is positive definite. Differentiating with respect to τ_i instead yields the second inequality. □

E.4 Proof of Lemma 4

Lemma 4. *The aggregate cost function $\bar{h}$ is given by*

$$\bar{h}(\Delta) = \frac{1}{2}(\Delta - \Delta^\circ)^\top \bar{H}(\Delta - \Delta^\circ), \tag{32}$$

where $\Delta^\circ := (\delta^\circ)^\top \tilde{\omega}$. Furthermore, the Hessian matrix of $\bar{h}$ is given by

$$\bar{H} = \left[\sum_{i=1}^N \tilde{\omega}_i H_i^{-1}\right]^{-1}, \tag{33}$$

where H_i is given by (24).

Proof. The Lagrangian of problem (31) is

$$\mathscr{L} = \sum_{i=1}^N \tilde{\omega}_i b_i(\delta_i) + \eta \sum_{i=1}^N \kappa_i g_i(\delta_i) - \sum_{m=1}^M \nu_m \left(\Delta_m - \mathbf{1}_m^\top \delta^\top \tilde{\omega}\right), \tag{76}$$

where ν_m is the Lagrange multiplier on the mth row of the constraint $\Delta = \delta^\top \tilde{\omega}$. The first-order condition with respect to δ_i is

$$\tilde{\omega}_i \nabla b_i + \eta \kappa_i \nabla g_i = -\nu \tilde{\omega}_i, \tag{77}$$

which implies that for all i, k,

$$\nabla b_i + \eta \frac{\kappa_i}{\tilde{\omega}_i} \nabla g_i = \nabla b_k + \eta \frac{\kappa_k}{\tilde{\omega}_k} \nabla g_k = \nabla \bar{h}, \tag{78}$$

where the last equality comes from the envelope theorem. Together with (5) and (4), this implies that

$$\delta_i = \delta_i^\circ + \left(B_i + \eta \frac{\kappa_i}{\tilde{\omega}_i} G_i\right)^{-1} \left(B_k + \eta \frac{\kappa_k}{\tilde{\omega}_k} G_k\right)(\delta_k - \delta_k^\circ).$$

Then, the constraint $\Delta = \delta^\top \tilde{\omega}$ can be rewritten as

$$\begin{aligned} \Delta = \sum_{i=1}^N \tilde{\omega}_i \delta_i = \Delta^\circ + \sum_{i=1}^N \tilde{\omega}_i \left(B_i + \eta \frac{\kappa_i}{\tilde{\omega}_i} G_i\right)^{-1} \left(B_k + \eta \frac{\kappa_k}{\tilde{\omega}_k} G_k\right)(\delta_k - \delta_k^\circ) \Leftrightarrow \\ \delta_k - \delta_k^\circ = H_k^{-1} \bar{H}(\Delta - \Delta^\circ). \end{aligned} \tag{79}$$

where H_k is given by (24) and $\bar{H}$ is given by (33). Plugging this last expression into (31) yields the result. □

E.5 Proof of Proposition 1

Proposition 1. *There exists a unique equilibrium, and its aggregate risk exposure Δ^* solves*

$$\mathcal{W}_{dist} := \max_{\Delta} \underbrace{\Delta^\top \mu - \bar{b}(\Delta) - \tilde{\omega}^\top \log(1+\tau) - \log \Gamma_L}_{\mathrm{E}[y]} - \frac{1}{2}(\rho-1)\underbrace{\Delta^\top \Sigma \Delta}_{\mathrm{V}[y]} - \bar{g}(\Delta), \tag{34}$$

where $\bar{b}(\Delta) + \bar{g}(\Delta) = \bar{h}(\Delta)$ is given by (32). Furthermore, without wedges ($\tau = 0$), the equilibrium is efficient.

Proof. We first show that the set of equilibrium allocations coincides with the set of solutions to a maximization problem. Consider the maximization problem

$$\max_{\delta} \tilde{\omega}^\top \delta\mu - \tilde{\omega}^\top b(\delta) - \tilde{\omega}^\top \log(1+\tau) - \log \Gamma_L - \frac{1}{2}(\rho-1)\tilde{\omega}^\top \delta\Sigma\delta^\top\tilde{\omega} + \log\left(V\left(\sum_{i=1}^N \frac{\tilde{\omega}_i(1+\tau_i)}{\omega_i} g_i(\delta_i)\right)\right). \tag{80}$$

Notice that we can write this problem as

$$\max_{\Delta} \Delta^\top \mu - \tilde{\omega}^\top \log(1+\tau) - \log \Gamma_L - \frac{1}{2}(\rho-1)\Delta^\top\Sigma\Delta - \min_{\delta \text{ s.t. } \Delta = \delta^\top\tilde{\omega}} \left\{ \tilde{\omega}^\top b(\delta) - \log\left(V\left(\sum_{i=1}^N \frac{\tilde{\omega}_i(1+\tau_i)}{\omega_i} g_i(\delta_i)\right)\right)\right\},$$

which is the maximization problem (34). Note also that the objective function of this problem is strictly concave if $\bar{h}(\Delta) = \bar{b}(\Delta) + \bar{g}(\Delta)$ is strictly convex. But recall from Lemma 4 that $\bar{h}$'s Hessian satisfies

$$\bar{H} = \left[\sum_{i=1}^N \tilde{\omega}_i H_i^{-1}\right]^{-1}. \tag{81}$$

Since H_i is positive definite, so is H_i^{-1} and so is the right-hand side of this equation. It follows that $\bar{H}$ is also positive definite, and $\bar{h}$ is therefore strictly convex. This implies that there is a unique solution to the maximization problems (80) and (34) and that the first-order conditions are sufficient to characterize it. To complete the proof, we will show that the equilibrium conditions coincide with these first-order conditions. This will imply that there is a unique equilibrium.

The first-order condition of the fictitious planner with respect to δ_{im} is

$$\tilde{\omega}_i\mu_m - \tilde{\omega}_i \frac{db_i(\delta_i)}{d\delta_{im}} - (\rho-1)\tilde{\omega}_i\tilde{\omega}^\top\delta\Sigma\mathbf{1}_m + \frac{V'\left(\sum_{i=1}^N \frac{\tilde{\omega}_i(1+\tau_i)}{\omega_i} g_i(\delta_i)\right)}{V\left(\sum_{i=1}^N \frac{\tilde{\omega}_i(1+\tau_i)}{\omega_i} g_i(\delta_i)\right)} \frac{\tilde{\omega}_i(1+\tau_i)}{\omega_i} \frac{dg_i(\delta_i)}{d\delta_{im}} = 0.$$

Because of V's exponential form, this expression simplifies to

$$\tilde{\omega}_i \mu_m - \tilde{\omega}_i \frac{db_i(\delta_i)}{d\delta_{im}} - (\rho - 1)\,\tilde{\omega}_i \tilde{\omega}^\top \delta \Sigma \mathbf{1}_m - \eta \frac{\tilde{\omega}_i (1+\tau_i)}{\omega_i} \frac{dg_i(\delta_i)}{d\delta_{im}} = 0. \tag{82}$$

The system of these $M \times N$ equations fully characterizes the unique solution to the fictitious planner's problem.

Now consider the equilibrium. Recall from (75) that the first-order condition of the firm implies that

$$\mathrm{E}[\varepsilon] - \nabla b_i(\delta_i) + \mathrm{Cov}(\varepsilon, \lambda) = \nabla g_i(\delta_i) \frac{W_R}{\omega_i \Gamma_L^{-1} / (1+\tau_i)}.$$

Then, using (56), we can write the stochastic discount factor as

$$\Lambda = \mathcal{U}'(Y) (V(R))^{1-\rho} \bar{P}^{-1} = \left(\Gamma_L^{-1} W_L L\right)^{-\rho} (V(R))^{1-\rho} \bar{P}^{\rho-1}. \tag{83}$$

Taking the log, we get

$$\lambda = \log\left(\left(\Gamma_L^{-1} W_L L\right)^{-\rho} (V(R))^{1-\rho}\right) - (\rho - 1)\,\beta^\top \tilde{\mathcal{L}}\left(\delta\varepsilon - b(\delta) - \log(1+\tau)\right), \tag{84}$$

where we have used the fact that $\bar{p} = \beta^\top p$ together with (16). The first-order condition therefore becomes

$$\mu - \nabla b_i(\delta_i) - (\rho - 1)\,\tilde{\omega}^\top \delta \Sigma = \nabla g_i(\delta_i) \frac{W_R}{\omega_i \Gamma_L^{-1} / (1+\tau_i)}.$$

Finally, from the definition of the labor share, we can write $\bar{P} Y = \Gamma_L^{-1} W_L L = \Gamma_L^{-1}$. It follows from (9) that $\eta \Gamma_L^{-1} = W_R$, and so the first-order conditions of the firms are equivalent to those of the fictitious planner (35), which completes the proof. □

E.6 Proof of Lemma 5

Lemma 5. *The equilibrium aggregate risk exposure Δ solves*

$$\mathcal{E}(\Delta) = \nabla \bar{h}(\Delta), \tag{35}$$

where the marginal value of aggregate risk exposure $\mathcal{E}$ can be written as

$$\mathcal{E} = \underbrace{\mu}_{\mathrm{E}[\varepsilon]} \underbrace{- (\rho - 1)\,\Sigma \Delta}_{\mathrm{Cov}[\lambda, \varepsilon]}. \tag{23}$$

Proof. The fictitious planner's first-order conditions are given by

$$\mu - (\rho - 1)\,\Delta \Sigma = \nabla \bar{h}(\Delta).$$

Combining the definition of $\mathcal{E}$ given by (23) with (84), we can write

$$\mathcal{E} = \mathrm{E}\left[\varepsilon\right] + \mathrm{Cov}\left[-\left(\rho-1\right)\tilde{\omega}^{\top}\delta\varepsilon, \varepsilon\right],$$

which with the definition of Δ yields the result. □

E.7 Proof of Proposition 2

Proposition 2. *The equilibrium aggregate risk-exposure decisions are given by*

$$\Delta = \Delta^{\circ} + \mathcal{H}^{-1}\mathcal{E}^{\circ}, \tag{38}$$

where $\mathcal{E}^{\circ} := \mu - \left(\rho-1\right)\Sigma\Delta^{\circ}$ *and where the* $M \times M$ *positive definite matrix* $\mathcal{H}^{-1}$ *is*

$$\mathcal{H}^{-1} := \left(\bar{H} + \left(\rho-1\right)\Sigma\right)^{-1}. \tag{39}$$

Proof. We can write the fictitious planner's first-order condition as

$$\Delta - \Delta^{\circ} = \bar{H}^{-1}\mathcal{E}\left(\Delta\right). \tag{85}$$

It follows that

$$\begin{gathered}\mathcal{E} - \bar{H}\left(\Delta - \Delta^{\circ}\right) = 0 \Leftrightarrow \Delta = \Delta^{\circ} + \bar{H}^{-1}\left(\mu - \left(\rho-1\right)\Sigma\Delta\right) \Leftrightarrow \\ \Delta = \mathcal{H}^{-1}\left(\bar{H}\Delta^{\circ} + \mu\right) = \Delta^{\circ} + \mathcal{H}^{-1}\mathcal{E}^{\circ},\end{gathered}$$

where $\mathcal{E}^{\circ} = \mu - \left(\rho-1\right)\Sigma\Delta^{\circ}$. □

E.8 Proof of Proposition 3

Proposition 3. *Let* γ *denote either the mean* μ_m *or an element* Σ_{mn} *of the covariance matrix. The response of the equilibrium aggregate risk exposure* Δ *to a change in* γ *is given by*

$$\frac{d\Delta}{d\gamma} = \mathcal{H}^{-1}\frac{\partial\mathcal{E}}{\partial\gamma}, \tag{40}$$

where $\partial\mathcal{E}/\partial\gamma$ *is given by* (37).

Proof. From (38), we can write

$$\frac{d\Delta}{d\mu} = \mathcal{H}^{-1}\frac{d\mathcal{E}^{\circ}}{d\mu} = \mathcal{H}^{-1}\frac{\partial\mathcal{E}}{\partial\mu}.$$

For Σ,

$$\begin{aligned}\frac{d\Delta}{d\Sigma_{mn}} &= \frac{d\mathcal{H}^{-1}}{d\Sigma_{mn}}\mathcal{E}^{\circ} + \mathcal{H}^{-1}\frac{d\mathcal{E}^{\circ}}{d\Sigma_{mn}} = -(\rho-1)\,\mathcal{H}^{-1}\frac{d\Sigma}{d\Sigma_{mn}}\mathcal{H}^{-1}\mathcal{E}^{\circ} - (\rho-1)\,\mathcal{H}^{-1}\frac{d\Sigma}{d\Sigma_{mn}}\Delta^{\circ} \\ &= -(\rho-1)\,\mathcal{H}^{-1}\frac{d\Sigma}{d\Sigma_{mn}}\Delta,\end{aligned}$$

and the result follows from (37). □

E.9 Proof of Corollary 2

Corollary 2. *An increase in the expected value μ_m of risk factor m leads to an increase in aggregate risk exposure Δ_m. An increase in the variance Σ_{mm} of risk factor m leads to a decrease in Δ_m if $\Delta_m > 0$ and to an increase in Δ_m if $\Delta_m < 0$.*

Proof. By Proposition 3, $\mathcal{H}$ is positive definite and, therefore, so is $\mathcal{H}^{-1}$. It follows that the diagonal elements of $\mathcal{H}^{-1}$ are positive and the result follows from (37). □

E.10 Proof of Proposition 4

Proposition 4. *The response of the equilibrium aggregate risk exposure Δ to a change in wedge τ_i is given by*

$$\frac{d\Delta}{d\tau_i} = \frac{d\mathcal{H}^{-1}}{d\tau_i}\mathcal{E}^{\circ}, \tag{41}$$

where

$$\frac{d\mathcal{H}^{-1}}{d\tau_i} = -\mathcal{H}^{-1}\left(\sum_{j=1}^{N}\frac{\partial\bar{H}}{\partial\kappa_j}\frac{d\kappa_j}{d\tau_i}\right)\mathcal{H}^{-1},$$

and where the response $d\kappa_j/d\tau_i$ of the efficiency gap is nonnegative, and $\partial\bar{H}/\partial\kappa_j$ is a positive definite matrix.

Proof. The first equation follows directly from (38). Differentiating $\mathcal{H}^{-1}$ with respect to τ_i, we find

$$\frac{d\mathcal{H}^{-1}}{d\tau_i} = -\mathcal{H}^{-1}\left(\sum_{j=1}^{N}\frac{\partial\bar{H}}{\partial\kappa_j}\frac{d\kappa_j}{d\tau_i}\right)\mathcal{H}^{-1}.$$

Differentiating $\kappa_i = (1+\tau_i)\,\frac{\tilde{\omega}_i}{\omega_i}$ with respect to τ_j yields

$$\frac{d\kappa_i}{d\tau_j} = \frac{\kappa_i}{\omega_i}\frac{\omega_j}{1+\tau_j}\mathcal{L}_{ji} \geq 0. \tag{86}$$

Finally, using (33), we find

$$\frac{\partial\bar{H}}{\partial\kappa_i} = \bar{H}\left(\eta H_i^{-1}G_iH_i^{-1}\right)\bar{H}. \tag{87}$$

Since $\bar{H}$, H_i^{-1}, and G_i are positive definite, so is $\frac{\partial \bar{H}}{\partial \kappa_i}$. □

E.11 Proof of Corollary 3

Corollary 3. *In a diagonal economy, a higher wedge τ_i increases Δ_m for all m such that $\mathcal{E}_m < 0$ (bad risks) and decreases Δ_m for all m such that $\mathcal{E}_m > 0$ (good risks).*

Proof. By Proposition 4,

$$\frac{d\Delta}{d\tau_i} = -\mathcal{H}^{-1}\left(\sum_{j=1}^{N} \bar{H}\left(\eta H_i^{-1} G_i H_i^{-1}\right)\bar{H}\frac{d\kappa_j}{d\tau_i}\right)\mathcal{H}^{-1}\mathcal{E}^{\circ}.$$

Recall that $\frac{d\kappa_j}{d\tau_i} \geq 0$ and $\frac{d\kappa_i}{d\tau_i} > 0$ by (86). If Σ, G_i, and H_i are diagonal for all i, then $\mathcal{H}^{-1}\left(\sum_{j=1}^{N} \bar{H}\left(\eta H_i^{-1} G_i H_i^{-1}\right)\bar{H}\frac{d\kappa_j}{d\tau_i}\right)\mathcal{H}^{-1}$ is a positive diagonal matrix. Therefore, the sign of $\frac{d\Delta_m}{d\tau_i}$ is the opposite of the sign of $\mathcal{E}_m^{\circ}$. Finally, from the first-order condition of the fictitious planner we have

$$\mathcal{E} - \bar{H}\left(\Delta - \Delta^{\circ}\right) = 0 \Leftrightarrow \mathcal{E} = \bar{H}\mathcal{H}^{-1}\mathcal{E}^{\circ}. \tag{88}$$

It follows that $\mathcal{E}_m$ and $\mathcal{E}_m^{\circ}$ have the same sign and the result follows. □

E.12 Proof of Lemma 6

Lemma 6. *Suppose that $\tau_j > 0$ for at least one firm j. Then $\left(\Delta - \Delta_{SP}\right)^{\top}\mathcal{E}^{\circ} < 0$, where Δ and Δ_{SP} are the aggregate risk exposure vectors in the equilibrium and the efficient allocation, respectively. Furthermore, in a diagonal economy the sign of $\Delta_i - \Delta_{SP,i}$ is the opposite of that of $\mathcal{E}_i^{\circ}$.*

Proof. We first show that $\mathcal{H}_{SP}^{-1} - \mathcal{H}^{-1}$ is positive definite. For any two positive definite matrices A and B, if $A - B$ is positive definite, so is $B^{-1} - A^{-1}$. It follows that $\mathcal{H}_{SP}^{-1} - \mathcal{H}^{-1}$ is positive definite if $\mathcal{H} - \mathcal{H}_{SP}$ is. From (39),

$$\mathcal{H} - \mathcal{H}_{SP} = \bar{H} - \bar{H}_{SP}.$$

It follows that $\mathcal{H}_{SP}^{-1} - \mathcal{H}^{-1}$ is positive definite if $\bar{H}_{SP}^{-1} - \bar{H}^{-1}$ is. From (33), we can write

$$\bar{H}_{SP}^{-1} - \bar{H}^{-1} = \sum_{i=1}^{N} \tilde{\omega}_i \left(\left(B_i + \eta\frac{1}{\tilde{\omega}_i}G_i\right)^{-1} - \left(B_i + \eta\frac{\kappa_i}{\tilde{\omega}_i}G_i\right)^{-1}\right).$$

Recall that $\kappa_i \geq 1$, with the equality holding for all firms simultaneously only if $\tau = 0$. Then $B_i + \eta\frac{\kappa_i}{\tilde{\omega}_i}G_i - B_i + \eta\frac{1}{\tilde{\omega}_i}G_i = \eta\frac{\kappa_i - 1}{\tilde{\omega}_i}G_i$ is a positive semidefinite matrix, and so is $\left(B_i + \eta\frac{1}{\tilde{\omega}_i}G_i\right)^{-1} - \left(B_i + \eta\frac{\kappa_i}{\tilde{\omega}_i}G_i\right)^{-1}$. Then $\bar{H}_{SP}^{-1} - \bar{H}^{-1}$ is positive definite, and so is $\mathcal{H}_{SP}^{-1} - \mathcal{H}^{-1}$. Using Proposition 2,

we can write

$$\left(\Delta - \Delta^{SP}\right)^{\top} \mathcal{E}^{\circ} = \left(\mathcal{H}^{-1}\mathcal{E}^{\circ} - \mathcal{H}_{SP}^{-1}\mathcal{E}^{\circ}\right)^{\top} \mathcal{E}^{\circ} = \left(\mathcal{E}^{\circ}\right)^{\top} \left(\mathcal{H}^{-1} - \mathcal{H}_{SP}^{-1}\right) \mathcal{E}^{\circ} < 0,$$

which is the result of the first part of this lemma.

Furthermore, if the risk factors are uncorrelated (diagonal Σ), and that individual risk exposures are neither complements nor substitutes in the cost functions $(b_1, \ldots, b_N)$ and $(g_1, \ldots, g_N)$ (diagonal B_i and G_i for all i), then $\mathcal{H}$ and $\mathcal{H}_{SP}$ are diagonal matrices. Therefore, the sign of $\Delta_i - \Delta_i^{SP}$ is the opposite of the sign of $\mathcal{E}_i^{\circ}$, which is the result. □

E.13 Proof of Corollary 4

Corollary 4. *Let $B_i = 0$ for all i. In a diagonal economy, the following statements hold.*

1. *The impact of an increase in μ_m on GDP satisfies*

$$sign\left(\frac{d\,\mathrm{E}\left[y\right]}{d\mu_m} - \frac{\partial\,\mathrm{E}\left[y\right]}{\partial\mu_m}\right) = sign\left(\mu_m\right) \quad and \quad sign\left(\frac{d\,\mathrm{V}\left[y\right]}{d\mu_m} - \frac{\partial\,\mathrm{V}\left[y\right]}{\partial\mu_m}\right) = sign\left(\Delta_m\right). \tag{43}$$

2. *The impact of an increase in Σ_{mm} on GDP satisfies*

$$sign\left(\frac{d\,\mathrm{E}\left[y\right]}{d\Sigma_{mm}} - \frac{\partial\,\mathrm{E}\left[y\right]}{\partial\Sigma_{mm}}\right) = -sign\left(\mu_m\Delta_m\right) \quad and \quad \frac{d\,\mathrm{V}\left[y\right]}{d\Sigma_{mm}} - \frac{\partial\,\mathrm{V}\left[y\right]}{\partial\Sigma_{mm}} < 0. \tag{44}$$

Proof. From the definition of $\bar{b}$ and using (79), we get

$$\bar{b}\left(\Delta\right) = \tilde{\omega}^{\top} b\left(\delta^{*}\left(\Delta\right)\right) = \frac{1}{2}\left(\Delta - \Delta^{\circ}\right)^{\top} \bar{H} \left(\sum_{i=1}^{N} \tilde{\omega}_i H_i^{-1} B_i H_i^{-1}\right) \bar{H}\left(\Delta - \Delta^{\circ}\right).$$

Therefore,

$$\nabla\bar{b} = \bar{H}\left(\sum_{i=1}^{N} \tilde{\omega}_i H_i^{-1} B_i H_i^{-1}\right) \mathcal{E},$$

where we have used (85). From (33), we see that given our assumptions, $\bar{H}$ is diagonal with positive entries. This implies that $\mathcal{H}^{-1}$ is also diagonal with positive entries. Combining the first equation in (42) with (40) and (37), we can write

$$\operatorname{sign}\left(\frac{d\,\mathrm{E}\left[y\right]}{d\mu_m} - \frac{\partial\,\mathrm{E}\left[y\right]}{\partial\mu_m}\right) = \operatorname{sign}\left(\left(\mu - \bar{H}\left(\sum_{i=1}^{N} \tilde{\omega}_i H_i^{-1} B_i H_i^{-1}\right)\mathcal{E}\right)^{\top} \mathcal{H}^{-1}\mathbf{1}_m\right).$$

This expression is general, but if there are no productivity costs of adjusting risk exposure, $B_i = 0$

$\forall i$, we get

$$\text{sign}\left(\frac{d\,\mathrm{E}\left[y\right]}{d\mu_m}-\frac{\partial\,\mathrm{E}\left[y\right]}{\partial\mu_m}\right)=\text{sign}\left(\mu_m\right).$$

If instead there are no resource costs of adjusting risk exposure, $G_i=0\ \forall i$, we get $H_i=B_i$ and so

$$\text{sign}\left(\frac{d\,\mathrm{E}\left[y\right]}{d\mu_m}-\frac{\partial\,\mathrm{E}\left[y\right]}{\partial\mu_m}\right)=\text{sign}\left(\Delta_m\right).$$

Combining the second expression in (42) with (40) and (37), we can write

$$\text{sign}\left(\frac{d\,\mathrm{V}\left[y\right]}{d\mu_m}-\frac{\partial\,\mathrm{V}\left[y\right]}{\partial\mu_m}\right)=\text{sign}\left(2\Delta^\top\Sigma\mathcal{H}^{-1}\mathbf{1}_m\right)=\text{sign}\left(\Delta_m\right).$$

We can follow the same procedure for Σ_{mm}. From the first equation in (42) together with (40) and (37), we find

$$\text{sign}\left(\frac{d\,\mathrm{E}\left[y\right]}{d\Sigma_{mm}}-\frac{\partial\,\mathrm{E}\left[y\right]}{\partial\Sigma_{mm}}\right)=\text{sign}\left(-\left(\mu-\bar{H}\left(\sum_{i=1}^{N}\tilde{\omega}_i H_i^{-1}B_iH_i^{-1}\right)\mathcal{E}\right)^\top\mathcal{H}^{-1}\left(\rho-1\right)\Delta_m\mathbf{1}_m\right).$$

If $B_i=0\ \forall i$, we have

$$\text{sign}\left(\frac{d\,\mathrm{E}\left[y\right]}{d\Sigma_{mm}}-\frac{\partial\,\mathrm{E}\left[y\right]}{\partial\Sigma_{mm}}\right)=\text{sign}\left(-\mu_m\Delta_m\right).$$

If $G_i=0\ \forall i$, we have

$$\text{sign}\left(\frac{d\,\mathrm{E}\left[y\right]}{d\Sigma_{mm}}-\frac{\partial\,\mathrm{E}\left[y\right]}{\partial\Sigma_{mm}}\right)=\text{sign}\left(-\Delta^\top\Sigma\mathcal{H}^{-1}\Delta_m\mathbf{1}_m\right)<0.$$

From the second expression in (42) together with (40) and (37), we can write

$$\frac{d\,\mathrm{V}\left[y\right]}{d\Sigma_{mm}}-\frac{\partial\,\mathrm{V}\left[y\right]}{\partial\Sigma_{mm}}=-2\left(\rho-1\right)\Delta^\top\Sigma\mathcal{H}^{-1}\Delta_m\mathbf{1}_m,$$

which is always negative given our assumptions. □

E.14 Proof of Corollary 5

Corollary 5. *Let $B_i=0$ for all i, and suppose that there is a single risk factor. Then*

$$sign\left(\frac{d\,\mathrm{E}\left[y\right]}{d\tau_i}-\frac{\partial\,\mathrm{E}\left[y\right]}{\partial\tau_i}\right)=-sign\left(\mu\mathcal{E}\right)\qquad and\qquad sign\left(\frac{d\,\mathrm{V}\left[y\right]}{d\tau_i}-\frac{\partial\,\mathrm{V}\left[y\right]}{\partial\tau_i}\right)=-sign\left(\Delta\mathcal{E}\right).\quad(45)$$

Proof. Suppose that there is only one risk factor. Then all $M\times M$ positive definite matrices are

simply positive scalars. From Proposition 5, we can write

$$\operatorname{sign}\left(\frac{d\,\mathrm{E}\left[y\right]}{d\tau_i}-\frac{\partial\,\mathrm{E}\left[y\right]}{\partial\tau_i}\right)=\operatorname{sign}\left(-\left(\mu-\bar{H}\left(\sum_{i=1}^{N}\tilde{\omega}_i H_i^{-1}B_i H_i^{-1}\right)\mathcal{E}\right)\mathcal{H}^{-1}\left(\sum_{j=1}^{N}\frac{\partial\bar{H}}{\partial\kappa_j}\frac{d\kappa_j}{d\tau_i}\right)\bar{H}^{-1}\mathcal{E}\right)=$$
$$=\operatorname{sign}\left(-\left(\mu-\bar{H}\left(\sum_{i=1}^{N}\tilde{\omega}_i H_i^{-1}B_i H_i^{-1}\right)\mathcal{E}\right)\mathcal{E}\right).$$

If there are no productivity costs of adjusting risk exposure, $B_i=0\ \forall i$, we get

$$\operatorname{sign}\left(\frac{d\,\mathrm{E}\left[y\right]}{d\tau_i}-\frac{\partial\,\mathrm{E}\left[y\right]}{\partial\tau_i}\right)=-\operatorname{sign}\left(\mu\mathcal{E}\right).$$

If there are no resource costs of adjusting risk exposure, $G_i=0\ \forall i$, we get $H_i=B_i$, and so

$$\operatorname{sign}\left(\frac{d\,\mathrm{E}\left[y\right]}{d\tau_i}-\frac{\partial\,\mathrm{E}\left[y\right]}{\partial\tau_i}\right)=-\operatorname{sign}\left(\Delta\mathcal{E}\right).$$

Combining the second expression in (42) with (41), we can write

$$\operatorname{sign}\left(\frac{d\,\mathrm{V}\left[y\right]}{d\tau_i}-\frac{\partial\,\mathrm{V}\left[y\right]}{\partial\tau_i}\right)=\operatorname{sign}\left(-2\Delta\Sigma\mathcal{H}^{-1}\left(\sum_{j=1}^{N}\frac{\partial\bar{H}}{\partial\kappa_j}\frac{d\kappa_j}{d\tau_i}\right)\bar{H}^{-1}\mathcal{E}\right),$$

and the second result follows. □

E.15 Proof of Lemma 7

Before proving Lemma 7, we formulate a few auxiliary results. Define

$$\bar{h}_V\left(\Delta\right):=\tilde{\omega}^{\top}b\left(\delta\left(\Delta\right)\right)-\log V\left(\sum_{i=1}^{N}g_i\left(\delta_i\left(\Delta\right)\right)\right),\tag{89}$$

where $\delta_i\left(\Delta\right)$ is the distorted equilibrium exposure given by (26). Unlike $\bar{h}$, the weights of the different firms in $\bar{h}_V$ are not themselves distorted by the efficiency gaps $(\kappa_1,\ldots,\kappa_N)$. As a result, $\bar{h}_V$ properly measures the disutility of equilibrium risk management decisions, and we can use it to compute welfare. The function $\bar{h}_V$ also differs from the planner's cost function $\bar{h}_{SP}$ since it uses the (distorted) equilibrium risk-exposure matrix δ as input instead of its efficient counterpart.

We can use $\bar{h}_V$ to describe how a change in the environment affects welfare. As before, we focus on the role played by the response of the exposure vector Δ by filtering out the fixed-exposure effect $\partial\mathcal{W}/\partial\chi$.

Proposition 7. *Let χ denote either μ_m, Σ_{mn}, or τ_i. Then the impact of a change in χ on welfare*

is given by

$$\frac{d\mathcal{W}}{d\chi} - \frac{\partial \mathcal{W}}{\partial \chi} = \left(\mathcal{E} - \nabla \bar{h}_V\right)^\top \frac{d\Delta}{d\chi} = \left(\nabla \bar{h} - \nabla \bar{h}_V\right)^\top \frac{d\Delta}{d\chi}, \tag{90}$$

where the partial derivatives indicate that Δ is kept fixed, and where $d\Delta/d\chi$ is given by (40) for $\chi = \mu_i$ or Σ_{mn}, and by (41) for $\chi = \tau_i$.

Proof. Combining the definition of $\bar{h}_V$ with (26) and (85), we can write

$$\bar{h}_V = \frac{1}{2}\left(\Delta - \Delta^\circ\right)^\top \bar{H} \left[\sum_{i=1}^N \tilde{\omega}_i H_i^{-1} \left(B_i + \eta \frac{1}{\tilde{\omega}_i} G_i\right) H_i^{-1}\right] \bar{H} \left(\Delta - \Delta^\circ\right).$$

This equation implies that

$$\nabla \bar{h}_V = \bar{H}_V \left(\Delta - \Delta^\circ\right),$$

where

$$\bar{H}_V = \bar{H} \left[\sum_{i=1}^N \tilde{\omega}_i H_i^{-1} \left(B_i + \eta \frac{1}{\tilde{\omega}_i} G_i\right) H_i^{-1}\right] \bar{H}.$$

Simple algebra implies that

$$\nabla \bar{h} - \nabla \bar{h}_V = \bar{H} \left(\sum_{i=1}^N \eta \left(\kappa_i - 1\right) H_i^{-1} G_i H_i^{-1}\right) \bar{H} \left(\Delta - \Delta^\circ\right). \tag{91}$$

We can write welfare as

$$\mathcal{W} = \mathcal{W}_{dist} + \bar{h}\left(\Delta\right) - \bar{h}_V\left(\Delta\right).$$

Differentiating with respect to χ yields

$$\frac{d\mathcal{W}}{d\chi} - \frac{\partial \mathcal{W}}{\partial \chi} = \left(\nabla \bar{h} - \nabla \bar{h}_V\right)^\top \frac{d\Delta}{d\chi},$$

where $\frac{d\mathcal{W}_{dist}}{d\chi} - \frac{\partial \mathcal{W}_{dist}}{\partial \chi} = 0$ by the envelope theorem. □

The first equality in (90) highlights the two channels through which a change in χ affects welfare via Δ. First, the response $d\Delta/d\chi$ of the risk-exposure vector triggers a change in the expected utility of consumption that is proportional to the marginal benefit of that exposure, as captured by $\mathcal{E}$. For instance, an increase in the expected value μ_m of a good risk factor triggers an increase in welfare through that channel. Second, the response of Δ also triggers a change in risk management costs that is proportional to $\nabla \bar{h}_V$. If the economy is already heavily exposed to ε_m ($\Delta_m > \Delta_m^\circ$), the same increase in μ_m leads to a large increase in costs, which reduces welfare.

The second equality in (90) follows since in equilibrium the benefit of exposure $\mathcal{E}$ must be equal to its marginal cost as it is *perceived* by the fictitious planner, $\nabla \bar{h}$. When there are distortions, this

perceived cost deviates from the *true* marginal cost of exposure $\nabla\bar{h}_V$, and welfare can be affected by Δ as a result. The gap $\nabla\bar{h} - \nabla\bar{h}_V$ therefore provides a measure of how important distortions are for the impact of risk exposure on welfare.

To further explore the impact of χ on welfare, we therefore need to characterize the gap between $\nabla\bar{h}$ and $\nabla\bar{h}_V$. In general, one can compute these objects and look at their difference. We can however sign that gap in a simple way in a diagonal economy.

Lemma 9. *Suppose that the economy is diagonal and that $\tau_j > 0$ for at least one firm j. Then the sign of $\left[\nabla\bar{h}\right]_i - \left[\nabla\bar{h}_V\right]_i$ is the same as the sign of $\mathcal{E}_i$.*

Proof. Combining (91) with (85), we can write

$$\nabla\bar{h} - \nabla\bar{h}_V = \bar{H}\left(\sum_{i=1}^{N}\eta\left(\kappa_i - 1\right)H_i^{-1}G_iH_i^{-1}\right)\mathcal{E}. \tag{92}$$

Under our assumptions, $\bar{H}\left(\sum_{i=1}^{N}\eta\left(\kappa_i - 1\right)H_i^{-1}G_iH_i^{-1}\right)$ is a positive definite diagonal matrix. Therefore, the sign of $\left[\nabla\bar{h}\right]_i - \left[\nabla\bar{h}_V\right]_i$ is the same as the sign of $\mathcal{E}_i$. □

Since in equilibrium we must have $\mathcal{E} = \nabla\bar{h}$, the gradient $\left[\nabla\bar{h}\right]_i$ must be positive if ε_i is a good risk factor and negative otherwise. Because of the wedges, however, the fictitious planner perceives that adjusting Δ is costlier than it really is, which implies that $\bar{h}_V$ is flatter than $\bar{h}$. It follows that for a good risk factor, the true marginal cost is lower than the perceived one, and $\left[\nabla\bar{h}\right]_i - \left[\nabla\bar{h}_V\right]_i > 0$. The opposite is true for a bad risk factor.

Lemma 7. *Suppose that the economy is diagonal and that $\tau_j > 0$ for at least one firm j. Then,*

$$sign\left(\frac{d\mathcal{W}}{d\mu_m} - \frac{\partial\mathcal{W}}{\partial\mu_m}\right) = sign\left(\mathcal{E}_m\right) \qquad and \qquad sign\left(\frac{d\mathcal{W}}{d\Sigma_{mm}} - \frac{\partial\mathcal{W}}{\partial\Sigma_{mm}}\right) = -sign\left(\Delta_m\mathcal{E}_m\right).$$

Proof. Using (92), we can write (90) as

$$\frac{d\mathcal{W}}{d\chi} - \frac{\partial\mathcal{W}}{\partial\chi} = \left(\frac{d\Delta}{d\chi}\right)^{\top}\bar{H}\left(\sum_{i=1}^{N}\eta\left(\kappa_i - 1\right)H_i^{-1}G_iH_i^{-1}\right)\mathcal{E},$$

From Proposition 3, we have

$$\frac{d\Delta}{d\gamma} = \mathcal{H}^{-1}\frac{\partial\mathcal{E}}{\partial\gamma}.$$

Setting $\chi = \gamma$ and combining the last two equations, we find

$$\frac{d\mathcal{W}}{d\gamma} - \frac{\partial\mathcal{W}}{\partial\gamma} = \left(\frac{\partial\mathcal{E}}{\partial\gamma}\right)^{\top}\mathcal{H}^{-1}\bar{H}\left(\sum_{i=1}^{N}\eta\left(\kappa_i - 1\right)H_i^{-1}G_iH_i^{-1}\right)\mathcal{E}.$$

If $\gamma = \mu_m$, this equation becomes

$$\frac{d\mathcal{W}}{d\mu_m} - \frac{\partial \mathcal{W}}{\partial \mu_m} = \mathbf{1}_m^\top \mathcal{H}^{-1} \bar{H} \left(\sum_{i=1}^{N} \eta \left(\kappa_i - 1 \right) H_i^{-1} G_i H_i^{-1} \right) \mathcal{E}.$$

Given our assumptions, the matrix $\mathcal{H}^{-1} \bar{H} \left(\sum_{i=1}^{N} \eta \left(\kappa_i - 1 \right) H_i^{-1} G_i H_i^{-1} \right)$ is diagonal with positive elements. It follows that the sign of $\frac{d\mathcal{W}}{d\mu_m} - \frac{\partial \mathcal{W}}{\partial \mu_m}$ is the same as the sign of $\mathcal{E}_m$. If instead $\gamma = \Sigma_{mm}$, the equation becomes

$$\frac{d\mathcal{W}}{d\Sigma_{mm}} - \frac{\partial \mathcal{W}}{\partial \Sigma_{mm}} = -\left(\rho - 1 \right) \Delta_m \mathbf{1}_m^\top \mathcal{H}^{-1} \bar{H} \left(\sum_{i=1}^{N} \eta \left(\kappa_i - 1 \right) H_i^{-1} G_i H_i^{-1} \right) \mathcal{E},$$

and so the sign of $\frac{d\mathcal{W}}{d\Sigma_{mm}} - \frac{\partial \mathcal{W}}{\partial \Sigma_{mm}}$ is the same as the sign of $-\Delta_m \mathcal{E}_m$. □

E.16 Proof of Lemma 8

Lemma 8. *In a diagonal economy, an increase in wedges is more detrimental to welfare when risk-exposure decisions can adjust, that is, $d\mathcal{W}/\partial\tau_i \leq \partial\mathcal{W}/\partial\tau_i$.*

Proof. Using (91) and (85), we can write (90) as

$$\frac{d\mathcal{W}}{d\chi} - \frac{\partial \mathcal{W}}{\partial \chi} = \left(\frac{d\Delta}{d\chi} \right)^\top \bar{H} \left(\sum_{i=1}^{N} \eta \left(\kappa_i - 1 \right) H_i^{-1} G_i H_i^{-1} \right) \mathcal{E}.$$

From Proposition 4 and (85), we have

$$\frac{d\Delta}{d\tau_i} = -\mathcal{H}^{-1} \bar{H} \left(\sum_{j=1}^{N} \eta H_i^{-1} G_i H_i^{-1} \frac{d\kappa_j}{d\tau_i} \right) \mathcal{E}.$$

Setting $\chi = \tau_i$ and combining the last two equations, we find

$$\frac{d\mathcal{W}}{d\tau_i} - \frac{\partial \mathcal{W}}{\partial \tau_i} = -\mathcal{E}^\top \left(\sum_{j=1}^{N} \eta H_i^{-1} G_i H_i^{-1} \frac{d\kappa_j}{d\tau_i} \right) \bar{H} \mathcal{H}^{-1} \bar{H} \left(\sum_{i=1}^{N} \eta \left(\kappa_i - 1 \right) H_i^{-1} G_i H_i^{-1} \right) \mathcal{E}.$$

Under our assumptions, the matrix between $\mathcal{E}^\top$and $\mathcal{E}$ is diagonal with positive elements, and the result follows. □

F Robustness, extensions, and additional analysis

In this appendix, we provide additional analysis of the benchmark model presented in the main text. We also show that our model can be extended in different ways.

F.1 Production network and aggregate risk exposure

The structure of the production network is a key determinant of the Domar weight vector. It therefore affects, among other things, how firms' individual exposure decisions contribute to Δ. The following proposition describes how a change in the network affects Δ when $\tau = 0$, in which case the forces at work are more transparent. The proof of the proposition also provides expressions for the general case.

Proposition 8. *Suppose that $\tau = 0$. Then the response of the equilibrium aggregate risk exposure Δ to a change in network connection α_{ij} is given by*

$$\frac{d\Delta}{d\alpha_{ij}} = \mathcal{H}^{-1}\bar{H}\sum_{k=1}^{N}\left(\frac{d\Delta^{\circ}}{d\tilde{\omega}_k} + \frac{d\bar{H}^{-1}}{d\tilde{\omega}_k}\mathcal{E}\right)\frac{d\tilde{\omega}_k}{d\alpha_{ij}}, \tag{93}$$

where the response of Domar weights to a change in α_{ij} is given by $\frac{d\tilde{\omega}_k}{d\alpha_{ij}} = \tilde{\omega}_i\tilde{\mathcal{L}}_{jk}$, the response of the natural exposure to a change in Domar weight is given by $\frac{d\Delta^{\circ}}{d\tilde{\omega}_k} = \delta_k^{\circ}$, and the response of the curvature of $\bar{h}$ to a change in Domar weight is given by $\frac{d\bar{H}^{-1}}{d\tilde{\omega}_k} = H_k^{-1} + \eta\frac{1}{\tilde{\omega}_k}H_k^{-1}G_kH_k^{-1}$.

Proof. Using (85), we can write

$$\Delta - \Delta^{\circ} - \bar{H}^{-1}\mathcal{E} = 0.$$

Using the implicit function theorem, we get

$$\frac{d\Delta}{d\alpha_{ij}} = \mathcal{H}^{-1}\bar{H}\left(\frac{d\Delta^{\circ}}{d\alpha_{ij}} + \frac{d\bar{H}^{-1}}{d\alpha_{ij}}\mathcal{E}\right). \tag{94}$$

Suppose that $\tau = 0$. We have $\frac{d\Delta^{\circ}}{d\alpha_{ij}} = \sum_{k=1}^{N}\frac{d\Delta^{\circ}}{d\tilde{\omega}_k}\frac{d\tilde{\omega}_k}{d\alpha_{ij}}$ and, from (33),

$$\frac{d\bar{H}^{-1}}{d\alpha_{ij}} = \sum_{k=1}^{N}\frac{d\tilde{\omega}_k}{d\alpha_{ij}}\left(H_k^{-1} + \eta\frac{1}{\tilde{\omega}_i}H_k^{-1}G_kH_k^{-1}\right).$$

By definition, $\Delta^{\circ} = \sum_{j=1}^{N}\delta_j^{\circ}\tilde{\omega}_j$ and $\tilde{\omega}_k = \beta^{\top}\tilde{\mathcal{L}}\mathbf{1}_k = \beta^{\top}(I-\alpha)^{-1}\mathbf{1}_k$. Then $\frac{d\Delta^{\circ}}{d\tilde{\omega}_k} = \delta_k^{\circ}$ and $\frac{d\tilde{\omega}_k}{d\alpha_{ij}} = \tilde{\omega}_i\tilde{\mathcal{L}}_{jk}$. Using these results, it is immediate to see that (93) follows from (94).

If $\tau \neq 0$, the only difference is that

$$\frac{d\bar{H}^{-1}}{d\alpha_{ij}} = \sum_{k=1}^{N}\frac{d\tilde{\omega}_k}{d\alpha_{ij}}\left(H_k^{-1} + \eta\frac{\kappa_i}{\tilde{\omega}_i}H_k^{-1}G_kH_k^{-1}\right) - \sum_{k=1}^{N}\frac{d\kappa_k}{d\alpha_{ij}}\eta H_k^{-1}G_kH_k^{-1},$$

where

$$\frac{d\kappa_k}{d\alpha_{ij}} = \frac{(1+\tau_k)}{\omega_k}\frac{d\tilde{\omega}_k}{d\alpha_{ij}} - \frac{\tilde{\omega}_k(1+\tau_k)}{\omega_k^2}\frac{d\omega_k}{d\alpha_{ij}},$$

and $\frac{d\omega_k}{d\alpha_{ij}} = \frac{\omega_i}{1+\tau_i}\mathcal{L}_{jk}$. Therefore,

$$\frac{d\bar{H}^{-1}}{d\alpha_{ij}} = \sum_{k=1}^{N}\left\{\frac{d\tilde{\omega}_k}{d\alpha_{ij}}H_k^{-1} + \eta\frac{\kappa_k}{\omega_k}\frac{d\omega_k}{d\alpha_{ij}}H_k^{-1}G_kH_k^{-1}\right\}.$$

□

Equation (93) shows that when $\tau = 0$, the impact of α_{ij} on the aggregate risk exposure Δ operates exclusively through its effect on the Domar weights $\tilde{\omega}$ (last term in (93)). Since $\frac{d\tilde{\omega}_k}{d\alpha_{ij}} = \tilde{\omega}_i\tilde{\mathcal{L}}_{jk}$, an increase in i's cost share of good j always leads to an increase in $\tilde{\omega}_k$. Indeed, a higher α_{ij} implies that firm j gains in importance as a supplier. It follows that the Domar weight of any firm k that supplies to j ($\tilde{\mathcal{L}}_{jk} > 0$) also benefits from the larger α_{ij}. The magnitude of that increase is larger when i is an important firm in the network (large $\tilde{\omega}_i$) or when k is an important supplier to j (large $\tilde{\mathcal{L}}_{jk}$).

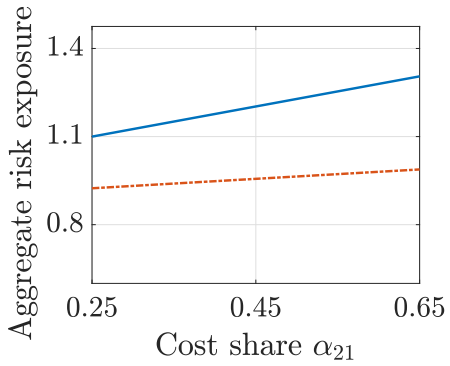
This increase in Domar weights means that firms find it more cost-effective to manage risk. Consequently, the aggregate cost function becomes less curved, in the sense that $\frac{d\bar{H}^{-1}}{d\tilde{\omega}_k} = H_k^{-1} + \eta\frac{1}{\tilde{\omega}_k}H_k^{-1}G_kH_k^{-1}$ is a positive definite matrix. Through that channel, exposure to good risks tends to increase, while exposure to bad risks tends to decrease. Furthermore, the increase in Domar weights implies that the natural aggregate risk-exposure vector Δ° also adjusts. Specifically, an increase in Domar weight $\tilde{\omega}_k$ means that k's natural exposure δ_k° matters more for Δ°. Both the adjustment in the curvature of $\bar{h}$ and in the natural aggregate risk exposure Δ° are propagated through the matrix $\mathcal{T} = \mathcal{H}^{-1}\bar{H}$ to shape the response of the aggregate risk exposure Δ to a change in α_{ij}.

Figure 20 shows the impact of increasing the importance of good 1 in the production of good 2 in the example economy of Figure 3. As α_{21} increases, firm 1 becomes a more important supplier and its Domar weight rises. Because of firm 1's rigid exposure to the good risk factor, Δ_1 increases. This triggers a response from firm 2, which reduces its own exposure to risk factor 1 to avoid creating too much correlated risk. Since κ_2 is parametrized so that both risk factors are substitutes, firm 2's exposure to risk 2 increases, and Δ_2 rises as a result.

F.2 Multiple risk management resources

In the baseline model, we assume that there is only one risk management resource. In this appendix, we show that this setup can be easily extended to handle multiple such resources. Specifically, suppose that there are K risk management resources, and that those can be supplied by the household at a cost in terms of utility. Denote by $R = (R_1, \ldots, R_K)$ the vector of these resources. We assume that the representative household's utility function is $\mathcal{U}(Y)\,\mathcal{V}(R)$, where, as in the baseline model, $U(Y) = \frac{Y^{1-\rho}}{1-\rho}$, and where $\mathcal{V}(R) = \exp\left((\rho-1)\sum_{k=1}^{K}\eta_k R_k\right)$.

Figure 20: Changes in network in the example economy



Notes. The structure of the economy is given in panel (a) of Figure 3. Blue solid line shows $\Delta_1(\alpha_{21})$, and red dashed line shows $\Delta_2(\alpha_{21})$. Initial parametrization is as follows. Household: $\rho = 2$ and $\beta_2 = 0.8$, $\beta_1 = \beta_3 = 0.1$. Network: $\alpha_{21} = \alpha_{23} = 0.25$, all other entries of α are zero. Beliefs: $\mu = (0.75, 0)$, Σ is diagonal with $\text{diag}(\Sigma) = (0.5, 0.5)$. Risk exposures: $\delta^\circ_{11} = \delta^\circ_{32} = 1$, $\delta^\circ_{22} = 1.9$, $\delta^\circ_{12} = \delta^\circ_{21} = \delta^\circ_{31} = 0$, $G_1 = G_3$ are diagonal with very large entries on the main diagonals; $G_{2,11} = G_{2,22} = 1$, $G_{2,12} = G_{2,21} = 0.75$. We set $B_i = 0$ for all i. α_{21} changes from 0.25 to 0.65.

The household's first-order conditions take the same form as before. Specifically, following the same steps as in Appendix A, we can derive

$$\Lambda = \left(\prod_{i=1}^{N} P_i^{\beta_i}\right)^{\rho-1} \Gamma_L^{\rho} \exp\left((\rho-1)\sum_{k=1}^{K}\eta_k R_k\right),$$

and

$$W_{R,k} = \eta_k \Gamma_L^{-1},$$

for all k, and where $W_{R,k}$ is the price of risk management resource k.

We assume that to achieve risk exposure δ_i, firm i must use $R_{ik} = g_{i,k}(\delta_i)$ units of resource k. As in the baseline model, we assume that $g_{i,k}(\delta_i)$ is quadratic with $g_{i,k}(\delta_i) = \frac{1}{2}(\delta_i - \delta_i^\circ)^\top G_{i,k}(\delta_i - \delta_i^\circ)$, where $G_{i,k}$ is a positive definite matrix. Consider now the problem of the firm. For a given risk exposure δ_i, the cost minimization characterization is the same as in the baseline model (see Section 2.5). When choosing its risk exposure, firm i then solves

$$\delta_i^* \in \arg\max_{\delta_i \in \mathcal{A}_i} \mathrm{E}\left[\Lambda\left[Q_i(P_i - K_i(\delta_i, P)) - \sum_{k=1}^{K} g_{i,k}(\delta_i) W_{R,k}\right]\right].$$

Following the same steps as in the baseline model (see, in particular, the proof of Proposition 1), we can show that the equilibrium exposure decisions follow

$$\delta_i - \delta_i^\circ = \left(B_i + \frac{1+\tau_i}{\omega_i}\sum_{k=1}^{K}\eta_k G_{i,k}\right)^{-1}\mathcal{E}. \tag{95}$$

The key difference between this extension and the model with a single risk management resource

is that $\sum_{k=1}^{K}\eta_k G_{i,k}$ replaces ηG_i. Since both objects are constant, we can think of the Hessian matrix G_i from the main text as capturing the aggregated substitution patterns generated by the resource-specific Hessians $G_{i,k}$.

We can again characterize the equilibrium using a distorted planner's problem. Consider the maximization problem

$$\max_{\delta}\tilde{\omega}^{\top}\delta\mu-\tilde{\omega}^{\top}b\left(\delta\right)-\tilde{\omega}^{\top}\log\left(1+\tau\right)-\log\Gamma_L-\frac{1}{2}\left(\rho-1\right)\tilde{\omega}^{\top}\delta\Sigma\delta^{\top}\tilde{\omega}-\sum_{i=1}^{N}\sum_{k=1}^{K}\eta_k\kappa_i g_{i,k}\left(\delta_i\right), \quad (96)$$

where $\kappa_i=\frac{\tilde{\omega}_i(1+\tau_i)}{\omega_i}$. Taking first-order conditions with respect to δ_i, we get

$$0=\tilde{\omega}_i\mu-\tilde{\omega}_i B_i\left(\delta_i-\delta_i^{\circ}\right)-\left(\rho-1\right)\tilde{\omega}_i\tilde{\omega}^{\top}\delta\Sigma-\sum_{k=1}^{K}\eta_k\frac{\tilde{\omega}_i\left(1+\tau_i\right)}{\omega_i}G_{i,k}\left(\delta_i-\delta_i^{\circ}\right)\Leftrightarrow$$

$$\left(\delta_i-\delta_i^{\circ}\right)=\left(B_i+\frac{\left(1+\tau_i\right)}{\omega_i}\sum_{k=1}^{K}\eta_k G_{i,k}\right)^{-1}\left(\mu-\left(\rho-1\right)\tilde{\omega}^{\top}\delta\Sigma\right),$$

which is equivalent to (95). It follows that any equilibrium must coincide with a solution to the optimization problem (96).

We can write (96) as

$$\max_{\Delta}\Delta^{\top}\mu-\tilde{\omega}^{\top}\log\left(1+\tau\right)-\log\Gamma_L-\frac{1}{2}\left(\rho-1\right)\Delta^{\top}\Sigma\Delta-\min_{\delta\text{ s.t. }\Delta=\delta^{\top}\tilde{\omega}}\left[\tilde{\omega}^{\top}b\left(\delta\right)+\sum_{i=1}^{N}\sum_{k=1}^{K}\eta_k\kappa_i g_{i,k}\left(\delta_i\right)\right],$$

or, equivalently, as

$$\mathcal{W}_{dist}:=\max_{\Delta}\underbrace{\Delta^{\top}\mu-\bar{b}\left(\Delta\right)-\tilde{\omega}^{\top}\log\left(1+\tau\right)-\log\Gamma_L}_{\mathrm{E}[y]}-\frac{1}{2}\left(\rho-1\right)\underbrace{\Delta^{\top}\Sigma\Delta}_{\mathrm{V}[y]}-\bar{g}\left(\Delta\right), \quad (97)$$

where

$$\bar{b}\left(\Delta\right):=\tilde{\omega}^{\top}b\left(\delta\left(\Delta\right)\right)\text{ and }\bar{g}\left(\Delta\right):=\sum_{i=1}^{N}\sum_{k=1}^{K}\eta_k\kappa_i g_{i,k}\left(\delta_i\left(\Delta\right)\right), \quad (98)$$

where $\delta\left(\Delta\right)$ solves

$$\bar{h}\left(\Delta\right)=\min_{\delta\text{ s.t. }\Delta=\delta^{\top}\tilde{\omega}}\left[\tilde{\omega}^{\top}b\left(\delta\right)+\sum_{i=1}^{N}\sum_{k=1}^{K}\eta_k\kappa_i g_{i,k}\left(\delta_i\right)\right]. \quad (99)$$

As in the baseline model, we can solve for $\bar{h}\left(\Delta\right)=\bar{b}\left(\Delta\right)+\bar{g}\left(\Delta\right)$.

Lemma 10. *The aggregate cost function $\bar{h}$ is given by*

$$\bar{h}\left(\Delta\right)=\frac{1}{2}\left(\Delta-\Delta^{\circ}\right)^{\top}\nabla^2\bar{h}\left(\Delta-\Delta^{\circ}\right), \quad (100)$$

where $\Delta^\circ = (\delta^\circ)^\top \tilde{\omega}$*, and where the Hessian matrix of* $\bar{h}$ *is given by*

$$\bar{H} = \left[\sum_{i=1}^{N} \tilde{\omega}_i H_i^{-1}\right]^{-1}, \text{ where } H_i = B_i + \frac{\kappa_i}{\tilde{\omega}_i} \sum_{k=1}^{K} \eta_k G_{i,k}. \tag{101}$$

Proof. The Lagrangian of problem (99) is

$$\mathscr{L} = \sum_{i=1}^{N} \tilde{\omega}_i b_i(\delta_i) + \sum_{i=1}^{N} \sum_{k=1}^{K} \eta_k \kappa_i g_{i,k}(\delta_i) - \sum_{m=1}^{M} \nu_m \left(\Delta_m - \mathbf{1}_m^\top \delta^\top \tilde{\omega}\right),$$

where ν_m is the Lagrange multiplier on the mth row of the constraint $\Delta = \delta^\top \tilde{\omega}$. The first-order condition with respect to δ_i is

$$\tilde{\omega}_i \nabla b_i(\delta_i) + \kappa_i \sum_{k=1}^{K} \eta_k \nabla g_{i,k}(\delta_i) = -\nu \tilde{\omega}_i,$$

which implies that

$$\nabla b_i(\delta_i) + \frac{\kappa_i}{\tilde{\omega}_i} \sum_{k=1}^{K} \eta_k \nabla g_{i,k}(\delta_i) = \nabla b_j(\delta_j) + \frac{\kappa_j}{\tilde{\omega}_j} \sum_{k=1}^{K} \eta_k \nabla g_{j,k}(\delta_j) = \nabla \bar{h},$$

for all i, j, where the last equality comes from the envelope theorem. Notice that

$$\left(B_i + \frac{\kappa_i}{\tilde{\omega}_i} \sum_{k=1}^{K} \eta_k G_{i,k}\right)(\delta_i - \delta_i^\circ) = \left(B_j + \frac{\kappa_j}{\tilde{\omega}_j} \sum_{k=1}^{K} \eta_k G_{j,k}\right)(\delta_j - \delta_j^\circ) \Leftrightarrow$$

$$\delta_i = \delta_i^\circ + \left(B_i + \frac{\kappa_i}{\tilde{\omega}_i} \sum_{k=1}^{K} \eta_k G_{i,k}\right)^{-1} \left(B_j + \frac{\kappa_j}{\tilde{\omega}_j} \sum_{k=1}^{K} \eta_k G_{j,k}\right)(\delta_j - \delta_j^\circ).$$

Then, the constraint $\Delta = \delta^\top \tilde{\omega}$ can be rewritten as

$$\begin{aligned}
\Delta &= \sum_{i=1}^{N} \tilde{\omega}_i \delta_i = \sum_{i=1}^{N} \tilde{\omega}_i \left(\delta_i^\circ + \left(B_i + \frac{\kappa_i}{\tilde{\omega}_i} \sum_{k=1}^{K} \eta_k G_{i,k}\right)^{-1} \left(B_j + \frac{\kappa_j}{\tilde{\omega}_j} \sum_{k=1}^{K} \eta_k G_{j,k}\right)(\delta_j - \delta_j^\circ)\right) \\
&= \Delta^\circ + \sum_{i=1}^{N} \tilde{\omega}_i \left(B_i + \frac{\kappa_i}{\tilde{\omega}_i} \sum_{k=1}^{K} \eta_k G_{i,k}\right)^{-1} \left(B_j + \frac{\kappa_j}{\tilde{\omega}_j} \sum_{k=1}^{K} \eta_k G_{j,k}\right)(\delta_j - \delta_j^\circ) \Leftrightarrow \\
\delta_j - \delta_j^\circ &= \left(B_j + \frac{\kappa_j}{\tilde{\omega}_j} \sum_{k=1}^{K} \eta_k G_{j,k}\right)^{-1} \bar{H}(\Delta - \Delta^\circ),
\end{aligned}$$

where $\bar{H}$ is given by (101). Combining this last expression with (99) yields the result. □

Since (97) and (98) fully characterize the equilibrium, it follows that the only difference between

the baseline model and the multiple resources model comes from the cost functions $\bar{h}$. In turn, the cost function in the baseline model is identical to (100) if we impose that

$$\eta G_i = \sum_{k=1}^{K} \eta_k G_{i,k}$$

for all i, k. Under that restriction, both models behave identically. It follows that we can simply think of the unique resource in the baseline model as an aggregate of various risk management resources, each with its own supply elasticity η_k.

F.3 General cost function

In this appendix, we relax our assumption on the functional form of the cost functions $b_i(\delta_i)$ and $g_i(\delta_i)$. Specifically, we assume that $b_i(\delta_i)$ and $g_i(\delta_i)$ are strictly convex—but not necessarily quadratic—functions. The Hessians of $b_i(\delta_i)$ and $g_i(\delta_i)$ are denoted by B_i and G_i, respectively, and both of them are positive definite matrices. Unlike in the baseline model, B_i and G_i are not necessarily constant, in the sense that they can depend on δ_i.

All the results of Sections 2, 3, (4), and 5 remain unchanged. However, (32) and (32) do not hold when b_i and g_i are non-quadratic. Nevertheless, the Hessian of the aggregate cost function $\bar{h}$ is still given by (33), as the following result shows.

Lemma 11. *The Hessian matrix of $\bar{h}$ is given by (33).*

Proof. Equations (76), (77), and (78) do not use the assumption that κ_i is quadratic and, hence, hold in this more general case. Differentiate (78) with respect Δ_j to get

$$\frac{d^2\bar{h}}{d\Delta_i d\Delta_j} = \mathbf{1}_i^\top \left(B_i + \eta \frac{\kappa_i}{\tilde{\omega}_i} G_i \right) \frac{d\delta_i}{d\Delta_j} \tag{102}$$

and

$$\left(B_k + \eta \frac{\kappa_k}{\tilde{\omega}_k} G_k \right) \frac{d\delta_k}{d\Delta_j} = \left(B_i + \eta \frac{\kappa_i}{\tilde{\omega}_i} G_i \right) \frac{d\delta_i}{d\Delta_j}$$

for all i, k.

Next, we can differentiate the constraint $\Delta = \delta^\top \tilde{\omega}$ with respect to Δ_j to find $\sum_{k=1}^{N} \tilde{\omega}_k \frac{d\delta_k}{d\Delta_j} = \mathbf{1}_j$. Combining this with the last equation, we find

$$\left(B_i + \eta \frac{\kappa_i}{\tilde{\omega}_i} G_i \right) \frac{d\delta_i}{d\Delta_j} = \left(\sum_{k=1}^{N} \tilde{\omega}_k \left(B_k + \eta \frac{\kappa_k}{\tilde{\omega}_k} G_k \right)^{-1} \right)^{-1} \mathbf{1}_j. \tag{103}$$

Plugging this into (102) yields (33). □

F.3.1 Risk exposures

Turning to the results of Section 6, we are no longer able to solve for Δ and δ in closed form. Nevertheless, we can still derive similar comparative statics results. We start by characterizing how δ depends on aggregate risk exposure Δ. We then characterize how Δ depends on various primitives of the model.

Proposition 9. *The response of firm i's risk exposure δ_i to a change in the aggregate risk exposure Δ_j is given by*

$$\frac{d\delta_i}{d\Delta_j} = H_i^{-1}\bar{H}\mathbf{1}_j, \tag{104}$$

where $H_i = B_i + \eta\frac{\kappa_i}{\tilde{\omega}_i}G_i$.

Proof. This result follows immediately from (103) and (33). □

In the model with general cost functions, (35) still implicitly defines the equilibrium Δ. Proposition 3 immediately follows by applying the implicit function theorem to (35). In what follows, we prove analogues of Propositions 4 and 8.

Proposition 10. *The response of the equilibrium aggregate risk exposure Δ to a change in wedge τ_i is given by*

$$\frac{d\Delta}{d\tau_i} = -\mathcal{H}^{-1}\bar{H}\sum_{j=1}^{M}\eta\frac{d\kappa_j}{d\tau_i}H_j^{-1}\left(\nabla g_j\right), \tag{105}$$

where $\mathcal{H}$ is given by (39), $\bar{H}$ is given by (33), $d\kappa_j/d\tau_i$ is given by (86), and ∇g_j is evaluated at the equilibrium δ_j.

Proof. By the implicit function theorem applied to (35),

$$\frac{d\Delta}{d\tau_i} = -\mathcal{H}^{-1}\frac{\partial\nabla\bar{h}}{\partial\tau_i} = -\mathcal{H}^{-1}\sum_{j=1}^{M}\frac{\partial\nabla\bar{h}}{\partial\kappa_j}\frac{d\kappa_j}{d\tau_i}, \tag{106}$$

where $\mathcal{H}$ is given by (39) and $\frac{d\kappa_j}{d\tau_i}$ is given by (86).

Next, from (78), we have $\frac{\partial\bar{h}}{\partial\kappa_j} = \eta g_j\left(\delta_j\right)$, and so

$$\frac{d}{d\Delta_k}\frac{\partial\bar{h}}{\partial\kappa_j} = \eta\left(\nabla g_j\right)^\top\frac{d\delta_j}{d\Delta_k}.$$

Using (104), we get

$$\frac{\partial\nabla\bar{h}}{\partial\kappa_j} = \eta\bar{H}H_j^{-1}\nabla g_j. \tag{107}$$

Plugging this into (106), we get

$$\frac{d\Delta}{d\tau_i} = -\mathcal{H}^{-1}\bar{H}\sum_{j=1}^{M}\eta H_j^{-1}\nabla g_j\frac{d\kappa_j}{d\tau_i},$$

which is the result. □

Next, we derive an analogue of Proposition 8.

Proposition 11. *The response of the equilibrium aggregate risk exposure Δ to a change in network connection α_{ij} is given by*

$$\frac{d\Delta}{d\alpha_{ij}} = -\mathcal{H}^{-1}\bar{H}\sum_{k=1}^{N}\left[\eta H_k^{-1}\left(\nabla g_k\right)\frac{d\kappa_k}{d\alpha_{ij}} - \left(\delta_k - \eta\frac{\kappa_k}{\tilde{\omega}_k}H_k^{-1}\nabla g_k\right)\frac{d\tilde{\omega}_k}{d\alpha_{ij}}\right], \tag{108}$$

where $\frac{d\kappa_k}{d\alpha_{ij}} = \frac{(1+\tau_k)}{\omega_k}\frac{d\tilde{\omega}_k}{d\alpha_{ij}} - \frac{\tilde{\omega}_k(1+\tau_k)}{\omega_k^2}\frac{d\omega_k}{d\alpha_{ij}}$, $\frac{d\tilde{\omega}_k}{d\alpha_{ij}} = \tilde{\omega}_i\tilde{\mathcal{L}}_{jk}$, $\frac{d\omega_k}{d\alpha_{ij}} = \frac{\omega_i}{1+\tau_i}\mathcal{L}_{jk}$, $\mathcal{H}$ *is given by* (39), $\bar{H}$ *is given by* (33), *and* ∇g_k *is evaluated at the equilibrium* δ_k.

Proof. By the implicit function theorem applied to (35),

$$\frac{d\Delta}{d\alpha_{ij}} = -\mathcal{H}^{-1}\sum_{k=1}^{N}\left(\frac{\partial\nabla\bar{h}}{\partial\kappa_k}\frac{d\kappa_k}{d\alpha_{ij}} + \frac{\partial\nabla\bar{h}}{\partial\tilde{\omega}_k}\frac{d\tilde{\omega}_k}{d\alpha_{ij}}\right), \tag{109}$$

where $\mathcal{H}$ is given by (39), $\frac{d\tilde{\omega}_k}{d\alpha_{ij}} = \tilde{\omega}_i\tilde{\mathcal{L}}_{jk}$,

$$\frac{d\kappa_k}{d\alpha_{ij}} = \frac{(1+\tau_k)}{\omega_k}\frac{d\tilde{\omega}_k}{d\alpha_{ij}} - \frac{\tilde{\omega}_k\left(1+\tau_k\right)}{\omega_k^2}\frac{d\omega_k}{d\alpha_{ij}},$$

and $\frac{d\omega_k}{d\alpha_{ij}} = \frac{\omega_i}{1+\tau_i}\mathcal{L}_{jk}$.

Next, from (76), we obtain $\frac{\partial\bar{h}}{\partial\tilde{\omega}_k} = b_k + \nu^\top\delta_k$, where $\nu = -\nabla\bar{h}$ from (78). Therefore,

$$\frac{d}{d\Delta_m}\frac{\partial\bar{h}}{\partial\tilde{\omega}_k} = \left(\nabla b_k\right)^\top\frac{d\delta_k}{d\Delta_m} - \mathbf{1}_m^\top\bar{H}\delta_k - \left(\nabla\bar{h}\right)^\top\frac{d\delta_k}{d\Delta_m}.$$

Using (104), this equation can be rewritten as

$$\frac{\partial\nabla\bar{h}}{\partial\tilde{\omega}_k} = -\bar{H}\delta_k - \bar{H}H_k^{-1}\left(\nabla\bar{h} - \nabla b_k\right). \tag{110}$$

Plugging (107), (78), and (110) into (109), we get (108). □

F.3.2 GDP and welfare

As can be seen from the analysis of the baseline model, we do not use the fact that b_i and g_i are quadratic in Section 7, with the exception of deriving the expression for $\nabla\bar{h} - \nabla\bar{h}_V$, which is

given by (92). Below, we derive an analogue of this equation under general cost functions. Using the definition of $\bar{h}_V$, we can derive

$$\frac{d\bar{h}_V}{d\Delta_m} = \sum_{i=1}^{N} \left(\tilde{\omega}_i \left(\nabla b_i \right)^\top + \eta \left(\nabla g_i \right)^\top \right) \frac{d\delta_i}{d\Delta_m}.$$

From (78), we have $\nabla b_i + \eta \frac{\kappa_i}{\tilde{\omega}_i} \nabla g_i = \nabla \bar{h}$. Furthermore, $\frac{d\delta_i}{d\Delta_m}$ is given by (104). Then we have

$$\nabla \bar{h} - \nabla \bar{h}_V = \left(\nabla^2 \bar{h} \right) \sum_{i=1}^{N} \eta \left(\kappa_i - 1 \right) H_i^{-1} \nabla g_i.$$

In our baseline model with quadratic cost functions, we have $\nabla g_i = G_i \left(\delta_i - \delta_i^\circ \right) = G_i H_i^{-1} \mathcal{E}$ (from (26)), such that in a diagonal economy, $\text{sign} \left(\nabla g_i \right) = \text{sign} \left(\mathcal{E} \right) \ \forall i$. In the model with general cost functions, it is not always true that the signs of ∇g_i and $\mathcal{E}$ coincide. Thus, Lemma 9 and Corollaries 7 and 8 will generally look different. One sufficient condition under which the signs of ∇g_i and $\mathcal{E}$ coincide is the following. Suppose that $b_i \left(\delta_i \right) = v_i g_i \left(\delta_i \right)$ where v_i is a non-negative constant. Then, from (78) and (35),

$$\left(v_i + \eta \frac{\kappa_i}{\tilde{\omega}_i} \right) \nabla g_i = \nabla \bar{h} = \mathcal{E},$$

such that the signs of ∇g_i and $\mathcal{E}$ coincide. Then the results of Section 7 hold even if b_i and g_i are non-quadratic.

F.4 Microfoundations for the risk management cost functions

In the main text, we treat the risk management cost functions b_i and g_i as exogenous. In reality, many firm decisions might affect their risk profiles (hiring, investment, research and development, location choices, etc.) and modeling all those individually would make the model intractable. Assuming general forms for b_i and g_i allows us to sidestep these issues, and lets us explore the impact of endogenous uncertainty on firm decisions and the macroeconomy in a simple way.

But for some applications, it might be helpful to consider specific risk-mitigation mechanisms, in which case one can microfound b_i and g_i based on the mechanisms of interest. Doing so could yield a greater understanding of the origin of the cost functions and of the features of the environment that affect them. In this appendix, we provide four possible microfoundations for b_i and g_i. These microfoundations are stylized and can certainly be generalized in many ways, at the cost of tractability. They are also only examples. It is possible to come up with many such microfoundations, especially if b_i and g_i are interpreted as approximations to more complicated functions.

F.4.1 Research and development

Consider a firm that seeks to gain expertise in a new technology that will become operational soon. The true quality ε of the technology is unknown, but it is expected to be distributed according to $\varepsilon \sim \mathcal{N}(\mu, \Sigma)$, where $\mu \gg 0$. Firms that gain expertise with the technology will be able to use it to increase their productivity by an amount proportional to ε.

To gain expertise, firms must complete tasks (studies, experiments, etc.) that familiarize its workforce with the new technology. Each task is associated with a difficulty level $t \geq 0$. Completing task t requires $\iota_g t$ units of resources and lowers the productivity of the firm by $\iota_b t$. Intuitively, resources might have to be purchased for experiments, and workers might need to be reallocated away from production to run those experiments, which lowers productivity. There is one unit mass of tasks of each type t, and the firm is free to complete as many tasks as it wants, with the understanding that tasks become increasingly difficult. Suppose that the firm chooses to complete a mass δ of tasks. It obviously starts with the easiest ones. Its resource cost is therefore

$$g_i(\delta) = \int_0^{\delta} \iota_g t \, dt = \iota_g \frac{\delta^2}{2},$$

and similarly for its productivity cost. The total productivity of the firm is therefore

$$a_i(\varepsilon, \delta_i) = \delta\varepsilon - \iota_b \frac{\delta^2}{2}.$$

As we can see, this simple task model of research and development fits into the general framework of the main text, and the cost functions also fit within the quadratic structure given by (4) and (5).

F.4.2 Natural disasters and transportation to market

Our second microfoundation involves a location decision by a firm. In the spirit of Hotelling (1929), suppose that there is a set of locations l on the unit interval $l \in [0, 1]$. A firm needs to locate its operation somewhere on $[0, 1]$. At $l = 1$ there is a city where customers live, but that city is located on the ocean shore and, as a result, is subject to potential natural catastrophes (flood, hurricane, climate changes, etc.). By being close to $l = 1$, the firm is closer to customers but also more subject to catastrophes. We model this by assuming that if the firm is located at $l = \delta$, the stochastic part of its TFP is given by $\delta\varepsilon$, where the catastrophe shock ε is distributed according to $\varepsilon \sim \mathcal{N}(\mu, \Sigma)$ with $\mu \ll 0$.[69] If the firm is away from $l = 1$, it must transport its goods to the city. Transportation involves costs (drivers, gasoline, vehicles, repairs, etc.) and we model those as a combination of a variable and a fixed costs. The variable cost scales with the amount of goods sold, and can therefore be modeled as a productivity shifter $\kappa_b(1 - \delta)$. The firm's productivity is

[69] Such a catastrophe risk can also be modeled as a disaster (rather than normal) shock, as in Section 9.5.

therefore

$$a_i(\varepsilon, \delta_i) = \delta\varepsilon - \kappa_b(1-\delta).$$

The fixed cost $\kappa_g(1-\delta)$ does not vary with the size of the firm and is therefore akin to the risk management resources of the main text. It is natural to think that the functions κ_b and κ_g are convex: longer trips might require an overnight stay, fresh produce are more likely to go bad, mechanical problems away from the city incur larger costs, etc. In this case, approximating κ_b and κ_g using quadratic functions like b_i and g_i in the main text would be reasonable. As we can see, this simple catastrophe and transportation model can provide a microfoundation for the general setup of the main text.

F.4.3 Workplace rules

Many organizations put rules into place to avoid bad outcomes. For instance, metal detectors are often installed at office entry points, access to computer networks requires two-factor authentication, installing new software requires the approval of the IT department, etc. While these policies can help mitigate catastrophic outcomes (security incident, hackers taking over the computer networks, etc.) they might also impede normal work activities and lower productivity.

Suppose for instance that a firm has a unit mass of workers, and that each worker is endowed with an unobservable skill $s \sim U([0,1])$. To produce, workers must complete tasks. There is a unit mass of tasks, each with a difficulty level $d \sim U([0,1])$ that is observable. Tasks are assigned to workers at random. If a task d is assigned to a worker with skill $s \geq d$, the task is completed successfully and generates ι units of productivity for the firm. If a task is assigned to a worker with skill $s < d$, the task does not generate any productivity.

Any attempted task increases the importance of a potential problem $\varepsilon \sim \mathcal{N}(\mu, \Sigma)$, with $\mu \ll 0$, for the firm. To lower the amount of risk it is exposed to, the firm can impose a *rule* δ. With a rule δ, the workers do not attempt any tasks with a difficulty level $d \geq \delta$. Without rules ($\delta = 1$), the full unit mass of tasks is attempted, so that the random problem lowers productivity by $1 \times \varepsilon$, and half of those tasks are successfully completed, which generates productivity $1/2\iota$. In that case,

$$a_i(\varepsilon, \delta_i) = 1 \times \varepsilon + \frac{1}{2}\iota.$$

Under a rule δ, the firm faces some productivity losses from the tasks that are not attempted but that high-skill workers would have otherwise successfully completed. There are $(1-\delta)^2/2$ such tasks.[70] On the positive side, the intensity of the random problem ε declines. Putting the pieces

[70] To get to that number, it helps to think of the square $[0,1] \times [0,1]$ of skills and difficulty. We need to compute the area of the upper triangle created by the horizontal line $d = \delta$ and the linear function $d = s$.

together, the productivity of the firm under rule δ is

$$a_i(\varepsilon, \delta_i) = \delta\varepsilon + \left(\frac{1}{2} - \frac{(1-\delta)^2}{2}\right)\iota.$$

As we can see, rules imposed by organizations can serve as a microfoundation for the TFP process used in the main text (2) and the productivity risk management cost b_i[71]

$$b_i(\delta_i) = \frac{(1-\delta)^2}{2}\iota,$$

which fits into our general quadratic formulation (4) with $\delta^\circ = 1$.

F.4.4 Skilled managers

To microfound the risk management cost functions, we can also assume that adjusting risk within a firm requires the use of managers. Managers differ in their skill level, and larger risk adjustments can only be done by managers with more advanced skills. In turn, high-skill managers are paid more, which leads to a convex relationship between the size of a risk adjustment and its cost.

For simplicity, suppose that there is a single risk factor ε and that firm i's has a natural exposure δ_i° to that factor. To set $\delta_i \neq \delta_i^\circ$ requires a manager with a skill level $s = |\delta_i - \delta_i^\circ|$. Such a manager is paid an equilibrium wage of $W_R(s)$. The firm's problem (11) is therefore

$$\max_{\delta_i} \mathrm{E}\left[\Lambda\left(Q_i\left(P_i - K_i(\delta_i)\right) - W_R(s)\right)\right],$$

where $s = |\delta_i - \delta_i^\circ|$.

It is straightforward to adapt the problem of the household to determine $W_R(s)$ in equilibrium. For instance, we can assume that the amount of total risk management resources R supplied by the household is given by

$$R = \int_0^\infty R(s)\,\varphi(s)\,ds,$$

where $\varphi(s)$ is the cost of supplying a manager of skill s, and $R(s)$ is the mass of such managers. It is natural to think that φ is increasing and convex. Indeed, while almost anyone can learn basic management techniques, teaching new skills to high-level executives might incur larger and larger costs. In that case, we can approximate $\varphi(s) \approx \frac{h}{2}s^2$, for some $h > 0$. For simplicity, we assume that this relationship holds with equality from now on. Household optimization then implies that $W_R(s) = \frac{h}{2}s^2\bar{W}_R$, where $\bar{W}_R$ is the average manager wage in the economy.

Taking the equilibrium manager wage function into account, we can rewrite the profit maxi-

[71] The constant 0.5ι in a_i can be thought of as a fixed exposure to a constant risk factor $\varepsilon_0 \sim \mathcal{N}(1, 0)$. It therefore fits into the general model of the main text.

mization problem of firm i as

$$\max_{\delta_i} \mathrm{E}\left[\Lambda\left(Q_i\left(P_i - K_i\left(\delta_i\right)\right) - \frac{h}{2}\left(\delta_i - \delta_i^{\circ}\right)^2 \bar{W}_R\right)\right].$$

Comparing this equation with (11), we see that we have microfounded the risk resource cost function as $g_i = \frac{h}{2}\left(\delta_i - \delta_i^{\circ}\right)^2$. Under this microfoundation changes in the training cost of new managers would affect the risk profile of the firms and the aggregate economy.

F.5 A general specification of the disutility from risk management

In the main model, we assume that the household's disutility is given by (6), where $\mathcal{V}(R) = \exp\left(-\eta\left(1-\rho\right)R\right)$. In this appendix, we consider a general disutility from supplying risk management resources. Specifically, the household's utility function is given by

$$\mathcal{U}\left(\frac{Y^{1-\rho}}{1-\rho}, R\right),$$

where $Y = \prod_{i=1}^{N}\left(\beta_i^{-1}C_i\right)^{\beta_i}$. Denote by $\mathcal{U}_1$ and $\mathcal{U}_2$ the derivatives of $\mathcal{U}\left(\cdot,\cdot\right)$ with respect to the first and second inputs, respectively. Following the same steps as in the baseline model (see Appendix A), we can derive

$$\mathrm{E}\left[\mathcal{U}_2\right] = -W_R\,\mathrm{E}\left[\Lambda\right], \tag{111}$$

where

$$\Lambda = \Gamma_L^{\rho}\mathcal{U}_1\left(\prod_{i=1}^{N} P_i^{\beta_i}\right)^{\rho-1}. \tag{112}$$

As in the baseline model, the labor share of income Γ_L is given by $\Gamma_L := \frac{W_L L}{PY} = 1 - \tau^{\top}\left(\mathrm{diag}\left(1+\tau\right)\right)^{-1}\omega$, and real GDP is given by $Y = \Gamma_L^{-1}\left(\prod_{i=1}^{N} P_i^{-\beta_i}\right)$ (see Lemma 1).

For a given risk-exposure decision, firms' cost minimization does not depend on the household's utility function. Therefore, the unit cost for firm i is given by (10), and the vector of prices is given by (16). The choice of risk exposure for firm i is described by (11), which can be rewritten as

$$\delta_i^* \in \arg\min_{\delta_i} \mathrm{E}\left[\Lambda\left[Q_i K_i\left(\delta_i, P\right) + g_i\left(\delta_i\right)W_R\right]\right].$$

Taking the first-order condition with respect to δ_{im}, we get

$$\mathrm{E}\left[\Lambda Q_i \frac{dK_i}{d\delta_{im}}\right] + \mathrm{E}\left[\Lambda\right]W_R\frac{dg_i}{d\delta_{im}} = 0.$$

Using the expression for the unit cost (10) together with equations (111) and (112), this expression

can be simplified as

$$\frac{dg_i}{d\delta_{im}} = -K_i Q_i \frac{\mathrm{E}\left[\Gamma_L^{\rho} \mathcal{U}_1 \left(\prod_{i=1}^N P_i^{\beta_i}\right)^{\rho-1} \left(\varepsilon_m - \frac{db_i}{d\delta_{im}}\right)\right]}{\mathrm{E}\left[\mathcal{U}_2\right]}.$$

Recall that in equilibrium, $K_i Q_i = \frac{\omega_i \Gamma_L^{-1}}{1+\tau_i}$. Using the expression for prices (16), we can rewrite the equation above as

$$\frac{dg_i}{d\delta_{im}} = \\ -\frac{\omega_i \Gamma_L^{\rho-1}}{1+\tau_i} \exp\left(-(\rho-1)\tilde{\omega}^{\top}\left(b(\delta) + \log(1+\tau)\right)\right) \frac{\mathrm{E}\left[\mathcal{U}_1 \exp\left(-(\rho-1)\tilde{\omega}^{\top}\delta\varepsilon\right)\left(\varepsilon_m - \frac{db_i}{d\delta_{im}}\right)\right]}{\mathrm{E}\left[\mathcal{U}_2\right]}. \tag{113}$$

The system of equations (113) for all i, m implicitly defines the equilibrium risk-exposure matrix δ. In our baseline model, we show that the equilibrium allocation is a solution to a distorted planner's problem (34). This distorted planner chooses the aggregate risk exposure Δ, and then chooses δ to minimize the aggregate cost function (31) subject to the $\Delta = \delta^{\top}\tilde{\omega}$ constraint. For a general utility function, we are not able to write the equilibrium allocation as a solution to a distorted planner's problem. Therefore, to characterize the equilibrium, we need to work with equation (113) directly. In what follows, we use the implicit function theorem to characterize how equilibrium δ changes in response to a change in the environment. It is then straightforward to characterize how the aggregate risk exposure $\Delta = \delta^{\top}\tilde{\omega}$ responds.

In general, one cannot compute the expectations on the right-hand side of (113) analytically. Therefore, it is generally infeasible to compute the derivatives of δ with respect to changes in the environment in closed form. We consider two special cases for which this is possible.

F.5.1 Multiplicative disutility

Suppose that the household's utility function is $\mathcal{U}\left(\frac{Y^{1-\rho}}{1-\rho}, R\right) = \frac{Y^{1-\rho}}{1-\rho} V(R)$, where $V(R)$ takes positive values and is an increasing convex function. One special case, considered in the baseline model, is $V(R) = \exp((\rho-1)\eta R)$. Under this specification, (113) becomes

$$\frac{dg_i}{d\delta_{im}} = (\rho-1)\frac{\omega_i}{1+\tau_i}\left(\mu_m - (\rho-1)\mathbf{1}_m^{\top}\Sigma\delta^{\top}\tilde{\omega} - \frac{db_i}{d\delta_{im}}\right)\frac{V(R)}{V'(R)}.$$

Market clearing implies that $R = \sum_{i=1}^N g_i(\delta_i)$. While we cannot solve for δ in closed form for a general $V(\cdot)$ function, we can characterize how δ changes when one of the parameter changes using

the implicit function theorem. Specifically, write the first-order condition for δ_{im} as

$$F_{im} = \frac{dg_i}{d\delta_{im}} - (\rho - 1)\frac{\omega_i}{1+\tau_i}\left(\mu_m - (\rho-1)\mathbf{1}_m^\top \Sigma \delta^\top \tilde{\omega} - \frac{db_i}{d\delta_{im}}\right)\frac{V\left(\sum_{i=1}^N g_i(\delta_i)\right)}{V'\left(\sum_{i=1}^N g_i(\delta_i)\right)} = 0.$$

Then, by the implicit function theorem,

$$\begin{pmatrix} \frac{d\delta_{11}}{d\chi} \\ \vdots \\ \frac{d\delta_{1M}}{d\chi} \\ \vdots \\ \frac{d\delta_{N1}}{d\chi} \\ \vdots \\ \frac{d\delta_{NM}}{d\chi} \end{pmatrix} = - \begin{pmatrix} \frac{\partial F_{11}}{\partial \delta_{11}} & \cdots & \frac{\partial F_{11}}{\partial \delta_{1M}} & \cdots & \frac{\partial F_{11}}{\partial \delta_{N1}} & \cdots & \frac{\partial F_{11}}{\partial \delta_{NM}} \\ \vdots & \ddots & \vdots & \ddots & \vdots & \ddots & \vdots \\ \frac{\partial F_{1M}}{\partial \delta_{11}} & \cdots & \frac{\partial F_{1M}}{\partial \delta_{1M}} & \cdots & \frac{\partial F_{1M}}{\partial \delta_{N1}} & \cdots & \frac{\partial F_{1M}}{\partial \delta_{NM}} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \frac{\partial F_{N1}}{\partial \delta_{11}} & \cdots & \frac{\partial F_{N1}}{\partial \delta_{1M}} & \cdots & \frac{\partial F_{N1}}{\partial \delta_{N1}} & \cdots & \frac{\partial F_{N1}}{\partial \delta_{NM}} \\ \vdots & \ddots & \vdots & \ddots & \vdots & \ddots & \vdots \\ \frac{\partial F_{NM}}{\partial \delta_{11}} & \cdots & \frac{\partial F_{NM}}{\partial \delta_{1M}} & \cdots & \frac{\partial F_{NM}}{\partial \delta_{N1}} & \cdots & \frac{\partial F_{NM}}{\partial \delta_{NM}} \end{pmatrix}^{-1} \begin{pmatrix} \frac{\partial F_{11}}{\partial \chi} \\ \vdots \\ \frac{\partial F_{1M}}{\partial \chi} \\ \vdots \\ \frac{\partial F_{N1}}{\partial \chi} \\ \vdots \\ \frac{\partial F_{NM}}{\partial \chi} \end{pmatrix}, \tag{114}$$

where

$$\begin{aligned} \frac{\partial F_{im}}{\partial \delta_{jl}} =& \frac{d^2 g_i}{d\delta_{im} d\delta_{il}} \mathbf{1}_{j=i} + \\ & (\rho-1)\frac{\omega_i}{1+\tau_i}\left((\rho-1)\Sigma_{ml}\tilde{\omega}_j + \frac{d^2 b_i}{d\delta_{im}d\delta_{jl}}\mathbf{1}_{j=i}\right)\frac{V(R)}{V'(R)} - \frac{dg_i}{d\delta_{im}}\frac{dg_j}{d\delta_{jl}}\left(\frac{V'(R)}{V(R)} - \frac{V''(R)}{V'(R)}\right), \end{aligned}$$

and $\frac{\partial F_{im}}{\partial \chi}$ is the derivative of F_{im} with respect to parameter χ. For example, if $\chi = \mu_j$, we have

$$\frac{\partial F_{im}}{\partial \mu_j} = -\mathbf{1}_{j=m}(\rho-1)\frac{\omega_i}{1+\tau_i}\frac{V(R)}{V'(R)}.$$

Similarly, one can compute

$$\begin{aligned} \frac{\partial F_{im}}{\partial \Sigma_{jl}} =& (\rho-1)^2 \frac{\omega_i}{1+\tau_i}\frac{1}{2}\left(\mathbf{1}_{j=m}\Delta_l + \mathbf{1}_{l=m}\Delta_j\right)\frac{V(R)}{V'(R)}, \\ \frac{\partial F_{im}}{\partial \tau_j} =& (\rho-1)\frac{1}{1+\tau_i}\frac{\omega_j}{1+\tau_j}\mathcal{L}_{ji}\left(\mu_m - (\rho-1)\mathbf{1}_m^\top\Sigma\Delta - \frac{db_i}{d\delta_{im}}\right)\frac{V(R)}{V'(R)}, \\ \frac{\partial F_{im}}{\partial \alpha_{jl}} =& -(\rho-1)\frac{\mathcal{L}_{li}}{1+\tau_i}\frac{\omega_j}{1+\tau_j}\left(\mu_m - (\rho-1)\mathbf{1}_m^\top\Sigma\Delta - \frac{db_i}{d\delta_{im}}\right)\frac{V(R)}{V'(R)} + \\ & (\rho-1)^2\frac{\omega_i}{1+\tau_i}\mathbf{1}_m^\top\Sigma\left(\sum_{k=1}^N \delta_k \tilde{\mathcal{L}}_{lk}\right)\tilde{\omega}_j\frac{V(R)}{V'(R)}, \end{aligned}$$

where we have used that $\frac{d\tilde{\omega}_k}{d\alpha_{ij}} = \tilde{\omega}_i \tilde{\mathcal{L}}_{jk}$ and $\frac{d\omega_k}{d\alpha_{ij}} = \frac{\omega_i}{1+\tau_i}\mathcal{L}_{jk}$ (see the proof of Proposition 8), and $\frac{d\left(\frac{\omega_j}{1+\tau_j}\right)}{d\tau_i} = -\frac{\mathcal{L}_{ij}}{1+\tau_j}\frac{\omega_i}{1+\tau_i}$.

F.5.2 Additive disutility

Suppose that the household's utility function is $\mathcal{U}\left(\frac{Y^{1-\rho}}{1-\rho}, R\right) = \frac{Y^{1-\rho}}{1-\rho} - V(R)$, where $V(R)$ is an increasing convex function. Under this specification, (113) becomes

$$\begin{aligned}\frac{dg_i}{d\delta_{im}} =& \frac{\omega_i \Gamma_L^{\rho-1}}{1+\tau_i} \exp\left(-(\rho-1)\tilde{\omega}^\top (b(\delta) + \log(1+\tau))\right) \mathrm{E}\left[\exp\left(-(\rho-1)\tilde{\omega}^\top \delta\varepsilon\right)\left(\varepsilon_m - \frac{db_i}{d\delta_{im}}\right)\right] \frac{1}{V'(R)} \\ =& \frac{\omega_i \Gamma_L^{\rho-1}}{1+\tau_i}\left(\mu_m - \frac{db_i}{d\delta_{im}} - (\rho-1)\mathbf{1}_m^\top \Sigma \delta^\top \tilde{\omega}\right) \times \\ & \exp\left(-(\rho-1)\tilde{\omega}^\top (\delta\mu - b(\delta) - \log(1+\tau)) + \frac{1}{2}(\rho-1)^2 \tilde{\omega}^\top \delta \Sigma \delta^\top \tilde{\omega}\right) \frac{1}{V'(R)}.\end{aligned}$$

Market clearing implies that $R = \sum_{i=1}^N g_i(\delta_i)$. As in the previous section, we cannot solve for δ in closed form. However, we can similarly characterize how δ changes when one of the parameter changes using the implicit function theorem. Specifically, write the first-order condition for δ_{im} as

$$\begin{aligned}F_{im} =& \frac{dg_i}{d\delta_{im}} - \frac{\omega_i \Gamma_L^{\rho-1}}{1+\tau_i}\left(\mu_m - \frac{db_i}{d\delta_{im}} - (\rho-1)\mathbf{1}_m^\top \Sigma \delta^\top \tilde{\omega}\right) \times \\ & \exp\left(-(\rho-1)\tilde{\omega}^\top (\delta\mu - b(\delta) - \log(1+\tau)) + \frac{1}{2}(\rho-1)^2 \tilde{\omega}^\top \delta \Sigma \delta^\top \tilde{\omega}\right) \frac{1}{V'\left(\sum_{i=1}^N g_i(\delta_i)\right)} = 0.\end{aligned}$$

We can again use the implicit function theorem to get (114), where

$$\begin{aligned}\frac{\partial F_{im}}{\partial \delta_{jl}} =& \frac{d^2 g_i}{d\delta_{im} d\delta_{il}} \mathbf{1}_{j=i} + \frac{dg_i}{d\delta_{im}} \times \\ & \left\{ \frac{\frac{d^2 b_i}{d\delta_{im} d\delta_{il}} \mathbf{1}_{j=i} + (\rho-1)\Sigma_{ml}\tilde{\omega}_j}{\mu_m - \frac{db_i}{d\delta_{im}} - (\rho-1)\mathbf{1}_m^\top \Sigma \delta^\top \tilde{\omega}} + \frac{V''(R)}{V'(R)}\frac{dg_j}{d\delta_{jl}} + (\rho-1)\tilde{\omega}_j\left(\mu_l - \frac{db_j}{d\delta_{jl}} - (\rho-1)\mathbf{1}_l^\top \Sigma \delta^\top \tilde{\omega}\right)\right\},\end{aligned}$$

and $\frac{\partial F_{im}}{\partial \chi}$ is the derivative of $F_{im}$with respect to parameter χ. For example, if $\chi = \mu_j$, we have

$$\frac{\partial F_{im}}{\partial \mu_j} = \left(-\frac{\mathbf{1}_{j=m}}{\mu_m - \frac{db_i}{d\delta_{im}} - (\rho-1)\mathbf{1}_m^\top \Sigma \delta^\top \tilde{\omega}} + (\rho-1)\tilde{\omega}^\top \delta \mathbf{1}_j\right)\frac{dg_i}{d\delta_{im}}.$$

Similarly, one can compute derivatives of F_{im} with respect to Σ_{jl}, τ_j and α_{jl} in a straightforward way.

F.6 Choice of numeraire

In the baseline model, we set $W_L = 1$ in all states of the world. In this appendix, we show that this normalization is innocuous. Specifically, we treat W_L as a random variable and show that the equilibrium risk-exposure decisions by firms δ, the aggregate supply of risk risk management resources R and, hence, all macroeconomic aggregates are the same as in the baseline model.

The household's choice of R is governed by the first-order condition (59), which we can write as

$$W_R = -V'(R)\left(V(R)\right)^{-\rho}\frac{\mathrm{E}\left[Y^{1-\rho}\right]}{\mathrm{E}\left[\Lambda\right]}, \tag{115}$$

where Λ is given by (56).

Next, given the Cobb-Douglas production function, the unit cost is given by

$$K_i\left(\delta_i, P, W_L\right) = \exp\left(-\delta_i^{\top}\varepsilon + b_i\left(\delta_i\right)\right) W_L^{1-\sum_{j=1}^{n}\alpha_{ij}} \prod_{j=1}^{N} P_j^{\alpha_{ij}}.$$

Since $P_i = (1+\tau_i)K_i$, this implies that

$$p - w_L = -\tilde{\mathcal{L}}\left(\delta\varepsilon - b\left(\delta\right) - \log\left(1+\tau\right)\right), \tag{116}$$

where $w_L = \log W_L$ and $\tilde{\mathcal{L}} = (I-\alpha)^{-1}$.

Firms' risk exposure solves (11). Taking first-order condition with respect to δ_{im}, we get

$$\frac{dg_i\left(\delta_i\right)}{d\delta_{im}} = \frac{1}{W_R}\frac{\mathrm{E}\left[\Lambda Q_i K_i\left(\varepsilon_m - \frac{db_i(\delta_i)}{d\delta_{im}}\right)\right]}{\mathrm{E}\left[\Lambda\right]}.$$

Combining this equation with (115), we get

$$\frac{dg_i\left(\delta_i\right)}{d\delta_{im}} = -\frac{1}{V'(R)\left(V(R)\right)^{-\rho}}\frac{\mathrm{E}\left[\Lambda K_i Q_i\left(\varepsilon_m - \frac{db_i(\delta_i)}{d\delta_{im}}\right)\right]}{\mathrm{E}\left[Y^{1-\rho}\right]}. \tag{117}$$

Next, from the market-clearing condition (13),

$$P_i Q_i = P_i C_i + \sum_j P_i X_{ji}.$$

Using (55) and the fact that $P_i = (1+\tau_i)K_i$, we can derive that $\frac{K_i Q_i}{W_L + W_R R + \Pi} = \frac{\omega_i}{1+\tau_i}$, where ω is non-stochastic vector of Domar weights, given by (15). Hence, (117) can be rewritten as

$$\frac{dg_i\left(\delta_i\right)}{d\delta_{im}} = -\frac{1}{V'(R)\left(V(R)\right)^{-\rho}}\frac{\omega_i}{1+\tau_i}\frac{\mathrm{E}\left[\Lambda\left(W_L + W_R R + \Pi\right)\left(\varepsilon_m - \frac{db_i(\delta_i)}{d\delta_{im}}\right)\right]}{\mathrm{E}\left[Y^{1-\rho}\right]}.$$

Combining this with (54), we get

$$\frac{dg_i(\delta_i)}{d\delta_{im}} = -\frac{V(R)}{V'(R)}\frac{\omega_i}{1+\tau_i}\frac{\mathrm{E}\left[U'(Y)\,Y\left(\varepsilon_m - \frac{db_i(\delta_i)}{d\delta_{im}}\right)\right]}{\mathrm{E}\left[Y^{1-\rho}\right]} = -\frac{V(R)}{V'(R)}\frac{\omega_i}{1+\tau_i}\frac{\mathrm{E}\left[Y^{1-\rho}\left(\varepsilon_m - \frac{db_i(\delta_i)}{d\delta_{im}}\right)\right]}{\mathrm{E}\left[Y^{1-\rho}\right]}. \tag{118}$$

From (57), we can write GDP Y as

$$Y = \left(\frac{W_L + W_R R + \Pi}{W_L}\right)\prod_{i=1}^{N}\left(\frac{P_i}{W_L}\right)^{-\beta_i}. \tag{119}$$

Furthermore, total profit in the economy is $\Pi = \sum_i \frac{\tau_i}{1+\tau_i} P_i Q_i - W_R R \Leftrightarrow \Pi + W_R R = (W_L + W_R R + \Pi)\sum_i \frac{\tau_i}{1+\tau_i}\omega_i$, which implies $W_R R + \Pi = W_L \frac{\sum_i \frac{\tau_i}{1+\tau_i}\omega_i}{1-\sum_i \frac{\tau_i}{1+\tau_i}\omega_i}$. Therefore, (119) can be written as

$$Y = \frac{1}{1-\sum_i \frac{\tau_i}{1+\tau_i}\omega_i}\prod_{i=1}^{N}\left(\frac{P_i}{W_L}\right)^{-\beta_i}.$$

Hence, using (116), (119) becomes

$$\frac{dg_i(\delta_i)}{d\delta_{im}} = -\frac{V(R)}{V'(R)}\frac{\omega_i}{1+\tau_i}\frac{\mathrm{E}\left[\exp\left(-(\rho-1)\,\Delta^\top\varepsilon\right)\left(\varepsilon_m - \frac{db_i(\delta_i)}{d\delta_{im}}\right)\right]}{\mathrm{E}\left[\exp\left(-(\rho-1)\,\Delta^\top\varepsilon\right)\right]}.$$

Using the properties of the normal distribution to simplify the right-hand side of this equation, it is straightforward to see that this equation coincides with (82). Therefore, setting $W_L = 1$ does not affect equilibrium risk exposures and hence moments of GDP.

F.7 Disaster risk

In this appendix, we explore the impact of disaster risk on the mechanisms of the model (Rietz, 1988; Barro, 2006; Barro and Ursúa, 2012).

F.7.1 Model

The model is the same as in the main text, except for a different distribution for the vector of fundamental risk ε. Specifically, we still assume that the elements $(\varepsilon_2, \ldots, \varepsilon_M)$ of ε are normally distributed, but we allow $\varepsilon_1 = \varepsilon_1^n + \varepsilon_1^d$ to be the sum of a normal component ε_1^n and a disaster component ε_1^d. Disasters ε_1^d follow a Bernoulli distribution such that

$$\varepsilon_1^d = \begin{cases} 0 & \text{with probability } 1-\pi_d, \\ -d & \text{with probability } \pi_d, \end{cases}$$

where π_d is the probability of a disaster and $d > 0$ is its size. As in the main text, the normal components of ε are jointly distributed, $(\varepsilon_1^n, \varepsilon_2, \ldots, \varepsilon_M) \sim \mathcal{N}(\mu, \Sigma)$. We further assume that ε_1^d is independent of all other shocks.

The presence of disaster changes the moments of the GDP process, such that

$$\mathrm{E}[y] = \Delta^\top \mu - \Delta_1 \pi_d d - \bar{b}(\Delta) - \tilde{\omega}^\top \log(1+\tau) - \Gamma_L, \tag{120}$$

$$\mathrm{V}[y] = \Delta^\top \Sigma \Delta + \Delta_1^2 (1-\pi_d)\pi_d d^2. \tag{121}$$

The expression for the expected value of GDP is similar to that in the main text, with the exception of the disaster term $\Delta_1 \pi_d d$. Unsurprisingly, high exposure to disasters Δ_1 lowers $\mathrm{E}[y]$. Frequent (high π_d) and large (high d) disasters also lower the mean of log GDP. Disasters also increase the variance of log GDP, and high Δ_1 or d make this increase larger. Disaster probability π_d close to one half also lead to higher volatility.

F.7.2 Equilibrium characterization

The equilibrium in the economy with disaster can still be characterized as the solution of a distorted planner problem. However, in this case, welfare can no longer be written as a function of the first two moments of log GDP. Specifically, if we define $\bar{h}(\Delta)$, $\bar{b}(\Delta)$ and $g(\Delta)$ as in the baseline model, we have the following result.

Proposition 12. *There exists a unique equilibrium, and its aggregate risk exposure Δ^* solves*

$$\begin{aligned}\mathcal{W}_{dist} := \max_{\Delta} \Delta^\top \mu - \bar{b}(\Delta) - \tilde{\omega}^\top \log(1+\tau) - \log \Gamma_L - \\ \frac{1}{2}(\rho-1)\Delta^\top \Sigma \Delta - \frac{1}{\rho-1}\log\left(1-\pi_d+\pi_d \exp\left((\rho-1)\Delta_1 d\right)\right) - \bar{g}(\Delta).\end{aligned} \tag{122}$$

Without wedges ($\tau = 0$), the equilibrium is efficient.

Compared to (34), we see that the presence of disaster risk lowers welfare. Unsurprisingly, that decline is larger when d and π_d are large, or when the economy is heavily exposed to the first risk factor (Δ_1 large).

The first-order condition associated with (122) is

$$\mathcal{E}(\Delta) = \nabla \bar{h}(\Delta), \tag{123}$$

where risk-adjusted TFP $\mathcal{E}$ now takes the form

$$\mathcal{E}(\Delta) = \mu - (\rho-1)\Sigma\Delta - \mathbf{1}_1 \frac{\pi_d d \exp\left((\rho-1)\Delta_1 d\right)}{1-\pi_d+\pi_d \exp\left((\rho-1)\Delta_1 d\right)}. \tag{124}$$

The last term in this expression captures the impact of disaster risk on $\mathcal{E}$. This term is always negative to exposure to disaster risk is always detrimental to the household's utility. That term is also an increasing function of Δ_1, so that more disaster exposure increases how painful disaster exposure is to the household's utility.

Since, by (123), $\nabla\bar{h}(\Delta) = \bar{H}(\Delta - \Delta^\circ)$, we can rewrite (123) as

$$\Delta - \Delta^\circ = \bar{H}^{-1}\mathcal{E}(\Delta). \tag{125}$$

F.7.3 How the environment affects risk exposure

In the presence of disaster risk, $\mathcal{E}(\Delta)$ is no longer a linear function of Δ, and we cannot solve for Δ in closed form. Nevertheless, we can use (125) and the implicit function theorem to derive how various changes in the environment affect equilibrium risk exposure.

Proposition 13. *The responses of the equilibrium aggregate risk exposure Δ to changes in μ_m, Σ_{mn}, π_d, τ_i are given by*

$$\begin{aligned}
\frac{d\Delta_k}{d\mu_m} &= \left(\mathcal{H}^{-1}\right)_{km} - \frac{\varphi\left(\mathcal{H}^{-1}\right)_{k1}\left(\mathcal{H}^{-1}\right)_{1m}}{1+\varphi\left(\mathcal{H}^{-1}\right)_{11}},\\
\frac{d\Delta_k}{d\Sigma_{mn}} &= -(\rho-1)\left(\mathbf{1}_k^\top\mathcal{H}^{-1} - \frac{\varphi\left(\mathcal{H}^{-1}\right)_{k1}\mathbf{1}_1^\top\mathcal{H}^{-1}}{1+\varphi\left(\mathcal{H}^{-1}\right)_{11}}\right)\left(\frac{1}{2}\mathbf{1}_m\Delta_n + \frac{1}{2}\mathbf{1}_n\Delta_m\right),\\
\frac{d\Delta_k}{d\pi_d} &= -\frac{\left(\mathcal{H}^{-1}\right)_{k1}}{1+\varphi\left(\mathcal{H}^{-1}\right)_{11}}\frac{d\exp\left((\rho-1)\Delta_1 d\right)}{\left[1-\pi_d+\pi_d\exp\left((\rho-1)\Delta_1 d\right)\right]^2},\\
\frac{d\Delta_k}{d\tau_i} &= -\mathbf{1}_k^\top\left(I - \frac{\varphi\left(\mathcal{H}^{-1}\right)_{\cdot 1}\mathbf{1}_1^\top}{1+\varphi\left(\mathcal{H}^{-1}\right)_{11}}\right)\mathcal{H}^{-1}\left(\sum_{j=1}^{N}\frac{\partial\nabla^2\bar{h}}{\partial\kappa_j}\frac{d\kappa_j}{d\tau_i}\right)\left(\nabla^2\bar{h}\right)^{-1}\mathcal{E},
\end{aligned}$$

where $\mathcal{H}$ is given by (39), $\left(\mathcal{H}^{-1}\right)_{\cdot 1} = \left(\left(\mathcal{H}^{-1}\right)_{11},\ldots,\left(\mathcal{H}^{-1}\right)_{11}\right)$, and

$$\varphi = \frac{(1-\pi_d)\pi_d d^2(\rho-1)\exp\left((\rho-1)\Delta_1 d\right)}{\left(1-\pi_d+\pi_d\exp\left((\rho-1)\Delta_1 d\right)\right)^2}. \tag{126}$$

This proposition provides expression for how the parameters of the economy affect the equilibrium aggregate risk exposure vector Δ. In general, how Δ depends on those parameters depends on the potentially complicated complementarity/substitutability patterns between risk factors embedded in matrices $\bar{H}$ and Σ. We can, however, make sharper analytical predictions when the economy is diagonal (see Definition 2)

Corollary 6. *In a diagonal economy, $\frac{d\Delta_k}{d\mu_k} > 0$, $sign\left(\frac{d\Delta_k}{d\Sigma_{kk}}\right) = -sign(\Delta_k)$, $\frac{d\Delta_1}{d\pi_d} < 0$, and $sign\left(\frac{d\Delta_k}{d\tau_i}\right) = -sign(\mathcal{E}_k)$.*

Unsurprisingly, an increase in the expected value μ_k of a risk factor leads to increased exposure

Δ_k to that factor. If a factor k becomes more risky (higher Σ_{kk}) the equilibrium exposure to that factor Δ_k moves toward zero. Indeed, recall from (121) that the impact of ε_k on the variance of GDP depends on the square of Δ. Corollary 6 also shows that if disasters become more likely (higher π_d), then the economy will become less exposed to them. Finally, the impact of wedges is as in the main text (Corollary 3).

To further characterize how the presence of the disaster risk affects the economy, we analyze how an increase in π_d affects the sensitivity of Δ to changes in the environment. To keep the analysis concise, we focus on a diagonal economy which is efficient ($\tau_i = 0, \forall i$).

Corollary 7. *Consider a diagonal efficient economy with no disaster risk, $\pi_d = 0$. An increase in the disaster probability π_d has the following effects on the sensitivity of Δ to μ_1 and Σ_{11}:*

$$\frac{d^2\Delta_1}{d\mu_1 d\pi_d} = -\left[\left(\mathcal{H}^{-1}\right)_{11}\right]^2 d^2\left(\rho-1\right)\exp\left(\left(\rho-1\right)\Delta_1 d\right),$$
$$\frac{d^2\Delta_1}{d\Sigma_{11} d\pi_d} = \left(\rho-1\right)\left[\left(\mathcal{H}^{-1}\right)_{11}\right]^2 d\left(d\left(\rho-1\right)\Delta_1+1\right)\exp\left(\left(\rho-1\right)\Delta_1 d\right).$$

All other cross-derivatives are zeros, i.e., $\frac{d^2\Delta_k}{d\mu_m d\pi_d} = \frac{d^2\Delta_k}{d\Sigma_{mn} d\pi_d} = 0$ unless $k = m = n = 1$.

Corollary 7 shows that disaster risk weakens the response of the economy to changes in the moments of the normal shocks. A higher π_d implies that the disaster component plays a larger role in $\mathcal{E}$, as is evident from (124). Therefore, the equilibrium exposure is affected relatively less by changes in μ and Σ. Consider for instance an increase in μ_1. In the absence of disasters, a higher μ_1 leads to a higher exposure to the first risk factor (Corollary 2). The presence of disasters attenuates this response. As π_d increases, a higher Δ_1 implies more disaster exposure. As a result, the same increase in μ_1 leads to a smaller increase in Δ_1. Similarly, without disasters a higher Σ_{11} implies a lower exposure to the first risk factor if $\Delta_1 > 0$ (Corollary 2). When disaster risk is present, the same decline in Δ_1 has a more positive effect on $\mathcal{E}$ because of the reduced disaster component. Hence, $\frac{d^2\Delta_1}{d\Sigma_{11} d\pi_d} > 0$ if $\Delta_1 > 0$.

F.7.4 Implications for GDP

Next, we analyze how parameters affect the moments of GDP and welfare through the endogenous risk exposure channel.

Proposition 14. *Let χ denote either μ_m, Σ_{mn}, π_d, or τ_i. Then the impact of a change in χ on the moments of log GDP is given by*

$$\frac{d\,\mathrm{E}\left[y\right]}{d\chi} - \frac{\partial\,\mathrm{E}\left[y\right]}{\partial\chi} = \left(\mu - \nabla\bar{b}\right)^\top \frac{d\Delta}{d\chi} - \pi_d d\frac{d\Delta_1}{d\chi}, \quad \frac{d\,\mathrm{V}\left[y\right]}{d\chi} - \frac{\partial\,\mathrm{V}\left[y\right]}{\partial\chi} = 2\Delta^\top\Sigma\frac{d\Delta}{d\chi} + 2\Delta_1\left(1-\pi_d\right)\pi_d d^2\frac{d\Delta_1}{d\chi}, \tag{127}$$

where partial derivatives indicate that Δ is kept fixed, and where $\frac{d\Delta}{d\chi}$ is given in Proposition 13.

Proof. The result follows directly from (120) and (121). □

Is it straightforward to combine these expressions with those of Proposition 13 to capture the full impact of χ on $\mathrm{E}[y]$ and $\mathrm{V}[y]$. Unsurprisingly, disaster risk introduces another channel for changes in the environment to affects the mean and the variance of log GDP. For instance, if a change in χ leads to an increase in Δ_1, this change will result in more destructive disasters and to a decline in $\mathrm{E}[y]$ and an increase in $\mathrm{V}[y]$. However, in normal times when disasters do not occur, this additional channel is not operative. Therefore, it is also useful to evaluate the moments of GDP conditional on no disasters. These are given by (19). In the following corollary, we analyze how the presence of disaster risk affects these moments through the endogenous risk exposure channel.

Corollary 8. *Consider a diagonal efficient economy with no disaster risk, $\pi_d = 0$. Suppose also that the TFP cost of adjusting risk exposure is zero, $b_i = 0$ for all i. An increase in the disaster probability π_d has the following effects on the sensitivity of the moments of log GDP to $\chi \in \{\mu_1, \Sigma_{11}\}$:*

$$\frac{d}{d\pi_d}\left(\frac{d\,\mathrm{E}[y|\textit{no disasters}]}{d\chi} - \frac{\partial\,\mathrm{E}[y|\textit{no disasters}]}{\partial\chi}\right) = \mu_1 \frac{d^2\Delta_1}{d\pi_d d\chi},$$
$$\frac{d}{d\pi_d}\left(\frac{d\,\mathrm{V}[y|\textit{no disasters}]}{d\chi} - \frac{\partial\,\mathrm{V}[y|\textit{no disasters}]}{\partial\chi}\right) = 2\Sigma_{11}\left(\frac{d\Delta_1}{d\pi_d}\frac{d\Delta_1}{d\chi} + \frac{d^2\Delta_1}{d\pi_d d\chi}\Delta_1\right).$$

To understand the intuition behind this Corollary, suppose that $\mu_1 > 0$ and $\Delta_1 > 0$. From Corollary 7, we know that the presence of disaster risk attenuates how changes in μ and Σ affect equilibrium risk exposure. This attenuation effect is also present for the moments of GDP. In particular, an increase in π_d implies a weaker positive response of Δ_1 to μ_1, $\frac{d^2\Delta_1}{d\pi_d d\mu_1} > 0$, resulting in a smaller increase in the expected log GDP. At the same time, the variance of log GDP increases by less (we have $\frac{d\Delta_1}{d\pi_d}\frac{d\Delta_1}{d\mu_1} < 0$ and $\frac{d^2\Delta_1}{d\pi_d d\mu_1}\Delta_1 < 0$). Next, $\frac{d^2\Delta_1}{d\pi_d d\Sigma_{11}} > 0$, which means that an increase in Σ_{11} leads to a smaller decline in the expected log GDP and a smaller increase in the variance of log GDP due to disasters.

Next, consider welfare implications.

Proposition 15. *Without wedges ($\tau = 0$), the impact of μ, Σ, and π_d on welfare is given by*

$$\frac{d\mathcal{W}}{d\mu_m} = \frac{\partial\mathcal{W}}{\partial\mu_m} = \Delta_m, \quad \frac{d\mathcal{W}}{d\Sigma_{mn}} = \frac{\partial\mathcal{W}}{\partial\Sigma_{mn}} = -\frac{1}{2}(\rho-1)\Delta_m\Delta_n,$$
$$\frac{d\mathcal{W}}{d\pi_d} = \frac{\partial\mathcal{W}}{\partial\pi_d} = -\frac{1}{\rho-1}\frac{\exp((\rho-1)\Delta_1 d) - 1}{1 - \pi_d + \pi_d\exp((\rho-1)\Delta_1 d)}.$$

Furthermore, the impact of the wedge τ_i on welfare is given by

$$\frac{d\mathcal{W}}{d\tau_i} = \frac{\partial\mathcal{W}}{\partial\tau_i} = 0.$$

Proof. When $\tau = 0$, the objective function $\mathcal{W}_{dist}$ of the fictitious planner, given by (122), coincides

with welfare. The result then follows from the envelope theorem. □

Proposition 15 shows that the impact of μ, Σ and τ in an economy with disaster risk is as in an economy without disasters (Proposition 6). It also shows that an increase in the likelihood of disasters reduces welfare.It is also straightforward to see how an increase in the disaster probability affects the sensitivity of welfare to changes in μ and Σ.

$$\frac{d^2\mathcal{W}}{d\mu_m d\pi_d} = \frac{d\Delta_m}{d\pi_d} \qquad \text{and} \qquad \frac{d^2\mathcal{W}}{d\Sigma_{mm} d\pi_d} = -(\rho - 1)\,\Delta_m \frac{d\Delta_m}{d\pi_d}.$$

In a diagonal economy, $\frac{d\Delta_1}{d\pi_d} < 0$ and $\frac{d\Delta_m}{d\pi_d} = 0$ if $m \neq 1$ (Corollary 6). A higher μ_1 leads to a higher welfare if the economy is positively exposed to the first risk factor. The disaster risk implies a lower exposure Δ_1, which leads to an attenuated response of $\mathcal{W}$ to μ_1. A higher Σ_{11}, in contrast, harms welfare if the economy is strongly exposed to the first risk factor, i.e., $|\Delta_1|$ is high. A higher π_d lowers Δ_1. Therefore, if $\Delta_1 > 0$, a higher π_d makes the impact of Σ_{11} on $\mathcal{W}$ less negative.

We can push this analysis further. Define $\bar{h}_V$ as in (89). Then Proposition 7 holds in the economy with disasters because, as in the main model, we have $\mathcal{W} = \mathcal{W}_{dist} + \bar{h}(\Delta) - \bar{h}_V(\Delta)$. Furthermore, Lemma 9 also holds, and its proof is unaffected by the presence of the disaster risk.[72]

Finally, the analogues of Corollaries 7 and 8 hold in this economy as well.

Corollary 9. *Suppose that the economy is diagonal and that $\tau_j > 0$ for at least one firm j. Then the following holds.*

$$sign\left(\frac{d\mathcal{W}}{d\mu_m} - \frac{\partial\mathcal{W}}{\partial\mu_m}\right) = sign\left(\mathcal{E}_m\right), \quad sign\left(\frac{d\mathcal{W}}{d\Sigma_{mm}} - \frac{\partial\mathcal{W}}{\partial\Sigma_{mm}}\right) = -sign\left(\Delta_m\mathcal{E}_m\right),$$
$$sign\left(\frac{d\mathcal{W}}{d\pi_d} - \frac{\partial\mathcal{W}}{\partial\pi_d}\right) = -sign\left(\mathcal{E}_1\right).$$

Furthermore, an increase in wedges is more detrimental to welfare when risk-exposure decisions can adjust, that is, $\frac{d\mathcal{W}}{d\tau_i} \leq \frac{\partial\mathcal{W}}{\partial\tau_i}$.

In addition to the results described in the main text, this corollary describes how π_d affects welfare when the equilibrium is inefficient. When disaster risk is bad ($\mathcal{E}_1 < 0$), an increase in the likelihood of disaster has a smaller impact on welfare when risk exposure is endogenous. Intuitively, firms limit their exposure to ε_1 when disaster become more destructive, which limits their adverse effects.

[72] In the proof of Lemma 9 , we use the fact that $\delta_i - \delta_i^\circ = \tilde{\omega}_i\left(\eta\kappa_i G_i + \tilde{\omega}_i B_i\right)^{-1}\left(\nabla^2\bar{h}\right)(\Delta - \Delta^\circ)$. This is also the case in the model with the disaster risk, as follows from (130). Given that, the proof is unaffected by the presence of disaster risk. In particular, the expression for $\nabla\bar{h} - \nabla\bar{h}_V$ is still given by (130), but $\mathcal{E}$ is now given by (124).

F.7.5 Asset pricing moments

As is well-known in the literature (Rietz, 1988; Barro, 2006), the presence of disaster risk can resolve the equity premium and risk-free rate puzzles in a standard endowment economy without an implausibly high risk aversion.[73] The ability of the model to match these moments is important because it implies that it accurately captures the attitude of the representative household to changes in uncertainty. In this section, we show that our model augmented with disaster risk can generate realistic a equity premium and risk-free rate.

Specifically, suppose that there is just one risk factor and firm i's TFP is given by (46). Then output growth is

$$y_t - y_{t-1} = \Delta\left(\varepsilon_t - \varepsilon_{t-1}\right) + \tilde{\omega}^\top\left(v_t - v_{t-1}\right) + \gamma,$$

where $\varepsilon_t = \varepsilon_t^n + \varepsilon_t^d$, $\varepsilon_t^n \sim \mathcal{N}(\mu, \Sigma)$ and ε_t^d takes values 0 with probability $1 - \pi_d$ and $-d$ with probability π_d; $v_{it} \sim \mathcal{N}(\mu_i^v, \Sigma_i^v)$. All shocks are iid. The Euler equation of the representative household is

$$1 = e^{-r}\mathbb{E}_t\left[\frac{D_{t+1} + S_{t+1}}{S_t}\frac{C_{t+1}^{-\rho}\exp\left(\eta\left(\rho-1\right)R_{t+1}\right)}{C_t^{-\rho}\exp\left(\eta\left(\rho-1\right)R_t\right)}\right],$$

where e^{-r} is the time discount factor, D_t is real dividend, S_t is the real asset price and R_t is the amount of risk management resources. Under our shock specification, risk exposures are time invariant, and hence $R_{t+1} = R_t$. Then the Euler equation simplifies to

$$1 = e^{-r}\mathbb{E}_t\left[\frac{D_{t+1} + S_{t+1}}{S_t}\left(\frac{C_{t+1}}{C_t}\right)^{-\rho}\right].$$

Consider first a risk-free asset that always pays out one, even if a disaster happens. Then the risk-free rate is

$$\log \mathbb{E}R_{f,t+1} = \log \mathbb{E}\left\{e^{-r}\mathbb{E}_t\left[\exp\left(-\rho\left(y_{t+1} - y_t\right)\right)\right]\right\}^{-1} = r + \rho\gamma.$$

Consider now a consumption claim, such that $D_t = C_t$.[74] Define the price-consumption ratio $\phi_t = \frac{S_t}{C_t}$. Then the Euler equation can be written as

$$1 = e^{-r}\mathbb{E}_t\left[\frac{1+\phi_{t+1}}{\phi_t}\exp\left(-\left(\rho-1\right)\left(\Delta\left(\varepsilon_{t+1} - \varepsilon_t\right) + \tilde{\omega}^\top\left(v_{t+1} - v_t\right) + \gamma\right)\right)\right]. \tag{128}$$

[73] Additional features, such as Epstein-Zin preferences, are needed to explain other asset pricing moments; see, e.g., Barro (2009) and Wachter (2013).

[74] Equivalently, we can consider a claim on the total firms' profits. The equity premium is going to be the same because real profits in our our model are proportional to aggregate output.

Conjecture that

$$\phi_t = \phi_0 \exp\left(\phi_\varepsilon \varepsilon_t + \phi_v^\top v_t\right).$$

Plugging this into (128), we can immediately see that $\phi_\varepsilon = (\rho-1)\Delta$ and $\phi_v = (\rho-1)\tilde{\omega}$. Then

$$\begin{aligned}\phi_0 =& \frac{\exp(-r-(\rho-1)\gamma)}{1-\exp(-r-(\rho-1)\gamma)}\mathbb{E}_t\left[\exp\left(-(\rho-1)\left(\Delta\varepsilon_{t+1}+\tilde{\omega}^\top v_{t+1}\right)\right)\right] = \\ =& \frac{\exp(-r-(\rho-1)\gamma)}{1-\exp(-r-(\rho-1)\gamma)}\exp\left(-(\rho-1)\left(\Delta\mu+\tilde{\omega}^\top\mu^v\right)+\frac{1}{2}(\rho-1)^2\left(\Delta^2\Sigma+\tilde{\omega}^\top\Sigma^v\tilde{\omega}\right)\right)\times \\ & \left[1-\pi_d+\pi_d\exp\left((\rho-1)\Delta d\right)\right].\end{aligned}$$

The expected return on the consumption claim is

$$\log\mathbb{E}R_{c,t+1} = \log\mathbb{E}\left(\frac{C_{t+1}+S_{t+1}}{S_t}\right) = \log\mathbb{E}\left(\frac{1+\phi_{t+1}}{\phi_t}\frac{C_{t+1}}{C_t}\right).$$

Omitting tedious yet straightforward derivations, we get

$$\begin{aligned}\log\mathbb{E}R_{c,t+1} =& \gamma+\rho\left(\Delta^2\Sigma+\tilde{\omega}^\top\Sigma^v\tilde{\omega}\right)+\log\left(1-\pi_d+\pi_d\exp(\rho\Delta d)\right) \\ & +\log\left\{\exp\left((\rho-1)\rho\left(\Delta^2\Sigma+\tilde{\omega}^\top\Sigma^v\tilde{\omega}\right)\right)\left(1-\pi_d+\pi_d\exp(-\rho\Delta d)\right)+\right. \\ & \left.\frac{1-\exp(-r-(\rho-1)\gamma)}{\exp(-r-(\rho-1)\gamma)}\frac{1-\pi_d+\pi_d\exp(-\Delta d)}{1-\pi_d+\pi_d\exp\left((\rho-1)\Delta d\right)}\right\}.\end{aligned}$$

Using parameters from the calibration of Section 9.5 ($\rho = 5$, $\Delta = 0.012$, $d = 24.6$, $\pi_d = 0.017$, $\sqrt{\Delta^2\Sigma+\tilde{\omega}^\top\Sigma^v\tilde{\omega}} = 0.017$) and setting $\gamma = 0$ (no real GDP per capita growth in Spanish economy in the past 20 years), $r = 0.02$ (so that risk-free rate is 2%), we get $\log\mathbb{E}R_{f,t+1} = 2\%$ and $\log\mathbb{E}R_{c,t+1} - \log\mathbb{E}R_{f,t+1} = 4.70\%$. Without disasters ($\pi_d = 0$), the equity premium becomes 0.68%.

F.7.6 Proofs

Proof of Proposition 12

Proposition 12. *There exists a unique equilibrium, and its aggregate risk exposure Δ^* solves*

$$\begin{aligned}\mathcal{W}_{dist} := \max_{\Delta}& \Delta^\top\mu-\bar{b}(\Delta)-\tilde{\omega}^\top\log(1+\tau)-\log\Gamma_L- \\ & \frac{(\rho-1)}{2}\Delta^\top\Sigma\Delta-\frac{1}{\rho-1}\log\left(1-\pi_d+\pi_d\exp\left((\rho-1)\Delta_1 d\right)\right)-\bar{g}(\Delta).\end{aligned} \tag{122}$$

Without wedges ($\tau = 0$), the equilibrium is efficient.

Proof. We first show that the set of equilibrium allocations coincides with the set of solutions to a maximization problem. Consider the maximization problem

$$\max_{\delta} \tilde{\omega}^\top \delta\mu - \tilde{\omega}^\top b(\delta) - \tilde{\omega}^\top \log(1+\tau) - \log\Gamma_L - \frac{1}{2}(\rho-1)\,\tilde{\omega}^\top \delta\Sigma\delta^\top\tilde{\omega} - \frac{1}{\rho-1}\log\left(1-\pi_d+\pi_d\exp\left((\rho-1)\,d\sum_{i=1}^{N}\delta_{i1}\tilde{\omega}_i\right)\right) + \log\left(V\left(\sum_{i=1}^{N}\frac{\tilde{\omega}_i(1+\tau_i)}{\omega_i}g_i(\delta_i)\right)\right). \quad (129)$$

Notice that we can write this problem as

$$\max_{\Delta} \Delta^\top\mu - \tilde{\omega}^\top\log(1+\tau) - \log\Gamma_L - \frac{1}{2}(\rho-1)\,\Delta^\top\Sigma\Delta - \frac{1}{\rho-1}\log\left(1-\pi_d+\pi_d\exp\left((\rho-1)\,\Delta_1 d\right)\right) - \min_{\delta \text{ s.t. } \Delta=\delta^\top\tilde{\omega}}\left\{\tilde{\omega}^\top b(\delta) - \log\left(V\left(\sum_{i=1}^{N}\frac{\tilde{\omega}_i(1+\tau_i)}{\omega_i}g_i(\delta_i)\right)\right)\right\},$$

which is the maximization problem (122). Note also that the objective function of this problem is strictly concave because $\bar{h}(\Delta)$ is strictly convex. This implies that there is a unique solution to the maximization problems (129) and (34) and that the first-order conditions are sufficient to characterize it. To complete the proof, we will show that the equilibrium conditions coincide with these first-order conditions. This will imply that there is a unique equilibrium.

The first-order condition of the fictitious planner with respect to δ_{im} is

$$\delta_i - \delta_i^\circ = H_i^{-1}\left[\mu - (\rho-1)\,\Sigma\Delta - \mathbf{1}_1\frac{\pi_d d\exp\left((\rho-1)\,\Delta_1 d\right)}{1-\pi_d+\pi_d\exp\left((\rho-1)\,\Delta_1 d\right)}\right], \quad (130)$$

where H_i is given by (24) and $\mathbf{1}_1 = (1,0,\ldots,0)$ is the $M\times 1$ basis vector. The system of these $M\times N$ equations fully characterizes the unique solution to the fictitious planner's problem.

Now consider the equilibrium. One can show that the firm's problem (21) is strictly convex, such that its first-order conditions are necessary and sufficient. As in the model of the main text, they are given by

$$\mathrm{E}\left[\left(\varepsilon - \nabla b_i(\delta_i)\right)K_i Q_i\right] + \mathrm{Cov}\left(\left(\varepsilon-\nabla b_i(\delta_i)\right)K_i Q_i, \frac{\Lambda}{\mathrm{E}[\Lambda]}\right) = \nabla g_i(\delta_i)\,W_R,$$

where we used the fact that $\partial K_i/\partial\delta_i = -\left(\varepsilon-\nabla b_i(\delta_i)\right)K_i$. Next, the definition of revenue-based Domar weights implies that $\omega_i = \frac{P_i Q_i}{PY} \Leftrightarrow Q_i = \frac{\omega_i\Gamma_L^{-1}}{P_i}$. Together with (12), we can therefore write $K_i Q_i = \omega_i\Gamma_L^{-1}/(1+\tau_i)$, such that $K_i Q_i$ is deterministic. The first-order conditions become

$$\frac{dg_i}{d\delta_{im}} = \frac{1}{W}\frac{\mathbb{E}\left[\Lambda\frac{\left(\mathcal{H}^{-1}\right)_{kM}\iota_M^\top\mathcal{H}^{-1}\frac{(1-\pi_d)\pi_d d^2(\rho-1)\exp\left((\rho-1)\Delta_M d\right)}{\left[1-\pi_d+\pi_d\exp\left((\rho-1)\Delta_M d\right)\right]^2}}{1+\left(\mathcal{H}^{-1}\right)_{MM}\frac{(1-\pi_d)\pi_d d^2(\rho-1)\exp\left((\rho-1)\Delta_M d\right)}{\left[1-\pi_d+\pi_d\exp\left((\rho-1)\Delta_M d\right)\right]^2}}K_i Q_i\left(\varepsilon_m - \frac{db_i}{d\delta_{im}}\right)\right]}{\mathbb{E}\Lambda}.$$

Using (84) for the expression of Λ, we can simplify this equation to (130), which completes the proof. □

Proof of Proposition 13

Proposition 13. *The responses of the equilibrium aggregate risk exposure Δ to changes in μ_m, Σ_{mn}, π_d,τ_i are given by*

$$\frac{d\Delta_k}{d\mu_m} = \left(\mathcal{H}^{-1}\right)_{km} - \frac{\varphi\left(\mathcal{H}^{-1}\right)_{k1}\left(\mathcal{H}^{-1}\right)_{1m}}{1+\varphi\left(\mathcal{H}^{-1}\right)_{11}}, \tag{131}$$

$$\frac{d\Delta_k}{d\Sigma_{mn}} = -\left(\rho-1\right)\left(\mathbf{1}_k^\top\mathcal{H}^{-1} - \frac{\varphi\left(\mathcal{H}^{-1}\right)_{k1}\mathbf{1}_1^\top\mathcal{H}^{-1}}{1+\varphi\left(\mathcal{H}^{-1}\right)_{11}}\right)\left(\frac{1}{2}\mathbf{1}_m\Delta_n + \frac{1}{2}\mathbf{1}_n\Delta_m\right), \tag{132}$$

$$\frac{d\Delta_k}{d\pi_d} = -\frac{\left(\mathcal{H}^{-1}\right)_{k1}}{1+\varphi\left(\mathcal{H}^{-1}\right)_{11}}\frac{d\exp\left(\left(\rho-1\right)\Delta_1 d\right)}{\left[1-\pi_d+\pi_d\exp\left(\left(\rho-1\right)\Delta_1 d\right)\right]^2}, \tag{133}$$

$$\frac{d\Delta_k}{d\tau_i} = -\mathbf{1}_k^\top\left(I - \frac{\varphi\left(\mathcal{H}^{-1}\right)_{\cdot 1}\mathbf{1}_1^\top}{1+\varphi\left(\mathcal{H}^{-1}\right)_{11}}\right)\mathcal{H}^{-1}\left(\sum_{j=1}^N \frac{\partial\nabla^2\bar{h}}{\partial\kappa_j}\frac{d\kappa_j}{d\tau_i}\right)\left(\nabla^2\bar{h}\right)^{-1}\mathcal{E}, \tag{134}$$

where $\mathcal{H}$ is given by (39), $\left(\mathcal{H}^{-1}\right)_{\cdot 1} = \left(\left(\mathcal{H}^{-1}\right)_{11},\ldots,\left(\mathcal{H}^{-1}\right)_{11}\right)$, and

$$\varphi = \frac{\left(1-\pi_d\right)\pi_d d^2\left(\rho-1\right)\exp\left(\left(\rho-1\right)\Delta_1 d\right)}{\left(1-\pi_d+\pi_d\exp\left(\left(\rho-1\right)\Delta_1 d\right)\right)^2}. \tag{135}$$

Proof. This proposition directly follows from the implicit function theorem applied to (125). Specifically, denote

$$\mathcal{F} = \Delta - \Delta^\circ - \bar{H}^{-1}\left(\mu - \left(\rho-1\right)\Sigma\Delta - \mathbf{1}_1\frac{\pi_d d\exp\left(\left(\rho-1\right)\Delta_1 d\right)}{1-\pi_d+\pi_d\exp\left(\left(\rho-1\right)\Delta_1 d\right)}\right).$$

The Jacobian of $\mathcal{F}$ is

$$\begin{aligned} J_{\mathcal{F}} &= I + \bar{H}^{-1}\left(\rho-1\right)\Sigma + \varphi\left(\bar{H}^{-1}\right)_{\cdot 1}\mathbf{1}_1^\top = . \\ &= \bar{H}^{-1}\left[\mathcal{H} + \varphi\mathbf{1}_1\mathbf{1}_1^\top\right], \end{aligned}$$

where φ is given by (135). Using the Sherman-Morrison formula, we get

$$J_{\mathcal{F}}^{-1} = \left(I - \frac{\varphi\left(\mathcal{H}^{-1}\right)_{\cdot 1}\mathbf{1}_1^\top}{1+\varphi\left(\mathcal{H}^{-1}\right)_{11}}\right)\mathcal{H}^{-1}\bar{H}.$$

Then, for a given parameter χ,

$$\frac{d\Delta_k}{d\chi} = -\mathbf{1}_k^\top J_{\mathcal{F}}^{-1}\frac{\partial\mathcal{F}}{\partial\chi}.$$

In particular, if χ is μ_m, Σ_{mn}, π_d, we have

$$\frac{d\Delta_k}{d\chi} = \mathbf{1}_k^\top \left(I - \frac{\varphi \left(\mathcal{H}^{-1}\right)_{\cdot 1} \mathbf{1}_1^\top}{1 + \varphi \left(\mathcal{H}^{-1}\right)_{11}} \right) \mathcal{H}^{-1} \frac{\partial \mathcal{E}}{\partial \chi},$$

where $\mathcal{E}$ is given by (124). Straightforward differentiating of $\mathcal{E}$ with respect to different parameters yields the result. If $\chi = \tau_i$, we have

$$\frac{d\Delta_k}{d\tau_i} = \mathbf{1}_k^\top \left(I - \frac{\varphi \left(\mathcal{H}^{-1}\right)_{\cdot 1} \mathbf{1}_1^\top}{1 + \varphi \left(\mathcal{H}^{-1}\right)_{11}} \right) \mathcal{H}^{-1} \bar{H} \frac{\partial \bar{H}^{-1}}{\partial \tau_i} \mathcal{E},$$

which immediately simplifies to (134). □

Proof of Corollary 6

Corollary 6. *In a diagonal economy,* $\frac{d\Delta_k}{d\mu_k} > 0$, $sign\left(\frac{d\Delta_k}{d\Sigma_{kk}}\right) = -sign\left(\Delta_k\right), \frac{d\Delta_1}{d\pi_d} < 0$, *and* $sign\left(\frac{d\Delta_k}{d\tau_i}\right) = -sign\left(\mathcal{E}_k\right).$

Proof. In a diagonal economy, $\mathcal{H}_{ij} = 0$ if $i \neq j$. Then

$$\frac{d\Delta_k}{d\mu_k} = \begin{cases} \left(\mathcal{H}^{-1}\right)_{kk} & k \neq 1, \\ \frac{\left(\mathcal{H}^{-1}\right)_{11}}{1+\varphi\left(\mathcal{H}^{-1}\right)_{11}} & k = 1. \end{cases}$$

Since $\mathcal{H}^{-1}$ is a positive definite matrix, $\frac{d\Delta_k}{d\mu_k} > 0$.

Next,

$$\frac{d\Delta_k}{d\Sigma_{kk}} = \begin{cases} -\left(\rho - 1\right)\left(\mathcal{H}^{-1}\right)_{kk} \Delta_k & k \neq 1, \\ -\left(\rho - 1\right) \frac{\left(\mathcal{H}^{-1}\right)_{11}}{1+\varphi\left(\mathcal{H}^{-1}\right)_{11}} \Delta_1 & k = 1, \end{cases}$$

and so $\operatorname{sign}\left(\frac{d\Delta_k}{d\Sigma_{kk}}\right) = -\operatorname{sign}\left(\Delta_k\right)$.

$\frac{d\Delta_1}{d\pi_d} < 0$ directly follows from (133). Finally,

$$\frac{d\Delta_k}{d\tau_i} = -\mathbf{1}_k^\top \begin{pmatrix} \frac{1}{1+\varphi\left(\mathcal{H}^{-1}\right)_{11}} & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 \end{pmatrix} \mathcal{H}^{-1} \left(\sum_{j=1}^{N} \frac{\partial \bar{H}}{\partial \kappa_j} \frac{d\kappa_j}{d\tau_i} \right) \bar{H}^{-1} \mathcal{E}. \tag{136}$$

In a diagonal economy, the matrix between $\mathbf{1}_k^\top$ and $\mathcal{E}$ is a diagonal matrix with a positive main diagonal. Hence, $\operatorname{sign}\left(\frac{d\Delta_k}{d\tau_i}\right) = -\operatorname{sign}\left(\mathcal{E}_k\right)$. □

Proof of Corollary 7

Corollary 7. *Consider a diagonal efficient economy with no disaster risk, $\pi_d = 0$. An increase in the disaster probability π_d has the following effects on the sensitivity of Δ to μ_1 and Σ_{11}:*

$$\frac{d^2\Delta_1}{d\mu_1 d\pi_d} = -\left[\left(\mathcal{H}^{-1}\right)_{11}\right]^2 d^2\left(\rho-1\right)\exp\left(\left(\rho-1\right)\Delta_1 d\right),$$
$$\frac{d^2\Delta_1}{d\Sigma_{11} d\pi_d} = \left(\rho-1\right)\left[\left(\mathcal{H}^{-1}\right)_{11}\right]^2 d\left(d\left(\rho-1\right)\Delta_1+1\right)\exp\left(\left(\rho-1\right)\Delta_1 d\right).$$

All other cross-derivatives are zeros, i.e., $\frac{d^2\Delta_k}{d\mu_m d\pi_d} = \frac{d^2\Delta_k}{d\Sigma_{mn} d\pi_d} = 0$ unless $k = m = n = 1$.

Proof. First, notice that $\varphi\left(\pi_d\right)$ is given by (126), which implies

$$\frac{d\varphi}{d\pi_d} \overset{\pi_d=0}{=} d^2\left(\rho-1\right)\exp\left(\left(\rho-1\right)\Delta_1 d\right).$$

Using this expression when differentiating (131) and (132) with respect to π_d, we get

$$\frac{d^2\Delta_1}{d\mu_1 d\pi_d} \overset{\pi_d=0}{=} -\left[\left(\mathcal{H}^{-1}\right)_{11}\right]^2 d^2\left(\rho-1\right)\exp\left(\left(\rho-1\right)\Delta_1 d\right).$$
$$\frac{d^2\Delta_1}{d\Sigma_{11} d\pi_d} \overset{\pi_d=0}{=} \left(\rho-1\right)\left[\left[\left(\mathcal{H}^{-1}\right)_{11}\right]^2 d^2\left(\rho-1\right)\exp\left(\left(\rho-1\right)\Delta_1 d\right)\Delta_1 - \left(\mathcal{H}^{-1}\right)_{11}\frac{d\Delta_1}{d\pi_d}\right] =$$
$$= \left(\rho-1\right)\left[\left(\mathcal{H}^{-1}\right)_{11}\right]^2 d\exp\left(\left(\rho-1\right)\Delta_1 d\right)\left(d\left(\rho-1\right)\Delta_1+1\right),$$

where $\frac{d\Delta_1}{d\pi_d}$ is given by (133). Clearly, in the diagonal economy all other cross-derivatives are zero because $\mathcal{H}$ is a diagonal matrix. □

Proof of Corollary 8

Corollary 8. *Consider a diagonal efficient economy with no disaster risk, $\pi_d = 0$. Suppose also that the TFP cost of adjusting risk exposure is zero, $b_i = 0$ for all i. An increase in the disaster probability π_d has the following effects on the sensitivity of the moments of log GDP to $\chi \in \{\mu_1, \Sigma_{11}\}$:*

Proof. From (19),

$$\frac{d\,\mathrm{E}\left[y|\text{no disasters}\right]}{d\chi} - \frac{\partial\,\mathrm{E}\left[y|\text{no disasters}\right]}{\partial\chi} = \mu^\top\frac{d\Delta}{d\chi},$$
$$\frac{d\,\mathrm{V}\left[y|\text{no disasters}\right]}{d\chi} - \frac{\partial\,\mathrm{V}\left[y|\text{no disasters}\right]}{\partial\chi} = 2\Delta^\top\Sigma\frac{d\Delta}{d\chi}.$$

In a diagonal economy, π_d has no effect on $\Delta_2, \ldots, \Delta_M$. Furthermore, Σ is diagonal. Then the result follows from a straightforward differentiation of the above expressions with respect to π_d. □

Proof of Corollary 9

Corollary 9. *Suppose that the economy is diagonal and that $\tau_j > 0$ for at least one firm j. Then the following holds.*

$$sign\left(\frac{d\mathcal{W}}{d\mu_m} - \frac{\partial \mathcal{W}}{\partial \mu_m}\right) = sign\left(\mathcal{E}_m\right), \quad sign\left(\frac{d\mathcal{W}}{d\Sigma_{mm}} - \frac{\partial \mathcal{W}}{\partial \Sigma_{mm}}\right) = -sign\left(\Delta_m \mathcal{E}_m\right),$$
$$sign\left(\frac{d\mathcal{W}}{d\pi_d} - \frac{\partial \mathcal{W}}{\partial \pi_d}\right) = -sign\left(\mathcal{E}_1\right).$$

Furthermore, an increase in wedges is more detrimental to welfare when risk-exposure decisions can adjust, that is, $\frac{d\mathcal{W}}{d\tau_i} \leq \frac{\partial \mathcal{W}}{\partial \tau_i}$.

Proof. Using (92), we can write (90) as

$$\frac{d\mathcal{W}}{d\chi} - \frac{\partial \mathcal{W}}{\partial \chi} = \left(\frac{d\Delta}{d\chi}\right)^{\top} \bar{H} \left(\sum_{i=1}^{N} \eta\left(\kappa_i - 1\right) H_i^{-1} G_i H_i^{-1}\right) \mathcal{E}. \tag{137}$$

From Proposition 13, in a diagonal economy we have $\frac{d\Delta_k}{d\mu_k} > 0$ and $\frac{d\Delta_k}{d\mu_m} = 0$ if $k \neq m$. Furthermore, $\frac{d\Delta_1}{d\pi_d} < 0$ and $\frac{d\Delta_k}{d\pi_d} = 0$ if $k \neq 1$. Finally, $\text{sign}\left(\frac{d\Delta_k}{d\Sigma_{kk}}\right) = \text{sign}\left(\Delta_k\right)$ and $\frac{d\Delta_k}{d\Sigma_{mm}} = 0$ if $k \neq m$. The result for μ_m,Σ_{mm}, and π_d then follows from the fact that $\bar{H}\left(\sum_{i=1}^{N} \eta\left(\kappa_i - 1\right) H_i^{-1} G_i H_i^{-1}\right)$ is a diagonal matrix with a positive main diagonal.

Finally, from 136 we know that in a diagonal economy

$$\left(\frac{d\Delta}{d\tau_i}\right)^{\top} = -\mathcal{E}^{\top} \bar{H}^{-1} \left(\sum_{j=1}^{N} \frac{\partial \bar{H}}{\partial \kappa_j} \frac{d\kappa_j}{d\tau_i}\right) \mathcal{H}^{-1} \begin{pmatrix} \frac{1}{1+\varphi(\mathcal{H}^{-1})_{11}} & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 \end{pmatrix}.$$

Combining this with (137), we get that $\frac{d\mathcal{W}}{d\tau_i} - \frac{\partial \mathcal{W}}{\partial \tau_i} \leq 0$ in a diagonal economy. □

References

Acemoglu, Daron and Fabrizio Zilibotti (1997). "Was Prometheus unbound by chance? Risk, diversification, and growth". *Journal of Political Economy* 105.4, pp. 709–751.

Ackerberg, Daniel A, Kevin Caves, and Garth Frazer (2015). "Identification properties of recent production function estimators". *Econometrica* 83.6, pp. 2411–2451.

Arellano, Cristina, Yan Bai, and Jing Zhang (2012). "Firm dynamics and financial development". *Journal of Monetary Economics* 59.6, pp. 533–549.

Baqaee, David Rezza and Emmanuel Farhi (2019). "Productivity and Misallocation in General Equilibrium". *Quarterly Journal of Economics* 135.1, pp. 105–163.

Barro, Robert J (2006). "Rare disasters and asset markets in the twentieth century". *Quarterly Journal of Economics* 121.3, pp. 823–866.

— (2009). "Rare disasters, asset prices, and welfare costs". *American Economic Review* 99.1, pp. 243–264.

Barro, Robert J and José F Ursúa (2012). "Rare macroeconomic disasters". *Annu. Rev. Econ.* 4.1, pp. 83–109.

Bigio, Saki and Jennifer La'O (2020). "Distortions in Production Networks". *Quarterly Journal of Economics* 135.4, pp. 2187–2253.

Blundell, Richard and Stephen Bond (2000). "GMM estimation with persistent panel data: an application to production functions". *Econometric reviews* 19.3, pp. 321–340.

Bond, Steve et al. (2021). "Some unpleasant markup arithmetic: Production function elasticities and their estimation from production data". *Journal of Monetary Economics* 121, pp. 1–14.

Burstein, Ariel, Vasco M Carvalho, and Basile Grassi (2024). "Bottom-up Markup Fluctuations".

De Loecker, Jan and Jan Eeckhout (2018). *Global market power*. Tech. rep. National Bureau of Economic Research.

De Loecker, Jan, Jan Eeckhout, and Gabriel Unger (2020). "The rise of market power and the macroeconomic implications". *Quarterly Journal of Economics* 135.2, pp. 561–644.

De Loecker, Jan and Frederic Warzynski (2012). "Markups and firm-level export status". *American Economic Review* 102.6, pp. 2437–2471.

De Ridder, Maarten, Basile Grassi, Giovanni Morzenti, et al. (2024). "The Hitchhiker's Guide to Markup Estimation".

Dell, Melissa, Benjamin F Jones, and Benjamin A Olken (2012). "Temperature shocks and economic growth: Evidence from the last half century". *American Economic Journal: Macroeconomics* 4.3, pp. 66–95.

Fama, Eugene F and Kenneth R French (1992). "The cross-section of expected stock returns". *Journal of Finance* 47.2, pp. 427–465.

Gopinath, Gita et al. (2017). "Capital allocation and productivity in South Europe". *Quarterly Journal of Economics* 132.4, pp. 1915–1967.

Gorodnichenko, Yuriy and Monika Schnitzer (2013). "Financial constraints and innovation: Why poor countries don't catch up". *Journal of the European Economic association* 11.5, pp. 1115–1152.

Hotelling, Harold (1929). "Stability in competition". *Economic Journal* 39.153, pp. 41–57.

Ibbotson, Roger G, Paul D Kaplan, and James D Peterson (1997). "Estimates of small-stock betas are much too low". *Journal of Portfolio Management* 23, pp. 104–111.

Kalemli-Özcan, Şebnem et al. (2024). "How to Construct Nationally Representative Firm-Level Data from the Orbis Global Database: New Facts on SMEs and Aggregate Implications for Industry Concentration". *American Economic Journal: Macroeconomics* 16.2, pp. 353–374.

Kelly, Bryan T, Seth Pruitt, and Yinan Su (2019). "Characteristics are covariances: A unified model of risk and return". *Journal of Financial Economics* 134.3, pp. 501–524.

— (2020). "Instrumented principal component analysis". *Available at SSRN 2983919*.

Lerner, Abba P (1934). "Monopoly and the measurement of monopoly power". *Review of Economic Studies* 1, pp. 157–175.

Levinsohn, James and Amil Petrin (2003). "Estimating production functions using inputs to control for unobservables". *Review of Economic Studies* 70.2, pp. 317–341.

Liu, Tianzi and Zebang Xu (2024). *The (Mis) Allocation Channel of Climate Change*. Tech. rep.

Midrigan, Virgiliu and Daniel Yi Xu (2014). "Finance and misallocation: Evidence from plant-level data". *American economic review* 104.2, pp. 422–458.

Olley, G Steven and Ariel Pakes (1996). "The dynamics of productivity in the telecommunications equipment industry". *Econometrica* 64.6, p. 1263.

Raval, Devesh (2023). "Testing the production approach to markup estimation". *Review of Economic Studies* 90.5, pp. 2592–2611.

Rietz, Thomas A (1988). "The equity risk premium a solution". *Journal of monetary Economics* 22.1, pp. 117–131.

Rovigatti, Gabriele and Vincenzo Mollisi (2018). "Theory and practice of total-factor productivity estimation: The control function approach using Stata". *The Stata Journal* 18.3, pp. 618–662.

Stock, James H and Mark W Watson (2002). "Forecasting using principal components from a large number of predictors". *Journal of the American statistical association* 97.460, pp. 1167–1179.

Wachter, Jessica A (2013). "Can time-varying risk of rare disasters explain aggregate stock market volatility?" *Journal of Finance* 68.3, pp. 987–1035.

Zhang, Peng et al. (2018). "Temperature effects on productivity and factor reallocation: Evidence from a half million Chinese manufacturing plants". *Journal of Environmental Economics and Management* 88, pp. 1–17.