The Origin of Risk*

November 11, 2024

Abstract

We propose a model in which risk, at both the micro and the macro levels, is endogenous. In the model, each firm can choose the mean and the variance of its productivity process, as well as how it covaries with the productivity of other firms. To study the aggregate impact of these decisions, we embed the firms into an otherwise standard production network economy. Through their impact on risk-taking decisions, distortions such as taxes and markups can make GDP more volatile in equilibrium. The theory also predicts that the productivity of larger firms and those with smaller markups is less volatile and less correlated with aggregate productivity. We find support for these predictions in the data. In a calibrated version of the model, removing distortions significantly reduces GDP volatility.

JEL Classifications: E32, C67, D57, D80, D85

*We thank Julieta Caunedo, Philipp Kircher, Narayana Kocherlakota, Ezra Oberfield, Giulio Trigilia and participants at various seminars and conferences for helpful comments.

1 Introduction

Risk influences the economy in many ways. At the micro level, it affects household decisions to consume and save, as well as firm choices to produce, invest and innovate. At the macro level, it manifests itself through aggregate fluctuations and influences decisions that are crucial for growth. Governments also enact policies to mitigate long-term risks like climate change.

Given its importance, understanding the origin and determinants of risk is crucial. Most of macroeconomics treats risk as exogenous by assuming that productivity, whether at the aggregate or firm level, follows a predefined stochastic process. 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 can be influenced by 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 who to hire, how to organize production, what projects to pursue, where to locate a plant, which markets to enter, etc. As an example, growing crops near the shore might provide a steady water supply for irrigation, but it also increases vulnerability to flooding. Conversely, growing the same crops inland might reduce flood risk at the expense of higher drought vulnerability. When aggregated, these individual exposure decisions shape the risk profile of the entire economy. For instance, if many firms choose to locate by the shore, crop yields become correlated, turning flooding into an aggregate risk factor.

Some stylized features of the data also 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 the size of markups. 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

[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 evidence that aggregate TFP volatility decreases with per-capita income and increases with the share of government expenditure in GDP in a cross-section of countries. In Section 7, we provide evidence that the productivity of larger firms and those with lower markups is less volatile and less correlated with aggregate TFP.

adjust their exposure to those risk factors.

In the model, managing a firm's risk exposure requires the use of resources. For instance, avoiding droughts by planting crops near the shore may involve renting an expensive piece of land, navigating a rapidly changing market might require skilled managers, etc. These resources are provided by a representative, risk-averse household at a cost in terms of utility. When making risk-exposure decisions, firms balance this cost with the benefit of exposure. Consequently, the presence of risk in this economy is a choice. There could be no aggregate risk at all, 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 so firms 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 sales 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, this leads to a form of diversification, with firms 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. 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 reduce the cost of managing risk. Because of the wedges, risk-exposure decisions are in general 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 the 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. Larger firms therefore 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 stock market. The mechanisms of the model provide 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 aggregate stock market, consistent 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 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 2.4% to 1.7%. This suggests that the quantitative impact of distortions such as taxes and markups on aggregate volatility can be substantial. Without endogenous risk decisions, the removal of wedges would have no impact on the volatility of GDP.

We also conduct a second experiment in which we mimic an uncertainty shock by doubling the volatility of the economy's underlying risk factor. If firms were unable to adjust their risk-taking decisions, this would increase GDP volatility by 70 basis points. In contrast, when firms are free to manage their risk, they strongly reduce their exposure in response to the large increase in fundamental risk, and the volatility of GDP increases by only 20 basis points. This last finding suggests that frictions and policies that impede firms' ability to manage risk might have sizable detrimental effects on aggregate volatility.

Literature review

The main contribution of this paper is to propose a theory of aggregate fluctuations driven by endogenous risk exposure, thus connecting incentives and risk within the economy. In contrast, standard representative agent models of the business cycles such as Kydland and Prescott (1982) assume that 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). 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 effectively 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 et al. (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 cycles 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 here. We also relate to the literature on wedges in production network economies, which includes Jones (2011), Baqaee and Farhi (2019a), 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.[2]

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

2 A model of endogenous risk

We study the origin of risk in an otherwise standard production network economy under uncertainty. The economy is populated by a set of firms, each producing a differentiated good that can be used either as intermediate input or for consumption. Firms can influence their productivity

[2]Our work also relates to the endogenous network literature as in Oberfield (2018), Acemoglu and Azar (2020) and Kopytov et al. (2024). In those papers, firms choose a technique that specifies their suppliers. In our paper, firms can be interpreted as selecting a technique that specifies their productivity processes.

process by using resources provided by a risk-averse household. Together, these decisions shape the risk profile of the aggregate economy.

2.1 Firms and production functions

There are N firms, indexed by $i \in \{1, \ldots, N\}$, each producing a differentiated good.[3] Firm i has access to the constant returns to scale Cobb-Douglas production function

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

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.[4] The matrix α, containing the input shares $0 < \alpha_{ij} < 1$, describes the production network in this economy.[5]

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

$$a_i(\varepsilon, \delta_i) = \delta_i^{\top} \varepsilon, \tag{2}$$

such that δ_{im} determines i's exposure to risk m. We do not restrict δ_i to be positive so that firms can be negatively exposed to a shock. We further assume that $\varepsilon \sim \mathcal{N}(\mu, \Sigma)$ is normally distributed. The vector μ captures the expected level of the risk factors and the positive definite covariance matrix Σ determines the uncertainty about individual elements of ε and how they covary.

By choosing δ_i, firm i implicitly selects the mean and the variance of a_i, as well as how it correlates with the productivity of other firms. As we will see, these correlations, in turn, create aggregate risk, so that firm decisions matter for the risk profile of the aggregate economy. For instance, if many firms decide to locate their operations in the same region, a localized earthquake might trigger an economic disaster (Barro, 2006; Carvalho et al., 2021).

Our specification of the sources of risk 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

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

[4] To simplify the unit cost expression, given by (9) below, we set $\zeta_i^{-1} = \left(1 - \sum_{j=1}^{N} \alpha_{ij}\right)^{1-\sum_{j=1}^{N} \alpha_{ij}} \prod_{j=1}^{n} \alpha_{ij}^{\alpha_{ij}}$.

[5] 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.

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.[6]

Firms must use risk management resources to adjust their exposure δ_i. In reality, these resources can take many forms. A firm might need to rent a piece of land away from the shore to reduce flood risk and vulnerability to global warming. 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 different resources are provided by the representative household at a cost in terms of utility. Different resources can have different supply elasticities, so that some, like land, might be hard to adjust. We show in Appendix E.2 that, under some conditions, these different resources can be aggregated into a single composite resource. For simplicity, we therefore write the model directly in terms of that composite risk management resource.

To achieve risk exposure δ_i, we assume that firm i requires $R_i = \kappa_i(\delta_i)$ units of risk management resources. We further let $\kappa_i(\delta_i)$ take the quadratic form,

$$\kappa_i(\delta_i) = \frac{1}{2}(\delta_i - \delta_i^\circ)^\top H_i(\delta_i - \delta_i^\circ), \tag{3}$$

where H_i is a positive-definite $M \times M$ matrix. Since firm i can achieve risk exposure $\delta_i = \delta_i^\circ$ without using any resources, we refer to δ_i° as firm i's *natural exposure* to the risk factors.[7] This specification is sufficiently rich to allow for both aggregate and firm-specific sources of risk. Indeed, there can be a risk factor 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. Similarly, the model can accommodate risks that cannot be mitigated by assuming that κ is sufficiently steep. Finally, the model can also accommodate a constant firm-specific TFP shifter by assuming that there exists a degenerate risk factor m with $\Sigma_{mm} = 0$ and that deviating from a fixed exposure to m is arbitrarily costly.

Events in this economy unfold in two stages. Before ε is realized, firms choose how much risk management resources to purchase and make their exposure decisions δ. The household also decides how many risk management resources to supply at that time. After ε is drawn, consumption, labor and intermediate inputs are chosen and their respective markets clear. This timing reflects the facts that risk management decisions (where to locate a plant, which technology to adopt, etc.) are long-lived and cannot be changed rapidly once the shocks are realized.

[6] Ideally, we would let firms pick their productivity distributions in an arbitrary way, but the fact that productivities are correlated across firms creates some issues. For instance, it is not clear whether firm i or j should 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.

[7] As we show in Appendix E.3, most of our results hold in some form without assuming that κ_i is quadratic, but the exploration of the model is much more complicated in that case and does not yield important additional insights.

2.2 Household preferences

A risk-averse representative household supplies one unit of production workers inelastically and a variable amount R of risk management resources.[8] It also owns the firms. The household values consuming a bundle $Y = \prod_{i=1}^{N} \left(\beta_i^{-1} C_i\right)^{\beta_i}$ of the different consumption goods. We impose that $\sum_{i=1}^{N} \beta_i = 1$, 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[9]

$$\mathcal{U}(Y)\,\mathcal{V}(R), \tag{4}$$

where $\mathcal{U}(Y) = \frac{Y^{1-\rho}}{1-\rho}$ is CRRA with relative risk aversion parameter $\rho \geq 1$, and $\mathcal{V}(R) = V(R)^{1-\rho}$ with $V(R) = \exp(-\eta R)$.[10] The elasticity $\eta > 0$ controls how costly it is for the household to supply additional risk management resources. The exponential form implies that the price of risk management resources is a constant fraction of nominal GDP. We adopt it for tractability but show in Appendix E.4 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

$$\sum_{i=1}^{N} P_i C_i \leq W_L + W_R R + \Pi, \tag{5}$$

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 $W_L = 1$.[11] We further define $\overline{P} = \prod_{i=1}^{N} P_i^{\beta_i}$ as the consumption price index.

There is a complete set of state-ε contingent claims in this economy. Since the household can trade these claims, their prices reflect the marginal utility of consumption in each state. As usual,

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

[9] 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. Several of our results also apply more generally, as we show in Appendix E.4.

[10] 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.

[11] 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 E.5 that we can normalize $W_L = 1$ in each state of the world.

this implies that firms use the household's stochastic discount factor,

$$\Lambda = \frac{1}{\overline{P}} \frac{d\left[\mathcal{U}\left(Y\right)\mathcal{V}\left(R\right)\right]}{dY}, \tag{6}$$

to discount profits in different states of the world (see Appendix A.1 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

$$-\bar{P}Y\frac{V'\left(R\right)}{V\left(R\right)} = W_R. \tag{7}$$

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

2.3 Firm problem

Once ε is realized, firms must choose how much labor to hire and how many intermediate inputs to purchase, given the risk-exposure vector δ_i that they previously chose. In that case, firm i's cost minimization problem is

$$K_i\left(\delta_i, P\right) = \min_{L_i, X_i}\left(L_i + \sum_{j=1}^{N} P_j X_{ij}\right), \qquad \text{subject to } F\left(\delta_i, L_i, X_i\right) \geq 1, \tag{8}$$

where $P = \left(P_1, \ldots, P_n\right)$ is the vector of prices. Since, for a fixed δ_i, firm i operates a constant returns to scale Cobb-Douglas technology, its unit cost of production is

$$K_i\left(\delta_i, P\right) = \frac{1}{e^{a_i\left(\varepsilon,\delta_i\right)}}\prod_{j=1}^{N} P_j^{\alpha_{ij}}. \tag{9}$$

The cost of producing one unit of good i is therefore equal to the geometric average of the individual input prices, weighted by their respective shares, and adjusted for total factor productivity, which depends on the firm's risk-taking decision.

Given (9), the first stage of the firm's problem involves choosing risk exposure δ_i to maximize expected discounted profits

$$\delta_i^* \in \arg\max_{\delta_i} \mathrm{E}\left[\Lambda\left[P_i Q_i - K_i\left(\delta_i, P\right)Q_i - \kappa_i\left(\delta_i\right)W_R\right]\right], \tag{10}$$

where Q_i is 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 cashflows 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.

Since firms behave competitively, they take as given the equilibrium objects P, Q_i, W_R and Λ when solving (10).[12] We can therefore rewrite that equation as

$$\delta_i^* \in \arg\min_{\delta_i} \mathrm{E}\left[K_i\left(\delta_i, P\right) Q_i\right] + \mathrm{Cov}\left(K_i\left(\delta_i, P\right) Q_i, \frac{\Lambda}{\mathrm{E}\left[\Lambda\right]}\right) + \kappa_i\left(\delta_i\right) W_R. \tag{11}$$

Naturally, when firm i chooses its risk exposure δ_i, it seeks to minimize the expected cost of goods sold $\mathrm{E}\left[K_i\left(\delta_i, P\right) Q_i\right]$ and risk management expenses $\kappa_i\left(\delta_i\right) W_R$. But, importantly, it also cares about the comovement of its cost with the stochastic discount factor. Intuitively, the firm prefers to have a low cost in states with high marginal utility of consumption or, equivalently, with low GDP. This creates a motive for diversification across firms. As we will see in the next section, GDP is mostly driven by the risk factors to which firms are exposed the most. As a result, any individual firm has incentives to move against the crowd and to gain exposure to less popular factors, or to those that are negatively correlated with GDP.

Finally, as $\mathrm{Cov}\left(K_i Q_i, \Lambda\right) = \mathrm{Corr}\left(K_i Q_i, \Lambda\right) \sqrt{\mathrm{V}\left[K_i Q_i\right]} \sqrt{\mathrm{V}\left[\Lambda\right]}$, the variance of the cost of goods sold $\mathrm{V}\left[K_i Q_i\right]$ also matters. Since GDP and TFP tend to move hand-in-hand, firm costs and GDP tend to move in opposite directions, so that $\mathrm{Corr}\left(K_i Q_i, \Lambda\right)$ is positive for the typical firm. In this case, the optimal risk exposure δ_i seeks to minimize $\mathrm{V}\left[K_i Q_i\right]$ as well. 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.

2.4 Equilibrium conditions

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

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

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 (9) and (12) allow us to fully characterize the price system as a function of the random productivity shocks ε.[13]

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

[12] 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.

[13] Due to constant returns to scale, risk management expenses lead to negative profit 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.

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 (10) given prices P^*, the price for risk management resources W_R^*, demand Q_i^*, and the stochastic discount factor Λ^* given by (6).

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 (8) 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 (4) subject to (5) given prices P^* and W_R^*.

4. (Unit cost pricing) For each $i \in \{1, \ldots, N\}$, P_i^* solves (12) where $K_i\left(\alpha_i^*, P^*\right)$ is given by (9).

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

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

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. These decisions, once aggregated, shape the risk structure of the macroeconomy.

We have kept the model parsimonious to make the exposition more 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 κ_i is 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 E. 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.

3 Equilibrium prices and GDP

Before describing the firms' risk-exposure decisions, we first characterize how prices and GDP behave in this economy. For that purpose, we follow Baqaee and Farhi (2019a) and 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 $\Omega_{ij} = \frac{\alpha_{ij}}{1+\tau_i}$. 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

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

While Ω_{ij} and $\tilde{\Omega}_{ij}$ capture the direct exposure of firm i to firm j as a share of its revenue and cost 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 share of its sales in nominal GDP, such that $\omega_i := \frac{P_i Q_i}{PY}$. 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: $\tilde{\omega}^\top := \beta^\top \tilde{\mathcal{L}} = \beta^\top (I - \alpha)^{-1}$.

Using the relation between Ω and $\tilde{\Omega}$, the revenue-based Domar weights can also be written in terms of primitives as

$$\omega^\top = \beta^\top \left(I - \left[\operatorname{diag}(1+\tau)\right]^{-1} \alpha\right)^{-1}, \tag{14}$$

where $\operatorname{diag}(1+\tau)$ is the diagonal matrix whose ith diagonal element is $1 + \tau_i$. One can show that ω is weakly decreasing in τ_j, such that higher wedges reduce revenue-based Domar weights, but leave their cost-based counterparts unchanged. Neither cost-based nor revenue-based Domar weights depend on risk-exposure decisions. Without input-output linkages ($\alpha = 0$), both Domar weight vectors are equal to β.

Combining these definitions with (9) and (12), we can write the vector of log prices as

$$p = -\tilde{\mathcal{L}} \left(\delta\varepsilon - \log(1+\tau)\right), \tag{15}$$

where $p = (\log P_1, \ldots, \log P_N)$, $\log(1+\tau) = (\log(1+\tau_1), \ldots, \log(1+\tau_N))$, 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 firm i is high if its productivity $a_i(\varepsilon_i, \delta_i) = \delta_i^\top \varepsilon$ is low, or if the productivity of one of its important suppliers is low, and so on. Equation (15) 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.

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 $\Delta_m = \sum_{i=1}^N \tilde{\omega}_i \delta_{im}$ of that vector 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*

$$y = \Delta^\top \varepsilon - \tilde{\omega}^\top \log(1+\tau) - \log \Gamma_L, \tag{16}$$

where the labor share of income Γ_L *is given by*[14]

$$\Gamma_L := \frac{W_L L}{\bar{P} Y} = 1 - \tau^\top \left(diag(1+\tau)\right)^{-1} \omega. \tag{17}$$

The first term in (16) shows that the contribution of factor ε_m to GDP is proportional to its aggregate exposure 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 (16) are deterministic and reflect the role of wedges. When $\tau = 0$, cost-based and revenue-based Domar weights coincide, such that (16) collapses to $y(\delta) = \Delta^\top \varepsilon = \omega^\top a(\varepsilon, \delta)$, and Hulten's (1978) theorem applies. As wedges grow, GDP is distorted from its efficient level.

Equation (16) implies that log GDP is normally distributed. Its moments are then

$$\mathrm{E}[y] = \Delta^\top \mu - \tilde{\omega}^\top \log(1+\tau) - \log \Gamma_L \qquad \text{and} \qquad \mathrm{V}[y] = \Delta^\top \Sigma \Delta. \tag{18}$$

The first equation shows that a marginal increase in aggregate exposure Δ_m raises expected log GDP by the expected productivity μ_m of that factor. Unsurprisingly, gaining exposure to a high-mean risk factor is beneficial for $\mathrm{E}[y]$. Similarly, an increase in μ_m has a beneficial impact on $\mathrm{E}[y]$ if the economy is positively exposed to it ($\Delta_m > 0$), but reduces expected log GDP otherwise.

The second equation in (18) 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

$$\frac{d\,\mathrm{V}[y]}{d\Delta_m} = 2\,\mathrm{Cov}[y, \varepsilon_m] = 2\sum_{n=1}^{M} \Delta_n \,\mathrm{Cov}[\varepsilon_n, \varepsilon_m]. \tag{19}$$

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 the other

[14] To be precise, Γ_L is the labor share of income gross of risk management expenses.

risk factors, and on the economy's exposure to those factors. If the economy is heavily exposed to some factor ε_n ($\Delta_n > 0$), and that 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.

Finally, (18) 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. We will explore that channel later on.

4 Equilibrium existence, uniqueness and efficiency

In the previous section, we characterized GDP as a function of aggregate risk exposure Δ, but Δ itself is an equilibrium object that is shaped by the decisions of the firms. In this section, we consider the full equilibrium mapping in which Δ is endogenous. We start by introducing the problem of a social planner and by describing the efficient allocation. We then provide a characterization of the equilibrium of the model. We also show that there exists a unique equilibrium and that it is in general inefficient.

4.1 A planning problem

Since there is only one household in this economy, the social planner simply maximizes the expected utility of that household subject to the resource constraints (13). To describe that problem, it is convenient to introduce the *aggregate cost function* $\bar{\kappa}_{SP}$, defined as the value function

$$\bar{\kappa}_{SP}(\Delta) := \min_{\delta} -\log V\left(\sum_{i=1}^{N} \kappa_i(\delta_i)\right), \qquad \text{subject to } \Delta = \delta^\top \tilde{\omega}. \tag{20}$$

For a given aggregate risk-exposure vector Δ, $\bar{\kappa}_{SP}(\Delta)$ corresponds to the smallest possible (log) utility cost required to achieve Δ. The minimization problem (20) also implicitly defines a function $\delta_{SP}(\Delta)$ that provides the individual risk-exposure matrix δ that minimizes the welfare cost of achieving Δ.

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

$$\mathcal{W}_{SP} := \max_{\Delta} \underbrace{\Delta^\top \mu}_{\mathrm{E}[y_{SP}]} - \frac{1}{2}(\rho - 1) \underbrace{\Delta^\top \Sigma \Delta}_{\mathrm{V}[y_{SP}]} - \bar{\kappa}_{SP}(\Delta), \tag{21}$$

where $\mathcal{W}_{SP}$ provides a measure of (log) welfare in the efficient allocation. This problem shows that the planner seeks to maximize the expected value of log GDP $\mathrm{E}[y_{SP}]$, while minimizing its volatility $\mathrm{V}[y_{SP}]$ and the risk management cost $\bar{\kappa}_{SP}$. Unsurprisingly, the importance of the variance term

increases with the risk aversion ρ of the household.[15] Since the objective function (21) 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 associated with this risk exposure.

4.2 A distorted planning problem

Wedges imply that the equilibrium allocation differs from the efficient allocation. To characterize the inefficient equilibrium, we rely on a distorted version of the planning problem (21) that accounts for these distortions. For that purpose, we introduce a distorted version of $\bar{\kappa}_{SP}$, defined as

$$\bar{\kappa}(\Delta) := \min_{\delta} -\log V\left(\sum_{i=1}^{N} g_i \kappa_i(\delta_i)\right), \qquad \text{subject to } \Delta = \delta^\top \tilde{\omega}, \tag{22}$$

and where $g_i := (1+\tau_i)\frac{\tilde{\omega}_i}{\omega_i} \geq 1$. We refer to g_i as the *efficiency gap* of firm i. As it multiplies κ_i in (22), a higher g_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, $\bar{\kappa}(\Delta) = \bar{\kappa}_{SP}(\Delta)$ for all Δ. When τ increases, however, so do the efficiency gaps. Specifically, we can show that

$$\frac{dg_i}{d\tau_j} = \frac{g_i}{\omega_i}\frac{\omega_j}{1+\tau_j}\mathcal{L}_{ji} > 0. \tag{23}$$

An increase in τ_j reduces firm j's size, as measured by its revenue-based Domar weight, and so its demand for inputs from its suppliers decreases as well. It follows that if firm i is a direct or indirect supplier to j ($\mathcal{L}_{ji} > 0$), its revenue-based Domar weight also shrinks, and its efficiency gap g_i increases. Without input-output linkages ($\alpha = 0$), the efficiency gaps take a particularly simple form, $g_i = 1 + \tau_i$, and a change in wedge τ_i only affects the efficiency gap of firm i.

The function $\bar{\kappa}$ will play an important role in our analysis, and so it is useful to derive some of its properties right away.

Lemma 2. *The aggregate cost function $\bar{\kappa}$ is given by*

$$\bar{\kappa}(\Delta) = \frac{1}{2}(\Delta - \Delta^\circ)^\top \nabla^2 \bar{\kappa}(\Delta - \Delta^\circ), \tag{24}$$

[15] Notice that (21) does not depend directly on the individual risk-exposure decisions δ. Given the definition of $\bar{\kappa}_{SP}(\Delta)$, the impact of δ on welfare operates only through the aggregate risk exposure Δ.

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

$$\nabla^2 \bar{\kappa} = \eta \left(\sum_{i=1}^{N} \frac{\tilde{\omega}_i^2}{g_i} H_i^{-1} \right)^{-1}. \tag{25}$$

This result shows that the distorted utility cost $\bar{\kappa}(\Delta)$ 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 $\bar{\kappa}(\Delta^\circ) = 0$. Whenever Δ departs from Δ°, the utility cost increases by an amount that depends on the curvature $\nabla^2 \bar{\kappa}$ of $\bar{\kappa}$. As (25) shows, this curvature is a weighted harmonic average of the underlying individual curvature matrices $\nabla^2 \kappa_i = H_i$. Intuitively, if all firms must pay a large cost 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. Having firms with higher cost-based Domar weights imply that a given change in $\Delta = \delta^\top \tilde{\omega}$ can be achieved via smaller, and therefore less costly, movements in δ, and so $\partial \bar{\kappa} / \partial \tilde{\omega}_i < 0$. On the other hand, allocating risk management resources to highly distorted firms (large g_i) is costly, and $\partial \bar{\kappa} / \partial \tau_i > 0$.

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

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

$$\mathcal{W}_{dist} := \max_{\Delta} \underbrace{\Delta^\top \mu - \tilde{\omega}^\top \log(1+\tau) - \log \Gamma_L}_{\mathrm{E}[y]} - \frac{1}{2}(\rho - 1) \underbrace{\Delta^\top \Sigma \Delta}_{\mathrm{V}[y]} - \bar{\kappa}(\Delta). \tag{26}$$

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 (26) by distorting expected log GDP $\mathrm{E}[y]$ and by increasing the cost of risk exposure embedded in $\bar{\kappa}$. Since, however, Δ does not interact with the wedges in $\mathrm{E}[y]$, only the distortions in $\bar{\kappa}$ affect the equilibrium exposure decisions.

Why are distortions in $\bar{\kappa}$ needed to replicate the equilibrium? From the planner's perspective, the impact on GDP of firm i's risk exposure δ_i is proportional to its cost-based Domar weight $\tilde{\omega}_i$. From the equilibrium's perspective, however, the firm makes decisions to minimize its production cost $K_i Q_i$, which, using the definition of revenue-based Domar weights, is proportional to $\omega_i / (1 + \tau_i)$. The effective cost of risk exposure must therefore be adjusted by the ratio of these quantities, the efficiency gap g_i, for problem (26) to yield the equilibrium allocation. When wedges vanish, the two functions $\bar{\kappa}$ and $\bar{\kappa}_{SP}$ become identical and the equilibrium coincides with the efficient allocation. The wedges τ are therefore the only source of inefficiency in this model.

5 Forces shaping risk-exposure decisions

In this section, we solve the model and provide expressions for the equilibrium risk-exposure decisions. We also explore how primitives of the economy, such as the moments of the risk factors and the wedges, affect those decisions.

Our analysis relies on the optimality conditions of the fictitious planner's problem.

Lemma 3. *The equilibrium aggregate risk exposure* Δ *solves*

$$\mathcal{E}(\Delta) = \nabla\bar{\kappa}(\Delta), \tag{27}$$

where the marginal value of aggregate risk exposure $\mathcal{E}$ *is given by*

$$\mathcal{E} := \mathrm{E}[\varepsilon] + \mathrm{Cov}[\lambda, \varepsilon], \tag{28}$$

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

This lemma follows directly from the first-order condition of the fictitious planner. It states that the marginal utility cost $\nabla\bar{\kappa}$ of aggregate exposure is equal to its marginal benefit to the household's expected utility of consumption, which is given by the vector $\mathcal{E}$. Unsurprisingly, the expression for $\mathcal{E}$ implies that exposure to factors with high expected value $\mathrm{E}[\varepsilon_m] = \mu_m$ is particularly valuable. Moreover, factors that are highly correlated with the stochastic discount factor λ also have high exposure value. These factors perform well when the household is poor, and they therefore provide insurance in high marginal utility states. Given $\mathcal{E}$'s definition, we say that a risk factor ε_m is *good* if $\mathcal{E}_m > 0$, and that it is *bad* if $\mathcal{E}_m < 0$.

Using the definition of the stochastic discount factor (6), we can simplify the expression for $\mathcal{E}$.

Lemma 4. *The equilibrium marginal value of exposure vector* $\mathcal{E}$ *can be written as*

$$\mathcal{E} = \mu - (\rho - 1)\Sigma\Delta. \tag{29}$$

To understand this equation, consider first an economy in which there is 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 1. That figure also shows the quadratic relationship between Δ and $\mathrm{V}[y]$ implied by (18). It is clear from this figure that if the economy is already heavily 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 high-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 (29) and Figure 1 show. When there is more than one risk factor, correlations matter as well. If

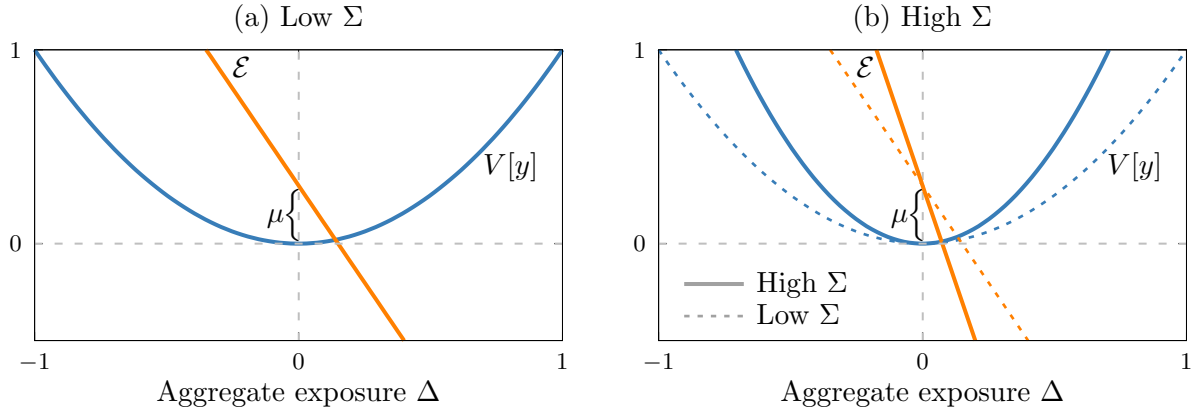
the economy is heavily exposed to factors that covary with ε_m, any increase in Δ_m would raise aggregate volatility. In that case, the marginal value of exposure to ε_m would be low.

The moments of a risk factor also influence $\mathcal{E}$. Let $\mathbf{1}_m$ be the mth standard basis vector, then[16]

$$\frac{\partial \mathcal{E}}{\partial \mu_m} = \mathbf{1}_m \qquad \text{and} \qquad \frac{\partial \mathcal{E}}{\partial \Sigma_{mn}} = -\frac{1}{2}(\rho - 1)(\Delta_m \mathbf{1}_n + \Delta_n \mathbf{1}_m). \tag{30}$$

Therefore, an increase in μ_m always makes factor m more attractive but leaves the marginal value of exposure to all other factors unchanged. The second equation in (30) describes how changes in the covariance matrix Σ affect the attractiveness of the risk factors. To understand the intuition behind it, consider the right panel of Figure 1 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 1 and equation (30) show. When there is more than one factor, a similar reasoning applies to changes in covariances. 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 (30) also shows that a higher risk aversion ρ magnifies the impact of Σ on $\mathcal{E}$.

Figure 1: Impact of Δ on $\mathrm{V}[y]$ and $\mathcal{E}$ in an economy with a single risk factor



5.1 Aggregate and individual risk exposure

From (27), we can solve for the equilibrium exposure decisions in closed-form.

[16] 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.

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

$$\delta_i = \delta_i^\circ + \frac{1}{\eta}\frac{\omega_i}{1+\tau_i} H_i^{-1}\mathcal{E} \qquad \textit{and} \qquad \Delta = \Delta^\circ + \mathcal{H}^{-1}\mathcal{E}^\circ, \tag{31}$$

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

$$\mathcal{H}^{-1} := \left(\nabla^2 \bar{\kappa} + (\rho - 1)\,\Sigma\right)^{-1}. \tag{32}$$

The first part of (31) 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 δ_i°. The size of these deviations depends on H_i^{-1}, which captures the elasticity of the cost function κ_i. Since H_i is positive definite, H_i^{-1} has a positive diagonal, and 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 the risk management of firm i. 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.

These substitution patterns can be used to describe physical constraints that a firm might face when managing risk. For instance, suppose that a firm can locate its plant in a single location: by the shore, where floods can happen, or in a plain, where droughts 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.

The size of the firm also matters for exposure decisions. Intuitively, since the impact of productivity scales with the input bundle, the benefit of managing risk increases with the cost of goods sold $K_i Q_i$. In contrast, its cost $\kappa_i(\delta_i) W_R$ does not. It follows that large firms find it more cost-efficient to spend on risk management. Furthermore, since $K_i Q_i \propto \omega_i/(1+\tau_i)$, firms with high revenue-based Domar weights and those with low markups have greater incentives to manage risk, as (31) shows. More concretely, we can show that

$$\frac{\partial \mathcal{E}^\top \delta_i}{\partial \omega_i} > 0 \qquad \text{and} \qquad \frac{\partial \mathcal{E}^\top \delta_i}{\partial \tau_i} < 0,$$

such that higher revenue-based Domar weight firms and those with lower markups tend to be more exposed to good risk factors ($\mathcal{E} > 0$) and less exposed to bad ones. This relationship between size and risk-taking behavior plays an important role in the mechanisms of the model.[17] In Section 7

[17] In Acemoglu and Zilibotti (1997) and Koren and Tenreyro (2013), incentives to diversify scale with size. In our model, large firms spend more on risk management than small firms, but these expenditures as a share of sales

we will see that it helps explain several features of the data.[18]

The supply elasticity of risk-management resources η also matters for risk-taking decisions. When η is large, the economy struggles to provide more resources when needed, and so δ_i tends to stay close to δ_i°. In contrast, when η is small, large adjustments can be easily achieved.

The second part of (31) provides a closed-form expression for aggregate risk exposure Δ. We see 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 (32). Since $\mathcal{H}^{-1}$ is positive definite, if a risk factor m is naturally attractive ($\mathcal{E}_m^\circ > 0$), firms and the aggregate economy tend to be 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 the matrix $\mathcal{H}^{-1}$. 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 $\mathcal{E}_m^\circ$ leads, all else equal, to an increase in Δ_n. If, instead, $\left[\mathcal{H}^{-1}\right]_{mn} < 0$, we say that m and n are *global substitutes*, in which case an increase in $\mathcal{E}_m^\circ$ leads to a decline in Δ_n.

Equation (32) implies that these global substitution patterns depend on the Hessian matrix $\nabla^2\bar{\kappa}$ of the aggregate cost function $\bar{\kappa}$. An element (m, n) of $\nabla^2\bar{\kappa}$ captures how an increase in exposure to factor m changes the *marginal* cost of gaining exposure to factor n. Recall from (25) that $\nabla^2\bar{\kappa}$ an average of the Hessians of the underlying individual cost functions $(\kappa_1, \ldots, \kappa_N)$. It follows that $\nabla^2\bar{\kappa}$ captures the global impact of the local substitution patterns embedded in $(\kappa_1, \ldots, \kappa_N)$. The global substitution patterns also depend on the response of the value of exposure to a change in Δ, which is given by $d\mathcal{E}/d\Delta = -(\rho - 1)\Sigma$. 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.

5.2 Beliefs and aggregate risk exposure

It is clear from the definition of $\mathcal{E}$ that the moments (μ, Σ) affect how attractive the risk factors are. The following result describes their impact on risk-taking decisions.

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*

$$\frac{d\Delta}{d\gamma} = \mathcal{H}^{-1}\frac{\partial \mathcal{E}}{\partial \gamma}, \tag{33}$$

where $\frac{\partial \mathcal{E}}{\partial \gamma}$ is given by (30).

decrease with firm size. In Appendix B.6, we provide evidence consistent with these patterns.

[18] When calibrating the model in Section 8 we allow H_i to depend on $\tilde{\omega}_i$ to allow more flexibility between risk management costs and firm size.

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 $\partial\mathcal{E}/\partial\gamma$ in (33). 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 1. *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 $\Delta_m < 0$.*

An increase in μ_m makes exposure to ε_m more attractive in terms of its contribution to the household's expected consumption utility without affecting the cost $\bar{\kappa}$ of achieving that exposure. Through the maximization problem of the fictitious planner (26), those changes lead to a higher Δ_m. The impact of Σ is more subtle. Recall from (18) that the sensitivity of the variance of GDP 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.

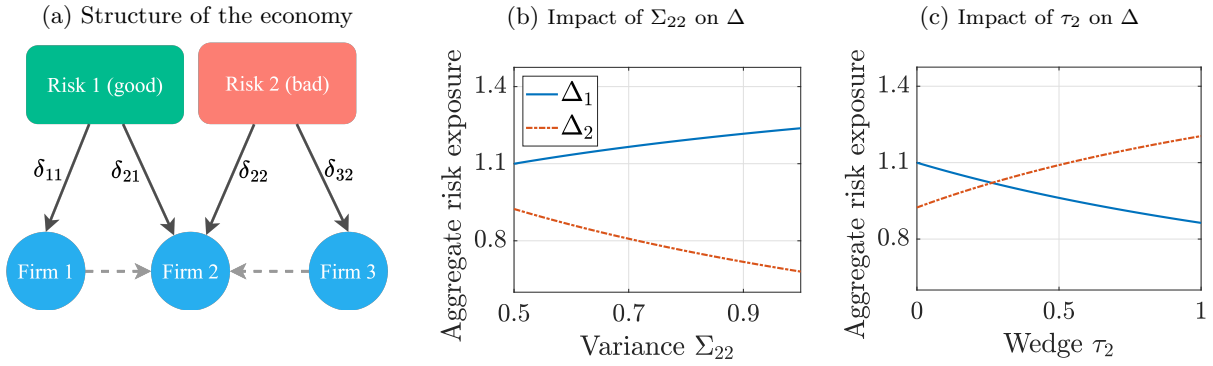
As in Proposition 2, the global substitution matrix $\mathcal{H}^{-1}$ controls how changes in the attractiveness of a risk factor ε_m affect exposure to another factor, say ε_n. To better understand how these substitution patterns work, consider a simple economy in which three firms can be exposed to two sources of risk, as in panel (a) of Figure 2. Risk 1 is a good risk with $\mathcal{E}_1 > 0$, while risk 2 is bad with $\mathcal{E}_2 < 0$. The cost functions κ are parametrized so that firm 1 is only exposed to the good risk and cannot adjust its exposure. Similarly, firm 3 is only exposed to the bad risk. Firm 2, in contrast, is free to change its exposure to the two risk factors, and its cost function κ_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 2 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 1 cannot move from region 1 and is thus only exposed to risk 1, while firm 3 cannot move from region 2 and is therefore exposed to risk 2 only. Firm 2, in contrast, can locate its plants in both regions. 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 2 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 2: 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: $\delta^\circ_{11} = \delta^\circ_{32} = 1$, $\delta^\circ_{22} = 1.9$, $\delta^\circ_{12} = \delta^\circ_{21} = \delta^\circ_{31} = 0$, $H_1 = H_3$ are diagonal with very large entries on the main diagonals; $H_{2,11} = H_{2,22} = 1$, $H_{2,12} = H_{2,21} = 0.75$. Wedges: $\tau = 0$. In panel (b), Σ_{22} changes from 0.5 to 1. In panel (c), τ_2 changes from 0 to 1.

5.3 Wedges and aggregate risk exposure

Wedges affect the size of firms and, through that channel, risk-exposure decisions.

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

$$\frac{d\Delta}{d\tau_i} = \frac{d\mathcal{H}^{-1}}{d\tau_i}\mathcal{E}^\circ, \tag{34}$$

where

$$\frac{d\mathcal{H}^{-1}}{d\tau_i} = -\mathcal{H}^{-1}\left(\sum_{j=1}^{N}\frac{\partial\nabla^2\bar{\kappa}}{\partial g_j}\frac{dg_j}{d\tau_i}\right)\mathcal{H}^{-1},$$

and where the response of the efficiency gap $dg_j/d\tau_i$ is given by (23), and $\partial\nabla^2\bar{\kappa}/\partial g_j$ is a positive definite matrix.[19]

Wedges affect the equilibrium risk exposure Δ through the curvature of the aggregate cost function $\bar{\kappa}$. Specifically, from (23), an increase in τ_i leads to a higher efficiency gap g_j for all j. This, in turn, means that firms find risk management less appealing, and the curvature of the aggregate cost function increases, in the sense that $\partial\nabla^2\bar{\kappa}/\partial g_j$ is a positive definite matrix. These changes in curvature then affect $\mathcal{H}^{-1}$ to influence Δ.

[19] We give an expression for $\partial\nabla^2\bar{\kappa}/\partial g_j$ in the proof of this proposition.

Proposition 4 provides a general characterization of the impact of τ, but we can get a sharper result under a particular parametrization of the covariance matrix and of the individual cost functions.

Definition 2. An economy is *diagonal* if the risk factors are uncorrelated (diagonal Σ), and the individual risk exposures are neither complements nor substitutes in the individual cost functions $(\kappa_1, \ldots, \kappa_N)$ (diagonal H_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. For instance, the following result describes the impact of τ in such a setting.

Corollary 2. *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 become less cost-effective: since the firm is smaller, managing its TFP is now less rewarding. As a result, firms spend less to get 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.

The last panel of Figure 2 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.

So far, we have characterized how risk-taking decisions are affected by the moments (μ, Σ) of the risk factors and the wedges τ. It is, however, straightforward from (30) to describe how any parameter affects Δ. For instance, in Appendix E.1 we also explore how changes in the production network α affect Δ. Our main finding there is that if a change in α_{ij} leads to an increase in revenue-based Domar weights, the affected firms manage risk more aggressively in response.

5.4 Equilibrium and efficient risk exposure

Proposition 4 implies that the distortions matter for risk-taking decisions. The following lemma shows that as a result the equilibrium Δ departs from its efficient level.

Lemma 5. *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.*

When there is a single risk factor, Lemma 5 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, Lemma 5 shows that these forces operate on average. For instance, if there are only 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.

In a diagonal economy, this over/under exposure result applies factor by factor, as the next corollary shows.

Corollary 3. *Suppose that the economy is diagonal and that $\tau_j > 0$ for at least one firm j. Then the sign of $\Delta_i - \Delta_{SP,i}$ is the opposite of the sign of $\mathcal{E}_i^\circ$.*

We will see in the next section how over/under exposure to risk factors matters for welfare.

6 GDP and welfare

As we have seen, changes in the environment affect the economy's risk profile. In this section, we study the impact of that mechanism on the moments of GDP and on welfare.

6.1 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 might, in turn, affect GDP. This is the main channel that 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 its impact on Δ.[20]

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*

$$\frac{d\,\mathrm{E}\left[y\right]}{d\chi} - \frac{\partial\,\mathrm{E}\left[y\right]}{\partial\chi} = \mu^\top \frac{d\Delta}{d\chi} \qquad \textit{and} \qquad \frac{d\,\mathrm{V}\left[y\right]}{d\chi} - \frac{\partial\,\mathrm{V}\left[y\right]}{\partial\chi} = 2\Delta^\top \Sigma \frac{d\Delta}{d\chi}, \tag{35}$$

[20] Propositions 5 and 7 also apply when χ denotes a network link α_{ij}. In that case, the term $d\Delta/d\chi$ is given by (83) in Appendix E.1.

where partial derivatives indicate that Δ is kept fixed, and where $\frac{d\Delta}{d\chi}$ is given by (33) for $\chi = \mu_i$ or $\chi = \Sigma_{mn}$, and by (34) for $\chi = \tau_i$.

Proof. The result follows directly from (18). □

The first expression in (35) shows that the response of $\mathrm{E}[y]$ to χ depends on the mean μ of the risk factors whose exposure responds to χ. For instance, recall that an increase in μ_m raises Δ_m (Corollary 1). If $\mu_m > 0$, this additional exposure has a beneficial impact on $\mathrm{E}[y]$. In contrast, expected GDP declines if $\mu_m < 0$. In general, a change in χ affects the economy's exposure to many risk factors, and so the first equation in (35) 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 (35) as

$$\frac{d\,\mathrm{V}[y]}{d\chi} - \frac{\partial\,\mathrm{V}[y]}{\partial\chi} = 2\,\mathrm{Cov}\left[y, \left(\frac{d\Delta}{d\chi}\right)^{\top}\varepsilon\right].$$

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.

Corollary 4. *In a diagonal economy, the following statements hold.*

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

$$sign\left(\frac{d\,\mathrm{E}[y]}{d\mu_m} - \frac{\partial\,\mathrm{E}[y]}{\partial\mu_m}\right) = sign(\mu_m) \quad and \quad sign\left(\frac{d\,\mathrm{V}[y]}{d\mu_m} - \frac{\partial\,\mathrm{V}[y]}{\partial\mu_m}\right) = sign(\Delta_m). \tag{36}$$

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

$$sign\left(\frac{d\,\mathrm{E}[y]}{d\Sigma_{mm}} - \frac{\partial\,\mathrm{E}[y]}{\partial\Sigma_{mm}}\right) = -sign(\mu_m\Delta_m) \quad and \quad \frac{d\,\mathrm{V}[y]}{d\Sigma_{mm}} - \frac{\partial\,\mathrm{V}[y]}{\partial\Sigma_{mm}} < 0. \tag{37}$$

The first result describes how an increase in μ_m affects GDP through its impact on Δ. As explained above, $\mathrm{E}[y]$ rises in response if $\mu_m > 0$ and decreases otherwise. The response of $\mathrm{V}[y]$, in contrast, depends on whether the economy is positively or negatively exposed to m. Suppose that $\Delta_m > 0$. Then 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 1, we know that an increase in Σ_{mm} leads to a decline in $|\Delta_m|$. This implies a beneficial impact on $\mathrm{E}[y]$ if $\mu_m \Delta_m < 0$, but a decline in $\mathrm{E}[y]$ 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. *Suppose that there is a single risk factor. Then*

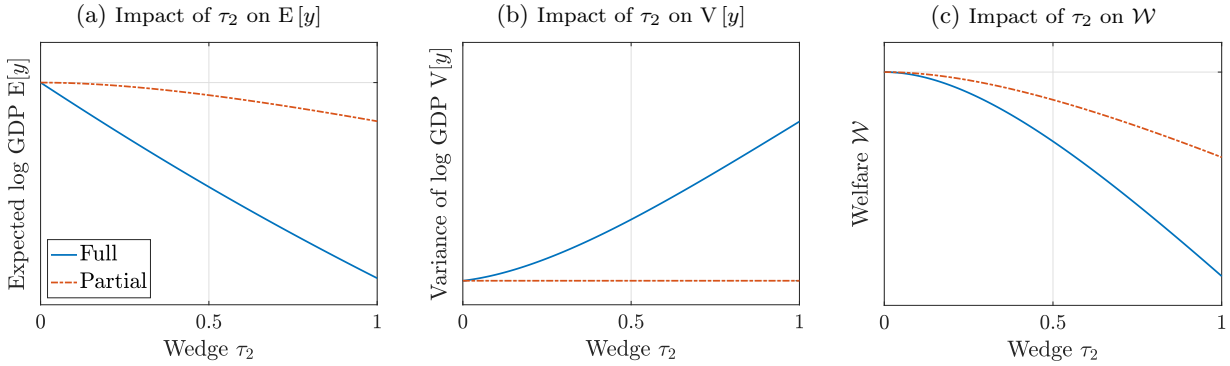
$$sign\left(\frac{d\,\mathrm{E}[y]}{d\tau_i} - \frac{\partial\,\mathrm{E}[y]}{\partial\tau_i}\right) = -sign\left(\mu\mathcal{E}\right) \qquad and \qquad sign\left(\frac{d\,\mathrm{V}[y]}{d\tau_i} - \frac{\partial\,\mathrm{V}[y]}{\partial\tau_i}\right) = -sign\left(\Delta\mathcal{E}\right). \quad (38)$$

To get some intuition for this result, it helps to fix ideas and suppose, for instance, that 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 bad risk exposures (Corollary 2), we know that increasing τ_i makes firms more exposed to the risk factor, that is, Δ increases. 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. As this example shows, the mechanisms of the model therefore provides novel channel through which wedges can increase GDP volatility.

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 2. The first panel of Figure 3 shows the impact of raising τ_2 on $\mathrm{E}[y]$. Recall from the second panel of Figure 2 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 the usual distortionary effects in network economies. 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 3 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 (18)), $\mathrm{V}[y]$ would be unaffected by wedges if Δ remained constant. This is illustrated by the red line in the panel.

Figure 3: 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 2. 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 $\mathrm{diag}(\Sigma) = (0.5, 0.5)$. Risk exposures: $\delta^\circ_{11} = \delta^\circ_{32} = 1$, $\delta^\circ_{22} = 1.9$, $\delta^\circ_{12} = \delta^\circ_{21} = \delta^\circ_{31} = 0$, $H_1 = H_3$ are diagonal with very large entries on the main diagonals; $H_{2,11} = H_{2,22} = 1$, $H_{2,12} = H_{2,21} = 0.75$. Wedges: $\tau = 0$. In all panels, τ_2 changes from 0 to 1. The red dot-dashed lines show how variables change with τ_2 holding $\Delta = \Delta(\tau = 0)$ fixed, and the blue solid lines take into account endogenous adjustments in Δ.

6.2 Welfare

In the previous section, we have characterized how changes in model primitives affect GDP through their impact on risk-exposure decisions. But what the household ultimately cares about is welfare, an object that combines the moments of GDP with the disutility of managing risk. In this section, we evaluate the consequences of risk management decisions for welfare.

We begin by characterizing how welfare reacts to changes in the environment when there are no distortions.

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

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

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.[21] The situation is, however, different in the presence of wedges, in which case the response of risk exposure can have a first-order effect.

To evaluate the response of welfare in a distorted equilibrium, it is helpful to introduce the function $\bar{\kappa}_V(\Delta)$ which captures the equilibrium utility loss associated with a given risk-exposure vector Δ. It is defined as

$$\bar{\kappa}_V(\Delta) := -\log V\left(\sum_{i=1}^{N} \kappa_i\left(\delta_i(\Delta)\right)\right), \tag{39}$$

where $\delta_i(\Delta)$ is the distorted equilibrium exposure given by the first expression in (31). Unlike $\bar{\kappa}$, the weights of the different firms in $\bar{\kappa}_V$ are not themselves distorted by the efficiency gaps $(g_1, \ldots, g_N)$. As a result, $\bar{\kappa}_V$ properly measures the disutility of equilibrium risk management decisions, and we can use it to compute welfare. The function $\bar{\kappa}_V$ also differs from the planner's cost function $\bar{\kappa}_{SP}$ since it uses the (distorted) equilibrium risk-exposure matrix δ as input instead of its efficient counterpart.

We can use $\bar{\kappa}_V$ 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*

$$\frac{d\mathcal{W}}{d\chi} - \frac{\partial\mathcal{W}}{\partial\chi} = \left(\mathcal{E} - \nabla\bar{\kappa}_V\right)^\top \frac{d\Delta}{d\chi} = \left(\nabla\bar{\kappa} - \nabla\bar{\kappa}_V\right)^\top \frac{d\Delta}{d\chi}, \tag{40}$$

where partial derivatives indicate that Δ is kept fixed, and where $d\Delta/d\chi$ is given by (33) for $\chi = \mu_i$ or Σ_{mn}, and by (34) for $\chi = \tau_i$.

The first equality in (40) 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 $\nabla\bar{\kappa}_V$. 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.

[21] 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, 2019a; Bigio and La'O, 2020).

The second equality in (40) 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, $\nabla\bar{\kappa}$. When there are distortions, this perceived cost deviates from the *true* marginal cost of exposure $\nabla\bar{\kappa}_V$, and welfare can be affected by Δ as a result. The gap $\nabla\bar{\kappa} - \nabla\bar{\kappa}_V$ 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 $\nabla\bar{\kappa}$ and $\nabla\bar{\kappa}_V$. 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 6. *Suppose that the economy is diagonal and that $\tau_j > 0$ for at least one firm j. Then the sign of $[\nabla\bar{\kappa}]_i - [\nabla\bar{\kappa}_V]_i$ is the same as the sign of $\mathcal{E}_i$.*

Since in equilibrium we must have $\mathcal{E} = \nabla\bar{\kappa}$, the gradient $[\nabla\bar{\kappa}]_i$ 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 $\bar{\kappa}_V$ is flatter than $\bar{\kappa}$. It follows that for a good risk factor, the true marginal cost is lower than the perceived one, and $[\nabla\bar{\kappa}]_i - [\nabla\bar{\kappa}_V]_i > 0$. The opposite is true for a bad risk factor.

We can combine this insight with Proposition 7 to characterize the impact of (μ, Σ) on welfare in a diagonal economy.

Corollary 6. *Suppose that the economy is diagonal and that $\tau_j > 0$ for at least one firm j. Then the following holds.*

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

From the previous discussion, we know that the distortions imply that for a good risk factor ($\mathcal{E}_m > 0$), the marginal benefit $\mathcal{E}_m$ of increasing Δ_m is larger than the *true* utility cost $\nabla\bar{\kappa}_V$ of that increase. It follows that higher exposure to that risk factor is beneficial. This explains why an increase in μ_m leads to an increase in welfare, on top of the fixed-exposure effect, if $\mathcal{E}_m > 0$. In contrast, if ε_m is a bad risk factor, the marginal benefit of exposure is lower than the true marginal cost, and the same increase in μ_m is detrimental to welfare, as the corollary shows. A similar result holds for an increase in Σ_{mm}. Suppose that ε_m is good, so that $\mathcal{E}_m - [\nabla\bar{\kappa}_V]_m > 0$ by Lemma 6, and higher exposure Δ_m would be beneficial. After an increase in Σ_{mm}, the fictitious planner reduces $|\Delta_m|$ to limit how vulnerable the economy is 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 more exposure Δ_m, which improves welfare.

A similar result holds for the impact of wedges.

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

an increase in wedges is more detrimental to welfare when risk-exposure decisions can adjust, that is, $\frac{d\mathcal{W}}{d\tau_i} \leq \frac{\partial \mathcal{W}}{\partial \tau_i}$.

This result shows that endogenous risk-exposure decisions by firms increase the welfare cost of wedges. Indeed, recall from Corollary 2 that a higher τ_i leads to an increase in exposure to bad risks and to a decline in exposure to good risks. Lemma 6, in turn, implies that additional exposure to bad risks is detrimental to welfare, and vice-versa for good risks, which explains the result.

The last panel of Figure 3 illustrates how these economic forces operate in the model economy of Figure 2. 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 2).

Overall, the model captures a novel channel through which distortions like taxes and markups are detrimental to welfare. Beyond their conventional distortionary effects, these wedges interfere with firms' risk-taking decisions, disrupting the macroeconomy's risk profile and creating larger welfare losses. This highlights the importance of accounting for endogenous risk taking when assessing the welfare implications of taxation and other policy interventions.

7 Reduced-form evidence

In the model, the characteristics of a firm affect its risk-exposure choices and, hence, its TFP's variance and covariance with GDP. In this section, we verify that some of the key predictions of the model are visible in the data. We focus on business cycle risk and, to simplify the exposition, assume that there is a unique aggregate risk factor $\varepsilon_t \sim \text{iid } \mathcal{N}(0, \Sigma)$. Specifically, we assume that firm i's TFP is

$$\log TFP_{it} = \delta_{it}\varepsilon_t + \gamma_i t + v_{it}, \tag{41}$$

where γ_i is a firm-specific deterministic trend to capture long-run changes in TFP, and where $v_{it} \sim \text{iid } \mathcal{N}(\mu_i^v, \Sigma_i^v)$ is a firm-specific shock to capture variation that might be uncorrelated with GDP.[22]

Given this setup, it is straightforward to see that firm i's risk-exposure decision is time-independent, such that $\delta_{it} = \delta_i$.[23] Proposition 2 further implies that

$$\delta_i = \delta_i^\circ + \frac{1}{\eta}\frac{\omega_i}{1+\tau_i} H_i^{-1}\mathcal{E}, \tag{42}$$

[22] For simplicity, we assume that firms cannot adjust their exposure to v_{it} or to the trend γ_i. The analysis is similar if they can, but requires additional assumptions on the cost function κ_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.

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

where $\mathcal{E} = -(\rho - 1)\Sigma\Delta$ since $\mathrm{E}[\varepsilon_t] = 0$. It is natural to think that the aggregate economy is positively exposed to the aggregate risk such that $\Delta > 0$.[24] In this case, business cycle risk is bad ($\mathcal{E} < 0$), and (42) implies that firms with larger revenue-based Domar weights ω_i or smaller wedges τ_i choose, all else equal, a lower risk exposure δ_i.

Through that channel, ω_i and τ_i affect the variance of firm-level TFP growth, which is given by

$$\mathrm{V}\left[\log TFP_{it} - \log TFP_{it-1}\right] = 2\delta_i^2\Sigma + 2\Sigma_i^v. \tag{43}$$

It follows that if $\delta_i > 0$, firms with larger Domar weights and smaller wedges have less volatile TFP growth.[25] Furthermore, ω_i and τ_i affect how the TFP of a firm covaries with GDP. Using (16), we can show that

$$\mathrm{Cov}\left[\log TFP_{it} - \log TFP_{it-1}, y_t - y_{t-1}\right] = 2\Delta\Sigma\delta_i + 2\tilde{\omega}_i\Sigma_i^v. \tag{44}$$

Together with (42), this equation implies that the productivity growth of firms with larger Domar weights and lower wedges covaries less with aggregate TFP growth.

To see whether 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 2018. Our sample contains 7,513,081 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, but include more details 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 $1 + \tau_{it} = \hat{\alpha}_{Li} \Big/ \left(\frac{\text{Wage Bill}_{it}}{\text{Sales}_{it}}\right)$, where $\hat{\alpha}_{Li}$ is the Levinsohn and Petrin (2003)'s estimate of labor elasticity in production (specific to each 2-digit sector) and $\frac{\text{Wage Bill}_{it}}{\text{Sales}_{it}}$ is the share of labor expenditure in firms' sales.[26] The median markup is 1.51 and the median Domar weight is 3.6×10^{-7} across all firm-year observations in our sample. Finally, we compute each firm's TFP as a markup-corrected deflated Solow residual.

With these data, we first explore how the volatility of a firm's TFP growth covaries with its Domar weight and markup. To do so, we compute the standard deviation of TFP growth for each firm, $\sigma_i\left(\Delta \log TFP_{it}\right)$, and the time-series average of its markup and Domar weight. We construct deciles based on average Domar weights and markups, and create a set of dummy variables, FE_{ji}^{Domar} and FE_{ji}^{Markup}, such that $FE_{ji}^{Domar} = 1$ if firm i's Domar weight is in decile j, and

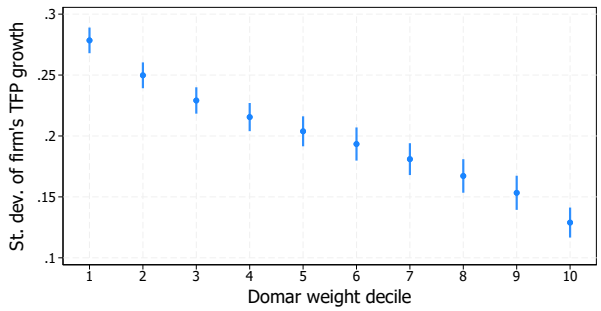
[24] Since $\mathrm{E}[\varepsilon_t] = 0$, this is essentially a normalization that implies that on average firms have positive exposure $\delta_{it} > 0$. Alternatively, we can consider an equivalent tilde economy in which $\tilde{\delta}^\circ = -\delta^\circ$, $\tilde{\delta}_{it} = -\delta_{it}$, $\tilde{\Delta} = -\Delta < 0$ and $\tilde{\mathcal{E}} = -\mathcal{E} > 0$. That economy would be indistinguishable from the original one in terms of the variance of firm-level TFP growth and its covariance with GDP, as (43) and (44) show below. Indeed, our results in Section 6 regarding 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 the signs of Δ and $\mathcal{E}$ separately.

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

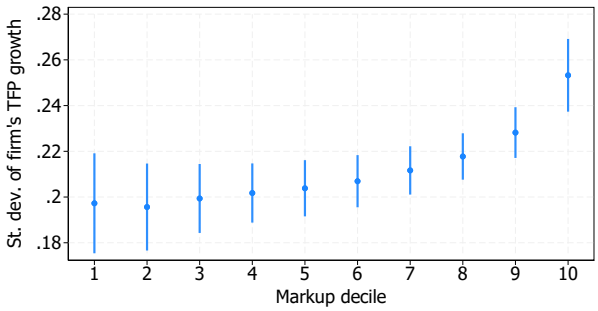
[26] In Appendix B.3, we show that our results are robust to using other measures of markups.

Figure 4: TFP volatility, Domar weights and markups

(a) TFP volatility by Domar weight decile



(b) TFP volatility by markup decile



Notes. Estimation results of (45) using a sample of Spanish firms from Orbis. 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. Sample construction is described in Appendix B.2.

analogously for markups. We then run the cross-sectional regression

$$\sigma_i\left(\Delta \log TFP_{it}\right) = \alpha + \sum_{j=1}^{10} \beta_j^{Domar} FE_{ji}^{Domar} + \sum_{j=1}^{10} \beta_j^{Markup} FE_{ji}^{Markup} + \varepsilon_i, \tag{45}$$

and plot β_j^{Domar} in panel (a) and β_j^{Markup} in panel (b) of Figure 4. We find that 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 15 p.p. less volatile than a firm in the bottom decile. In contrast, firms in the top decile of the markup distribution are about 6 p.p. more volatile than firms in the bottom decile.[27] Our theory is consistent with these findings. In the model, firms with small Domar weights and high markups are less aggressive in managing risk, and are therefore more volatile.

To confirm the empirical findings of Figure 4, we also run a simple cross-sectional regression of the standard deviation of firm-level TFP growth on Domar weights and markups. The results are reported in Table 2 in Appendix B.3. There, we see that a 1% increase in a firm's Domar weight is associated with a decline in TFP growth volatility of 0.032 log points, and a 1% increase in a firm's markup is associated with an increase in TFP growth volatility of 0.030 log points (in our sample, average TFP growth volatility is 20.8%). The results are statistically significant. In Appendix B.3, we conduct an analogous analysis using a sample of US firms from Compustat. Despite this data having much fewer firms and covering a significantly smaller fraction of the US economy, we find similar results.

Next, we turn to our model's prediction that the TFP of firms with higher Domar weights and lower markups should covary less with GDP. Again, we construct a set of dummy variables,

[27] See Stanley et al. (1996) and Yeh (2023) for evidence that sales growth volatility declines with firm size.

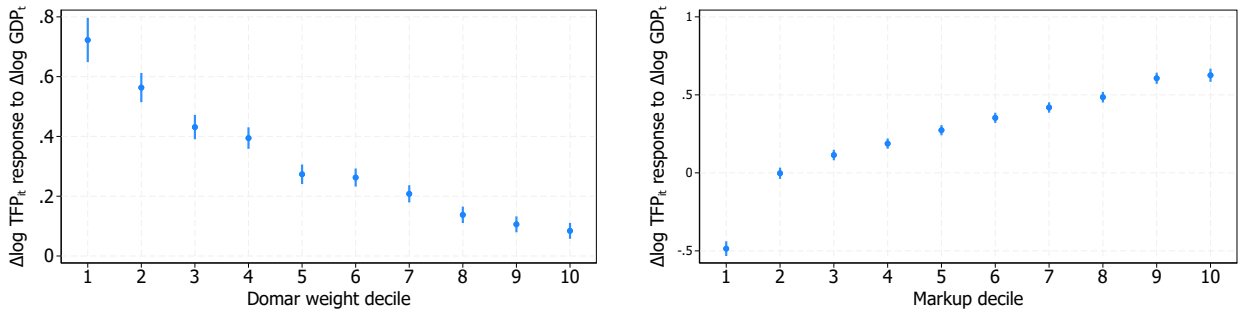
FE^{Domar}_{jit} and FE^{Markup}_{jit}, such that $FE^{Domar}_{jit} = 1$ if firm i's Domar weight is in decile j in year t, and analogously for markups. Notice that these variables are time dependent for this exercise. We then run the following panel regression,

$$\Delta \log TFP_{it} = \sum_{j=1}^{10} \beta_j^{Domar} \left(FE^{Domar}_{jit} \times \Delta \log GDP_t \right) + \sum_{j=1}^{10} \beta_j^{Markup} \left(FE^{Markup}_{jit} \times \Delta \log GDP_t \right)$$
$$+ \alpha + \beta_0 \Delta \log GDP_t + \sum_{j=1}^{10} FE^{Domar}_{jit} + \sum_{j=1}^{10} FE^{Markup}_{jit} + \varepsilon_{it}, \tag{46}$$

where $\Delta \log TFP_{it}$ is the annual growth of firm i's log TFP and $\Delta \log GDP_t$ is the annual growth of Spanish log GDP. The coefficients of interest, β_j^{Domar} and β_j^{Markup}, are reported in Figure 5. Panel (a) shows that firms with Domar weights in the top decile covary substantially less with GDP than firms in the bottom decile. The difference is economically large, about 64 p.p., and statistically significant. Panel (b), in contrast, shows that firms with high markups covary more with GDP than their low-markup counterparts. The relationship is economically and statistically significant.

Figure 5: 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 (46) using a sample of Spanish firms from Orbis. 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. Sample construction is described in Appendix B.2.

In Appendix B.3, we investigate the relationship between these quantities through simple panel regressions. Controlling for various sets of fixed effects, we still find that firms with high Domar weights and low markups comove less strongly with the business cycles. Again, these regressions and the results of Figure 5 are consistent with our theory. In the model, those firms manage their risk exposure more aggressively and decide to be less exposed to aggregate fluctuations.[28]

[28] Large firm sales also covary less with GDP than those of small firms in US data (Crouzet and Mehrotra, 2020).

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, provided that productivity increases generally benefit 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.

Overall, the findings of this section suggest that key predictions of our theory are in line with firm-level productivity data and stock market return data.[29] Of course, other mechanisms could also be at work (see Yeh (2023) for a discussion of the literature on firm size and volatility). 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.

8 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 already introduced in the previous section to identify the parameters of the model. With the calibrated model in hand, we investigate how the presence of wedges affects aggregate volatility, and how the economy handles an increase in the variance of a risk factor. We present an overview of our calibration strategy below and include more detail in Appendix C.

8.1 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 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. The model that we present in this section is dynamic but firms make the same risk-exposure decisions every period. We will therefore analyze the model in a single period.

There are S sectors. In each sector s, there is an aggregator that converts the output of the N_s

[29] 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 richer countries have access to better technologies, including 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 larger governments feature more volatile TFP. Larger governments imply more taxes, and insofar as some of those taxes affect the gap between prices and production costs, our theory would suggest that firms in these countries are less aggressive with their risk management, in line with the data.

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 $\sum_{i=1}^{N_s} \theta_{si} = 1$, $z_s \sim \text{iid}\, \mathcal{N}(\mu_s^z, \Sigma_s^z)$ are sectoral productivity shocks, and where Q_{si} denotes the output of firm i in industry s. This firm, in turn, produces according to the production function

$$Q_{si} = e^{\delta_{sit}\varepsilon_t + \gamma_{si}t + v_{sit}} \zeta_{si} L_{si}^{1-\sum_{s'=1}^{S} \hat{\alpha}_{ss'}} \prod_{s'=1}^{S} X_{si,s'}^{\hat{\alpha}_{ss'}}, \tag{47}$$

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.[30] Finally, the household only consumes goods produced by the sectoral aggregators so that GDP is given by $Y = \prod_{s=1}^{S} \left(\beta_s^{-1} C_s\right)^{\beta_s}$, where $\sum_{s=1}^{S} \beta_s = 1$.

As in Section 7, the TFP of firm i is the sum of a trend $\gamma_{si}t$, an idiosyncratic shock $v_{sit} \sim \text{iid}\, \mathcal{N}(\mu_{si}^v, \Sigma_{si}^v)$, and an aggregate component $\varepsilon_t \sim \text{iid}\, \mathcal{N}(0, \Sigma)$ to which firms can adjust their exposure. The stationary structure of the model implies that firms make the same risk-exposure decisions in every period, and so from now on we drop the subscript t when writing 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

$$\mathrm{V}\left[y_t - y_{t-1}\right] = 2\Sigma\Delta^2 + 2\tilde{\omega}_f^\top \Sigma^v \tilde{\omega}_f + 2\tilde{\omega}_s^\top \Sigma^z \tilde{\omega}_s, \tag{48}$$

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 properly-weighted 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. We can rewrite (48) as

$$\mathrm{V}\left[y_t - y_{t-1}\right] = \sum_{s=1}^{S} \sum_{i=1}^{N_s} \tilde{\omega}_{si} \,\mathrm{Cov}\left[y_t - y_{t-1}, \log \mathrm{TFP}_{si,t} - \log \mathrm{TFP}_{si,t-1}\right] + 2\tilde{\omega}_s^\top \Sigma^z \tilde{\omega}_s. \tag{49}$$

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.

[30] 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.

8.2 Calibration strategy

Our calibration strategy aims at replicating the firm-level Orbis dataset into our model economy. Our calibrated economy therefore features 62 sectors and 492,917 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%.

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 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 7. We therefore rely on our markup estimates from that section to pin down the firm-level markups τ_{si}. Our reduced-form analysis also yielded 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 (43) and (44) to identify the risk exposure δ_{si} and the variance Σ^v_{si} of each firm in the economy. We provide more details about the procedure that we use in Appendix C.3.

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.[31]

Finally, we can estimate the parameters of the firm-level risk-management cost function, δ°_{si} and H^{-1}_{si}. We rely on Proposition 2, which states that

$$\delta_i = \delta^\circ_i + \frac{\omega_i}{1+\tau_i}\left(\frac{1}{\eta}H_i^{-1}\right)\mathcal{E}. \tag{50}$$

We proceed through a curve-fitting exercise. In our baseline calibration, we assume that $\frac{1}{\eta}H^{-1}_{si}$ is the sum of a constant and a power function of the firm's Domar weight. We allow $\frac{1}{\eta}H^{-1}_{si}$ to potentially vary with $\tilde{\omega}_i$ to capture the fact that bigger firms might face different costs of managing risk. Both the constant and the parameters of the power function are free to change across sectors. We also assume that δ°_{si} is the sum of a sector-specific term and a firm-specific residual. We estimate these terms sector by sector by minimizing the sum of the squares of the residual.[32] In Appendix C.5, we show that our results are robust to assuming that H^{-1}_{si} is constant by sector instead.

[31] 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 also increases by 1%.

[32] That is, we assume $\delta_{si} = \delta^\circ_s + \frac{\omega_i}{1+\tau_i}\frac{1}{\eta}H^{-1}_{si}\mathcal{E} + u_{si}$, where $\frac{1}{\eta}H^{-1}_{si} = a_s\tilde{\omega}^{b_s}_{si} + c_s$, and where the parameters $a_s \geq 0$, $b_s \leq 0$ and $c_s \geq 0$ are sector-specific constants to be estimated. For a few sectors, the estimated $\frac{1}{\eta}H^{-1}_{si}$ is slightly negative. We set those estimates to zero. These are sectors for which adjusting risk is infinitely costly in our estimation.

Of the five biggest sectors in the economy, "public administration and defense" and "health" are the ones for which adjusting risk exposure is the most costly. In contrast, adjusting exposure in "accommodation and food services" and "retail trade" is cheap.

Our calibration procedure cannot distinguish between $\frac{1}{\eta}$ and H_{si}^{-1}, and only provides an estimate for their product $\frac{1}{\eta}H_{si}^{-1}$. But since these quantities enter the model equations only through that product, there is no need for us to pick values for $\frac{1}{\eta}$ and H_{si}^{-1} separately. We therefore remain agnostic about the supply elasticity η of risk management resources.

Overall, our calibration procedure ensures that we match certain features of the data exactly. That is the case for the 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.4 provides more details about the calibrated economy.

8.3 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 when firms must keep their previous risk exposure δ fixed, and when they can adjust δ in response to the change.

When Σ doubles and δ remains fixed, (48) 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 2.4% in the calibrated economy to 3.1%. Now that aggregate risk is more important, exposure value $\mathcal{E}$ 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 lower their risk exposure, they decide to do so, and Δ falls from 0.014 to 0.011 in response. This decline in Δ leads to a decline in aggregate volatility, and the standard deviation of GDP is reduced to 2.6%. 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.

We can also see the impact of the change in Σ in the firm-level micro data. The left panel in Figure 6 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.[33] Since on average firms are positively

[33] This can be seen more readily in (44). 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 somewhat to mitigate that effect without offsetting it fully.

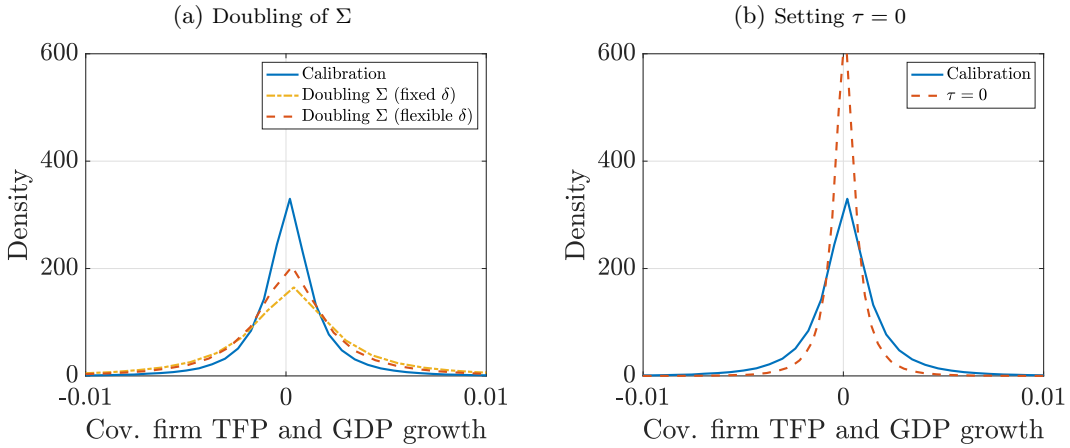
exposed to ε_t ($\Delta > 0$), this widening leads to an increase in the average covariance and, by (49), to an increase in aggregate volatility. When δ is flexible, firms respond to this large increase in aggregate volatility by managing their risk more aggressively, which undoes some of the widening, and mitigates the overall increase in aggregate risk.

Table 1: Exposure and GDP volatility in different environments

	Calibration	Doubling Σ		No wedges	
		Fixed δ	Flexible δ	Fixed δ	Flexible δ
Std. Dev. of GDP growth	2.4%	3.1%	2.6%	2.4%	1.7%
Aggregate risk exposure Δ	0.014	0.014	0.011	0.014	0.007
Exposure value $\mathcal{E}$	−0.06	−0.11	−0.09	−0.06	−0.03
Share of endogenous vol.	69%	81%	74%	69%	35%

Notes: Share of endogenous volatility is $2\Sigma\Delta^2 / \left(\mathrm{V}\left[\log \mathrm{GDP}_t - \log \mathrm{GDP}_{t-1}\right]\right)$.

Figure 6: 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 $\tau = 0$.

8.4 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 (48) 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 are free to adjust their risk exposure. Since the removal of the wedges makes firms larger, they manage their risk exposure more aggressively and, from (50), this leads to a decline in δ. As a result, aggregate exposure Δ falls from 0.014 to 0.007 (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 firms control how much risk they are exposed to.

Again, we can look at the impact of the removal of the wedges in firm-level micro data. The right panel of Figure 6 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 (44), 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 (49).

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 they capture some of the key mechanisms at work in reality, and that they provide a rough estimate of their importance for the economy.

9 Conclusion

This paper proposes a parsimonious theory of endogenous risk in which firms are free to pick 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 would be worth exploring. 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).

Another promising extension would involve a fully dynamic business cycle model. Indeed, it seems natural that today's risk management decisions would affect future payoffs. These decisions might interact with investment, with potential consequences for business cycle dynamics.

References

Acemoglu, D. and P. D. Azar (2020): "Endogenous Production Networks," *Econometrica*, 88, 33–82.

Acemoglu, D., V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi (2012): "The Network Origins of Aggregate Fluctuations," *Econometrica*, 80, 1977–2016.

Acemoglu, D. and F. Zilibotti (1997): "Was Prometheus unbound by chance? Risk, diversification, and growth," *Journal of Political Economy*, 105, 709–751.

Ackerberg, D. A., K. Caves, and G. Frazer (2015): "Identification properties of recent production function estimators," *Econometrica*, 83, 2411–2451.

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

Alesina, A., S. Özler, N. Roubini, and P. Swagel (1996): "Political instability and economic growth," *Journal of Economic Growth*, 1, 189–211.

Baqaee, D. R. and E. Farhi (2019a): "Productivity and Misallocation in General Equilibrium," *Quarterly Journal of Economics*, 135, 105–163.

——— (2019b): "Replication Data for: 'Productivity and Misallocation in General Equilibrium'," Replication package.

Barro, R. J. (2006): "Rare disasters and asset markets in the twentieth century," *Quarterly Journal of Economics*, 121, 823–866.

Bigio, S. and J. La'O (2020): "Distortions in Production Networks," *Quarterly Journal of Economics*, 135, 2187–2253.

Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta-Eksten, and S. J. Terry (2018): "Really Uncertain Business Cycles," *Econometrica*, 86, 1031–1065.

Carvalho, V. M., M. Nirei, Y. U. Saito, and A. Tahbaz-Salehi (2021): "Supply Chain Disruptions: Evidence from the Great East Japan Earthquake," *Quarterly Journal of Economics*, 136, 1255–1321.

Clementi, G. L. and B. Palazzo (2016): "Entry, exit, firm dynamics, and aggregate fluctuations," *American Economic Journal: Macroeconomics*, 8, 1–41.

Cole, H. L., J. Greenwood, and J. M. Sanchez (2016): "Why doesn't technology flow from rich to poor countries?" *Econometrica*, 84, 1477–1521.

Crouzet, N. and N. R. Mehrotra (2020): "Small and large firms over the business cycle," *American Economic Review*, 110, 3549–3601.

De Loecker, J., J. Eeckhout, and G. Unger (2020): "The rise of market power and the macroeconomic implications," *Quarterly Journal of Economics*, 135, 561–644.

De Loecker, J. and F. Warzynski (2012): "Markups and firm-level export status," *American Economic Review*, 102, 2437–2471.

Fama, E. F. and K. R. French (1992): "The cross-section of expected stock returns," *Journal of Finance*, 47, 427–465.

Fan, J. (2024): "Talent, geography, and offshore R&D," *Review of Economic Studies*, rdae044.

Gabaix, X. (2011): "The Granular Origins of Aggregate Fluctuations," *Econometrica*, 79, 733–772.

Gopinath, G., Ş. Kalemli-Özcan, L. Karabarbounis, and C. Villegas-Sanchez (2017): "Capital allocation and productivity in South Europe," *Quarterly Journal of Economics*, 132, 1915–1967.

Greenwood, J. and B. Jovanovic (1990): "Financial development, growth, and the distribution of income," *Journal of political Economy*, 98, 1076–1107.

Grossman, G. M. and E. Helpman (1991): *Innovation and growth in the global economy*, MIT press.

Hulten, C. R. (1978): "Growth Accounting with Intermediate Inputs," *Review of Economic Studies*, 45, 511–518.

Ibbotson, R. G., P. D. Kaplan, and J. D. Peterson (1997): "Estimates of small-stock betas are much too low," *Journal of Portfolio Management*, 23, 104–111.

Jensen, M. C. and W. H. Meckling (1976): "Theory of the firm: Managerial behavior, agency costs and ownership structure," *Journal of Financial Economics*, 3, 305–360.

Jones, C. I. (1995): "R&D-based models of economic growth," *Journal of Political Economy*, 103, 759–784.

——— (2011): "Intermediate Goods and Weak Links in the Theory of Economic Development," *American Economic Journal: Macroeconomics*, 3, 1–28.

Kalemli-Özcan, Ş., B. E. Sørensen, C. Villegas-Sanchez, V. Volosovych, and S. Yeşiltaş (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, 353–374.

Khan, A. and J. K. Thomas (2008): "Idiosyncratic shocks and the role of nonconvexities in plant and aggregate investment dynamics," *Econometrica*, 76, 395–436.

King, R. G., C. I. Plosser, and S. T. Rebelo (1988): "Production, growth and business cycles: I. The basic neoclassical model," *Journal of Monetary Economics*, 21, 195–232.

Klette, T. J. and S. Kortum (2004): "Innovating firms and aggregate innovation," *Journal of Political Economy*, 112, 986–1018.

Kopytov, A., B. Mishra, K. Nimark, and M. Taschereau-Dumouchel (2024): "Endogenous Production Networks under Supply Chain Uncertainty," Working paper.

Koren, M. and S. Tenreyro (2007): "Volatility and development," *Quarterly Journal of Economics*, 122, 243–287.

——— (2013): "Technological diversification," *American Economic Review*, 103, 378–414.

Kydland, F. E. and E. C. Prescott (1982): "Time to Build and Aggregate Fluctuations," *Econometrica*, 50, 1345–1370.

Levinsohn, J. and A. Petrin (2003): "Estimating production functions using inputs to control for unobservables," *Review of Economic Studies*, 70, 317–341.

Liu, E. (2019): "Industrial Policies in Production Networks," *Quarterly Journal of Economics*, 134, 1883–1948.

Long, J. B. and C. I. Plosser (1983): "Real Business Cycles," *Journal of Political Economy*, 91, 39–69.

Oberfield, E. (2018): "A Theory of Input-Output Architecture," *Econometrica*, 86, 559–589.

Olley, G. S. and A. Pakes (1996): "The dynamics of productivity in the telecommunications equipment industry," *Econometrica*, 64, 1263.

Pellet, T. and A. Tahbaz-Salehi (2023): "Rigid production networks," *Journal of Monetary Economics*, 137, 86–102.

Ramey, G. and V. A. Ramey (1995): "Cross-Country Evidence on the Link Between Volatility and Growth," *American Economic Review*, 85, 1138–1151.

Romer, P. M. (1990): "Endogenous technological change," *Journal of Political Economy*, 98, S71–S102.

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

Stanley, M. H., L. A. Amaral, S. V. Buldyrev, S. Havlin, H. Leschhorn, P. Maass, M. A. Salinger, and H. E. Stanley (1996): "Scaling behaviour in the growth of companies," *Nature*, 379, 804–806.

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

Online Appendix

A Additional derivations

This appendix contains the derivation of some expressions featured in the main text.

A.1 First-order conditions of the household

In this appendix, we derive equations (6) and (7).

The household chooses how much to consume after uncertainty about ε is realized but how many risk management resources 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 R is taken as given here. 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 $U(Y) = \frac{Y^{1-\rho}}{1-\rho}$. Summing over i on both sides of this equation and using the binding budget constraint, we find

$$\mathcal{U}'(Y)\, Y\, (V(R))^{1-\rho} = \Lambda\left(W_L + W_R R + \Pi\right). \tag{51}$$

Combining with the first-order condition implies

$$P_i C_i = \beta_i \left(W_L L + W_R R + \Pi\right). \tag{52}$$

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

$$\begin{aligned} Y = \prod_{i=1}^{N} \left(\beta_i^{-1} C_i\right)^{\beta_i} = \prod_{i=1}^{N} \left(\beta_i^{-1} \frac{\beta_i \mathcal{U}'(Y)\, Y\, (V(R))^{1-\rho}}{\Lambda P_i}\right)^{\beta_i} \Leftrightarrow \\ \Lambda = \mathcal{U}'(Y)\, (V(R))^{1-\rho} \prod_{i=1}^{N} P_i^{-\beta_i}, \end{aligned} \tag{53}$$

which we can combine with (51) to find

$$Y = \left(W_L L + W_R R + \Pi\right) \prod_{i=1}^{N} P_i^{-\beta_i}. \tag{54}$$

This last equation implicitly defines a price index,

$$\bar{P} = \prod_{i=1}^{N} P_i^{\beta_i}, \tag{55}$$

such that $\bar{P}Y = W_L L + W_R R + \Pi$. Therefore, (53) is equivalent to (6).

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

$$\mathrm{E}\left[Y^{1-\rho}\right] (V(R))^{-\rho} V'(R) + \mathrm{E}\left[\Lambda\right] W_R = 0. \tag{56}$$

Using (53), we find

$$\mathrm{E}\left[Y^{1-\rho}\right] \frac{V'(R)}{V(R)} + \mathrm{E}\left[\frac{Y^{1-\rho}}{\bar{P}Y}\right] W_R = 0. \tag{57}$$

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, (57) simplifies to (7).

B Appendix for Section 7

This appendix contains details about the reduced-form results of Section 7.

B.1 Derivation of (43) and (44)

From (16), we can write log real GDP in the model of Section 7 as

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

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

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

where μ^v is the vector of expected values 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 (22).

Next, the TFP process for firm i is given by $\log TFP_i = \delta_i \varepsilon_t + \gamma_i t + v_{it}$, where $v_{it} \sim$ iid $\mathcal{N}(\mu_i^v, \Sigma_i^v)$

and $\varepsilon_t \sim \text{iid } \mathcal{N}(0, \Sigma)$. It follows that

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

Similarly, for the covariance we have

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

Those are the equations reported in Section 7.

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-universe (covering more than 95% of total industry gross output after 2010) coverage of firms as detailed in Gopinath et al. (2017) and Kalemli-Özcan et al. (2024). We use the sample period 1995-2018 for our analysis.[34]

Sample cleaning Following the procedure of Kalemli-Özcan et al. (2024), we link multiple vintages of Orbis products over time and link the 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.

2. In a given year, firms might have multiple values of reported sales from different sources (local registry, annual report, or others), for consolidated or unconsolidated accounts. Following

[34] 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 2018 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). We also verify this by noting that the number of firms peaked in 2018, suggesting that data collection for 2019 has not yet been finalized in our version of data.

Fan (2024), we keep the unconsolidated accounts as the consolidated account might include other firms in the conglomerate which can lead to double counting.

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 drop firms if their average revenue products of any input (fixed asset, wage bills, and material costs) are larger than the 99.9 percentile of the distribution in any year.

Variable definitions, firm-level markup and TFP estimation 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 consider a Cobb-Douglas production function of the following form:

$$\log Q_{it} = \alpha_{Li} \log L_{it} + \alpha_{Mi} \log M_{it} + \alpha_{Ki} \log K_{it} + \varepsilon_{it}, \tag{58}$$

where the Q_{it}, L_{it}, M_{it} and K_{it} are deflated values of, respectively, sales, wage bills, material costs and fixed asset for each firm i in calendar year t. Sales, wage bills and material costs are deflated using the industry-specific gross output price indices. We deflate fixed assets using the capital price indices. All 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 et al. (2015) 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 et al. (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.

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,[35]

$$\begin{aligned}\Delta \log \text{TFP}_{it} =& \Delta \log Q_{it} - \alpha_{Li} \Delta \log L_{it} - \alpha_{Mi} \Delta \log M_{it} - \alpha_{Ki} \Delta \log K_{it} \\ & - \left(\Delta \log \left(1 + \tau_{it}\right) - \Delta \log \left(1 + \tau_{s(i)t}\right) \right).\end{aligned} \tag{59}$$

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 change in firm-specific nominal price that are not taken into account by the sector-level price deflator.[36] We thus move closer to a quantity-based interpretation of our TFP measure.

For all reduced-form exercises, we restrict our sample to only include firms with more than 10 non-missing observations of measured TFP growth ($\Delta \log \text{TFP}_{it}$) such that TFP volatility can be precisely estimated. Additionally, for all cross-sectional regressions, we only keep the firms with 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 the firm-level Domar weight and markup and (2) the WRDS Beta Suite database that provides stocks' loading on the aggregate return, i.e. stock market betas.

CRSP/Compustat Merged Our sample selection and cleaning procedure follows standard practices in the literature, particularly those carried out in Baqaee and Farhi (2019a).

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. We use the data from 1979 to 2015 due to the availability of industry-level price indices.[37]

2. We use PERMNO as our firm identifier to allow for easy merge with stock market information.

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

[35] 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.

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

[37] Throughout, we use the KLEMS industry price indices obtained from the replication package by Baqaee and Farhi (2019b).

4. As in Baqaee and Farhi (2019b), we drop firm-year observations with COGS-to-sales and XSGA-to-sales ratios in the top and bottom 2.5% of the corresponding year-specific 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_i for each firm-month pair is obtained from the time series regression:

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

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 set to be `monthly`; (2) both the estimation window and the minimum window is set to be 60 months; (3) risk model is set to be `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 of the calendar year.

Variable definitions, firm-level markup and TFP estimation 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 U.S. Bureau of Economic Analysis).

Our production function estimation procedure for the Compustat sample largely follows Baqaee and Farhi (2019a). To implement the estimation procedure, we consider a Cobb-Douglas production function of the following form:

$$\log Q_{it} = \alpha_{Vit} \log V_{it} + \alpha_{Kit} \log K_{it} + \varepsilon_{it}, \tag{60}$$

where Q_{it}, V_{it}, and K_{it} are deflated values of, respectively, sales, costs of goods sold, and fixed asset for each firm i in calendar year t.[38] We deflate sales and costs of goods sold using the gross output price indices and deflate fixed assets using the capital price indices.

The output elasticities are estimated using the Olley and Pakes (1996) methodology with the Ackerberg et al. (2015) correction. We use fixed assets as the "state" variable, costs of goods sold as the "free" variable and investment (`capx` deflated by the capital price indices) as the "proxy" variable. We estimate the production function for each sector and each year.[39] As panel data are required in the Olley and Pakes (1996) approach, we use 5-year rolling windows such that the elasticity estimates for year t is obtained using data from year $t-2$ to year $t+2$. In the

[38] As discussed in De Loecker et al. (2020), we use "costs of goods sold" as an input because material costs and wage bills are not separately reported in Compustat.

[39] We use the same 14-sector partition as Baqaee and Farhi (2019a), which is based on SIC 2-digit codes.

estimation, we also control for markups using firms' sales shares at the SIC 3-digit and 4-digit industry level. Treating "costs of goods sold" 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 $\frac{\text{Costs of Goods Solds}_{it}}{\text{Sales}_{it}}$ is the share of variable input expenditure in firms' sales. We compute firm-level TFP growth as

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

where $\Delta \log(1+\tau_{s(i)t})$ is the sector-level markup growth and $s(i)$ represents the SIC 2-digit sector that firm i belongs to. For all reduced-form exercises, we only consider the firms with more than 10 observations of measured TFP growth and positive average log markup. Moreover, for the analysis of equity beta in Section B.4, we further restrict our sample to firm-year observations that have valid Domar weights, markups, and 5-year CAPM beta estimates.

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 10 or 11 such that we only keep the firms that are traded in the NYSE to construct breakpoints.

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 various additional results related to the reduced-form exploration of Section 7.

Volatility of TFP growth First, we run a cross-sectional regression of firm-level standard deviation TFP growth on average log Domar weights and log markups,

$$\sigma_i\left(\Delta \log TFP_{it}\right) = \alpha + \beta^{Domar} \log\left(\text{Domar Weight}_i\right) + \beta^{Markup} \log\left(\text{Markup}_i\right) + \varepsilon_i. \tag{61}$$

The results are given in column (1) of Table 2. Consistent with our theory, we find $\beta^{Domar} < 0$ and $\beta^{Markup} > 0$ for both the Orbis Spain sample and the Compustat sample.

Table 2: Volatility of TFP growth, Domar weights and markups

	Dependent variable: Volatility of firm-level TFP growth	
	(1): Orbis Spain	(2): Compustat
$\log(\text{Domar Weight}_i)$	-0.032^{***}	-0.0025^{***}
	(0.001)	(0.0003)
$\log(\text{Markup}_i)$	0.030^{***}	0.0098^{***}
	(0.007)	(0.0032)
Observations	241,557	3,589
R^2	0.098	0.122

Notes: The table presents estimation results of (61). The sample includes all Orbis Spain (column 1) and Compustat (column 2) firms with 1) at least 10 non-missing observations of TFP growth, and 2) positive average log markup. The sample is trimmed at the top and bottom 0.5% of observations of average Domar weight, average markup, and the standard deviation of TFP growth. Standard errors (in parentheses) are clustered at the NACE 4-digit industry level for the Orbis Spain sample in column (1), and at the NAICS 2-digit industry for the Compustat sample in column (2). *,** ,*** indicate significance at the 10%, 5%, and 1% levels, respectively.

Sensitivity of firm-level TFP to GDP Next, 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:

$$\Delta \log TFP_{it} = \beta_1 \log(\text{Domar Weight}_{it}) \times \Delta \log GDP_t + \beta_2 \log(\text{Markup}_{it}) \times \Delta \log GDP_t$$
$$\alpha + \beta_0 \Delta \log GDP_t + \beta_3 \log(\text{Domar Weight}_{it}) + \beta_4 \log(\text{Markup}_{it}) + FE + \varepsilon_{it}. \quad (62)$$

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 3 shows the results of estimating (62) without any fixed effects. In column (2), we add firm and year fixed effects, and in column (3), we use industry×year fixed effects. In all specifications, we find the statistically significant coefficients whose signs are consistent with our theory.

Robustness with an alternative markup measure In this section, we verify that the our main results from Table 2 and 3 are robust to an alternative markup measure. We compute this measure as $1 + \tau_{it} = \hat{\alpha}_{Mi} \Big/ \left(\frac{\text{Material Costs}_{it}}{\text{Sales}_{it}}\right)$, where we use the same set of production function estimates as before but now treat materials as the variable input. Using the material-based markup measure, we further compute a consistently defined measure of firm-level TFP growth using (59). Results from estimating (61) and (62) with this alternative markup measure are reported in columns (2) and (4) of Table 4, alongside our benchmark estimates in columns (1) and (3). The results are qualitatively similar: smaller firms and those with higher markups show higher TFP volatility and higher sensitivity of firm-level TFP to GDP.

Table 3: Sensitivity of firm-level TFP to GDP

	Dependent variable: Firm-level TFP growth		
	(1)	(2)	(3)
$\Delta \log GDP_t$	-0.770^{***}		
	(0.052)		
$\log(\text{Domar Weight}_{it}) \times \Delta \log GDP_t$	-0.064^{***}	-0.138^{***}	-0.147^{***}
	(0.004)	(0.004)	(0.004)
$\log(\text{Markup}_{it}) \times \Delta \log GDP_t$	0.400^{***}	0.428^{***}	0.416^{***}
	(0.009)	(0.009)	(0.010)
Firm FE	No	Yes	Yes
Year FE	No	Yes	No
Industry × Year FE	No	No	Yes
Observations	4,141,982	4,141,982	4,053,751
R^2	0.024	0.141	0.164

Notes: Table presents the results of estimation of (62) using a sample of Spanish firms from Orbis. For each firm, we compute its TFP as a markup-corrected deflated Solow residual, following the approach of Baqaee and Farhi (2019a). Firms' revenue-based Domar weights are computed as ratios of their nominal sales to Spain's nominal GDP. Firms' markups are estimated using the control function approach of De Loecker and Warzynski (2012) by treating wage bill as a flexible input. The estimation sample is trimmed at the top and bottom 0.5% of observations of TFP growth, Domar weight and markup. Standard errors (in parentheses) are clustered at the firm 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 covary more weakly 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 Baqaee and Farhi (2019a) and De Loecker et al. (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, we construct deciles based on a firm's Domar weight and its markup each year.[40] We then construct a set of dummy variables, FE_{jit}^{Domar} and FE_{jit}^{Markup}, such that $FE_{jit}^{Domar} = 1$ if firm i's Domar weight is in decile j in year t, and analogously for markups. We

[40] Following Fama and French (1992), we use NYSE breakpoints when defining the Domar weight deciles. The size distribution of firms in the 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.

Table 4: Results with alternative markup measures

Dependent variable	Firm-level TFP growth		Vol. of firm-level TFP gr.	
Markups estimated with	(1) Labor	(2) Materials	(3) Labor	(4) Materials
$\Delta \log GDP_t$	-0.770^{***}	-0.301^{***}		
	(0.052)	(0.053)		
$\Delta \log GDP_t$	-0.064^{***}	-0.022^{***}		
$\times \log(\text{Domar Weight}_{it})$	(0.004)	(0.004)		
$\Delta \log GDP_t$	0.400^{***}	0.149^{***}		
$\times \log(\text{Markup}_{it})$	(0.009)	(0.013)		
$\log(\text{Domar Weight}_i)$			-0.032^{***}	-0.028^{***}
			(0.001)	(0.004)
$\log(\text{Markup}_i)$			0.030^{***}	0.141^{***}
			(0.007)	(0.009)
Observations	4,141,982	4,144,831	241,557	131,742
R^2	0.024	0.039	0.098	0.224

Notes: Table presents the results of estimation of (62) (columns (1) and (2)) and (61) (columns (3) and (4)) using a sample of Spanish firms from Orbis. In columns (1) and (2), the dependent variable in the regression is the TFP growth for each firm, and the independent variables are the Domar weight and markup for each firm-year observation. The estimation sample for columns (1) and (2) is trimmed at the top and bottom 0.5% of observations of TFP growth, Domar weight and markup. In columns (3) and (4), the dependent variable in the regression is the standard deviation of TFP growth for each firm, and the independent variables are constructed using the time-series averages of Domar weight and markup for each firm. The estimation sample for columns (3) and (4) is trimmed at the top and bottom 0.5% of observations of average Domar weight and average markup. We also trim at the top and bottom 0.5% before computing the standard deviation of TFP growth. Across all specifications, we only keep the firms with more than 10 non-missing observations of measured TFP growth and positive average log markup. In columns (1) and (3), a firm's markup and markup-corrected TFP growth are computed using labor/wage bill-based markups. In columns (2) and (4), a firm's markup and markup-corrected TFP growth are computed using material-based markups. Standard errors (in parentheses) are clustered at the firm level. *,** ,*** indicate significance at the 10%, 5%, and 1% levels, respectively.

then run the following regression,

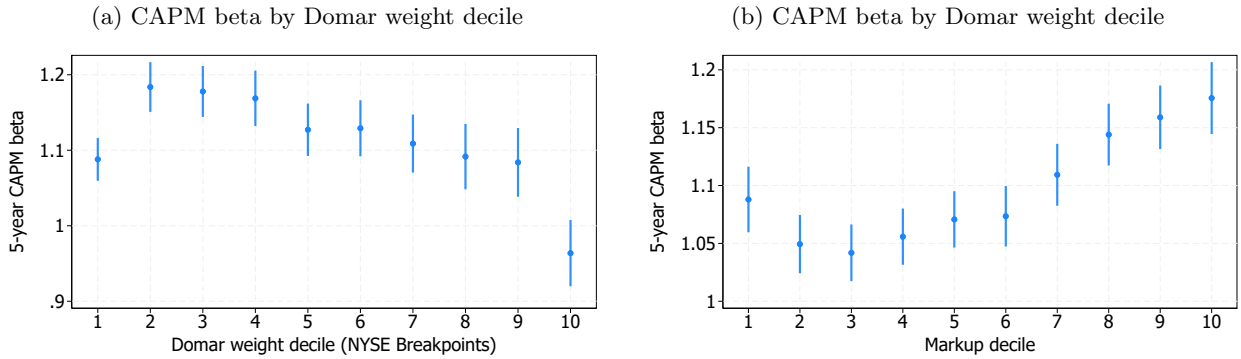
$$\beta_{it}^{CAPM} = b_0 + \sum_{j=1}^{10} b_j^{Domar} FE_{jit}^{Domar} + \sum_{j=1}^{10} b_j^{Markup} FE_{jit}^{Markup} + \varepsilon_{it}, \tag{63}$$

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 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 et al., 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.

B.5 Cross-country evidence

In this appendix, we show that countries with higher GDP per capita or a lower share of government in GDP tend to have lower TFP growth volatility. To do so, we use data from the Penn World Table (version 10.01). Our sample includes 6,171 country-year observations (unbalanced panel of 118 countries from 1956 to 2019). For each country, we compute the standard deviation

Figure 7: CAPM betas, Domar weights and markups



Notes. Estimation results of (63) using a sample of US firms from Compustat. Panel (a): $b_j^{Domar} + b_0$ by Domar weight decile j ($j = 1$ is lowest Domar weight, $j = 10$ is highest Domar weight). Panel (b): $b_j^{Markup} + b_0$ 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. Sample construction is described in Appendix B.2.

of TFP growth σ_{it}^{TFP} using 5-year rolling windows (we use TFP at constant national prices). The cross-country average of σ_{it}^{TFP} in our sample is 3.37%. Following Acemoglu and Zilibotti (1997), we then run the following regression:

$$\log \sigma_{it}^{TFP} = \alpha + \beta^{GDP} \log \mathrm{GDP}_{\mathrm{pc},it} + \beta^{Gov} \log gov_{it} + FE_i + FE_t + \varepsilon_{it}, \tag{64}$$

where $\log(\mathrm{GDP}_{\mathrm{pc},it})$ is log of GDP per capita of country i in year t (we use expenditure-side real GDP divided by total population), gov_{it} is the share of government consumption in GDP. The variables FE_i and FE_t are country- and year-fixed effects. The estimated coefficients are given in Table 5, and the associated binscatters are shown in Figure 8.

Our model can rationalize these findings. Countries with larger government consumption shares are taxed more heavily. Insofar as some of those taxes 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 H_i or η) which would explain the lower aggregate volatility.

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

Figure 9 shows that the share of management expenditure (measured as `other_operating_expenses` in the Orbis dataset) declines with a firm's Domar weight. Insofar as some of this expenditure involves managing risk, this would be consistent with the model's implication that risk management expenditures as a share of sales decrease with firm size.

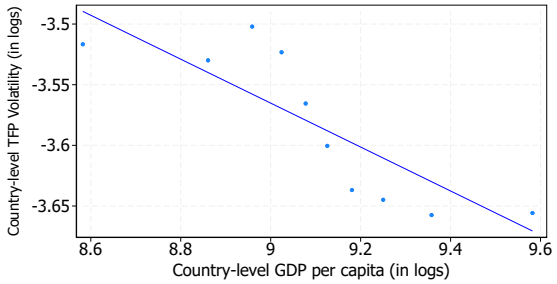
Table 5: Cross-country TFP volatility, GDP per capita and government share

	Dependent variable: country-specific TFP volatility		
	(1)	(2)	(3)
$\log \text{GDP}_{\text{pc},it}$	-0.181^{**}		-0.176^{**}
	(0.075)		(0.072)
$\log gov_{it}$		0.114^{**}	0.123^{**}
		(0.056)	(0.054)
Firm FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Observations	5,701	5,817	5,701
R^2	0.646	0.639	0.649

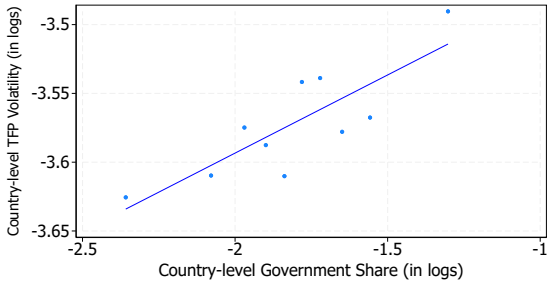
Notes: Results of estimating (64) using a cross-country sample from Penn World Tables. σ_{it} is year t volatility of country i's annual TFP growth, computed using 5-year rolling windows; $\log \text{GDP}_{\text{pc},it}$ is expenditure-side real GDP at chained PPPs per capita; gov_{it} is share of government consumption at curent PPPs. Standard errors (in parentheses) are clustered at the country level. *,** ,*** indicate significance at the 10%, 5%, and 1% levels, respectively.

Figure 8: Binscatters plots of cross-country TFP volatility, GDP per capita and government share

(a) TFP volatility and GDP per capita

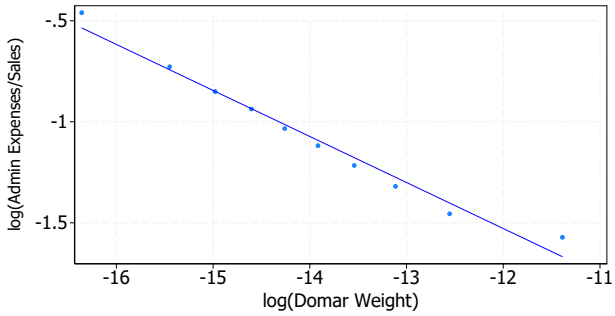


(b) TFP volatility and government share



Binned scatter plots of cross-country TFP volatility by the log of GDP per capita (left panel) and the log of government share (right panel). Cross-country TFP volatility is year t volatility of firm i's annual TFP growth, computed using 5-year rolling windows; GDP per capita is expenditure-side real GDP at chained PPPs per capita; government share is share of government consumption at current PPPs. The estimation sample is trimmed at the top and bottom 1% of observations of the country-level TFP volatility, GDP per capita and government share.

Figure 9: Administrative expenditure share and firm size



Binned scatter plots of the (log) administrative expenditure share by (log) Domar weight. The administrative expenditure share is defined as the ratio between "other operating expenses" and sales, and the Domar weight is computed as the ratio of the firm's nominal sales to Spain's nominal GDP. The estimation sample is trimmed at the top and bottom 0.5% of observations of the constructed expenditure share and Domar weight.

C Appendix for Section 8

This appendix contains details about the calibration of Section 8.

C.1 Construction of the sample

This appendix describes the datasets used in the quantitative exercise 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 1995-2018 period.

2. We calibrate the sectoral parameters using the 2010 input-output table from the Annual Spanish National Accounts. The 2010 input-output table partitions the Spanish economy into 62 sectors which are usually defined at the 2-digit NACE industry level.[41] 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. 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 the firm-level objects from the Orbis sample. After steps 1-4 in Section B.2.1, we follow a few additional steps.

[41] 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.

(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.[42]

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

(c) We drop firm-year observations if markup growth, real sales growth, real input (capital, labor, material) growth, or the level of markups exceed the top or bottom 0.1% of the distribution of all firm-year observations. We compute TFP growth using (59) 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 and revenue-based Domar weights, which correspond to $1+\tau_{si}$ and ω_{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.

C.2 Stationarity of the risk-exposure decision

In the model of Section 8, log real GDP is given by

$$\tilde{y} = \Delta\varepsilon + \left(\tilde{\omega}^f\right)^\top (v + \gamma t) + \left(\tilde{\omega}^s\right)^\top z - \left(\tilde{\omega}^f\right)^\top \log(1+\tau) - \log\Gamma_L, \tag{65}$$

where $\tilde{\omega}^f$ and $\tilde{\omega}^s$ are the cost-based Domar weights 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

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

where 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 (22).

[42] Trimming the data at the 0.1% level instead has only a minimal impact on the counterfactual exercises of Section 8 but leads to a greater number of extreme firms being dropped during the calibration procedure (see last part of Appendix C.3).

C.3 Identifying δ_{si} and Σ^v_{si}

Combining firm i's TFP from (47) with log real GDP from (65), we can write

$$\begin{aligned} A_{si} := \mathrm{V}\left[\log TFP_{si,t} - \log TFP_{si,t-1}\right] &= \mathrm{V}\left[\delta_{si}\varepsilon_t + v_{sit} + \gamma_{si}t - \delta_{si}\varepsilon_{t-1} - v_{sit-1} - \gamma_{si}\left(t-1\right)\right] \\ &= \mathrm{V}\left[\delta_{si}\left(\varepsilon_t - \varepsilon_{t-1}\right) + v_{si,t} - v_{si,t-1}\right] \\ &= 2\delta_{si}^2\Sigma + 2\Sigma^v_{si}. \end{aligned} \tag{66}$$

Similarly, for the covariance, we have

$$\begin{aligned} B_{si} := \mathrm{Cov}\left[y_t - y_{t-1}, \log TFP_{si,t} - \log TFP_{si,t-1}\right] = \mathrm{Cov}\Big[&\Delta_t\varepsilon_t - \Delta_{t-1}\varepsilon_{t-1} + \left(\omega^f\right)^\top\left(v_t - v_{t-1}\right) \\ &+ \left(\omega^s\right)^\top\left(z_t - z_{t-1}\right), \delta_{si}\varepsilon_t - \delta_{si}\varepsilon_{t-1} + v_{si,t} - v_{si,t-1}\Big] \\ = 2\Delta\Sigma\delta_{si} &+ 2\tilde{\omega}_{si}\Sigma^v_{si}. \end{aligned} \tag{67}$$

In these equations, we define A_{si} and B_{si} to simplify the notation in what follows. Notice also that (66) and (67) are the same as in Section 7, 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 (66), we can write

$$\Sigma^v_{si} = \frac{A_{si}}{2} - \delta_{si}^2\Sigma. \tag{68}$$

Combining with (67), we find a quadratic equation in δ_{si}, whose solutions are given by

$$\delta_{si} = \frac{\Delta}{2\tilde{\omega}^f_{si}}\left(1 \pm \sqrt{1 - 2\tilde{\omega}^f_{si}\frac{B_{si} - \tilde{\omega}^f_{si}A_{si}}{\Delta^2\Sigma}}\right). \tag{69}$$

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

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

or

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

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 (69), we find δ_{si} (we pick the negative root because it corresponds to a positive Δ, in line with our normalization). Finally, combining with

(68) gives us Σ^v_{si}.

We use this procedure to pin down the risk exposure δ_{si} and the idiosyncratic volatility Σ^v_{si} 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}.[43] 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 measurement (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.

C.4 Calibrated economy

The calibrated economy matches several features of the data perfectly. Figure 10 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),[44] 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 11. 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 midsize firms. The right panel of Figure 11 shows that the calibrated distribution of cost-based Domar weights, which we do not observe in the data, follow a similar shape.

[43] This happens when a firm covaries strongly with GDP while having a small variance. From (67) the strong covariance implies that δ_i must be large, but then a negative Σ^v_{si} might be required to match the variance through (66).

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

Figure 10: Data distributions that the calibration matches exactly

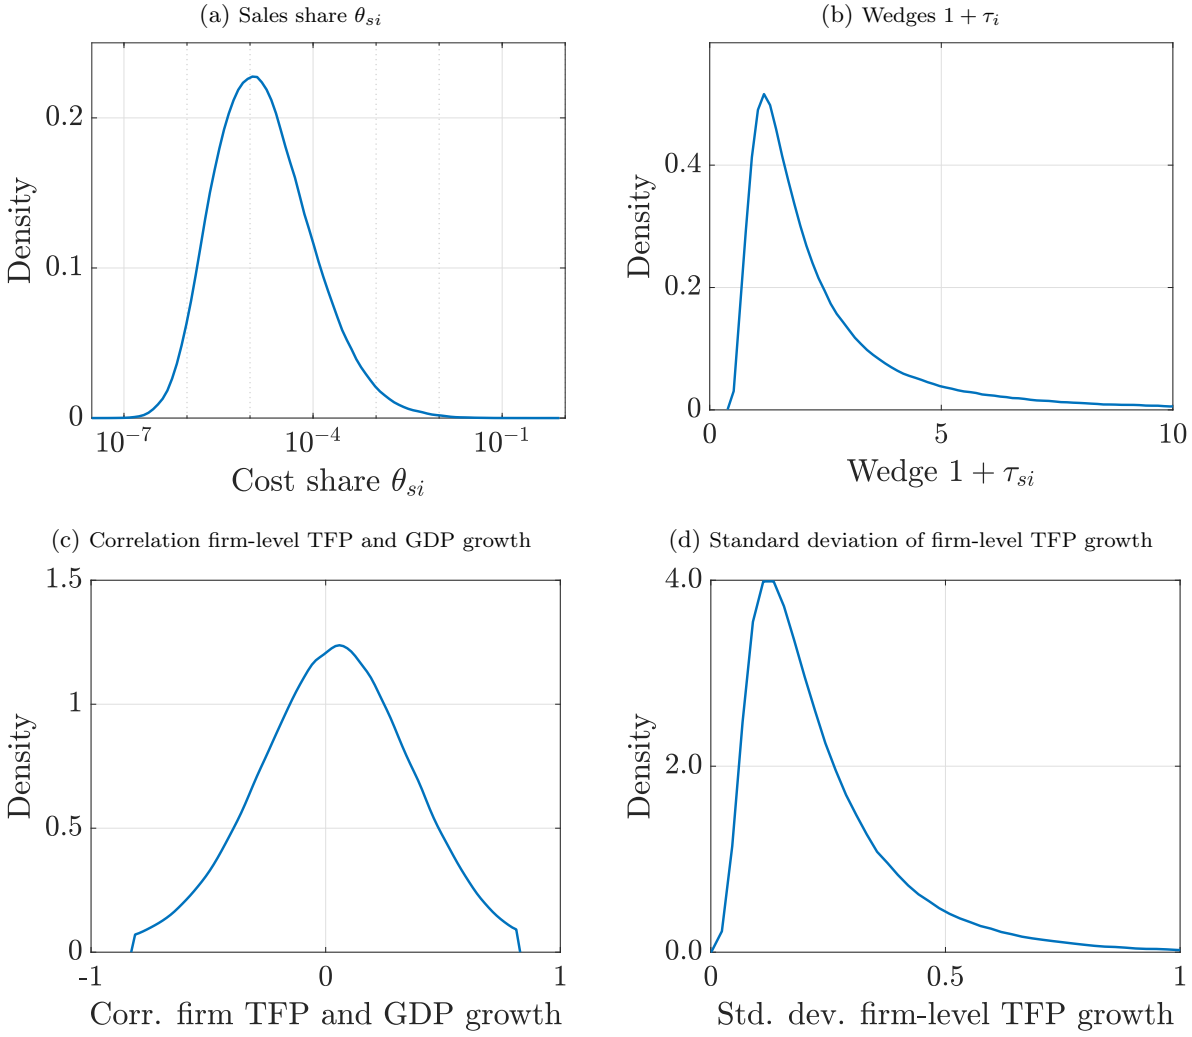


Figure 11: Domar weights of the firms in the data and in the model

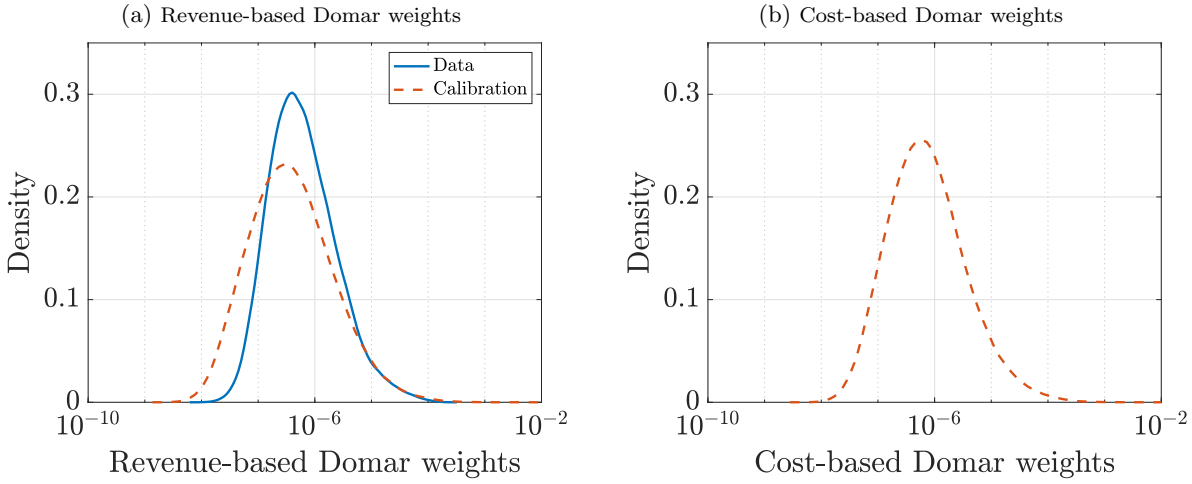
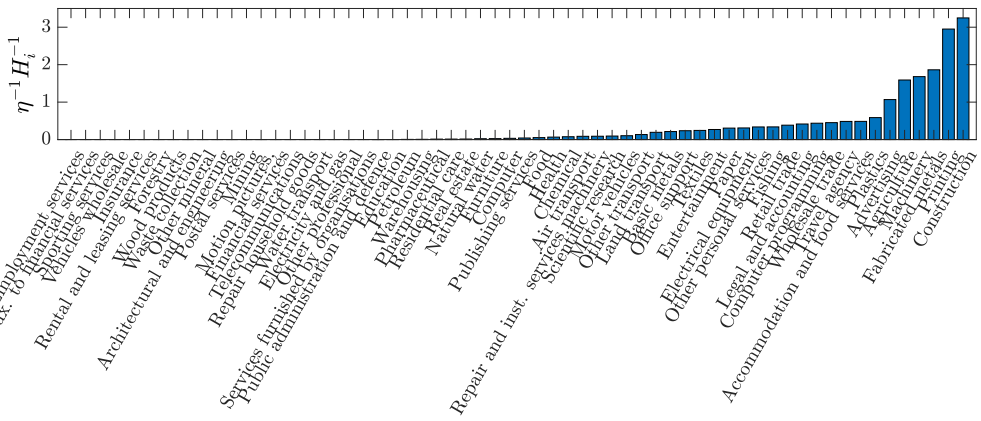


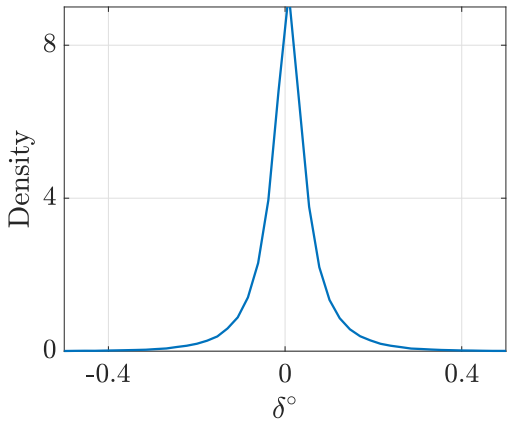
Figure 12 shows the estimated value of $\frac{1}{\eta}H_i^{-1}$ for each sector. The sectors to the left of "Electricity and gas" have an estimated value of $\frac{1}{\eta}H_i^{-1}$ of zero, implying that their risk-taking behavior is rigid. Figure 13 reports the distribution of the estimated firm-level natural risk exposure δ_i°.

Figure 12: Estimated value of $\frac{1}{\eta}H_i^{-1}$ for each sector.



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

Figure 13: 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.

C.5 Robustness

In this appendix, we provide several robustness exercises in our calibrated economy.

C.5.1 Constant risk-exposure technology

In the baseline calibration, we pick δ° and H_i^{-1} in the risk-exposure equation (50) by assuming that H_i^{-1} is, in each sector, the sum of a constant and a power function of $\tilde{\omega}_i$ (see footnote (32)).

In this appendix, we assume instead that H_i^{-1} is constant across firms in a sector. The outcome of the counterfactual exercises of Section 8 are presented in Figure 6. As we can see, they do not vary much compared to the baseline results of Table 1.

Table 6: Exposure and GDP volatility in different environments

	Calibration	Doubling Σ		No wedges	
		Fixed δ	Flexible δ	Fixed δ	Flexible δ
Aggregate risk exposure Δ	0.014	0.014	0.013	0.014	0.009
Exposure value $\mathcal{E}$	-0.06	-0.11	-0.10	-0.06	-0.04
Std. Dev. of GDP growth	2.4%	3.1%	2.9%	2.4%	1.9%
Share of endogenous vol.	69%	81%	79%	67%	51%

Notes: Share of endogenous volatility is $2\Sigma\Delta^2/\left(\mathrm{V}\left[\log \mathrm{GDP}_t - \log \mathrm{GDP}_{t-1}\right]\right)$.

D Proofs

D.1 Proof of Lemma 1

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

$$y\left(\delta\right) = \Delta^{\top}\varepsilon - \tilde{\omega}^{\top}\log\left(1+\tau\right) - \log\Gamma_L, \tag{16}$$

where the labor share of income Γ_L is given by

$$\Gamma_L := \frac{W_L L}{\bar{P}Y} = 1 - \tau^{\top}\left(diag\left(1+\tau\right)\right)^{-1}\omega. \tag{17}$$

Proof. Total profit in this economy is

$$\Pi = \sum_{i=1}^{N}\Pi_i = \sum_{i=1}^{N}\tau_i K_i Q_i - W_R\sum_{i=1}^{N}\kappa_i\left(\delta_i\right) = \sum_{i=1}^{N}\frac{\tau_i}{1+\tau_i}\omega_i P^{\top}C - W_R R,$$

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

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

Since $\Gamma_L + \Gamma_{\Pi} = 1$, (17) follows. Next, from the definition of the labor share, we can write $P^{\top}C = \Gamma_L^{-1}W_L L$. Taking the log and combining with (15) and (55), we find

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

Since $W_L L = 1$ given our choice of numeraire, we find (16).

Finally, differentiating (17) with respect to τ_i together with (14) yields

$$\frac{d\Gamma_L}{d\tau_i} = -\frac{\omega_i}{1+\tau_i}\left(1 - \tau^\top \left(\operatorname{diag}(1+\tau)\right)^{-1}\left(I - \alpha^\top \left[\operatorname{diag}(1+\tau)\right]^{-1}\right)^{-1}\mathbf{1}_i\right).$$

Notice that the expression in parentheses is the labor share in an economy with $\beta = \mathbf{1}_i$, i.e.,

$$1 - \tau^\top \left(\operatorname{diag}(1+\tau)\right)^{-1}\left(I - \alpha^\top \left[\operatorname{diag}(1+\tau)\right]^{-1}\right)^{-1}\mathbf{1}_i = \Gamma_L|_{\beta=\mathbf{1}_i} > 0.$$

Therefore, $\frac{d\Gamma_L}{d\tau_i} < 0$.

$$\frac{d\Gamma_L}{d\tau_i} = -\tilde{\omega}_i\left(1 - \tau^\top \left(\operatorname{diag}(1+\tau)\right)^{-1}\left(I - \alpha^\top \left[\operatorname{diag}(1+\tau)\right]^{-1}\right)^{-1}\mathbf{1}_i\right).$$

$$\begin{aligned}
\frac{dy(\delta)}{d\tau_i} &= -\tilde{\omega}_i\frac{1}{1+\tau_i} - \frac{1}{\Gamma_L}\frac{d\Gamma_L}{d\tau_i} \\
&= \frac{1}{1+\tau_i}\left(-\tilde{\omega}_i + \frac{1}{\Gamma_L}\omega_i\left(1 - \tau^\top \left(\operatorname{diag}(1+\tau)\right)^{-1}\left(I - \alpha^\top \left[\operatorname{diag}(1+\tau)\right]^{-1}\right)^{-1}\mathbf{1}_i\right)\right) \\
&= \frac{1}{1+\tau_i}\left(-\tilde{\omega}_i + \frac{1}{\Gamma_L}\omega_i - \frac{1}{\Gamma_L}\omega_i\tau^\top \left(\operatorname{diag}(1+\tau)\right)^{-1}\left(I - \alpha^\top \left[\operatorname{diag}(1+\tau)\right]^{-1}\right)^{-1}\mathbf{1}_i\right)
\end{aligned}$$

□

D.2 Proof of Lemma 2

Lemma 2. *The aggregate cost function $\bar{\kappa}$ is given by*

$$\bar{\kappa}(\Delta) = \frac{1}{2}(\Delta - \Delta^\circ)^\top \nabla^2\bar{\kappa}(\Delta - \Delta^\circ), \tag{24}$$

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

$$\nabla^2\bar{\kappa} = \eta\left(\sum_{i=1}^N \frac{\tilde{\omega}_i^2}{g_i}H_i^{-1}\right)^{-1}. \tag{25}$$

Proof. The Lagrangian of problem (22) is

$$\mathscr{L} = \eta\sum_{i=1}^N g_i\kappa_i(\delta_i) - \sum_{m=1}^M \nu_m\left(\Delta_m - \mathbf{1}_m^\top\delta^\top\tilde{\omega}\right), \tag{70}$$

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

$$\eta g_i \nabla \kappa_i = -\nu \tilde{\omega}_i, \tag{71}$$

which implies that for all i, k,

$$\eta \frac{g_i}{\tilde{\omega}_i} \nabla \kappa_i = \eta \frac{g_k}{\tilde{\omega}_k} \nabla \kappa_k = \nabla \bar{\kappa}, \tag{72}$$

where the last equality comes from the envelope theorem. Using (3), we can write

$$\frac{g_i}{\tilde{\omega}_i} H_i \left(\delta_i - \delta_i^\circ\right) = \frac{g_k}{\tilde{\omega}_k} H_k \left(\delta_k - \delta_k^\circ\right) \Leftrightarrow \delta_i = \delta_i^\circ + \frac{\tilde{\omega}_i}{g_i} \frac{g_k}{\tilde{\omega}_k} H_i^{-1} H_k \left(\delta_k - \delta_k^\circ\right).$$

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

$$\Delta = \sum_{j=1}^N \tilde{\omega}_j \delta_j = \sum_{j=1}^N \tilde{\omega}_j \left(\delta_j^\circ + \frac{\tilde{\omega}_j}{g_j} \frac{g_k}{\tilde{\omega}_k} H_j^{-1} H_k \left(\delta_k - \delta_k^\circ\right)\right) = \Delta^\circ + \left(\sum_{j=1}^N \frac{\tilde{\omega}_j^2}{g_j} H_j^{-1}\right) \frac{g_k}{\tilde{\omega}_k} H_k \left(\delta_k - \delta_k^\circ\right) \Leftrightarrow$$

$$\delta_k - \delta_k^\circ = \frac{1}{\eta} \frac{\tilde{\omega}_k}{g_k} H_k^{-1} \nabla^2 \bar{\kappa} \left(\Delta - \Delta^\circ\right), \tag{73}$$

where $\nabla^2 \bar{\kappa}$ is given by (25). Combining this last expression with (22) yields the result. □

D.3 Proof of Proposition 1

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

$$\mathcal{W}_{dist} = \max_{\Delta} \underbrace{\Delta^\top \mu - \tilde{\omega}^\top \log\left(1+\tau\right) - \log \Gamma_L}_{\mathrm{E}[y]} - \frac{1}{2}\left(\rho - 1\right) \underbrace{\Delta^\top \Sigma \Delta}_{\mathrm{V}[y]} - \bar{\kappa}\left(\Delta\right). \tag{26}$$

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

$$\max_{\delta} \tilde{\omega}^\top \delta \mu - \tilde{\omega}^\top \log\left(1+\tau\right) - \log \Gamma_L - \frac{1}{2}\left(\rho - 1\right) \tilde{\omega}^\top \delta \Sigma \delta^\top \tilde{\omega} + \log\left(V\left(\sum_{i=1}^N \frac{\tilde{\omega}_i \left(1+\tau_i\right)}{\omega_i} \kappa_i\left(\delta_i\right)\right)\right). \tag{74}$$

Notice that we can write this problem as

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

which is the maximization problem (26). Note also that the objective function of this problem is

strictly concave if $\bar{\kappa}$ is strictly convex. But recall from Lemma 2 that $\bar{\kappa}$'s Hessian satisfies

$$\left(\nabla^2 \bar{\kappa}\right)^{-1} = \frac{1}{\eta} \sum_{i=1}^{N} \frac{\tilde{\omega}_i^2}{g_i} H_i^{-1}. \tag{75}$$

Since H_i is positive definite, so is H_i^{-1} and so is the right-hand side of this equation. It follows that $\nabla^2 \bar{\kappa}$ is also positive definite, and $\bar{\kappa}$ is therefore strictly convex. This implies that there is a unique solution to the maximization problems (74) and (26) 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

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

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

$$\tilde{\omega}_i \mu_m - (\rho - 1)\, \tilde{\omega}_i \tilde{\omega}^\top \delta \Sigma \mathbf{1}_m - \eta \frac{\tilde{\omega}_i \left(1+\tau_i\right)}{\omega_i} \frac{d\kappa_i\left(\delta_i\right)}{d\delta_{im}} = 0. \tag{76}$$

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 (11) is strictly convex, such that its first-order conditions are necessary and sufficient. They are given by

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

where we used the fact that $da_i/d\delta_i = \varepsilon$. 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 $K_i Q_i = \omega_i \Gamma_L^{-1} / \left(1+\tau_i\right)$, such that $K_i Q_i$ is deterministic. The first-order conditions become

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

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

Then, using (53), we can write the stochastic discount factor as

$$\Lambda = \mathcal{U}'\left(Y\right) \left(V\left(R\right)\right)^{1-\rho} \bar{P}^{-1} = \left(\Gamma_L^{-1} W_L L\right)^{-\rho} \left(V\left(R\right)\right)^{1-\rho} \bar{P}^{\rho-1}. \tag{77}$$

Taking the log, we get

$$\lambda = \log\left(\left(\Gamma_L^{-1} W_L L\right)^{-\rho} \left(V\left(R\right)\right)^{1-\rho}\right) - \left(\rho - 1\right) \beta^\top \tilde{\mathcal{L}} \left(\delta\varepsilon - \log\left(1 + \tau\right)\right), \tag{78}$$

where we have used the fact that $\bar{p} = \beta^\top p$ together with (15). The first-order condition therefore becomes

$$\mathrm{E}\left[\varepsilon\right] - \left(\rho - 1\right) \tilde{\omega} \Sigma \delta^\top = \frac{W_R}{\omega_i \Gamma_L^{-1} / \left(1 + \tau_i\right)}.$$

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 (7) that $\eta \Gamma_L^{-1} = W_R$, and so the first-order conditions of the firms are equivalent to those of the factious planner (27), which completes the proof. □

D.4 Proof of Lemma 3

Lemma 4. *The equilibrium aggregate risk exposure Δ solves*

$$\mathcal{E}\left(\Delta\right) = \nabla \bar{\kappa}\left(\Delta\right), \tag{27}$$

where the marginal value of aggregate risk exposure $\mathcal{E}$ is given by

$$\mathcal{E} := \mathrm{E}\left[\varepsilon\right] + \mathrm{Cov}\left[\lambda, \varepsilon\right], \tag{28}$$

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

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

$$\mu - \left(\rho - 1\right) \Delta \Sigma = \nabla \bar{\kappa}\left(\Delta\right).$$

Combining the definition of $\mathcal{E}$ given by (28) with (78), 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. □

D.5 Proof of Lemma 4

Lemma 4. *The equilibrium marginal value of exposure vector $\mathcal{E}$ can be written as*

$$\mathcal{E} = \mu - \left(\rho - 1\right) \Sigma \Delta. \tag{29}$$

Proof. This is an intermediary result that we proved in the proof of Lemma 3. □

D.6 Proof of Proposition 2

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

$$\delta_i = \delta_i^\circ + \frac{1}{\eta}\frac{\omega_i}{1+\tau_i} H_i^{-1}\mathcal{E} \qquad \textit{and} \qquad \Delta = \Delta^\circ + \mathcal{H}^{-1}\mathcal{E}^\circ, \tag{31}$$

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

$$\mathcal{H}^{-1} := \left(\nabla^2\bar{\kappa} + (\rho-1)\,\Sigma\right)^{-1}. \tag{32}$$

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

$$\Delta - \Delta^\circ = \left[\nabla^2\bar{\kappa}\right]^{-1}\mathcal{E}\left(\Delta\right).$$

It follows that

$$\begin{gathered}\mathcal{E} - \nabla^2\bar{\kappa}\left(\Delta - \Delta^\circ\right) = 0 \Leftrightarrow \Delta = \Delta^\circ + \left(\nabla^2\bar{\kappa}\right)^{-1}\left(\mu - (\rho-1)\,\Sigma\Delta\right) \Leftrightarrow \\ \Delta = \mathcal{H}^{-1}\left(\nabla^2\bar{\kappa}\Delta^\circ + \mu\right) = \Delta^\circ + \mathcal{H}^{-1}\mathcal{E}^\circ,\end{gathered}$$

where $\mathcal{E}^\circ = \mu - (\rho-1)\,\Sigma\Delta^\circ$. Combining $\mathcal{E} = \nabla^2\bar{\kappa}\left(\Delta - \Delta^\circ\right)$ with (73) yields the equation for δ. □

D.7 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*

$$\frac{d\Delta}{d\gamma} = \mathcal{H}^{-1}\frac{\partial\mathcal{E}}{\partial\gamma}, \tag{33}$$

where $\frac{\partial\mathcal{E}}{\partial\gamma}$ *is given by* (30).

Proof. From (31), we can write

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

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 (30). □

D.8 Proof of Corollary 1

Corollary 1. *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 (30). □

D.9 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*

$$\frac{d\Delta}{d\tau_i} = \frac{d\mathcal{H}^{-1}}{d\tau_i}\mathcal{E}^{\circ}, \tag{34}$$

where

$$\frac{d\mathcal{H}^{-1}}{d\tau_i} = -\mathcal{H}^{-1}\left(\sum_{j=1}^{N}\frac{\partial\nabla^2\bar{\kappa}}{\partial g_j}\frac{dg_j}{d\tau_i}\right)\mathcal{H}^{-1},$$

and where the response of the efficiency gap $dg_j/d\tau_i$ is given by (23), and $\partial\nabla^2\bar{\kappa}/\partial g_j$ is a positive definite matrix.

Proof. The first equation follows directly from (31). Differentiating $\mathcal{H}^{-1}$ with respect to τ_i, we find

$$\frac{d\mathcal{H}^{-1}}{d\tau_i} = -\mathcal{H}^{-1}\left(\sum_{j=1}^{N}\frac{\partial\nabla^2\bar{\kappa}}{\partial g_j}\frac{dg_j}{d\tau_i}\right)\mathcal{H}^{-1}.$$

Next, using (75), we find

$$\frac{\partial\nabla^2\bar{\kappa}}{\partial g_i} = \left(\nabla^2\bar{\kappa}\right)\left(\frac{1}{\eta}\frac{\tilde{\omega}_i^2}{g_i^2}H_i^{-1}\right)\left(\nabla^2\bar{\kappa}\right). \tag{79}$$

Since $\nabla^2\bar{\kappa}$ and H_i^{-1} are positive definite, so is $\frac{\partial\nabla^2\bar{\kappa}}{\partial g_i}$. □

D.10 Proof of Corollary 2

Corollary 2. *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,

$$\frac{d\Delta}{d\tau_i} = -\mathcal{H}^{-1}\left(\sum_{j=1}^{N}\frac{1}{\eta}\frac{\tilde{\omega}_j^2}{g_j^2}\nabla^2\bar{\kappa}H_j^{-1}\nabla^2\bar{\kappa}\frac{dg_j}{d\tau_i}\right)\mathcal{H}^{-1}\mathcal{E}^{\circ}.$$

Recall that $\frac{dg_j}{d\tau_i} > 0$ by (23). If Σ and H_i are diagonal for all i, then $\mathcal{H}^{-1}\left(\sum_{j=1}^{N}\frac{1}{\eta}\frac{\tilde{\omega}_j^2}{g_j^2}\nabla^2\bar{\kappa}H_j^{-1}\nabla^2\bar{\kappa}\frac{dg_j}{d\tau_i}\right)\mathcal{H}^{-1}$ 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

$$\mathcal{E} - \nabla^2\bar{\kappa}\left(\Delta - \Delta^\circ\right) = 0 \Leftrightarrow \mathcal{E} = \nabla^2\bar{\kappa}\mathcal{H}^{-1}\mathcal{E}^\circ. \tag{80}$$

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

D.11 Proof of Lemma 5

Lemma 5. *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.*

Proof. We first show that $\mathcal{H}_{SP}^{-1} - \mathcal{H}^{-1}$ is positive definite. For any two positive definite matrix 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 (32),

$$\begin{aligned}\mathcal{H} - \mathcal{H}_{SP} &= \left(\nabla^2\bar{\kappa} - \frac{d\mathcal{E}}{d\Delta}\right) - \left(\nabla^2\bar{\kappa}_{SP} - \left(\frac{d\mathcal{E}}{d\Delta}\right)_{SP}\right)\\ &= \nabla^2\bar{\kappa} - \nabla^2\bar{\kappa}_{SP},\end{aligned}$$

where we have used the fact that $\frac{d\mathcal{E}}{d\Delta} = -\left(\rho - 1\right)\Sigma$ is the same in the equilibrium and in the efficient allocation. It follows that $\mathcal{H}_{SP}^{-1} - \mathcal{H}^{-1}$ is positive definite if $\left(\nabla^2\bar{\kappa}_{SP}\right)^{-1} - \left(\nabla^2\bar{\kappa}\right)^{-1}$ is. From (25), we can write

$$\left(\nabla^2\bar{\kappa}_{SP}\right)^{-1} - \left(\nabla^2\bar{\kappa}\right)^{-1} = \frac{1}{\eta}\sum_{k=1}^{N}\tilde{\omega}_k^2\left(1 - \frac{1}{g_k}\right)H_k^{-1}.$$

Recall that $g_k \geq 1$, with the equality holding only if $\tau = 0$. Furthermore, H_k is positive definite. Then $\left(\nabla^2\bar{\kappa}_{SP}\right)^{-1} - \left(\nabla^2\bar{\kappa}\right)^{-1}$ is also 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. □

D.12 Proof of Corollary 3

Corollary 3. *Suppose that the economy is diagonal and that $\tau_j > 0$ for at least one firm j. Then the sign of $\Delta_i - \Delta_{SP,i}$ is the opposite of the sign of $\mathcal{E}_i^\circ$.*

Proof. Using Proposition 2, we can write

$$\Delta - \Delta^{SP} = \left(\mathcal{H}^{-1} - \mathcal{H}_{SP}^{-1}\right) \mathcal{E}^{\circ}.$$

From the proof of Lemma 5, we know that $\mathcal{H}_{SP}^{-1} - \mathcal{H}^{-1}$ is positive definite. Furthermore, if the risk factors are uncorrelated (diagonal Σ), and that individual risk exposures are neither complements nor substitutes in the cost functions $(\kappa_1, \ldots, \kappa_N)$ (diagonal H_i for all i), then $\mathcal{H}$ and $\mathcal{H}_{SP}$ are diagonal matrices. Therefore, the sign $\Delta_i - \Delta_i^{SP}$ is the opposite of the sign of $\mathcal{E}_i^{\circ}$, which is the result. □

D.13 Proof of Corollary 4

Corollary 4. *In a diagonal economy, the following statements hold.*

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

$$sign\left(\frac{d\,\mathrm{E}\,[y]}{d\mu_m} - \frac{\partial\,\mathrm{E}\,[y]}{\partial\mu_m}\right) = sign\left(\mu_m\right) \quad and \quad sign\left(\frac{d\,\mathrm{V}\,[y]}{d\mu_m} - \frac{\partial\,\mathrm{V}\,[y]}{\partial\mu_m}\right) = sign\left(\Delta_m\right). \tag{36}$$

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

$$sign\left(\frac{d\,\mathrm{E}\,[y]}{d\Sigma_{mm}} - \frac{\partial\,\mathrm{E}\,[y]}{\partial\Sigma_{mm}}\right) = -sign\left(\mu_m\Delta_m\right) \quad and \quad \frac{d\,\mathrm{V}\,[y]}{d\Sigma_{mm}} - \frac{\partial\,\mathrm{V}\,[y]}{\partial\Sigma_{mm}} < 0. \tag{37}$$

Proof. From (25), we see that given our assumptions, $\nabla^2\bar{\kappa}$ is diagonal with positive entries. This implies that $\mathcal{H}^{-1}$ is also diagonal with positive entries. Combining the first equation in (35) with (33) and (30), we can write

$$\operatorname{sign}\left(\frac{d\,\mathrm{E}\,[y]}{d\mu_m} - \frac{\partial\,\mathrm{E}\,[y]}{\partial\mu_m}\right) = \operatorname{sign}\left(\mu^{\top}\mathcal{H}^{-1}\mathbf{1}_m\right) = \operatorname{sign}\left(\mu_m\right).$$

Combining the second expression in (35) with (33) and (30), we can write

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

We can follow the same procedure for Σ_{mm}. From the first equation in (35) together with (33) and (30), we find

$$\operatorname{sign}\left(\frac{d\,\mathrm{E}\,[y]}{d\Sigma_{mm}} - \frac{\partial\,\mathrm{E}\,[y]}{\partial\Sigma_{mm}}\right) = \operatorname{sign}\left(-\mu^{\top}\mathcal{H}^{-1}\left(\rho - 1\right)\Delta_m\mathbf{1}_m\right) = \operatorname{sign}\left(-\mu_m\Delta_m\right).$$

From the second expression in (35) together with (33) and (30), we can write

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

which is always negative given our assumptions. □

D.14 Proof of Corollary 5

Corollary 5. *Suppose that there is a single risk factor. Then*

$$sign\left(\frac{d\,\mathrm{E}\,[y]}{d\tau_i} - \frac{\partial\,\mathrm{E}\,[y]}{\partial\tau_i}\right) = -sign\,(\mu\mathcal{E}) \qquad and \qquad sign\left(\frac{d\,\mathrm{V}\,[y]}{d\tau_i} - \frac{\partial\,\mathrm{V}\,[y]}{\partial\tau_i}\right) = -sign\,(\Delta\mathcal{E})\,. \quad (38)$$

Proof. When there is only one risk factor, we have$\mathcal{H}^{-1} = \left(\nabla^2\bar{\kappa} + (\rho-1)\,\Sigma\right)^{-1} > 0$, where the inequality follows since $\nabla^2\bar{\kappa}$ and Σ are positive scalars given our assumptions. Furthermore, $\frac{\partial\nabla^2\bar{\kappa}}{\partial g_i}$, given by (79), is a positive scalar. Combining the first expression in (35) with (34) and (80), we can write

$$\text{sign}\left(\frac{d\,\mathrm{E}\,[y]}{d\tau_i} - \frac{\partial\,\mathrm{E}\,[y]}{\partial\tau_i}\right) = \text{sign}\left(-\mu\mathcal{H}^{-1}\left(\sum_{j=1}^{N}\frac{\partial\nabla^2\bar{\kappa}}{\partial g_j}\frac{dg_j}{d\tau_i}\right)\left(\nabla^2\bar{\kappa}\right)^{-1}\mathcal{E}\right),$$

and the first result follows since $\frac{dg_j}{d\tau_i} > 0$ by (23). Combining the second expression in (35) with (34), we can write

$$\text{sign}\left(\frac{d\,\mathrm{V}\,[y]}{d\tau_i} - \frac{\partial\,\mathrm{V}\,[y]}{\partial\tau_i}\right) = \text{sign}\left(-2\Delta\Sigma\mathcal{H}^{-1}\left(\sum_{j=1}^{N}\frac{\partial\nabla^2\bar{\kappa}}{\partial g_j}\frac{dg_j}{d\tau_i}\right)\left(\nabla^2\bar{\kappa}\right)^{-1}\mathcal{E}\right),$$

and the second result follows. □

D.15 Proof of Proposition 7

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*

$$\frac{d\mathcal{W}}{d\chi} - \frac{\partial\mathcal{W}}{\partial\chi} = \left(\mathcal{E} - \nabla\bar{\kappa}_V\right)^{\top}\frac{d\Delta}{d\chi} = \left(\nabla\bar{\kappa} - \nabla\bar{\kappa}_V\right)^{\top}\frac{d\Delta}{d\chi}, \quad (40)$$

where partial derivatives indicates that Δ is kept fixed, and where $d\Delta/d\chi$ is given by (33) for $\chi = \mu_i$ or Σ_{mn}, and by (34) for $\chi = \tau_i$.

Proof. Combining the definition of $\bar{\kappa}_V$ with (31), we can write

$$\bar{\kappa}_V = \frac{1}{2}\left(\Delta - \Delta^{\circ}\right)^{\top}\nabla^2\bar{\kappa}\left(\sum_i \frac{1}{\eta}\left(\frac{\tilde{\omega}_i}{g_i}\right)^2 H_i^{-1}\right)\nabla^2\bar{\kappa}\left(\Delta - \Delta^{\circ}\right).$$

This equation implies that

$$\nabla\bar{\kappa}_V = \nabla^2\bar{\kappa}_V\left(\Delta - \Delta^\circ\right),$$

where

$$\nabla^2\bar{\kappa}_V = \nabla^2\bar{\kappa}\left(\sum_i \frac{1}{\eta}\left(\frac{\tilde{\omega}_i}{g_i}\right)^2 H_i^{-1}\right)\nabla^2\bar{\kappa}.$$

Simple algebra implies that

$$\nabla\bar{\kappa} - \nabla\bar{\kappa}_V = \nabla^2\bar{\kappa}\left(\frac{1}{\eta}\sum_{i=1}^N \frac{\tilde{\omega}_i^2}{g_i}\left(1 - \frac{1}{g_i}\right)H_i^{-1}\right)\nabla^2\bar{\kappa}\left(\Delta - \Delta^\circ\right). \tag{81}$$

We can write welfare as

$$\mathcal{W} = \mathrm{E}\left[y\right] - \frac{1}{2}\left(\rho - 1\right)\mathrm{V}\left[y\right] - \bar{\kappa}_V.$$

Differentiating it with respect to χ yields

$$\frac{d\mathcal{W}}{d\chi} = \frac{\partial\,\mathrm{E}\left[y\right]}{\partial\chi} + \left(\frac{d\Delta}{d\chi}\right)^\top \frac{d\,\mathrm{E}\left[y\right]}{d\Delta} - \frac{1}{2}\left(\rho - 1\right)\left(\frac{\partial\,\mathrm{V}\left[y\right]}{\partial\chi} + \left(\frac{d\Delta}{d\chi}\right)^\top \frac{d\,\mathrm{V}\left[y\right]}{d\Delta}\right) - \left(\frac{\partial\bar{\kappa}_V}{\partial\chi} + \left(\frac{d\Delta}{d\chi}\right)^\top \frac{d\bar{\kappa}_V}{d\Delta}\right)$$

or

$$\begin{aligned}\frac{d\mathcal{W}}{d\chi} - \frac{\partial\mathcal{W}}{\partial\chi} &= \left(\frac{d\Delta}{d\chi}\right)^\top \frac{d\,\mathrm{E}\left[y\right]}{d\Delta} - \frac{1}{2}\left(\rho - 1\right)\left(\frac{d\Delta}{d\chi}\right)^\top \frac{d\,\mathrm{V}\left[y\right]}{d\Delta} - \left(\frac{d\Delta}{d\chi}\right)^\top \frac{d\bar{\kappa}_V}{d\Delta} \\ &= \left(\frac{d\Delta}{dx}\right)^\top \left(\nabla\bar{\kappa} - \nabla\bar{\kappa}_V\right),\end{aligned}$$

where we have used the definition of $\mathcal{E}$ given by (29) and the first-order condition (27). □

D.16 Proof of Lemma 6

Lemma 6. *Suppose that the economy is diagonal and that $\tau_j > 0$ for at least one firm j. Then the sign of $\left[\nabla\bar{\kappa}\right]_i - \left[\nabla\bar{\kappa}_V\right]_i$ is the same as the sign of $\mathcal{E}_i$.*

Proof. Combining (81) with (80), we can write

$$\nabla\bar{\kappa} - \nabla\bar{\kappa}_V = \nabla^2\bar{\kappa}\left(\frac{1}{\eta}\sum_{i=1}^N \frac{\tilde{\omega}_i^2}{g_i}\left(1 - \frac{1}{g_i}\right)H_i^{-1}\right)\mathcal{E}. \tag{82}$$

Under our assumptions, $\nabla^2\bar{\kappa}\left(\frac{1}{\eta}\sum_{i=1}^N \frac{\tilde{\omega}_i^2}{g_i}\left(1 - \frac{1}{g_i}\right)H_i^{-1}\right)$ is a positive definite diagonal matrix. Therefore, the sign of $\left[\nabla\bar{\kappa}\right]_i - \left[\nabla\bar{\kappa}_V\right]_i$ is the same as the sign of $\mathcal{E}_i$. □

D.17 Proof of Corollary 6

Corollary 6. *Suppose that the economy is diagonal and that $\tau_j > 0$ for at least one firm j. Then the following holds.*

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

Proof. Using (81) and (80), we can write (40) as

$$\frac{d\mathcal{W}}{d\chi} - \frac{\partial \mathcal{W}}{\partial \chi} = \left(\frac{d\Delta}{d\chi}\right)^\top \nabla^2 \bar{\kappa} \left(\frac{1}{\eta} \sum_{i=1}^N \frac{\tilde{\omega}_i^2}{g_i} \left(1 - \frac{1}{g_i}\right) H_i^{-1}\right) \mathcal{E}.$$

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

$$\frac{d\mathcal{W}}{d\gamma} - \frac{\partial \mathcal{W}}{\partial \gamma} = \left(\frac{\partial \mathcal{E}}{\partial \gamma}\right)^\top \mathcal{H}^{-1} \nabla^2 \bar{\kappa} \left(\frac{1}{\eta} \sum_{j=1}^N \frac{\tilde{\omega}_j^2}{g_j} \left(1 - \frac{1}{g_j}\right) H_j^{-1}\right) \mathcal{E}.$$

If $\gamma = \mu_m$, this equation becomes

$$\frac{d\mathcal{W}}{d\mu_m} - \frac{\partial \mathcal{W}}{\partial \mu_m} = \mathbf{1}_m^\top \mathcal{H}^{-1} \nabla^2 \bar{\kappa} \left(\frac{1}{\eta} \sum_{j=1}^N \frac{\tilde{\omega}_j^2}{g_j} \left(1 - \frac{1}{g_j}\right) H_j^{-1}\right) \mathcal{E}.$$

Given our assumptions, the matrix $\mathcal{H}^{-1} \nabla^2 \bar{\kappa} \left(\frac{1}{\eta} \sum_{j=1}^N \frac{\tilde{\omega}_j^2}{g_j} \left(1 - \frac{1}{g_j}\right) H_j^{-1}\right)$ 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

$$\frac{d\mathcal{W}}{d\Sigma_{mm}} - \frac{\partial \mathcal{W}}{\partial \Sigma_{mm}} = -\left(\rho - 1\right) \Delta_m \mathbf{1}_m^\top \mathcal{H}^{-1} \nabla^2 \bar{\kappa} \left(\frac{1}{\eta} \sum_{j=1}^N \frac{\tilde{\omega}_j^2}{g_j} \left(1 - \frac{1}{g_j}\right) H_j^{-1}\right) \mathcal{E},$$

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 $-\Delta_m \mathcal{E}_m$. □

D.18 Proof of Corollary 7

Corollary 7. *Suppose that the economy is diagonal and that $\tau_j > 0$ for at least one firm j. Then 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 (81) and (80), we can write (40) as

$$\frac{d\mathcal{W}}{d\chi} - \frac{\partial \mathcal{W}}{\partial \chi} = \left(\frac{d\Delta}{d\chi}\right)^{\top} \nabla^2 \bar{\kappa} \left(\frac{1}{\eta} \sum_{i=1}^{N} \frac{\tilde{\omega}_i^2}{g_i} \left(1 - \frac{1}{g_i}\right) H_i^{-1} \right) \mathcal{E}.$$

From Proposition 4 and (80), we have

$$\frac{d\Delta}{d\tau_i} = -\mathcal{H}^{-1} \nabla^2 \bar{\kappa} \left(\sum_{j=1}^{N} \frac{1}{\eta} \frac{\tilde{\omega}_j^2}{g_j^2} H_j^{-1} \frac{dg_j}{d\tau_i} \right) \mathcal{E}.$$

Setting $\chi = \tau_i$ and combining the last two equations, we find

$$\frac{d\mathcal{W}}{d\tau_i} - \frac{\partial \mathcal{W}}{\partial \tau_i} = -\left(\Delta - \Delta^{\circ}\right)^{\top} \nabla^2 \bar{\kappa} \left(\sum_{j=1}^{N} \frac{1}{\eta} \frac{\tilde{\omega}_j^2}{g_j^2} H_j^{-1} \frac{dg_j}{d\tau_i} \right)^{\top} \nabla^2 \bar{\kappa} \mathcal{H}^{-1} \nabla^2 \bar{\kappa} \left(\frac{1}{\eta} \sum_{j=1}^{N} \frac{\tilde{\omega}_j^2}{g_j} \left(1 - \frac{1}{g_j}\right) H_j^{-1} \right) \mathcal{E}.$$

Using (80)again, this equation becomes

$$\frac{d\mathcal{W}}{d\tau_i} - \frac{\partial \mathcal{W}}{\partial \tau_i} = -\mathcal{E}^{\top} \left(\sum_{j=1}^{N} \frac{1}{\eta} \frac{\tilde{\omega}_j^2}{g_j^2} H_j^{-1} \frac{dg_j}{d\tau_i} \right)^{\top} \nabla^2 \bar{\kappa} \mathcal{H}^{-1} \nabla^2 \bar{\kappa} \left(\frac{1}{\eta} \sum_{j=1}^{N} \frac{\tilde{\omega}_j^2}{g_j} \left(1 - \frac{1}{g_j}\right) H_j^{-1} \right) \mathcal{E}.$$

Under our assumptions, the matrix between $\mathcal{E}^{\top}$and $\mathcal{E}$ is diagonal with positive elements, and the result follows. □

E 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 that model can be extended in different ways.

E.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 firm's 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*

$$\frac{d\Delta}{d\alpha_{ij}} = \mathcal{H}^{-1}\left[\nabla^2\bar{\kappa}\right]\sum_{k=1}^{N}\left(\frac{d\Delta^\circ}{d\tilde{\omega}_k} + \frac{d\left[\nabla^2\bar{\kappa}\right]^{-1}}{d\tilde{\omega}_k}\mathcal{E}\right)\frac{d\tilde{\omega}_k}{d\alpha_{ij}}, \tag{83}$$

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* $\bar{\kappa}$ *to a change in Domar weight is given by* $\frac{d\left[\nabla^2\bar{\kappa}\right]^{-1}}{d\tilde{\omega}_k} = \frac{2}{\eta}\tilde{\omega}_k H_k^{-1}$.

Proof. Using (31) and (80), we can write

$$\Delta - \Delta^\circ = \left[\nabla^2\bar{\kappa}\right]^{-1}\mathcal{E}\left(\Delta\right).$$

Using the implicit function theorem, we get

$$\frac{d\Delta}{d\alpha_{ij}} = \mathcal{H}^{-1}\left[\nabla^2\bar{\kappa}\right]\left(\frac{d\Delta^\circ}{d\alpha_{ij}} + \frac{d\left[\nabla^2\bar{\kappa}\right]^{-1}}{d\alpha_{ij}}\mathcal{E}\right). \tag{84}$$

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 (75), $\frac{d\left[\nabla^2\bar{\kappa}\right]^{-1}}{d\alpha_{ij}} = \sum_{k=1}^{N}\frac{d\tilde{\omega}_k}{d\alpha_{ij}}\frac{d\left[\nabla^2\bar{\kappa}\right]^{-1}}{d\tilde{\omega}_k} = \sum_{k=1}^{N}\frac{d\tilde{\omega}_k}{d\alpha_{ij}}\frac{2}{\eta}\tilde{\omega}_k H_k^{-1}$. By definition, $\Delta^\circ = \sum_{j=1}^{N}\delta_j^\circ\tilde{\omega}_j$ and $\tilde{\omega}_k = \beta^\top\tilde{\mathcal{L}}\mathbf{1}_k = \beta^\top\left(I-\alpha\right)^{-1}\mathbf{1}_k$. 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 (83) follows from (84).

If $\tau \neq 0$, the only difference is that

$$\frac{d\left[\nabla^2\bar{\kappa}\right]^{-1}}{d\alpha_{ij}} = -\frac{1}{\eta}\sum_{k=1}^{N}\frac{\tilde{\omega}_k^2}{g_k^2}\frac{dg_k}{d\alpha_{ij}}H_k^{-1} + \frac{2}{\eta}\sum_{i=1}^{N}\frac{d\tilde{\omega}_k}{d\alpha_{ij}}\frac{\tilde{\omega}_k}{g_k}H_k^{-1},$$

where

$$\frac{dg_k}{d\alpha_{ij}} = \frac{\left(1+\tau_k\right)}{\omega_k}\frac{d\tilde{\omega}_k}{d\alpha_{ij}} - \frac{\tilde{\omega}_k\left(1+\tau_k\right)}{\omega_k^2}\frac{d\omega_k}{d\alpha_{ij}},$$

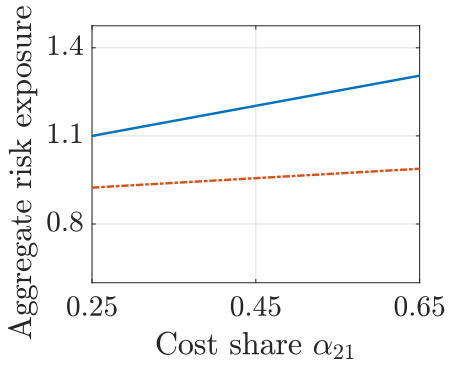
and $\frac{d\omega_k}{d\alpha_{ij}} = \frac{\omega_i}{1+\tau_i}\mathcal{L}_{jk}$. Therefore, $\frac{d\left[\nabla^2\bar{\kappa}\right]^{-1}}{d\alpha_{ij}} = \frac{1}{\eta}\sum_{k=1}^{N}\frac{\tilde{\omega}_k}{g_k}\left(\frac{d\omega_k}{d\alpha_{ij}}\frac{\tilde{\omega}_k}{\omega_k} + \frac{d\tilde{\omega}_k}{d\alpha_{ij}}\right)H_k^{-1}$. □

Equation (83) 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 (83)). 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 $\frac{\partial\left[\nabla^2\bar{\kappa}\right]^{-1}}{\partial\tilde{\omega}_k} = \frac{2}{\eta}\tilde{\omega}_i H_i^{-1}$ is a positive definite matrix. Through that channel, exposure to good risks tends to increase, while exposure to bad risk 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 $\bar{\kappa}$ and in the natural aggregate risk exposure Δ° are propagated through the matrix $\mathcal{T}$ to shape the response of the aggregate risk exposure Δ to a change in α_{ij}.

Figure 14 shows the impact of increasing the importance of good 1 in the production of good 2 in the example economy of Figure 2. 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.

Figure 14: Changes in network in the example economy



Notes. The structure of the economy is given in panel (a) of Figure 2. 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: $\delta_{11}^\circ = \delta_{32}^\circ = 1$, $\delta_{22}^\circ = 1.9$, $\delta_{12}^\circ = \delta_{21}^\circ = \delta_{31}^\circ = 0$, $H_1 = H_3$ are diagonal with very large entries on the main diagonals; $H_{2,11} = H_{2,22} = 1$, $H_{2,12} = H_{2,21} = 0.75$. α_{21} changes from 0.25 to 0.65.

E.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 $\mathcal{V}(R) = \exp\left((\rho-1)\sum_{k=1}^K \eta_k R_k\right)$.

The household's first-order conditions take the same form as before. Specifically, following the same steps as in Appendix A.1, 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 $R_{ik} = \kappa_{i,k}(\delta_i)$ units of resource k. As in the baseline model, we assume that $\kappa_{i,k}(\delta_i)$ is quadratic with $\kappa_{i,k}(\delta_i) = \frac{1}{2}(\delta_i - \delta_i^\circ)^\top H_{i,k}(\delta_i - \delta_i^\circ)$, where $H_{i,k}$ 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 2.3). When choosing its risk exposure, firm i then solves

$$\delta_i^* \in \arg\max_{\delta_i \in \mathcal{A}_i} \mathrm{E}\left[\Lambda\left[Q_i(P_i - K_i(\delta_i, P)) - \sum_{k=1}^{K}\kappa_{i,k}(\delta_i) W_{R,k}\right]\right].$$

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

$$\delta_i - \delta_i^\circ = \frac{\omega_i \Gamma_L^{-1}}{1+\tau_i}\left(\sum_{k=1}^{K} W_{R,k} H_i^k\right)^{-1} \mathcal{E} = \frac{\omega_i}{1+\tau_i}\left(\sum_{k=1}^{K}\eta_k H_i^k\right)^{-1}\mathcal{E}. \tag{85}$$

The key difference between this extension and the model with a single risk management resource is that $\sum_{k=1}^{K}\eta_k H_i^k$ replaces H_i^{-1}. Since both objects are constant, we can think of the Hessian matrix H_i from the main text has capturing the aggregated substitution patterns generated by the resource-specific Hessians H_i^k.

We can again characterize the equilibrium using a distorted planner's problem. Consider the maximization problem

$$\max_{\delta} \tilde{\omega}^\top \delta\mu - \tilde{\omega}^\top \log(1+\tau) - \log\Gamma_L - \frac{1}{2}(\rho-1)\tilde{\omega}^\top \delta\Sigma\delta^\top\tilde{\omega} - \sum_{i=1}^{N}\sum_{k=1}^{K}\eta_k g_i \kappa_{i,k}(\delta_i), \tag{86}$$

where $g_i = \frac{\tilde{\omega}_i(1+\tau_i)}{\omega_i}$. Taking first-order conditions with respect to δ_i, we get

$$\tilde{\omega}_i\mu - (\rho-1)\,\tilde{\omega}_i\tilde{\omega}^\top\delta\Sigma - \sum_{k=1}^{K}\eta_k\frac{\tilde{\omega}_i\,(1+\tau_i)}{\omega_i}H_{i,k}\,(\delta_i-\delta_i^\circ) = 0 \Leftrightarrow$$

$$\delta_i - \delta_i^\circ = \frac{\omega_i}{1+\tau_i}\left(\sum_{k=1}^{K}\eta_k H_i^k\right)^{-1}\left(\mu - (\rho-1)\,\tilde{\omega}^\top\delta\Sigma\right),$$

which is equivalent to (85). It follows that any equilibrium must coincide with a solution to the optimization problem (86).

We can write (86). as

$$\max_{\Delta}\Delta^\top\mu - \tilde{\omega}^\top\log(1+\tau) - \log\Gamma_L - \frac{1}{2}(\rho-1)\,\Delta^\top\Sigma\Delta - \min_{\delta\text{ s.t. }\Delta=\delta^\top\tilde{\omega}}\sum_{i=1}^{N}\sum_{k=1}^{K}\eta_k g_i\kappa_{i,k}\,(\delta_i),$$

or, equivalently, as

$$\mathcal{W}_{dist} := \max_{\Delta}\underbrace{\Delta^\top\mu - \tilde{\omega}^\top\log(1+\tau) - \log\Gamma_L}_{\mathrm{E}[y]} - \frac{1}{2}(\rho-1)\underbrace{\Delta^\top\Sigma\Delta}_{\mathrm{V}[y]} - \bar{\kappa}\,(\Delta), \tag{87}$$

where

$$\bar{\kappa}\,(\Delta) := \min_{\delta}\sum_{i=1}^{N}\sum_{k=1}^{K}\eta_k g_i\kappa_{i,k}\,(\delta_i), \tag{88}$$

subject to $\Delta = \delta^\top\tilde{\omega}$.

As in the baseline model, we can solve for $\bar{\kappa}\,(\Delta)$.

Lemma 7. *The aggregate cost function $\bar{\kappa}$ is given by*

$$\bar{\kappa}\,(\Delta) = \frac{1}{2}(\Delta-\Delta^\circ)^\top\nabla^2\bar{\kappa}\,(\Delta-\Delta^\circ), \tag{89}$$

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

$$\nabla^2\bar{\kappa} = \left(\sum_{i=1}^{N}\frac{\tilde{\omega}_i^2}{g_i}\left(\sum_{k=1}^{K}\eta_k H_i^k\right)^{-1}\right)^{-1}. \tag{90}$$

Proof. The Lagrangian of problem (88) is

$$\mathscr{L} = \sum_{i=1}^{N}\sum_{k=1}^{K}\eta_k g_i\kappa_{i,k}\,(\delta_i) - \sum_{m=1}^{M}\nu_m\left(\Delta_m - \mathbf{1}_m^\top\delta^\top\tilde{\omega}\right),$$

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

$$g_i \sum_{k=1}^{K} \eta_k \nabla \kappa_{i,k}(\delta_i) = -\nu \tilde{\omega}_i,$$

which implies that

$$\eta \frac{g_i}{\tilde{\omega}_i} \sum_{k=1}^{K} \eta_k \nabla \kappa_{i,k}(\delta_i) = \eta \frac{g_j}{\tilde{\omega}_j} \sum_{k=1}^{K} \eta_k \nabla \kappa_{j,k}(\delta_j) = \nabla \bar{\kappa},$$

for all i, j, where the last equality comes from the envelope theorem. Notice that

$$\frac{g_i}{\tilde{\omega}_i} \sum_{k=1}^{K} \eta_k H_{i,k}(\delta_i - \delta_i^\circ) = \frac{g_j}{\tilde{\omega}_j} \sum_{k=1}^{K} \eta_k H_{j,k}\left(\delta_j - \delta_j^\circ\right) \Leftrightarrow \delta_i = \delta_i^\circ + \frac{\tilde{\omega}_i}{g_i} \frac{g_j}{\tilde{\omega}_j} \left(\sum_{k=1}^{K} \eta_k H_{i,k}\right)^{-1} \left(\sum_{k=1}^{K} \eta_k H_{j,k}\right)(\delta_k - \delta_k^\circ).$$

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

$$\begin{aligned}
\Delta &= \sum_{i=1}^{N} \tilde{\omega}_i \delta_i = \sum_{i=1}^{N} \tilde{\omega}_i \left(\delta_i^\circ + \frac{\tilde{\omega}_i}{g_i} \frac{g_j}{\tilde{\omega}_j} \left(\sum_{k=1}^{K} \eta_k H_{i,k}\right)^{-1} \left(\sum_{k=1}^{K} \eta_k H_{j,k}\right)(\delta_k - \delta_k^\circ) \right) \\
&= \Delta^\circ + \left(\sum_{i=1}^{N} \frac{\tilde{\omega}_i^2}{g_i} \left(\sum_{k=1}^{K} \eta_k H_{i,k}\right)^{-1} \right) \frac{g_j}{\tilde{\omega}_j} \left(\sum_{k=1}^{K} \eta_k H_{j,k}\right)\left(\delta_j - \delta_j^\circ\right) \Leftrightarrow \\
\delta_j - \delta_j^\circ &= \frac{\tilde{\omega}_j}{g_j} \left(\sum_{k=1}^{K} \eta_k H_{j,k}\right)^{-1} \nabla^2 \bar{\kappa}(\Delta - \Delta^\circ),
\end{aligned}$$

where $\nabla^2 \bar{\kappa}$ is given by (90). Combining this last expression with (88) yields the result. □

Since (87) and (88) fully characterize the equilibrium, it follows that the only difference between the baseline model and the multiple resources model comes from the cost functions $\bar{\kappa}$. In turn, the cost function in the baseline model is identical to (89) if we impose that

$$H_i^{-1} = \sum_{k=1}^{K} \eta_k H_i^k$$

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.

E.3 General cost function

In this appendix, we relax our assumption on the functional form of the cost function $\kappa_i(\delta_i)$. Specifically, we assume that $\kappa_i(\delta_i)$ is a strictly convex—but not necessarily quadratic—function. The Hessian of $\kappa_i(\delta_i)$ is denoted by H_i. It is a positive definite matrix. Unlike in the baseline model, H_i is not necessarily constant, in the sense that it can depend on δ_i.

All the results of Sections 2, 3, and 4 remain unchanged. However, Equation (24) does not hold when κ_i is non-quadratic. Nevertheless, the Hessian of the aggregate cost function $\bar{\kappa}$ is still given by (25).

Lemma 8. *The Hessian matrix of $\bar{\kappa}$ is given by (25).*

Proof. Equations (70), (71), and (72) do not use the assumption that κ_i is quadratic and, hence, hold in this more general case. Differentiate (72) with respect Δ_j to get

$$\frac{d^2\bar{\kappa}}{d\Delta_i d\Delta_j} = \eta \frac{g_i}{\tilde{\omega}_i} \mathbf{1}_i^\top H_i \frac{d\delta_i}{d\Delta_j} \tag{91}$$

and

$$\frac{g_k}{\tilde{\omega_k}} H_k \frac{d\delta_k}{d\Delta_j} = \frac{g_i}{\tilde{\omega}_i} H_i \frac{d\delta_i}{d\Delta_j}$$

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

$$\frac{g_i}{\tilde{\omega}_i} H_i \frac{d\delta_i}{d\Delta_j} = \left(\sum_{k=1}^N \frac{\tilde{\omega}_k^2}{g_k} H_k^{-1} \right)^{-1} \mathbf{1}_j. \tag{92}$$

Plugging this into (91) yields (25). □

E.3.1 Risk exposures

Going to the results of Section 5, 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*

$$\frac{d\delta_i}{d\Delta_j} = \left(\frac{1}{\eta} \frac{\tilde{\omega}_i}{g_i} H_i^{-1} \right) \nabla^2 \bar{\kappa} \mathbf{1}_j. \tag{93}$$

Proof. This result follows immediately from (92) and (25). □

In the model with general cost functions, (27) still implicitly defines equilibrium Δ. Proposition 3 immediately follows by applying the implicit function theorem to (27). 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*

$$\frac{d\Delta}{d\tau_i} = -\mathcal{H}^{-1}\nabla^2\bar{\kappa}\left(\sum_{j=1}^{N}\frac{1}{\eta}\frac{\tilde{\omega}_j^2}{g_j^2}H_j^{-1}\frac{dg_j}{d\tau_i}\right)\mathcal{E}, \tag{94}$$

where $\mathcal{H}$ is given by (32), $\nabla^2\bar{\kappa}$ is given by (25), and $\frac{dg_j}{d\tau_i}$ is given by (23).

Proof. By the implicit function theorem applied to (27),

$$\frac{d\Delta}{d\tau_i} = -\mathcal{H}^{-1}\sum_{j=1}^{N}\frac{d\nabla\bar{\kappa}}{dg_j}\frac{dg_j}{d\tau_i}, \tag{95}$$

where $\mathcal{H}$ is given by (32) and $\frac{dg_j}{d\tau_i}$ is given by (23).

Next, applying the envelope theorem to (70) implies that $\frac{d\bar{\kappa}}{dg_j} = \eta\kappa_j(\delta_j)$, and so

$$\frac{d}{d\Delta_k}\frac{d\bar{\kappa}}{dg_j} = \eta\left(\nabla\kappa_j\right)^\top\frac{d\delta_j}{d\Delta_k}.$$

Using (72) and (93), we get

$$\frac{d\nabla\bar{\kappa}}{dg_j} = \nabla^2\bar{\kappa}\left(\frac{1}{\eta}\frac{\tilde{\omega}_j^2}{g_j^2}H_j^{-1}\right)\nabla\bar{\kappa}. \tag{96}$$

Recall that from (27), $\nabla\bar{\kappa} = \mathcal{E}$. Therefore, plugging the equation above into (95), we get 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*

$$\frac{d\Delta}{d\alpha_{ij}} = -\mathcal{H}^{-1}\nabla^2\bar{\kappa}\sum_{k=1}^{N}\left[\frac{dg_k}{d\alpha_{ij}}\left(\frac{1}{\eta}\frac{\tilde{\omega}_k^2}{g_k^2}H_k^{-1}\right)\mathcal{E} - \frac{d\tilde{\omega}_k}{d\alpha_{ij}}\left[\delta_k + \left(\frac{1}{\eta}\frac{\tilde{\omega}_k}{g_k}H_k^{-1}\right)\mathcal{E}\right]\right], \tag{97}$$

where $\frac{dg_k}{d\alpha_{ij}} = \frac{(1+\tau_k)}{\omega_k}\frac{d\tilde{\omega}_k}{d\alpha_{ij}} - \frac{\tilde{\omega}_k(1+\tau_k)}{\omega_k^2}\frac{d\omega_k}{d\alpha_{ij}}$, $\frac{d\tilde{\omega}_k}{d\alpha_{ij}} = \tilde{\omega}_i\tilde{\mathcal{L}}_{jk}$, $\frac{d\omega_k}{d\alpha_{ij}} = \frac{\omega_i}{1+\tau_i}\mathcal{L}_{jk}$, $\mathcal{H}$ is given by (32), and $\nabla^2\bar{\kappa}$ is given by (25).

Proof. By the implicit function theorem applied to (27),

$$\frac{d\Delta}{d\alpha_{ij}} = -\mathcal{H}^{-1}\sum_{k=1}^{N}\left(\frac{d\nabla\bar{\kappa}}{dg_k}\frac{dg_k}{d\alpha_{ij}} + \frac{d\nabla\bar{\kappa}}{d\tilde{\omega}_k}\frac{d\tilde{\omega}_k}{d\alpha_{ij}}\right), \tag{98}$$

where $\mathcal{H}$ is given by (32), $\frac{d\tilde{\omega}_k}{d\alpha_{ij}} = \tilde{\omega}_i \tilde{\mathcal{L}}_{jk}$,

$$\frac{dg_k}{d\alpha_{ij}} = \frac{(1+\tau_k)}{\omega_k}\frac{d\tilde{\omega}_k}{d\alpha_{ij}} - \frac{\tilde{\omega}_k(1+\tau_k)}{\omega_k^2}\frac{d\omega_k}{d\alpha_{ij}}$$

and $\frac{d\omega_k}{d\alpha_{ij}} = \frac{\omega_i}{1+\tau_i}\mathcal{L}_{jk}$.

From (96) and (76),

$$\frac{d\nabla\bar{\kappa}}{dg_k} = \nabla^2\bar{\kappa}\left(\frac{1}{\eta}\frac{\tilde{\omega}_k^2}{g_k^2}H_k^{-1}\right)\mathcal{E}, \tag{99}$$

where $\nabla^2\bar{\kappa}$ is given by (25).

Next, applying the envelope theorem to (70) implies that $\frac{d\bar{\kappa}}{d\tilde{\omega}_k} = \nu^\top \delta_k$, where $\nu = -\nabla\bar{\kappa}$ from (72). Therefore,

$$\frac{d}{d\Delta_m}\frac{d\bar{\kappa}}{d\tilde{\omega}_k} = -\mathbf{1}_m^\top \nabla^2\bar{\kappa}\delta_k - (\nabla\bar{\kappa})^\top \frac{d\delta_k}{d\Delta_m}.$$

Using (93), this equation can be rewritten as

$$\frac{d\nabla\bar{\kappa}}{d\tilde{\omega}_k} = -\nabla^2\bar{\kappa}\delta_k - \nabla^2\bar{\kappa}\left(\frac{1}{\eta}\frac{\tilde{\omega}_k}{g_k}H_k^{-1}\right)\mathcal{E}. \tag{100}$$

Plugging (99) and (100) into (98), we get (97). □

E.3.2 GDP and welfare

The results of Section 6 hold under general cost functions κ_i. As can be seen from the analysis of the baseline model, we do not use the fact that κ_i is quadratic in this section, with the exception of deriving the expression for $\nabla\bar{\kappa} - \nabla\bar{\kappa}_V$, which is given by (82). This equation holds under general cost functions κ_i. Indeed, using the definition of $\bar{\kappa}_V$, we can derive

$$\frac{d\bar{\kappa}_V}{d\Delta_m} = \eta\sum_{i=1}^N (\nabla\kappa_i)^\top \frac{d\delta_i}{d\Delta_m}.$$

From (72), we have $\nabla\kappa_i = \frac{1}{\eta}\frac{\tilde{\omega}_i}{g_i}\nabla\bar{\kappa}$. Furthermore, $\frac{d\delta_i}{d\Delta_m}$ is given by (93). Then we have

$$\nabla\kappa_V = \nabla^2\bar{\kappa}\left(\sum_{i=1}^N \frac{1}{\eta}\frac{\tilde{\omega}_i^2}{g_i^2}H_i^{-1}\right)\nabla\bar{\kappa}.$$

Using the definition of $\nabla^2\bar{\kappa}$, given by (25), and the first-order condition (76), it is straightforward to see that (82) holds.

E.4 A general specification of the disutility from risk management

In the main model, we assume that the household's disutility is given by (4), where $\mathcal{V}(R) = \exp(-\eta(1-\rho)R)$. 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 $\mathcal{U}(\cdot,\cdot)$ with respect to the first and second inputs, respectively. Following the same steps as in the baseline model (see Appendix A.1), we can derive

$$\mathrm{E}[\mathcal{U}_2] = -W_R \mathrm{E}[\Lambda], \tag{101}$$

where

$$\Lambda = \Gamma_L^{\rho}\mathcal{U}_1\left(\prod_{i=1}^{N}P_i^{\beta_i}\right)^{\rho-1}. \tag{102}$$

As in the baseline model, the labor share of income Γ_L is given by $\Gamma_L := \frac{W_L L}{P^\top C} = 1 - \tau^\top\left(\mathrm{diag}(1+\tau)\right)^{-1}\omega$, 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 (9), and the vector of prices is given by (15). The choice of risk exposure for firm i is described by (10), which can be rewritten as

$$\delta_i^* \in \arg\min_{\delta_i\in\mathcal{A}_i}\mathrm{E}\left[\Lambda\left[Q_iK_i(\delta_i,P) + \kappa_i(\delta_i)W_R\right]\right].$$

Taking the first-order condition with respect to δ_{im}, we get

$$\mathrm{E}\left[\Lambda Q_i\frac{dK_i}{d\delta_{im}}\right] + \mathrm{E}[\Lambda]W_R\frac{d\kappa_i}{d\delta_{im}} = 0.$$

Using the expression for the unit cost (9) together with equations (101) and (102), this expression can be simplified as

$$\frac{d\kappa_i}{d\delta_{im}} = -K_iQ_i\frac{\mathrm{E}\left[\Gamma_L^{\rho}U_1\left(\prod_{i=1}^{N}P_i^{\beta_i}\right)^{\rho-1}\varepsilon_m\right]}{\mathrm{E}[U_2]}.$$

Recall that in equilibrium, $K_iQ_i = \frac{\omega_i\Gamma_L^{-1}}{1+\tau_i}$. Using the expression for prices (15), we can rewrite the equation above as

$$\frac{d\kappa_i}{d\delta_{im}} = -\frac{\omega_i\Gamma_L^{\rho-1}}{1+\tau_i}\exp\left(-(\rho-1)\tilde{\omega}^\top\log(1+\tau)\right)\frac{\mathrm{E}\left[\mathcal{U}_1\exp\left(-(\rho-1)\tilde{\omega}^\top\delta\varepsilon\right)\varepsilon_m\right]}{\mathrm{E}[\mathcal{U}_2]}. \tag{103}$$

The system of equations (103) 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 (26). This distorted planner chooses the aggregate risk exposure Δ, and then chooses δ to minimize the aggregate cost function (22) 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 (103) 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 (103) 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.

E.4.1 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, (103) becomes

$$\begin{aligned}\frac{d\kappa_i}{d\delta_{im}} &= (\rho-1)\frac{\omega_i}{1+\tau_i}\frac{\mathbb{E}\left[\exp\left(-(\rho-1)\tilde{\omega}^\top\delta\varepsilon\right)\varepsilon\right]}{\mathbb{E}\left[\exp\left(-(\rho-1)\tilde{\omega}^\top\delta\varepsilon\right)\right]}\frac{V(R)}{V'(R)} \\ &= (\rho-1)\frac{\omega_i}{1+\tau_i}\left(\mu_m - (\rho-1)\mathbf{1}_m^\top\Sigma\delta^\top\tilde{\omega}\right)\frac{V(R)}{V'(R)}.\end{aligned}$$

Market clearing implies that $R = \sum_{i=1}^N \kappa_i(\delta_i)$. 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

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

Then, by the implicit function theorem,

$$\begin{pmatrix} \frac{d\delta_{11}}{d\chi} \\ \vdots \\ \frac{d\delta_{1M}}{d\chi} \\ \vdots \\ \frac{d\delta_{N1}}{d\chi} \\ \vdots \\ \frac{d\delta_{NM}}{d\chi} \end{pmatrix} = - \begin{pmatrix} \frac{\partial F_{11}}{\partial \delta_{11}} & \cdots & \frac{\partial F_{11}}{\partial \delta_{1M}} & \cdots & \frac{\partial F_{11}}{\partial \delta_{N1}} & \cdots & \frac{\partial F_{11}}{\partial \delta_{NM}} \\ \vdots & \ddots & \vdots & \ddots & \vdots & \ddots & \vdots \\ \frac{\partial F_{1M}}{\partial \delta_{11}} & \cdots & \frac{\partial F_{1M}}{\partial \delta_{1M}} & \cdots & \frac{\partial F_{1M}}{\partial \delta_{N1}} & \cdots & \frac{\partial F_{1M}}{\partial \delta_{NM}} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \frac{\partial F_{N1}}{\partial \delta_{11}} & \cdots & \frac{\partial F_{N1}}{\partial \delta_{1M}} & \cdots & \frac{\partial F_{N1}}{\partial \delta_{N1}} & \cdots & \frac{\partial F_{N1}}{\partial \delta_{NM}} \\ \vdots & \ddots & \vdots & \ddots & \vdots & \ddots & \vdots \\ \frac{\partial F_{NM}}{\partial \delta_{11}} & \cdots & \frac{\partial F_{NM}}{\partial \delta_{1M}} & \cdots & \frac{\partial F_{NM}}{\partial \delta_{N1}} & \cdots & \frac{\partial F_{NM}}{\partial \delta_{NM}} \end{pmatrix}^{-1} \begin{pmatrix} \frac{\partial F_{11}}{\partial \chi} \\ \vdots \\ \frac{\partial F_{1M}}{\partial \chi} \\ \vdots \\ \frac{\partial F_{N1}}{\partial \chi} \\ \vdots \\ \frac{\partial F_{NM}}{\partial \chi} \end{pmatrix}, \tag{104}$$

where

$$\frac{\partial F_{im}}{\partial \delta_{jl}} = \frac{\partial^2 \kappa_i}{\partial \delta_{im} \partial \delta_{il}} 1_{j=i} + \frac{\omega_i}{1+\tau_i} (\rho - 1)^2 \Sigma_{ml} \tilde{\omega}_j \frac{V(R)}{V'(R)} - \frac{\partial \kappa_i}{\partial \delta_{im}} \frac{\partial \kappa_j}{\partial \delta_{jl}} \left(\frac{V'(R)}{V(R)} - \frac{V''(R)}{V'(R)} \right),$$

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

$$\frac{\partial F_{im}}{\partial \mu_j} = -1_{j=m} (\rho - 1) \frac{\omega_i}{1+\tau_i} \frac{V(R)}{V'(R)},$$

Similarly, one can compute

$$\begin{aligned}
\frac{\partial F_{im}}{\partial \Sigma_{jl}} &= (\rho - 1)^2 \frac{\omega_i}{1+\tau_i} \frac{1}{2} (1_{j=m} \Delta_l + 1_{l=m} \Delta_j) \frac{V(R)}{V'(R)}, \\
\frac{\partial F_{im}}{\partial \tau_j} &= (\rho - 1) \frac{1}{1+\tau_i} \frac{\omega_j}{1+\tau_j} \mathcal{L}_{ji} \left(\mu_m - (\rho - 1) \mathbf{1}_m^\top \Sigma \Delta \right) \frac{V(R)}{V'(R)}, \\
\frac{\partial F_{im}}{\partial \alpha_{jl}} &= -(\rho - 1) \frac{\mathcal{L}_{li}}{1+\tau_i} \frac{\omega_j}{1+\tau_j} \left(\mu_m - (\rho - 1) \mathbf{1}_m^\top \Sigma \Delta \right) \frac{V(R)}{V'(R)} + (\rho - 1)^2 \frac{\omega_i}{1+\tau_i} \mathbf{1}_m^\top \Sigma \left(\sum_{k=1}^N \delta_k \tilde{\mathcal{L}}_{lk} \right) \tilde{\omega}_j \frac{V(R)}{V'(R)},
\end{aligned}$$

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 $\frac{d\left(\frac{\omega_j}{1+\tau_j}\right)}{d\tau_i} = -\frac{\mathcal{L}_{ij}}{1+\tau_j} \frac{\omega_i}{1+\tau_i}$.

E.4.2 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, (103) becomes

$$\begin{aligned}
\frac{d\kappa_i}{d\delta_{im}} &= \frac{\omega_i \Gamma_L^{\rho-1}}{1+\tau_i} \exp\left(-(\rho - 1) \tilde{\omega}^\top \log(1+\tau) \right) \mathrm{E}\left[\exp\left(-(\rho - 1) \tilde{\omega}^\top \delta \varepsilon \right) \varepsilon_m \right] \frac{1}{V'(R)}, \\
&= \frac{\omega_i \Gamma_L^{\rho-1}}{1+\tau_i} \left(\mu_m - (\rho - 1) \mathbf{1}_m^\top \Sigma \delta^\top \tilde{\omega} \right) \exp\left(-(\rho - 1) \tilde{\omega}^\top (\delta \mu - \log(1+\tau)) + \frac{1}{2} (\rho - 1)^2 \tilde{\omega}^\top \delta \Sigma \delta^\top \tilde{\omega} \right) \frac{1}{V'(R)}.
\end{aligned}$$

Market clearing implies that $R = \sum_{i=1}^{N} \kappa_i(\delta_i)$. 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

$$\begin{aligned} F_{im} =& \frac{\partial \kappa_i}{\partial \delta_{im}} - \frac{\omega_i \Gamma_L^{\rho-1}}{1+\tau_i}\left(\mu_m - (\rho-1)\,\mathbf{1}_m^\top \Sigma \delta^\top \tilde{\omega}\right) \times \\ & \exp\left(-(\rho-1)\,\tilde{\omega}^\top(\delta\mu - \log(1+\tau)) + \frac{1}{2}(\rho-1)^2\,\tilde{\omega}^\top \delta\Sigma\delta^\top\tilde{\omega}\right)\frac{1}{V'\left(\sum_{i=1}^{N}\kappa_i(\delta_i)\right)} = 0. \end{aligned}$$

We can again use the implicit function theorem to get (104), where

$$\frac{\partial F_{im}}{\partial \delta_{jl}} = \frac{\partial^2 \kappa_i}{\partial \delta_{im}\partial \delta_{il}}\mathbf{1}_{j=i} + \frac{\partial \kappa_i}{\partial \delta_{im}}\left\{\frac{(\rho-1)\,\Sigma_{ml}\tilde{\omega}_j}{\mu_m - (\rho-1)\,\mathbf{1}_m^\top\Sigma\delta^\top\tilde{\omega}} + \frac{V''(R)}{V'(R)}\frac{\partial \kappa_j}{\partial \delta_{jl}} + + (\rho-1)\,\tilde{\omega}_j\left(\mu_l - (\rho-1)\,\mathbf{1}_l^\top\Sigma\delta^\top\tilde{\omega}\right)\right\},$$

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

$$\frac{\partial F_{im}}{\partial \mu_j} = \left(-\mathbf{1}_{j=m}\frac{1}{\mu_m - (\rho-1)\,\mathbf{1}_m^\top\Sigma\delta^\top\tilde{\omega}} + (\rho-1)\,\tilde{\omega}^\top\delta\mathbf{1}_j\right)\frac{\partial \kappa_i}{\partial \delta_{im}}.$$

Similarly, one can compute derivatives of F_{im} with respect to Σ_{jl}, τ_j and α_{jl} in a straightforward way.

E.5 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 (56), which we can write as

$$W_R = -V'(R)\,(V(R))^{-\rho}\,\frac{\mathrm{E}\left[Y^{1-\rho}\right]}{\mathrm{E}\left[\Lambda\right]}, \tag{105}$$

where Λ is given by (53).

Next, given the Cobb-Douglas production function, the unit cost is given by

$$K_i(\delta_i, P, W_L) = \exp\left(-\delta_i^\top\varepsilon\right)W_L^{1-\sum_{j=1}^{n}\alpha_{ij}}\prod_{j=1}^{N}P_j^{\alpha_{ij}}.$$

Since $P_i = (1+\tau_i)\,K_i$,this implies that

$$p - w_L = -\tilde{\mathcal{L}}\left(\delta\varepsilon - \log(1+\tau)\right), \tag{106}$$

where $w_L = \log W_L$ and $\tilde{\mathcal{L}} = (I - \alpha)^{-1}$.

Firms' risk exposure solves (10). Taking first-order condition with respect to δ_{im}, we get

$$\frac{d\kappa_i(\delta_i)}{d\delta_{im}} = \frac{1}{W_R}\frac{\mathrm{E}[\Lambda Q_i K_i \varepsilon_m]}{\mathrm{E}[\Lambda]}.$$

Combining this equation with (105), we get

$$\frac{d\kappa_i(\delta_i)}{d\delta_{im}} = -\frac{1}{V'(R)(V(R))^{-\rho}}\frac{\mathrm{E}[\Lambda K_i Q_i \varepsilon_m]}{\mathrm{E}[Y^{1-\rho}]}. \tag{107}$$

Next, from the market-clearing condition (13),

$$P_i Q_i = P_i C_i + \sum_j P_i X_{ji}.$$

Using (52) 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 (14). Hence, (107) can be rewritten as

$$\frac{d\kappa_i(\delta_i)}{d\delta_{im}} = -\frac{1}{V'(R)(V(R))^{-\rho}}\frac{\omega_i}{1+\tau_i}\frac{\mathrm{E}[\Lambda(W_L + W_R R + \Pi)\varepsilon_m]}{\mathrm{E}[Y^{1-\rho}]}.$$

Combining this with (51), we get

$$\frac{d\kappa_i(\delta_i)}{d\delta_{im}} = -\frac{V(R)}{V'(R)}\frac{\omega_i}{1+\tau_i}\frac{\mathrm{E}[U'(Y) Y \varepsilon_m]}{\mathrm{E}[Y^{1-\rho}]} = -\frac{V(R)}{V'(R)}\frac{\omega_i}{1+\tau_i}\frac{\mathrm{E}\left[Y^{1-\rho}\varepsilon_m\right]}{\mathrm{E}[Y^{1-\rho}]}. \tag{108}$$

From (54), we can write GDP Y as

$$Y = \left(\frac{W_L + W_R R + \Pi}{W_L}\right)\prod_{i=1}^{N}\left(\frac{P_i}{W_L}\right)^{-\beta_i}. \tag{109}$$

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, (109) 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 (106), (109) becomes

$$\frac{d\kappa_i(\delta_i)}{d\delta_{im}} = -\frac{V(R)}{V'(R)}\frac{\omega_i}{1+\tau_i}\frac{\mathrm{E}\left[\exp\left(-(\rho-1)\Delta^\top\varepsilon\right)\varepsilon_m\right]}{\mathrm{E}\left[\exp\left(-(\rho-1)\Delta^\top\varepsilon\right)\right]}.$$

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 (76). Therefore, setting $W_L = 1$ does

not affect equilibrium risk exposures and hence moments of GDP.