From Just in Time, to Just in Case, to Just in *Worst*-Case:

Simple models of a Global Supply Chain under Uncertain Aggregate Shocks. [‡§]

Bomin Jiang — MIT, IDSS

Daniel Rigobon — Princeton, ORFE

Roberto Rigobon — MIT and NBER

October 12, 2020

Abstract

Covid-19 highlighted the weaknesses in the supply chain. Many have argued that a more resilient or robust supply chain is needed. But what does a robust supply chain mean? And how do firms' decisions change when taken that approach? This paper studies a very stylized model of a supply chain, where we study how the decision of a multinational corporation changes in the presence of uncertainty. We find that a robust supply chain implies concentrating on the worst-case, which is not achieved by simply increasing the size of shocks. Our model rationalizes or explains the well known "probability matching" behavior observed in the experimental literature.

JEL Codes: E7, F02, F12, F13, L15

Keywords: Global Supply Chain, International Trade, Robustness, International Trade Policy

[‡] We are grateful to the IMF for inviting us to present at 21st Jacques Polak Annual Research Conference - November 5-6, 2020. All errors are our own. Data and code will be made publicly available.

[§] Jiang: bominj AT mit.edu. Rigobon, Daniel: drigobon AT princeton.edu. Rigobon, Roberto: rigobon AT mit.edu

1 Introduction

"Just-in-Time" (JIT) manufacturing was introduced in Japan during the late 1940s and early 1950s by Toyota, with the purpose of reducing inventories, reducing setup times, and saving costs in other aspects of the supply chain. The reasons why it started in Japan are not completely clear. Some have argued that it was a combination of limited natural resources, a lack of physical space to hold the inventories, and financial constraints that the Japanese industry faced at the end of the war[1]. The cost reduction and efficiency gains of JIT became well known worldwide. Indeed, it became the standard of excellence in a short period of time and was adopted by many corporations. In fact, the globalization of the manufacturing of goods and services that started in the 1980s was, for the most part, inspired by JIT premises. Interestingly, even from the very beginning, Toyota suggested that the main risk of this strategy was its excessive reliance on suppliers — which could be less resilient and flexible than Toyota itself. Hence, a successful JIT implementation required a large emphasis on supplier development. Toyota argued that the JIT's biggest weakness was its vulnerability.

In response to those risks and seeking resilience, companies have explored other alternatives. These strategies, dubbed "Just-in-Case", usually recommend actions such as larger inventories, diversification of the production network, and parts harmonization. In the end, the advice is simple: to develop flexibility and redundancy in the supply chain. Regardless of all these efforts, the world's supply chains were proved to be unreliable during the Covid-19 pandemic. Either international trade frictions, quarantine restrictions, large shifts in demand (such as protective equipment), or even panic demand purchases of some products (such as toilet paper and disinfectant wipes in the US) highlighted the fragility of supply chains in the world. Over the past months, many countries experienced supply disruptions for various products. The collapse is due to a design flaw: while the just-in-case approach might be appropriate for idiosyncratic shocks, it seems to have failed in the presence of an aggregate shock.

The need to redefine the internationalization of production was eloquently captured by the Financial Times headline: From "Just-in-Time" to "Just-in-Case"[2]. Many have argued that resilience and robustness are needed. However, what exactly does it mean to have a robust supply chain? How does it differ from assuming a larger shock? In this paper, we make the case that just-in-case manufacturing strategies were not enough to patch the existing vulnerabilities. We believe it is now time to approach supply chain optimization through robust control. Robustness implies more than just assuming larger shocks; it implies

[1]See the Toyota Production System, where one of the two pillars for production is just-in-time: a type of production where "only the necessary products, at the necessary time, in the necessary quantity are manufactured, and in addition, the stock on hand is held to a minimum." Also see Plenert (2007).

[2]See The FT Editorial Board (2020), Evans (2020) and Long (2020)

concentrating on the worst possible outcome; it implies shifting to "Just-in-*Worst*-Case".

In this paper, we present a very simple model. It introduces the idea that, in the pursuit of efficiency, a decentralized supply chain could become vulnerable to aggregate shocks. In other words, our setup reflects the well-known trade-off between efficiency and robustness. We are interested in understanding the consequences of natural disasters, environmental shocks, and pandemics on the global supply chain. These shocks tend to be substantial and infrequent, but also widespread — affecting many countries and regions at the same time. Just-in-Case's standard supply chain analysis suggests that increased inventories or diversified production protect against large and frequent idiosyncratic shocks. In the presence of aggregate shocks, however, the right notion of diversification needs to be expanded.

We study the survival of a multinational firm that purchases from small global suppliers. The suppliers decide where to locate, and locations are subject to aggregate shocks. We compare two types of global supply chain arrangements and two different natures of shocks. First, the small suppliers individually decide their location, and the multinational purchases the surviving suppliers' products. Second, we study the case when the multinational can choose all its suppliers' locations, thereby internalizing the location decision (i.e., the suppliers are subsidiaries of the multinational firm). From the perspective of the shocks, we compare the situation of risk versus uncertainty through two settings of random aggregate shocks. In the first case, we assume the distribution is known, while in the second case, only the distribution's support us known.

In addition to these two aspects, our model has four additional ingredients: possible locations for production facilities, aggregate shock characteristics, prices of intermediate and final products, and growth of the number of suppliers.

Our basic model has two locations: the Mountain and the Valley[3]. The two locations differ only in the probability that an aggregate shock hits. We assume the aggregate shock is extreme, such that all suppliers in the affected location perish. We also assume that the shock is more likely to occur in the Mountain than in the Valley. Therefore, if an aggregate shock occurs, then the conditional probability of the shock affecting the Mountain (or the Valley) is θ (or $1 - \theta$), where $\theta > 1/2$.

Each supplier produces a single unit, or a "part". The cost to produce and price per part are the same: in both locations, for all producers, and in any state of the world. In particular, we assume that prices are independent of both the realization of the aggregate shock, and the number of surviving suppliers[4]. However, in canonical macroeconomic and international

[3]This example is inspired by Lo (2017) discussion on probability matching.

[4]This assumption is motivated by price gouging laws and consumer anger in response to natural disasters. For consumer anger see Rotemberg (2002, 2011)) and for price gouging laws Executive Order 13910 of March

models, this would not be the case. The demand is usually chosen such that when quantities tend to zero, prices increase and can even tend to infinity. Models based on Cobb-Douglas or CES functions have this feature, and the pricing system reflects scarcity. Nonetheless, if firms are concerned about the consequences of increasing their prices after natural disasters, or there is a law that does not allow prices to increase after such an event, then the price required to achieve the efficient allocation may never be realized. If this were known ex-ante, then it would affect the willingness of suppliers to diversify into risky locations. Our model takes an extreme assumption — prices are fixed — to capture this feature of regulations and institutions. Moreover, below we include some anecdotal evidence of law enforcement and consumers' negative reactions to price gouging during the Covid-19 pandemic.

The number of suppliers grows and depends on the number of surviving suppliers. We treat each suppliers' part as a different intermediate good, and the multinational purchases as many parts as possible to sell the final product internationally. The product is more desirable the more parts it has — but it can be manufactured with a subset of the parts. This setting is equivalent to assume that the quality of the item increases with the number of parts it has[5].

The parameters in the model are such that each supplier's location decision has a dominant strategy, which is not collectively optimal when the probability of global survival is taken into account. This discrepancy comes from the inability of the pricing system to compensate firms properly for moving into the Mountain. In this setting, the multinational wants to ensure its suppliers survive. Putting this differently, the multinational cares about survival, while the suppliers do not. This difference implies that the multinational might be willing to set production facilities in the Mountain to insure itself against an aggregate shock in the Valley. This part of the model captures a simple externality and the need for a diversified supply chain, but not yet a robust one.

This is where the nature of the shock matters. When the shocks have known distributions — what is known as risk or the *nominal* model — the multinational will exhibit behavior that takes into account all the sources of risk. This setting implies a desire for diversification, and one of the implications, for example, is that the multinational's optimal allocation of firms to the Mountain depends on the number of suppliers that have survived. There is both a marginal benefit and cost of diversification, and in general, an internal solution is found (at least under our assumptions). The multinational's policy changes dramatically when the shock has bounded uncertainty — meaning that the distribution is unknown — in the *robust*

23, 2020, *Preventing Hoarding of Health and Medical Resources To Respond to the Spread of COVID-19*. For a list of price gouging laws in the US, see King and Spalding (2020)

[5]It is very common in supply chain management to assume that if one good or part is missing, the whole product can't be manufactured. By relaxing this assumption, we can eliminate the typical assumptions behind just-in-case theories and concentrate on the robustness aspect.

model. In this setting, utility optimization requires a robust control approach. For example, the optimal allocation of firms is independent of the number of surviving firms. The robust supply chain decision, therefore, looks very different from the nominal model.

Our model replicates a behavior from evolutionary games known as the probability matching. Even though, for individual firms, the rational optimal decision is to always locate in the Valley, the socially optimal allocation implies a higher degree of diversification. In the case of uncertainty, firms' optimal allocation is independent of the number of suppliers that exist– very much in the spirit of probability matching. However, the optimal solution does not coincide exactly with the probability matching from the experimental literature.

This paper includes several theoretical results worth highlighting. First, we compare the centralized and decentralized solutions to the model. We show a corner solution of the decentralized allocation (all firms locate themselves in the Valley), exposing the multinational to an aggregate shock to the Valley. This result contrasts with the internal solution (a proportion larger than zero of firms in the Mountain) when a centralized approach is used. Individual suppliers maximize efficiency (or productivity) while the multinational maximizes survival.

Second, we study the implications of risk and uncertainty in the probability distribution of the aggregate shock. We depart from the assumption of the first part of our paper: that the value of θ, the probability the aggregate shock affects the Mountain, is known. We first study what occurs when θ is risky; for example assuming that it is distributed between $[\bar{\theta} - \Delta, \bar{\theta} + \Delta]$ with some known distribution. This exercise represents a risky setup — the nominal model. We compare it to the uncertain setup where the multinational only knows that $\theta \in [\bar{\theta} - \Delta, \bar{\theta} + \Delta]$, but the distribution is unknown. This case represents uncertainty, and the optimal control problem implies a *robust* approach. The individual suppliers and the multinational maximize assuming the worst-case scenario. For example, if $\bar{\theta} > 1/2$ but the interval contains $1/2$ then the robust control makes the choice as if $\theta = 1/2$. Robustness implies a degree of diversification in the supply chain that is not achieved when the multinational faces only risk.

Finally, we study what happens when prices and costs of production differ across locations. We compare three settings for the global supply chain: (i) the decentralized myopic setting that always chooses a corner solution, except when prices and costs are at the knife-edge when the value of both locations is the same; (ii) the probability matching heuristic where the allocation of firms coincides with the probabilities of survival in each location; (iii) and the optimal allocation by the multinational. We show that the decentralized solution can replicate the centralized solution when the value in the Mountain is equal to the value in the Valley. Our model justifies why governmental subsidies can help the decentralized economy achieve the robust solution. Our discussion was motivated by Japan's policy actions during the summer of 2020, where Japan set up a fund to compensate firms that diversify out of China[6]. Of course, the decision could have been driven by political aspects not considered in this paper, but our model at least rationalizes economically why such industrial policy makes sense.

In each of the different simple models, we draw some policy implications that have been motivated by the recent experience with the pandemic, which is summarized in the end. All of them, however, have a simple message: robustness is under-supplied.

1.1 Price Gouging

One of the important assumptions of our model is that prices do not adjust to the aggregate shocks. This has been motivated by price gouging laws in many countries — either because of fairness considerations or consumer anger. This subsection discusses the origin and evidence of such an assumption.

When facing large aggregate shocks, the price of essential goods cannot float freely as assumed by the classic general equilibrium model. Take Covid-19 as an example; during the pandemic, the demand for personal protective equipment, foods, and other essential supplies has risen dramatically, which raises the concern of price gouging with both the regulators and the general public.

On the side of regulators, Executive Order 13910 of March 23, 2020, *Preventing Hoarding of Health and Medical Resources To Respond to the Spread of COVID-19* was issued by the US to deal with the threat of price gouging. Individual States in the US, referring to the laws of King and Spalding (2020), are very active in controlling companies pricing for products related to the pandemic.

[6]See Bloomberg News (2020)

Figure 1.1: Law enforcements' efforts to control price gouging.

As shown in Figure 1.1, many State Attorney General Offices created procedures to deal with price gouging complaints. The purpose of these laws is quite "benevolent"; regulators seek to prevent hoarding and ensure that the prices of essential goods do not increase beyond that which is considered "reasonable or fair". However, in practice, it is not possible to distinguish which part of the price increase is reasonable (e.g., a price increase which generates profits to compensate for the cost of diversification in normal times) and which part is not reasonable (e.g., price increase due to hoarding).

During Covid-19, the companies that had previously diversified their production, and were hence able to keep producing essential products during the pandemic, were not rewarded with higher profit. Instead, they were penalized by the increasing litigation risk of anti price gouging enforcement. As a result, diversifying and preparing for aggregate shocks may not a financially optimal decision for companies in classic market equilibrium models. Theoretically, the problem described above is inefficient allocation due to market incompleteness. Our model, with individual producers and the fixed-price assumption, reflects this issue.

The general public was also paying attention to price gouging. For example, a Google Images search on October 5th for "price gouging" presents the following results.

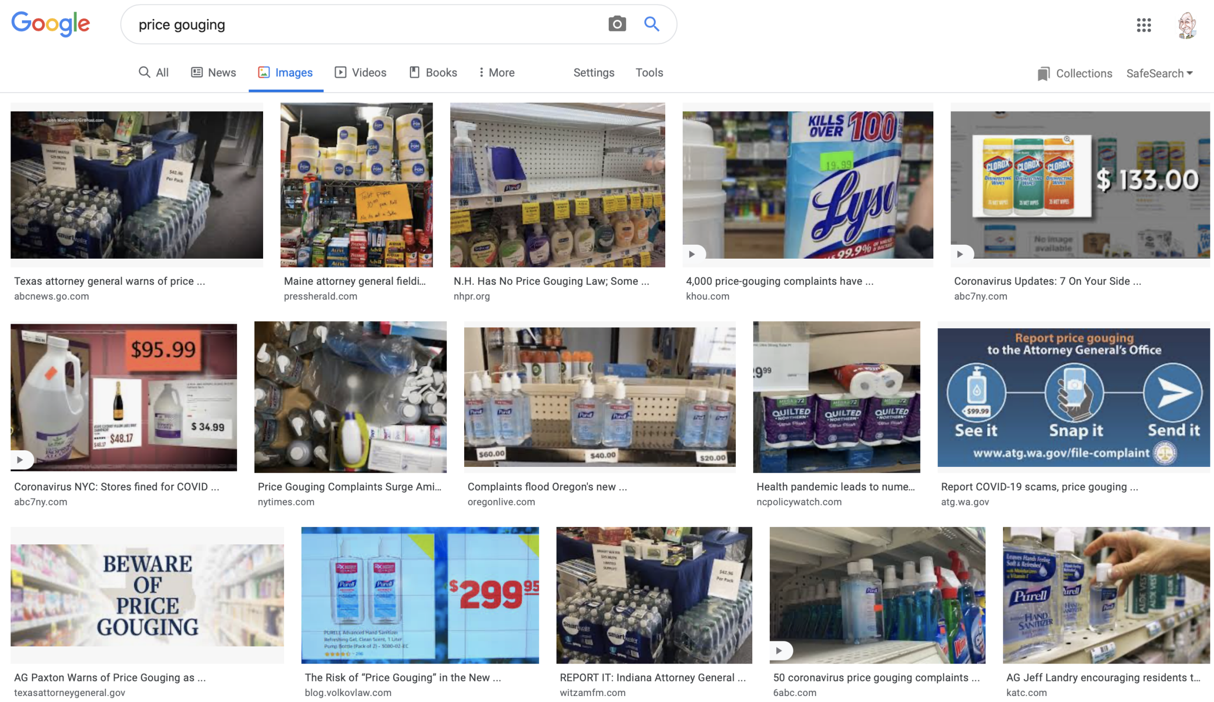

Figure 1.2: Google Search results for Price Gouging.

Some of the images are shown in detail in figure 1.3: consumers compare the price of isopropyl alcohol with a a bottle of champagne, a convenience store clarifies the price of toilet paper is not a joke, and water purchases are restricted.

When the consumers and law enforcement are so concerned with price increases after a natural disaster, is reasonable to expect that firms will prefer not to supply than to face the public relations nightmare that would require explaining the price needed to sell.

The examples we show here are only in the US, but Europe experienced a similar search intensity of firms violating price gouging laws. See Cary et al. (2020) and UK Competition Authority (2020) for a discussion of the recent law enforcement efforts regarding complains of price abuses in many developed nations.

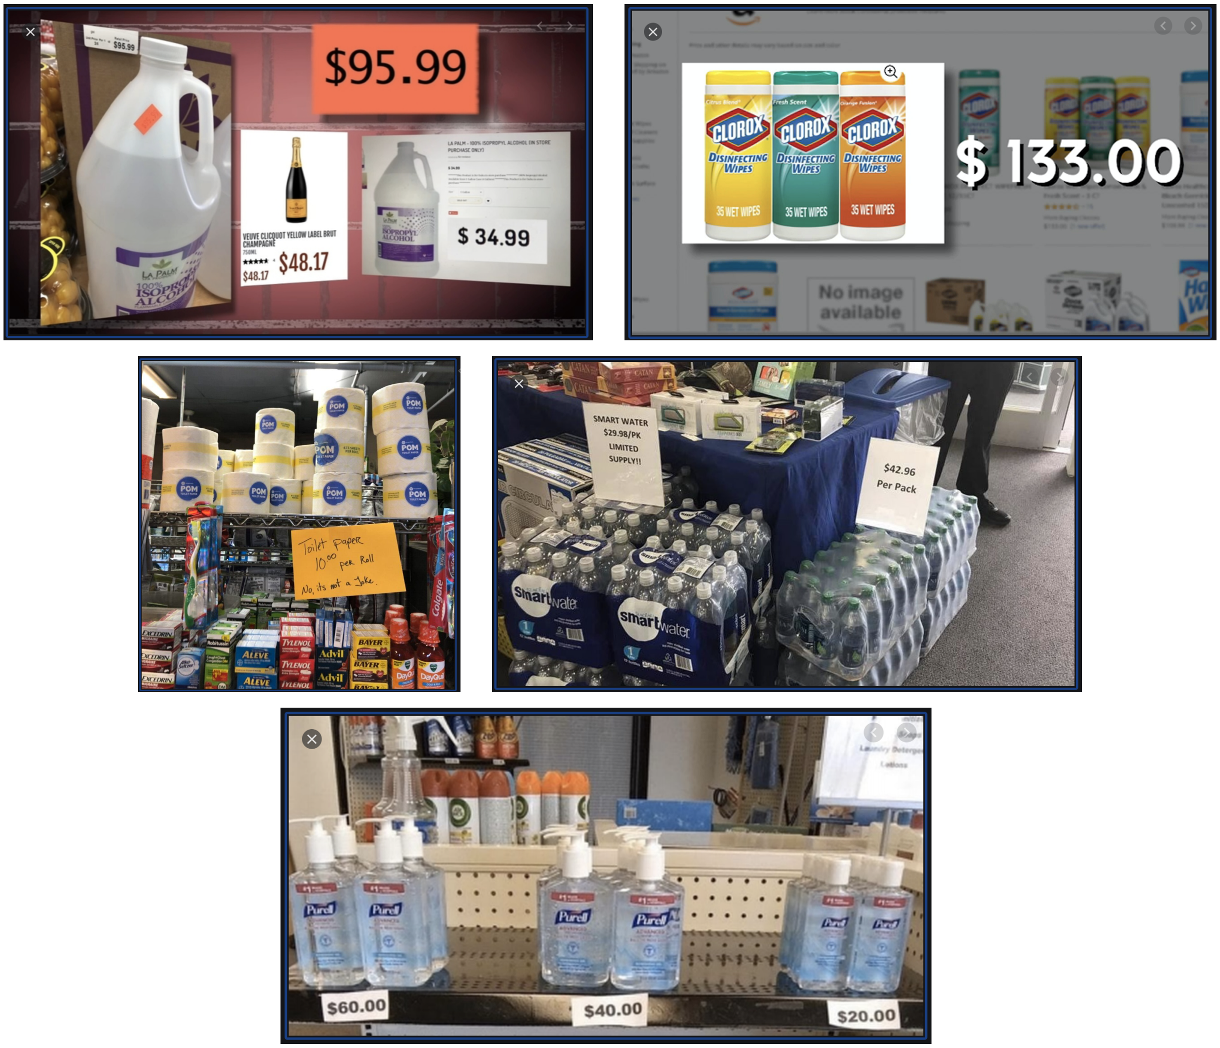

Figure 1.3: Price gouging examples in stores and online.

On the side of the general public, price gouging receives immediate attention at the start of the pandemic. And even though the images constitute anecdotal evidence, we can provide evidence on the intensity people were searching.

Figure 1.4 shows the Google search frequencies of topics "Coronavirus" and "Price Gouging" in different regions (US, UK, and worldwide) as well as different languages (English and Spanish). It is evident that the public's awareness of "price gouging" rises almost simultaneously with awareness of "Coronavirus" itself.

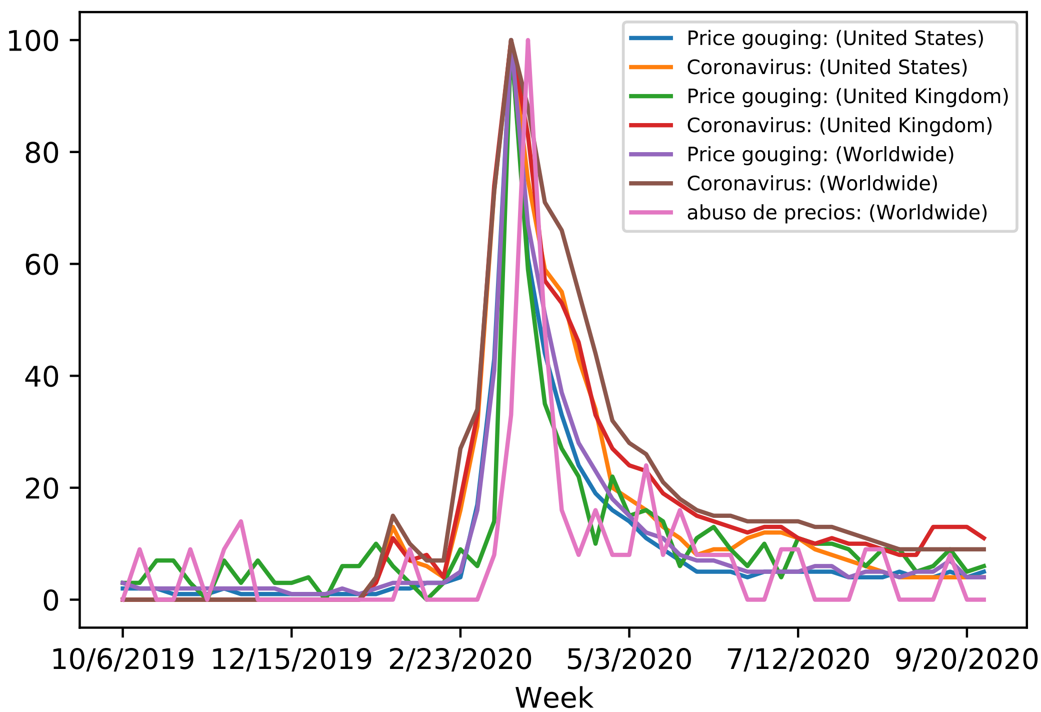

Figure 1.4: Google search frequencies of topics "Coronavirus" and "Price Gouging". Numbers are normalized by 100 at maximum values. *Data Source: https://trends.google.com/trends/*

1.2 Literature Review

Probability Matching: Our model provides a control problem rationalization of probability matching. When a game similar to ours is presented to individuals, experiments show that individuals tend to locate at the Valley roughly with probability θ, and locate at the Mountain roughly with probability $1-\theta$. This experimental result is known as probability matching in the literature[7]. Many different theoretical approaches to behavior are developed to explain this phenomenon that humans prefer probability matching over rational expected utility maximization. Some early work suggests that it is a behavioral limitation due to bounded rationality, but more recent literature attributes that to learning strategies. Vulkan (2000) and Gaissmaier and Schooler (2008) argue that people would consistently try to learn patterns of the outcome series in a repeated game even when they are informed that the series is completely i.i.d. As a result, if the outcome series is truly i.i.d. as many of the earlier models assume, then probability matching seems irrational. On the other hand, if there is, in fact, a pattern in the series, probability matching claims a higher expected reward in the long

[7]See, for example, Fiorina (1971), Morse and Runquist (1960), Vulkan (2000), and Brennan and Lo (2011)

run by gradually learning the patterns. Other literature suggests that probability matching is related to the growth pattern of a group of individuals. For example, Brennan and Lo (2011) concludes that if two choices result in similar growth rates, then deterministic decision rule prevails. On the other hand, if two choices result in drastically different growth rates, then probability matching gives an evolutionary advantage over the deterministic decision rules.

In our model, the intuition of the mechanism is simple: what is individually optimal is not globally optimal. The reason is that the survival of firms in our model has a global externality from the inability of firms to coordinate — the increase in the number of firms that could produce — that given our set up the decentralized market does not take into account. When the social planner solves the model, however, it internalizes this effect, and it forces firms to do something that looks locally irrational. As said before, locating firms in the Mountain provides insurance to the multinational when an aggregate shock to the Valley takes place. That insurance is extremely valuable when uncertainty is present.

Supply Chain: The literature on the supply chain is extensive and impossible to summarize in a few paragraphs. There are, however, aspects that have been discussed in the literature that are relevant to the model we present here.

Aggregate shocks like Covid-19 has a significant implication of global supply chain risk management. Earlier empirical research describes this as supply chain flexibility, see Vickery et al. (1999). That paper defines flexibility as the ability to adapt to aggregate shocks. It shows through correlation analysis that supply chain flexibility is critical to the long-run survival of an organization. On the other hand, flexibility may affect the immediate competitiveness of the firm in the short run.

The vulnerability of the global supply chain to identical suppliers has already raised some concerns in the industry. Wagner and Bode (2006) studied questionnaires from company executives in Germany and concludes that a firms dependence on single type customers and suppliers is the largest contributor to a firms exposure to supply chain risk.

In recent years, the question of whether to integrate suppliers or not has been receiving more and more attention to supply management. The existing empirical literature has been studying this issue by looking at the elasticity of substitution of produces, see Antràs and Chor (2013) and Alfaro et al. (2019). When the demand for the final product is elastic, and inputs are not substitutable, firms choose not to integrate upstream suppliers. On the other hand, when the demand for the final product is inelastic, and inputs are substitutable, firms choose to integrate upstream suppliers. This finding shows that firms supply chain decision is optimal for the deterministic case, but not necessary when aggregate shock hits.

Apart from the works highlighted above, two groups of literature align with the spirit of this paper.

First, there is ample literature discussing the organization of the supply chain. Following Antràs (2020), global supply chains can be viewed through different lenses. First, the value-added approach where firms allocate production internationally, and each stage of production contributes to the final product. In general, this literature concentrates on countries and industries as the unit of analysis. Second, the firm-level perspective — started by Melitz (2003) — offers an alternative to the aggregate view of the first approach. In this literature, the firms are the unit of analysis, and they are the ones that decide whether or not to participate in global supply chains. Both of these approaches assume there is no informational problem. This is relaxed by the relational view of supply chains. In this view, firms and suppliers face contracting problems — moral hazard or incomplete contracts — and therefore enter in relation to solve the informational problem. The main question it addresses is the organizational structure of the firm. The boundary of the firm in the global supply chain started with the seminal contribution of Antras (2003). The author discusses how incomplete contracts determine whether a firm should be integrated internationally versus enter in arms-length negotiations[8]. Finally, Yeaple (2003) studies the vertical and horizontal integration of multinationals.

Second, the literature on supply resilience highlights that resilience can be obtained by organizational robustness or organizational flexibility. See, for instance, Ambulkar et al. (2015), Töyli et al. (2013), Zhao and You (2019), Saenz et al. (2015), Durach and Machuca (2018), Helpman et al. (2004) and the references therin. Most of this literature, however, has two features. One is very related to our model — the literature advises that a robust supply chain can be achieved by working closely with the suppliers. In the spirit of our model, that is equivalent to when the multinational decides the global allocation problem. The second aspect is that most of these papers think about the robustness of a supply chain in response to shocks to the firms — i.e. the robustness to idiosyncratic shocks.

2 Model

In this section, we present a firm-location-problem model that highlights the vulnerability of the global supply chain to aggregate shocks. It is a simple survival model in which individual firms fail to take into consideration the impact they have on the aggregate — a standard externality argument — and whose decisions change quite substantially once uncertainty is taken into account.

We assume two different forms of organizing the world supply chain. In the first,

[8]See Antràs and De Gortari (2020), Antràs and Chor (2018) and Antras (2015).

a multinational asks already established firms (factories) to independently offer the parts required to produce a final product. In this case, the factories decide where to locate themselves. We identify this structure with a global supply chain of *Independent Suppliers* or the *Decentralized* economy. The second organization is one in which the multinational allocates its production facilities — which are the subsidiaries of the multinational. We identify this structure as *Multinational Subsidiaries* or as the *Centralized* economy.

As said before, a second ingredient in our model is the difference between risk and uncertainty. Optimization under risk produces a policy function that is very different from that derived under uncertainty. Our model is a single firm, partial equilibrium model, which concentrates on the existence and response of the supply chain to aggregate shocks. The Covid-19 pandemic was an obvious aggregate shock. However, natural and environmental disasters becoming more prevalent implies that we need a different approach to the understanding of resilience and robustness of the supply chain. A distinct feature of these shocks is their aggregate nature, but also how uncertain they are; we might know that sea level will be rising, but the extent of the damage has tremendous uncertainty, and the distribution itself is likely to be unknown.

We use both ingredients — the internalization of survival (the externality) and the robust approach (uncertainty) — to rationalize supply chains whose level of diversification are an order of magnitude larger than what we observe in practice.

2.1 Baseline Model

A product sold by a multinational is comprised of many "parts", each produced by a factory, and each factory can be located in two different regions. Time occurs in discrete steps, and the discount rate is β.

Assume there are N_t firms at the start of period t. Each factory/supplier has a location decision: for simplicity, we will identify the locations as the Mountain and the Valley. Factories choose one of these two locations at time t where they set up production. Each factory only produces one unit of the part, which has a constant cost c. Suppliers sell the part to the multinational, who produces the final good. The cost is paid before production takes place.

Production is uncertain. In each period, one of two locations might suffer an aggregate shock with arrival probability γ. Conditional on such a shock, and before production occurs, all firms in the Mountain or Valley perish with probability θ or $1-\theta$, respectively. We assume that $\theta >> 1/2$. In other words, the Mountain is significantly riskier than the Valley. With probability $1-\gamma$ there is no aggregate shock. Production takes place only by the surviving firms, and the multinational produces the final product with the parts it has access to[9].

At the end of each period, the number of subsidiaries can grow. The growth rate is given by

$$N_{t+1} = A \cdot (N_t^s)^{1-\mu} \tag{2.1}$$

where N_t^s denotes the number of firms that have survived the aggregate shock. Notice that this growth process has a fixed point at $N^* = A^{1/\mu}$. The timing is denoted in Figure 2.1.

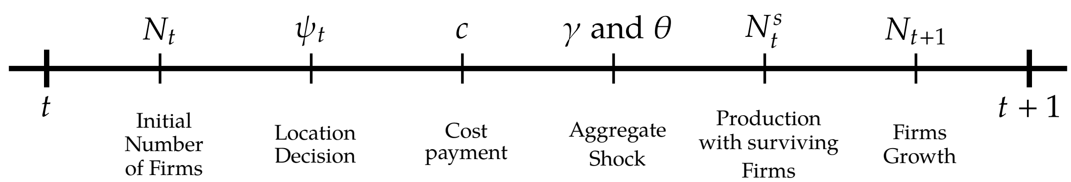

Figure 2.1: Model timing.

The multinational aggregates parts from all the suppliers and produces a final good. The complexity of the final good depends on the number of parts included. This model has an extremely simple demand side; we assume that the final product's revenue when sold is linear in the number of parts it includes. Furthermore, we add a *Survival Constraint* to this model by assuming that the firm needs at least 1 part to be able to produce the final product. In other words, $N_t^s \geq 1$ for the firm to be able to continue operating.

$$\Pi_t = \begin{cases} pN_t^s & \text{if } N_t^s \geq 1 \\ 0 & \text{o.w.} \end{cases}$$

The price per part, p, is constant and independent of production and the state of the world. This is equivalent to assuming that the multinational is a price taker, but it also is capturing the fact that prices rarely move freely after natural disasters. This is a very simple demand that simplifies the exposition.

One crucial reason why prices might not adjust after aggregate shocks is the presence of price gouging. In many countries, there are price gouging regulations that limit the extent to which the pricing system helps ameliorate the supply chain problem. Therefore, the pricing system cannot finance the supply chain reforms required to reestablish production, and instead other actions (such as time or government subsidies) are required to recover the supply chain. During the Covid-19 pandemic, it has been clear that fairness arguments have dominated the discussion. For example, see Executive Order 13910 of March 23, 2020,

[9]The reason behind this assumption is that in the canonical model of complementary inputs (e.g., Kremer (1993)) an idiosyncratic shock has macroeconomic consequences. In our model, we want idiosyncratic shocks to be "harmless" and concentrate on the role of aggregate shocks.

Preventing Hoarding of Health and Medical Resources To Respond to the Spread of COVID-19, or King and Spalding (2020) for a list of price gouging laws in the US. A more detailed discussion of price gouging during Covid-19 was given earlier in Subsection 1.1.

Denote ψ_t the proportion of firms that are located in the Valley. The evolution of firms is given by

$$N_{t+1} = \begin{cases} A \cdot (N_t)^{1-\mu} & \text{w/p} \quad (1-\gamma) \\ A \cdot (\psi_t N_t)^{1-\mu} & \text{w/p} \quad \gamma\theta \\ A \cdot ((1-\psi_t)N_t)^{1-\mu} & \text{w/p} \quad \gamma(1-\theta) \end{cases}$$

where the top realization occurs when there is no aggregate shock, and the second (third, respectively) one is when the aggregate shock hits the Mountain (Valley, respectively). The growth of the firms has two components: multiplicative and exponential. As can be seen, the growth of suppliers depends on the total number of surviving suppliers in the world. We assume that the new suppliers are distributed according to the existing number of surviving firms in each location, but that the growth rate depends on the total number of existing firms. This assumption in the basic model is innocuous, but it is essential if the model is extended to introduce adjustment costs — or switching costs. We leave this interesting application for future research.

2.1.1 Independent Producers

In the independent producers setting, the suppliers decide their location individually, and then the multinational contracts with the firms. We assume that all the revenue from the multinational is transferred to the suppliers — i.e. the multinational has zero profits. The total revenue is equally shared among the surviving suppliers.

Suppliers are small and they do not take into account the impact their decision has on the decision of the location of others (ψ_t). As we mentioned before, there is no cost of switching between locations. Therefore the suppliers are solving a static problem — the continuation value is exactly the same for all firms. Firms are maximizing the expected value of Mountain versus Valley and given our assumptions Valley dominates for all firms. Then, the value at time t of locating in the Valley or the Mountain is given by

$$V_t^v = ((1-\gamma)+\gamma\theta)p \qquad - c + \frac{1}{1+\beta}((1-\gamma)+\gamma\theta)V_{t+1} \tag{2.2}$$

$$V_t^m = ((1-\gamma)+\gamma(1-\theta))p - c + \frac{1}{1+\beta}((1-\gamma)+\gamma(1-\theta))V_{t+1}. \tag{2.3}$$

The continuation value for each supplier, conditional on having survived the aggregate shock, is independent of the location. This is a feature of the zero cost of relocation. Therefore, the value of locating in the Valley is always larger than the value of locating in the Mountain. Formally,

$$V_t^v - V_t^m = \gamma(2\theta - 1)\left(p + \frac{1}{1+\beta}V_{t+1}\right) > 0 \tag{2.4}$$

for $\theta > 1/2$, and hence $\psi_t = 1$.

2.1.2 Multinational Subsidiaries

Assume now that the multinational has all the decision power and it allocates the production units. Two aspects now matter for the multinational firm that were not relevant for the independent suppliers: the multinational takes into account the distribution of firms, and it takes into account the expected value of continuation in all states of the world.

The problem of the multinational firm can be written as follows

$$V(N_t) = \max_{\psi_t}\left\{\left[\begin{array}{cccc} (1-\gamma) & \cdot & (pN_t & +\frac{1}{1+\beta}V\left(A\cdot(N_t)^{1-\mu}\right) \quad) \\ +\gamma\theta & \cdot & (p\psi_t N_t & +\frac{1}{1+\beta}V\left(A\cdot(\psi_t N_t)^{1-\mu}\right) \quad) \\ +\gamma(1-\theta) & \cdot & (p(1-\psi_t)N_t & +\frac{1}{1+\beta}V\left(A\cdot((1-\psi_t)N_t)^{1-\mu}\right) \quad) \end{array}\right] - cN_t\right\} \tag{2.5}$$

where

$$\lim_{N\to 1^-} V(N) = 0 \tag{2.6}$$

is the value matching constraint.

As before, the top line represents the value when the aggregate shock does not occur, and the second (third, respectively) when the aggregate shock hits the Mountain (Valley, respectively). Although the cost is the same cN_t, the revenue depends on the number of surviving firms. Recall that the cost of production is paid irrespectively of the aggregate shock.

The first-order condition (FOC) with respect to ψ_t, after simplifying, is

$$(2\theta-1)p = A(1-\mu)\frac{N_t^{-\mu}}{1+\beta}\left\{(1-\theta)(1-\psi_t)^{-\mu}V'\left(A((1-\psi_t)N_t)^{1-\mu}\right) - \theta\psi_t^{-\mu}V'\left(A(\psi_t N_t)^{1-\mu}\right)\right\}. \tag{2.7}$$

We simulate the discrete time version of the model to characterize the solution[10]. The parameters used in the simulation are: $\beta = 0.02$, $\gamma = 0.2$, $\theta = 0.6$, $\mu = 0.5$, $A = 5$, $p_m = p_v = 1$ and $c_m = c_v = 0.5$.

In terms of the number of suppliers in equilibrium, the choice of $\mu = 0.5$ implies a

[10] We present the derivation of a continuous time version of the model in Appendix A.

fixed point of $N_t = 25$, in the absence of aggregate shocks. We initialize all simulations in this fixed point.

In Figure 2.2, we present the proportion of firms in the Valley as a function of the total number of suppliers (horizontal axis). The orange line represents the decentralized allocation — the individual rationality solution. The blue line indicates the probability matching solution for survival. Finally, The green line indicates the optimal solution of the multinational.

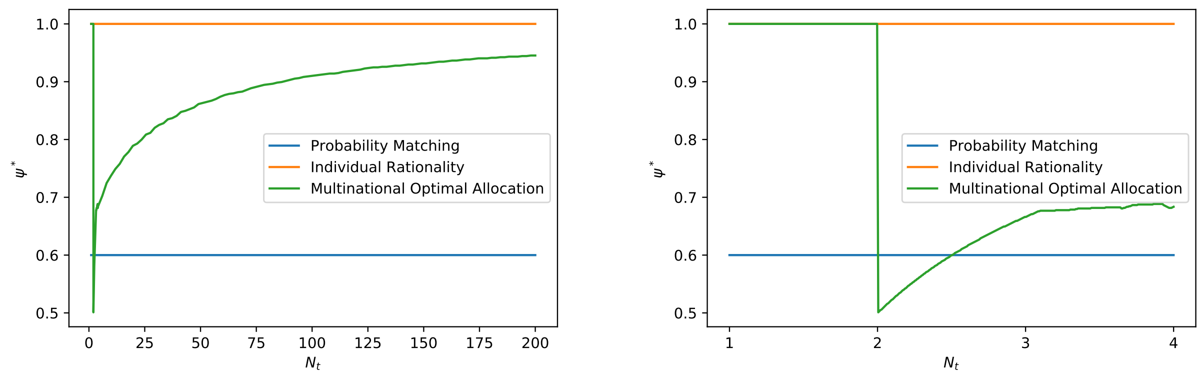

Figure 2.2: Optimal ψ^* as a function of N_t.

The multinational has a tradeoff between instantaneous profits (what the individual suppliers maximize) and the probability of survival. The right panel of Figure 2.2 is a closer view of the left panel, but concentrating on relatively small N.

As shown in Figure 2.2, the multinational's optimal ψ is a function of the number of production units N_t, and has three distinct phases. In the first phase, when $N_t \in [1, 2)$, the multinational's optimal choice is a corner solution, which coincides with individuals' optimum. This occurs because when $N_t < 2$, losing one unit will discontinue the multinational's operation. Hence, there is no way to ensure survival and to realize the benefit of continuation value. In the second phase, when $N_t \in [2, 3)$, the multinational will allocate exactly one production unit to the Mountain to take advantage of the continuation value. As a result, the optimal allocation is given by $\psi^* = 1 - 1/N_t$. In the third phase, when $N_t \geq 3$, the optimal ψ is a concave increasing function of N_t. It is increasing because, with guaranteed survival, it is optimal to allocate a greater percentage of production units to the Valley to maximize profit. It is concave because the function $\psi^*(N_t)$ asymptotically approaches a constant less than 1.

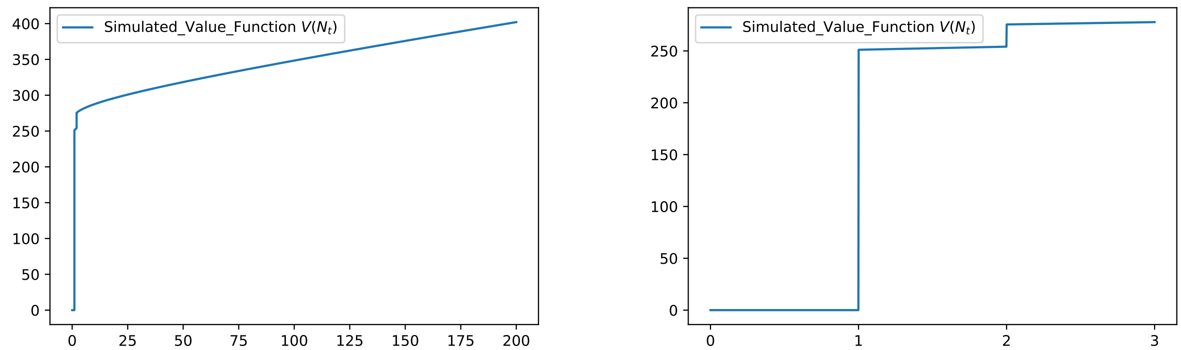

Figure 2.3: Value function $V(N_t)$ with optimal ψ.

The value function of the multinational, with the optimal ψ, is presented in Figure 2.3. The concavity of the value function comes from the concavity of the growth function of the suppliers, and also from the value matching constraint. If the suppliers grow at a constant rate, the value function would be linear with respect to the number of suppliers — and therefore, even the solution of the multinational would be at a corner.

The value function has two discontinuities points at $N_t = 1$ and $N_t = 2$. The discontinuity at $N_t = 1$ is trivial due to the constraint that $V(N_t) = 0, \forall N_t < 1$. On the other hand, the discontinuity at $N_t = 2$ has an important implication about continuation value. When the number of production units is less than 2, the multinational firm cannot ensure survival. When the number of production units is greater than or equal to 2, the multinational can allocate one production unit on the Mountain to ensure survival. This ensured survival creates a continuation value for the multinational, which is responsible for the jump of value function at $N_t = 2$.

2.2 Probability Matching

As was shown in Figure 2.2, the optimal proportion of firms in the Valley varies with the number of suppliers. It is interesting to compare the effectiveness of the optimal firm allocation with respect to a naive strategy — a constant psi independent of the number of firms. For instance, assuming the probability matching strategy is adopted, the value function of the multinational with constant $\psi = 0.6$ is presented in Figure 2.4.

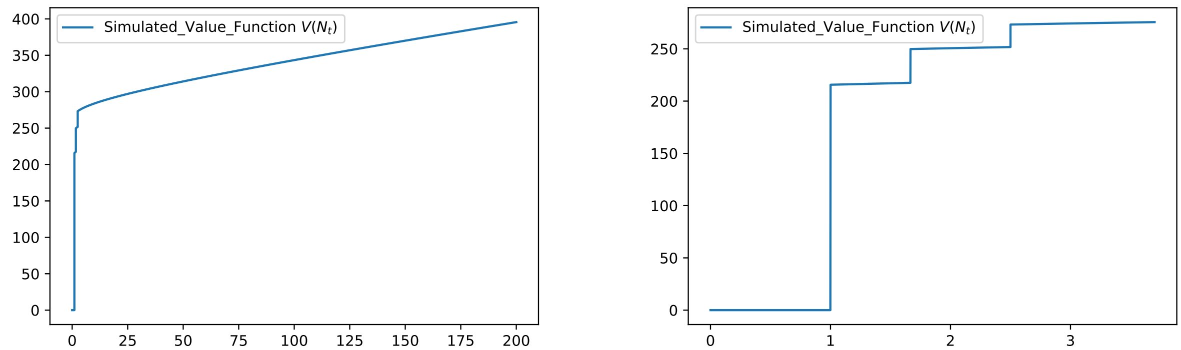

Figure 2.4: Value function $V(N_t)$ with constant $\psi = 0.6$.

Apart from the trivial discontinuity at $N_t = 1$, the value function has two other discontinuous points at $N_t = 5/3$ and $N_t = 2.5$. Below 5/3, the constant $\psi = 0.6$ does not place a single production unit in either location, and the multinational fails when any aggregate shock occurs. As a result, the discontinuity at 5/3 represents the continuation value of firms in the Valley. On the other hand, the discontinuity at $N_t = 2.5$ represents the continuation value of firms in the Mountain. When the number of available production units is less than 2.5, the multinational places fewer than one unit in the Mountain, and will not survive an aggregate shock to the Valley. However, when the number of production units is greater than or equal to 2.5, the multinational can allocate at least one production unit on the Mountain, thereby ensuring survival. The jump of the value function at $N_t = 2.5$ reflects this guaranteed survival.

The difference between the value function using the optimal strategy, and the value function following the probability matching strategy is small. We compare the optimal-strategy value function with the probability-matching value function (when ψ is constant and equal to 0.60). Figure 2.5 shows the percentage increase in the value function when the firm switches between probability matching to optimal.

The x-axis is the number of firms on a logarithmic scale, and for comparison purposes, we concentrated on $N_t > 3$. The discrete jumps in the value function for smaller N_t swamp any possible comparison outside that region. On the y-axis is the percentage difference between the two value functions.

The relationship, as expected, is increasing. The reason is that the optimal ψ increases with the number of surviving firms; therefore, the loss incurred by fixing it at 0.60 is also increasing. Having said this, notice that the magnitudes are small: between one and two percent.

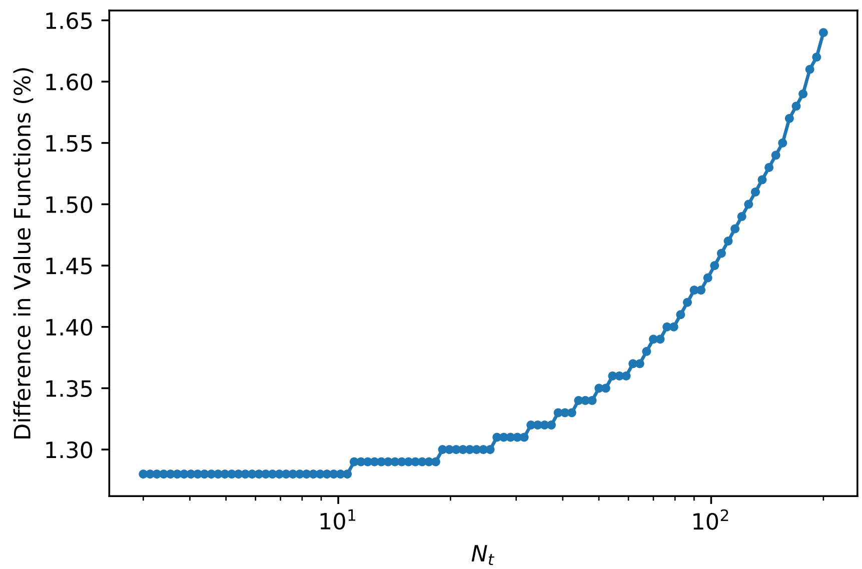

Figure 2.5: Difference between Value function $V(N_t)$ with optimal ψ and Value function with constant $\psi = 0.6$.

One interesting question to ask is how the probability of survival is affected by different possible allocation strategies by the multinational. Recall that Figure 2.2 indicates that the optimal proportion of firms in the Valley is a function of the total number of suppliers that exist. However, we here study a simple, naive allocation strategy in the spirit of probability matching models. For instance, assume the multinational chooses a fixed proportion of suppliers in the Valley regardless of the total number of suppliers that exist.

In Figure 2.6, we present the probability of survival for various fixed values of ψ over different horizons. We define the probability of survival as one minus the probability that the number of suppliers is smaller than 1 for any time step within the horizon, and initialize simulations with the same parameters as above.

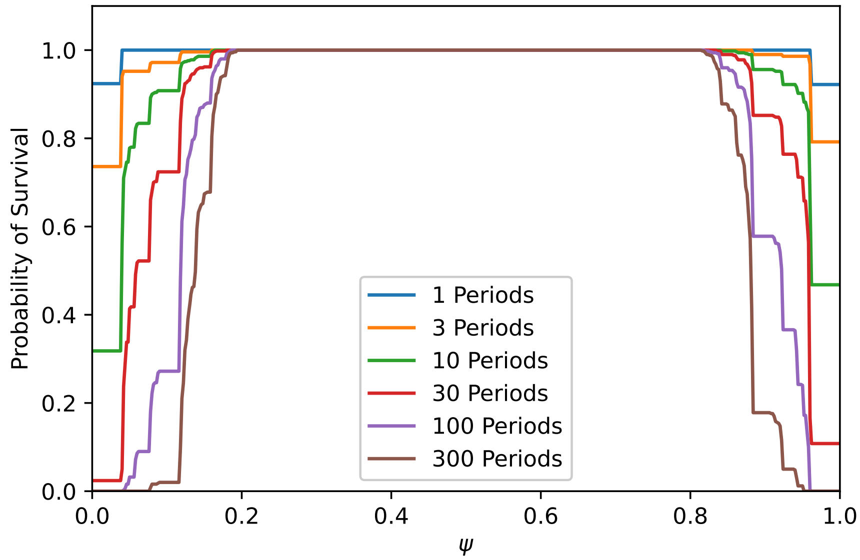

Figure 2.6: Probability of survival

For large time horizons, the probability is either one or zero; interestingly the break-points include the probability matching proportion ($\psi = 0.6$).

2.3 Discussion

This model presents a simple contrast between three possible strategies: the decentralized allocation in which firms do not take into account the survival probability of the multinational, the centralized allocation in which the multinational internalizes the decision, and the behavioral response that would use simple probability matching heuristics concentrating on the maximization of the probability of survival.

The behavioral finance literature points out many cases in which individuals will tend to chose the third strategy. In our model, indeed, such a strategy will guarantee the survival of the multinational. However, it is inefficient. A dynamic allocation increases profits, for example, and also guarantees the survival of the firm in equilibrium.

Many questions arise from this framework that we explore further in this paper and some that are left for future research.

First, how can the probability matching behavior be rationalized in this setting? As shown in figure 2.2 the optimal allocation in the Valley is an increasing function of the number of surviving firms - after $N_t > 2$. So, the optimal solution is an internal solution, and it is dependent on the number of firms. As we discussed before, some jumps happen at small numbers, which are the result of the constraint at which the multinational shuts down. Probability matching implies a constant proportion of firms regardless of the number of surviving firms, which contradicts this feature of the multinational's optimal allocation. In Section 3 we will introduce uncertainty and show how the robust optimal response is indeed very close in its spirit to the probability matching.

Second, given the externality, is there something that a government could do such that the decentralized economy could reproduce the centralized allocation? The answer to this question is usually yes, and it either requires taxes or subsidies. We study this point in Section 4. We show that in the context of differences in prices, there exists a governmental policy in which the decentralized economy achieves the centralized outcome — or at least — the firms are indifferent between locating in the Mountain versus the Valley, so the centralized allocation is available. The price or cost differences can be interpreted as an *ex-ante* tax or subsidy to the allocation of a firm in a particular destination.

Although trivial, it is worth highlighting that if prices are allowed to adjust, the decentralized allocation will replicate the centralized one. In our model, prices are not allowed to change, and therefore it is impossible to compensate the firms in the Mountain when a shock to the Valley has taken place. If the prices were to adjust, then once an aggregate shock takes place, the revenues of the surviving suppliers would increase. Moreover, because there are fewer firms in the Mountain, the price increase when a shock to the Valley takes place would increase the price of parts more than when the shock occurred in the Mountain. Because survival is very important, any usual demand function — CES for example — implies that the expected value of firms in the Mountain and the Valley are equalized. Under those circumstances, the decentralized economy reproduces the centralized one. The assumption of prices NOT adjusting is crucial. We do believe it is a reasonable assumption when aggregate shocks occur, justified through the prevention of price gouging.

Third, a simplifying assumption in our model is that the growth of firms is related to the total number of surviving firms regardless of where the firms were located. Also, we are assuming that there is no cost of reallocation. These are simplifying, but unreasonable assumptions. Further research should look into the implications when the growth of firms is specific to the location, and there are adjustment costs. We leave this extension to future research.

Finally, our supply chain structure is extremely simple. In reality supply chains look like complex networks[11]. Future research should look at the implications of robustness in a more complex structure.

[11]See Yeaple (2003)

3 The *Nominal* and the *Robust* Models

The model in the previous section only deals with risk. In this section, we explore the implication of adding uncertainty into the model. In particular, we will assume the probability of the aggregate shock in the Valley (θ) is uncertain. As has been said before, this section shows that the optimal robust strategy is exactly in the spirit of probability matching: a constant proportion regardless of the number of firms that exist.

We study two cases: in the first, we consider a risky θ. Here we assume that θ is uniformly distributed in $[\bar{\theta}-\Delta, \bar{\theta}+\Delta]$. Due to the linearity of the value function with respect to θ, we show that the optimal choices for both the individual producers and the multinationals are the same as the baseline case. Though the result is trivial in this case, it allows us to set up the comparison with our second case, where we assume model uncertainty: θ can be any value between $[\bar{\theta}-\Delta, \bar{\theta}+\Delta]$ and agents deploy a robust decision rule by solving a minimax problem.

3.1 Risk: the *Nominal* Model

Let us assume that $\theta \in [\bar{\theta}-\Delta, \bar{\theta}+\Delta]$ where Δ is small enough to guarantee that the support of θ is contained in $[0, 1]$. We assume that the distribution is uniform and known by all agents. This setting is identified in our discussion as the *nominal* model. We will continue to assume that the Mountain is riskier, therefore, $\bar{\theta} > 1/2$. Because individual suppliers are risk-neutral and the aggregate shock enters linearly, the decentralized equilibrium is identical: all the suppliers choose to locate in the Valley.

Similarly for the multinational, because θ enters linearly to the value function, the profit maximization problem and its solution is unchanged. More specifically

$$
\begin{aligned}
V(N_t) &= \max_{\psi_t(N_t)} \int_{\bar{\theta}-\Delta}^{\bar{\theta}+\Delta} \left\{ \begin{array}{lcll} (1-\gamma) & \cdot & \left(pN_t + \frac{1}{1+\beta} V\left(A \cdot (N_t)^{1-\mu}\right)\right. & \left.\right) \\ +\gamma(\theta) & \cdot & \left(p\psi_t N_t + \frac{1}{1+\beta} V\left(A \cdot (\psi_t N_t)^{1-\mu}\right)\right. & \left.\right) - cN_t \\ +\gamma(1-\theta) & \cdot & \left(p(1-\psi_t)N_t + \frac{1}{1+\beta} V\left(A \cdot ((1-\psi_t)N_t)^{1-\mu}\right)\right. & \left.\right) \end{array} \right\} d\theta \\
&= \max_{\psi_t(N_t)} \left\{ \begin{array}{lcll} (1-\gamma) & \cdot & \left(pN_t + \frac{1}{1+\beta} V\left(A \cdot (N_t)^{1-\mu}\right)\right. & \left.\right) \\ +\gamma(\bar{\theta}) & \cdot & \left(p\psi_t N_t + \frac{1}{1+\beta} V\left(A \cdot (\psi_t N_t)^{1-\mu}\right)\right. & \left.\right) - cN_t \\ +\gamma(1-\bar{\theta}) & \cdot & \left(p(1-\psi_t)N_t + \frac{1}{1+\beta} V\left(A \cdot ((1-\psi_t)N_t)^{1-\mu}\right)\right. & \left.\right) \end{array} \right\}
\end{aligned}
$$

Therefore, under the assumption of risk the solutions of the decentralized and decentralized economy are identical to the baseline model. Of course this is a feature of the assumptions we have chosen to make and where the parameter risk was introduced. We have made these choices for simplicity.

3.2 Uncertainty: The *Robust* Model

The second case we study is the case of uncertainty in the sense of robust control. Assume that all agents know that $\theta \in [\bar{\theta} - \Delta, \bar{\theta} + \Delta]$, but they do not know the distribution.

Optimal control implies that the optimization maximizes the worst possible case. For individual producers, the values in the Valley and the Mountain are

$$V_t^v = \min_{\delta \in [-\Delta, \Delta]} ((1-\gamma) + \gamma(\bar{\theta} + \delta))p \quad - c + \frac{1}{1+\beta}((1-\gamma) + \gamma(\bar{\theta} + \delta))V_{t+1} \tag{3.1}$$

$$V_t^m = \min_{\delta \in [-\Delta, \Delta]} ((1-\gamma) + \gamma(1 - \bar{\theta} - \delta))p - c + \frac{1}{1+\beta}((1-\gamma) + \gamma(1 - \bar{\theta} - \delta))V_{t+1} \tag{3.2}$$

Because nature will choose a value of Δ that minimize producers' value, the worst case scenario for the Valley is when $\delta = -\Delta$, and the worst case scenario for the Mountain is when $\delta = \Delta$. The difference in value between the Valley and Mountain is given by

$$V_t^v - V_t^m = \gamma(2\bar{\theta} - 1)\left(p + \frac{1}{1+\beta}V_{t+1}\right) > 0 \tag{3.3}$$

for $\bar{\theta} > 1/2$. Note this is identical to the baseline case (Equation 2.4 in Section 2. The worst case for each individual firm still implies that the worst case in the Valley is better than the worst case in the Mountain.

The problem of the multinational firm can be written as follows

$$V(N_t) = \max_{\psi_t(N_t)} \min_{\delta \in [-\Delta, \Delta]} \left\{ \begin{bmatrix} (1-\gamma) & \cdot & \left(pN_t \right. & + \frac{1}{1+\beta} V\left(A \cdot (N_t)^{1-\mu}\right) & \left.\right) \\ +\gamma(\bar{\theta} + \delta) & \cdot & \left(p\psi_t N_t \right. & + \frac{1}{1+\beta} V\left(A \cdot (\psi_t N_t)^{1-\mu}\right) & \left.\right) \\ +\gamma(1 - \bar{\theta} - \delta) & \cdot & \left(p(1-\psi_t)N_t \right. & + \frac{1}{1+\beta} V\left(A \cdot ((1-\psi_t)N_t)^{1-\mu}\right) & \left.\right) \end{bmatrix} - cN_t \right\}$$

subject to

$$\lim_{N \to 1^-} V(N) = 0.$$

The above value function has the following characteristics: First, given any fixed δ, the value function is upper semi-continuous and concave in ψ, for all $N_t \geq 2$. For $N_t < 2$ there is a unique corner solution. Second, given any fixed ψ, the value function is linear (therefore continuous and convex) in δ. Finally, $\psi \in [0, 1]$ and $\delta \in [-\Delta, \Delta]$ are chosen from compact sets.

As a result, according to Sion's Minimax Theorem Sion et al. (1958), the maximization and minimization are interchangable, and the minimax problem has at least one solution.

The optimal δ (meaning the choice that produces the worst possible case for the multinational) is given by

$$\delta^* = \begin{cases} -\Delta & \text{if } \bar{\theta} - \Delta > 1/2 \\ -\bar{\theta} + 1/2 & \text{if } \bar{\theta} - \Delta <= 1/2 \end{cases} \tag{3.4}$$

This then implies that the multinational's optimal response is the ψ from Figure 2.2, but where the shock probability is given by $\theta = \bar{\theta} + \delta^*$.

3.3 Efficiency vs. Robustness

It may seem too conservative to always considering the worst-case scenario, especially for multinationals trying to maximize profit. Even though the worst-case scenario is known to be $\delta = -\Delta$ for small enough Δ, this scenario may not be considered by the agents in the system who seek efficiency. Naturally, an efficiency versus robustness tradeoff emerges.

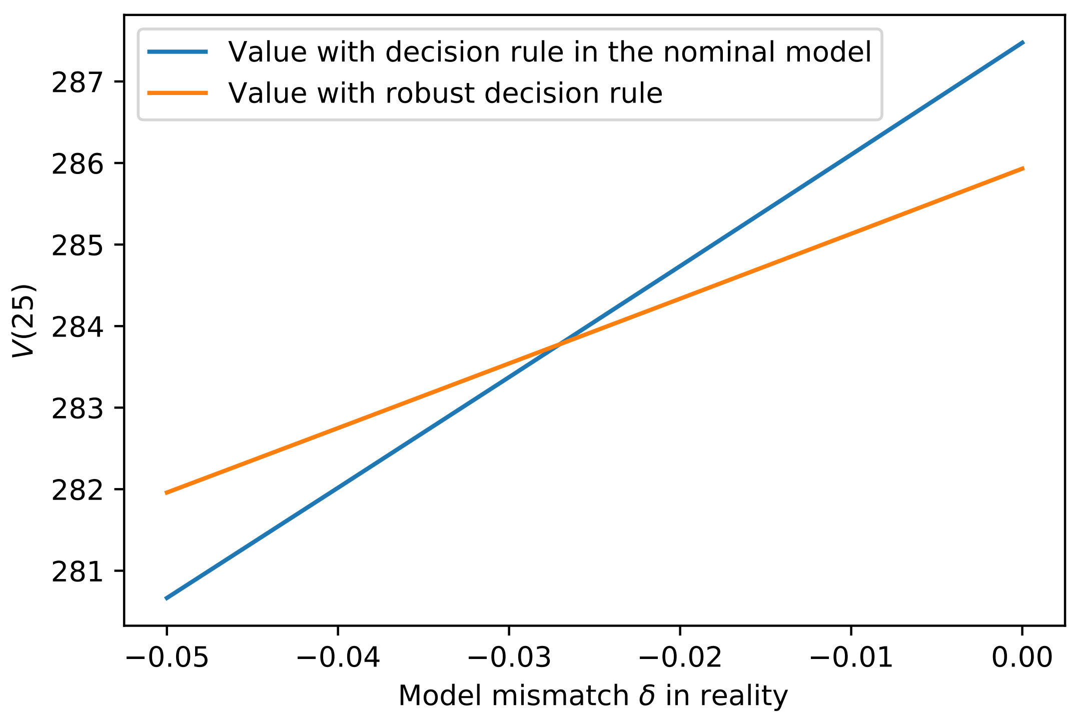

(a) Δ=0.05

Figure 3.1: Efficiency versus Robustness. $\Delta = 0.05$ Efficiency is represented by the value in nominal model, and robustness is represented by the value in worse-case model. The number of firms N_t is 25.

Figure 3.1 compares the case of the nominal and worst-case models. The blue line represents the value function for the nominal model, and the orange the value function for the robust model. The simulations are constructed assuming that the uncertainty parameter $\Delta = 0.05$, and that $\bar{\theta} = 0.6$. So, the range for θ is between $[.55, .65]$. The nominal model optimizes as if $\delta = 0$, and we plot what the realized value function for $\delta < 0$ instead being zero. In other words, when $\delta = 0$ the blue line reaches the maximum because the real value of theta is exactly the one used by the multinational to optimize. on the other extreme (left) when the value of $\delta = -0.05$ the multinational makes choices thinking that the relative shock parameter is $\bar{\theta}$ when it actually is $\bar{\theta} - 0.05$. Therefore, the nominal value function is subject to potential losses.

The orange line is the robust model. Given the assumption of the bounded range, we know that the multinational assumes that the $\theta = \bar{\theta} - 0.05$. In this case, notice that the orange line is flatter, and the worst-case is better than when the nominal model chooses. In fact, the robust model is optimal when $\delta = -0.05$.

Figure 3.2 depicts the case when $\Delta = 0.15$. In this case, the size of the uncertainty is large enough that the case of $\theta = 0.5$ is in the support. According to equation 3.4 the robust approach assumes that $\theta = 0.5$.

Notice that the blue line — the nominal model — behaves similarly as in the previous case, except that a larger range implies bigger potential losses. Again, in the extreme left, the multinational assumes that $\theta = 0.6$ when it actually is 0.45. In contrast, the orange line assumes a $\theta = 0.5$ and produces a much flatter value function.

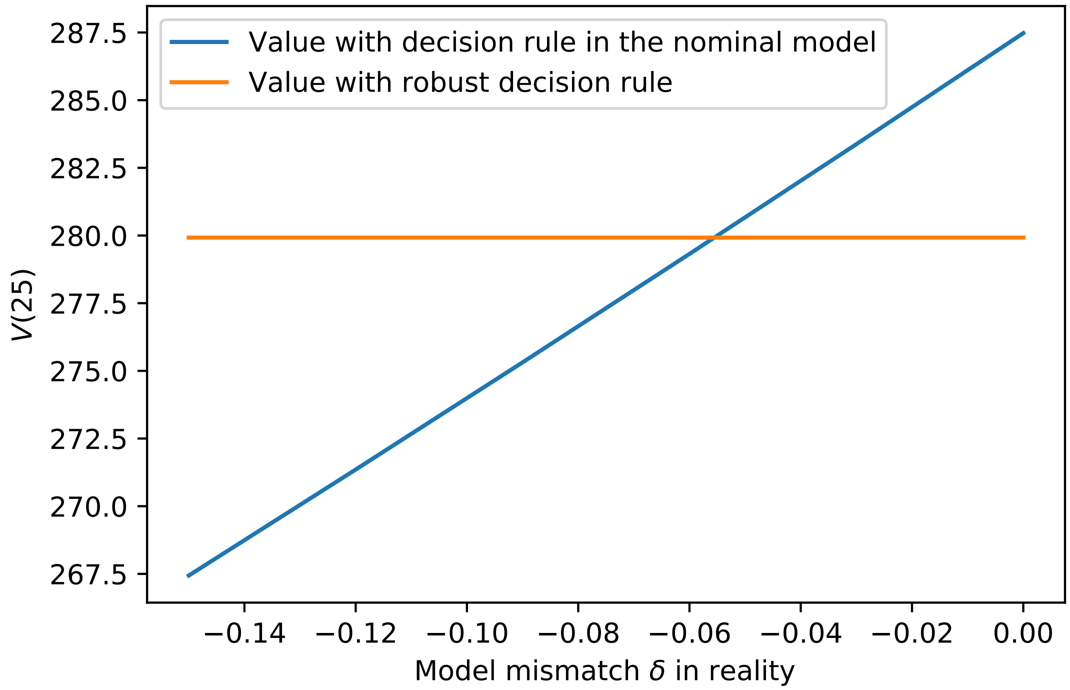

(b) Δ=0.15

Figure 3.2: Efficiency versus Robustness. $\Delta = 0.15$. Efficiency is represented by the value in nominal model, and robustness is represented by the value in worse-case model. The number of firms N_t is 25.

These two figures highlight the implications of applying robustness to a decision problem. Robustness is needed when the agents do not know the distribution of the shock they are facing, and therefore need to prepare for the worst. Robustness, then, serves to find a policy that reduces the differences over all possible states of nature. In the limit, the most robust action is one in which the outcomes are identical in all states of nature (as shown in Figure 3.2)

3.4 Discussion

Robustness is often compared to the solution of a model with large risk aversion. In some applications, that is indeed the case. In our model, robustness changes the nature of the solution — the optimal allocation is independent of the number of firms, and the allocation in Mountain and Valley is symmetric. Our model, implicitly, rationalizes probability matching, which is considered a behavioral result in finance.

This result needs two ingredients that are quite specific to our model: robustness, and a centralized decision maker that internalizes survival probabilities. Table 3.1 summarizes the relationships between the modeling choices and the characteristics of the policy function.

Table 3.1: Relationship between modeling choices and characteristics of the policy function.

	Risk	Uncertainty
Decentralized	Corner Solution	Corner Solution
Centralized	Internal Solution $\psi(N_t)$	Probability Matching $\psi'(N_t) = 0$

When the economy is populated with decentralized decision makers, the optimal policy choices are independent of the nature of the shock. In other words, individual suppliers will choose a corner solution (e.g. locating exclusively in the Valley) irrespectively of whether they are facing risk or uncertainty. A centralized decision maker, on the other hand, tends to prefer internal solutions. When they face risk, the optimal allocation of firms is a function of the number of firms – which contradicts with the probability matching patterns that are often found in experimental literature. It is when both uncertainty and centralization are present that the solution is an internal and fixed ratio; very much in the spirit of probability matching.

This result has important implications for the supply chain. In our model, irrespectively of how bad the Mountain is relative to the Valley, there is a level of uncertainty for which the multinational allocates half the firms in the Mountain. A robust supply chain is one in which the survival probability is maximized, and where the firm will be able to produce even in the worst of circumstances. Of course, this is a result that depends on the underlying assumptions of the model, but the intuition should be easy to extend to more realistic circumstances: if the supply chains in the world would have been prepared to supply goods in the worst possible circumstance, then the Covid-19 shock should have produced zero stockouts. A supply chain that deals with risk but optimizes using the "expected" value is found to be ill-prepared to handle an aggregate shock.

4 Price and Cost Differences

Up to now, we have not allowed prices to change depending on the state of the world. This is clearly a simplification that has allowed us to characterize the solution of the model and study the price gouging case. This section, in practice, relaxes the price gouging assumption[12]. The conclusion so far is that even when prices do not adjust — making the decision to locate

[12] See section 1.1 for the justification.

the suppliers in the Mountain a less profitable decision — the values of continuation and robustness are enough for the multinational to allocate firms in the Mountain. In this section we study the implications of allowing prices and costs in the two locations to be different.

4.1 Model

We assume that there are heterogeneous costs $c_v \neq c_m$ and prices $p_v \neq p_m$ in the two locations. We use this model to address many different questions. First, can the government define an intervention (either reducing the cost of the Mountain, or increasing its price) for which the decentralized economy achieves the social optimum - even in the case of uncertainty? Second, what is the profit margin of the Mountain at which the robust control strategy ceases to diversify the supply chain? In other words, when is robustness undesirable?

The value functions for the individual suppliers are given by

$$V_t^v = ((1-\gamma)+\gamma\theta)p_v \qquad - c_v + \frac{1}{1+\beta}((1-\gamma)+\gamma\theta)V_{t+1} \tag{4.1}$$

$$V_t^m = ((1-\gamma)+\gamma(1-\theta))p_m - c_m + \frac{1}{1+\beta}((1-\gamma)+\gamma(1-\theta))V_{t+1} \tag{4.2}$$

The continuation values still are identical in each of the two locations because there is no cost of relocation of suppliers. The difference between the two locations is given by

$$V_t^v - V_t^m = (((1-\gamma)+\gamma\theta)p_v - c_v) - (((1-\gamma)+\gamma(1-\theta))p_m - c_m) + \gamma(2\theta-1)\left(\frac{1}{1+\beta}V_{t+1}\right) \tag{4.3}$$

There is an expected markup in the Mountain larger than the markup at the Valley at which the firms are indifferent in their location. Intuitively, it is not enough for the markups of the Mountain and Valley to be the same. We need to compensate individuals going to the Mountain for their lower probability of survival. Therefore, given p_v, c_v and c_m, there exists a cutoff p_m^* at which $V_t^v = V_t^m = V_{t+1}$. Substituting in Equations 4.1 and 4.2, the transition occurs when

$$\frac{(1-\gamma+\gamma\theta)p_v - c_v}{1-\frac{1}{1+\beta}(1-\gamma+\gamma\theta)} = \frac{(1-\gamma+\gamma(1-\theta))p_m^* - c_m}{1-\frac{1}{1+\beta}(1-\gamma+\gamma(1-\theta))}. \tag{4.4}$$

Given the parameters of our simulation, $p_m^* \approx 1.24$. Then, for any $p_m < p_m^*$ all individual firms locate in the Valley, and for $p_v > p_m^*$ all locate in the Mountain. The intuition of Equation 4.4 is simple; it states that the expected markups adjusted by the survival probabilities need to be equated in the two locations.

In this case, the individual allocation implies multiple equilibria due to the indifference between the two locations. Below that markup the dominant strategy is to locate in the Valley, and above it the optimal decision is to locate in the Mountain. We compare this solution to the one chosen by the multinational in the exact same setting.

The problem of the multinational firm can be written as

$$V(N_t) = \max_{\psi_t} \left\{ \begin{bmatrix} (1-\gamma) & \left((p_v\psi_t + p_m(1-\psi_t))N_t \right. & \left. + \frac{1}{1+\beta} V\left(A \cdot (N_t)^{1-\mu}\right) \right) \\ +\gamma\theta & \left(p_v\psi_t N_t \right. & \left. + \frac{1}{1+\beta} V\left(A \cdot (\psi_t N_t)^{1-\mu}\right) \right) \\ +\gamma(1-\theta) & \left(p_m(1-\psi_t)N_t \right. & \left. + \frac{1}{1+\beta} V\left(A \cdot ((1-\psi_t)N_t)^{1-\mu}\right) \right) \end{bmatrix} - (c_v\psi_t + c_m(1-\psi_t))N_t \right\} \tag{4.5}$$

subject to the same boundary condition we have imposed before:

$$\lim_{N \to 1^-} V(N) = 0.$$

The instantaneous profits are linear in prices and costs, so we have decided to keep costs constant and only change the Mountain's price (p_m) in our simulations. Our first result studies how the multinational's optimal policy $\psi^*(N)$ changes with the Mountain's price p_m.

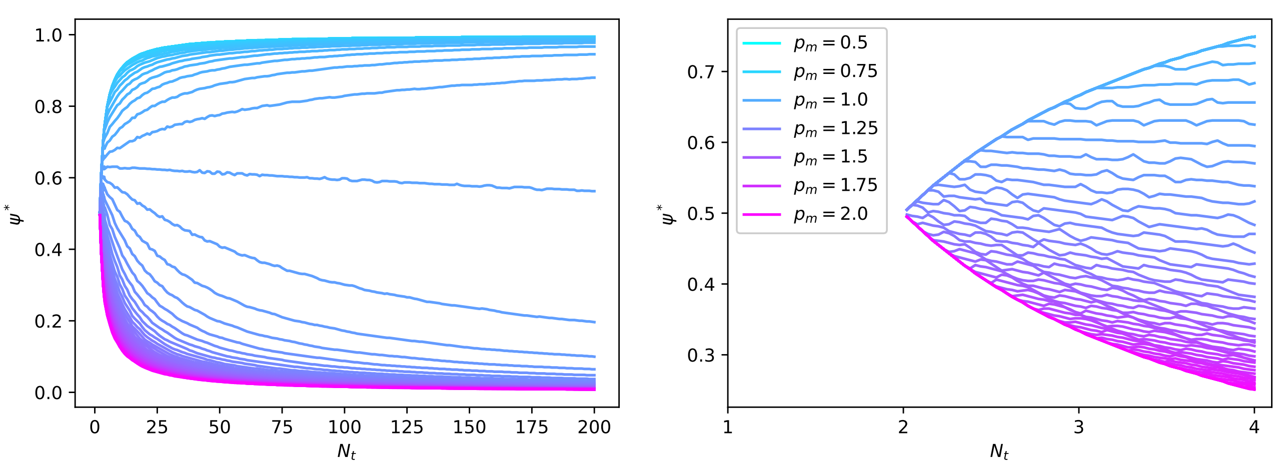

Figure 4.1: Optimal policy as cost of Mountain changes

In Figure 4.1 we present two panels. The left panel shows the optimal allocation for different number of suppliers, and the panel on the right just zooms into the case when there are few suppliers available. Each colored line indicates a different price level, and prices in the Mountain were varied in the range $[0.5, 2.0]$.

As in the baseline model, when there are two firms the multinational allocates one in each location to ensure survival. We decided to start the figure at that point because nothing particularly new occurs for $N_t = 1$. As N_t increases, the multinational allocates firms depending on the prices and the number of available suppliers.

When the prices are low, the optimal allocation is biased toward the Valley, and when the prices are high, the allocation is biased toward the Mountain. Interestingly, there is a price at which the allocation is virtually flat.

Figure 4.2 shows the optimal allocation in the Valley for a given N, but different values of p_m. The plot has been drawn for $N_t = 10$.

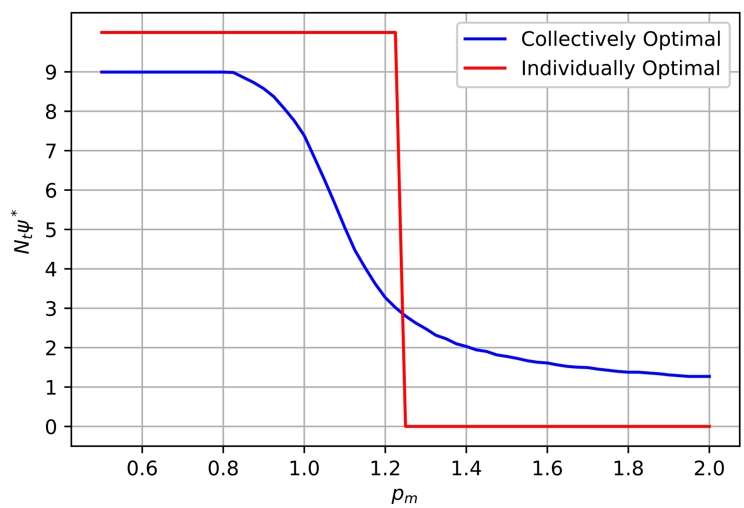

Figure 4.2: Individual vs. Multinational Optimal with Heterogeneous prices.

The multinational transitions smoothly between extreme values of ψ^* as the price of the Mountain varies, but the decentralized equilibrium instantaneously shifts at a critical value of $p^*_m \approx 1.24$. This phase transition from $\psi^* = 0$ to $\psi^* = 1$ occurs when the values of the Mountain and Valley given by Equation 4.1 are equal.

There are two aspects worth highlighting from this simulation. First, the price at which the multinational is indifferent between Mountain and Valley — the point at which it allocates half the firms in the Valley — occurs between 1 and 1.2 (when $N_t\psi = 5$). Notice that the price at this indifferent point is much lower than the price at which the individuals are indifferent. The reason is that the value of continuation is marginally improved when ψ is lowered from 1, and therefore the marginal contribution for the multinational is larger than for the individual firm.

Second, there is a kink at the top left for the multinational. This is the place where the optimal allocation in the Mountain would have implied less than one firm. However, because of the value matching condition, the multinational allocates a maximum of 9 firms to the Valley. A similar kink occurs on the right side of the graph, but the prices required to reach it are large - swamping the details of the graph presented.

4.2 Tax Incentives and Relaxing Price Gouging

As shown in Figure 4.1, there is a price of the Mountain at which the optimal policy implies diversification similar to what is implied by the robust policy in the presence of uncertainty. This implies a subsidy that can be given to the individual firms can achieve the robust policy — or close enough.

This conclusion is important because governments and the private sector might experience risk differently. Typically, governments are the residual claimant in case of natural disasters. Therefore, governments are more likely to prefer a "robust approach" than the private sector would. For instance, if the cost of a natural disaster is very asymmetric, the government is more likely to pay attention to the worst-case than the private sector would. If that is the situation, a government can align private incentives with a small subsidy to the costly location.

Japan is indeed doing so right as a response to Covid-19. In August 2020, Japan set up a fund to compensate firms that diversify out of China — see Bloomberg News (2020). This is not the only case where countries are trying to diversify the supply chain. Moreover, countries are also trying to change other aspects of the global economy. For example, Russia announced a tax break for companies that diversify away from dollar-denominated exporting contracts. In particular, if firms contract in Euros for their exports, they save the domestic sales tax. The need to reduce the dependence on the dollar as the international currency is an aspect that could also be understood through the lens of robustness.

5 Conclusions

Covid-19 was an aggregate shock that highlighted the weaknesses on the supply chain. Many products suffered disruptions: from personal protective equipment, to toilet paper, and beer. It is clear that the supply chains of the world were not prepared for this event. Many are imploring that "future" supply chains become more resilient or robust — but what exactly is a robust supply chain? and how exactly do firms' decisions change when taking that approach?

This paper studies a very stylized model of a supply chain. A company producing a product benefits from many suppliers providing parts, but those firms might not choose the best allocation of resources when aggregate shocks are present. Our model discusses how different arrangements of the supply chain emerge in different settings. In particular, we concentrate on two factors: (i) the internalization of the survival probability — in the spirit of the usual externalities; and (ii) the nature of the shocks when we contrast risk and uncertainty.

We find that a robust supply chain implies concentrating on the worst-case. Robustness yields a strategy that seems to maximize survival probability in the worst case, and therefore our model rationalizes or explains the well known "probability matching" behavior observed in experimental literature.

This solution needs a centralized decision maker who faces uncertainty. Our paper finishes with the discussion of a policy in which the government can create circumstances for which the robust strategy is viable in the decentralized setting. We briefly discuss the recent policies in Japan to subsidize diversification and reduce reliance on China, as an example.

As we highlighted in this paper, robustness is not equivalent to assuming that shocks are larger; it is a different strategy that tries to minimize the losses of the worst case outcome. Increasing the variance of shocks is equivalent to arguing that the supply chain moves from Just-in-Time to Just-in-Case. Robustness, on the other hand, implies that the supply chain moves to Just-in-Worst-Case.

Appendices

A Continuous-Time Model

In the main body of this paper, we study a discrete-time model of supply chains and aggregate shocks. This section formulates a continuous-time model and studies the basic model, the effects of robustness, and finally price differences.

To derive a similar model to that of Section 2, but in continuous time, we need to define the stochastic process N_t, and the value function $v(n)$. We treat shocks to the Mountain or Valley as two independent Poisson counting processes M_t and V_t, with intensity $\gamma\theta$ and $\gamma(1-\theta)$, respectively. Finally, we will assume that firms grow continuously in time, and discontinuities occur only when a shock arrives. Therefore, the number of firms to evolve according to the following.

$$dN_t = b(N_t)dt - \psi_t N_t dV_t - (1-\psi_t)N_t dM_t$$

We can write this stochastic differential equation in a general setting, but for consistency with Section 2, assume that $b(N_t) = AN_t^{1-\mu} - N_t$. Notice that our choice of the drift will yield a mean-reverting process.

The value function depends on a discount rate β and a running reward $r(N_t)$. Similar to the discrete time model, we will let $r(N_t) = pN_t - cN_{t-}$, where N_{t-} is the left-most limit, so that the cost to produce is paid before any production occurs. The value function is then the integral of discounted profit:

$$v(n) = \max_{(\psi_s)_{s\geq 0}} \mathbb{E}\left[\int_0^\infty e^{-\beta t} r(N_t)dt \mid N_0 = n\right], \tag{A.1}$$

where we require that $v(x) = 0,\ \forall x \in [0,1)$ as the boundary condition, as in Section 2.

Due to the fact that v does not depend on time, and the process N_t is Markovian, a dynamic programming argument means that we can re-write the value function as

$$v(x) = \max_{(\psi_s)_{s\geq 0}} \mathbb{E}\left[\int_0^t e^{-\beta s} r(N_s)ds + e^{-\beta t} v(N_t) \mid N_0 = x\right]. \tag{A.2}$$

Next, we work to find the HJB differential equation for the value function. First, we apply Ito's Lemma to $f(t,x) = e^{-\beta t}v(x)$, which gives

$$f(t,N_t) - f(0,N_0) = \int_0^t \frac{df}{dt}(s,N_s)ds + \int_0^t \frac{df}{dx}(s,N_s)dN_s$$
$$+ \sum_{s\le t, dV_s=1} \left[f(s,N_s) - f(s,N_{s-}) + \psi_s N_s \frac{df}{dx}(s,N_s) \right]$$
$$+ \sum_{s\le t, dM_s=1} \left[f(s,N_s) - f(s,N_{s-}) + (1-\psi_s) N_s \frac{df}{dx}(s,N_s) \right]$$

After some simplification, using the definition of $f(t,x)$ and N_t, we arrive at the next step in our derivation:

$$e^{-\beta t} v(N_t) - v(N_0) = \int_0^t e^{-\beta s} \left[b(N_s) v'(N_s) - \beta v(N_s) \right] ds$$
$$+ \sum_{s\le t, dV_s=1} e^{-\beta s} (v(N_s) - v(N_{s-}))$$
$$+ \sum_{s\le t, dM_s=1} e^{-\beta s} (v(N_s) - v(N_{s-}))$$

The left-hand side of this expression is present in Equation A.2, which when plugged in yields

$$0 = \max_{(\psi_s)_{s\ge 0}} \mathbb{E} \left[\int_0^t e^{-\beta s} \left(r(N_s) + b(N_s) v'(N_s) - \beta v(N_s) \right) ds \right.$$
$$\left. + \int_0^t e^{-\beta s} \left(1_{dV_s=1} \left(v((1-\psi_s)N_s) - v(N_s) \right) + 1_{dM_s=1} (v(\psi_s N_s) - v(N_s)) \right) ds \right| N_0 = n \Bigg].$$

Now we use the definitions of $r(N_s)$ and $b(N_s)$:

$$0 = \max_{(\psi_s)_{s\ge 0}} \mathbb{E} \left[\int_0^t e^{-\beta s} \left((p-c)N_s + \left(A N_s^{1-\mu} - N_s \right) v'(N_s) - \beta v(N_s) \right) ds \right.$$
$$\left. + \int_0^t e^{-\beta s} \left(1_{dV_s=1} \left(v((1-\psi_s)N_s) - v(N_s) - p\psi_s N_s \right) + 1_{dM_s=1} \left(v(\psi_s N_s) - v(N_s) - p(1-\psi_s)N_s \right) \right) ds \right| N_0 = n$$

Finally, we divide both sides by t and take the limit as $t \to 0$. Making use of the intensities of the counting processes M_t and V_t, we find the HJB equation for the value function:

$$
\begin{aligned}
0 &= (p-c)x + \left(Ax^{1-\mu} - x\right) v'(x) - (\gamma+\beta)v(x) \\
&\quad + \max_{\psi} \gamma \left[(1-\theta)\left(v((1-\psi)x) - p\psi x\right) + \theta\left(v(\psi x) - p(1-\psi)x\right)\right] \\
v(x) &= 0, \ \forall x \in [0,1)
\end{aligned}
\tag{A.3}
$$

This expression is a nonlinear first-order ordinary differential equation. If the maximization problem is concave, we can solve for the first-order condition to find the optimal ψ^*, which will satisfy:

$$
\frac{1-\theta}{\theta} = \frac{p + v'(\psi^* x)}{p + v'((1-\psi^*)x)}
\tag{A.4}
$$

as long as $\psi^* x \geq 1$ and $(1-\psi^*)x \geq 1$.

The optimal allocation satisfies a 'probability matching' of sorts: the ratio of short-term profit (p) plus long-term marginal benefit (v') must equal the ratio of aggregate shock probabilities. In particular, if $v'(1) = \infty$, then the multinational will choose an interior solution.

A.1 Robustness

We can also study the robust minimax problem of Section 3. Again, we assume that $\theta \in [\bar{\theta} - \Delta, \bar{\theta} + \Delta]$. Now, the multinational's value function is given by

$$
v(n) = \max_{(\psi_s)_{s\geq 0}} \min_{\theta \in [\bar{\theta}-\Delta, \bar{\theta}+\Delta]} \mathbb{E}\left[\int_0^\infty e^{-\beta t} r(N_t) dt \mid N_0 = n\right],
\tag{A.5}
$$

and the corresponding HJB is

$$
\begin{aligned}
0 &= (p-c)x + \left(\frac{1}{1-\mu} x^{1-\mu} - x\right) v'(x) - (\gamma+\beta)v(x) \\
&\quad + \max_{\psi} \min_{\theta \in [\bar{\theta}-\Delta, \bar{\theta}+\Delta]} \gamma \left[(1-\theta)\left(v((1-\psi)x) - p\psi x\right) + \theta\left(v(\psi x) - p(1-\psi)x\right)\right] \\
v(x) &= 0, \ \forall x \in [0,1)
\end{aligned}
\tag{A.6}
$$

First, we study the first-order condition for the minimization problem. For a fixed ψ, the function to minimize is linear in θ:

$$
\left(v((1-\psi)x) - p\psi x\right) + \theta\left[v(\psi x) - p(1-\psi)x - v((1-\psi)x) + p\psi x\right].
$$

Nature's optimal choice of θ will be an extreme value. If the quantity in the brackets is negative (or positive), the optimal θ^* is $\bar{\theta} + \Delta$ (or $\bar{\theta} - \Delta$, respectively). Formally,

$$\theta^* = \begin{cases} \bar{\theta} + \Delta & \text{if } v(\psi N_t) - v((1-\psi)N_t) < p(1-2\psi)N_t \\ \bar{\theta} - \Delta & \text{if } v(\psi N_t) - v((1-\psi)N_t) > p(1-2\psi)N_t \end{cases} \tag{A.7}$$

Notice that the choice $\psi = 0.5$ makes nature indifferent among all θ. Here lie the effects of robustness: if Δ is large, nature can significantly punish the multinational for choosing $\psi \neq 0.5$, but if Δ is small, the robust optimum may choose a larger ψ and lose fewer firms. The message is the same — uncertainty rewards optimizing the worst-case outcome.

Finally, the first-order conditions for ψ will give the same expression as in the previous section, but with $\theta = \theta^*$.

A.2 Heterogeneous Prices

We now study the analogue of Section 4 by allowing prices to be different. Due to linearity of our system, heterogeneous prices will give the same qualitative results as heterogeneous costs.

Let us consider how the optimal solution changes if $p_m > p_v$. Now, the only difference is that

$$r(N_t) = N_t \left(p_v \psi (1 - dV_t) + p_m (1-\psi)(1 - dM_t) - c \right).$$

Following the derivation from before, we find that the HJB for the value function is

$$\begin{aligned} 0 &= -cx + \left(Ax^{1-\mu} - x\right) v'(x) - (\gamma + \beta) v(x) \\ &\quad + \max_{\psi} \left(p_v \psi + p_m(1-\psi)\right) x + \gamma \left[(1-\theta)\left(v((1-\psi)x) - p_v \psi x\right) + \theta \left(v(\psi x) - p_m(1-\psi)x\right) \right] \\ v(x) &= 0, \ \forall x \in [0,1) \end{aligned} \tag{A.8}$$

Under similar concavity conditions, the optimal ψ^* will satisfy:

$$\frac{p_m - p_v}{\gamma} = \theta \left[v'(\psi^* x) + p_m \right] - (1-\theta) \left[v'((1-\psi^*)x) + p_v \right].$$

Rearranging, the first-order condition becomes:

$$\frac{1-\theta}{\theta} = \frac{\left(1-\frac{1}{\gamma\theta}\right)p_m + v'(\psi^* x)}{\left(1-\frac{1}{\gamma(1-\theta)}\right)p_v + v'((1-\psi^*)x)} \tag{A.9}$$

This is again a probability matching solution. Note that both coefficients of p_m and p_v are negative, and therefore if p_m is large, then $v'(\psi^* x)$ must also be large, which we expect implies a small ψ. This indicates that we may be able to re-create any allocation through the choice of a suitable p_m.

Additionally, the arrival probability of the aggregate shock plays a role in this expression. If γ is close to one, then the multinational's decision is heavily dictated by the ratio of marginal values of firms, and the price differences are ignored. However, if γ is small, then the ratio of prices is more important.

References

Alfaro, L., D. Chor, P. Antras, and P. Conconi (2019). Internalizing global value chains: A firm-level analysis. *Journal of Political Economy 127*(2), 508–559.

Ambulkar, S., J. Blackhurst, and S. Grawe (2015). Firm's resilience to supply chain disruptions: Scale development and empirical examination. *Journal of operations management 33*, 111–122.

Antras, P. (2003). Firms, contracts, and trade structure. *The Quarterly Journal of Economics 118*(4), 1375–1418.

Antras, P. (2015). *Global production: Firms, contracts, and trade structure*. Princeton University Press.

Antràs, P. (2020). *Conceptual aspects of global value chains*. The World Bank.

Antràs, P. and D. Chor (2013). Organizing the global value chain. *Econometrica 81*(6), 2127–2204.

Antràs, P. and D. Chor (2018). On the measurement of upstreamness and downstreamness in global value chains. *World Trade Evolution: Growth, Productivity and Employment*, 126.

Antràs, P. and A. De Gortari (2020). On the geography of global value chains. *Econometrica 88*(4), 1553–1598.

Bloomberg News (2020, July). Japan starts paying firms to cut reliance on chinese factories. *www.bloomberg.com*.

Brennan, T. J. and A. W. Lo (2011). The origin of behavior. *The Quarterly Journal of Finance 1*(01), 55–108.

Cary, G. S., M. J. F. M. Dolmans, B. Hoffman, T. Graf, L. Brannon, R. Pepper, H. Mostyn, A. R. B. Lazda, S. Haynes, K. Georgieva, and J. Przerwa (2020, April). Exploitative abuses, price gouging & covid-19: The cases pursued by eu and national competition authorities. *www.concurrences.com*.

Durach, C. F. and J. A. Machuca (2018). A matter of perspective–the role of interpersonal relationships in supply chain risk management. *International Journal of Operations & Production Management*.

Evans, D. (2020, Feb). Coronavirus shows that supply chains are outdated and unfit for modern manufacturing. *www.forbes.com*.

Fiorina, M. P. (1971). A note on probability matching and rational choice. *Behavioral Science 16*(2), 158–166.

Gaissmaier, W. and L. J. Schooler (2008). The smart potential behind probability matching. *Cognition 109*(3), 416–422.

Helpman, E., M. J. Melitz, and S. R. Yeaple (2004). Export versus fdi with heterogeneous firms. *American economic review 94*(1), 300–316.

King and Spalding (2020). Covid-19 survey of federal and state price gouging laws. *www.kslaw.com*.

Kremer, M. (1993). The o-ring theory of economic development. *The Quarterly Journal of Economics 108*(3), 551–575.

Lo, A. W. (2017). *Adaptive markets*. Princeton University Press.

Long, S. (2020, May). Coronavirus pandemic exposes fatal flaws of the 'just-in-time' economy. *www.abc.net.au*.

Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica 71*(6), 1695–1725.

Morse, E. B. and W. N. Runquist (1960). Probability-matching with an unscheduled random sequence. *The American journal of psychology 73*(4), 603–607.

Plenert, G. (2007). *Reinventing lean: introducing lean management into the supply chain*. Oxford, UK: Butterworth-Heinemann.

Rotemberg, J. J. (2002). Customer anger at price increases, time variation in the frequency of price changes and monetary policy. Technical report, National Bureau of Economic Research.

Rotemberg, J. J. (2011). Fair pricing. *Journal of the European Economic Association 9*(5), 952–981.

Saenz, M. J., X. Koufteros, C. F. Durach, A. Wieland, and J. A. Machuca (2015). Antecedents and dimensions of supply chain robustness: a systematic literature review. *International Journal of Physical Distribution & Logistics Management*.

Sion, M. et al. (1958). On general minimax theorems. *Pacific Journal of mathematics 8*(1), 171–176.

The FT Editorial Board (2020, April). Companies should shift from just in time to just in case. *Financial Times*.

Töyli, J., H. Lorentz, L. Ojala, A. Wieland, and C. M. Wallenburg (2013). The influence of relational competencies on supply chain resilience: a relational view. *International Journal of Physical Distribution & Logistics Management*.

UK Competition Authority (2020, April). The uk competition authority publishes update on its covid-19 taskforce, including a number of excessive pricing complaints and investigations. *www.concurrences.com*.

Vickery, S. n., R. Calantone, and C. Dröge (1999). Supply chain flexibility: an empirical study. *Journal of supply chain management 35*(2), 16–24.

Vulkan, N. (2000). An economists perspective on probability matching. *Journal of economic surveys 14*(1), 101–118.

Wagner, S. M. and C. Bode (2006). An empirical investigation into supply chain vulnerability. *Journal of purchasing and supply management 12*(6), 301–312.

Yeaple, S. R. (2003). The complex integration strategies of multinationals and cross country dependencies in the structure of foreign direct investment. *Journal of international economics 60*(2), 293–314.

Zhao, S. and F. You (2019). Resilient supply chain design and operations with decision-dependent uncertainty using a data-driven robust optimization approach. *AIChE Journal 65*(3), 1006–1021.