

Essays on Trade Costs, Supply Chain Uncertainty and CEO Compensation

A Dissertation

Presented for the

Doctor of Philosophy

Degree

The University of Tennessee, Knoxville

Valentina Kozlova

August 2013

© by Valentina Kozlova, 2013
All Rights Reserved.

DEDICATION

To my mother, Anna Kozlova,
for always believing in me.

Acknowledgements

First of all, I would like to thank my advisor Georg Schaur for all the effort he put in to guide me through the difficult process of writing a dissertation. It is thanks to his support, patience, hard work and enduring willingness to help that I was able to finish my graduate studies. Georg was always there to give feedback and to work on problems, and for that I am sincerely grateful. I have benefited greatly from the advice of my remaining committee members: Don Clark, Mohammed Mohsin and Matt Myers. I would also like to thank my coauthor and mentor William Neilson. He has always provided feedback that was critical but incredibly encouraging at the same time, an impossible task to do in my opinion. I am sincerely grateful to Dr. Neilson for everything that he has done for me over the years, and special thanks for the hooding ceremony that he agreed to participate in and made unforgettable. I am also indebted to my boss and friend Edward Taylor, as well as Dr. John Moore, for the support and warmth that they gave me, especially in the last semester of my program.

On a personal note, I would also like to acknowledge the support that came from my family and friends. In particular, my friend Kirill Kacharov should receive special recognition. He was able to help even though he always lived thousands of miles away. Kirill was there every time I needed to call and complain about difficulties, and I am immensely grateful for that. Natalia Gritsko is my friend of 10 years and a colleague, a very powerful combination. I am grateful for Natalia's presence in my

life, and especially for her sense of humor and all of our inside jokes that were the best stress reliever.

During grad school years I met a lot of special people. I am grateful to Cristina Reiser, Steve Ogden, Kelly Padden, Kate Harper, Kara Smith Mitchell, Caleb Siladke and many others for their friendship and invaluable advice, as well as for the words of support that always made a difference. I would also like to thank my close friend Bruno Wichmann for giving me the encouragement when it was most needed. He was close to me in the hardest moments of the process and every time he was able to lift my spirits up and brighten my day. His talent, intelligence and dedication were always inspiring to me, and his faith in me was life-changing.

I would not have been able to complete this work without my most important supporters, my family. My grandparents Galina and Vasiliy Tantsura are my biggest cheerleaders. I appreciate my two aunts, Nina and Natasha Kozlova, for their care and encouragement they gave me every time. I would like to mention my coolest cousins, Masha and Ruslan. I am grateful to them for loving me from such distance, I always felt it. Finally, I can't even begin to say how much I appreciate the most special woman in my life, my mom. She is the best person in the world, without her none of this would be possible. Everything I achieve in life is because of my mom, her care, her hard work, her support and her unconditional love.

Abstract

This dissertation consists of two chapters that examine high managerial pay and supply chain uncertainty.

Chapter 1 constructs a game-theoretic model in which high CEO pay emerges as the outcome of an arms race, with each firm paying its CEO highly to protect its competitive position against rivals who also pay highly. For an arms race to emerge, highly-paid CEOs must generate idiosyncratic, privately-known internal effects on profit, and CEO pay disparities must also generate asymmetric profit differences from external effects beyond the simple differences in pay. If the distribution of internal effects satisfies a key uniformity condition, an arms race emerges as the only equilibrium of the game.

Chapter 2 examines the impact of supply chain uncertainty and ordering costs on trade. Importers hold safety stock to hedge against delays in delivery. An increase in supply chain uncertainty raises safety stocks, increases inventory costs, and reduces imports from locations with high delivery time uncertainty. An increase in order costs reduces a firm's shipping frequency and increases average inventory holding cost for the firm's base inventory stock. As a result, firms import less from locations with high ordering costs to reduce average inventory holding costs. Detailed data on actual and expected arrival times of vessels at U.S. ports serve to measure supply chain uncertainty consistent with the theory. Combined with detailed data on U.S. imports, freight charges and unit values, a 10 percent increase in supply-chain uncertainty lowers imports by as much as 3.7 percent. This is evidence that delivery uncertainty

imposes a cost on imports according to the management of safety stocks. A one percent increase in ordering costs lowers imports by as much as 1.2 percent. Ordering costs impact the intensive margin of trade due to the management of base inventory stocks.

Table of Contents

1	CEO Arms Races	1
1.1	Introduction	1
1.2	Conceptual Framework	4
1.3	The CEO Hiring Game	6
1.4	The CEO Arms Race	10
1.5	Concluding Remarks	19
2	Supply Chain Uncertainty as a Trade Barrier	23
2.1	Introduction	23
2.2	Theory	27
2.2.1	Import Demand	27
2.2.2	Importer's Costs	28
2.3	Data, Estimation and Results	34
2.3.1	Specification	34
2.3.2	Data	37
2.3.3	Results	40
2.3.4	Robustness Checks	43
2.3.5	Structural Estimation	47
2.4	Conclusions	48
	Bibliography	60

Appendix	67
A Jensen's inequality and relative demands	68
B Log-linearization	69
C Units vs. kilograms	71
Vita	73

List of Tables

2.1	Table of summary statistics of uncertainty	53
2.2	Descriptive regressions	54
2.3	Summary statistics	55
2.4	The effect of uncertainty on log of ocean-shipped weight	56
2.5	Alternative measures of uncertainty	57
2.6	The impact of air shipment, import from Mexico and Canada and source country's income on the uncertainty elasticity of vessel imports	58
2.7	The effect of uncertainty on log of ocean-shipped weight by category	59

List of Figures

1.1	The Theorem Illustrated	22
2.1	Inventory Process Illustrated	50
2.2	Difference between the actual date of arrival and the estimated date of arrival	51
2.3	Uncertainty for district d and country j	52

Chapter 1

CEO Arms Races

1.1 Introduction

Over the past decade both journalists and academics have written about the level of chief executive officer (CEO) pay both in the United States and in Europe. In 1989, the average CEO total compensation in the United States was \$2.3 million, increased to \$8.4 million in 1999, and reached \$11.4 million in 2009.¹ At the same time, average annual worker income in the United States rose from only \$32,438 to \$38,376.² Recently, this type of earnings discrepancy has attracted the interest of politicians, the media, and the general public, and is the driving force behind movements such as *Occupy Wall Street*.³ Because of the large increase in CEO pay, and because it is so high relative to the pay of ordinary workers, much of the recent academic literature on executive compensation has sought to rationalize these trends.

The conventional approach to this issue entails investigating whether increased CEO pay is matched by increased value to the firm. To this end, Finkelstein and Boyd (1998) find that pay is positively correlated with managerial discretion, and Murphy

¹Source: Forbes Annual Executive Compensation Reports. Data are in 2008 dollars.

²Source: United States Census Bureau. Data are in 2008 dollars.

³Occupy Wall Street is a movement that began on September 17, 2011 in Liberty Square in Manhattan's Financial District, and has spread to over 100 cities in the United States. Their goal is "to expose how the richest 1% of people are writing the rules of an unfair global economy". Source: <http://occupywallst.org/about/> (Accessed on November 07, 2011).

and Zábojník (2004, 2007) argue that increased CEO pay follows from increases in the general (as opposed to specific) human capital required for the job. Rajgopal et al. (2006) find that increased demand for talented CEOs, caused by market-wide shocks, forces firms to set CEO pay high in order to retain their more talented CEOs. Geletkanycz et al. (2001) go on to demonstrate that CEOs receive premiums for strategically valuable external networks, and Westphal et al. (2006) document that these networks include informal (friendship) networks with top executives of both suppliers and financial institutions. Kroll et al. (1997); Schaefer (1998); and Gabaix and Landier (2008) show that CEO compensation increases with the size of the average firm in the economy. Most recently, Chang et al. (2010) test whether firm performance and high pay reflect CEO ability or something else like other firm assets, either physical or human, or simply luck. They reject their other explanations in favor of high CEO pay reflecting high CEO ability.

An alternative approach looks for reasons why CEO pay might be higher than would be consistent with the usual labor-market approach. Bebchuk et al. (2002) consider the fact that CEOs have some ability to set their own pay, allowing them to extract rents. Dow and Raposo (2005) explore the implications of CEO contracts being different from typical principal-agent contracts because the former change over time in accordance with firm’s performance and strategic direction. Aggarwal and Samwick (1999) argue that CEO compensation is positively sensitive to rival firms’ performance. More recently, Malmendier and Tate (2009) find that firms pay award-winning CEOs more, even though they do not perform better than their non-award-winning matches, suggesting alternative benefits of having a highly paid “superstar” CEO. Hayes and Schaefer (2009) find a “Lake Wobegon effect”, in which firms pay their CEOs highly to signal that their chief executives are above average.⁴

We add to this second stream of literature by providing an additional reason for high CEO pay - an arms race. Firms desire highly-paid CEOs because their rivals

⁴The Lake Wobegon effect is derived from Garrison Keillor’s radio program, *A Prairie Home Companion*. The town of Lake Wobegon is described as a place where “all the women are strong, all the men are good looking, and all the children are above average.”

have highly-paid CEOs. In fact, in our model the *only* reason most firms hire highly-paid CEOs is because they believe their rivals will, and only a small fraction of firms would choose to hire highly-paid CEOs if they thought their rivals would hire low-paid CEOs. Even so, this small fraction willing to hire highly-paid CEOs causes a ripple effect that courses through the entire set of firms, with all firms choosing to hire highly-paid CEOs in equilibrium.⁵

The arms race literature began during the 1960s motivated by the Cold War. The first mathematical model of an arms race was formulated by Richardson (1960) using a system of linear differential equations.⁶ More recent approaches use game theory, with models of arms races between nations based on Cournot competition (Dumas, 1979; Okuguchi, 1981) and on prisoner’s dilemma games (Brams, 1985; Brams and Kilgour, 1988). These arms race games have been applied to settings outside of actual weapons build-up, including hospital technology investments (Foote, 1992; James, 2002), and household vehicle size choices (White, 2004).

The most recent game-theoretic model of an arms race was developed by Baliga and Sjöström (2004), and we rely heavily upon their framework to construct a model of a CEO arms race. In their model each of two countries can choose to build new weapons or not, and countries benefit from having more weapons than their rivals. Nations differ in their costs of arming, with higher-cost nations preferring not to arm when they believe that their rivals will not build new weapons, but lower-cost nations having a dominant strategy of always building new weapons. No matter what the cost, though, nations prefer to build more weapons when their rivals do. Baliga and Sjöström establish that if nations believe that there is even a small probability of facing a dominant strategy rival, a multiplier effect occurs and the unique Bayes-Nash equilibrium involves an arms race with probability one. We build on Baliga and

⁵This result is in line with an important insight from Carlsson and Van Damme (1993). They demonstrate that, under asymmetric information about the opponent’s preferences, a non-zero probability that opponents will choose a particular action a^* may cause every player to follow suit.

⁶Refer to Anderton (1989) for a comprehensive discussion on Richardson-type arms race models and applications.

Sjöström’s model by making firms heterogeneous in the benefit, not cost, from arming. We also generalize the model to the n firms case and demonstrate the conditions under which an arms race occurs in this general scenario.

The paper is organized as follows. Section 2 describes the conceptual framework for understanding the CEO pay game. Section 3 utilizes a two-firm setting to introduce the game firms play and identify the internal and external benefits of hiring a highly-paid CEO. Section 4 generalizes the game to the n firms case and also contains the main result, i.e. the conditions under which both the dominant-strategy firms and the non-dominant-strategy firms hire highly-paid CEOs. Section 5 offers some conclusions.

1.2 Conceptual Framework

We construct a model in which each firm has rivals, and we abstract away from whether this rivalry occurs in product, input, employment, or capital markets. Firms choose what type of CEO to hire. They can hire a highly-paid (high-type) CEO or a low-paid (low-type) CEO. A highly-paid CEO can benefit the firm in two ways, one internal and one external.

The internal benefit is separate from the rivalry, and arises as the higher-paid CEO runs the firm in a more efficient, more profitable way. This may occur, for example, through improved positioning in capital markets, improved bargaining power with suppliers or customers, increased social networks, or other factors that make the firm more profitable. The internal benefit is therefore anything that the conventional CEO pay literature describes as a reason for high CEO pay.⁷

The external benefit arises from a higher-paid CEO providing a competitive advantage over rivals with a lower-paid CEO. This may occur, for example, because the higher-paid CEO provides more credibility in capital markets, or signals higher quality to customers. This benefit could arise, for example, from the superstar effect

⁷See the references cited in the second paragraph of the introduction.

of Malmendier and Tate (2009) or the Lake Wobegon effect of Hayes and Schaefer (2009). Firms with a lower-paid CEO facing rivals with a higher-paid CEO observe negative external effect. This can result from low pay being a signal of low firm value or low managerial quality (Hayes and Schaefer, 2009) which makes investors prefer high-paying firms over their low-paying rivals.

We assume that the external effect is identical and commonly known, but that the internal effect is private information. The nature of the game has every firm preferring a highly-paid CEO when their rivals have highly-paid CEOs, but this may or may not be a dominant strategy. When the internal effect is large enough that the benefit the highly-paid CEO provides exceeds the wage difference, having a highly-paid CEO dominates. On the other hand, when the internal effect is smaller than the wage difference, firms would prefer a low-paid CEO when their rivals also have low-paid CEOs.

We show that if there are some firms with a large enough internal effect to make hiring a highly-paid CEO a dominant strategy, and a sufficiently uniform distribution of internal effects across firms, an arms race ensues and every firm chooses a highly-paid CEO. The reasoning is as follows, and matches the contagion story of Baliga and Sjöström (2004). When the internal effect is sufficiently large to make hiring a highly-paid CEO a dominant strategy, the firm chooses a highly-paid CEO. Now consider a firm whose internal effect is just slightly lower than this. The firm will prefer to hire a low-paid CEO in response to its rivals hiring a low-paid CEO, while hiring a highly-paid CEO is still a best response to other firms' high hiring because of the external rivalry effect. This firm does not possess a dominant strategy, but we can think of such a firm as an "almost-dominant strategy" firm because its internal effect is almost as large as that of a dominant strategy firm.

For these "almost-dominant strategy" firms, hiring a highly-paid CEO is still a best response to hiring a highly-paid CEO, but hiring a low-paid CEO is a best response to hiring a low-paid CEO. Choosing to hire a highly-paid CEO, though, entails only a small loss if the rival hires a low-paid CEO, whereas choosing to hire a

low-paid CEO results in large losses when the rival hires a highly-paid CEO. When there is a sufficiently positive (but still small) probability that other firms have highly-paid CEOs, the “almost dominant” firm will choose to hire a highly-paid CEO to avoid the large losses. Thus, the firm becomes infected and mimics the dominant strategy firms. The same reasoning then applies to firms with the next lowest internal benefits, and such contagion could eventually affect all firms. In this case, an arms race occurs.

Based on these arguments, the arms-race approach to CEO pay leads to a new understanding of why CEO pay is so high. It only takes a few firms hiring highly-paid CEOs to make everyone hire highly-paid CEOs, because these few firms are enough to start an “epidemic” of high hiring. Importantly, though, this prevalence of highly-paid CEO arises from equilibrium behavior, and no firm can gain from unilaterally lowering pay. Consequently, a typical firm’s correct rationale for hiring a highly-paid CEO amounts to “because everyone else does.”

1.3 The CEO Hiring Game

Two firms compete in some relevant market, either a product market, capital market, or labor market. Both firms maximize profit, and to do so choose what type of CEO to hire. Firms can draw CEOs from a market of high-type CEOs or from a market of low-type CEOs. We assume that firms pay competitive wages in both markets, however, high-type CEOs demand a wage premium (i.e. high market clearing wage) while low-type CEOs are paid standard wages (i.e. low market clearing wage). The premium (or wage differential) is fixed at $w > 0$. For the baseline case in which both firms choose to hire low-paid CEOs, we normalize profit to be 0.

Hiring a highly-paid CEO does two things. First, it adds a firm-specific net increment $x_i \in [\underline{x}, \bar{x}]$ to the firm’s profit.⁸ This increment is private information to the firm and drawn independently from the distribution $F(x)$. This is the *internal effect*. We require $\bar{x} > 0$, but $\underline{x}$ may be positive or negative. Second, the highly-paid

⁸One can interpret x_i as firm i ’s net benefit above a firm-specific hiring cost.

CEO provides the firm with a competitive advantage if its rival chooses a low-paid CEO, with the highly-paid CEO firm earning an extra $B > 0$ in gross profit and the low-paid CEO firm losing $N > 0$. This is the *external effect*. It is assumed to be identical across firms, which is consistent with the notion that the impact of having a higher-paid CEO than one's rival is determined by the nature of the market pairing rather than the characteristics of the individual firm. The resulting payoffs generate the following normal form game.

		Firm j	
		Highly-Paid	Low-Paid
		(H)	(L)
Firm i	Highly-Paid (H)	$x_i - w, x_j - w$	$x_i + B - w, -N$
	Low-Paid (L)	$-N, x_j + B - w$	$0, 0$

To understand the game, consider first a situation in which both x_i and x_j are zero. If firm j hires a highly-paid CEO, firm i 's choice is between hiring a highly-paid CEO and earnings of $-w$, or a low-paid CEO and earnings of $-N$. If $w < N$ hiring a highly-paid CEO is a best response to hiring a highly-paid CEO. On the other hand, if firm j hires a low-paid CEO, firm i faces a choice between a highly-paid CEO and earnings of $B - w$, or a low-paid CEO and earnings of 0. If $B < w$ hiring a low-paid CEO is a best response to hiring a low-paid CEO. Consequently, when $B < w < N$ there are two pure strategy equilibria in the special case where it is known that $x_i = x_j = 0$, and in these two pure strategy equilibria both firms hire the same type of CEO.

Now suppose instead that firm i 's internal profitability increment x_i satisfies $x_i + B > w$. Hiring a highly-paid CEO remains a best response to hiring a highly-paid CEO, but firm i 's best response to firm j hiring a low-paid CEO is for firm i to hire a highly-paid CEO, because the combination of the productivity increment and the competitive advantage more than offset the extra CEO pay. Consequently, hiring

a highly-paid CEO is a dominant strategy when the firm's private increment x_i is sufficiently high.

We restrict the components of the game so that:

$$w < N + \underline{x} \tag{1.1}$$

$$\underline{x} + B < w < \bar{x} + B \tag{1.2}$$

Condition (1.1) states that for firms with the smallest internal effect, hiring a highly-paid CEO is a best response to the rival hiring a highly-paid CEO. Higher values of the internal effect x_i only enhance this incentive, so highly-paid being a best response to highly-paid holds for all values of $x_i \in [\underline{x}, \bar{x}]$. The first inequality in condition (1.2) states that for the firm with the lowest internal effect hiring a low-paid CEO is a best response to hiring a low-paid CEO, while the second inequality, coupled with condition (1.1), says that for the firm with the highest internal effect hiring a highly-paid CEO is a dominant strategy.

These two inequalities imply that $B < N$, that is, the benefit to the winner of the CEO-hiring competition is smaller than the loss to the loser. Such a condition is consistent with findings in the literature. According to Hayes and Schaefer (2009), CEO pay serves as a signal of managerial quality and firm value to stock market participants. The different magnitudes might reflect that positive signals attract less capital than negative signals deter. It might also reflect the countervailing effect caused by institutional investors moving away from firms with higher-paid CEOs, as documented by David et al. (1998) and Hartzell and Starks (2003).

Conditions (1.1) and (1.2) imply that for low values of x both types of hiring are rationalizable but for high values of x hiring a highly-paid CEO becomes a dominant strategy. If there is a positive probability of some players being dominant strategy firms, a contagion effect comes into play, and can lead to a CEO arms race, with all

firms hiring highly-paid CEOs even though for most of them it is not a dominant strategy.

To see how this contagion comes about, consider firm i 's decision of whether to hire a highly-paid CEO or a low-paid CEO. Firm i knows that if i finds it beneficial to hire a highly-paid CEO, any firm with a larger internal effect than x_i will also find it beneficial. Consequently, firm i realizes that if it chooses a highly-paid CEO, the probability that it faces a rival with a highly-paid CEO is at least $1 - F(x_i)$. If the firm faces a high-hiring rival, hiring high itself generates a gain of $x_i - w + N$ compared to hiring low. If, instead, the firm faces a low-hiring rival, hiring high generates a gain of $x_i - w + B$. Because the probability of facing a high-hiring rival, conditional on i choosing to hire a highly-paid CEO itself, is at least $1 - F(x_i)$, the expected gain from hiring a highly-paid CEO satisfies

$$E(H) - E(L) \geq [1 - F(x_i)](x_i - w + N) + F(x_i)(x_i - w + B),$$

where $E(H)$ is the expected payoff from hiring a highly-paid CEO (i.e. the expected payoff from action H) and $E(L)$ is the expected payoff from hiring a low-paid CEO (i.e. the expected payoff from action L). The inequality holds because $N > B$ by conditions (1) and (2). Firm i chooses to hire the highly-paid CEO if $E(H) - E(L) \geq 0$, which can be rearranged to

$$1 - F(x_i) \geq \frac{w - x_i - B}{N - B}. \quad (1.3)$$

According to this expression, firm i chooses to hire a highly-paid CEO when there are “enough” firms with higher internal effects than it has. Essentially, firm i hires the high-type CEO to defend itself against all the other firms hiring high-types. If

expression (1.3) holds for every value of x in $[\underline{x}, \bar{x}]$ then every firm hires highly-paid CEOs, and the arms race ensues.

Note that $N > B$ by conditions (1) and (2), and that $w - x < N$ for all $x \in [\underline{x}, \bar{x}]$ by condition (1), so that the right-hand side of (3) lies between 0 and 1. Also, condition (2) implies the existence of a value $\hat{x} \in [\underline{x}, \bar{x}]$ such that $w - \hat{x} = B$. Because the right-hand side of (3) is decreasing in x and equal to 0 when $x = \hat{x}$, condition (3) is automatically satisfied for firm with internal values in $[\hat{x}, \bar{x}]$ but not for those in $[\underline{x}, \hat{x})$.

1.4 The CEO Arms Race

We begin this section by extending the above analysis to the case of more than two firms. To that end, suppose that $n \geq 2$ profit-maximizing firms compete in the relevant market. Since it is a static game, the number of firms is fixed. Each firm chooses between hiring either a high-type (highly-paid) CEO or a low-type (low-paid) CEO, with the wage premium $w > 0$ fixed across firms. If every firm chooses to hire the low-type CEO, profits are normalized to be 0.

Consider firm i and suppose that h_i of the other $n - 1$ firms hire high-types and $n - h_i - 1$ hire low-types. When firm i hires a low-type, it will lose the CEO-hiring competition and have a payoff of $-r(h_i)N$, where $r(h_i)$ is a function that determines firm i 's share of the total loss N . When firm i hires a high-type, its payoff is $x_i + s(n - h_i - 1)B - w$, where $s(n - h_i - 1)$ determines i 's share of total benefit.

Formally, the function $r : [0, n - 1] \rightarrow [0, 1]$ is a *loss-sharing rule* if it is nondecreasing and onto. In particular, this requires $r(0) = 0$ so that a firm does not experience any external loss if it plays L when all other firms do, and $r(n - 1) = 1$ so that a firm suffers the entire external loss if it is the only firm to play L when everyone

else plays H . Similarly, the function $s : [0, n - 1] \rightarrow [0, 1]$ is a *benefit-sharing rule* if it is nonincreasing and onto. This requires that $s(n - 1) = 0$ so that a firm gains no external benefit when it plays H and so do all other firms, and it requires that $s(0) = 1$ so that a firm gains the entire external benefit by being the only firm to hire the highly-paid CEO. Requiring that r be a loss-sharing rule and s a benefit-sharing rule allows for unique identification of the external loss and benefit parameters N and B from the cases where one firm acts differently from all of the others. Furthermore, other than the requirement that the functions be strictly monotone, the rules say nothing about whether total external losses and benefits increase or decrease as the number of firms hiring high-types changes.

One result below relies on *proportional* scaling functions:

$$r(h) = \frac{h}{n - 1} \tag{1.4}$$

$$s(h) = \frac{n - h - 1}{n - 1} \tag{1.5}$$

These functions have intuitive appeal. Firm i 's payoff would be $-Nh_i/(n - 1)$ in the low-type case, and $x_i + B(n - h_i - 1)/(n - 1) - w$ in the high-type case. Using these functions, if firm i is the only one to hire a low-type, then $h_i = n - 1$ and its payoff is $-N$ just as in the two-firm case. As fewer firms choose to hire a high-type CEO, though, i 's external loss is scaled back proportionally so that if only one firm chooses to hire a high-type, i 's payoff is $-N/(n - 1)$. Similarly, if i is the only firm to hire a high-type its profit is $x_i + B - w$, as in the two-firm case, but as more firms hire a high-type the external benefit is scaled back. If all firms hire the same type of CEO the external effect is zero and profit is either $x_i - w$ when every firm hires a high-type, or 0 when every firm hires a low-type.

As before, assume that firms draw the internal effect x_i independently from $F(x)$ with support $[\underline{x}, \bar{x}]$. Conditions (1.1) and (1.2) now have different interpretations. With more than two firms, condition (1.1) states that the negative external effect is large enough to make a firm with the smallest internal effect hire a high-type when all other firms hire a high-type. The first inequality in condition (1.2) states that for the firm with the lowest internal effect hiring a low-type is a best response to hiring a low-type, even if no other firm hires a highly-paid CEO. The second inequality makes the firm with the highest internal effect hire a highly-paid CEO even when it is the only one to do so.

Proposition 1.1. *Suppose that conditions (1.1) and (1.2) hold. Then for any pair of loss- and benefit-sharing rules there exists a Bayes-Nash equilibrium in which every firm chooses to hire a highly-paid CEO, but there does not exist a Bayes-Nash equilibrium in which all firms choose to hire a low-paid CEO.*

Proof. To show that there exists a Bayes-Nash equilibrium in which every firm plays H , note that that if everyone else plays H then the lowest internal effect firm $\underline{x}$ prefers H to L if and only if

$$\underline{x} - w \geq -r(n-1)N = -N, \quad (1.6)$$

where the last equality comes from the requirement that r be a loss-sharing rule. Condition (1.1) then guarantees that the lowest internal effect firm plays H in response to the other $n-1$ firms playing H . However, if (1.6) holds for the firm with $\underline{x}$, it also holds for any other firm with $x > \underline{x}$. Thus, H is a best response to H for all firms and so there exists a Bayes-Nash equilibrium in which all firms play H .

To show that there is no Bayes-Nash equilibrium in which all firms play L , note that if $h = 0$ a firm with internal effect x plays H if and only if

$$0 \leq x + s(0)B - w = x + B - w. \quad (1.7)$$

Condition (1.2) guarantees that a strict version of the inequality (1.7) holds for $x = \bar{x}$, and firms with internal effect in the interval $[w - B, \bar{x}]$ play H even when no other firms do. Therefore any Bayes-Nash equilibrium must have at least some firms playing H , and consequently cannot have every firm playing L . $\square$

The above proposition establishes that equilibria of the game exist, and begins to narrow down the types of behavior that can occur in equilibrium. The next proposition characterizes equilibrium behavior further.

Proposition 1.2. *Suppose that conditions (1) and (2) hold. Then for any pair of loss- and benefit-sharing rules, in every Bayes-Nash equilibrium there exists an $x^* \in [\underline{x}, \bar{x})$ such that every firm with internal effect in the interval $[x^*, \bar{x}]$ chooses to hire a highly-paid CEO and every firm with internal effect in the interval $[\underline{x}, x^*)$ chooses to hire a low-paid CEO.*

Proof. In a Bayes-Nash equilibrium, strategies must respond optimally with their beliefs. Beliefs, in turn, must be consistent with equilibrium strategies. For each firm with internal effect x , then, Bayes-Nash equilibrium specifies an action from $\{H, L\}$ and a belief $\tilde{h}$ about the random variable governing the number of firms playing H . Because internal effects are drawn independently, no firm's x provides any information about the posterior distribution of drawn internal effects. Consequently, $\tilde{h}$ is independent of x .

To prove the assertion suppose, to the contrary, that there exist two firms such that firm 1 with internal effect x_1 plays H and firm 2 with internal effect $x_2 > x_1$ plays L . Given the common beliefs $\tilde{h}$, firm 1 plays H if and only if

$$E \left[x_1 + s(n - \tilde{h} - 1)B - w \right] \geq E \left[-r(\tilde{h})N \right]$$

where the expectation is taken over $\tilde{h}$. Rearranging yields

$$x_1 \geq E \left[-r(\tilde{h})N \right] - E \left[s(n - \tilde{h} - 1)B - w \right]. \quad (1.8)$$

Note that the right-hand side of (1.8) is independent of type. Consequently,

$$x_2 > x_1 \geq E \left[-r(\tilde{h})N \right] - E \left[s(n - \tilde{h} - 1)B - w \right]$$

and firm 2 strictly prefers to play H . This provides a contradiction, thereby establishing the assertion. $\square$

The above result shows that in every equilibrium the firms hiring high-type CEOs are those with large internal effects, and that firms hiring low-type CEOs have smaller internal effects. This is intuitively plausible. It also states that in every equilibrium there is at least a partial arms race, as the firms with the largest internal effects choose to hire highly-paid CEOs and those with slightly smaller internal effects choose to hire highly-paid CEOs in defense. The next proposition establishes circumstances under which an arms race must impact the entire population. It also establishes that, under those circumstances, a CEO arms race *must* occur because it is the only Bayes-Nash equilibrium.

Proposition 1.3. *If conditions (1.1) and (1.2) hold, if $r(\tilde{h}) = \tilde{h}/(n-1)$ and $s(\tilde{h}) = (n - \tilde{h} - 1)/(n-1)$, and if*

$$F(x) < \frac{N - (w - x)}{N - B} \quad (1.9)$$

for all $x \in [\underline{x}, \bar{x}]$, then the Bayes-Nash equilibrium is unique and all firms choose to hire a highly-paid CEO.

Proof. We first show that using proportional scaling functions (1.4) and (1.5) condition (1.8) under which firm i chooses to play H reduces to a simpler expression. Notice that $E[\tilde{h}] = (n-1)(1 - F(x^*))$. Then:

$$\begin{aligned} E[-r(\tilde{h})] &= -\frac{(n-1)(1 - F(x^*))}{n-1} = F(x^*) - 1 \\ E[s(n - \tilde{h} - 1)] &= \frac{n-1 - (n-1)(1 - F(x^*))}{n-1} = F(x^*). \end{aligned}$$

Inequality (1.8) for any firm i becomes:

$$x_i \geq (F(x^*) - 1)N - F(x^*)B + w. \quad (1.10)$$

Rearranging we get:

$$F(x^*) \leq \frac{N - (w - x_i)}{N - B}.$$

For this special case of proportional scaling functions we show that if the distribution of internal effects satisfies (1.9) there is no equilibrium in which L is played with positive probability. Suppose not, so that firm i plays L with positive probability. Because H is a best response to H , for i to play L with positive probability i must place positive probability on some of its rivals also playing L . Because the x_i s are drawn independently and players must follow a cutoff strategy, there exists an interval $[\underline{x}, \xi_i)$ such that i believes that rivals who draw values $x \in [\underline{x}, \xi_i)$

play L and those who draw $x \in [\xi_i, \bar{x}]$ play H . Furthermore, there exists an interval $[\underline{x}, x_i^*)$ such that player i plays L if and only if $x_i \in [\underline{x}, x_i^*)$. Then for player i

$$x_i^* = w - N + F(\xi_i)(N - B).$$

Suppose that $x_i^* \geq \xi_i$. Then

$$x_i^* = w - N + F(\xi_i)(N - B) \leq w - N + F(x_i^*)(N - B)$$

which implies that

$$F(x_i^*) \geq \frac{N - (w - x_i^*)}{N - B}$$

and violates condition (1.9). Consequently, it must be the case that $x_i^* < \xi_i$ for firm i to play L with positive probability. This same condition must hold for every player, which means that every player believes that its rivals use a higher cutoff point than they do. Such beliefs cannot be consistent with Bayes-Nash equilibrium. This establishes a contradiction to the hypothesis that firms play L with positive probability. Thus, when (1.9) holds there is a unique Bayes-Nash equilibrium and every firm chooses H .

□

The three propositions establish conditions under which an arms race for CEO hiring occurs, and under which it is the only equilibrium. In particular, the uniqueness result of Proposition 1.3 establishes conditions under which firms cannot avoid a CEO arms race. To help understand their meaning, it is beneficial to see how the different assumptions impact the outcome. The first assumption was the existence of idiosyncratic internal effects, captured by x_i . Suppose instead that x_i is fixed at x for all i , with the added assumption that $x > w - N$, consistent with condition (1.1). This

condition implies that H is a best response to H , which we had assumed to hold for all types. Two possibilities emerge. If $x \geq w - B$, then H is a (weak) best response to L , so that H is a (weakly) dominant strategy. No arms race explanation is required here, though, because high-type CEO hiring arises entirely from the internal effect, that is, highly-paid CEOs add enough to the firm's profit to justify the high wage regardless of what rivals do. The second case has $x < w - B$, in which case L is a best response to L . The game has two pure strategy equilibria, in which case high-type CEO hiring occurs not because of an arms race but simply as a result of the solution to a coordination problem. Idiosyncratic internal effects prove necessary for the explanation of high-type CEO hiring to come from an arms race. Furthermore, firms must face the possibility that some of their rivals have a sufficiently large internal effect to make hiring a highly-paid CEO a dominant strategy. In this case, if idiosyncratic internal effects are distributed according to condition (1.9), even the small possibility of facing a dominant-strategy opponent will cause contagion. The small group of dominant-strategy players will induce the rest of the firms to hire highly-paid CEOs.

The second assumption was that the external effect consisted of a larger loss from having the lower-paid CEO than the gain from having the higher-paid CEO, that is, $N > B$.⁹ If this fails, condition (1.9) would require that $F(x)$ be negative, which is impossible. As Proposition 1.3 proves, when condition (1.9) fails no arms race occurs. Consequently, asymmetric impacts of CEO type differences, with losers losing more than winners gain, prove crucial for the existence of an arms race in CEO hiring.

We introduce the loss- (benefit-) sharing rule functions that only need to be nondecreasing (nonincreasing) and onto for Propositions 2 and 1.2 to hold. For Proposition 1.3 we reduce the set of possible functions to linear form in order to

⁹This second assumption precludes the possibility of no external effect. Without the external effect there is no reason for firms to care what type of CEO their rivals hire, and the denominator of condition (1.9) is zero. Obviously, some sort of external effect is vital.

be able to calculate their expectations more easily. However, the inequality (1.10) holds for a more general set of concave functions.¹⁰ Concavity means that losses (benefits) of firms that pay low wage (high wage) exceed losses (benefits) that are generated by linear proportional scaling functions.

Finally, Proposition 1.3 links a CEO arms race to condition (1.9), about which more must be said. The left-hand side, $F(x)$, measures the fraction of firms that have internal effects no larger than x , while the right-hand side is an increasing function of x , and both functions are depicted in Figure 1. The key to the condition is the right-hand-side function $(N - (w - x))/(N - B)$. This function takes the value zero when $w - x = N$. Condition (1.1) states that $w - x < N$ for all $x \in [\underline{x}, \bar{x}]$, and so the function $(N - (w - x))/(N - B)$ crosses the horizontal axis to the left of $\underline{x}$. The function takes the value 1 when $w - x = B$, and by (1.2) this occurs for some value of x between $\underline{x}$ and $\bar{x}$, as shown. Thus, the “support” of the function $(N - (w - x))/(N - B)$ is shifted to the left of the support of $F(x)$. As the figure shows, condition (1.9) is satisfied as long as $F(x)$ does not place too much weight in its left tail, and the figure is drawn for a bell-shaped density. Furthermore, in the special case in which x is uniformly distributed, so that $F(x)$ is linear, condition (1.9) is automatically satisfied. Consequently, one can interpret condition (1.9) as saying that the distribution of the internal effect must be “close” to uniform without too much weight, or too large of a spike, in its lower tail. Any distribution that shifts mass rightward from the uniform distribution will automatically satisfy the condition as well. As Proposition 1.3 proves, such distributions of internal effect lead to a CEO arms race.

¹⁰By Jensen’s inequality, $x_i + s(n - E[\tilde{h}] - 1)B - w \geq E[x_i + s(n - \tilde{h} - 1)B - w]$ and $E[-r(\tilde{h})N] \geq -r(E[\tilde{h}])N$, so $x_i + F(x^*)B - w \geq (F(x^*) - 1)N$, which is the same as inequality (1.10).

1.5 Concluding Remarks

This paper identifies circumstances under which firms might engage in an arms race to have higher paid CEOs. Such an arms race requires both an internal effect and an external effect from hiring a highly-paid CEO. The internal effect arises when marginal improvements in own CEO type (or quality) generate marginal increases in own profit, holding rival CEO type constant. The external effect provides a profit boost to firms with higher-paid CEOs and a loss to those with lower-paid CEOs. Arms races occur when the following circumstances hold: (i) the magnitude of the internal effect is heterogeneous, private information, and sufficiently uniform, (ii) there is positive probability that some firms have large enough internal effects that they find it worthwhile to hire high-type (highly-paid) CEOs regardless of what their rivals do, and (iii) the external effect is homogeneous and causes firms to lose more from having the lower-paid CEO than they would gain from having the higher-paid CEO.

For CEOs, the internal effect arises as a higher-paid CEO runs the firm in a more efficient, more profitable way, perhaps because of improved access to capital markets or increased bargaining power with suppliers or labor. The external effect occurs because a higher-paid CEO provides a competitive advantage over a rival with a lower-paid CEO, perhaps because of signaling to customers. Under appropriate distributional assumptions, these two effects can lead to an arms race for CEO pay. The arms race itself arises out of a kind of contagion, albeit one that occurs instantly rather than through time. The firms with the highest internal effects have an incentive to hire highly-paid CEOs, which provides the firms with the next highest internal effects with an incentive to also hire a highly-paid CEO, and so on, until every firm hires a highly-paid CEO.

One implication of the model is that firms hire highly-paid CEOs regardless of their internal effects. Instead, these CEOs are hired not because of what they can do for the firm, but what their type, in and of itself, does for the firm. Essentially, the firm needs a highly-paid CEO to protect its competitive position. This may shed light on one particularly prominent instance of high CEO pay. In 2001 the legendary Jack Welch retired from his CEO position at General Electric. GE chose his successor from among three well-qualified, internal candidates: Jeffrey Immelt, Jim McNerney, and Robert Nardelli. GE promoted Immelt to CEO, while McNerney left to become CEO of 3M and Nardelli left to become CEO of Home Depot. GE and 3M have many similarities, using technological innovation to develop new products in many diverse markets, so one can imagine why McNerney would generate a large, positive internal effect at 3M. It is much harder for one to see how the skills from running GE Power Systems would translate into strong internal effects in the world of big-box retailing at Home Depot. The arms race model predicts that Nardelli's pay would be high anyway, and it was.

The increased relevance of external effects may be one of the drivers why CEO pay increased so dramatically over the last twenty years. An event that possibly made external effects more relevant was the adoption of amendments to the executive compensation disclosure approach by the Securities and Exchange Commission in 1992. The amendments required highly formatted disclosure to facilitate the comparison of annual compensation among companies. Examples of external effects given earlier in the paper included higher-paid CEOs having more credibility in capital markets or signaling higher quality to customers. For these channels to operate outsiders must be able to verify the CEO's pay, and the new disclosure requirements allow this channel to open.

An identical setting, with identical results, could be applied to markets other than that for CEOs. Perhaps the most obvious is the market for mutual fund managers. A highly-paid mutual fund manager might be able to generate higher return on the portfolio, which is the internal effect. At the same time, a highly-paid manager might attract customers away from a rival, whereas a fund with a low-paid manager might lose investors not just to a single rival but to the industry as a whole. This creates an asymmetry for the external effect, with the losses from having a low-paid manager outpacing the gains from having a highly-paid one. Given these circumstances, one might expect to see an arms race for mutual fund managers.

The arms race analysis cannot, however, explain high salaries in professional sports, or for that matter high salaries for collegiate coaches. Hiring marquee athletes or coaches may generate an external effect as the fan base and team revenue may respond to the league's distribution of superstars. At the same time, individual athletes or coaches certainly have an internal effect, but one would be hard pressed to argue that these internal effects are private information given the very public nature of the profession. Consequently, player salaries are unlikely to be driven by an arms race.

Comparing these settings leads to an empirical approach. The paper's main proposition states that when the conditions for an arms race arise, the *only* outcome is an arms race. An appropriate empirical strategy, then, concentrates on the conditions for an arms race rather than an arms race itself, identifying both internal and external benefits of hiring highly-paid CEOs.




Figure 1.1: The Theorem Illustrated

Chapter 2

Supply Chain Uncertainty as a Trade Barrier

2.1 Introduction

Goods traded over long distances are subject to unexpected delays in delivery. At U.S. ports over forty percent of vessels arrive one or more days late. In 2007 the average vessel arriving from China at U.S. ports was 3 days late and the average vessel from Europe missed the arrival date by 2.7 days. Vessels also often arrive early, requiring storage until the date of delivery to a customer or absorption in a production process. This matters for importers that rely on timely delivery, because late arrivals may result in lost demand and early arrivals increase storage costs.

This paper identifies the impact of supply chain uncertainty and inventory management on international trade. The theory examines how inventory management impacts import demand if importers hedge against bad arrival shocks by holding safety stock. The empirics identify the impact of supply chain uncertainty on import

demand and quantify the inventory management costs associated with managing supply side risk.

The theory employs a stochastic inventory model from the logistics and economics literature (Song et al. (2009), Wisner et al. (2005), Eppen and Martin (1988), Baumol and Vinod (1970)) to derive two testable hypotheses. First, to avoid disruptions due to unexpected delays in delivery, importers respond to an increase in delivery time uncertainty with an increase in their safety stock. This increase raises inventory holding costs and reduces import demand relative to locations with a lower degree of supply chain uncertainty. Second, to minimize base-stock inventory costs, an importer trades off a higher order frequency at a fixed ordering cost for lower average inventory holding costs. As a consequence, an importer responds to an increase in order costs with a reduction in shipping frequency which implies an increase in average inventory holding costs. The importer responds to this increase in average inventory holding costs by reducing demand from locations with high ordering costs relative to other locations subject to lower ordering costs.

To test these two predictions we combine several sources of information. To measure supply chain uncertainty across source countries for U.S. imports and districts of entry we employ expected vessel arrival dates filed by shippers and actual arrival dates at U.S. ports for 2007-2009 from Import Genius¹. We obtain ordering costs from Doing Business and import, freight charge and unit value data from the U.S. Census Imports of Merchandise. With these data sources at hand, we construct a panel data set of imports that arrive in the U.S. via ocean vessel reported by district of unloading, source country, year of entry and HS 10 product. This rich source of variation allows us to account for several variables that are suggested by the theory, but are not directly observable. Exploiting variation across source countries of imports

¹importgenius.com

within district-by-product pairs the results imply that a 10 percent increase in supply chain uncertainty lowers trade by up to 4.2 percent. A one percent increase in the ordering costs lowers imports by as much as 1.2 percent. This is evidence that supply chain uncertainty and ordering costs impact the intensive margin of trade consistent with the inventory process from the theory.

A back of the envelope structural model shows that a one standard deviation increase in supply chain uncertainty from the mean raises inventory holding costs by \$627,694 per year if the total shipping quantity is unchanged. This cost is based on an estimated daily per unit inventory cost of about \$0.72. In other words, supply chain uncertainty raises the costs of importing and is therefore a friction to trade.

Identifying sources of trade costs that are otherwise difficult to observe is an ongoing area of research. Hummels and Schaur (2012) quantify the impact of transit time as a trade barrier. Djankov et al. (2010) identify the impact of time delays within countries on trade and Carballo et al. (2012) examine the impact of trade and customs delay on firm level imports and exports. Harrigan and Venables (2006) show that timeliness imposes trade costs that are qualitatively different from monetary costs because of demand or supply uncertainty. Blonigen and Wilson (2008) identify the impact of port efficiency on trade. While all of these articles have examined some aspect of the supply chain in determining trade flows, to our knowledge we are the first to identify supply chain uncertainty as a trade barrier.

In our empirical application we account for several other mechanisms that firms may use to mitigate the impact of uncertainty. However, contrary to the existing literature, we consider how these mechanisms may mitigate supply side as opposed to demand side shocks. Hummels and Schaur (2010) show that firms subject to demand uncertainty speed up their supply chain by substituting into expensive air transport. Also with demand uncertainty in mind, Evans and Harrigan (2005) provide

evidence that firms move closer to the destination market to speed up delivery. We provide evidence that inventories are an alternative means to manage uncertainty. Nevertheless, we also provide evidence that air transport as well as sourcing from close by markets such as Canada and Mexico are relevant strategies to manage supply chain uncertainty.

Alessandria et al. (2010) examine how firms use inventories to respond to demand uncertainty. In their calibration exercise they find that the volatility necessary to explain the large inventory holdings found in the data is about 5 times bigger than the real data-based demand volatility found in Khan and Thomas (2007) and therefore must include other sources of uncertainty not specified in the model. We provide evidence that supply side shocks in the form of unexpected delays in delivery are an additional source that contributes to a firm's inventory. Alessandria et al. (2010) build a dynamic model that examines a firm's optimal inventory and pricing behavior. We abstract from optimal price adjustments when firms run out of inventory. The logistics literature suggests that firms hold safety stock such that they satisfy between 95-99 percent of demand.² While Alessandria et al. (2010) show that price adjustments are important, we use this fact to assume that firms hold enough safety stock to satisfy all of their demand. Therefore, our theory focuses on the year to year problems and costs imposed by inventory management as opposed to the extreme event when firms run out of safety stock.

This paper is also related to literature on trade in intermediate inputs. This type of trade is growing and comprises 40 to 60 percent of total international trade in the modern world (Ramanarayanan (2006)). About one-third of all international trade is intra-firm trade and for the U.S. this number is about 46%, as shown by Antràs (2003). Recent research derives intermediate input demand from a CES production

²For example, see Dullaert et al. (2007) and Fortuin (1980)

functions that exhibits love of variety; firms import all available varieties to lower their production costs (e.g. Kasahara and Lapham (2012) or Amiti and Davis (2012)). We also derive our import demand from a CES production function and we allow for heterogeneity in productivity of importing firms. However, assuming homogeneity in productivity the import demand is akin to the import demand in a standard Melitz (2003) type model. Therefore, our model can be easily interpreted in terms of final goods or intermediate inputs.

Section 2.2 derives the import demand as a function of supply chain uncertainty and ordering costs to derive the main predictions for the empirics. Section 2.3 derives the empirical specifications, details the data construction and identification approach to discuss results and robustness checks. Section 2.4 finishes with some broad conclusions and ideas for future research.

2.2 Theory

This section derives a firm's import demand taking into account that the importer holds inventory to smooth supply chain uncertainty. An increase in the uncertainty of the arrival time of ordered products requires firms to hold a larger amount of safety stock to hedge against bad arrival shocks. As a result, an increase in supply chain uncertainty increases inventory costs and lowers the import demand.

2.2.1 Import Demand

Consider an importer indexed by i who sells a final bundle of goods Q_{it} on the home market. To produce the final good the importer orders products q_{ijt} from international markets indexed by $j \in J$, where J is an exogenous set of source countries³ and

³Similar to the existing literature we do not solve for the endogenous number of source countries. For example, Amiti and Davis (2012) assume that firms import inputs from all available markets,

combines them to the product bundle $Q_{it} = \varphi_i \left(\sum_j q_{ijt}^\rho \right)^{\frac{1}{\rho}}$.⁴ Importers differ in productivity φ_i . For a given optimal bundle Q_{it} the firm wants to supply on the home market, the importer minimizes the costs of importing $\sum_j (m_{ijt} q_{ijt} + F_{jt})$ such that $Q_{it} = \varphi_i \left(\sum_j q_{ijt}^\rho \right)^{\frac{1}{\rho}}$, where m_{ijt} is a constant marginal cost of importing and storing and F_{jt} is a fixed cost of importing from country j . The optimal import demand is then

$$q_{ijt} = \frac{m_{ijt}^{\frac{1}{\rho-1}}}{\left(\sum_j m_{ijt}^{\frac{\rho}{\rho-1}} \right)^{\frac{1}{\rho}}} \cdot \frac{Q_{it}}{\varphi_i} \quad (2.1)$$

Without the productivity parameter φ_i this import demand is similar to Melitz (2003). Therefore, we can interpret q_{ijt} as the demand for the variety i from the aggregate consumption bundle Q_{it} demanded by the representative consumer in any period t . We now derive the constant import cost m_{ijt} as a function of factory gate prices, transit costs and inventory costs.

2.2.2 Importer's Costs

The importer holds inventory to serve gradually arriving demand on the home market. Assume that there is no uncertainty in the arrival of ordered products. In that case the firm trades off ordering costs with inventory holding costs to determine the optimal amount of inventory and the cost minimizing number of shipments within a planning period t . Figure (2.1) shows an importer who orders half of his yearly imports in the beginning of the year and uses it up gradually until the next shipment arrives. More frequent shipping lowers the average amount stored in inventory, but comes with additional ordering costs. Let ordering costs r_{jt} represent all the expenses associated

due to the CES production function's love of variety. For the empirics this assumption is innocuous, as we absorb the aggregate variable with fixed effects.

⁴Finally, $0 < \rho < 1$ is a parameter that determines the elasticity of substitution between imported goods $\theta = 1/(1 - \rho) > 1$.

with ordering a shipment from country j . Total ordering costs per year are then $r_{jt}n_{ijt}$, where n_{ijt} is the total number of orders the firm places during the planning period t . Let w denote the unit annual inventory cost. Because the withdrawal from inventory is linear, the average amount of inventory is $\frac{1}{2} \frac{q_{ijt}}{n_{ijt}}$ and the base stock annual inventory cost is $\frac{1}{2} w \frac{q_{ijt}}{n_{ijt}}$.

Now suppose that the firm holds safety stock to hedge against delays in the arrival time of ordered products.⁵ Let l_{sjt} denote the lead time, the time that passes between ordering and receiving a shipment s in a year t . Let v_{ijt} denote the daily inventory a firm withdraws to supply the home market. If the firm does not hold safety stock, then it will stock out (run out of inventory), if $l_{sjt} > \bar{l}_{jt}$, where $\bar{l}_{jt}$ is the expected delivery time. The logistics literature suggests that firms hold safety stock to keep the probability of stocking out between 1-5 percent.⁶ How much safety stock must a firm hold to stock out with a probability of 1 percent? Let l_{jt}^{99} be a threshold such that $l_{sjt} > l_{jt}^{99}$ with a probability of one percent, $P(l_{sjt} > l_{jt}^{99}) = 0.01$. Then, if the firm can cover the potential wait time $l_{jt}^{99} - \bar{l}_{jt}$ at a withdrawal rate of v_{ijt} , the firm stocks out with a probability of 1 percent. To guarantee a 1 percent stock-out probability, safety stock must then be $(\text{safety stock})^{1\%} = (l_{jt}^{99} - \bar{l}_{jt})v_{ijt}$. The logistics literature shows that $(l_{jt}^{99} - \bar{l}_{jt})v_{ijt} = k\sigma_{ijt}v_{ijt}$, for a normal distribution of the lead time in days with a standard deviation σ_{ijt} . The exogenously given parameter k is called a service factor such that for a 1 percent probability of stocking out k is the 99th percentile of the standard normal distribution. For daily withdrawal rates $v_{ijt} = q_{ijt}/365$ the total expected costs of base and safety stocks are then the standard expected inventory

⁵We assume that varieties imported from different locations are not substitutable in the short run and therefore firms hold buffer stock. Even though goods are substitutable according to the elasticity of substitution framework over longer planning periods, we assume that it is too costly to substitute varieties imported from, say, Germany with varieties from Russia in the short run.

⁶See Dullaert et al. (2007) and Fortuin (1980)

cost⁷

$$\text{IC} = r_{jt}n_{ijt} + \frac{1}{2}w\frac{q_{ijt}}{n_{ijt}} + w\frac{k}{365}\sigma_{ijt}q_{ijt}. \quad (2.2)$$

The expected amount of inventory is calculated as the average amount of base stock over the year plus the safety stock. This is an approximation, because in case of a late shipment the firm will use safety stock to satisfy demand, so the amount of safety stock will decrease. In case of an early shipment the safety stock will increase. With the assumption that the lead time is distributed normally, the average amount of safety stock over the year is equal to $\frac{k}{365}\sigma_{ijt}q_{ijt}$. Johnson and Montgomery (1974) show that this approximation is good if the stock-out time is small relative to the cycle length. When a shipment is extremely late the firm stocks out completely and does not have any imports in its inventory. However, if the firm chooses its service factor k to be large, it only stocks out with a small probability. Additional discussion of this issue can be found in Appendix A.

Inventory cost function (2.2) reflects the key issues of inventory management. First, for a given stock-out probability and annual quantity q_{ijt} , an increase in the lead time uncertainty raises inventory holding costs due to an increase in the safety stock. Second, an increase in the number of shipments raises ordering costs but lowers average inventory holding costs, reflecting a trade-off between ordering costs and the base stock. To obtain the optimal shipping frequency as a function of q_{ijt} , minimize (2.2) with respect to n_{ijt} and solve for $n_{ijt}(q_{ijt}) = \sqrt{\frac{wq_{ijt}}{2r_{jt}}}$. Substitute the number of shipments into the inventory cost (2.2) to obtain the equilibrium inventory cost

$$\text{IC} = \sqrt{2r_{jt}wq_{ijt}} + w\frac{k}{365}\sigma_{ijt}q_{ijt}. \quad (2.3)$$

⁷For examples in the logistics literature see: Baumol and Vinod (1970), Tyworth and O'Neill (1997) or Ray et al. (2005).

First examine the special case where per unit inventory costs are $w_{ijt} = 1/q_{ijt}$. Then the inventory costs are not a function of q_{ijt} and are fully captured by the fixed costs of importing F_{jt} . However, across countries j , the fixed costs of importing vary due to uncertainty in the delivery time and ordering costs. In a standard trade model based on firm level selection into export markets such as Melitz (2003), this suggests that inventory costs impact the extensive margin of trade.

Now examine the impact of an increase in ordering costs and supply chain uncertainty if the per unit inventory holding costs are constant. Note that an increase in the ordering costs r_{jt} increases the inventory costs for any given q_{ijt} . However, the ordering cost is fixed per shipment, but the impact of an increase in the ordering costs on the total inventory costs depends on q_{ijt} . To obtain intuition for this result note that firms that ship a larger quantity spread this quantity over more shipments. As a result, the same increase in the ordering costs impacts a firm that imports a large quantity more than a firm that imports a small quantity. For this reason, a firm with a larger q_{ijt} lowers the number of shipments faster in response to an increase in r_{jt} , raising the average inventory holding costs by more compared to a firm with a smaller q_{ijt} . Therefore, an increase in ordering costs has a larger impact on costs for firms that ship a larger quantity. It is straightforward to see that an increase in supply chain uncertainty raises inventory costs due to an increase in the safety stock.

For large q_{ijt} , the derivative $\frac{\partial IC}{\partial q_{ijt}} = \sqrt{2r_{jt}w/q_{ijt}} + w\frac{k}{365}\sigma_{ijt} \approx w\frac{k}{365}\sigma_{ijt}$. In other words, for large quantities the total inventory cost is linear in q_{ijt} and we can approximate it well with the first order Taylor approximation around some constant c :

$$IC \approx \sqrt{\frac{r_{jt}wc}{2}} + \left(\sqrt{\frac{r_{jt}w}{2c}} + w\frac{k}{365}\sigma_{ijt} \right) q_{ijt}$$

In this case the importer's fixed cost, F_{jt} , depends on ordering and per unit inventory holding costs, but the marginal inventory cost is a function of supply chain uncertainty

as well as ordering costs. Assuming that the importer takes any varieties' factory gate price, p_{ijt} , and per-unit freight rate, f_{ijt} , as given, the constant marginal cost of importing then equals

$$m_{ijt} = p_{ijt} + f_{ijt} + \sqrt{\frac{r_{jt}w}{2c}} + w\frac{k}{365}\sigma_{ijt}.$$

All else equal, supply chain uncertainty and ordering costs impact the constant marginal cost of importing and storing a variety. Therefore, ordering and per unit inventory costs impact the intensive margin of trade. Substitute m_{ijt} into the firms import demand (2.1) to obtain

$$q_{ijt} = \frac{\left(p_{ijt} + f_{ijt} + \sqrt{\frac{r_{jt}w}{2c}} + w\frac{k}{365}\sigma_{ijt}\right)^{\frac{1}{\rho-1}}}{\left(\sum_{jt} m_{ijt}^{\frac{\rho}{\rho-1}}\right)^{\frac{1}{\rho}}} \cdot \frac{Q_{it}}{\varphi_i}. \quad (2.4)$$

Supply chain uncertainty and ordering costs impact the intensive margin of trade. To derive this import demand we make an additional simplifying assumption. We assume that firms do not take into account the small chance that they run out of safety stock when they derive the import demand. In other words, we are violating some form of Jensen's inequality. This assumption is based on the fact from the logisitics literature that firms hold safety stock to satisfy 95-99% of their demand. Therefore, firms act as if they are not really planing to run out of inventory. In theory, we can always increase k to lower the probabiliy of running out, and Appendix A shows on a simplified example that for relative import demands this probability does not matter if one uses CES demand. This assumption emphasizes that we are concerned with the average long run cost of inventory as opposed to the issue of short term adjustments when firms run out of inventory as discussed by Alessandria et al. (2010).

As shown in equation (2.4), delivery time uncertainty represents an additive demand shifter. However, it can enter the import demand equation in a muktiplicative way. Assume that ordering costs also depend on the per unit inventory holding costs, so that the IC function from equation (2.2) becomes:

$$\text{IC} = wr_{jt}n_{ijt} + \frac{1}{2}w\frac{q_{ijt}}{n_{ijt}} + w\frac{k}{365}\sigma_{ijt}q_{ijt}.$$

Then the optimal number of shipments is $n_{ijt}(q_{ijt}) = \sqrt{\frac{q_{ijt}}{2r_{jt}}}$ and assuming freight rates away the marginal costs are

$$m_{ijt} = p_{ijt} + w \left(\sqrt{\frac{r_{jt}}{2c}} + \frac{k}{365}\sigma_{ijt} \right).$$

If the per unit inventory holding costs are proportional to the unit price: $w = ap_{ijt}$, the import demand function from equation (2.4) becomes:

$$q_{ijt} = \frac{\left[p_{ijt} \left(1 + a \left(\sqrt{\frac{r_{jt}}{2c}} + \frac{k}{365}\sigma_{ijt} \right) \right) \right]^{\frac{1}{\rho-1}}}{\left(\sum_{jt} m_{ijt}^{\frac{\rho}{\rho-1}} \right)^{\frac{1}{\rho}}} \cdot \frac{Q_{it}}{\varphi_i},$$

and uncertainty is a multiplicative demand shifter.

The following two predictions summarize the main theoretical result describing the impact of supply chain uncertainty and ordering costs on import demand.

Prediction 1. *All else equal, an increase in the supply chain uncertainty for country j decreases a firm's imports from country j relative to its imports from all other countries.*

The intuition is that an increase in the supply chain uncertainty leads firms to increase their safety stock to hedge against bad arrival shocks. Therefore, an increase in the supply chain uncertainty of country j raises inventory holding costs for imports

from country j relative to other source countries and lowers the firms' import demand relative to all other source countries.

The ordering costs impact the intensive margin of trade as follows.

Prediction 2. *All else equal, an increase in the ordering costs for country j decreases a firm's imports from country j relative to its imports from all other countries.*

The intuition follows the discussion above. Firms that ship a large total quantity split this quantity over a larger number of shipments to economize on average inventory holding costs. As a result, the same increase in ordering costs has a greater impact on the inventory costs of firms that ship a large quantity. As a consequence, for large quantities ordering costs are linear in the quantity and an increase in the ordering costs raises the cost of importing.

2.3 Data, Estimation and Results

To test the predictions we need information on imports, uncertainty, ordering and delivery costs, and prices. This section describes the specification, the data sources and how the variables were constructed. We finish with presenting results as well as robustness checks.

2.3.1 Specification

We employ a highly disaggregated panel dataset of U.S. imports for three years with four dimensions of variation: across districts of entry d , across commodities h , across source countries j and across time t . We do not observe data at the firm level. For all regressions we assume that every district-commodity combination represents one firm. In other words, the firm indicator i in the theory gets replaced with a $h - d$ couple in the empirical section.

Note that the main equation of interest (2.4) is nonlinear in the variables of interest. This poses a challenge for the estimation. While non-linear estimators are available, the difficulty is that identification requires that we absorb several unobserved variables with fixed effects. Given the large number of products and countries this is difficult with a non-linear estimator and we work with two linear approximations instead. First approximate the sum of trade costs in (2.4) by

$$\begin{aligned} \frac{1}{\rho - 1} \ln(p_{ijt} + f_{ijt} + \sqrt{\frac{r_{jt}w}{2c}} + w \frac{k}{365} \sigma_{ijt}) \\ \approx \beta_0 + \beta_1 \ln(p_{hijt}) + \beta_2 \ln(f_{jt}) + \beta_3 \ln(r_{jt}) + \beta_4 \ln(\sigma_{djt}). \end{aligned}$$

Then we obtain the empirical model

$$\ln(q_{hijt}) = \beta_0 + \beta_1 \ln(p_{hijt}) + \beta_2 \ln(f_{jt}) + \beta_3 \ln(r_{jt}) + \beta_4 \ln(\sigma_{djt}) + \epsilon_{hijt} \quad (2.5)$$

where

$$\epsilon_{hijt} = \ln \frac{Q_{hdt}}{\left(\sum_{jt} m_{hijt}^{\frac{\rho}{\rho-1}} \right)^{\frac{1}{\rho}} \varphi_{hd}} + u_{hijt} = v_{hijt} + u_{hijt}$$

captures approximation error, unobserved variables such as productivity differences and aggregate demand. The rich variation in our panel data set allows us to accommodate several assumptions about the disturbance ϵ_{hijt} . If ϵ_{hijt} is not correlated with the trade costs, then an OLS estimator provides consistent estimates for $\beta_0, \dots, \beta_4$. If firms and demands are symmetric across districts but vary across commodities, then $v_{hijt} = \delta_h$, a commodity fixed effect. If demands vary systematically across districts due to differences in productivity or market size, then $v_{hijt} = \delta_{hd}$. If aggregate demands vary by districts, commodities and time then $v_{hijt} = \delta_{hdt}$. This specification of the disturbance accounts for variation in import demands driven by the financial

crisis over the sample period. Even though we do not model quality differences explicitly, differences in quality of a product h across source countries may change the relative import demands. To account for this unobserved variation we specify $\epsilon_{hajt} = \delta_{haj} + u_{hajt}$ where δ_{haj} is a product-by-district-by-exporter fixed effect. In summary, the rich variation of the data allows us to account for several sources of unobserved variation that may impact the identification of the parameters with a dummy variable estimator that pools the data over all dimensions. Specification (2.5) has several advantages. The coefficients are easy to interpret in terms of elasticities and fixed effects are easy to accommodate in the log-linear model. A further advantage is that it solves an identification problem related to the fact that we only observe import weights instead of quantities. Let q_{hajt} be the imports measured in weights and suppose that a constant conversion factor λ_h translates weights to quantities⁸. Then, we can convert variables measured in kg to quantities by multiplying by λ_h . However, given the log separability of the model we obtain

$$\ln(q_{hajt}) = \beta_0 + \beta_1 \ln(p_{hajt}) + \beta_2 \ln(f_{jt}) + \beta_3 \ln(r_{jt}) + \beta_4 \ln(\sigma_{hajt}) + \gamma \ln(\lambda_h) + \epsilon_{hajt}$$

where $\gamma = 1 - \beta_1 - \beta_2$. Therefore, the product level fixed effect accounts for this conversion problem in the log-linear model. The main disadvantage of the log-linear model is that the coefficient estimates do not have a structural interpretation. To obtain a “back-of-the-envelope” structural model, we can log-linearize $\ln(q_{hajt})$ around the variable means based on equation (2.4). Let $B = \bar{p} + \bar{f} + \sqrt{\frac{\bar{r}w}{2c}} + w \frac{k}{365} \bar{\sigma}$. Then we obtain

$$\ln(q_{hajt}) = \beta_0 + \beta_1 \ln(p_{hajt}) + \beta_2 \ln(f_{jt}) + \beta_3 \ln(r_{jt}) + \beta_4 \ln(\sigma_{hajt}) + \epsilon_{hajt} \quad (2.6)$$

⁸See Appendix C for a more detailed discussion

where $\beta_1 = \frac{B}{\rho-1}\bar{p}$, $\beta_2 = \frac{B}{\rho-1}\bar{f}$, $\beta_3 = \frac{B}{\rho-1}\sqrt{\frac{w}{2c\bar{r}}}\bar{r}$ and $\beta_4 = \frac{B}{\rho-1}\frac{wk}{365}\bar{\sigma}$.⁹ The constant, β_0 , absorbs all constant terms around which the linearization is centered. Combining the coefficient estimate β_4 with the coefficient estimate β_2 we can back out $wk/365$ and quantify changes in the total costs of holding safety stock $\frac{wk}{365}\sigma q$.

2.3.2 Data

U.S. Import Data

For most of our variables we rely on U.S. Imports of Merchandise dataset. These data report imports into the U.S. at monthly frequencies disaggregated by district of entry, HS10 product, mode of transportation (air or ocean) and country of origin. Constrained by the data sources we discuss below, we use the data for 2007-2009. We have quantities (kg), the total value of the shipment (U.S.\$) and the total freight charges (U.S.\$). Our theory applies to products that can be stored and managed in inventories. Therefore we focus the identification on manufacturing imports and exclude all other commodities from the sample. We also drop all shipments that come from Canada and Mexico, because most of these shipments use ground transportation.

Let q_{hdt}^o be the total quantity of good h imported via ocean transport from country j in year t arriving in district d . Price per kg is denoted as p_{hdt}^o and is calculated as total value of the shipment divided by weight. Let f_{hdt}^o denote the unit ocean freight rate associated with the shipment computed as the total freight charge associated with the shipment divided by the shipping weight, q_{hdt}^o . Similarly for air shipments we obtain q_{hdt}^a , f_{hdt}^a and p_{hdt}^a .

⁹For a detailed derivation of this linearization see Appendix B.

Shipment level data

We now describe how we generate a proxy for supply chain uncertainty. The data source we employ is a transaction level import database that includes vessel arrival information purchased for the years 2007 to 2009 from Import Genius¹⁰. The main dataset contains information on every import that arrives at a U.S. district by day of entry. For each import we observe the vessel that delivered the product, the country of origin, the last foreign port the vessel cleared and the expected arrival date. All information is collected from the electronic bills of lading filed by the shipper. Import Genius receives this information via a U.S. customs feed and compiles the information. Similar to the imports of merchandise data we focus on manufacturing goods. Eliminating non-manufacturing goods from the sample is more difficult because the data does not report HS10 product codes. To solve this problem we drop all products that include bulk shipments and liquid-carrying containers. This reduces the number of observations by about 12%.¹¹

We compute a district-by-country-by-year measure of supply chain uncertainty.¹² This measure captures the idea that for a given exporter it may be more difficult to be on time at different districts. Across the east and west coasts this may be due to longer travel times and the necessity to cross the Panama Canal. At a given coast, it may be more difficult to be on time at certain districts due to weather and possibly congested ports.

Let S_{dtj} be the total number of vessels that arrive in district d in year t that unload imports sourced from origin country j . Let s_{dtj} identify a unique vessel arrival in S_{dtj}

¹⁰importgenius.com

¹¹In the robustness section we experiment with a sample that includes only intermediate inputs.

¹²The data does not include HS10 codes. Therefore it is difficult to combine any product level information from the Import Genius data with the imports of merchandise. Unfortunately we cannot generate a product specific measure of supply chain uncertainty and merge it with the imports. With our calculation of uncertainty we assume that firms in a certain district are faced with the same variation in the delivery time from a certain country.

and let $AD(s_{djt})$ and $ED(s_{djt})$ denote the actual and estimated arrival day of the shipment. Let $\Delta(s_{djt}) = AD(s_{djt}) - ED(s_{djt})$ denote the number of days the actual arrival deviates from the expected arrival date. Figure 2.2 plots the distribution of $\Delta(s_{djt})$ over all shipments.

We define supply chain uncertainty as the standard deviation of the difference between the actual and expected arrival date:

$$\sigma_{djt} \equiv \sqrt{\frac{\sum_{s \in S_{djt}} \Delta(s_{djt})^2}{S_{djt}}}$$

In the strictest sense this variable measures how punctual vessel arrivals are. On the other hand, if export countries that have difficulty keeping a timely shipping schedule are likely subject to other sources of supply chain uncertainty that we do not observe, then our measure of uncertainty is a proxy for overall supply chain uncertainty.

See histogram 2.3 and Table 2.1 for the statistics of the uncertainty variable. Initially we have 1096770 observations for uncertainty, 3% of them are zeros, so we lose them when making a log transformation. We lose an additional 10% of observations because of merges with other variables.

Uncertainty seems to be lower for higher-income countries, it is very high for Central America and Africa, and the share of shipments that are early is higher whenever the share of late shipments is low. Distance seems to be positively correlated with uncertainty. Also, uncertainty is higher for 2008 and 2009 compared to 2007. All these correlations are supported by the descriptive regression results presented in Table 2.2.

Additional Control Variables

The theory asks that we control for ordering costs r_{jt} to identify the impact of supply chain uncertainty on trade. We obtain proxies for ordering costs from the Doing Business database. They include the costs for documents, administrative fees for customs clearance and technical control, customs broker fees, terminal handling charges and inland transport. This is more than just per shipment fixed costs of ordering as our theory requires, however the variable ordering costs such as inland transport charges are impossible to separate. In addition we collect GDP, and GDP per capita from the World Development Indicators to account for a source countries' level of development. Both variables are in 2000 constant U.S. dollars, foreign currencies converted using the official exchange rate. See Table 2.3 for the summary statistics of all variables used.

2.3.3 Results

Table 2.4 reports the coefficient estimates for specification (2.5) with standard errors clustered by commodity-country-year. Column 1 reports benchmark OLS estimates. Across columns 2 to 6 we include alternative fixed effects to account for unobserved variation that impact the identification. Across all columns, the coefficient estimates on the supply chain uncertainty are consistent with Prediction 1, an increase in the supply chain uncertainty lowers imports. Column 2 introduces commodity fixed effects, columns 3 and 4 introduce commodity-district fixed effects, column 5 presents commodity-district country fixed effects, and column 6 finishes with commodity-district year fixed effects. Consistent with our assumption that product-district pairs identify firms, this means that the coefficients are identified within firms by the variation in uncertainty across importer's sources of products. In other words,

the coefficient estimates suggest that all else equal, a given firm imports relatively less from locations that exhibit a greater degree of supply chain uncertainty.

After accounting for commodity fixed effects, the coefficient estimates are relatively stable across the specifications. The one outlier is the estimate of Model 5 that introduces commodity-by-district-by-exporter fixed effects. This specification is robust with respect to unobserved variation in quality of imports by product across export locations, and it is comforting that even at this level of rigor the estimates are still consistent with Prediction 1. However, the particular fixed effect regiment of this specification eliminates the main source of identifying variation; the fixed effects absorb the cross-exporter-cross-district variation and rely only on variation across time within countries to identify the coefficients. If supply chain uncertainty is driven for example by port infrastructure, then we would not expect to see a large amount of useful variation that helps identify the impact of uncertainty on imports over time. Therefore, while this specification is robust to several sources of unobserved information, we prefer the specifications that exploit the main source of identifying variation such as the specification presented in Model 6.

Next we turn to Prediction 2. Across all but one model, an increase in order costs lowers imports. This results is surprising, as we identify an impact of a variable on the intensive margin of trade that usually is thought of as a fixed cost and therefore mainly operates on the extensive margin of trade. A caveat is necessary. It is difficult to find information that separates fixed ordering from variable trade costs. While several costs reported in the World Bank's doing business data reflect costs that an importer incurs by order, some reflect more of a variable cost. This means that our estimates are not precise about separating fixed ordering costs from variable costs. However, most of the costs are fixed per order in nature and according to the theory they should be included in the specification. This suggests that specifications that

identify import demand omit an important source of trade costs if they do not include fixed ordering costs and the fact that importers hold inventory.

An increase in per unit transport costs and factory gate prices lowers the import demand. This is what we would expect from the theory, but several identification issues need to be addressed. First, as of now we are not really interested in the coefficients on these variables but are mainly concerned with absorbing the variation to identify the impact of supply chain uncertainty and ordering costs on international trade. However, if firms are small and supply curves are flat such as with constant marginal costs, then importers take the factory gate prices as given and the impact of a price change on imports is identified. Similarly, if individual importers are small such that they do not impact the per unit shipping charges set by shippers that serves many firms within a given period, then also the impact of an increase in the per unit freight charge on import demand is identified. Hummels and Schaur (2012) discuss this identification assumption and perform robustness checks using prices and freight rates from lagged periods. They do not find evidence that freight rates are endogenous even at the higher level of aggregation of HS6 commodities.

Columns 4-6 augment the specification with export country GDP and GDP per capita to control for country characteristics. Most traditional models of trade include these parameters. Both variables have a significant impact on imports as we would expect from a long list of gravity estimation. However, including this information in our model does not change our conclusions about supply chain uncertainty or ordering costs.

All of the specifications in Table 2.4 include some sort of year effect. Accounting for systematic differences across years is important due to the potential trade effects of the financial crisis. While models 1 to 5 allow for systematic changes in the overall average import demand, specification allows for time effects that are specific to commodities

and districts. Therefore, model 6 is robust with respect to shocks driven by the financial crisis that vary across industries and districts.

2.3.4 Robustness Checks

Alternative Measures of Supply Chain Uncertainty

First we use alternative measures of uncertainty and report estimation results in Table 2.5. We re-estimate Model 6 of the previous tables, but examine different ways of computing supply chain uncertainty. The estimated date of arrival is supposed to be the same for all shipments on a vessel, and it is for about 70% of observations. Other vessels happen to have different estimated dates of arrival for different shipments due to data entering error or other reasons. To solve this problem we take the most common estimated date for the vessel (mode), obtaining an alternative measure of the difference between the actual and estimated dates of arrival. The uncertainty obtained this way is denoted $\sigma_{djt}^{\text{mode}}$ and its impact on trade is still negative.

Another indicator of uncertainty in the supply chain is the share of shipments that are late or early, share_not_{djt} . It can be interpreted as a probability of a shipment from a certain country to be not on time. To obtain that variable we calculate the share of vessel trips for district d from country j in year t that were late or early out of the total number of vessel trips for that district-country pair. As expected, the results of Model 6.2 show that an increase in the probability of late arrival lowers the import demand.

The third measure of uncertainty comes from the Logistics Performance Index developed by the World Bank. Specifically we employ the timeliness component. This timeliness parameter is very similar to what we estimate as delivery time uncertainty: “timeliness of shipments in reaching destination within the scheduled or expected

delivery time”. It is an index that ranges from 1.38 for Somalia to 4.48 for Germany. Since this measure is increasing with better timeliness, the expected sign on the LPI variable is positive. Column 3 shows that an increase in timeliness raises imports.

Alternative Theories

We now explore mechanisms that firms employ to smooth uncertainty other than inventory management. The literature provides two competing theories. Evans and Harrigan (2005) show evidence that firms move closer to the destination market to ensure timely delivery if there are demand shocks. Hummels and Schaur (2010) provide evidence that firms substitute into fast transport to hedge against demand shocks. We add another possible explanation To identify the possibility of switching between ocean and air transport we introduce the variable rate-mile_{ht}^a . This variable captures the average unit freight rate for air shipping per mile traveled for a particular commodity h . Distance is measured as a straight line from the capital of country j to the U.S. capital.

$$\text{rate-mile}_{it}^a = \sqrt{\frac{\sum_{j=1}^J \frac{\sum_{d=1}^D f_{ijdt}^a}{\text{distance}}}{J}}$$

This variable is an indicator of how expensive it is to ship via air as opposed to shipping by ocean. Next denote the share of trade value that comes from Mexico and Canada as mcshare_{it} for each commodity i .

Table 2.6 provides evidence for the mechanisms introduced by Evans and Harrigan and Hummels and Schaur based on the specification including commodity-district-year fixed effects. To test Hummels and Schaur (2010) we include an interaction term between uncertainty and air unit charges per mile. Products with higher air unit charges make it more difficult to substitute for faster transport, therefore the expected result is that uncertainty in ocean shipping would matter more for those

products. The coefficient on the interaction term is negative, which makes the overall effect of uncertainty larger in absolute value.

To examine the mechanism proposed by Evans and Harrigan we interact the trade share that comes from Canada and Mexico with the supply chain uncertainty. A higher share means that importers source a large amount of a commodity from close-by countries allowing for a fast response time in case an importer runs out of inventory instead of waiting for ocean shipments to arrive from far away. The expected result is that for goods with a high share of trade coming from Canada and Mexico uncertainty in the ocean transit time matters less. The coefficient on the interaction term between the uncertainty and import share from Mexico and Canada is positive. Therefore, supply chain uncertainty has less of an impact for commodities that are heavily sourced from close by locations.

In conclusion, while we still find that supply chain uncertainty has a negative impact on imports as predicted by the theory, this impact is heterogeneous in an importer's ability to source from close by locations or air ship.

Uncertainty and per capita income

Imports from rich countries are more attractive than imports from poor countries. Holding everything else constant, we expect that higher uncertainty will hurt imports from poor countries more than imports from rich countries. In this section we show that uncertainty matters more for products ordered from countries with smaller per capita income. An interaction term between the delivery time uncertainty and a country of origin's GDP per capita has a positive coefficient as shown in column 3 of Table 2.6, which means that lower income leads to a higher impact of uncertainty on imports.

Uncertainty impact by product category

Our estimates of the uncertainty impact are averaged across all commodities. That average hides a lot of variation in how imports of different commodity groups respond to changes in the delivery time uncertainty. One important classification is the differentiated versus homogeneous products. It should be less harmful for an importer of a product to be faced with high delivery shocks if the product is easily substitutable. The first two rows of Table 2.7 show that this idea is supported by our data based on the preferred specification with commodity-country-year fixed effects. Using Rauch (1999) classification we find that uncertainty impact is higher than average for the differentiated goods and lower than average for homogeneous goods.

We also employ the U.S. Census end use classification to study the different impact of uncertainty based on the products' end use characteristics. Imports of goods that are more difficult to store in inventory, such as foods and beverages, respond to uncertainty in a less significant way. Industry supplies, capital goods and machinery are associated with a smaller uncertainty impact than consumer products and automotive vehicles. This could be explained by better organized supply chains for the goods needed in manufacturing, since those goods are especially time sensitive. It could also be that firms that import inputs hold inventory in final output as opposed to inputs and the import cost due to inventory is average over all sources of inputs. Furthermore, Hummels and Schaur (2012) provide evidence that inputs are among the most time sensitive commodities in trade. As a result, firms that require timely delivery may have other sourcing strategies such as purchasing inputs from close by locations or flying them in by more expensive air transport. We examine these alternative channels in section 2.3.4. Final products might also be ordered by a larger number of importers, which means uncertainty will impact more entities and have a greater cumulative effect. For example, imports of footwear go down by 9.8% in

response to a 10% increase in uncertainty, which is a much bigger than average impact. Uncertainty in the delivery of chemicals that are part of the industrial supplies and materials group is associated with a much smaller decrease in trade of about 3%.

Ordering costs have a higher impact on the homogeneous goods, however when split by the end use category, ordering costs impact follows the pattern similar to uncertainty. Imports of consumer goods respond more to the changes in the per order cost than imports of intermediate products represented by industry supplies or capital goods. Import demand for differentiated goods such as footwear goes down by 20% if the costs to order increase by 10%, whereas the demand for homogeneous products like chemicals only goes down by 5%.

2.3.5 Structural Estimation

Finally we estimate the “structural model” equation (2.6). This specification coincides with the reduced form, so we use the results presented in Table 2.4¹³. We pick the coefficients from the last two models because model 6 is our preferred specification and model 5 gives the smallest estimate of the uncertainty impact. The coefficient on σ divided by the coefficient on the freight rate is equal to

$$w_1 \frac{k}{365} \frac{\bar{\sigma}}{\bar{f}} = 0.13$$

$$w_2 \frac{k}{365} \frac{\bar{\sigma}}{\bar{f}} = 1.67$$

For $k = 2.33$, per kilogram inventory holding costs are equal to $w_1 = \$20.11/\text{kg}$ or $w_2 = \$261.75/\text{kg}$. Then, for a one standard deviation increase in uncertainty ($sd(\sigma)$)

¹³Structural specification does not include variables GDP and GDP per capita. We estimated model 6 without those variables, and the final calculations of inventory costs did not change in a significant way

= 0.83), the safety stock costs (2.3) will increase by

$$0.83w_1 \frac{k}{365} \bar{q}_{ijt} = 0.83 \cdot 17.55 \cdot \frac{2.33}{365} \cdot 452606 = 48225.14$$

$$0.83w_2 \frac{k}{365} \bar{q}_{ijt} = 0.83 \cdot 307.39 \cdot \frac{2.33}{365} \cdot 452606 = 627694.24$$

In other words, the cost of additional safety stock due to a one standard deviation increase in σ , holding the total inventory quantity fixed, is about 627,694 dollars per year, if we use our preferred specification estimate. This means that if a firm orders the same quantity from two different countries, it has to spend 627,694 dollars more per year to store products obtained from the country with a one standard deviation higher uncertainty. Inventory holding costs represent the annual cost of adding one more unit to the safety stock. Our estimate shows that these marginal costs are rather large: 262 dollars per kg per year. Freight charges are equal to only 1.1 dollars/kg for ocean shipping or 6.1 dollars/kg for air shipping. Inventory costs are also high compared to unit values, U.S.\$ 36 on average. However, it is important to note that this is the cost if the unit is stored for an entire year. At daily withdrawal rates, the cost of inventory per day is $261.75/365 \approx 0.72$ U.S.\$.

2.4 Conclusions

This paper investigates the impact of supply chain uncertainty and ordering costs on trade flows. Previous studies have been preoccupied with the effects of demand uncertainty on trade. An increase in supply chain uncertainty raises safety stocks, increases inventory holding costs, and reduces imports from locations with high delivery uncertainty. High ordering costs reduce shipping frequency, increase average inventory holding cost for a firm's base inventory stock, and result in fewer imports

from locations where ordering costs are high. Supply chain uncertainty is measured using detailed data on actual and expected arrival times of vessels at U.S. ports. Results indicate delivery time delay significantly reduces trade volumes. A 10 percent increase in supply chain uncertainty reduces imports by as much as 4.2 percent. When ordering costs rise one percent, imports fall by up to 1.2 percent.

Late shipments may be due to factors beyond shipping lines' control. Included here are bad weather, labor strikes, fires, ship collisions, groundings, and delays at previous ports of call. A large nation, like the U.S., imports most shipments over direct trade routes. Trade costs associated with supply chain uncertainty are more important for lower income countries with inadequate port facilities that ship through multiple ports of call. A container ship that misses its contractually negotiated birthing window affects both berth and yard planning at seaport terminals, leading to port congestion. High trade costs associated with supply chain uncertainty suggest much can be gained from reducing port congestion. Countries can reduce congestion at ports by investing in additional container handling capacity and by improving infrastructure. Reductions in supply chain uncertainty can stimulate trade and lead to significant cost savings for the shipper, importing firm, and final consumers.




Figure 2.1: Inventory Process Illustrated




Figure 2.2: Difference between the actual date of arrival and the estimated date of arrival




Figure 2.3: Uncertainty for district d and country j

Table 2.1: Table of summary statistics of uncertainty

		Δ_{vdjt}				Late shipments only				Early shipments only			
		N	avg	med	sd	share	avg	med	sd	share	avg	med	sd
all observations		566453	0.58	0	1.96	0.34	2.3	1.53	2.09	0.13	1.6	1	1.98
years	2007	236781	0.7	0	1.96	0.38	2.35	2	2	0.12	1.55	1	1.85
	2008	171290	0.56	0	2.06	0.33	2.4	1.64	2.24	0.14	1.71	1	2.13
	2009	158382	0.41	0	1.82	0.3	2.1	1.02	2.05	0.14	1.53	1	1.97
regions	n america	10937	0.45	0	1.5	0.27	2.07	1	1.75	0.075	1.6	1	1.82
	c america	78462	0.53	0	1.8	0.31	2.28	1.09	2	0.11	1.64	1	1.78
	s america	47273	0.85	0	2.22	0.39	2.69	2	2.27	0.12	1.68	1	1.84
	europa	156878	0.71	0	2.05	0.37	2.49	2	2.1	0.14	1.56	1	1.85
	asia	247074	0.45	0	1.88	0.32	2.06	1	2.04	0.14	1.59	1	2.13
	australia	10998	0.55	0	1.87	0.35	2.25	1.88	1.86	0.16	1.45	1	1.86
	africa	12985	1	0	2.29	0.44	2.82	2	2.25	0.1	1.76	1	1.76
	puerto rico	1846	0.29	0	1.08	0.26	1.52	1	1.19	0.075	1.32	1	1.43
coasts	east coast	385502	0.66	0	1.96	0.35	2.4	1.98	2.11	0.12	1.56	1	1.84
	west coast	156240	0.38	0	1.92	0.32	2.04	1	2.02	0.17	1.6	1	2.15
income	low income	6278	0.71	0	1.85	0.37	2.34	2	1.95	0.097	1.6	1	1.67
	lower-middle income	193622	0.63	0	2.02	0.35	2.38	1.67	2.19	0.12	1.64	1	2.04
	upper-middle income	73977	0.53	0	1.86	0.33	2.24	1.14	1.98	0.13	1.6	1	1.85
	higher income non-oecd	65420	0.44	0	1.84	0.31	2.1	1	2.01	0.13	1.62	1	2.14
	higher income oecd	207325	0.61	0	1.98	0.35	2.35	1.91	2.06	0.14	1.57	1	1.9

Table 2.2: Descriptive regressions

LHS variable	log of uncertainty	log of share not on time
log of distance	0.158 (0.035)***	0.094 (0.029)***
log of GDP	0.006 (0.011)	-0.006 (0.008)
2008 dummy	0.212 (0.027)***	0.120 (0.028)***
2009 dummy	0.100 (0.042)**	0.035 (0.035)
Constant	-0.869 (0.418)**	-1.392 (0.328)***
R-squared	0.135	0.059
N	343	343

Note: The unit of observation is at the country-year level. Estimator: OLS with robust standard errors that are clustered by country.

Table 2.3: Summary statistics

Variable		Source	Mean	Std. Dev.	N
σ_{djt}	uncertainty by district	Import Genius	1.92	0.828	946087
$\sigma_{djt}^{\text{mode}}$	uncertainty by district, mode	Import Genius	1.928	0.832	946087
share_not $_{djt}$	share not on time	Import Genius	0.47	0.23	933113
lpi	logistics performance index	World Bank	3.991	0.346	944219
q_{ijdt}^o	ocean weight (tons)	US Census	452.606	9106.514	946087
p_{hdjt}^o	value per kg	US Census	35.487	1069.581	946087
f_{ijdt}^a	air unit freight rates	US Census	6.08	29.503	293405
f_{ijdt}^o	ocean unit freight rates	US Census	1.154	32.567	946087
rate-mile $_{it}^a$	air freight rates per mile traveled	US Census	0.001	0.004	293390
share $_{it}^{MC}$	Mexico and Canade trade share	US Census	0.176	0.213	946080
r_{jt}	ordering costs	Doing Business	768.189	290.79	946087
GDP $_{jt}$	GDP (bln \$ US)	World Development Indicators	1347.469	1341.533	946087
GDPC $_{jt}$	GDP per capita (\$ US)	World Development Indicators	13733.095	13065.34	946087

Table 2.4: The effect of uncertainty on log of ocean-shipped weight

Variable (in logs)	(1)	(2)	(3)	(4)	(5)	(6)
uncertainty	-0.114 (0.004)***	-0.262 (0.004)***	-0.405 (0.008)***	-0.390 (0.008)***	-0.025 (0.010)***	-0.417 (0.010)***
price	-1.045 (0.003)***	-1.096 (0.003)***	-1.104 (0.004)***	-1.047 (0.004)***	-0.869 (0.005)***	-1.067 (0.004)***
ordering cost	-0.728 (0.011)***	-1.014 (0.010)***	-1.170 (0.012)***	-0.705 (0.010)***	-0.064 (0.029)	-0.735 (0.012)***
freight rate	-0.207 (0.002)***	-0.131 (0.002)***	-0.131 (0.002)***	-0.143 (0.002)***	-0.117 (0.003)***	-0.150 (0.003)***
year8	0.100 (0.009)***	0.102 (0.008)***	0.107 (0.008)***	0.042 (0.007)***	-0.042 (0.005)***	
year9	-0.024 (0.010)***	-0.005 (0.008)***	-0.003 (0.008)***	-0.126 (0.008)***	-0.296 (0.007)***	
GDP				0.374 (0.003)***	-0.005 (0.273)	0.391 (0.003)***
GDPC				-0.247 (0.003)***	0.984 (0.276)***	-0.253 (0.003)***
Constant	15.950 (0.073)***	18.098 (0.065)***	19.220 (0.079)***	8.074 (0.092)***	2.899 (5.021)	7.891 (0.105)***
Commodity FE	no	yes	no	no	no	no
Comm-district FE	no	no	yes	yes	no	no
Comm-dist-country FE	no	no	no	no	yes	no
Comm-dist-year FE	no	no	no	no	no	yes
R-squared	0.314	0.439	0.489	0.520	0.858	0.466
N	946087	946087	946087	946087	946087	946087

Note: The dependent variable in all specifications is the log of imports in kgs. The unit of observation is at the commodity-district-country-year level. The standard errors are robust and clustered by commodity-country-year. Estimator: Pooled OLS

Table 2.5: Alternative measures of uncertainty

Variable (in logs)	(1)	(2)	(3)
uncertainty (mode)	-0.440 (0.010)***		
share not on time		-0.052 (0.007)***	
LPI			0.426 (0.061)***
price	-1.069 (0.004)***	-1.070 (0.004)***	-1.074 (0.004)***
ordering cost	-0.728 (0.012)***	-0.739 (0.012)***	-0.728 (0.012)***
freight rate	-0.149 (0.003)***	-0.149 (0.003)***	-0.148 (0.003)***
GDP	0.395 (0.003)***	0.394 (0.003)***	0.392 (0.003)***
GDPC	-0.256 (0.003)***	-0.249 (0.003)***	-0.264 (0.004)***
Constant	7.790 (0.107)***	7.535 (0.107)***	7.126 (0.121)***
R-squared	0.464	0.462	0.462
N	931409	931409	931409

Note: The dependent variable in all specifications is the log of imports in kgs. The unit of observation is at the commodity-district-country-year level. The standard errors are robust and clustered by commodity-country-year. Estimator: Pooled OLS with commodity-district-year fixed effects

Table 2.6: The impact of air shipment, import from Mexico and Canada and source country's income on the uncertainty elasticity of vessel imports

Variable (in logs)	(1)	(2)	(3)
uncertainty	-0.935 (0.060)***	-0.477 (0.013)***	-1.297 (0.035)***
uncertainty · air rate per mile	-0.079 (0.009)***		
air rate per mile	dropped		
uncertainty · MexicoCanada share		0.312 (0.045)***	
MexicoCanada share		dropped	
uncertainty · GDPC			0.098 (0.004)***
price	-1.066 (0.004)***	-1.066 (0.004)***	-1.064 (0.004)***
ordering cost	-0.745 (0.012)***	-0.744 (0.012)***	-0.739 (0.012)***
freight rate	-0.151 (0.003)***	-0.151 (0.003)***	-0.151 (0.003)***
GDP	0.394 (0.003)***	0.394 (0.003)***	0.398 (0.003)***
GDPC	-0.255 (0.003)***	-0.255 (0.003)***	-0.303 (0.004)***
Constant	7.899 (0.106)***	7.882 (0.106)***	8.165 (0.105)***
R-squared	0.461	0.461	0.462
N	932405	932405	932405

Note: The dependent variable in all specifications is the log of imports in kgs. The unit of observation is at the commodity-district-country-year level. The standard errors are robust and clustered by commodity-country-year. Estimator: Pooled OLS with commodity-district-year fixed effects

Table 2.7: The effect of uncertainty on log of ocean-shipped weight by category

Product category	Uncertainty impact	Ordering costs impact	Obs N
Homogeneous goods	-0.138 (0.027)***	-0.263 (0.028)***	152035
Differentiated goods	-0.501 (0.012)***	-0.884 (0.015)***	655109
End use category (1 digit)			
Foods and beverages	-0.348 (0.210)*	-0.097 (0.241)	1964
Industry supplies and materials	-0.145 (0.018)***	-0.290 (0.020)***	289877
Capital goods, except automotive	-0.223 (0.019)***	-0.522 (0.021)***	231358
Automotive vehicles, parts and engines	-0.465 (0.043)***	-0.582 (0.049)***	43230
Consumer goods	-0.655 (0.015)***	-1.302 (0.020)***	360741
Nondurable goods	-0.831 (0.020)***	-1.375 (0.026)***	206384
Durable goods	-0.304 (0.023)***	-1.315 (0.031)***	152290
Other goods	-0.171 (0.054)***	-0.105 (0.061)*	10372
End use category examples (5 digit)			
Chemicals (12540)	-0.331 (0.080)***	-0.446 (0.071)***	23527
Footwear (40040)	-0.985 (0.107)***	-2.054 (0.117)***	13517

Note: The dependent variable in all specifications is the log of imports in kgs. The unit of observation is at the commodity-district-country-year level. The standard errors are robust and clustered by commodity-country-year. Estimator: Pooled OLS with commodity-district-year fixed effects

Bibliography

- Aggarwal, R. K. and Samwick, A. A. (1999). Executive compensation, strategic competition, and relative performance evaluation: Theory and evidence. *Journal of Finance*, 54(6):1999–2043. 2
- Alessandria, G., Kaboski, J. P., and Midrigan, V. (2010). Inventories, lumpy trade, and large devaluations. *American Economic Review*, 100(5):2304–39. 26, 32
- Amiti, M. and Davis, D. R. (2012). Trade, Firms, and Wages: Theory and Evidence. *Review of Economic Studies*, 79(1):1–36. 27
- Anderton, C. H. (1989). Arms-race modeling - problems and prospects. *Journal of Conflict Resolution*, 33(2):346–367. 3
- Antràs, P. (2003). Firms, contracts, and trade structure. *Quarterly Journal of Economics*, 118(4):1375–1418. 26
- Baliga, S. and Sjöström, T. (2004). Arms races and negotiations. *The Review of Economic Studies*, 71(2):351–369. 3, 5
- Baumol, W. and Vinod, H. (1970). Inventory Theoretic Model of Freight Transport Demand. *Management Science Series A-Theory*, 16(7):413–421. 24, 30
- Bebchuk, L. A., Fried, J. M., and Walker, D. I. (2002). Managerial power and rent extraction in the design of executive compensation. *University of Chicago Law Review*, 69(3):751–846. 2
- Blonigen, B. and Wilson, W. (2008). Port efficiency and trade flows. *Review of International Economics*. 25
- Brams, S. J. (1985). *Superpower Games: Applying Game Theory to Superpower Conflict*. Yale University Press. 3

- Brams, S. J. and Kilgour, D. M. (1988). *Game Theory and National Security*. Basil Blackwell. 3
- Carballo, J., Graziano, A., Schaur, G., and Volpe, C. (2012). Decomposing the impact of customs and trade delays on exports: Firm level evidence. *Working Paper*. 25
- Carlsson, H. and Van Damme, E. (1993). Global games and equilibrium selection. *Econometrica: Journal of the Econometric Society*, pages 989–1018. 3
- Chang, C., Dasgupta, S., and Hilary, G. (2010). CEO ability, pay, and firm’s performance. *Management Science*, 56:1633–1652. 2
- David, P., Kochhar, P., and Levitas, E. (1998). The effect of institutional investors on the level and mix of CEO compensation. *Academy of Management Journal*, 41(2):200–208. 8
- Djankov, S., Freund, C., and Pham, C. (2010). Trading on time. *Review of Economics and Statistics*, 92(1):166–173. 25
- Dow, J. and Raposo, C. C. (2005). CEO compensation, change, and corporate strategy. *Journal of Finance*, 60(6):2701–2727. 2
- Dullaert, W., Vernimmen, B., Aghezzaf, E.-H., and Raa, B. (2007). Revisiting service-level measurement for an inventory system with different transport modes. *Transport Reviews*, 27(3):273–283. 26, 29
- Dumas, L. J. (1979). Armament, disarmament and national security - theoretical duopoly model of the arms-race. *Journal of Economics Studies*, 6(1):1–38. 3
- Eppen, G. and Martin, R. (1988). Determining safety stock in the presence of stochastic lead time and demand. *Management Science*, 34(11):1380–1390. 24

- Evans, C. and Harrigan, J. (2005). Distance, time, and specialization: Lean retailing in general equilibrium. *American Economic Review*, 95(1):292–313. 25
- Finkelstein, S. and Boyd, B. K. (1998). How much does the CEO matter? The role of managerial discretion in the setting of CEO compensation. *The Academy of Management Journal*, 41(2):179–199. 1
- Foote, S. B. (1992). *Managing the Medical Arms Race: Innovation and Public Policy in the Medical Device Industry*. University of California Press. 3
- Fortuin, L. (1980). Five popular probability density functions: A comparison in the field of stock-control models. *The Journal of the Operational Research Society*, 31(10):pp. 937–942. 26, 29
- Gabaix, X. and Landier, A. (2008). Why has CEO pay increased so much? *Quarterly Journal of Economics*, 123(1):49–100. 2
- Geletkanycz, M. A., Boyd, B. K., and Finkelstein, S. (2001). The strategic value of CEO external directorate networks: Implications for CEO compensation. *Strategic Management Journal*, 22(9):889–898. 2
- Harrigan, J. and Venables, A. J. (2006). Timeliness and agglomeration. *Journal of Urban Economics*, 59(2):300 – 316. 25
- Hartzell, J. C. and Starks, L. T. (2003). Institutional investors and executive compensation. *Journal of Finance*, 58(6):2351–2374. 8
- Hayes, R. M. and Schaefer, S. (2009). CEO pay and the Lake Wobegon effect. *Journal of Financial Economics*, 94(2):280–290. 2, 5, 8
- Hummels, D. L. and Schaur, G. (2010). Hedging price volatility using fast transport. *Journal of International Economics*, 82(1):15 – 25. 25

- Hummels, D. L. and Schaur, G. (2012). Time as a trade barrier. NBER 17758. 25, 42, 46
- James, P. (2002). Is the medical arms race still present in today's managed care environment? *SSRN eLibrary*. 3
- Johnson, L. A. and Montgomery, D. C. (1974). *Operations research in production planning, scheduling, and inventory control*. New York. 30
- Kasahara, H. and Lapham, B. (2012). Productivity and the decision to import and export: Theory and evidence. forthcoming in *Journal of International Economics*. 27
- Khan, A. and Thomas, J. K. (2007). Inventories and the business cycle: An equilibrium analysis of (s, s) policies. *The American Economic Review*, 97(4):pp. 1165–1188. 26
- Kroll, M., Wright, P., Toombs, L., and Leavell, H. (1997). Form of control: A critical determinant of acquisition performance and CEO rewards. *Strategic Management Journal*, 18(2):85–96.
- Malmendier, U. and Tate, G. (2009). Superstar CEOs. *Quarterly Journal of Economics*, 124(4):1593–1638. 2, 5
- Melitz, M. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, 71(6):1695–1725. 27, 28
- Murphy, K. J. and Zábojník, J. (2004). New developments in human-capital theory - CEO pay and appointments: A market-based explanation for recent trends. *American Economic Review*, 94(2):192–196. 1

- Murphy, K. J. and Zábajník, J. (2007). Managerial capital and the market for CEOs. *SSRN eLibrary*. 2
- Okuguchi, K. (1981). Stability of the arms-race models. *Zeitschrift Fur Nationalökonomie - Journal of Economics*, 41(3-4):353–360. 3
- Rajgopal, S., Shelvin, T., and Zamora, V. (2006). CEOs’ outside employment opportunities and the lack of relative performance evaluation in compensation contracts. *The Journal of Finance*, 61(4):1813–1844. 2
- Ramanarayanan, A. (2006). International trade dynamics with intermediate inputs. Job Market Paper. 26
- Rauch, J. (1999). Networks versus markets in international trade. *Journal of international Economics*, 48(1):7–35. 46
- Ray, S., Li, S., and Song, Y. (2005). Tailored supply chain decision making under price-sensitive stochastic demand and delivery uncertainty. *Management Science*, 51(12):pp. 1873–1891. 30
- Richardson, L. F. (1960). *Arms and Insecurity: A Mathematical Study of the Causes and Origins of War*. The Boxwood Press. 3
- Schaefer, S. (1998). The dependence of pay-performance sensitivity on the size of the firm. *Review of Economics and Statistics*, 80(3):436–443. 2
- Song, J.-S., Zhang, H., Hou, Y., and Wang, M. (2009). The effect of lead time and demand uncertainties in (r, q) inventory systems. *Operations Research*. 24
- Tyworth, J. and O’Neill, L. (1997). Robustness of the normal approximation of lead-time demand in a distribution setting. *Naval Research Logistics*, 44(2):165–186. 30

Westphal, J. D., Boivie, S., and Ming Chng, D. H. (2006). The strategic impetus for social network ties: reconstituting broken CEO friendship ties. *Strategic Management Journal*, 27:425–445. 2

White, M. J. (2004). The “arms race” on American roads: The effect of sport utility vehicles and pickup trucks on traffic safety. *Journal of Law & Economics*, 47(2):333–355. 3

Wisner, J. D., Tan, K.-C., and Leong, G. K. (2005). *Principles of supply chain management a balanced approach*. Thomson South-Western. 24

Appendix

A Jensen's inequality and relative demands

We show that ignoring probabilities of a stock-out does not make a difference for the relative import demands prediction.

Assume the CES production function for output Q and two inputs q_1 and q_2 from two different countries: $Q = (q_1^\alpha + q_2^\alpha)^{\frac{1}{\alpha}}$. For a simplified example let π be the probability that all of q_i arrives, $(1 - \pi)$ is the probability that only half of the ordered quantity arrives. Import costs have the general form of $m_i q_i$. There are four possibilities: both shipments arrive in full, one gets in reduced by half or both arrive reduced. Then the objective function becomes:

$$\begin{aligned} \min_{q_1, q_2} \quad & m_1 q_1 + m_2 q_2 - \lambda (\pi^2 (q_1^\alpha + q_2^\alpha)^{\frac{1}{\alpha}} + \pi(1 - \pi) (q_1^\alpha + \frac{q_2^\alpha}{2})^{\frac{1}{\alpha}} + \\ & + (1 - \pi)\pi (\frac{q_1^\alpha}{2} + q_2^\alpha)^{\frac{1}{\alpha}} + (1 - \pi)^2 (\frac{q_1^\alpha}{2} + \frac{q_2^\alpha}{2})^{\frac{1}{\alpha}}) \end{aligned} \quad (7)$$

First order conditions for imports from two countries are:

$$\begin{aligned} \frac{m_1}{\lambda} = & \pi^2 q_1^{\alpha-1} (q_1^\alpha + q_2^\alpha)^{\frac{1}{\alpha}-1} + \pi(1 - \pi) q_1^{\alpha-1} (q_1^\alpha + \frac{q_2^\alpha}{2})^{\frac{1}{\alpha}-1} + \\ & + \frac{1}{2} (1 - \pi) \pi q_1^{\alpha-1} (\frac{q_1^\alpha}{2} + q_2^\alpha)^{\frac{1}{\alpha}-1} + \frac{1}{2} (1 - \pi)^2 q_1^{\alpha-1} (\frac{q_1^\alpha}{2} + \frac{q_2^\alpha}{2})^{\frac{1}{\alpha}-1} \end{aligned} \quad (8)$$

$$\begin{aligned} \frac{m_2}{\lambda} = & \pi^2 q_2^{\alpha-1} (q_1^\alpha + q_2^\alpha)^{\frac{1}{\alpha}-1} + \pi(1 - \pi) q_2^{\alpha-1} (q_1^\alpha + \frac{q_2^\alpha}{2})^{\frac{1}{\alpha}-1} + \\ & + \frac{1}{2} (1 - \pi) \pi q_2^{\alpha-1} (\frac{q_1^\alpha}{2} + q_2^\alpha)^{\frac{1}{\alpha}-1} + \frac{1}{2} (1 - \pi)^2 q_2^{\alpha-1} (\frac{q_2^\alpha}{2} + \frac{q_1^\alpha}{2})^{\frac{1}{\alpha}-1} \end{aligned} \quad (9)$$

Dividing (8) by (9) clearly shows that probabilities cancel out and the import demand from country 1 relative to import demand from country 2 depends only on the respective import costs:

$$\left(\frac{q_1}{q_2} \right)^{\alpha-1} = \frac{m_1}{m_2},$$

and does not depend on the probabilities of a stock-out.

B Log-linearization

For simplification let

$$A = \frac{Q_{it}}{\varphi_i \left(\sum_{jt} m_{ijt}^{\frac{\rho}{\rho-1}} \right)^{\frac{1}{\rho}}},$$

$$B = \bar{p} + \bar{f} + \sqrt{\frac{\bar{r}w}{2c}} + w \frac{k}{365} \bar{\sigma}$$

Rewrite the main equation (2.4) such that every endogenous variable x is replaced with $e^{\ln(x)}$, and take logs of both sides:

$$\ln [e^{\ln(q_{ijt})}] = \ln \left[A \left(e^{\ln(p_{ijt})} + e^{\ln(f_{jt})} + \sqrt{\frac{e^{\ln(r_{jt})}w}{2c}} + w \frac{k}{365} e^{\ln(\sigma_{ijt})} \right)^{\frac{1}{\rho-1}} \right] \quad (10)$$

Using first-order Taylor approximation of equation (10) for all endogenous variables around $(\ln(\bar{q}), \ln(\bar{p}), \ln(\bar{f}), \ln(\bar{r}), \ln(\bar{\sigma}))$, we obtain:

$$\begin{aligned} \text{LHS} &= \ln(\bar{q}) + \frac{1}{\bar{q}} e^{\ln(\bar{q})} (\ln(q_{ijt}) - \ln(\bar{q})) = \ln(q_{ijt}) \\ \text{RHS} &= \ln \left(AB^{\frac{1}{\rho-1}} \right) + \frac{1}{AB^{\frac{1}{\rho-1}}} AB^{\frac{\rho}{\rho-1}} \frac{1}{\rho-1} e^{\ln(\bar{p})} (\ln(p_{ijt}) - \ln(\bar{p})) + \\ &\quad + \frac{1}{AB^{\frac{1}{\rho-1}}} AB^{\frac{\rho}{\rho-1}} \frac{1}{\rho-1} e^{\ln(\bar{f})} (\ln(f_{jt}) - \ln(\bar{f})) + \\ &\quad + \frac{1}{AB^{\frac{1}{\rho-1}}} AB^{\frac{\rho}{\rho-1}} \frac{1}{\rho-1} \sqrt{\frac{w}{2c\bar{r}}} e^{\ln(\bar{r})} (\ln(r_{jt}) - \ln(\bar{r})) + \\ &\quad + \frac{1}{AB^{\frac{1}{\rho-1}}} AB^{\frac{\rho}{\rho-1}} \frac{1}{\rho-1} \frac{wk}{365} e^{\ln(\bar{\sigma})} (\ln(\sigma_{ijt}) - \ln(\bar{\sigma})) \end{aligned}$$

We then combine the LHS and the RHS and rearrange to get:

$$\begin{aligned}
\ln(q_{ijt}) = & \ln A + \frac{1}{\rho-1} \ln B - \frac{B}{\rho-1} \bar{p} \ln(\bar{p}) - \frac{B}{\rho-1} \bar{f} \ln(\bar{f}) - \\
& - \frac{B}{\rho-1} \sqrt{\frac{w}{2c\bar{r}}} \bar{r} \ln(\bar{r}) - \frac{B}{\rho-1} \frac{wk}{365} \bar{\sigma} \ln(\bar{\sigma}) + \frac{B}{\rho-1} \bar{p} \ln(p_{ijt}) + \\
& + \frac{B}{\rho-1} \bar{f} \ln(f_{jt}) + \frac{B}{\rho-1} \sqrt{\frac{w}{2c\bar{r}}} \bar{r} \ln(r_{jt}) + \frac{B}{\rho-1} \frac{wk}{365} \bar{\sigma} \ln(\sigma_{ijt})
\end{aligned}$$

Now rename the coefficients:

$$\begin{aligned}
\beta_0 = & \ln A + \frac{1}{\rho-1} \ln B - \frac{B}{\rho-1} \bar{p} \ln(\bar{p}) - \frac{B}{\rho-1} \bar{f} \ln(\bar{f}) - \\
& - \frac{B}{\rho-1} \sqrt{\frac{w}{2c\bar{r}}} \bar{r} \ln(\bar{r}) - \frac{B}{\rho-1} \frac{wk}{365} \bar{\sigma} \ln(\bar{\sigma}) \\
\beta_1 = & \frac{B}{\rho-1} \bar{p} \\
\beta_2 = & \frac{B}{\rho-1} \bar{f} \\
\beta_3 = & \frac{B}{\rho-1} \sqrt{\frac{w}{2c\bar{r}}} \bar{r} \\
\beta_4 = & \frac{B}{\rho-1} \frac{wk}{365} \bar{\sigma}
\end{aligned}$$

Change the subscript from i to hd to match the level of aggregation in our data and obtain:

$$\ln(q_{hdjt}) = \beta_0 + \beta_1 \ln(p_{hdjt}) + \beta_2 \ln(f_{jt}) + \beta_3 \ln(r_{jt}) + \beta_4 \ln(\sigma_{hdjt}) + \epsilon_{hdjt},$$

which is our structural estimation of the model, equation (2.6).

C Units vs. kilograms

Assume a switching parameter γ that transforms units into kilograms, with γ being constant within one 10-digit industry code: $q_{ijt} = \gamma q'_{ijt}$, where q_{ijt} is quantity in units and q'_{ijt} is quantity in kgs. Similarly the purchasing price per unit and freight charges per unit transform into price per kg and charges per kg as $p_{ijt} = \frac{1}{\gamma} p'_{ijt}$ and $f_{jt} = \frac{1}{\gamma} f'_{jt}$. Then import demand equation (2.4) becomes:

$$\gamma q'_{ijt} = \frac{\left(\frac{1}{\gamma} p'_{ijt} + \frac{1}{\gamma} f'_{ijt} + \sqrt{\frac{r_{jt} w}{2c}} + w \frac{k}{365} \sigma_{ijt} \right)^{\frac{1}{\rho-1}}}{\left(\sum_{jt} m_{ijt}^{\frac{\rho}{\rho-1}} \right)^{\frac{1}{\rho}}} \cdot \frac{Q_{it}}{\varphi_i}.$$

Denote

$$A = \frac{Q_{it}}{\varphi_i \left(\sum_{jt} m_{ijt}^{\frac{\rho}{\rho-1}} \right)^{\frac{1}{\rho}}},$$

$$C = \frac{1}{\gamma} \bar{p} + \frac{1}{\gamma} \bar{f} + \sqrt{\frac{\bar{r} w}{2c}} + w \frac{k}{365} \bar{\sigma}$$

After log-linearizing as described in Appendix B we get:

$$\begin{aligned} \ln(q'_{ijt}) = & \ln A + \frac{1}{\rho-1} \ln C - \frac{1}{\gamma} \frac{C}{\rho-1} \bar{p} \ln(\bar{p}) - \frac{1}{\gamma} \frac{C}{\rho-1} \bar{f} \ln(\bar{f}) - \\ & - \frac{C}{\rho-1} \sqrt{\frac{w}{2c\bar{r}}} \bar{r} \ln(\bar{r}) - \frac{C}{\rho-1} \frac{wk}{365} \bar{\sigma} \ln(\bar{\sigma}) + \frac{1}{\gamma} \frac{C}{\rho-1} \bar{p} \ln(p_{ijt}) + \\ & + \frac{1}{\gamma} \frac{C}{\rho-1} \bar{f} \ln(f_{jt}) + \frac{C}{\rho-1} \sqrt{\frac{w}{2c\bar{r}}} \bar{r} \ln(r_{jt}) + \frac{C}{\rho-1} \frac{wk}{365} \bar{\sigma} \ln(\sigma_{ijt}) \end{aligned}$$

Combining all the constant terms into the intercept and renaming coefficients, the structural estimation becomes:

$$\ln(q'_{hajt}) = \beta_0 + \beta_1 \frac{1}{\gamma} \ln(p'_{hajt}) + \beta_2 \frac{1}{\gamma} \ln(\gamma f'_{jt}) + \beta_3 \ln(r_{jt}) + \beta_4 \ln(\sigma_{hajt}) + \epsilon_{hajt}$$

Estimates of β_4 and β_2 will give us the following:

$$\frac{\beta_4}{\beta_2} = \frac{\frac{C}{\rho-1} \frac{wk}{365} \bar{\sigma}}{\frac{C}{\rho-1} \frac{1}{\gamma} \bar{f}} = \frac{wk}{365} \frac{\bar{\sigma}}{\bar{f}} \gamma,$$

which we estimate to be equal to 1.67. From here we find that the per kg inventory holding costs $w\gamma$ are equal to \$262.

The switching parameter does not change the interpretation of the structural estimates.

Vita

Valentina Kozlova was born in Krasnoyarsk-26, Russia. Valentina graduated with honors from Liceum 103 “Harmony” in 2002. She got her undergraduate degree in Economics from the Siberian Federal University in Krasnoyarsk, Russia. In 2007 Valentina started her graduate studies at the University of Tennessee in Knoxville. Valentina obtained a Master of Arts and a Doctor of Philosophy degrees from UT. She has accepted a position of an Assistant Professor of Economics at the State University of New York at Oswego to start in August 2013.