

How Importers Hedge Demand Uncertainty: Inventory, Dual Sourcing and the Effects of Trade Policy¹

Chris Muris² Horst Raff³ Nicolas Schmitt⁴ Frank Stähler⁵

August 24, 2019

¹We would like to thank the editor Stephen Yeaple, two anonymous referees, Dennis Novy and seminar audiences for helpful comments and suggestions. We would also like to thank Max Cojahmetov, Jasleen Grewal, Gigi Lin and Brandon Wong for their research assistance. Financial support from SSHRC (Grant 435-2016-0354) and from the SFU VPR USRA program is gratefully acknowledged.

²Department of Economics, University of Bristol, BS8 1TU Bristol, UK. Email: chris.muris@bristol.ac.uk

³Department of Economics, University of Kiel, 24118 Kiel, Germany; Kiel Centre for Globalization; CESifo; Email: raff@econ-theory.uni-kiel.de

⁴Department of Economics, Simon Fraser University, Burnaby BC V5A 1S6, Canada; CESifo; Email: schmitt@sfu.ca (Corresponding author)

⁵University of Tübingen, University of Adelaide, CESifo and NoCeT, Mohlstr. 36 (V4), D-72074 Tübingen, Germany. Email: frank.staehler@uni-tuebingen.de

Abstract

This paper formulates a dynamic model of inventory investment to examine how importers may hedge demand uncertainty when importing intermediate goods takes time and orders have to be placed before the realization of demand is known. We consider two hedging strategies—building up inventory of imported goods, and using expensive domestic supplies to cover demand surges—and show how tariff protection affects hedging and hence inventory investment, the probability of stockouts and the import share. We find robust empirical evidence for these predictions from high-frequency transaction data for a U.S. steel wholesaler experiencing an episode of U.S. Section-201 tariffs on steel imports. In response to the tariff, inventories decline by up to 90%, accompanied by a significant decrease in the import share.

JEL classification: F12, L81.

Keywords: International trade, trade policy, import tariff, dual sourcing, inventory, demand uncertainty

1 Introduction

This paper develops a dynamic model of optimal inventory investment to study how an importer may hedge demand uncertainty when importing takes time and imports have to be ordered before the realization of demand is known. We consider two strategies to hedge this uncertainty: building up inventory of imported goods, and relying on dual sourcing, i.e., using more expensive, but quickly available domestic supplies to cover demand surges. In this setting, we examine how the firm’s optimal hedging strategy—and hence its inventory levels, stockout probability and import share—is affected by trade policy and by changes in the level of demand uncertainty. We test some of the model’s predictions using daily data on purchases, sales and inventories of 104 different steel products by an anonymous U.S. steel wholesaler over a period starting on February 11, 2000 and ending on July 8, 2004.¹ Trade policy effects are identified by considering U.S. Section-201 tariffs on steel imports imposed on March 20, 2002 and lifted on December 5, 2003.

We find robust empirical support for the economic mechanisms explored in our model. In particular, the imposition of the steel tariffs has statistically significant and economically sizeable effects on inventory, the stockout probability and the import share. In the first year following the imposition of a 30% steel tariff, inventories of products hit by the tariff decline by up to 90% (for products with a high initial import share), compared to products in a control group. In particular, measuring inventory in terms of average monthly sales, we find that inventory for products in the high initial-import-share category declines by the equivalent of 3.22 months of sales from a mean of 3.59 months. These effects persist in the remaining tariff period, and they are significant even for products with a low initial import share, for which inventory still declines by 13% on average.

The model predicts that the reduction in inventory is accompanied either by an increased stockout probability (for products that have no or only very expensive domestic substitutes), or by a switch toward domestic sourcing, i.e., a fall in the foreign sourcing share (for products with only moderately more expensive domestic substitutes). Indeed we find that for most products the wholesaler switches to domestic substitutes when the tariff hits, while reducing the stockout probability. But, for products with the highest initial import share, the probability that the product is out of stock on any given day

¹We thank George Hall for making the data available to us. The original data set covers more products and a longer period of time. See Section 5 and Hall and Rust (2000) for further details.

remains high at around 14%. As products are increasingly sourced domestically when they are exposed to the tariff, the import share falls substantially, for instance by over 27.6% on average for products with a high initial import share, and by even more in the case of products with a lower initial import share. Finally, we find evidence for the model's prediction that demand uncertainty magnifies the effect of tariff protection, as products experience a significantly larger decrease in inventory when hit by the tariff the greater is the level of demand uncertainty.

As we explain in detail in what follows, the general message of our paper is that tariffs may have severe consequences especially for firms that are exposed to demand uncertainty. First, as firms cut inventories to absorb the tariff, they may be more likely to run out of imported inputs thereby foregoing potentially profitable sales opportunities. Second, if firms want to avoid risking a higher frequency of stockouts, they have to switch to domestic substitutes that come at a substantial premium.

The long time delays involved in shipping goods across countries and clearing customs have been carefully documented (Hummels and Schaur, 2013).² These delays are especially problematic for importers exposed to demand uncertainty, because they mean that imports frequently have to be ordered before the firm knows the realization of demand and hence before it knows the quantity of inputs it needs. If demand turns out to be bigger than expected, the firm risks a stockout, which may mean idling production capacity and/or foregoing profitable sales opportunities. In case of lower-than-expected demand, the firm may accumulate stocks of unused goods that may be costly to store or subject to depreciation or spoilage. This suggests a clear trade-off for the firm when ordering inputs involves long lead times, namely between investing in greater inventory to prevent stockouts, and keeping inventory small but then risking either a stockout or having to resort to quicker but typically more expensive sources of supply, such as domestic suppliers or expensive air shipments. How best to deal with this trade-off and thereby hedge demand uncertainty is an important practical challenge for many importers and a key research question in operations management (see Jain et al. (2014), and Cachon and Terwiesch

²Ocean shipping times between various ports around the world can be downloaded from several webpages, such as, <https://www.searates.com/de/reference/portdistance> or <https://www.championfreight.co.nz/times.pdf>. The World Bank provides information on the time required for border and documentary compliance when goods are exported: <http://www.doingbusiness.org/data/exploretopics/trading-across-borders>.

(2019)).³

We develop a dynamic inventory model with dual sourcing that captures some of the essential elements found in the operations management literature (see, for instance, Allon and Van Mieghem (2010), and Boute and Van Mieghem (2015) for a literature review). The focus of that literature, however, is different from ours, as it is often interested in providing asymptotically optimal solutions to sophisticated sourcing and inventory management problems, whereas our interest lies in studying the effects of trade policy on an importer’s optimal hedging strategy. By keeping our model sufficiently simple we obtain analytically tractable solutions that allow us to formulate testable hypotheses about the effects of trade policy and demand uncertainty on inventory, the likelihood of stockouts and the domestic sourcing share, among other things.

We test some of these hypotheses using high-frequency data on purchases, sales, and inventory from a U.S. steel wholesaler. These data are ideally suited for our purposes as they allow us to overcome three main challenges in studying the inventory investment and dual sourcing strategies of importers. First, inventory data are typically not available at the firm-product level. By contrast, we observe inventory data for each product on a daily basis in terms of both quantity and value. Other studies of inventory investment by importers have to rely on much more aggregated data. Jain et al. (2014), for instance, rely on quarterly observations of the value of firm-level inventory of publicly traded U.S. wholesalers and retailers. Alessandria et al. (2010a) use yearly data of the value of inventory held by Chilean manufacturing plants only broken down into materials and goods in process.

Second, dual sourcing is difficult to identify in the data because, unlike international trade transactions, domestic purchases are often not recorded and because it is often unclear whether goods sourced from different origins are indeed close substitutes. Our data exhibit considerable variation in sourcing patterns across products and time, with many products being both imported and sourced domestically. We can also be fairly certain that for the products for which we observe dual sourcing the variants that are sourced domestically and those imported from abroad are indeed very close substitutes,

³Note that dual sourcing in the sense that we use it, namely ordering the same input from an inflexible but cheap and from a flexible but expensive source to hedge demand uncertainty, is different from “multi-sourcing” analyzed, for instance, by Gervais (2018). Multi-sourcing in Gervais (2018) refers to risk-averse firms sourcing the same input from multiple suppliers to hedge idiosyncratic supply shocks. Firms are assumed not to hold any inventory.

because the wholesaler itself treats them as perfect substitutes in its purchase, inventory and sales statistics.

Third, since importing and sourcing are, of course, endogenous we require a trade policy shock exogenous to the firm to identify the effects. Our sample period includes such a trade policy shock, namely the imposition of Section-201 tariffs by the United States on 272 separate 10-digit HS steel products between March 20, 2002 and December 4, 2003. Among the steel products imported by the steel wholesaler we observe a group of products on which the tariff was imposed and a group of products that were subject to no or only a smaller import tariff.

Our partial equilibrium model is closely related to models previously used to study how importers or exporters deal with demand uncertainty. Hummels and Schaur (2010) develop a partial equilibrium model, in which an exporter may hedge demand uncertainty by resorting to expensive but fast air shipments to cover demand surges, while relying on slow but inexpensive ocean shipping for baseline demand (see also Aizenman, 2004). While this mechanism is essentially the same as our dual sourcing mechanism, Hummels and Schaur consider only one demand period (or, equivalently, assume that inventory perishes after one demand period), which is enough to allow them to estimate the premium a firm is willing to pay for air transport. In our dynamic inventory model, by contrast, the firm may also build up inventory and hold this inventory across demand periods. If the trade cost is sufficiently low, the firm in our model only uses inventory to hedge demand uncertainty and may choose to stock out rather than source domestically. We require this dynamic element of our model to derive testable hypotheses about the effect of an import tariff on inventory, stockouts and domestic sourcing.

In the partial equilibrium model of Alessandria et al. (2010a), like in our model and similar to Hummels and Schaur (2010), importers face demand uncertainty and have to wait a period before import orders arrive. Alessandria et al. allow an importer to build up inventory to hedge demand uncertainty, but do not consider dual sourcing. Similarly, in the dynamic inventory model of Novy and Taylor (2019), a firm may source an intermediate good either domestically or from abroad but not both. High-price domestic inputs arrive without delay, but low-price imports arrive only with a time lag; hence the necessity in the case of imported inputs to hedge demand uncertainty by building up inventory. By contrast, dual sourcing is an essential element of our model, given that we want to formulate hypotheses about how the steel wholesaler responds to an import tariff. The

reason is that the steel wholesaler engages in dual sourcing for many of the products it carries, making it an important margin of adjustment to the tariff. Of course, we also know from the operations management literature cited above that dual sourcing is a strategy used extensively by importers across many industries in practice.⁴

On the empirical side, our paper is related to Jain et al. (2014) who study how much more inventory the average US retailer or wholesaler has to hold when it imports goods instead of sourcing them domestically. They argue that an increase in the firm's global sourcing share raises total inventory, whereas dual sourcing is associated with a significant reduction in inventory holdings compared to pure foreign sourcing. Our paper differs from Jain et al. in three important respects. First, unlike Jain et al., we construct an explicit theoretical model of the firm's inventory and dual sourcing decisions that makes clear how these decisions depend on the relative cost of imports versus domestic goods, which allows us to derive testable hypotheses about the effects of a tariff. We also use the model to guide the empirical implementation. Second, while their sample contains 177 firms and detailed information on import transactions by country of origin assembled from bills of lading, Jain et al. observe inventory only at the firm level, on a value basis and at quarterly frequency. By contrast, while we consider only one firm, we are able to trace inventory, both in terms of weight and value, on a daily basis for 104 distinct products over more than four years. One instance where this obviously makes a difference is when it comes to examining stockouts, which is an integral part of hedging demand uncertainty and can only be observed in high-frequency, product-level inventory data. Third, whereas Jain et al. regress inventory on an (endogenous) global sourcing share and (endogenous) supplier dispersion index, we base our identification strategy on an exogenous trade policy shock.

In the next section, we propose a simple dynamic model of inventory investment and dual sourcing that we use, in Section 3, to derive the firm's optimal inventory. In Section 4,

⁴Kropf and Sauré (2014) consider inventory investment by importing firms, but the motive behind it is to save on fixed costs per shipment rather than to avoid stockouts, as demand in their model is deterministic. They also do not consider dual sourcing.

Evans and Harrigan (2005) show that retailers tend to source items subject to uncertain demand from nearby countries, even if production costs there are higher than in alternative, more distant locations. They do not consider inventory investment, as they focus on how retailers source apparel that quickly goes out of fashion and can therefore not be stored.

Hall and Rust (2000) model the optimal inventory strategy of a durable good intermediary, but do not distinguish between domestic and foreign sourcing. They use the data for the US steel wholesaler to verify how closely the model fits the data. In particular, they calibrate the model to simulate how well it explains observed inventory dynamics.

we examine how changes in trade policy and thus in the price differential between imported and domestic inputs affects inventory investment, as well as the probability of a stockout and domestic sourcing. We also study what an increase in demand uncertainty means for the effect of trade policy and for optimal inventory. Section 5 contains a description of the data, and our empirical analysis is in Section 6. Section 7 concludes. In the Appendix, we collect proofs of our results.

2 Model

Consider a firm facing the inverse linear demand $p_t = a + \epsilon_t - bq_t$ for its product in period t , where ϵ_t is a random shock uniformly distributed on $[-\Delta, \Delta]$, q_t denotes output in period t , and $\Delta < a$ holds such that the random shock is small relative to the size of the market. For each unit of output, the firm needs one unit of a homogeneous intermediate good that it can obtain from two possible sources: a domestic source or a foreign source.⁵

Domestic sourcing is immediate in the sense that intermediates can be ordered and delivered after the demand in that period has been revealed; hence we may think of domestic orders as involving just-in-time delivery. The domestic order in period t is associated with the domestic unit cost w_t and quantity y_t . Foreign sourcing is inflexible, because it takes time for goods to be delivered; a foreign order therefore has to be placed before the realization of demand is known. In particular, an intermediate input purchase made (and paid) in period $t - 1$, at the foreign unit cost v_{t-1} and involving a quantity denoted by m_{t-1} , can only be used in production in period t or later. We interpret the foreign unit cost as including the purchase price as well as tariffs, transport and other transaction costs involved in purchasing the input. In our empirical analysis, the increase in the foreign unit cost corresponds to the tariff increase on steel. Since domestic and foreign intermediates used in a given period are not bought during the same period, we take into account the discount factor, $\delta < 1$, so that the appropriate comparison of the two input costs is between w_t and v_{t-1}/δ . We consider the case where the foreign, inflex-

⁵Note that in order to focus on dual sourcing and, thus, on the substitutability of domestic and imported intermediates, we abstract from possible substitutability between intermediates and other inputs, such as labor. Our model can easily accommodate other inputs, especially if these are perfect complements to intermediates. For example, if producing a unit of output also requires l units of labor, so that the unit labor cost is given by $c = l\omega$, where ω denotes the wage, we can simply define the new demand intercept as $a = A - c$ where A is the original demand intercept.

ible source is cheaper than the domestic, flexible source (i.e., $w_t \geq v_{t-1}/\delta$). The trade-off between the two sources is clear: the foreign source is cheap but forces the firm to commit to it before demand is known, and it can only be used in production next period, whereas the domestic source is expensive but ‘immediate’ in the sense that an order can be placed and received once current demand is known.⁶

Notice that we implicitly assume that the firm can purchase domestic and foreign goods on spot markets at exogenously given prices. According to Hall and Rust (2000), this is indeed the case for the steel wholesaler that we examine in the empirical section. However, the trade-off between sourcing from a cheap but inflexible or an expensive but flexible source does not depend on goods being purchased at exogenously given prices. For instance, we could endogenize w_t and v_{t-1} by assuming that they are set by (or negotiated with) upstream producers with market power. An important complication that this would entail is that we would then have to specify how upstream and downstream firms deal with the horizontal and vertical price and inventory externalities that typically arise in such settings.⁷

In any period, the firm may end up not using the entire quantity of inputs that it purchased. We denote the volume of these unsold units in period $t - 1$ by z_{t-1}^0 and they become part of the available inputs to be used in t . We refer to z_{t-1}^0 as *excess inventory* in period $t - 1$, and assume that storing a unit of excess inventory for one period costs $\gamma \geq 0$. We refer to z_t as *inventory* in period t and define it as the volume of inputs available for use at the beginning of the period.⁸ Thus inventory is equal to $z_t = m_{t-1} + z_{t-1}^0$, that is, the sum of imports purchased in period $t - 1$ and arriving at the beginning of period t , and the excess inventory inherited from period $t - 1$. We assume that z_{t-1}^0 is known when

⁶Our analysis could also accommodate the case where w_t is not known in $t - 1$, but is the expected domestic cost. Then it may happen that the realized w_t is smaller than v_{t-1}/δ .

⁷Pricing externalities include, for instance, double marginalization. Inventory externalities arise as downstream firms may hold more or less inventory than desired by upstream producers. See Krishnan and Winter (2007) for a study of these externalities in a static setting. Qu et al. (2018) examine vertical price and inventory externalities arising under different downstream market structures in an intertemporal inventory model.

⁸Notice that we implicitly assume that the firm only holds inventory of intermediate goods. This makes sense for a firm like our steel wholesaler who carries out very little transformation of inputs into output. A manufacturer, however, would typically also hold inventory of goods in process or of finished goods. If the purpose of holding inventory is to hedge demand uncertainty, it would not matter in which form this inventory is held. Thus the model could be extended to include a production process that allows for different forms of inventory from intermediate to finished goods.

the firm chooses m_{t-1} .⁹

It must be clear from above that, at the beginning of period t , the cost of the foreign inputs ordered in $t-1$ is sunk. The cost of ordering domestic inputs in period t is avoidable. Thus, a firm always prefers using its entire available inventory before buying from the domestic source so that, in any period, it faces three possibilities: (i) it uses less than its inventory z_t and does not order from the domestic source; (ii) it uses its entire inventory but does not order from the domestic source; or (iii) it uses its entire inventory and buys, as well as uses, y_t from the domestic source.

If the firm does not use all of its inventory z_t , what is the value of the excess inventory z_t^0 that the firm takes into period $t+1$? Clearly the firm values a unit of excess inventory at its opportunity cost, that is at what it would cost to import it instead for delivery in $t+1$, v_t , minus any cost of storing it, γ , provided that $v_t \geq \gamma$. Denoting the opportunity cost by ρ_t , we have $\rho_t = \max\{0, v_t - \gamma\}$. If $v_t < \gamma$, the firm does not accumulate excess inventory, but rather discards any unsold units. We assume that $\rho_t < v_{t-1}/\delta \leq w_t$. Thus, while the firm places a positive value on unsold units, this value is not so high that it would voluntarily accumulate excess inventory.¹⁰

Consider the firm's profit from sales in period t , denoted by $\pi_t(q_t)$:

$$\pi_t(q_t) = \begin{cases} (a + \epsilon_t - bq_t)q_t + \rho_t(z_t - q_t) & \text{if } q_t \leq z_t, \\ (a + \epsilon_t - bq_t)q_t - w_t(q_t - z_t) & \text{if } q_t > z_t. \end{cases} \quad (1)$$

The optimal output in t is therefore given by:

$$q_t^*(\epsilon_t) = \begin{cases} \frac{a + \epsilon_t - \rho_t}{2b} & \text{if } \frac{a + \epsilon_t - \rho_t}{2b} \leq z_t, \\ z_t & \text{if } \frac{a + \epsilon_t - w_t}{2b} < z_t < \frac{a + \epsilon_t - \rho_t}{2b}, \\ \frac{a + \epsilon_t - w_t}{2b} & \text{if } \frac{a + \epsilon_t - w_t}{2b} \geq z_t. \end{cases} \quad (2)$$

Eq. (2) shows that demand, and thus optimal output, must be high enough relative to available inventory z_t before the firm buys from the domestic source, simply because the marginal revenue has to exceed the domestic unit cost. If it does not, the firm uses at

⁹This implies that, when a foreign order is placed, the firm knows the realization of the demand of the current period. If orders were placed before z_{t-1}^0 is known, then $m_{t-1} = z_t - E(z_{t-1}^0)$, where $E(z_{t-1}^0)$ is the expected excess inventory.

¹⁰Our model could accommodate the case where $\rho_t > v_{t-1}/\delta$, but this is not the focus of our paper.

most its available inventory depending on the comparison between the marginal revenue of using one unit now and the value of holding on to it, ρ_t . Not surprisingly the lower ρ_t , the greater is the incentive to use this unit today. Hence a lower ρ_t and a higher domestic unit cost w_t imply a wider range of inventory levels over which the firm decides at time t to use up its entire available inventory, i.e., to stock out, without purchasing any domestic inputs.

From (2), we can derive the critical demand realizations where the firm is indifferent between using all inventory or not, $\underline{\epsilon}(z_t)$, and where it is indifferent between buying an additional unit from the domestic source or not, $\bar{\epsilon}(z_t)$:

$$\underline{\epsilon}(z_t) = 2bz_t - a + \rho_t, \quad \bar{\epsilon}(z_t) = 2bz_t - a + w_t, \quad (3)$$

with $\underline{\epsilon}(z_t) < \bar{\epsilon}(z_t)$ from our earlier assumptions. Thus (2) can be rewritten as:

$$q_t^*(\epsilon_t) = \begin{cases} \frac{a+\epsilon_t-\rho_t}{2b} & \text{if } \epsilon_t \leq \underline{\epsilon}(z_t), \\ z_t & \text{if } \underline{\epsilon}(z_t) < \epsilon_t < \bar{\epsilon}(z_t), \\ \frac{a+\epsilon_t-w_t}{2b} & \text{if } \epsilon_t \geq \bar{\epsilon}(z_t). \end{cases} \quad (4)$$

Eq. (4) is useful for three reasons. First, it makes clear that the firm's sourcing strategy depends on demand realizations. Demand can be: (i) low enough so that the firm does not use its entire inventory and therefore accumulates excess inventory, $z_t - q_t^*(\epsilon_t)$, that it may use next period; (ii) in an intermediate range such that it uses up its entire foreign source inventory, $q_t^*(\epsilon_t) = z_t$, but does not order from the domestic source; or (iii) high enough that it uses inputs from both foreign and domestic sources. In the latter case, the purchase of domestic inputs is equal to:

$$y_t^*(\epsilon_t) = \frac{a + \epsilon_t - w_t}{2b} - z_t. \quad (5)$$

Second, (4) makes clear that, in order for a firm to effectively face these three options, the realizations of demand must be feasible given the support $[-\Delta, \Delta]$. In particular, (4) is consistent with $[-\Delta, \Delta]$ provided that $-\Delta < \underline{\epsilon}(z_t) < \bar{\epsilon}(z_t) < \Delta$ which, using (3), requires that $\bar{\epsilon}(z_t) - \underline{\epsilon}(z_t) = w_t - \rho_t < 2\Delta$. Hence, given w_t and ρ_t , (4) requires a relatively high degree of demand uncertainty.

Third, it makes it easy to characterize graphically the possible cases that may arise.

Figure 1, where $[\underline{\epsilon}(z_t), \bar{\epsilon}(z_t)]$ is fully contained in $[-\Delta, \Delta]$, illustrates the possible realizations of demand consistent with (4). If demand uncertainty is low as in Figure 2, however, and thus if $[-\Delta, \Delta]$ is fully contained in $[\underline{\epsilon}(z_t), \bar{\epsilon}(z_t)]$, then the firm's optimal output can only be $q_t^*(\epsilon_t) = z_t$ irrespective of the realization of demand.




Figure 1: High demand uncertainty




Figure 2: Low demand uncertainty

3 Optimal Inventory

We can now proceed to the determination of optimal inventory z_t^* . We formally derive the Bellman equation for this dynamic optimization problem in Appendix A.1; there we also show how the Bellman principle can be used to derive the optimal z_t^* in a more general setting.¹¹ In the following, we focus on the intertemporal optimization condition obtained from the Bellman equation for the specific setting at hand. Notice that the cases illustrated by Figures 1 and 2 are just two possible outcomes among several that we have to examine. Since going through all of these cases is obviously repetitive, we show only one case here and refer the reader to Appendix A.2 for a derivation of the optimal inventory in the other cases. We focus on the case in which the firm's expected marginal revenue in period t is consistent with the optimal output q_t^* as provided by (4) and thus for realizations of demand consistent with Figure 1. The expected marginal revenue in

¹¹In Appendix A.1 we also show that our implicit assumption that the firm orders imports every period is not restrictive. Rather, the model can easily accommodate the case, in which imports are ordered only every n^{th} period, with $n > 1$.

period t , $E[MR_t]$, from a unit of input sourced from abroad in period $t - 1$ is equal to:

$$\delta E[MR_t] = \delta \left[\int_{-\Delta}^{\underline{\epsilon}(z_t)} \rho_t \frac{d\epsilon_t}{2\Delta} + \int_{\underline{\epsilon}(z_t)}^{\bar{\epsilon}(z_t)} (a + \epsilon_t - 2bz_t) \frac{d\epsilon_t}{2\Delta} + \int_{\bar{\epsilon}(z_t)}^{\Delta} w_t \frac{d\epsilon_t}{2\Delta} \right]. \quad (6)$$

The marginal revenue is equal to ρ_t for low demand realizations ($-\Delta \leq \epsilon_t \leq \underline{\epsilon}(z_t)$) and thus when the firm holds on to units for the next period; it is equal to $a + \epsilon_t - 2bz_t$ when the entire inventory z_t is used; and is equal to w_t when the demand realizations are sufficiently high ($\bar{\epsilon}(z_t) \leq \epsilon_t \leq \Delta$) that purchasing from the domestic source is required. Using (3) to evaluate (6), equating the outcome to the foreign price v_{t-1} , and solving for z_t , we obtain as optimal inventory:

$$z_{t,123}^* = \frac{2a - (w_t + \rho_t)}{4b} + \frac{\Delta(w_t + \rho_t)}{2b(w_t - \rho_t)} - \frac{\Delta v_{t-1}/\delta}{b(w_t - \rho_t)}. \quad (7)$$

Here $z_{t,123}^*$ denotes the optimal inventory in regime (123), in which all three ranges of demand realizations are feasible: accumulating excess inventory due to a low demand realization (labeled range 1); stocking out, that is, using up total inventory but without any domestic sourcing (labeled range 2); and using up total inventory but avoiding a stockout by sourcing additional inputs domestically (labeled range 3). In what follows, the subscript will denote the relevant regime, i.e., the feasible range(s).

Other regimes than $z_{t,123}^*$ are possible, all involving a subset of the three ranges of demand realizations. Thus, $z_{t,23}^*$ (valid for $\underline{\epsilon}(z_{t,23}^*) < -\Delta < \bar{\epsilon}(z_{t,23}^*) < \Delta$) refers to the optimal inventory which is always fully used but with some demand realizations requiring domestic sourcing; $z_{t,12}^*$ (valid for $-\Delta < \underline{\epsilon}(z_{t,12}^*) < \Delta < \bar{\epsilon}(z_{t,12}^*)$) is the optimal inventory when domestic sourcing never takes place and where the available inventory might not be entirely used. Finally, $z_{t,2}^*$ (valid for $\underline{\epsilon}(z_{t,2}^*) < -\Delta$ and $\bar{\epsilon}(z_{t,2}^*) > \Delta$) is the optimal inventory when it is always entirely used and no domestic sourcing takes place (illustrated by Figure 2).

Lemma 1 summarizes results derived in Appendix A.2.

Lemma 1. 1. In addition to $z_{t,123}^*$, the feasible inventory levels are:

$$\begin{aligned} z_{t,2}^* &= \frac{a - v_{t-1}/\delta}{2b}; & z_{t,12}^* &= \frac{a + \Delta - \rho_t - 2\sqrt{\Delta(v_{t-1}/\delta - \rho_t)}}{2b}; \\ z_{t,23}^* &= \frac{a - w_t - \Delta + 2\sqrt{\Delta(w_t - v_{t-1}/\delta)}}{2b}. \end{aligned}$$

2. The conditions under which $z_{t,2}^*$ and $z_{t,123}^*$ are obtained are mutually exclusive.

3. Two inventory levels are not feasible: $z_{t,1}^*$ and $z_{t,3}^*$.

Proof. See Appendix A.2. □

Two comments are in order. First, two cases never arise: $z_{t,1}^*$ when inventories always exceed needs (requiring $\underline{\epsilon}(z_{t,1}^*) > \Delta$), and $z_{t,3}^*$ involving systematic domestic sourcing irrespective of demand realizations (requiring $\bar{\epsilon}(z_{t,3}^*) < -\Delta$). The former would imply a permanent excess inventory build-up, which is not an equilibrium strategy in our model: a firm never systematically chooses to order so much from abroad that it would accumulate excess inventory for any demand realization. The latter would imply that a firm always sources at least some inputs at home for any demand realization, even if $v_{t-1}/\delta < w_t$.¹² Second, the conditions under which $z_{t,123}^*$ and $z_{t,2}^*$ hold are mutually exclusive because they depend only on the comparison of the degree of uncertainty (2Δ) with $(w_t - \rho_t)$, with $z_{t,123}^*$ requiring $2\Delta > w_t - \rho_t$, and $z_{t,2}^*$ requiring $2\Delta < w_t - \rho_t$.

4 Comparative Statics

We are now ready to determine how a tariff change, through its effect on the foreign unit cost, affects (i) the firm's optimal inventory, (ii) the probability of stocking out, and (iii) the likelihood of engaging in domestic sourcing. We also want to examine what a change in the level of demand uncertainty implies for the effectiveness of the tariff and for the firm's optimal inventory.

In our dynamic setting, we may distinguish between three trade policy scenarios. The first scenario is a temporary change in tariff protection. This corresponds to a change in

¹²In Appendix A.2, we show that domestic sourcing becomes the primary and only source (i.e. $z_{t,3}^* = 0$), if $v_{t-1}/\delta > w_t$.

v_{t-1} , but leaves v_t (and thus ρ_t) unaffected. The second scenario is a permanent change in the import tariff. This scenario corresponds to a simultaneous change in v_{t-1} and v_t (and thus ρ_t). The third scenario is an anticipated change in the tariff in period t , i.e., a change in v_t (and thus ρ_t).

4.1 Effects of a Tariff on Inventory

Consider first a temporary change in the tariff and thus in the foreign unit cost v_{t-1} , leaving v_t unchanged. Letting τ denote the tariff, we hence assume that $dv_{t-1}/d\tau = 1$ and $dv_t/d\tau = 0$. We observe that the optimal inventory levels in the different regimes summarized in Lemma 1 all decrease monotonically as the tariff rises (i.e., $\partial z_t^*/\partial\tau < 0$). In addition, we have to take into account that, as the tariff increases, the inventory regime switches. At the tariff, respectively foreign-unit-cost levels, at which these regime switches occur, inventory is continuous in the tariff. Therefore, a temporary increase in the tariff leads to a monotonic and continuous decrease in inventory.

A permanent increase in tariff protection has v_{t-1} and v_t both rising to the same extent so that $dv_{t-1}/d\tau = dv_t/d\tau = 1$. Since v_t affects ρ_t , the only changes relative to the case of temporary protection arise in regimes (12) and (123). As we show in the Appendix, a tariff increase reduces inventory in both of these regimes. We hence find:

Proposition 1. *A temporary or permanent increase in the tariff leads to a monotonic and continuous decrease in the optimal inventory z_t^* .*

Proof. See Appendix A.3. □

What happens if tariff protection is anticipated to come into effect in period t ? In this case, $dv_{t-1}/d\tau = 0$ and $dv_t/d\tau = 1$, i.e., v_{t-1} stays constant, but the increase in v_t raises ρ_t because $\partial\rho_t/\partial v_t = 1$ as long as ρ_t is positive to begin with. A change in ρ_t affects only $z_{t,12}^*$ and $z_{t,123}^*$ as all other regimes do not include inventory build-up. We may therefore state:

Proposition 2. *An anticipated increase in the tariff in period t (i) raises the optimal inventory z_t^* when the foreign unit cost is so low that the firm never sources domestically, i.e., in regime (12), and (ii) has an ambiguous or no effect in all other regimes.*

Proof. See Appendix A.4. □

An anticipated increase in the tariff raises the cost of ordering imports in period t . It thereby raises the value of excess inventory that is left over at the end of period $t - 1$ and carried into period t . In regime (12), in which the firm has a positive probability of carrying excess inventory from period $t - 1$ to period t , this induces the firm to build up more inventory in $t - 1$. This effect is counteracted in regime (123) by the fact that optimal hedging may involve dual sourcing, as domestic sourcing in this regime is not that expensive relative to importing. In all other regimes, an anticipated tariff increase has no effect, as the firm does not accumulate any excess inventory.

4.2 Effects of a Tariff on the Probability of a Stockout and Domestic Sourcing

Now consider how a temporary or permanent increase in the tariff affects the probability of a stockout and domestic sourcing. To do this we have to determine how a tariff change affects the switches between regimes. In the case of a temporary tariff increase, notice that $\bar{\epsilon}(z_t) - \underline{\epsilon}(z_t) = w_t - \rho_t$ is independent of v_{t-1} , but $\bar{\epsilon}(z_t)$ and $\underline{\epsilon}(z_t)$ depend on v_{t-1} through z_t . In particular, given (3), we have:

$$\frac{\partial \underline{\epsilon}(z_t)}{\partial v_{t-1}} = \frac{\partial \bar{\epsilon}(z_t)}{\partial v_{t-1}} = 2b \frac{\partial z_t}{\partial v_{t-1}} < 0. \quad (8)$$

This means that, graphically, an increase in the tariff shifts the interval $\bar{\epsilon}(z_t) - \underline{\epsilon}(z_t) = w_t - \rho_t$ from right to left relative to a fixed support of demand shock realizations of length 2Δ centered around zero. This implies two separate inventory regime paths: one conditional on $\bar{\epsilon}(z_t) - \underline{\epsilon}(z_t) = w_t - \rho_t < 2\Delta$ (see Figure 3) and the other conditional on $\bar{\epsilon}(z_t) - \underline{\epsilon}(z_t) = w_t - \rho_t > 2\Delta$ (see Figure 4). In particular, if $2\Delta > \bar{\epsilon}(z_t) - \underline{\epsilon}(z_t) = w_t - \rho_t$, there is a unique path from regime (12) to regime (123) to regime (23) to regime (3) with $z_{t,12}^* \geq z_{t,123}^* \geq z_{t,23}^* \geq z_{t,3}^* = 0$. If $2\Delta < \bar{\epsilon}(z_t) - \underline{\epsilon}(z_t) = w_t - \rho_t$, there is a unique path from regime (12) to regime (2) to regime (23) to regime (3) to with $z_{t,12}^* \geq z_{t,2}^* \geq z_{t,23}^* \geq z_{t,3}^* = 0$.

Consider the first path and assume that the tariff is sufficiently low so that $v_{t-1}/\delta < w_t$. A very low tariff implies that $\bar{\epsilon}(z_t) > \Delta$ which leads to an optimal inventory level $z_{t,12}^*$. In this case, the firm hedges demand uncertainty through inventory only. It stocks out when demand turns out to be high, and accumulates inventory when demand happens to be




Figure 3: Trade protection path with high demand uncertainty

low. Since a tariff increase reduces optimal inventory, the firm is more likely to stock out. That is, the probability of a stockout, $(\Delta - \underline{\epsilon}(z_t))/2\Delta$, is increasing in the tariff, which follows directly from (8).

As the tariff increases, we have $\bar{\epsilon}(z_t) < \Delta$ and $\underline{\epsilon}(z_t) > -\Delta$ so that $z_{t,123}^*$ becomes the optimal inventory level. Now, the firm accumulates inventory if demand is low, stocks out if demand is in an intermediate range, and begins dual sourcing by purchasing domestically when demand turns out to be high. The probability of a stockout is given by $(\bar{\epsilon}(z_t) - \underline{\epsilon}(z_t))/2\Delta = (w_t - \rho_t)/2\Delta$, which is independent of the tariff. But as the firm reacts to a tariff increase by reducing its optimal inventory, it is more likely now to resort to domestic sourcing. Formally, the probability of engaging in domestic sourcing is given by $(\Delta - \bar{\epsilon}(z_t))/2\Delta$, which is increasing in the tariff.

As the tariff rises even more, we obtain $\bar{\epsilon}(z_t) > -\Delta$ and $\underline{\epsilon}(z_t) < -\Delta$ so that $z_{t,23}^*$ becomes the optimal inventory level. Imports are now so expensive that the firm never accumulates excess inventory. If demand turns out to be low, it stocks out. In the case of high demand realizations the firm turns to domestic sourcing, i.e., it turns to dual sourcing to hedge demand uncertainty instead of stocking out. The stockout probability, which is given by $(\bar{\epsilon}(z_t) + \Delta)/2\Delta$, is therefore decreasing in the tariff. The probability of engaging in domestic sourcing is still $(\Delta - \bar{\epsilon}(z_t))/2\Delta$, which is increasing in the tariff.

The second possible path is the same as the first one except that, since $2\Delta < \bar{\epsilon}(z_t) -$

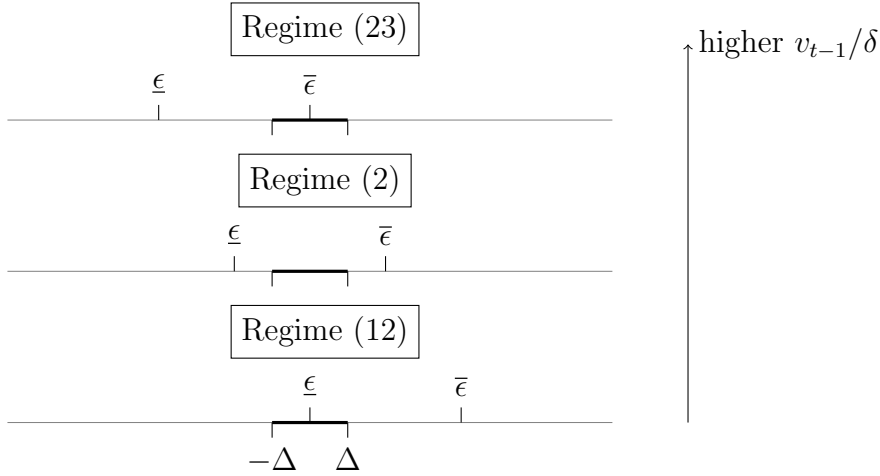


Figure 4: Trade protection path with low demand uncertainty

$\underline{\epsilon}(z_t) = w_t - \rho_t$, the optimal inventory level $z_{t,123}^*$ never arises but is replaced by $z_{t,2}^*$ instead, which is the optimal inventory when the firm prefers to stock out rather than to source domestically or to carry unsold units into the subsequent period. This corresponds to the case where the firm does not hedge demand uncertainty at all, but only ever sells exactly what it has in inventory. In other words, in regime (2) the probability of a stockout is equal to one and the probability of domestic sourcing is equal to zero. Both are hence independent of the tariff.

In the case of a permanent increase in the tariff, the preceding analysis remains unchanged except that $\bar{\epsilon}(z_t) - \underline{\epsilon}(z_t) = w_t - \rho_t$ becomes smaller as v_t (and thus ρ_t) increases. This tends to encourage a regime switch from (2) to (123), which means that as the firm reduces inventory in response to a tariff increase, it is more likely to rely on domestic sourcing when demand turns out to be high rather than to stock out.

We may summarize our results as follows:

Proposition 3. *A temporary or permanent increase in the tariff (i) raises the probability of a stockout if the foreign unit cost is so low that the firm never sources domestically, i.e., in regime (12); and (ii) leaves unchanged or reduces the probability of a stockout and raises the likelihood of domestic sourcing whenever the likelihood of domestic sourcing is positive, i.e., in regimes (123) and (23).*

4.3 Effects of Demand Uncertainty

Demand uncertainty is an essential feature of our model. In this section we therefore examine what a mean-preserving increase in demand uncertainty, equivalent to an increase in Δ , implies for the effect of trade policy, and what it means for the optimal inventory and thus imports.

As we know from the previous section, a temporary or permanent increase in the tariff reduces inventory. Except in regime (2), in which the firm always stocks out and never engages in domestic sourcing, this effect of a tariff on inventory is stronger the greater is Δ and thus the degree of demand uncertainty. Differentiating the comparative static effects in the proof of Proposition 1 with respect to Δ yields:

Proposition 4. *A mean-preserving increase in demand uncertainty magnifies the negative effect of a temporary or permanent increase in the tariff on inventory, except when the firm does not hedge demand uncertainty, i.e. except in regime (2).*

Thus a tariff increase is more effective in the sense of reducing inventory, and consequently imports, by more the greater is the level of demand uncertainty. There are two reasons for this. First, when the firm relies on a large inventory of imported goods to hedge demand uncertainty, a tariff makes this hedging strategy more costly and induces the firm to stock out more frequently the greater is the degree of demand uncertainty. Second, when the firm hedges demand uncertainty through dual sourcing, an increase in the tariff makes the firm more willing to switch to expensive domestic sourcing the higher is the degree of demand uncertainty.

Now consider the direct effect of a mean-preserving increase in demand uncertainty on inventory. In principle, the firm has two ways of hedging this uncertainty, namely building up sufficient inventory, or relying on domestic sourcing to meet high demand realizations. Which option is best for the firm obviously depends on the foreign unit cost. If it is sufficiently low, the firm reacts to greater demand uncertainty by raising its inventory level and accepting a greater expected inventory accumulation. Thus, more uncertainty enhances imports. If the foreign unit cost is sufficiently high, the firm responds to more uncertainty by doing the opposite, namely by reducing the level of inventory and thus imports. This is because the firm finds it optimal to rely more on immediate domestic sourcing. We find:

Proposition 5. *A mean-preserving increase in demand uncertainty raises the optimal inventory, and thus imports, if the foreign unit cost is sufficiently low, and lowers optimal inventory and imports, if the foreign unit cost is sufficiently high.*

Proof. See Appendix A.5. □

This result follows directly from the fact that the firm's optimal hedging strategy is endogenously determined and depends in particular on the foreign unit cost level. It is not very surprising that inventories and imports rise with uncertainty when the foreign unit cost is sufficiently low; after all, the foreign unit cost is lower than the domestic unit cost and units that would be left unsold with low demand realizations are still valuable. The more surprising part is that the firm may prefer responding to an increase in demand uncertainty by reducing inventory and imports even if the foreign unit cost is lower than the domestic cost, as indeed happens when the cost difference is small. In this case, the firm puts a greater premium on immediacy, which implies that it relies more on expensive domestic sourcing.

5 Data

The empirical analysis requires us to distinguish between domestic and foreign purchases preferably at the transaction level and to have inventory at the detailed product level. This is what the data set first used by Hall and Rust (2000) and based on confidential data from an anonymous US steel wholesaler offers. Specifically, the data set collects detailed information on purchases, sales, and stocks of steel products between 1997 and 2006 for one steel wholesaler. Each product has a code (e.g. PL075) identifying the group of products to which it belongs (e.g. 3/4 inch thick plates) and a long description (e.g. 96x240 PLATE 3/4 A-36) identifying more finely a product in that group (its dimension and other characteristics). For each transaction (sales, purchases), we observe the date, unit price, total weight, number of units (when relevant), and total value involved in the transaction. The stock data provides inventory measures at specific dates throughout the sample period. Based on the stock data, the wholesaler deals with a total of 889 unique

product codes and 3082 unique products described by their long descriptions.¹³

One advantage of the data set (used previously by Alessandria et al., 2010a) is that purchase transactions have a code indicating whether they are domestic, foreign or of unknown origin. Although a significant fraction of the transactions is without a known origin, we exploit this difference between domestic- and foreign-sourced purchases. Note however that the exact foreign source is not known.¹⁴ Another advantage of the data (explored by Hall and Rust, 2000) is that we can measure inventory on a daily basis. To benefit from both advantages, we select a sub-set of the products, such that each product satisfies the following criteria: (i) at least 20 purchase transactions are observed at the long description level; (ii) at least 25% of the purchase transactions have an identified origin; (iii) there are at least 20 stock observations, and (iv) the sales are not predominantly identified as adjustment transactions.¹⁵ This produces a sample of 83 products at the long-description level, all belonging to the group of the most heavily transacted products, at least conditional on having a high enough share of identified purchases and without many ‘adjustments’-sales. Because some sub-periods have missing observations for sales and/or purchases, we restrict the sample period to February 10, 2000 - July 8, 2004. This restricted sample period therefore includes a period of about 2 years preceding the Bush safeguard steel tariffs, the entire period during which the Bush tariffs was in place (March 20, 2002 - December 4, 2003), and it includes a period of about 18 months after the Bush tariffs.

In order to generate a daily measure of inventory for each product at the long description

¹³The number of products differs with respect to sales and purchases. The sales data contain 1310 unique codes and 8253 unique long descriptions; the purchase data contain 1211 unique product codes and 8500 unique long descriptions. Many of the long description products have just one transaction, which could be due to misspelling and coding errors. We consider the stock data as being representative of the product dimension of the data.

¹⁴According to Hall and Rust (2000) imports are ordered up to 12 weeks in advance, while some domestic purchases are made with a one or two day notice. Most sales orders are filled within 24 hours of commitment, 95 percent of sales orders filled within 5 days. Back-orders occur only occasionally. Customers expect the firm to have product on hand. Also note that the firm is able to price discriminate across customers, so it appears to have some market power with respect to customers, as assumed in our model. We do not model price discrimination, since we have not information on the identity of customers.

¹⁵In the data, adjustment transactions are sales transactions coming from a ‘transformation’ of a source-product. This criterion eliminates 25 different coil products (product codes HRC and FPC), that is, products bought by coil, but generally not sold as such but as sheets of various lengths cut from a steel coil. We keep these 25 products separate from the others.

level, we compute for each product:

$$Inv_t = Inv_{t-1} + Purchases_{t-1} - Sales_{t-1} - Adj_{t-1}, \quad (9)$$

where Inv_t is the inventory by weight at date t , and $Purchases_{t-1}$, $Sales_{t-1}$ and Adj_{t-1} are respectively purchases, sales, and adjustments (sales coming from transformations) at date $t - 1$ by weight. We thus evaluate inventory at the ‘start of the day’.

We use stock observations in two ways. First, the stock data allow us to anchor the inventory measure (February 10, 2000, the first date of our sample, is one of the dates for which we have stock observations in the data). Second, we use them to check how Inv_t generated by (9) matches with ‘stock’ observations at up to 13 dates within our 2000-04 sub-sample period. For most products we obtain an exact match.¹⁶

An important component of the econometric analysis is the use of the Section-201 tariffs on steel products introduced in 2002. The details can be found in President Bush’s Proclamation 7529 published on March 5, 2002 (‘To Facilitate Positive Adjustment to Competition From Imports of Certain Steel Products’). A total of 272 10-digit HTS steel products were subject to the new tariff, effective on March 20, 2002: 185 steel products received a 30% tariff, 60 received a 15% tariff, 15 received a 13% tariff, and 7 received an 8% tariff.¹⁷ Although designed to be in place for three years, with gradually lower tariff rates after each 12-month period, the tariffs were suddenly cancelled with an effective date of December 5, 2003, shortly after the US lost the case at the WTO. It is useful to point out that, in addition to country-exclusions included in the Proclamation (countries benefiting from a preferential trade agreement with the US such as Canada and Mexico, and developing countries (collectively representing less than 9% of the US steel import market)), several rounds of exclusion took place during the first year of these tariffs. These are firm- and product-specific exclusions regarding very specialized products that were in too short supply in the United States. These exclusions were not extended to other firms in the same country, let alone to other countries (Bown, 2013). The products that we use are standard steel products and are very unlikely to be subject to these firm-specific exclusions. It is also worth pointing out that a number of antidumping and countervailing

¹⁶When the match is not exact, we are generally missing only a few units. These cases typically arise for the most frequently sold products. The discrepancies may come from mis-labelled sales transactions.

¹⁷See also Bown (2013), Hufbauer and Goodrich (2003), and Read (2005).

orders affecting the US steel industry were in place during the same period.

In order to link these tariffs with the steel products, we assign an HTS code (version 2002) to each steel product. It turns out that all 83 products selected above are subject to the same initial safeguard tariff of 30%. This is not surprising because our selection mechanism captures the most transacted products, mostly steel plates and sheets. To establish a control group of steel products subject to lower or no initial safeguard tariffs, we concentrate on steel pipes and tubes. These products have either no safeguard tariffs or, when they do, the initial rate is 15%. To include them in our sample we have to relax the four inclusion criteria mentioned above. In particular, we consider all pipes and tubes products having at least 5 stock measures during the sample period as well as some sales and identified purchases. This results in 21 pipes and tubes products that we may use as a control group.

The total number of products considered in the empirical analysis is thus 104. Based on the 2000-04 sample period, these products have the following characteristics regarding the purchase transactions. First, 67 products of the 104 products (64.4%) exhibit some dual sourcing. Among the 37 products that do not, 17 are only bought domestically, while 20 are only sourced abroad. Thus for a majority of products, the wholesaler engages in dual sourcing and considers domestic and imported variants as perfect substitutes. Second, concentrating on the 67 products with dual sourcing, Table 1 summarizes information (averages and standard deviations) about the relative frequency of foreign transactions, the domestic price premium (defined at the product level as $(p_i^d - p_i^f) / p_i^f$, where p_i^d and p_i^f are the domestic, respectively foreign prices of product i), and the weight (in pounds) of (identified) foreign and domestic purchase transactions. Despite domestic and foreign variants of a product being considered as perfect substitutes, domestic prices are on average 10% higher than foreign ones. This is consistent with the fact that domestic products may be more immediately available than foreign ones. Foreign purchase transactions, representing on average 43.7% of the identified purchase transactions, are on average 2.9 times bigger than domestic transactions. Obviously there is a lot of variation, including about the importance of foreign purchase transactions across products, as the standard deviations are generally high.

	$\frac{\#Foreign}{\#Total\ Identified}$	Domestic Price Premium	Transaction Weight Domestic	Transaction Weight Foreign
Average	0.437	0.101	35,738	100,550
S.D.	0.312	0.110	33,937	108,364

Table 1: Purchase Characteristics across Products and Time

6 Empirical Analysis

We now turn to the empirical tests of our model. We focus on testing those propositions that are based on changes in the import tariff, because identification follows from the exogenous imposition of the steel tariffs and because they are directly policy relevant. We do not explore the effects predicted in Proposition 5. Identification of these effects would require an exogenous change in demand uncertainty that, moreover, is not affected by the tariff. But the sample period is too short to compute a proxy for demand uncertainty that has a sufficient amount of variation over time; this is even more true if we consider only the pre-tariff period.

We are also skeptical about being able to identify anticipation effects of a tariff, as predicted by Proposition 2. In particular, the announcement period prior to imposition of steel tariffs was too short—only two weeks from March 5 to 20, 2002—to have any effect on inventory; and the premature removal of the tariff on December 5, 2003 probably came as a surprise to the firm, given that the tariff was initially announced to last for three years. The reduction in the tariff rate following the first year of the tariff was anticipated, as it was announced when the tariff was initially introduced. While it would be difficult to separate any anticipation effects in year one of the tariff from the direct effect of the tariff, below we nevertheless consider effects separately for period 1 of the tariff (March 20, 2002 to March 19, 2003) and period 2 (March 20, 2003 - December 4, 2003).

We hence seek to test the results of Propositions 1, 3 and 4. We proceed by formulating a testable hypothesis for Proposition 1, presenting the empirical approach, estimation results and robustness checks, and then quantifying the effects before moving on to the next hypothesis.

6.1 The effect of a tariff on inventory

According to Proposition 1, a temporary or permanent tariff leads to a reduction in inventory for all products that, at any time t , have a positive probability of being imported, i.e., that fall into regimes (12), (2), (123), or (23). Observing products over time, those in regimes (12) and (2) are products that, prior to the imposition of the tariff, are only imported. We refer to them as having an initial import share (IIS) equal to one. Products in regimes (123) and (23) exhibit dual sourcing and therefore have an IIS strictly between zero and one, whereas products in regime (3) are only sourced domestically and hence have an IIS equal to zero.

After the imposition of the tariff, products may switch to regimes that involve less foreign sourcing, which further boosts the impact of the tariff on inventory as the firm either accepts a greater probability of stockouts or shifts toward domestic sourcing. The most basic hypothesis we can test is therefore:

Hypothesis 1 *Consider products with a positive initial import share. Products subject to the tariff experience a decline in inventory relative to products unaffected by the tariff.*

We test this hypothesis using a difference-in-difference (DiD) approach that compares the change in inventory (tariff-period inventory minus pre-tariff-period inventory) across certain sets of products. We conduct two complementary DiDs. In both, the treatment group consists of a set of products that are known to be affected by the Bush steel tariff. Where the two DiDs differ is in the control group. As we already explained in the preceding section, creating a control group from the pipes and tubes products, which were not or less affected by the tariff, meant applying softer criteria for sample selection than were applied to the treatment group.

In the first analysis (DiD-IIS), we therefore do not use tubes and pipes as the control group but rather a subset of products that are subject to the tariff, but that have an initial import share so close to zero that we would expect only minimal effects from the tariff. We form multiple treatments based on the value of IIS, which is computed for each product as the share of weight bought abroad out of total weight bought, in the pre-tariff period. Then, we stratify IIS into $J + 1$ categories $\{0, \dots, J\}$ and denote by IIS_{ij} the binary variable that equals 1 if product i is in category j . Category 0 is the control group.

IIS category	Y_{it}	
	mean	s.d.
$[0, 0.333]$	2.41	1.50
$(0.333, 0.656]$	2.52	1.69
$(0.656, 0.889]$	2.12	2.34
$(0.889, 1]$	3.59	4.16

Table 2: Inventory by IIS - Summary statistics

In our baseline specification, $J = 3$ and the categories are defined through the intervals determined in a way that ensures an equal number of products in each category. The corresponding intervals for the value of IIS are: $[0, 0.333)$, $[0.333, 0.656)$, $[0.656, 0.889)$, $[0.889, 1]$. If $IIS_{i0} = 1$, the product is part of the control group. If $IIS_{i2} = 1$, then product i 's initial import share is between 0.656 and 0.889.

As the dependent variable in these regressions, Y_{it} , we consider the ratio of inventory for product i on day t to the average monthly sales of that product. In other words, it is a measure of inventory expressed as the number of months of sales it could fulfill. This provides a natural way to convert the inventory measurements to a common scale, facilitating our interpretation of the regression coefficients below. Table 2 provides some summary statistics about the dependent variable, separated by IIS category.

The regression equation corresponding to DiD-IIS, with $T = 0$ denoting the pre-tariff period, and $T = 1, 2$ the tariff periods is:¹⁸

$$Y_{it} = \phi_i + \lambda_m + \beta_T + \sum_{j=1}^J \sum_{s=1}^2 \gamma_{j,s} \times IIS_{ij} \times 1_{T=s} + u_{it}, \quad (10)$$

where ϕ_i is a product fixed effect, λ_m is a month fixed effect to control for seasonality, β_T is a treatment period fixed effect, and $\gamma_{j,T}$ are the parameters of interest for each one of the three treatment groups.

In DiD-IIS, we use the $[0, 0.333)$ -IIS group as a control group, meaning that we implicitly assume that this group is unaffected by the tariff. Under standard assumptions for DiD (including the parallel paths assumption), $\gamma_{j,T}$ is the effect of the tariff on products in

¹⁸Period 0 is the baseline, pre-tariff period. We use two years of daily data leading up to March 20, 2002. Period 1 is the first year during which the tariff was in place, i.e. March 20, 2002 - March 19, 2003. Period 2 is the second period during which the tariff was in place, i.e. March 20, 2003 - December 4, 2003.

category j in period T . Hypothesis 1 says: $\gamma_{j,T} < 0$ for all j and $T = 1, 2$.

Our second DiD analysis uses steel pipes and tubes as control group. However, a standard DiD based on this control group is suspect because the parallel paths assumption is unlikely to hold in this case. The reason is as follows. We suspect that the tariff, being a safeguard measure, was imposed because steel imports were rising. But with steel imports going up, inventory for the treated products would have increased relative to that of the control group products even in the absence of a tariff, which clearly violates the assumption that the two groups show parallel trends when there is no treatment.

For this reason, our second DiD exercise must control for the IIS and make a comparison across time, across product groups, and then across values of the IIS. This is a difference-in-difference-in-difference (DDD) analysis. Letting $D_i = 1$ if the product is treated, and $D_i = 0$ otherwise, the regression equation corresponding to this DDD analysis is:

$$\begin{aligned}
Y_{it} = & \phi_i + \lambda_m + \beta_T + \sum_{j=1}^J \sum_{s=1}^2 \delta_{j,s} \times D_i \times IIS_{ij} \times 1_{T=s} \\
& + \sum_{s=1}^2 \kappa_s \times D_i \times 1_{T=s} + \sum_{j=1}^J \sum_{s=1}^2 \Gamma_{j,s} \times IIS_{ij} \times 1_{T=s} + u_{it},
\end{aligned} \tag{11}$$

where the last two terms (other than u_{it}) are part of a correctly specified triple-diff regression equation, and the $\delta_{j,T}$ are the parameters of interest. The DDD specification is a different approach to estimating the effect of the tariff. Just like the DiD-IIS, it uses the low-IIS group to control for time trends that affect all products in the treated group. It additionally uses a control group of pipes and tubes, not affected by the tariff, to control for time trends specific to IIS category. It requires that the inventory of pipes and tubes are unaffected by the tariff. Hypothesis 1 says: $\delta_{j,T} < 0$ for all j and $T = 1, 2$.

Table 3 presents the results where, for each case, the first number is the estimate and the second one is the standard error. Column (a) presents our baseline findings. Our findings are in line with Hypothesis 1. DiD-IIS finds negative and statistically significant treatment effects for each positive level of the IIS and both treatment periods. The effects are quantitatively large, judging by the means reported in Table 2. Consider, for example, the effect for IIS category 2 in phase 1 of the tariff, that is in the year immediately following the imposition of the tariff. The coefficient of $\hat{\gamma}_{2,1} = -0.92$ implies a reduction of inventory in the amount of 0.92 months of sales, when the average inventory for a

	Baseline		Era3		IIS1		Inventory		One year		3 cat	
	(a)		(b)		(c)		(d)		(e)		(f)	
$\gamma_{1,1}$	-0.32	0.05	-0.44	0.04	-0.30	0.05	74678	2815	0.29	0.05	-0.85	0.04
$\gamma_{2,1}$	-0.92	0.06	-1.25	0.05	-0.85	0.05	-11569	3297	-0.94	0.05	-2.54	0.05
$\gamma_{3,1}$	-3.22	0.06	-2.99	0.05	-2.51	0.05	-54832	3360	-1.94	0.06		
$\gamma_{1,2}$	-0.60	0.05			-0.25	0.06	-14511	3213	-0.01	0.05	-0.90	0.05
$\gamma_{2,2}$	-1.72	0.06			-1.34	0.06	-43093	3741	-1.77	0.06	-1.98	0.06
$\gamma_{3,2}$	-2.67	0.06			-1.50	0.06	-33708	3815	-1.42	0.06		
$\delta_{1,1}$	-0.04	0.16	-0.46	0.14	-0.47	0.10	82149	8954	0.92	0.16	-0.57	0.16
$\delta_{2,1}$	-1.44	0.17	-1.43	0.14	0.43	0.13	-14913	9087	-1.38	0.16	-1.82	0.20
$\delta_{3,1}$	-2.31	0.13	-2.01	0.11	-3.84	0.19	-42206	6811	-0.72	0.12		
$\delta_{1,2}$	-1.14	0.20			0.30	0.12	-12885	10820	0.00	0.19	-1.43	0.20
$\delta_{2,2}$	-1.51	0.20			1.94	0.15	-39350	10958	-1.28	0.19	-1.07	0.16
$\delta_{3,2}$	-1.66	0.16			-3.26	0.22	-18322	8628	0.06	0.15		

Table 3: Hypothesis 1 - Results

product in this category is 2.12 months of sales. In other words, a 30% tariff on average reduces inventory of products in that IIS category by around 43%. Unsurprisingly, the impact of the tariff is even stronger in IIS category 3, which collects products with an IIS greater than 0.889. Here a coefficient of $\hat{\gamma}_{3,1} = -3.22$ means that a 30% tariff leads to an average reduction in inventory of close to 90% in the first year, calculated as a decrease in inventory by 3.22 months of sales from a mean of 3.59 months. Effects in period 2 of the tariff are of comparable magnitude.

The DDD specification finds effects with the same sign and magnitude, although the category 1-phase 1 effect is quantitatively small and not statistically significant.

Several sensitivity checks corroborate our main findings. In column (b), we consider only one treatment period, merging the two phases. In column (c), we use a prediction algorithm to impute the origin of purchases where they are missing, and recompute IIS on this basis.¹⁹ In column (d), we use total inventory (in pounds) as outcome variable, rather than in months of sales. In column (e), we only use one year of inventory data before the tariff. In column (f), we discretize IIS into three categories instead of four. With few exceptions, the qualitative findings of our baseline specification are in line with these robustness tests.

¹⁹We use an off-the-shelf machine-learning algorithm (H20’s autoML) that achieves a training set AUC (area under the curve) of 0.98 and a test set AUC of 0.95, suggesting the algorithm achieves a high degree of accuracy.

IIS category	stockout	
	mean	s.d.
[0, 0.333]	0.091	0.151
(0.333, 0.656]	0.066	0.139
(0.656, 0.889]	0.147	0.150
(0.889, 1]	0.136	0.210

Table 4: Stockout by IIS - Summary statistics

6.2 The effect of a tariff on the stockout probability

We can seek further empirical evidence for the hedging strategies explored in the model and thus the economic mechanism behind Hypothesis 1 by examining how the tariff affects the probability that a product stocks out. According to Proposition 3, products that are only imported are more likely to experience a stockout when hit by a tariff, whereas products with a small initial import share are predicted to experience no change or even a decrease in the stockout probability, as the firm switches to domestic sourcing to avoid stockouts. The treatment effect may hence be non-monotonic depending on the product's initial import share. Thus we may state:

Hypothesis 2 *Consider products with a positive initial import share. The imposition of a tariff raises the stockout probability of products with an initial import share sufficiently close to one, but lowers the stockout probability of products with a low initial import share.*

To investigate this, we follow an empirical approach identical to that used for Hypothesis 1 with one exception: the dependent variable Y_{it} is now defined as a binary variable that equals 1 if the inventory of product i on day t is non-positive. We take a small tolerance of 10 weight units (pounds), so that stockout is defined as 1 if $inventory \leq 10$, 0 if $inventory > 10$. This allows for a bit of rounding or timing error in the inventory reconstruction.

Table 4 provides summary statistics by IIS category. Notice that we observe a considerable number of stockouts in the data. In IIS category 2, for example, the probability that a product is stocked out on any given day is 14.7%, and thus much higher than in the IIS category 1, where the probability of a stockout is only 6.6% on average.

According to Hypothesis 2, we may observe a non-monotonic effect of the tariff on the stockout probability, with the highest IIS category potentially exhibiting a positive

coefficient, and low IIS categories exhibiting a negative coefficient. Our baseline DiD-IIS findings (see Table 5) are in line with this prediction. In particular, we estimate that the effects of the tariff are $\hat{\gamma}_1 = -0.038(0.005)$, $\hat{\gamma}_2 = -0.017(0.006)$, $\hat{\gamma}_3 = 0.004(0.006)$ for the three categories. For the low IIS categories we find a negative coefficient as expected, indicating that the firm switches to hedging demand uncertainty by relying more on domestic sourcing. Note the positive but statistically insignificant effect for the top category. It is hard to identify the positive effect governing products with an IIS close to 1 for at least two reasons. First, the variation in the stockout variable is considerably less than in the inventory variable we used for Hypothesis 1. Second, category three is likely a mixture of products experiencing a negative effect (further away from $IIS = 1$) and products experiencing a positive effect (close to $IIS = 1$).

Since the coefficients reported in Table 5 are percentage points, they can be directly compared to the means in Table 4. We find substantial effects for products in IIS category 1. The tariff reduces the average stockout probability of products in this category by more than half from a mean value of 6.6% to 2.8%. The effect for category 2 is smaller, but the stockout probability still decreases from 14.7% to 13% on average. Even ignoring the statistically insignificant increase in stockout probability category 3, products in this category have the highest average stockout probability after imposition of the tariff, which is expected from the model.

We also find some support for Hypothesis 2 in the DDD specification. We find a much larger effect for category 1, and find a negative effect for the highest category.

We conduct the same robustness checks as for Hypothesis 1 in Table 3 except for column (d) which is meaningless in the current setting. The robustness checks largely corroborate our baseline findings for Hypothesis 2, with only a few cells yielding conflicting signs.

The evidence on how the tariff affects the stockout probability suggests that the firm is switching to domestic sourcing for most products rather than incurring a greater stockout probability. This is indeed what the model predicts for products that have domestic substitutes that are not too expensive, which are exactly the products for which the IIS should be small to start with. We can now check if the firm indeed switches to domestic sourcing to hedge demand uncertainty, which should reduce the foreign sourcing share.

	Baseline		Era3		IIS1		One year		3 cat	
	(a)		(b)		(c)		(d)		(e)	
$\gamma_{1,1}$	-0.038	0.005	-0.042	0.004	-0.036	0.005	-0.042	0.006	-0.001	0.005
$\gamma_{2,1}$	-0.017	0.006	-0.022	0.005	-0.049	0.005	-0.059	0.007	0.016	0.005
$\gamma_{3,1}$	0.004	0.006	-0.026	0.005	0.049	0.006	-0.029	0.007		
$\gamma_{1,2}$	-0.048	0.006			-0.093	0.006	-0.050	0.006	0.025	0.005
$\gamma_{2,2}$	-0.029	0.007			-0.038	0.006	-0.070	0.007	-0.018	0.006
$\gamma_{3,2}$	-0.068	0.007			-0.102	0.006	-0.100	0.008		
$\delta_{1,1}$	-0.258	0.017	-0.168	0.015	-0.059	0.011	-0.150	0.020	-0.221	0.017
$\delta_{2,1}$	-0.016	0.018	-0.032	0.015	-0.297	0.014	-0.082	0.021	0.018	0.021
$\delta_{3,1}$	-0.062	0.013	-0.065	0.012	-0.247	0.020	-0.108	0.015		
$\delta_{1,2}$	-0.046	0.021			-0.056	0.012	0.059	0.024	-0.033	0.013
$\delta_{2,2}$	-0.057	0.021			-0.234	0.016	-0.127	0.024	-0.024	0.016
$\delta_{3,2}$	-0.069	0.017			0.151	0.023	-0.112	0.019		

Table 5: Hypothesis 2 - Results

6.3 The effect of a tariff on domestic versus foreign sourcing

Whereas we have so far relied on inventory data to examine the economic mechanisms in our model, we now switch to using purchasing transactions data. Associated with the reduction in inventory following the imposition of a tariff is a change in the firm's sourcing strategy. In particular, Proposition 3 predicts that the firm will reduce foreign relative to domestic sourcing and thus experience a decline in its import share. We may hence test the following hypothesis:

Hypothesis 3 *Consider products with an initial import share strictly between zero and one. The imposition of a tariff leads to a reduction in the import share.*

Recall that we constructed the IIS variable as the share of total weight bought abroad to total weight bought in the pre-tariff period. To test Hypothesis 3, we construct the analogous variable for the time period during which the tariff was active, and call it the tariff import share (TIS). We also use the purchase predictions from the machine learning algorithm to compute the shares.²⁰ As before, we cut the IIS variable into four categories. The intervals for this version with imputations are $[0, 0.303]$, $(0.303, 0.476]$, $(0.476, 0.736]$, $(0.736, 1]$.

²⁰The results are qualitatively similar when we use only the purchase transactions with known origin.

The analysis for this hypothesis deviates from that in the previous sections, owing to the nature of the outcome variable. First, in the previous sections, we worked with daily inventory data. In this section, the outcome variable is constructed from individual purchase data. Because the frequency of purchase transactions is low for some products, this outcome variable is best computed as an average over significant time periods. As a result, we have only one or two measurements of the outcome variable during the tariff period, which a priori complicates the identification of effects. Second, we consider one tariff period rather than splitting the tariff period in two. In this dimension, the analysis in this section is therefore like the analysis in the sensitivity analysis “era3” in the previous sections. Third, we use a prediction algorithm to impute the origin of purchases where they are missing which is especially the case during the post-tariff period. Fourth, we do not exclude coil products that were excluded from the analysis in previous sections. Those products were excluded for reasons related to the sales data. In this section, we use only the purchases data.

Figure 5 presents the main result. Each point in the graph corresponds to a product. The lines are nonparametric regression lines for each product group: treated and untreated (the pipes and tubes control group). A significant fraction of products are close to one of two negatively sloped 45-degree lines. Those products changed their import share from their IIS to zero (full domestic sourcing after the imposition of the tariff, lower 45 degree line) or from their IIS to 1 (full foreign sourcing after the imposition of the tariff, upper 45 degree line). If a product decreased its share of foreign sourcing and thus its TIS, it appears below the horizontal line at 0.

Comparing the two regression lines, it is clear that there is a decrease in the TIS of the treatment group (the lower line) relative to the control group. Furthermore, the shape of the regression line for the treated products suggests that there is a decrease in the TIS at all but the lowest levels of the IIS.

A linear regression confirms and quantifies these results. We run the following regression, which corresponds to the DDD specification, but presented in a way that is consistent with Figure 5:

$$TIS_i - IIS_i = \beta_0 + \beta_1 D_i + \sum_{j=1}^J \gamma_j \times IIS_{ij} + \sum_{j=1}^J \delta_j \times D_i \times IIS_{ij} + u_i. \quad (12)$$




Figure 5: Scatterplot for Hypothesis 3.

We also run the DiD-IIS, using regression equation

$$TIS_i - IIS_i = \beta_0 + \sum_{j=1}^J \delta_j \times IIS_{ij} + u_i \quad (13)$$

without the pipes and tubes, using the low-IIS category as a control group.

Table 6 presents the results. In the DDD specification, β_1 captures the effect of a tariff on the products subjected to a tariff, conditional on any value of IIS. We estimate that the import share decreases by 0.43. This is substantial since the average IIS in the treated group is 0.49, implying that imports are simply wiped out for many products. The estimates for the treatment effects come with large standard errors, which is related to the small number of control units, a noisy estimate for the IIS for those products, and the small number of time periods (recall that we cannot use daily inventory fluctuations as in Hypotheses 1 and 2).

The DiD-IIS estimates are more precise, because the treatment effects are estimated without invoking the pipes and tubes control group. We find large effects of the tariff for

Import share	All (a)	Treated only (b)
β_1	-0.4264 (0.1590)	
γ_1	-0.0161 (0.1788)	
γ_2	-0.1271 (0.2385)	
γ_3	-0.3291 (0.3577)	
δ_1	-0.3372 (0.2025)	-0.3958 (0.0914)
δ_2	-0.1771 (0.2542)	-0.2212 (0.0923)
δ_3	0.0725 (0.3679)	-0.2760 (0.0914)

Table 6: Hypothesis 3 - Results

all IIS categories. For example, products in IIS category 3 (with an IIS between 0.736 and 1) experience an average decrease of 0.276 due to the tariff, implying a decrease in the import share of more than 27.6% on average. The import share falls by even more for the lower IIS categories.

6.4 The effect of demand uncertainty on the effectiveness of protection

Our model predicts that an increase in demand uncertainty magnifies the effect of the tariff (Proposition 4). This can be examined empirically by checking whether products exhibiting different degrees of demand uncertainty, *ceteris paribus*, exhibit different declines in inventory after the tariff is imposed. Specifically, we may state:

Hypothesis 4 *The imposition of a tariff leads to a larger decrease in inventory the greater is the product's level of demand uncertainty.*

We start by constructing a measure of demand uncertainty for each product in three steps. First, using the entire sample period, we aggregate the sales transactions to monthly data. Second, we fit a time series model with only year and month dummies. Third, we compute the standard error of regression for this variable, and divide it by the average monthly sales for this product. Next, we discretize the demand uncertainty measure by dividing products into four groups of equal size. The resulting four categories are defined through the intervals $[0,0.557]$, $(0.557,0.687]$, $(0.687,0.806]$, $(0.806,0.992]$, $(0.992,1.4]$. Then we define the dummy variables $DU_{ij} = 1$ if the demand uncertainty variable for product

i is in category $j = 0, 1, 2, 3$; $DU_{ij} = 0$ otherwise. Note that this is analogous to the construction of the discretized initial import share variable used in Hypotheses 1 and 2.

Replacing the IIS_{ij} with the DU_{ij} in (10), our regression equation is:

$$Y_{it} = \phi_i + \lambda_m + \beta_T + \sum_{j=1}^J \sum_{s=1}^2 \gamma_{j,s} \times DU_{ij} \times 1_{T=s} + u_{it}, \quad (14)$$

where Y_{it} is inventory, and the other terms are as in (10). For Hypothesis 4, we use inventory rather than inventory normalized by average monthly sales. Normalized inventory would have the same denominator as the main explanatory variable, and we are concerned that regression coefficients may reflect this mechanical relationship between LHS and RHS variables. Furthermore, the product fixed effect will pick up the product-specific scale.

We are reluctant to interpret the $\gamma_{j,T}$ as causal effects of the tariff. First, we have argued previously that the parallel paths assumptions is suspect without controlling for IIS. In the current exercise, we assume that the IIS of product i is controlled for by the product fixed effects, and we only consider how the impact of the tariff differs across demand uncertainty categories. This may not be sufficient to control for potential heterogeneity of the tariff effect due to IIS. Second, there is no natural reference group for DU. For example, the low-DU group can be expected to be affected by the tariff as well. In the case of Hypothesis 1, we regarded the IIS category zero as the control group since, given its very small initial import share, it was essentially unaffected by the tariff. For these reasons, we view the $\gamma_{j,T}$ as descriptive. They tell us the change in inventory before/after the tariff, split out by values of demand uncertainty. Hypothesis 4 suggests that the change in inventory should increase with the level of demand uncertainty, i.e. $0 > \gamma_{1,T} > \gamma_{2,T} > \gamma_{3,T}$.

The results are presented in Table 7, with our preferred specification in column (a)-Baseline. The DiD-DU panel for that column says that, following the imposition of the tariff, the inventory for demand uncertainty category 1 is lower by 17205 pounds on average compared to inventory in category 0. The effect is monotonic in DU, as predicted by the model, as inventories in categories 2 and 3 decrease by 31091 and 57903 pounds, respectively compared to inventory in category 0. All effects are statistically significant at all conventional significance levels and none of the sensitivity checks overturn this result.

All in all, using the exogenous tariff shock for identification, we find considerable empirical evidence to support the predictions of our theoretical model. In particular, we find

Inventory		Baseline	3 cat	5 cat	era4
		(a)	(b)	(c)	(d)
DiD-DU	γ_1	-17205 (2679)	-11588 (2403)	-69888 (2823)	-60374.4 (3135.7)
	γ_2	-31091 (2760)	-45974 (2293)	-36539 (3069)	-76453.5 (3222.7)
	γ_3	-57903 (2630)	-	-68675 (3228)	-98087.4 (3084.9)
	γ_4	-	-	-90443 (2852)	-
<i>period 2</i>	γ_1	-	-	-	44174.1 (3552.8)
	γ_2	-	-	-	34080.3 (3697.7)
	γ_3	-	-	-	-816.4 (3488.9)

Table 7: Hypothesis 4 - Results

that a trade policy shock has statistically significant and economically sizeable effects on inventory investment and hence also on the probability of stockouts and the domestic sourcing share. Moreover, as predicted by the model, demand uncertainty strengthens the protective effect of the tariff in the sense that products featuring greater demand uncertainty show a greater decline in inventory when hit by the tariff.

7 Conclusions

Optimizing inventory investment and sourcing strategies in the face of uncertain demand is a problem of considerable importance for many importers. How importers deal with this problem, in turn, plays a crucial role for the volume and the dynamics of international trade. For instance, assuming that importers hold substantially more inventory than firms purchasing intermediates domestically, Alessandria et al. (2010b) and Novy and Taylor (2019) argue that inventory investment has played an essential role in generating the great trade collapse accompanying the global financial crisis in 2008/09. Inventory adjustment following the negative demand shock represented by the financial crisis, they argue, is responsible for the strong decline in international trade relative to production and its subsequent strong recovery. In particular, firms drew down their inventories of imported goods following the negative demand shock, and only re-order imports once inventories have become sufficiently low.

Despite its obvious importance, micro-level evidence on the inventory investment of importers is extremely rare. The current paper intends to fill that gap. In order to address this issue in a precise way, we examine two alternative strategies for hedging demand

uncertainty, namely (i) building up inventory of imported goods, and (ii) dual sourcing, i.e., resorting to more expensive, but immediately available domestic substitutes to cover demand surges.

Our theoretical model predicts that the firm chooses dual sourcing for products for which the price differential between domestic and foreign sourcing is small. This strategy is chosen because, even if the domestic price is higher than the foreign one, the immediacy of domestic sourcing makes this strategy particularly valuable to adjust output in response to high demand. The impact of a tariff then is for the firm to reduce inventory and increase its domestic sourcing share.

If the price differential is big, the firm hedges domestic demand uncertainty by keeping a large inventory of imported goods to avoid stockouts in the case of a high demand, even if it implies accumulating excess inventory should demand turn out to be low. If the firm is exposed to an import tariff, it reduces inventory and accepts a larger probability of stockouts, rather than resort to expensive domestic sourcing.

When the price differential is in an intermediate range, the firm may choose not to hedge demand uncertainty at all. In essence, the immediate domestic sourcing is too expensive to be used when the demand is high and, even if foreign sourcing is the primary source, buying a large volume to satisfy a demand surge is not worth it either as it might end up as unsold inventory with too low a future value when the demand is low. In this case, the firm is content with buying a limited quantity making sure that it can use it entirely irrespective of the demand realization. Ordering less even if it means stocking out can thus also be an optimal strategy and even more so, if the firm has to pay an import tariff.

We explore the functioning of these hedging mechanisms using high-frequency data for a U.S. steel wholesaler and find considerable support for our model's predictions. Our empirical analysis also shows that the magnitudes of these effects are not trivial. In particular, products that are hit by a 30% tariff and have a high initial import share on average show a decrease in inventory of around 90% from 3.59 months worth of sales to just 0.37 months worth of sales in the first year following the imposition of the steel tariffs compared to products in a control group. These effects also show up in the remaining tariff period, and they are sizeable even for products with a relatively low initial import share.

As products are increasingly sourced domestically when they are exposed to the tariff, the import share falls substantially, for instance by over 27.6% on average for products

with a high initial import share, and by even more in the case of products with a lower initial import share. The stockout probability is also reduced as a result of the tariff especially for the products with relatively low initial import shares. This is less the case for products with a high initial import share and, with the tariff, these products end up with the highest average stockout probability at around 14%.

According to the model, products with higher demand uncertainty should be hit harder by the imposition of a tariff than products with smaller demand uncertainty. Our empirical analysis suggests that it is indeed the case.

Since our empirical analysis is for only one importing firm—a wholesaler—, an obvious question is to what extent our findings generalize. We expect this specific steel wholesaler to be fairly representative of durable goods wholesalers in general. According to Hall and Rust (2000), not only are there about 5000 similar steel wholesalers (or steel service centers as they are called in the industry) in the United States handling 23% of steel consumed in the United States in 1998, but the inventory and supply behavior of our steel wholesaler is consistent with observations at the U.S. industry level, specifically sector SIC 505 (wholesale trade: metal and minerals, except petroleum).

We would also argue that examining a wholesaler instead of, say, a manufacturer is not problematic when it comes to identifying mechanisms to hedge demand uncertainty and evaluating how they are impacted by trade policy. The reason is that hedging demand uncertainty through inventory and dual sourcing is not restricted to the wholesale sector, but is considered by the operations management literature to be a fundamental problem facing firms in manufacturing and retailing, too. For concreteness, Allon and Van Mieghem (2010) discuss the example of a U.S. manufacturer of wireless transmission components facing the problem of how much to source from assembly plants in China and Mexico, with China being the cheap but inflexible and Mexico the expensive but flexible source.

Finally the fact that source-countries are not known and that countries like Canada and Mexico have been exempted from the safeguard tariff may be interpreted as an issue for the validity of our econometric results. Total annual US imports of iron and steel products show that Canada and Mexico represent a combined import share of about 22% in the US in both 2002 and 2004, climbing to 42% in 2003. But this significantly higher share in 2003 does not come from the fact that these two countries replaced other source-countries but originates mainly from a sharp decrease in imports of steel products from other foreign sources (-38% with respect to 2002 and 2004). Thus either Canada and

Mexico predominantly sell different steel products to the US than the rest of the world or, more likely, they do not have enough combined capacity to significantly replace these other foreign countries. As a result, the impacts of higher tariffs and their consequences on imports and on inventory as described in our analysis are reasonable despite the significant exclusion of these two countries.

Our analysis could be extended in several directions. For instance, it could be extended to other trade policies. Antidumping orders, in particular, have the potential of affecting the hedging strategy in a much stronger way than the safeguard tariffs that we study in our empirical exercise. The reason is that antidumping orders generally eliminate the price difference between domestic and foreign sources (the ‘dumping margin’) and thereby completely eliminate the attractiveness of the hedging strategy through inventory build-up. This means that, where hedging strategies matter due to uncertain demand, antidumping orders have the potential of having very strong protectionist effects.²¹

The normative aspects of hedging strategies are also potentially interesting. For instance, if it is the case that inventory costs can be interpreted as barriers to trade, does this imply that a free-trade regime that makes inventory build-up attractive (as would be the case with a big difference between domestic and foreign unit costs) could be sub-optimal from a social welfare point of view compared to a protectionist regime that involves domestic sourcing and discourages inventory build-up?

Finally, while we have focused on how importers may hedge demand uncertainty, the methods developed in our paper, both theoretical and empirical, may also prove useful in examining the effects of trade policy uncertainty. Whether associated with Brexit or with the U.S. trade policies, trade policy uncertainty is undoubtedly higher today than in the past. We also know that it has significant inventory effects. For instance, Hasbro, a U.S.-based toymaker which outsources a large fraction of its production to Asia, not only has to deal with the demand uncertainty associated with its toys during the critical Christmas shopping season but also with trade policy uncertainties. Since it has essentially no domestic sources in the United States able to supply close substitutes, it has to rely on imports with their consequences not only in regard to the level of inventory it wants to hold but also in regard to the timing of its orders (New York Times, Aug. 15, 2019, ‘Trump delays a holiday tax, but toymakers are still worried’).

²¹We thank one referee for suggesting a link with antidumping orders.

Appendix

A.1 The general model

Let $R(x_t + y_t, \epsilon_t)$ denote the revenue the firm realizes in period t that depends on how much it sells from existing inventory (x_t), how much it sells by sourcing from the domestic supplier (y_t), and on the demand realization that is distributed according to the c.d.f. $F(\epsilon)$. In each period t , the firm chooses x_t and y_t to maximize $R(x_t + y_t, \epsilon_t) - w_t y_t$ s.t. $x_t \leq z_t$, where z_t denotes the size of the inventory and w_t denotes the unit cost of domestic sourcing. Rewriting the maximization problem as a Lagrange function $L(x_t, y_t, \lambda_t) = R(x_t + y_t, \epsilon_t) - w_t y_t + \lambda_t(z_t - x_t)$, we obtain the following first-order conditions:

$$\begin{aligned} \frac{\partial L(x_t^*, y_t^*, \lambda_t^*)}{\partial x_t} &= \frac{\partial R(x_t^* + y_t^*, \epsilon_t)}{\partial x_t} - \lambda_t^* \leq 0, \lambda_t^* \geq 0, \\ &\left(\frac{\partial R(x_t^* + y_t^*, \epsilon_t)}{\partial x_t} - \lambda_t^* \right) (z_t - x_t^*) = 0 \end{aligned} \quad (\text{A.1})$$

and

$$\begin{aligned} \frac{\partial L(x_t^*, y_t^*, \lambda_t^*)}{\partial y_t} &= \frac{\partial R(x_t^* + y_t^*, \epsilon_t)}{\partial y_t} - w_t \leq 0, y_t^* \geq 0, \\ &\left(\frac{\partial R(x_t^* + y_t^*, \epsilon_t)}{\partial y_t} - w_t \right) y_t^* = 0. \end{aligned} \quad (\text{A.2})$$

(A.1) and (A.2) imply optimal sales $x_t^*(z_t, \epsilon_t)$ and $y_t^*(z_t, \epsilon_t)$ such that maximized profits in period t are equal to:

$$\pi^* [x_t^*(z_t, \epsilon_t), y_t^*(z_t, \epsilon_t), m_t, \epsilon_t] = R(x_t^*(z_t, \epsilon_t) + y_t^*(z_t, \epsilon_t), \epsilon_t) - w_t y_t^*(z_t, \epsilon_t) - v_t m_t,$$

where m_t denotes the quantity of imports ordered and paid in period t that arrive in period $t + 1$. Inventory evolves such that $z_{t+1} = z_t - x_t^*(z_t, \epsilon_t) + m_t$.

The firm knows the demand realization ϵ_0 in the initial period when it decides on x_0^* and y_0^* according to (A.1) and (A.2), respectively, but only its distribution for all subsequent periods. In order to determine the optimal imports m_0^* in the initial period (and thus the optimal inventory z_1^*), we may apply Bellman's principle of optimality and consider the sum of expected discounted profits:

$$\begin{aligned}
\hat{\Pi}(m_0) &= \pi^* [x_0^*(z_0, \epsilon_0), y_0^*(z_0, \epsilon_0), m_0, \epsilon_0] \\
&+ \delta \int_{\underline{\epsilon}}^{\bar{\epsilon}} \pi^* [x_1^*(z_1, \epsilon_1), y_1^*(z_1, \epsilon_1), m_1^*, \epsilon_1] dF(\epsilon_1) \\
&+ \sum_{t=2}^{\infty} \delta^t \int_{\underline{\epsilon}}^{\bar{\epsilon}} \pi^* [x_t^*(z_t^*, \epsilon_t), y_t^*(z_t, \epsilon_t), m_t^*, \epsilon_t] dF(\epsilon_t),
\end{aligned} \tag{A.3}$$

where we have used the optimal import levels m_t^* for all periods after the initial period 0. For the initial period, z_0 is given, so m_0 determines z_1 , and the resulting optimal inventories for all periods $t \geq 2$ are determined by the optimal import levels m_t^* for $t \geq 1$. Note that we confine the analysis to cases in which the firm plans to import in each period such that $m_t^* > 0$ for $t \geq 1$, and consequently m_0 affects z_1 only and $\partial z_1 / \partial m_0 = 1$ holds. The optimal import level m_0^* is determined by $d\hat{\Pi}(m_0^*)/dm_0 = 0$, or equivalently by:

$$-v_0 + \delta \int_{\underline{\epsilon}}^{\bar{\epsilon}} \frac{\partial \pi^* [x_1^*(z_1, \epsilon_1), y_1^*(z_1, \epsilon_1), m_1^*, \epsilon_1]}{\partial z_1} dF(\epsilon_1) = 0, \tag{A.4}$$

where v_0 is the foreign unit cost in period 0. As the same problem arises anew in period 1 and in every period thereafter, we may use the equivalent of (A.4) to derive the optimal import level for each period, and thus to determine z_t^* , which gives the *ex ante* optimal size of the inventory for period t . Due to our assumptions on revenues $R(\cdot)$ and the pdf $F(\epsilon)$, we can arrive at closed form solutions for (A.4), in particular for the expected marginal revenue $E[\text{MR}_1(\cdot)] = (\partial \pi^* [x_1^*(z_1^*, \epsilon_1), y_1^*(z_1^*, \epsilon_1), m_1^*, \epsilon_1] / \partial x_1) dF(\epsilon_1)$; see, for instance, (7) for regime (123).

Our model can easily be extended to the case where imports will take place only every second period (or even less often), and (A.4) would have to be rewritten as

$$\begin{aligned}
-v_0 &+ \delta \int_{\underline{\epsilon}}^{\bar{\epsilon}} \frac{\partial \pi^* [x_1^*(z_1, \epsilon_1), y_1^*(z_1, \epsilon_1), 0, \epsilon_1]}{\partial z_1} dF(\epsilon_1) \\
&+ \delta^2 \int_{\underline{\epsilon}}^{\bar{\epsilon}} \frac{\partial \pi^* [x_2^*(z_2, \epsilon_2), y_2^*(z_2, \epsilon_2), m_2^*, \epsilon_2]}{\partial z_2} dF(\epsilon_2) = 0.
\end{aligned}$$

Note that this outcome would require a substantial differential between foreign and domestic costs such that (A.2) implies a corner solution for all possible ϵ -realizations in every second period. Our results, however, do not change, as we can collapse the two (or more) periods into one by defining

$$\begin{aligned}
E[\text{MR}_1(\cdot)] &= \int_{\underline{\epsilon}}^{\bar{\epsilon}} \frac{\partial \pi^* [x_1^*(z_1, \epsilon_1), y_t^*(z_1, \epsilon_1), 0, \epsilon_1]}{\partial z_1} dF(\epsilon_1) \\
&+ \delta \int_{\underline{\epsilon}}^{\bar{\epsilon}} \frac{\partial \pi^* [x_2^*(z_2, \epsilon_2), y_t^*(z_2, \epsilon_2), m_2^*, \epsilon_2]}{\partial z_2} dF(\epsilon_2)
\end{aligned}$$

such that $-v_0 + \delta E[\text{MR}_1(\cdot)] = 0$ holds as before.

A.2 Proof of Lemma 1

Whenever it exists, the optimal inventory is computed by finding z_t such that $\delta E[\text{MR}_t(\cdot)] = v_{t-1}$ (i.e., such that the discounted expected marginal revenue is equal to the marginal cost of a foreign-sourced unit). Consider $z_{t,2}^*$ corresponding to the case where sales are always equal to inventory and thus where there is no excess inventory and no domestic sourcing. The discounted expected marginal revenue is:

$$\delta E[\text{MR}_t] = \delta \int_{-\Delta}^{\Delta} (a + \epsilon_t - 2bz_t) \frac{d\epsilon_t}{2\Delta}.$$

Setting it equal to the foreign unit cost, v_{t-1} , yields the optimal inventory:

$$z_{t,2}^* = \frac{a - v_{t-1}/\delta}{2b}. \quad (\text{A.5})$$

This case requires $\underline{\epsilon}(z_{t,2}^*) < -\Delta$ and $\bar{\epsilon}(z_{t,2}^*) > \Delta$ which can be rewritten as:

$$\Delta < \text{Min} \{v_{t-1}/\delta - \rho_t, w_t - v_{t-1}/\delta\},$$

which in turn implies $2\Delta < w_t - \rho_t$. The optimal foreign inventories $z_{t,123}^*$ and $z_{t,2}^*$ are mutually exclusive.

Inventory level $z_{t,12}^*$ corresponds to the case where the firm never sources domestically. The discounted expected marginal revenue is given by:

$$\delta E[\text{MR}_t] = \delta \left(\int_{-\Delta}^{\underline{\epsilon}(z_t)} \rho_t \frac{d\epsilon_t}{2\Delta} + \int_{\underline{\epsilon}(z_t)}^{\Delta} (a + \epsilon_t - 2bz_t) \frac{d\epsilon_t}{2\Delta} \right),$$

leading to the optimal inventory:

$$z_{t,12}^* = \frac{a + \Delta - \rho_t - 2\sqrt{\Delta(v_{t-1}/\delta - \rho_t)}}{2b}. \quad (\text{A.6})$$

This case requires $\underline{\epsilon}(z_{t,12}^*) > -\Delta$ and $\bar{\epsilon}(z_{t,12}^*) > \Delta$ which can be rewritten as:

$$w_t - \rho_t > 2\sqrt{\Delta(v_{t-1}/\delta - \rho_t)}, \quad \Delta > v_{t-1}/\delta - \rho_t,$$

which in turn implies $w_t + \rho_t > 2v_{t-1}/\delta$.

Regime (23) corresponds to the case where the wholesaler never wants to accumulate excess inventory but could source domestically. The discounted expected marginal revenue is:

$$\delta E [MR_t] = \delta \left(\int_{-\Delta}^{\bar{\epsilon}(z_t)} (a + \epsilon_t - 2bz_t) \frac{d\epsilon_t}{2\Delta} + \int_{\bar{\epsilon}(z_t)}^{\Delta} w_t \frac{d\epsilon_t}{2\Delta} \right),$$

leading to the optimal foreign inventory:

$$z_{t,23}^* = \frac{a - w_t - \Delta + 2\sqrt{\Delta(w_t - v_{t-1}/\delta)}}{2b}. \quad (\text{A.7})$$

Since this case requires $\bar{\epsilon}(z_{t,23}^*) < \Delta$ and $\underline{\epsilon}(z_{t,23}^*) < -\Delta$, which, given $z_{t,23}^*$, can be written as:

$$\Delta > (w_t - v_{t-1}/\delta), \quad w_t - \rho_t > 2\sqrt{\Delta(w_t - v_{t-1}/\delta)},$$

which in turn implies that $w_t - \rho_t > 2(w_t - v_{t-1}/\delta)$.

Given our assumption that $w_t \geq v_{t-1}/\delta$, there is no interior solution for $z_{t,3}$, which corresponds to the case where the firm always engages in domestic sourcing ($\bar{\epsilon}(z_t) < -\Delta$). However, not surprisingly, if we allow $w_t < v_{t-1}/\delta$, then the firm never sources abroad so that $z_{t,3}^* = 0$. Regime (1) where sales would always be smaller than inventory ($\underline{\epsilon}(z_t) > \Delta$) so that the firm would want to accumulate excess inventory can be excluded since we assume $\rho_t < v_{t-1}/\delta$.

A.3 Proof of Proposition 1

Consider a temporary change in the tariff characterized by $dv_{t-1}/d\tau = 1$ and $dv_t/d\tau = d\rho_t/d\tau = 0$. We find:

$$\left. \frac{\partial z_{t,12}^*}{\partial \tau} \right|_{dv_t=0} = -\frac{\sqrt{\Delta}}{2b\delta\sqrt{v_{t-1}/\delta - \rho_t}} < 0, \quad (\text{A.8})$$

$$\left. \frac{\partial z_{t,123}^*}{\partial \tau} \right|_{dv_t=0} = -\frac{\Delta}{b\delta(w_t - \rho_t)} < 0, \quad (\text{A.9})$$

$$\left. \frac{\partial z_{t,2}^*}{\partial \tau} \right|_{dv_t=0} = -\frac{1}{2b\delta} < 0, \quad (\text{A.10})$$

$$\left. \frac{\partial z_{t,23}^*}{\partial \tau} \right|_{dv_t=0} = -\frac{\sqrt{\Delta}}{2b\delta\sqrt{w_t - v_{t-1}/\delta}} < 0. \quad (\text{A.11})$$

A permanent change in the tariff implies $dv_{t-1}/d\tau = dv_t/d\tau = d\rho_t/d\tau = 1$. We find

$$\frac{\partial z_{t,12}^*}{\partial \tau} = -\frac{\frac{(1-\delta)\sqrt{\Delta}}{\delta\sqrt{v_{t-1}/\delta - \rho_t}} + 1}{2b} < 0, \quad (\text{A.12})$$

$$\frac{\partial z_{t,123}^*}{\partial \tau} = -\frac{\frac{4\Delta(w_t - v_{t-1} - (1-\delta)\rho_t)}{\delta(w_t - \rho_t)^2} + 1}{4b} < 0, \quad (\text{A.13})$$

$$\frac{\partial z_{t,2}^*}{\partial \tau} = \left. \frac{\partial z_{t,2}^*}{\partial \tau} \right|_{dv_t=0} < 0, \quad (\text{A.14})$$

$$\frac{\partial z_{t,23}^*}{\partial \tau} = \left. \frac{\partial z_{t,23}^*}{\partial \tau} \right|_{dv_t=0} < 0. \quad (\text{A.15})$$

A.4 Proof of Proposition 2

In case of an anticipated change in the tariff in period t we have $dv_{t-1}/d\tau = 0$ and $dv_t/d\tau = 1$. Therefore

$$\left. \frac{\partial z_{t,12}^*}{\partial \tau} \right|_{dv_{t-1}=0} = \frac{\frac{\Delta}{\sqrt{v_{t-1}/\delta - \rho_t}} - 1}{2b} > 0,$$

as $z_{t,12}^*$ requires $\Delta > v_{t-1}/\delta \Leftrightarrow \Delta/\sqrt{\Delta}(v_{t-1}/\delta - \rho_t)$.

$$\left. \frac{\partial z_{t,123}^*}{\partial \tau} \right|_{dv_{t-1}=0} = \frac{\frac{4\Delta(w_t - v_{t-1}/\delta)}{\delta(w_t - \rho_t)^2} - 1}{4b},$$

which can be positive or negative.

A.5 Proof of Proposition 5

Given (A.7), (A.6) and 7), we have:

$$\frac{\partial z_{t,23}^*}{\partial \Delta} = \frac{1}{2b} \left(-1 + \sqrt{(w_t - v_{t-1}/\delta) / \Delta} \right) < 0,$$

provided $\Delta > (w_t - v_{t-1}/\delta)$, which is precisely one condition for $z_{t,23}^*$ to be valid (see point 3 in Appendix A.2).

Similarly,

$$\frac{\partial z_{t,12}^*}{\partial \Delta} = \frac{1}{2b} \left(1 - \sqrt{(v_{t-1}/\delta - \rho_t) / \Delta} \right) > 0,$$

provided $\Delta > v_{t-1}/\delta - \rho_t$. This is also precisely one of the conditions for $z_{t,12}^*$ to be valid (see point 2 in Appendix A.2).

Finally, $\partial z_{t,2}^* / \partial \Delta = 0$ and

$$\frac{\partial z_{t,123}^*}{\partial \Delta} = \frac{1}{2b(w_t - \rho_t)} (w_t + \rho_t - 2v_{t-1}/\delta)$$

is positive or negative depending on the values of w_t , v_{t-1}/δ and ρ_t with the constraint that $w_t > v_{t-1}/\delta > \rho_t$ as assumed.

To summarize, a mean-preserving marginal change in Δ has the following effects on the firm's optimal inventory level:

$$\frac{\partial z_t^*}{\partial \Delta} = \begin{cases} < 0 & \text{if } z_t^* = z_{t,23}^*, \\ >< 0 & \text{if } z_t^* = z_{t,123}^*, \\ = 0 & \text{if } z_t^* = z_{t,2}^*, \\ > 0 & \text{if } z_t^* = z_{t,12}^*. \end{cases} \quad (\text{A.16})$$

Noting that $z_{t,23}^* \leq z_{t,123}^* \leq z_{t,12}^*$ for $2\Delta > w_t - \rho_t$, (and $z_{t,23}^* \leq z_{t,2}^* \leq z_{t,12}^*$ for $2\Delta < w_t - \rho_t$), completes the proof.

References

- [1] Aizenman, J., 2004. Endogenous pricing to market and financing costs. *Journal of Monetary Economics* 51, 691–712.
- [2] Alessandria, G., Kaboski, J., Midrigan, V., 2010a. Inventories, lumpy trade and large devaluations. *American Economic Review* 100, 2304–39.
- [3] Alessandria, G., Kaboski, J., Midrigan, V., 2010b. The Great Trade Collapse of 2008-09: an inventory adjustment?. *IMF Economic Review* 58, 254–94.
- [4] Allon, G., Van Mieghem, J., 2010. Global dual sourcing: tailored base-surge allocation to near- and offshore production. *Management Science* 56, 110–124.
- [5] Boute, R., Van Mieghem, J., 2015. Global dual sourcing and order smoothing: the impact of capacity and lead times. *Management Science* 61, 2080–99.
- [6] Bown, C., 2013. How different are safeguards from antidumping? Evidence for US trade policies toward steel. *Review of Industrial Organization* 42, 449–81
- [7] Cachon, G., Terwiesch, C., 2019. *Matching Supply with Demand: An Introduction to Operations Management*. McGraw-Hill Education, New York.
- [8] Evans, C., Harrigan, J., 2005. Distance, time, and specialization: lean retailing in general equilibrium. *American Economic Review* 95, 292–313.
- [9] Gervais, A., 2018. Uncertainty, risk aversion and international trade. *Journal of International Economics* 115, 145–158.
- [10] Hall, G., Rust, J., 2000. An empirical model of inventory investment by durable commodity intermediaries. *Carnegie-Rochester Conference Series on Public Policy* 52, 171–214.
- [11] Hufbauer, G., Goodrich, B., 2003. Steel policy: The good, the bad and the ugly. *International Economics Policy Briefs PB03-1*. Institute for International Economics, Washington DC.
- [12] Hummels, D., Schaur, G., 2013. Time as a trade barrier. *American Economic Review* 103, 2935–59.
- [13] Hummels, D., Schaur, G., 2010. Hedging price volatility using fast transport. *Journal of International Economics* 82, 15–25.
- [14] Jain, N., Girotra, K., Netessine, S., 2014. Managing global sourcing: Inventory performance. *Management Science* 60, 5, 1202–22.
- [15] Krishnan H., Winter R., 2007. Vertical control of price and inventory. *American Economic Review* 97, 1840–57.

- [16] Kropf, A., Sauré, P., 2014. Fixed costs per shipment. *Journal of International Economics* 92, 166–184.
- [17] Novy, D., Taylor, A., 2019. Trade and uncertainty, *Review of Economics and Statistics*, forthcoming.
- [18] Qu, Z., Raff, H., Schmitt, N., 2018. Incentives through inventory control in supply chains. *International Journal of Industrial Organization* 59, 486–513.
- [19] Read, R., 2005. The political economy of trade protection: the determinants and welfare impact of the 2002 US emergency steel safeguard measures. *The World Economy* 28, 1119–37.