Spread Too Thin: The Impact of Lean Inventories*

Julio L. Ortiz †
Federal Reserve Board

December 2024

Abstract

Widespread adoption of just-in-time (JIT) production has reduced inventory holdings. This paper measures the frequency of JIT adoption among public firms and quantifies a trade-off created by JIT between firm profitability and vulnerability to supply disruptions. Empirically, JIT adopters experience higher sales and less volatility on average while also exhibiting heightened sensitivity to aggregate supply conditions and weather events faced by their suppliers. I explain these facts in a structurally estimated general equilibrium model of JIT production. Relative to a counterfactual economy without JIT, the baseline model implies higher firm profitability in normal times but a deeper contraction amid a supply disruption. A transition to an equilibrium with less JIT and larger inventory stocks leads to a 4% output loss.

Keywords: Inventory investment. Heterogeneous firms. Supply disruption. Just-in-time production.
JEL Codes: D22, D25, E22, E23, G31

*I would like to thank Adam Guren, Matteo Iacoviello, Nils Lehr, Hyunseung Oh, Pascual Restrepo, and Stephen Terry for their valuable insights and suggestions. I would also like to thank discussants George Alessandria, Ryan Charhour, and Jun Nie, and the participants at many seminars and conferences. The views expressed in this paper are solely the responsibility of the author and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System.

†Federal Reserve Board, 20th and C Street NW, Washington, DC 20551; Email: julio.l.ortiz@frb.gov.

1 Introduction

Up to 70% of manufacturers have reportedly adopted just-in-time (JIT) production, a management philosophy that aims to minimize the time between orders.[1] Firms adopt JIT to cut costs associated with managing large material purchases and storing idle stocks. Instead, these firms commit to placing smaller more frequent orders from suppliers.[2] Consequently, lean inventory management has contributed to the approximately 20% decline in the aggregate inventory-to-sales ratio since 1970.[3]

Do improvements in inventory management matter for macroeconomic fluctuations? Theoretically, in general equilibrium, inventories have been found to be immaterial for aggregate dynamics (Khan and Thomas, 2007; Iacoviello et al., 2011). Empirically, some find that inventory management improvements decreased aggregate volatility (Davis and Kahn, 2008) while others (Stock and Watson, 2002) find that it was broadly inconsequential.

This paper offers a different perspective on the role of lean inventories in driving aggregate fluctuations, finding that it can create macroeconomic fragility in the face of unexpected supply disruptions such as those experienced from the onset of COVID-19. I document evidence of a trade-off from a dataset of JIT firms and quantitatively assess the role that lean production plays at the aggregate level in a structurally estimated heterogeneous firms model.

I first measure the frequency of JIT adoption among public firms by comparing firm inventory holdings to historical industry-level inventory holdings.[4] Based on this measure, JIT adoption increased in popularity from 1980 through the late 2000s. Using my measure of JIT adoption, I provide firm-level evidence linking the JIT adoption decision to higher firm sales and lower firm volatility. This provides motivating evidence as well as a set of moments that I use when struc-

[1] In 2015, the Compensation Data Manufacturing & Distribution Survey found that 71% of surveyed firms employ lean manufacturing. Similarly, in 2007, the Industry Week/MPI Census of Manufacturers found that 70% of respondents had implemented lean manufacturing.

[2] Ohno (1988) provides a detailed history of JIT which started with Toyota's Kanban system.

[3] U.S. Bureau of Economic Analysis, Ratios of nonfarm inventories to final sales of domestic business [A812RC2Q027SBEA], retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/A812RC2Q027SBEA.

[4] I validate this approach by comparing JIT firms identified from this approach to a narrower set of JIT firms identified in the accounting literature through recorded public announcements.

turally estimating the model. Within firms, JIT adoption is associated with a 25% increase in sales, and sales per worker, and a 55% increase in earnings. In addition, JIT firms experience between a 4% to 18% decline in sales growth volatility, employment growth volatility, and earnings growth volatility. These empirical results, though not causal, are consistent with positive selection into JIT which subsequently yields firm-level efficiency gains as in my model.

I then exploit variation external to the firm and document that JIT adopters are more exposed to supply disruptions as proxied by fluctuations in supply chain pressures according to the New York Fed Global Supply Chain Pressure Index. At the firm level, sales among JIT firms decline more strongly than their non-JIT counterparts when there is an increase in supply chain pressures. JIT firms experience an additional 2.5% decline in sales compared to non-JIT firms. In addition, JIT adopters experience a sharper drop in sales when their suppliers are faced with adverse weather events. My analysis points to heightened sensitivity among JIT firms upon the realization of external supply shocks, indicating that an economy composed of more JIT producers is less resilient to such disturbances.

In light of these empirical facts, I build and structurally estimate a dynamic general equilibrium model of JIT production. The model features a distribution of firms that differ in idiosyncratic productivity, inventory holdings, and inventory management strategy. Materials, needed for production, can be acquired subject to a stochastic fixed order cost. JIT firms draw order costs from a distribution that is first order stochastically dominated by those of non-JIT firms. Implementing JIT requires incurring an initial adoption cost and a smaller continuation cost thereafter. In a given period, firms must choose their JIT status, how much to order, and how much to produce.

I numerically solve and structurally estimate my baseline model via the simulated method of moments (SMM) using data from 1980 to 2019. Relative to a counterfactual economy without JIT, the baseline economy features lower overall inventory holdings and higher output. In addition, because JIT adoption leads to a reduction in fixed order costs, JIT adopters can better align their material input usage with realized productivity. As a result, firm-level volatility is lower than in the counterfactual economy as JIT firms can smooth out their ordering cycles.

Whereas individual adopters benefit from JIT in normal times, an economy comprised of more lean producers is more vulnerable to supply disruptions. I model a shock to fixed order costs, calibrated to match the drop in U.S. real GDP during the onset of the COVID-19 pandemic, and find that the baseline economy experiences a decline in output along the transition that is 60% larger than in the counterfactual economy. An unexpected spike in fixed order costs causes firms' ordering inaction regions to expand, leading to a decline in orders which reduces inventory investment. With fewer material inputs on hand, firm sales also fall. At the aggregate level, the economy therefore experiences both a decline in final sales and inventory investment. These effects are more pronounced in the baseline economy where firms carry fewer stocks to begin with. In Appendix D, I show that this result also holds in a version of the model in which there is aggregate uncertainty about fixed order costs.

Finally, in light of recent discussions around re-shoring and achieving supply chain "resilience," I examine how the JIT economy would transition to a new steady state that features less JIT adoption and higher inventory holdings. I find that along the transition to an equilibrium that features 25% less JIT adoption, the stock of inventories increases by about 20% while output falls by 4% and consumption-equivalent welfare declines by -2.4%.

In short, my empirical and theoretical analysis quantifies a trade-off between long-run gains and macroeconomic vulnerability to supply disruptions. Firms benefit in normal times from pursuing a lean inventory strategy, however upon the realization of a supply disruption, an economy populated by more JIT firms experiences a deeper contraction.

Inventory investment has long been of interest as a potential source of macroeconomic volatility.[5] Seminal contributions developed production smoothing models (Ramey and Vine, 2004; Eichenbaum, 1984), stock out avoidance models (Kahn, 1987), and (S,s) models (Scarf, 1960; Caplin, 1985) of inventory investment. Khan and Thomas (2007) elegantly models inventories in general equilibrium and finds that they play little to no role in amplifying or dampening business

[5] See for instance Ahmed et al. (2004), McConnell and Perez-Quiros (2000), McCarthy and Zakrajsek (2007), Irvine and Schuh (2005), and McMahon and Wanengkirtyo (2015).

cycles.[6] My model is similar though I introduce an endogenous JIT adoption decision and analyze a shock to fixed order costs rather than an aggregate productivity shock. Another related paper is Alessandria et al. (2023), which develops a rich two-country general equilibrium model and studies unexpected shocks to domestic and international shipping delays. I focus on analyzing the firm-level JIT adoption decision and its implications in a closed-economy model. The fixed order cost shock that I study encompasses shipping delays among other factors that may shift the probability of placing an order.

This paper also speaks to the management literature that focuses on assessing the gains to JIT. Kinney and Wempe (2002) finds that JIT adopters outperform non-adopters, primarily through profit margins.[7] Gao (2018) examines the role of JIT production in corporate cash hoarding. My paper provides a bridge between evidence documented in the management literature and the rich literature on inventories in macroeconomics by highlighting how JIT production can matter for aggregate outcomes.

Furthermore, this paper relates to the literature on supply chain disruptions. On the empirical front, I adopt a strategy similar to Barrot and Sauvagnat (2016) to determine whether JIT producers are disproportionately exposed to unexpected weather events. Other empirical work has assessed how shocks propagate through a network of firms.[8] Similarly, Cachon et al. (2007) assesses empirical evidence of the bullwhip effect along the supply chain. From a theoretical perspective, my paper relates to models of heterogeneous firms, sunk costs, and supply chains. As previously noted, Alessandria et al. (2023) study delays in a general equilibrium model. Furthermore, Meier (2020) models supply chain disruptions in the context of time to build. My paper explicitly links supply disruptions to an important source of investment at the aggregate level, inventory accumulation.

The rest of the paper is organized as follows. Section 2 documents evidence that is consistent with the stabilizing effects of JIT at the firm level along with the exposure to unexpected shocks that it engenders at the macro level. Sections 3 and 4 develop the general equilibrium model of lean

[6] Iacoviello et al. (2011) comes to a similar conclusion through a different model. On the other hand, Wen (2011) builds a stock out avoidance model and finds that inventories can stabilize aggregate fluctuations.

[7] Nakamura et al. (1998) as well as Roumiantsev and Netessine (2008) find similar evidence.

[8] For instance, Carvalho et al. (2021) does this in the context of the 2011 Japanese earthquake.

production. I estimate the model in Section 5 . Section 6 quantifies the aforementioned trade-off associated with JIT, and Section 7 concludes.

2 Empirical Patterns Among JIT Firms

I start by describing my approach to measuring JIT producers. I then document empirical evidence indicating that these JIT producers are more efficient and yet more exposed to supply disruptions. I use this as motivating evidence for the model outlined in Section 3. This analysis will also provide moments and external validation to the model once I structurally estimate it.

2.1 A New Measure of JIT Adoption

I develop a measure of lean production among public firms by comparing firm-level inventory holdings to historical industry inventory holdings. For a given two-digit NAICS sector, I compute the median inventory-to-sales ratio from 1971 to 1979. Then, for each year from 1980 to 2019, I define a lean producer as a firm whose inventory-to-sales ratio is below the pre-1980 industry median inventory-to-sales ratio.

Ideally, one would measure JIT by observing which firms actually implement it or announce their intention to do so. In practice, this is a challenging undertaking as firms may not always explicitly disclose their intention or decision to adopt JIT. However, since a hallmark of JIT is the commitment to reducing or eliminating inventories, it stands to reason that if we observe significant declines in inventory holdings among firms in an industry, then it this must be due to the adoption of JIT and related technologies. In Appendix A, I study the JIT adoption decision among a narrower set of manufacturing firms whose decision to adopt JIT is identified based on public announcements. Using this alternative definition of JIT, I verify that inventory-to-sales ratios decline within the firm following the adoption of JIT. In addition, I verify that the majority of the firms in this narrower set of JIT producers are also identified as JIT producers by the approach that I take here.

Figure 1 plots the measured frequency of JIT adoption based on my measure. This measure

Figure 1: Frequency of JIT Adoption

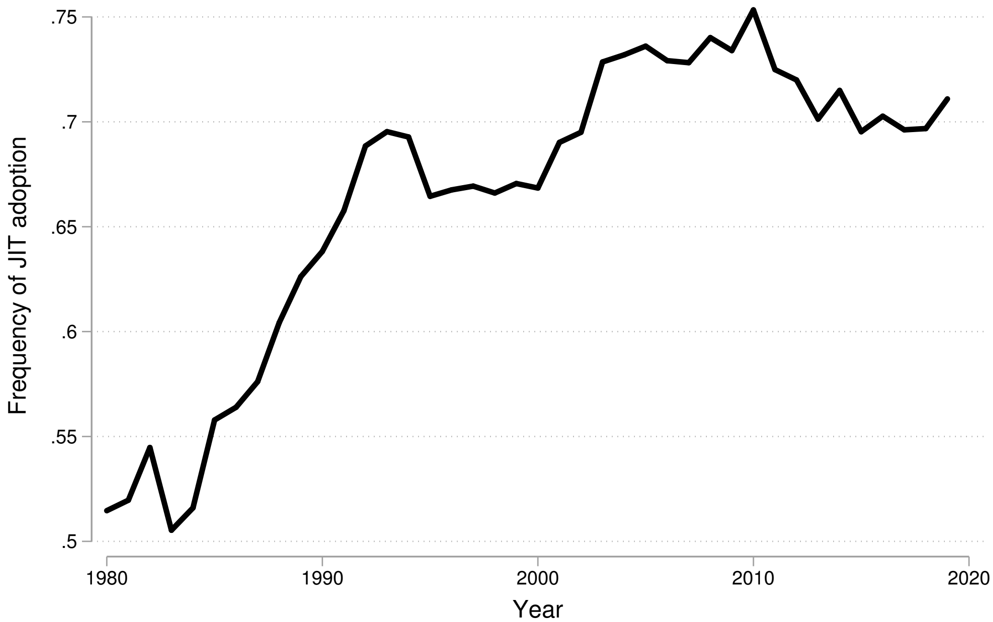

Note: The figure plots the measured frequency of JIT adoption over time.

implies that the frequency of JIT production trended up through the 1980s and part of the 1990s, consistent with evidence in the operations management literature which finds that JIT was popularized in the 1970s and 1980s following the success of the Toyota Kanban system. Around the mid 1990s, the frequency of adoption continued to increase, though at a slower pace, and peaked around 2010. The leveling off in the popularity of JIT observed from the mid-to-late 2000s could be attributed to uncertainty during the Great Recession, which perhaps activated precautionary inventory holding motives, or due to the formation of intricate and geographically vast supply chains which may have resulted in longer lead times. Over the 2010s, the frequency of adoption declined slightly though remained at above 70% in my sample.

Table 1: JIT Adoption and Firm Performance

	Sales	Sales per worker	Earnings
JIT	0.257***	0.254***	0.552***
	(0.016)	(0.012)	(0.057)
Fixed effects	Firm, Sector × Year	Firm, Sector × Year	Firm, Sector × Year
Firms	4,821	4,821	4,821
Observations	45,477	45,477	45,477

Note: The table reports panel regression results based on regression (1). The dependent variables are log sales, log sales per worker, and the inverse hyperbolic sine of earnings. Earnings are defined as income before extraordinary items. Two-digit NAICS codes are specified in the sector-by-year fixed effects. Standard errors are double clustered at the firm and fiscal year levels. The standard deviations of the dependent variables are 2.34, 0.87, and 3.25, respectively. *** denotes 1% significance, ** denotes 5% significance, and * denotes 10% significance.

2.2 Empirical Evidence

I use my measure of JIT adoption along with other firm-level balance sheet information from Compustat Fundamentals Annual data over the aforementioned years to examine differences in outcomes between JIT and non-JIT firms. To complete certain exercises in this section, I merge my sample with additional information such as the New York Fed's Global Supply Chain Pressure Index and county-level weather events from the National Oceanic and Atmospheric Administration (NOAA).

My final sample consists of an unbalanced panel of about 5,000 unique manufacturing firms spanning the years 1980 to 2019. Appendix A provides summary statistics.

I document four sets of facts about JIT adopters. First, JIT adoption is associated with higher sales, higher sales per worker, and higher earnings, the latter of which is measured as operating income.[9] I estimate regressions of the following form:

$$y_{ijt} = \gamma \text{JIT}_{ijt} + \mathbf{X}'_{ijt}\beta + \delta_{jt} + \delta_i + \nu_{ijt}, \tag{1}$$

where y_{ijt} is an outcome variable for firm i belonging to sector j in year t. The regressor of interest, JIT_{ijt}, is a time-varying indicator for whether a firm is a JIT adopter in a given year. I specify firm

[9] This finding is consistent with Fullerton and McWatters (2001) and Cua et al. (2001).

Table 2: JIT Adoption and Firm Volatility

	Sales growth volatility	Employment growth volatility	Earnings growth volatility
JIT	-0.180***	-0.043*	-0.164***
	(0.029)	(0.024)	(0.048)
Fixed effects	Sector × Year	Sector × Year	Sector × Year
Firms	2,452	2,449	2,452
Observations	16,333	16,274	16,317

Note: The table reports panel regression results based on regression (2). The dependent variables are rolling five-year standard deviations of firm sales growth, employment growth, and earnings growth. Lagged log capital stock and age in sample are specified as a controls. Two-digit NAICS codes are specified in the sector-by-year fixed effects. Standard errors are double clustered at the firm and fiscal year levels. *** denotes 1% significance, ** denotes 5% significance, and * denotes 10% significance.

and sector-by-year fixed effects in these regressions. Table 1 reports the regression results. The first column implies that JIT adopters experience a roughly 25% increase in sales following adoption. In addition, firms experience an estimated 25% increase in sales per worker and a 55% increase in earnings following JIT adoption. The results imply changes of 10%, 30%, and 17% of one standard deviation in the outcomes, respectively. The regression results allude to the benefits of JIT in my model. Facing lower fixed order costs, adopters hold fewer inventories in favor of placing smaller more frequent orders. Upon shrinking their inventory stocks, adopters also incur fewer carrying costs. These cost reductions enable JIT firms to allocate more resources to production, allowing them to generate more sales. As a result, these coefficients reflect both selection and treatment effects.

Second, JIT adopters experience less micro volatility. I estimate the following regression:

$$y_{ijt} = \gamma \text{JIT}_{ijt} + \mathbf{X}'_{ijt}\beta + \delta_{jt} + \eta_{ijt}, \tag{2}$$

where y_{ijt} now denotes rolling 5-year standard deviations of sales growth or employment growth for firm i in sector j in year t. Table 2 reports the results. Adopters see a roughly 18%, 4%, and 16% decline in sales growth volatility, employment growth volatility, and earnings growth volatility,

Table 3: JIT Adoption and Supply Chain Pressures

	(1) Sales	(2) Sales
Supply chain pressure	0.018	
	(0.012)	
Supply chain pressure × JIT	-0.025**	-0.024**
	(0.009)	(0.009)
Fixed effects	Firm	Firm, Sector × Year
Firms	2,766	2,766
Observations	20,608	20,608

Note: The table reports panel regression results from regression (3). The dependent variable is the log of firm sales. Lagged firm assets and lagged finished goods inventories are specified as controls as well as firm age and contemporaneous unemployment rate, real GDP growth, and manufacturing PPI inflation. Two-digit NAICS codes are specified in the sector-by-year fixed effects. Standard errors are double clustered at the firm and fiscal year levels. *** denotes 1% significance, ** denotes 5% significance, and * denotes 10% significance.

respectively. This is consistent with the stabilizing role that JIT plays in my model. Due to the lower fixed ordering costs, firms are able to more easily time their orders and can smooth out their ordering cycles which moderates the variability of other outcomes as well.

I next document facts relating to firm-level exposure brought on by JIT, exploiting aggregate variation and examining sensitivity to aggregate supply conditions and weather events. The regression results accord with the model in that adopters are less insured against unanticipated supply disruptions, and an economy with more JIT firms is more exposed to such events.

Third, JIT adopters tend to be more sensitive to aggregate supply disruptions. I merge my data with aggregate data from the New York Fed Supply Chain Pressure Index and estimate regressions that interact adoption with supply chain pressures. Supply pressures spiked following the onset of COVID-19 as bottlenecks and other disruptions that hampered order fulfillment. If JIT firms are more sensitive to supply disruptions, then we should observe more adverse outcomes for these firms relative to non-JIT producers when supply chain pressures rise. I therefore estimate the following regression:

$$y_{ijt} = \gamma_1 \text{JIT}_{ijt-1} + \gamma_2 \text{GSCPI}_t + \gamma_3 \big[\text{JIT}_{ijt-1} \times \text{GSCPI}_t\big] + \mathbf{X}'_{ijt}\beta + \text{FE} + \varepsilon_{ijt}, \tag{3}$$

where GSCPI denotes the global supply chain pressure index, which I standardize, and $\mathbf{X}$ reflects a set of controls which includes firm age in the sample, a lag of total assets and finished goods inventories, and contemporaneous unemployment rate, real GDP growth, and manufacturing PPI inflation. The coefficient γ_3 measures the extent to which JIT firms exhibit more or less sensitivity to increases in supply chain pressures.

Table 3 reports the regression results. Based on column (1), a one standard deviation increase in the supply chain pressure index is associated with a roughly 2.5% stronger decline in sales among JIT firms. Turning to column (2), when controlling for sector-by-year fixed effects, which subsumes the second term of equation (3), I find that the magnitude of the excess sensitivity of JIT firms is similar, implying that JIT firm sales decline by about 2.4% more than non-JIT firms.

Finally, JIT adopters are more sensitive to weather events faced by their suppliers. I examine this by merging my data with county-level weather events from NOAA using the Compustat Segment Files and links from Barrot and Sauvagnat (2016). I then estimate the following regression:

$$y_{ist} = \psi_1 \text{JIT}_{it-1} + \psi_2 \text{WeatherEvent}_{st} + \psi_3 \big[\text{JIT}_{it-1} \times \text{WeatherEvent}_{st}\big] + \mathbf{X}'_{ist}\beta + \text{FE} + \omega_{ist}. \tag{4}$$

I consider two ways of defining the "WeatherEvent" regressor: (i) as an indicator for a weather event occurring in the zip code where supplier s is headquartered in a given year and (ii) as the dollar value of property damage caused by the weather event. I collect information on county-level weather events from NOAA and link these events to public firm headquarter zip codes via the aforementioned Barrot and Sauvagnat (2016) links.

Ideally, one would want to link upstream weather events to the zip codes in which suppliers' production takes place. The Compustat data is limited in this respect since once cannot necessarily assume that production occurs at or near a firm's headquarters. Nonetheless, weather events may

Table 4: JIT Adoption and Weather Events

	(1)	(2)	(3)	(4)
Weather event indicator	-0.003			
	(0.036)			
Weather event indicator × JIT	-0.007	-0.195***		
	(0.038)	(0.069)		
Property damage			-0.0002	
			(0.002)	
Property damage × JIT			-0.0003	-0.010***
			(0.002)	(0.003)
Fixed effects	Firm, Supplier, Year	Firm, Supplier×Year	Firm, Supplier, Year	Firm, Supplier×Year
Firms	196	68	196	68
Observations	1885	317	1885	317

Note: The table reports panel regression results based on regression (4). The dependent variable is log sales. The control variables specified include lagged JIT indicator, lagged log capital stock, and firm age in sample. Standard errors are double clustered at the customer-supplier level. *** denotes 1% significance, ** denotes 5% significance, and * denotes 10% significance.

disrupt other relevant operations which might take place at a firm's headquarters such as logistics. Overall, I interpret this data limitation as a form of measurement error which likely biases my estimates toward zero.

Table 4 provides four sets of results which summarize the estimated sales response to supplier weather events. The first two columns report the effect of a weather event on sales when specifying the weather event indicator variable. The final two columns instead report the property damage with respect to weather events. The point estimates on the interaction between a supplier weather event and a JIT customer are negative across all specifications and are statistically significant when controlling for supplier-by-year fixed effects. The latter estimates, reported in columns (2) and (4), control for upstream time-variation which includes year-specific supplier characteristics such as age and size, as well as other unobserved shocks that suppliers face in a given year. Given the more robust set of controls specified in these regressions, columns (2) and (4) reflect my preferred specifications.

Through the series of links required to estimate these regressions, the sample size is reduced considerably.[10] Nonetheless, in my preferred specifications, I find that on average supplier weather

[10]Building this sample requires linking weather events to firm (supplier) headquarters in Compustat, then linking these suppliers to their customers (through the Segment files), and finally linking the customers to their JIT adoption

events in my sample predicts an additional 19.5% decline in JIT firm sales, and a 1% increase in the property damage caused by a given weather event is associated with a 1% excess sales contraction among JIT firms relative to non-JIT firms.

Taken together, the data suggest that JIT adopters benefit from more sales and smoother outcomes. At the same time, adoption is associated with heightened exposure to aggregate supply conditions and supply disruptions as proxied by local weather events. My model of heterogeneous firms with inventories, fixed ordering costs, and an endogenous JIT adoption decision can explain these patterns. The model also allows me to quantitatively assess the impact of JIT amid an unanticipated aggregate supply disruption, something that cannot be easily captured by firm level regressions.

3 A Model of Just-in-Time Production

Having illustrated the essence of the trade-off in the data, I next build a dynamic general equilibrium model which will provide quantitative statements about the implications of JIT. The model is similar in spirit to Khan and Thomas (2007) and Alessandria and Choi (2007).

A representative household has preferences over consumption and leisure. The household supplies its labor frictionlessly to the two sectors of the economy: the intermediate goods sector and the final goods sector. A representative intermediate goods firm produces materials by using labor and capital. In addition, a continuum of heterogeneous final goods firms make use of labor and materials to produce using a decreasing returns to scale technology. Final goods producers are heterogeneous in idiosyncratic productivity, inventory stocks, and JIT adoption status. All markets are perfectly competitive.

The representative household is endowed with one unit of time in each period and values consumption and leisure according to the following preferences:[11]

$$U(C_t, H_t) = \log(C_t) + \phi(1 - H_t),$$

status.

[11] Rogerson (1988) microfounds these preferences in a model of indivisible labor and lotteries.

where $\phi > 0$ denotes the household's labor disutility. Total hours worked is denoted by H_t and labor is paid wage, w_t. In addition to wage income, the household earns a dividend each period from ownership of firms, D_t, and chooses savings on a one period riskless bond, B_{t+1}, given interest rate R_{t+1}. The representative household, facing no aggregate uncertainty, maximizes its utility:

$$\max_{C_t, H_t, B_{t+1}} \sum_{t=0}^{\infty} \beta^t U(C_t, H_t),$$

subject to its budget constraint which holds for all t,

$$C_t + B_{t+1} \leqslant R_t B_t + w_t H_t + D_t.$$

The parameter $\beta \in (0, 1)$ is the household's subjective discount factor.

The representative intermediate goods firm produces materials using capital K_t and labor L_t according to:

$$F(K_t, L_t) = K_t^{\alpha} L_t^{1-\alpha}.$$

Taking prices as given, the problem of the intermediate goods firm is:

$$\max_{K_t, L_t} q_t F(K_t, L_t) - w_t L_t - R_t K_t,$$

where q_t denotes the price of the intermediate good.

Finally, a continuum of final goods firms produce using materials, m_t, and labor, n_t, according a decreasing returns to scale technology:

$$y_t = z_t m_t^{\theta_m} n_t^{\theta_n}, \quad \theta_n + \theta_m < 1,$$

where idiosyncratic productivity evolves as an AR(1) in logs:

$$\log(z_{t+1}) = \rho_z \log(z_t) + \sigma_z \varepsilon_t, \quad \varepsilon_t \sim N(0, 1).$$

Figure 2: Decisions of Final Goods Firms

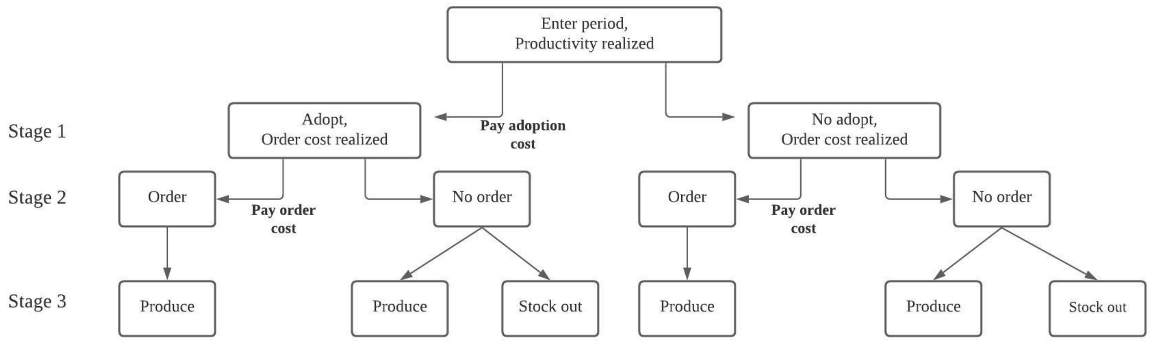

Note: The figure summarizes the order of the decisions made by final goods firms within a period.

Materials are drawn from the firm's existing inventory stock, s_t, to use in production. Final goods firms procure new materials from the intermediate goods firm subject to a stochastic fixed order cost drawn from a known distribution.

Figure 2 details the final goods producers' decision-making timeline. Each period consists of three stages. A producer enters the period with realized productivity, z_t, inventory stock, s_t, and adoption status, a_t. In the first stage, the producer decides whether or not to adopt JIT. If a producer does not enter the period as a continuing adopter, it must pay c_s in order to initially adopt. Alternatively, if the producer enters the period as an adopter, it must pay a smaller continuation cost $0 < c_f < c_s$ in order to maintain its status as a JIT producer.

Intuitively, adopting JIT requires that a plant repurpose its shop floor, enter into long-term contracts with suppliers to fulfill orders in a timely fashion, and possibly even purchase new technologies to facilitate information sharing with suppliers. The sunk setup cost encompasses all of these one-time costs. The continuation cost embodies smaller costs for suppliers to participate in timely delivery, costs of training labor on JIT best practices, and greater attention or communication required to share information with suppliers.

In the next stage, producers learn their order costs, $\xi \sim F(\xi)$, and decide whether or not to place an order, o_t. JIT producers face a more favorable order cost distribution, $\mathbb{E}(\xi_A) \leqslant \mathbb{E}(\xi_{NA})$. Lastly, following the adoption and the order decisions, final goods producers decide how much to produce.

I characterize the final goods firms' problem in terms of inventory stocks rather than specific order or material input choices. In particular, if a firm enters the period with inventory stock s_t, its target inventory stock is denoted by s_t^*. This means that any orders, if placed, are defined as $o_t = s_t^* - s_t$. Following the order decision, suppose that inventory stock $\widetilde{s}_t$ is carried into the production stage.[12] Materials used in production are then defined as $m_t = \widetilde{s}_t - s_{t+1}$ where s_{t+1} refers to the inventory stock carried forward into the next period. In the recursive formulation of the final goods firm problem that follows, I suppress the time subscript and instead denote next period variables with a prime.

Stage 1: Adoption Decision

A final goods producer begins the period with (z, s, a), faces labor-denominated adoption costs $\{c_s, c_f\}$, and endogenous prices, p, q, and w. The firm first decides whether to adopt JIT. Note that the adoption status is a binary outcome. The value of adopting is:

$$V^A(z,s,a) = \max\left\{ -pwc(a) + \int V^O(z,s,1,\xi)dF(\xi_A), \int V^O(z,s,0,\xi)dF(\xi_{NA})\right\}, \tag{5}$$

where

$$c(a) = \begin{cases} c_s & \text{if no JIT } (a=0) \\ c_f & \text{if JIT } (a=1), \end{cases}$$

and $V^O(z, s, a, \xi)$ refers to the firm's value in the second stage. Order costs are assumed to be distributed uniformly: $F(\xi) = U(\underline{\xi}, \overline{\xi})$.[13] The firm's optimal adoption policy, $a'(z, s, a)$, solves (5).

[12] $\widetilde{s} = s$ if no order is placed and $\widetilde{s} = s^*$ if an order is placed.

[13] As in Khan and Thomas (2007), I assume uniformly distributed order costs. In my context, uniformly distributed order costs are appealing because they strengthen the firms' precautionary inventory holding motive since order costs are not clustered around a central region as with, for instance, a normal distribution. To the extent that my later results expose a vulnerability associated with firms carrying too few inventories, this assumption should be relatively conservative.

Stage 2: Order Decision

Given the firm's order cost draw, ξ, also denominated in units of labor, it then decides whether to place an order, o. If the firm is an adopter, its order cost distribution is first order stochastically dominated by those of non-adopters. The value in the second stage is

$$V^O(z, s, a, \xi) = \max\Big\{ - pw\xi + V^*(z, s, a), V^P(z, s, a)\Big\}, \tag{6}$$

where the value of placing an order is[14]

$$V^*(z, s, a) = \max_{s^* \geqslant s}\Big[- pq(s^* - s) + V^P(z, s^*, a)\Big], \tag{7}$$

and $V^P(z, s, a)$ is defined below. The firm's order problem delivers a threshold rule. In particular, a firm places an order if and only if the order cost draw is lower than a threshold order cost: $\xi < \xi^*(z, s, a)$ where

$$\widetilde{\xi}(z, s, a) = \frac{V^*(z, s, a) - V^P(z, s, a)}{pw}, \tag{8}$$

and

$$\xi^*(z, s, a) = \min\big(\max\big(\underline{\xi}, \widetilde{\xi}(z, s, a)\big), \overline{\xi}\big). \tag{9}$$

Stage 3: Production Decision

Upon choosing its JIT status, deciding whether to place an order, and potentially selecting an order size, the firm then makes a production decision. Suppose that a firm enters the production stage with inventory stock $\widetilde{s}$ such that:

$$\widetilde{s} = \begin{cases} s^*\big(z, s, a'(z, s, a)\big) & \text{if order placed} \\ s & \text{if no order placed.} \end{cases}$$

[14]The constraint on the ordering decision allows for only positive orders. In other words, this model abstracts away from inventory liquidation.

In the production stage, the firm selects labor, $n(z, \widetilde{s}, s', a)$, and materials, $(\widetilde{s} - s')$, to maximize profits. Its value function in the production stage is:

$$V^P(z, \widetilde{s}, a) = \max_{s' \in [0, \widetilde{s}]} \pi(z, \widetilde{s}, s', a) + \beta \mathbb{E}\big[V^A(z', s', a')\big] \tag{10}$$

where

$$\pi(z, \widetilde{s}, s', a) = p\Big[zn(z, \widetilde{s}, s', a)^{\theta_n}(\widetilde{s} - s')^{\theta_m} - wn(z, \widetilde{s}, s', a) - c_m s'\Big] \tag{11}$$

are period profits. The end of period inventory stock is denoted by s', and c_m is a linear carrying cost of storing unused inventory.

A final goods producer is said to stock out if it enters the period with no inventories, $s = 0$, and chooses to not place an order. Without any inventories, the firm has no material inputs to draw from when making its production decision. As a result, the firm forgoes production in that period. The producer can flexibly restart production in the future conditional on a favorable productivity realization and order cost draw. I formally define the model equilibrium in Appendix B.

4 Analyzing the Model

The endogenous adoption decision allows the model to replicate important features of the data, namely, higher profitability and reduced micro volatility among JIT firms. Since implementing JIT comes at a relatively large sunk cost, not all firms optimally choose to adopt JIT. Figure 3 plots the adoption frontiers for JIT and non-JIT producers. The shaded area in the lower right corner represents the region of the state space in which non-JIT firms choose to adopt JIT. This illustrates the positive selection into adoption implied by the model. Moreover, the scope for initiating adoption is decreasing in inventory stocks as the value of adopting is higher among firms that are closer to their ordering thresholds.

At the same time, a producer is likely to remain an adopter conditional on already being one. This is because the continuation cost of retaining JIT is smaller than the initial sunk cost. Hence,

Figure 3: Adoption Frontiers

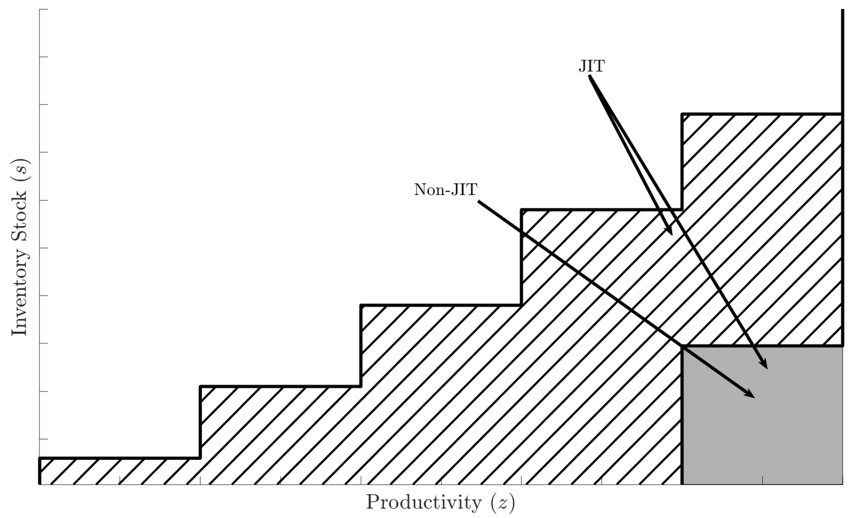

Note: The figure plots the adoption frontier among JIT and non-JIT firms. The solid gray area depicts the region of the state space in which non-JIT firms select into adoption. The striped area and the gray area jointly denote the region of the state space in which existing JIT firms choose to remain adopters.

the endogenous adoption decision exhibits persistence. The larger striped area in Figure 3 confirms this intuition. Only the least productive adopters will opt to abandon JIT. Furthermore, the scope for exiting adoption is increasing in inventory holdings. The positive selection detailed here could contribute to the patterns among JIT firms documented in the data. In particular, the decision to adopt JIT reflects a favorable productivity realization which, when coupled with lower average order costs, leads firms to reduce their inventory stocks and generate more sales such that there is an increase in sales per worker.

Figure 4 shows the probability of placing an order as a function of inventories. Consistent with the decision to select into adoption, order probabilities are increasing in productivity and decreasing in inventory holdings. The benefits of JIT adoption can be understood by comparing the two panels. Across both inventory levels, the probability of placing an order is higher for adopters since they face lower average order costs.

Finally, Figure 5 plots the optimal order size for JIT and non-JIT producers, holding the inventory stock constant. Order sizes among JIT firms never exceed those of non-JIT firms because the

Figure 4: Ordering Probabilities

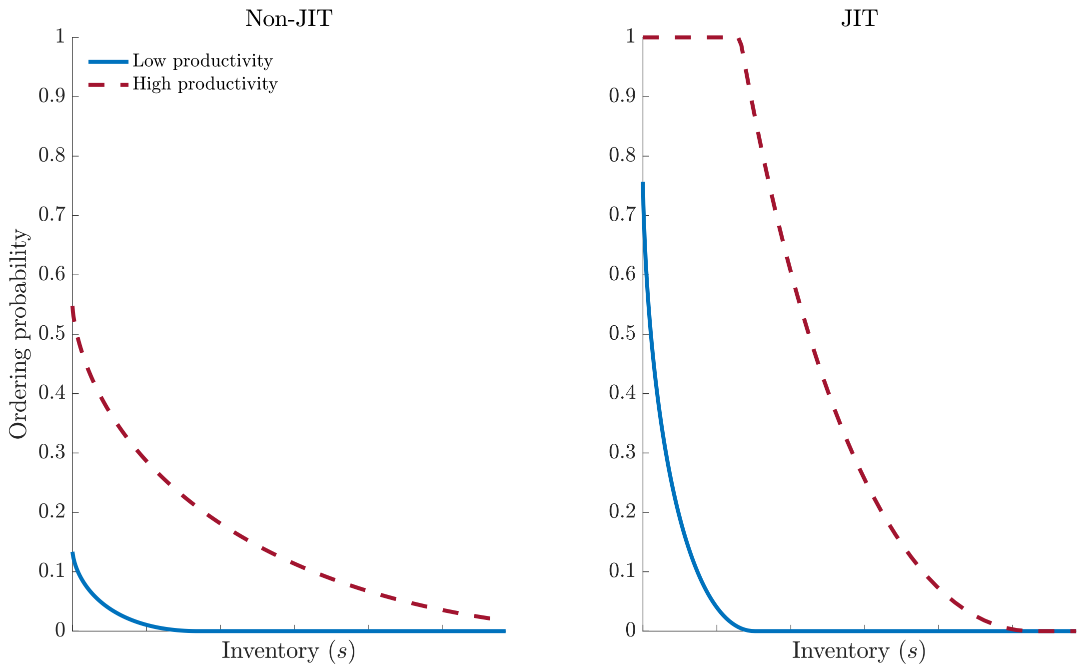

Note: The figure plots the probability of placing an order in the ordering stage as a function of inventories. The left panel plots the probabilities among non-adopters and the right panel plots the probabilities for adopters. The solid blue line reflects a low productivity establishment in the model while the dashed red line reflects a high productivity establishment.

former face lower average fixed order costs. At low and high levels of productivity, JIT and non-JIT firms optimally place similarly sized orders because they are either less profitable and require fewer material inputs since they generate less sales, or they are more profitable and are able to place larger orders irrespective of the fixed order costs. At intermediate levels of productivity, the differences in the fixed order cost distributions faced by JIT and non-JIT firms is important and is where JIT firms place smaller-sized orders.

5 Structural Estimation

I structurally estimate the model using the micro data analyzed in Section 2. The estimated model captures important features of the firm-level data including the levels of and covariances between

Figure 5: Order Size

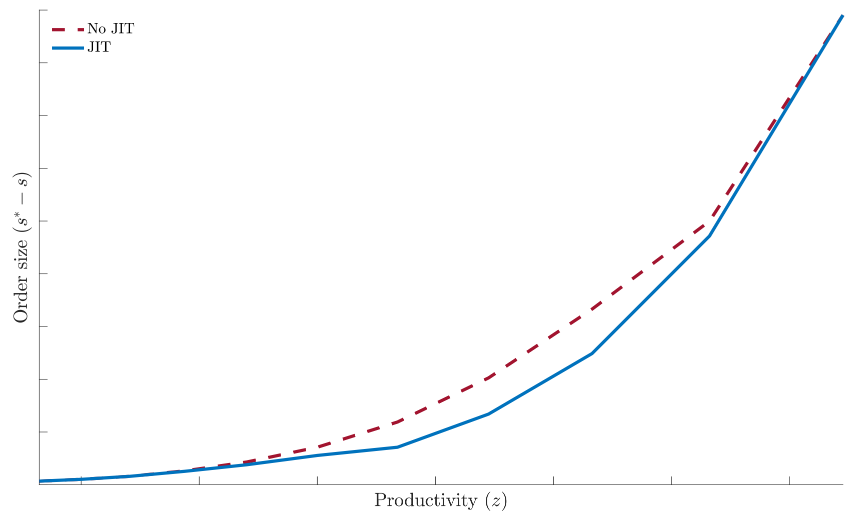

Note: The figure plots the order size for different realizations of idiosyncratic productivity. The solid blue line reflects a JIT producer while the dashed red line reflects a non-JIT producer.

inventories and sales as well as the frequency of JIT adoption and mode switching.

There are 14 parameters in the model. I externally fix seven parameters to match standard targets in the literature. Table 5 details the annual calibration. The discount factor, β is set to be consistent with a real rate of 4%. The material share, θ_m, is set to match the material share in the NBER-CES database, and the capital share, α, is fixed to match the capital-output ratio. The parameter θ_n is set to match an economy-wide labor share of 0.65. The leisure preference is calibrated so that the household works for one-third of total hours. Finally, I set the order cost lower bounds to zero for both JIT adopters and non-adopters.

5.1 Simulated Method of Moments

I estimate the remaining seven parameters. The parameter vector to be estimated is

$$\theta = \left(\rho_z \ \ \sigma_z \ \ \overline{\xi}_{NA} \ \ \overline{\xi}_A \ \ c_s \ \ c_f \ \ c_m\right)'.$$

These parameters residing in θ govern the exogenous productivity process, order costs, adoption costs, and the carrying cost. The model has no closed form solution, so I solve it using standard

Table 5: External Parameterization

Description	Parameter	Value
Discount Factor	β	0.962
Material share	θ_m	0.520
Capital share	α	0.390
Labor share	θ_n	0.280
Labor disutility	ϕ	2.450
Order cost lower bound	$\underline{\xi}_{NA}$	0.000
Order cost lower bound	$\underline{\xi}_{A}$	0.000

Note: The table reports the seven calibrated model parameters.

numerical dynamic programming techniques detailed in Appendix B. To parameterize the model, I employ SMM (Duffie and Singleton, 1993; Bazdresch et al., 2018). This is done by computing a set of targeted moments in the model and minimizing the weighted distance between the empirical moments and their model-based analogs.

Specifically, I target 10 moments to estimate the seven parameters. My estimator is therefore an overidentified SMM estimator. Of the ten moments, four are specific to JIT firms and four to non-JIT firms. These four moments, which are the same across both types of firms, are: the mean inventory-to-sales ratio and the covariance matrix of log sales and log inventories, the latter of which delivers three moments. I final two moments are the observed frequency of JIT adoption and the frequency of switching out of JIT adoption.[15] I specify the asymptotically efficient weighting matrix which is the inverse of the covariance matrix of the moments. Appendix C provides a discussion of the relationship between the moments and parameters which offers some intuition behind the identification of the model parameters.

Table 6 reports the estimated baseline model parameters, all of which are precisely estimated.[16] The technology parameters, ρ_z and σ_z, are consistent with parameterizations in the literature (Khan and Thomas, 2013; Khan et al., 2020; Hennessy and Whited, 2007).

[15]The empirical moments are listed in Table 7.

[16]A test of overidentifying restrictions delivers a J-statistic of 3.16 with a p-value of 0.37 for the baseline model. As a result, I fail to reject that the baseline model is misspecified, lending further support to the validity of my estimates.

Table 6: Estimated Baseline Parameters

Description	Parameter	Estimate	Standard error
Productivity shock persistence	ρ_z	0.861	0.0001
Productivity shock dispersion	σ_z	0.108	0.0002
Order cost upper bound (non-JIT)	$\overline{\xi}_{NA}$	0.740	0.0009
Order cost upper bound (JIT)	$\overline{\xi}_{A}$	0.102	0.0009
Sunk cost of adoption	c_s	0.963	0.0015
Continuation cost of adoption	c_f	0.067	0.0003
Carrying cost	c_m	0.158	0.0004

Note: The table reports the point estimates and standard errors for the seven estimated parameters. Standard errors obtained via numerical differentiation.

The upper support of the order cost distribution among non-adopters is 0.740. On the other hand, I find that the upper bound of order costs for JIT firms is 0.102. The average order costs implied by these estimates amount to 10% and 3% of value added. Furthermore, the adoption cost estimates suggest a meaningful amount of hysteresis in the adoption decision. In particular, firms pay a continuation cost that is about one quarter of the original sunk cost. Conditional on being an adopter, the probability of remaining an adopter is 98%. For reference, this estimate is higher than estimates of the sunk cost of exporting, which place the probability of remaining an exporter conditional on already being one at 87% (Alessandria and Choi, 2007). In equilibrium, economy-wide carrying costs are about 2% of value added, a non-negligible amount that prevents firms from storing too many inventories across time.

Table 7: Baseline Model vs. Empirical Moments

Moment	Model	Data
Mean(inventory-sales ratio$\mid$non-adopter)	0.283	0.280
		(0.004)
Mean(inventory-sales ratio$\mid$adopter)	0.128	0.071
		(0.001)
Var(log sales $\mid$non-adopter)	0.444	0.476
		(0.015)
Cov(log sales, log inventories$\mid$non-adopter)	0.363	0.305
		(0.012)
Var(log inventories$\mid$non-adopter)	0.380	0.401
		(0.008)
Var(log sales $\mid$adopter)	0.381	0.312
		(0.008)
Cov(log sales, log inventories$\mid$adopter)	0.127	0.244
		(0.010)
Var(log inventories$\mid$adopter)	0.463	0.429
		(0.024)
Frequency of adoption	0.662	0.660
		(0.010)
Frequency of switch out of adoption	0.020	0.057
		(0.001)

Note: The table reports model-based and empirical moments along with standard errors of the empirical moments.

5.2 Model Fit

Table 7 shows that the model is broadly successful in fitting the empirical moments. The model is able to reproduce lower average inventory-to-sales ratios, less sales dispersion among JIT firms relative to non-JIT firms, and empirically relevant adoption frequencies.

To further assess the baseline model's ability to match the patterns present in the data, I estimate the empirical regressions reported in Tables 1 and 2 based on a panel of simulated firms from the estimated baseline model. The results are reported in Table 8. Because the model abstracts away from an extensive margin of employment, I do not report model-based estimates for output per worker or employment growth volatility.

Following adoption, firms in the baseline model reduce their inventory-to-sales ratios and experience a 28% increase in sales, similar to the 25% increase estimated the data. The model also

Table 8: Empirical and Simulated Regressions

Panel A: Levels	Sales		Earnings	
	Model	Data	Model	Data
	0.282	0.257	0.184	0.552
	(0.004)	(0.016)	(0.003)	(0.057)
Panel B: Volatility	Sales growth		Earnings growth	
	Model	Data	Model	Data
	-0.281	-0.180	-0.347	-0.164
	(0.005)	(0.029)	(0.005)	(0.048)

Note: The table reports model-based and empirical regression coefficients with standard deviations and standard errors in parentheses.

predicts an increase in earnings following adoption, however the magnitude is smaller than in the data. Furthermore, the baseline model predicts reductions in firm sales volatility of 28% and a 35% decline in earnings growth volatility among JIT firms, similar though stronger than the 18% and 16% declines estimated in the data, respectively. With precisely estimated parameters delivering a broadly successful fit to targeted and non-targeted moments in the data, I can now exploit this structure as a laboratory for quantitative experiments.

6 Quantifying the Aggregate Effects of JIT

I next quantify the aggregate effects of the firm-level decision to adopt JIT. I begin by highlighting the vulnerability to unanticipated supply disruptions engendered by JIT. I then explore the aggregate implications of transitioning from a lean economy to a more "resilient" economy that features less JIT and, consequently, larger inventory stocks.

Figure 6: Deeper Crisis with More Adoption

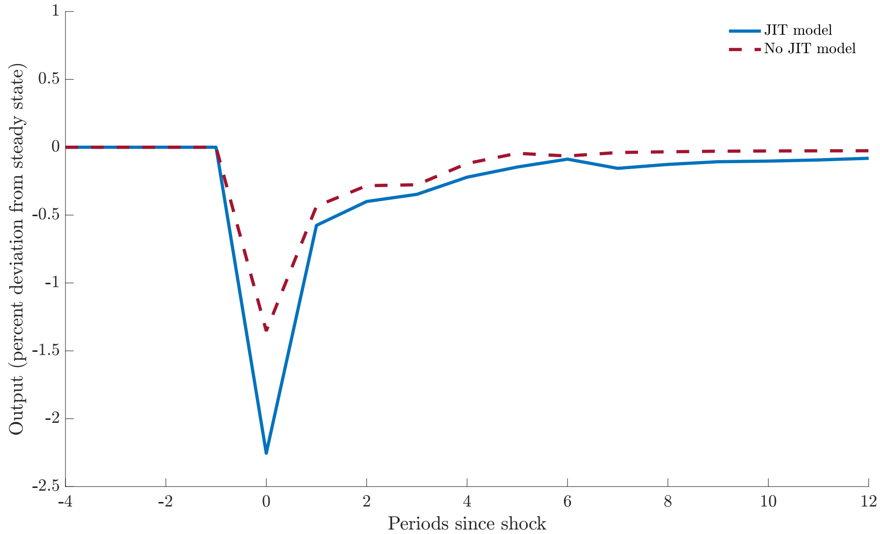

Note: The figure plots the output response to a fixed order cost shock that matches the 2.20% annual decline in real GDP in 2020. The persistence of the shock is set to 0.50.

6.1 Effects of an Unanticipated Supply Disruption

A natural benchmark against which to compare the estimated model is a world in which JIT adoption is not possible. I define such a counterfactual by solving a version of the estimated model with adoption cost parameters c_s and c_f fixed to be prohibitively large such that no adoption takes place.

Despite enjoying higher profits and smoother sales, an economy populated by lean producers is more vulnerable to an unexpected supply disruption. To quantify this supply side vulnerability, I consider an unexpected increase in economy-wide fixed order costs and assume that it evolves deterministically according to $\zeta_{t+1} = \rho_\zeta^\xi \zeta_t$ where $\rho_\zeta^\xi = 0.50$ and $\zeta_0 > 0$.[17] This shock shifts the average fixed order cost distribution of JIT and non-JIT producers:

$$\underline{\xi}_t = \underline{\xi} + \zeta_t \quad \text{and} \quad \overline{\xi}_t = \overline{\xi} + \zeta_t.$$

I calibrate the size of the order cost shock to reproduce a 2.2% GDP contraction in the baseline

[17] Consistent with the literature modeling COVID-19, I model the episode as an unanticipated event (Arellano et al., 2023; Espino et al., 2020).

Figure 7: Sources of the Stronger Decline in the Baseline Model

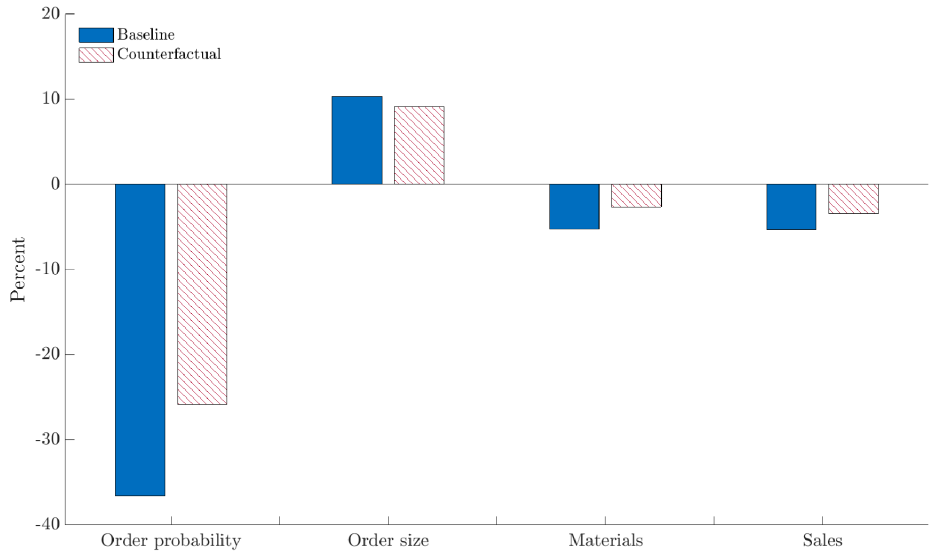

Note: The figure plots endogenous responses to a fixed order cost shock that matches the 2.2% annual decline in real GDP in 2020. The persistence of the shock is set to 0.50. The bars reflect cumulative responses along the transition back to steady state.

JIT model, in line with the annual contraction observed in U.S. GDP in 2020. I then introduce the same shock to the counterfactual model and compare the endogenous outcomes across the two economies. Figure 6 displays the output response to this unexpected shock. The JIT economy sees a roughly 1.7 percentage point excess output contraction on impact. The total output loss in the baseline economy, including along the transition back to steady state, is 5.5% while it is 3.4% in the counterfactual economy.

Figure 7 reports the key differences in endogenous responses between the two models. Amid the supply disruption, order-placing probabilities decline more in the baseline model relative to the counterfactual model. As an optimal response to the decline in ordering probabilities, firms in both economies increase their order sizes. Order sizes, however, rise more in the baseline economy, mirroring the stronger decline in ordering probabilities.

Despite the increase in order sizes, the extensive margin of ordering dominates so that aggregate orders decline in both economies though more so in the baseline model. Since inventory investment is the equal to the value of orders less materials, the stronger decline in aggregate orders relative to materials, and the general equilibrium decline in the price of orders in both economies, leads to a decline in inventory investment. This decline is more pronounced in the baseline economy.[18] Hence, from the perspective of the following identity:

$$\text{Output} = \text{Final sales} + \text{Inventory investment},$$

the excess output contraction in the baseline model comes from both a relatively stronger decline in final sales and a stronger fall in inventory investment. The relative contributions of final sales and inventory investment to GDP growth in the baseline model are consistent with what was observed amid the onset of COVID-19.[19]

A seemingly minor difference in inventory management strategies across the two models delivers a substantial difference in the extent to which the economy falls into crisis amid a supply disruption. The excess output loss amounts to slightly more than \$130 billion, a figure comparable to the funds appropriated to support state and local governments amid the pandemic.[20] Lean inventory management therefore can play a meaningful role in determining the vulnerability of the economy to unanticipated supply disruptions. During these episodes, the extent to which inventories can serve as a stabilizing force is economically significant.

In Appendix D, I explore a similar shock to fixed order costs but under aggregate uncertainty. My finding that the JIT economy is more vulnerable to supply disruptions, modeled as a shock to fixed order costs, continues to hold in this environment.

6.2 Understanding the Source of Supply Vulnerabilities

Firms in the baseline economy are more sensitive to the supply disruption because they carry fewer inventories. Figure 8 plots the relevant policies for a firm, at an intermediate level of firm produc-

[18] Amid the contraction, some JIT producers rethink their inventory management practices altogether. As a result, part of the stronger decline in ordering probabilities in the baseline model relative to the counterfactual model reflects mode switching since firms that return to being non-JIT producers face higher fixed ordering costs.

[19] In the data, final sales of domestic product declined by -1.69% in 2020. In the model, final sales decline by 1.77%.

[20] Coronavirus Aid, Relief, and Economic Security Act, H.R. 748, 116th Congress (2020).

Figure 8: Ordering During Supply Disruption

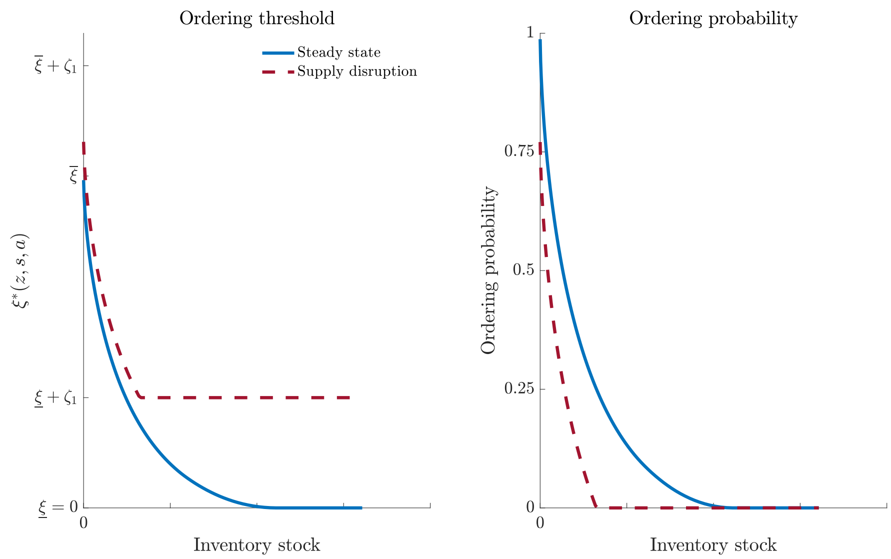

Note: The left panel plots a producer's ordering threshold as a function of its inventory stock in the steady state and amid a shock to fixed order costs at time $t = 1$. The right panel plots the probability of placing an order as a function of inventories in the steady state and amid the shock to fixed order costs at time $t = 1$.

tivity, as a function of inventories. Each panel plots two policies: one which reflects the steady state and the other which reflects the first period in which the supply disruption is realized. The left panel plots the ordering threshold policy function, $\xi^*(z, s, a)$ and the right panel plots the ordering probability.

The downward sloping region of the threshold order policy reflects the relative value of placing an order, $\widetilde{\xi}(z, s, a)$, as defined in equation (8). Recall that the ordering threshold is $\xi^*(z, s, a) = \min(\max(\underline{\xi}, \widetilde{\xi}(z, s, a)), \bar{\xi})$. Given its idiosyncratic state, a producer places an order if the relative value of doing so exceeds the fixed order cost. At higher levels of inventories, the relative value of ordering additional materials declines until it eventually falls below the lower bound of the order cost distribution, where $\xi^*(z, s, a) = \underline{\xi}$. Here, the ordering threshold is flat and the probability of placing an order is zero.

In response to an increase in fixed order costs, the threshold order policy shifts up and ordering probabilities shift down as the jump in fixed order costs exceeds the increase in the value of placing an order. Firms with larger stocks of inventories in the steady state, however, are less sensitive to the supply disruption as evidenced by the smaller downward shift in the order-placing probabilities at higher levels of inventories. This is because at higher levels of inventories, the value of placing an order is already relatively low. Since the counterfactual no JIT model features more firms operating with larger stocks of inventories, order-placing probabilities do not decline as much in the counterfactual model, and more firms are able to continue to produce without running out of materials.

Overall, the sensitivity of the economy to unanticipated supply disruptions depends on the level of inventories held by firms. The decision of how many inventories to hold in turn depends on these order thresholds which are themselves shaped by the decision to adopt JIT.

6.3 Transition to "Resilience"

Finally, I examine how the JIT economy would transition to a "new normal" that features higher fixed ordering costs. A new steady state with higher fixed ordering costs is intended to capture a greater risk of supply disruptions such as those experienced during and after the COVID-19 pandemic as well as amid heightened geopolitical risks. A prominent example of a firm which moved away from lean inventories is Toyota following the Fukushima earthquake.[21]

I consider an increase in economy-wide fixed order costs that generates a roughly 25% decline in the frequency of adoption in the new steady state. Figure 9 plots four panels which illustrate the transition from the lean steady state to the new, higher inventory steady state. The top left panel shows that, along the transition the frequency of JIT adoption gradually declines and ultimately falls by about 25%. Similarly, the probability of placing an order gradually declines by about 15%. Faced with higher fixed ordering costs, which represent elevated risks of disrupted order fulfillment, firms optimally increase their inventory holdings in an effort to buffer against such risks. The aggregate stock of inventories increases by about 20% along the transition. Meanwhile, output gradually declines over time and ultimately falls by 4% in the new steady state. Consumption-

[21] Learning from this episode, Toyota stocked up on its inventory of semiconductors which reportedly allowed it to weather the worst of the pandemic-related supply disruptions.

Figure 9: Transition to Higher Inventory Steady State

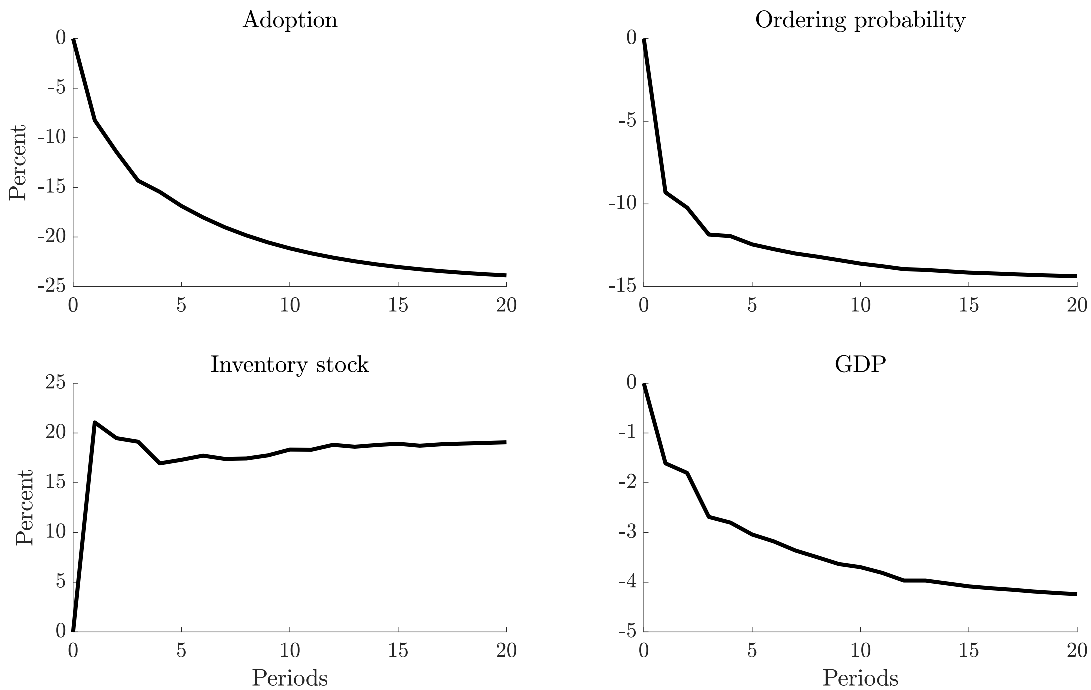

Note: The figure plots the transition of various endogenous outcomes to a new steady state that features a higher level of inventories.

equivalent welfare declines as well, falling by -2.4%. This exercise quantifies the potential costs associated with transitioning to a higher inventory steady state. A deeper analysis about the possible sources of more frequent supply disruptions after the COVID-19 pandemic, and an exploration of the incentives to abandon JIT, is beyond the scope of this paper and would be a fruitful avenue for future research.

7 Conclusion

In normal times, it pays to be lean. I provide empirical evidence of the benefits of JIT inventory management among publicly traded manufacturers. Upon adopting JIT, firms hold fewer inventories, and observe higher sales and smoother outcomes. JIT firms, however, are more susceptible to micro and macro supply disruptions. In a rich model of JIT production, firms that adopt JIT enjoy an increase in earnings and experience less unconditional micro volatility. At the same time, JIT elevates firm vulnerability to supply disruptions due to low inventory buffers. Amid an unexpected supply disruption, output in the estimated JIT economy contracts substantially more than a counterfactual economy with less JIT. Adoption, therefore, gives rise to an important trade-off which implies that inventories can matter for aggregate fluctuations. Economists interested in understanding fluctuations within firms, and the responsiveness of the economy to aggregate shocks, particularly supply disruptions, should play close attention to inventories and inventory management practices.

References

Ahmed, Shaghil, Andrew Levin, and Beth Anne Wilson (2004), "Recent U.S. Macroeconomic Stability: Good Policies, Good Practices or Good Luck?" *Review of Economics and Statistics*, 86, 824–832.

Alessandria, George and Horag Choi (2007), "Do Sunk Costs of Exporting Matter for Net Export Dynamics?" *The Quarterly Journal of Economics*, 122, 289–336.

Alessandria, George A, Shafaat Y Khan, Armen Khederlarian, Carter B Mix, and Kim J Ruhl (2023), "The Aggregate Effects of Global and Local Supply Chain Disruptions: 2020-2022." *Journal of International Economics*, 146.

Andrews, Isaiah, Matthew Gentzkow, and Jesse M. Shapiro (2017), "Measuring the Sensitivity of Parameter Estimates to Estimation Moments." *The Quarterly Journal of Economics*, 132, 1553–1592.

Arellano, Cristina, Yan Bai, and Gabriel Mihalache (2023), "Deadly Debt Crises: COVID-19 in Emerging Markets." *The Review of Ecoomic Studies*.

Barrot, Jean-Noël and Julien Sauvagnat (2016), "Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks." *The Quarterly Journal of Economics*, 131, 1543–1592.

Bazdresch, Santiago R., Jay Kahn, and Toni M. Whited (2018), "Estimating and Testing Dynamic Corporate Finance Models." *Review of Financial Studies*, 31, 322–361.

Cachon, Gerard, Taylor Randall, and Glen Schmidt (2007), "In Search of the Bullwhip Effect." *Manufacturing and Service Operations Management*, 4, 457–479.

Caplin, Andrew S. (1985), "The Variability of Aggregate Demand with (S,s) Inventory Policies." *Econometrica*, 53, 1395–1409.

Carvalho, Vasco, Makoto Nirei, Yukiko Saito, and Alrieza Tahbaz-Salehi (2021), "Supply Chain Disruptions: Evidence from the Great East Japan Earthquake." *Quarterly Journal of Economics*, 136.

Cua, Kristy O., Kathleen E. McKone, and Roger G. Schroeder (2001), "Relationships Between Implementation of TQM, JIT, and TPM, and Manufacturing Performance." *Journal of Operations Management*, 19, 675–694.

Davis, Steven and James Kahn (2008), "Interpreting the Great Moderation: Changes in the Volatility of Economic Activity at the Macro and Micro Levels." *Journal of Economic Perspectives*, 22, 155–180.

Den Haan, Wouter J. (2010), "Assessing the Accuracy of the Aggregate Law of Motion in Models with Heterogeneous Agents." *Journal of Economic Dynamics and Control*, 34, 79–99.

Duffie, Darrell and Kenneth J. Singleton (1993), "Simulated Moments Estimation of Markov Models of Asset Prices." *Econometrica*, 61, 929–952.

Eichenbaum, Martin S. (1984), "Rational Expectations and the Smoothing Properties of Inventories of Finished Goods." *Journal of Monetary Economics*, 14, 71–96.

Espino, Emilio, Julian Kozlowski, Fernando M. Martin, and Juan M. Sanchez (2020), "Seigniorage and Sovereign Default: The Response of Emerging Markets to COVID-19." Working Paper.

Fullerton, Rosemary R. and Cheryl S. McWatters (2001), "The Production Performance Benefits from JIT Implementation." *Journal of Operations Management*, 19, 81–96.

Gao, Xiaodan (2018), "Corporate Cash Hoarding: The Role of Just-in-Time Adoption." *Management Science*, 64, 4471–4965.

Hennessy, Christopher A. and Toni M. Whited (2007), "How Costly is External Financing? Evidence from a Structural Estimation." *Journal of Finance*, 62, 1705–1745.

Iacoviello, Matteo, Fabio Schiantarelli, and Scott Schuh (2011), "Input and Output Inventories in General Equilibrium." *International Economic Review*, 52, 1179–1213.

Irvine, Owen F. and Scott Schuh (2005), "Inventory Investment and Output Volatility." *International Journal of Production Economics*, 93-94, 75–86.

Kahn, James A. (1987), "Inventories and the Volatility of Production." *American Economic Review*, 77, 667–679.

Khan, Aubhik, Tatsuro Senga, and Julia Thomas (2020), "Default Risk and Aggregate Fluctuations in an Economy with Production Heterogeneity." Working Paper.

Khan, Aubhik and Julia Thomas (2007), "Inventories and the Business Cycle: An Equilibrium Analysis of (S,s) Policies." *American Economic Review*, 97, 1165–1188.

Khan, Aubhik and Julia Thomas (2013), "Credit Shocks and Aggregate Fluctuations in an Economy with Production Heterogeneity." *Journal of Political Economy*, 121, 1055–1107.

Kinney, Michael R and William F Wempe (2002), "Evidence on the Extent and Origins of JIT's Profitability Effects." *The Accounting Review*, 77, 203–225.

Koop, Gary, M Hashem Pesaran, and Simon M Potter (1996), "Impulse Response Analysis in Nonlinear Multivariate Models." *Journal of Econometrics*, 74, 119–147.

Krusell, P and A A Smith (1998), "Income and Wealth Heterogeneity in the Macroeconomy." *Journal of Political Economy*, 106, 867–896.

MacQueen, James B. (1966), "A Modified Dynamic Programming Method for Markovian Decision Problems." *Journal of Mathematical Analysis and Applications*, 14, 38–43.

McCarthy, Jonathan and Egon Zakrajsek (2007), "Inventory Dynamics and Business Cycles: What Has Changed?" *Journal of Money, Credit and Banking*, 39, 591–613.

McConnell, Margaret M. and Gabriel Perez-Quiros (2000), "Output Fluctuations in the United States: What Has Changed Since the Early 1980's?" *American Economic Review*, 90, 1464–1476.

McMahon, Michael and Boromeus Wanengkirtyo (2015), "Beyond Inventory Management: The Bullwhip Effect and the Great Moderation." Working Paper.

Meier, Matthias (2020), "Supply Chain Disruptions, Time to Build and the Business Cycle." Working Paper.

Nakamura, Maso, Sadao Sakakibara, and Roger Schroeder (1998), "Adoption of Just in Time Manufacturing Methods at U.S. and Japanese-Owned Plants: Some Empirical Evidence." *IEEE Transactions on Engineering Management*, 45, 230–240.

Ohno, Taiichi (1988). In *Toyota Production System: Beyond Large-Scale Production*, Productivity Press.

Porteus, Evan L. (1971), "Some Bounds for Discounted Sequential Decision Processes." *Management Science*, 18, 7–11.

Ramey, Valerie A. and Daniel J. Vine (2004), "Tracking the Source of the Decline in GDP Volatility: An Analysis of the Automobile Industry." NBER Working Paper 10384.

Rogerson, Richard (1988), "Indivisible Labor, Lotteries and Equilibrium." *Journal of Monetary Economics*, 21, 3–16.

Roumiantsev, Serguuei and Serguei Netessine (2008), "Should Inventory Policy Be Lean or Responsive? Evidence for US Public Companies." Working Paper.

Scarf, Herbert E. (1960), "The Optimality of (S,s) Policies in the Dynamic Inventory Problem." In *Mathematical Methods in the Social Sciences*, 196–202, Stanford University Press, Stanford.

Stock, James and Mark Watson (2002), "Has the Business Cycle Changed and Why?" *NBER Macroeconomics Annual*, 17, 159–224.

Tauchen, George (1986), “Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions.” *Economics Letters*, 20, 177–181.

Terry, Stephen (2017), “Alternative Methods for Solving heterogeneous Firm Models.” *Journal of Money, Credit, and Banking*, 49, 1081–111.

Wen, Yi (2011), “Input and Output Inventory Dynamics.” *American Economic Journal: Macroeconomics*, 3, 181–212.

Young, Eric R (2010), “Solving the Incomplete Markets Model with Aggregate Uncertainty Using the Krusell-Smith Algorithm and Non-Stochastic Simulations.” *Journal of Economic Dynamics and Control*, 34, 36–41.

SUPPLEMENTARY MATERIAL

Appendix A Empirics

This section provides summary statistics of the data used in Section 2. The section also includes further details on the JIT adopters sample, the weather regression results, and an alternative measure of JIT among public firms.

A.1 Sample Construction

My data come from three sources. First, I make use of annual Compustat data to obtain information on firm-level inventory holdings, sales, and other outcomes. Second, I use the New York Fed's Global Supply Chain Pressure index in the regressions that estimate the JIT sensitivity to aggregate supply conditions. Lastly, for the weather regressions, I collect county-level weather event data from NOAA and map them to firm headquarter zip codes.

Compustat Data

I make use of Compustat Fundamentals Annual data from 1971-2019. I keep only manufacturers (four-digit SIC codes between 2000-4000). In addition, I drop firm years in which acquisitions exceed 5% of total assets to avoid the influence of large mergers. To mitigate for any measurement error, I keep only those firms with non-missing and positive book value of assets, number of employees, inventories, and sales. All variables are winsorized at the top and bottom 0.5% of the empirical distribution.

Because the focus of the paper is on JIT, a concept that relates primarily to input inventories, I define the relevant measure of inventories to be the sum of raw material and works in process (`invrm+invwip`). This empirical definition also accords with the structural model developed in the main text in which producers carry stocks of inputs across time. My final sample consists of 5,912 unique firms. Table A1 reports summary statistics for the variables used.

Table A1: Compustat Summary Statistics

Variable name	Compustat code	Mean	Median	Standard Deviation	25%	75%
Earnings growth	$\frac{\Delta \texttt{ib}}{\texttt{ib}_{-1}}$	-0.285	-0.069	6.083	-0.805	0.449
Employment growth	$\frac{\Delta \texttt{emp}_t}{\texttt{emp}_{t-1}}$	-0.001	0.000	0.294	-0.083	0.100
Log employment	$\log(\texttt{emp}_t)$	-0.594	-0.693	2.023	-2.071	0.825
IHS earnings	$ihs(\texttt{ib}_t)$	0.839	0.850	3.248	-1.578	3.201
Log sales	$\log(\texttt{sale}_t)$	4.381	4.373	2.338	2.824	5.977
Sales growth	$\frac{\Delta \texttt{sale}_t}{\texttt{sale}_{t-1}}$	0.067	0.055	0.413	-0.062	0.178
Sales per worker	$\frac{\texttt{sale}_t}{\texttt{emp}_t}$	4.978	4.982	0.868	4.454	5.516
JIT adoption		0.662	1.000	0.473	0.000	1.000

Note: The table reports summary statistics for the relevant variables in the main text. The sample is constructed from Compustat Fundamentals Annual files for 1980-2019. Sample consists of 5,912 unique firms.

A.2 Validating Inventory-to-Sales-based Measure of JIT

In the main text, I measure JIT adoption based on historical declines of inventory-to-sales ratios. In this section, I compare this measure of JIT with an alternative measure of JIT that identifies the adoption of lean production based on financial news reports, press releases, and Form 10-K filings. These data were kindly provided to me by William Wempe, from his joint work with Michael Kinney, and Xiaodan Gao. See Kinney and Wempe (2002) and Gao (2018) for further details. This measure identifies JIT adoption for 177 manufacturing firms in my Compustat sample.

I start by regressing inventory-to-sales ratios on this alternative measure of JIT adoption to verify that inventory-to-sales ratios decline in the years following adoption for these firms. I estimate the following regression,

$$y_{ijt+h} = \gamma \text{adopt}_{ijt} + \delta_{jt} + \delta_i + \varepsilon_{ijt},$$

where the outcome of interest is the inventory-to-sales ratio, and adopt_{ijt} is an indicator taking on a value of one only in the recorded year of adoption. Industry-by-year and firm fixed effects are specified. The figure plots 95% confidence intervals for a three-year window around the recorded date of adoption, and shows that inventory holdings decline in the year of adoption and over the subsequent two years.

Figure A1: Validation of JIT Indicator

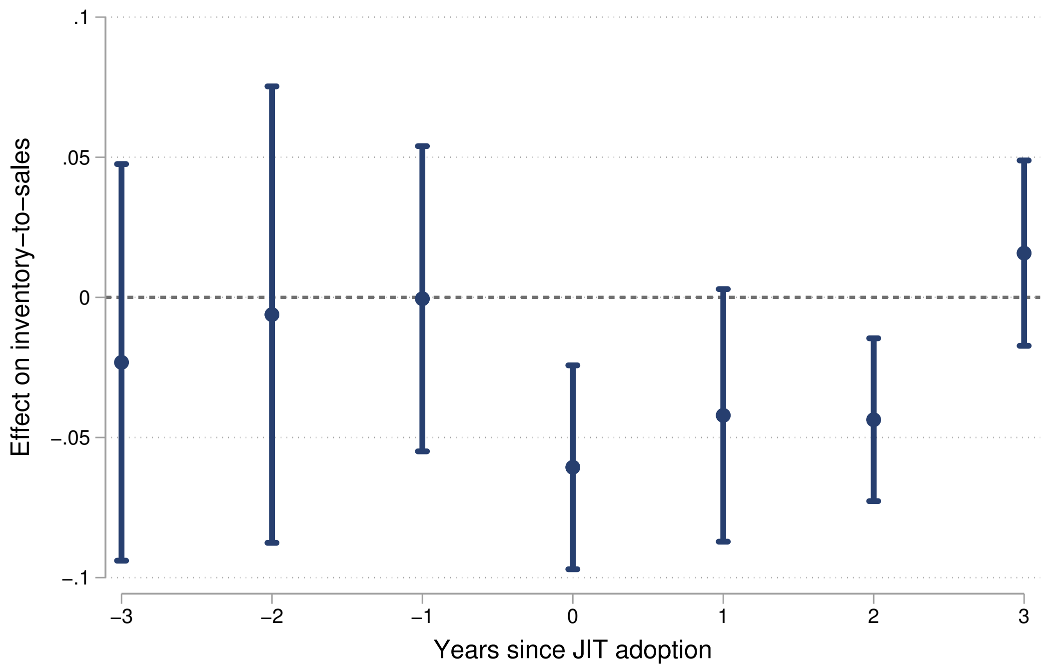

Note: The figure plots the estimated effect of JIT adoption on the level of inventory-to-sales. 95% confidence bands are displayed alongside point estimates.

Second, I compare the two measure of JIT adoption to verify that the measure of JIT used in the main text, which is able to capture more JIT firms, identifies those firms identified from the text-based approach introduced in this section. Of the 177 JIT adopters identified through this text-based approach, the measure of JIT based on historical inventory-to-sales ratios identified 152 of these firms as JIT producers.

Appendix B Model

B.1 Final Goods Firm

Firm profits are

$$\max_{s',n'} z n^{\theta_n}(s-s')^{\theta_m} - c_m s'$$

Maximizing out for n, we obtain

$$n = \left[\frac{\theta_n}{w} z(s-s')^{\theta_m}\right]^{\frac{1}{1-\theta_n}}.$$

Substituting this into the maximization problem, we obtain

$$\max_{s'}(1-\theta_n)\left[z\left(\frac{\theta_n}{w}\right)^{\theta_n}(s-s')^{\theta_m}\right]^{\frac{1}{1-\theta_n}} - c_m s',$$

or

$$\max_{s'}(1-\theta_n)y(z,s,s') - c_m s'$$

B.2 Intermediate Goods Firm

The intermediate goods firm problem is:

$$\max_{K,L} p\Big[qK^{\alpha}L^{1-\alpha} - RK - wL\Big]$$

The assumption that the intermediate goods firm utilizes a Cobb-Douglas production technology to produce implies that the intermediate goods firm's value can be expressed as a linear function of the aggregate capital stock. As a result, one can solve for q analytically. The price of the intermediate good is:

$$q = \left(\frac{1+r}{\alpha}\right)^{\alpha}\left(\frac{w}{1-\alpha}\right)^{1-\alpha}$$

B.3 Equilibrium

An equilibrium is a set of functions,

$$\{V^A, V^O, V^*, V^P, s^*, s', \xi^*, a', K, L, p, w, q, \Gamma_\mu\},$$

such that:

1. The household's first order conditions hold:

$$p = \frac{1}{C}, \quad w = \phi C.$$

2. The intermediate goods firm first order conditions hold:

$$w = (1-\alpha)q\left(\frac{K}{L}\right)^{\alpha} \quad R = \alpha q\left(\frac{L}{K}\right)^{1-\alpha}.$$

3. V^A, V^O, V^*, V^P solve the final goods firms' problem.

4. The market for final goods clears:

$$\begin{aligned} C = &\int\int y(z, s^*, s', a, \xi) dF(\xi^*) d\mu(z, s, a) + \int\int y(z, s, s', a, \xi)[1 - dF(\xi^*)] d\mu(z, s, a) \\ &- c_m\left(\int\int s'(z, s*, a) dF(\xi^*) d\mu(z, s, a) + \int\int s'(z, s, a)[1 - dF(\xi^*)] d\mu(z, s, a)\right) - K. \end{aligned}$$

5. The market for orders clears:

$$\left(\frac{(1-\alpha)q}{w}\right)^{\frac{1-\alpha}{\alpha}} K = \int\int [s^*(z, s, a) - s] dF(\xi^*) d\mu(z, s, a),$$

where the left hand side denotes the supply of orders, $K^\alpha L^{1-\alpha}$.

6. The market for labor clears:

$$H = \int\int n(z, s^*, s', \xi)dF(\xi^*)d\mu(z, s, a) + \int\int n(z, s, s', a, \xi)[1 - dF(\xi^*)]d\mu(z, s, a)$$

$$+ \int \left[\int_0^{\xi^*(z,s,a)} \xi dF(\xi) \right] d\mu(z, s, a) + \int a'(z, s, a)[(1-a)c_s + ac_f]d\mu(z, s, a) + \left(\frac{(1-\alpha)q}{w} \right)^{\frac{1}{\alpha}} K.$$

On the right hand side, the first two terms refer to labor demand from the final goods firms, the third term refers to the labor-denominated order cost, the fourth term refers to the labor-denominated adoption costs, and the final term refers to labor demand from the orders producer, L.

7. The evolution of the distribution of firms is consistent with individual decisions:

$$\Gamma_\mu(z, s, a) = \int\int\int 1_{\mathbb{A}} d\mu(z, s, a)dF(\xi)d\Phi(\varepsilon_z)$$

$$\mathbb{A}(z', s', a', \xi, \varepsilon_z; \mu) = \{(z, s, a) | s'(z, s, a, \xi; \mu) = s', z' = \rho_z + \sigma_z \varepsilon_z, a'(z, s, a, \xi; \mu) = a'\}$$

$$\Phi(x) = \mathbb{P}(\varepsilon_z \leqslant x),$$

B.4 Numerical Solution

The model is solved using methods that are standard in the heterogeneous firms literature. The exogenous productivity process is discretized following Tauchen (1986) which allows me to express the AR(1) process for log firm productivity as a Markov process. I select $N_z = 11$ grid points for idiosyncratic productivity and $N_s = 200$ grid points for the endogenous inventory holding state. Considering the binary adoption state, this implies that the discretized model has 4,400 grid points.

I solve for the policy functions via value function iteration which is accelerated by the use of the MacQueen-Porteus error bounds (MacQueen, 1966; Porteus, 1971). This acceleration method makes use of the contraction mapping theorem to obtain bounds for the true (infinite horizon) value function. These bounds are used to produce a better update of the value function. The ergodic distribution of firms is obtained via nonstochastic simulation as in Young (2010). This histogram-based method overcomes sampling error issues associated with simulating individual firms in order to obtain the stationary cross-sectional distribution.

Operationally, I solve the model by initiating a guess of the final goods price, p_0. Using the household and order producer's optimality conditions, I then obtain the implied wage and orders price, w_0 and q_0, given the guess p_0. From here, I solve the firm's problem via value function iteration and then obtain the ergodic distribution. Using the policies and ergodic distribution, I compute aggregates and the associated market clearing error from the household's optimality condition. I update the price based on this error via bisection.

For the unexpected shock exercises, I implement a standard shooting algorithm used to model deterministic dynamics. I fix the duration of the transition to a predetermined length T so that the model reaches steady state at $T+1$. I then solve the final goods firms problem backwards, obtaining a set of time-indexed policy functions. Using these policies, I push the distribution of final goods firms forward. With the time-indexed policies and weights in hand, I compute aggregates at each point in time and iterate on prices until the final goods market clears in each period, $\frac{1}{p_t} = C(p_t)$.

Appendix C Estimation

In this section, I detail the estimation of the model and provide additional results relating to identification.

C.1 Simulated Method of Moments

The parameter vector to be estimated is $\theta = \big(\rho_z \;\; \sigma_z \;\; \overline{\xi}_{NA} \;\; \overline{\xi}_A \;\; c_s \;\; c_f \;\; c_m\big)'$. Estimating θ requires making a guess, θ_0, solving and simulating the model, and computing the different moments. I collect the targeted empirical moments in a stacked vector $m(X)$ which comes from my Compustat sample. I next stack the model-based moments, which depend on θ, in the vector $m(\theta)$. Finally I search the parameter space to find the $\widehat{\theta}$ that minimizes the following objective

$$\min_{\theta} \big(m(\theta) - m(X)\big)' W \big(m(\theta) - m(X)\big)$$

where W is the optimal weighting matrix, defined to be the inverse of the covariance matrix of the moments. I obtain the covariance matrix via a clustered bootstrap, allowing for correlation within firms. I estimate the parameter vector via particle swarm, a standard stochastic global optimization solver.

The limiting distribution of the estimated parameter vector $\widehat{\theta}$ is

$$\sqrt{N}(\widehat{\theta} - \theta) \xrightarrow{d} N(0, \Sigma)$$

where

$$\Sigma = \left(1 + \frac{1}{S}\right)\left[\left(\frac{\partial m(\theta)}{\partial \theta}\right)' W \left(\frac{\partial m(\theta)}{\partial \theta}\right)\right]^{-1}$$

and S is the ratio of the number of observations in the simulated data to the number of observations in the sample.[22] I obtain the standard errors by computing the secant approximation to the partial derivative of the simulated moment vector with respect to the parameter vector. Given the

[22] S is set to be approximately 8.

discontinuities induced by the discretized state space, I select a step size of 1%.

C.2 Identification

While the targeted moments jointly determine the parameters to be estimated, there are nonetheless moments that are especially important for pinning down certain parameters. I discuss their informativeness in turn.

Idiosyncratic productivity persistence mostly informs the covariance between log inventory and log sales. Moreover, the dispersion of idiosyncratic productivity shocks mostly affects variances.

The fixed order costs are strongly related to the mean inventory-to-sales ratios. An increase in the upper bound of fixed order costs for non-JIT adopters leads to an increase in inventory-to-sales for non-JIT producers as these firms stock more inventories in order to incur the higher fixed cost less frequently. On the other hand, an increase in the fixed order cost upper bound for JIT firms leads to a decrease in the inventory-to-sales ratio for JIT firms. When fixed order costs rise among JIT firms, less productive JIT producers will abandon JIT and transition to being non-JIT firms. The remaining JIT producers are leaner and more productive than the firms that switched out of JIT. As a result, the average inventory-to-sales ratio declines among the pool of continuing JIT producers.

An increase in the sunk cost of adoption leads to less mode switching between JIT and non-JIT. This is because only relatively more productive firms will choose to select into JIT when the sunk cost is larger. At the same time, when deciding whether to abandon JIT, producers will recognize that they must pay a higher sunk cost to re-adopt JIT in the future. As a result, the firm is less likely to switch out of JIT today.

An increase in the continuation cost of adoption, on the other hand, causes the frequency of adoption and the probability of switching out of JIT to move in opposite directions. A higher continuation cost of adoption makes JIT more costly to maintain. As a result, the probability of switching out of adoption increases. At the same time, this makes non-adopters less likely to initiate JIT given that these higher continuation costs would have to be incurred in the future.

Finally, the storage cost also affects the distribution of inventory-to-sales ratios. A higher storage cost raises the marginal cost of carrying inventories across time, and therefore reduces average inventory-to-sales ratios across JIT and non-JIT firms alike.

Figure C1: Monotonic Relationships

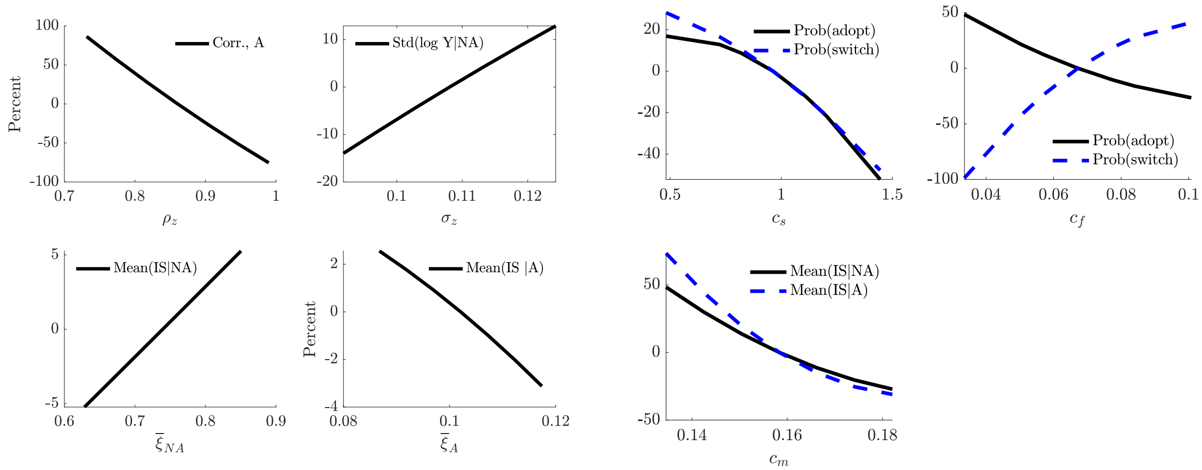

Note: The figure plots the changes in selected moments to changes in the parameters, in percent relative to moment at estimated parameter values.

The ten moments jointly determine the seven parameters that reside in vector θ. Figure C1 reports the monotone relationships between selected moments and parameters. Figure C2 reports the sensitivity of each of the seven parameters to changes in each of the moments. These results come from an implementation of Andrews et al. (2017). The sensitivity of $\widehat{\theta}$ to $m(\theta)$ is

$$\Lambda = -\left[\left(\frac{\partial m(\theta)}{\partial \theta}\right)' W \left(\frac{\partial m(\theta)}{\partial \theta}\right)\right]^{-1} \left(\frac{\partial m(\theta)}{\partial \theta}\right)' W$$

I then transform this matrix so that the coefficients reflect the effect on each parameter of a one standard deviation change in the respective moments.

Figure C2: Sensitivity

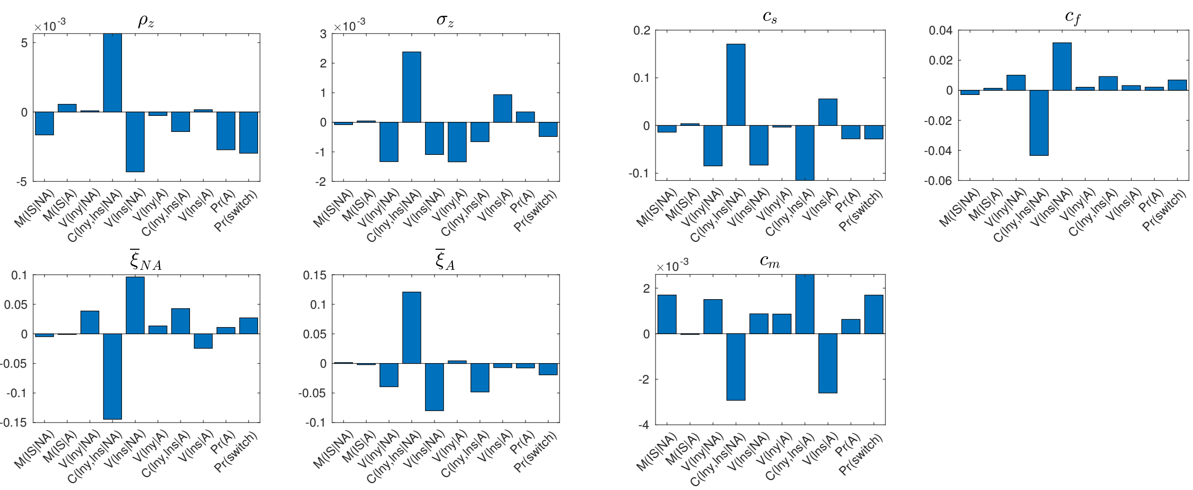

Note: The figure plots sensitivity estimates as in Andrews et al. (2017). These estimates describe the changes in each of the seven parameters to a one standard deviation increase in each moment.

Appendix D Robustness

In this section I further examine the sensitivity of JIT to supply disruptions by exploring the sensitivity of the baseline economy to a fixed order cost shock based on alternative parameterizations, and under aggregate uncertainty.

D.1 Sensitivity Analysis

Table D1: Alternative Parameterizations

Description	Parameter	Value	Value
Idiosyncratic productivity persistence	ρ_z	0.82	0.88
Idiosyncratic productivity volatility	σ_z	0.08	0.12
Order cost upper bound (non-adopters)	$\overline{\xi}_{NA}$	0.69	0.78
Order cost upper bound (adopters)	$\overline{\xi}_{A}$	0.05	0.15
Sunk cost of adoption	c_s	0.91	1.01
Continuation cost of adoption (adopters)	c_f	0.05	0.07
Carrying cost	c_m	0.12	0.21

Note: The table reports the alternative parameterizations chosen to compute the excess sensitivity to supply disruptions associated with JIT.

I start by analyzing the robustness of the unanticipated shock exercise in the main text to different parameterizations. Table D1 reports a number of different parameter specifications. I vary all parameters in different directions. Figure D1 plots the excess contraction amid a supply disruption between the JIT and counterfactual economy. Across all specifications, output contracts more sharply in the JIT economy than in the counterfactual, as shown by the dots residing below the 45-degree line.

Figure D1: JIT Vulnerabilities with Alternative Parameters

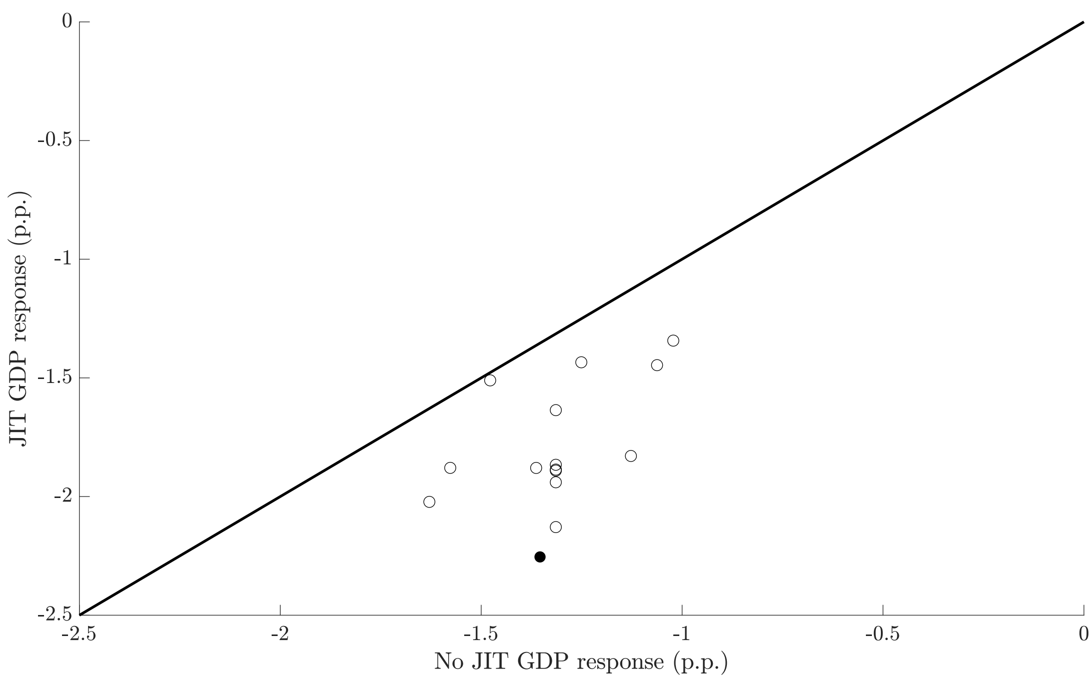

Note: The figure plots the GDP contraction in response to an unanticipated fixed order cost shock in the baseline economy (vertical axis) against the GDP contraction in the counterfactual model (horizontal axis) for a variety of different parameter specifications. The black dot denotes the baseline parameterization estimated in the main text, and the hollow dots represent the alternative parameterizations. The different parameterizations are detailed in Table D1.

D.2 Incorporating Aggregate Uncertainty

The supply disruption that I model in the main text takes the form of an unanticipated fixed order cost shock. After the realization of the shock, agents have perfect foresight about the transition back to steady state. In this section, I show that my findings extend to an environment in which agents form expectations over an aggregate shock to fixed order costs.

Aggregate Uncertainty About Fixed Order Costs

I assume that the support of fixed order costs is time varying,

$$\underline{\xi}_t = \underline{\xi} + x_t \quad \text{and} \quad \overline{\xi}_t = \overline{\xi} + x_t.$$

In addition, I assume that x_t is a two-dimensional state, $x_t = [0.0 \;\; 0.1]$, where $x_t = 0.0$ reflects normal times and $x_t = 0.1$ reflects a supply disruption. I set the transition matrix to be,

$$\Pi(x'|x) = \begin{bmatrix} 0.95 & 0.05 \\ 0.25 & 0.75 \end{bmatrix},$$

which implies that supply disruptions are infrequent and that a transition from a supply disruption back to normal times is relatively quick. The aggregate state space is now comprised of x and μ, the distribution of firms, and we can denote it as $\Psi = (x, \mu)$.

Solving the JIT model with aggregate shocks requires tracking prices and the distribution of firms, an infinite-dimensional object, across time. Following Krusell and Smith (1998), I solve the model by assuming that agents exhibit bounded rationality and use aggregate material usage to track the distribution across time. I define a log linear mapping between prices and aggregate materials,

$$\log \widehat{M}'(x) = \beta_0^M(x) + \beta_1^M(x) \log M(x)$$
$$\log \widehat{p}(x) = \beta_0^p(x) + \beta_1^p(x) \log M(x).$$

Below I summarize the steps taken to solve the model:

1. Simulate $\{X_t\}_{t=1}^T$ for some large T and guess an initial set of coefficients: $\beta_0^{M(0)}$, $\beta_1^{M(0)}$, $\beta_0^{p(0)}$, and $\beta_1^{p(0)}$.

2. For each iteration i:

   (a) Solve the final goods firm problem on a grid based on $\widehat{p}^{(i)}$ and $\widehat{M}'^{(i)}$, where $\widehat{p}^{(i)}$ and $\widehat{M}'^{(i)}$ are obtained using the coefficients $\beta_0^{M(i)}, \beta_1^{M(i)}, \beta_0^{p(i)}, \beta_1^{p(i)}$.

   (b) Simulate the model, using $\{X_t\}_{t=1}^T$, and obtain a time series $\{p_t, M_t\}_{t=1}^T$, where p_t reflects the equilibrium price which is obtained by clearing markets in each period (not by using the pricing rule, $\widehat{p}$).

   (c) Based on the simulated data, update the forecast rules via OLS to obtain $\beta_0^{M(i+1)}$, $\beta_1^{M(i+1)}$, $\beta_0^{p(i+1)}$, and $\beta_1^{p(i+1)}$. If the coefficients are sufficiently close, up to a pre-specified tolerance, then exit. Otherwise, update the each coefficient as a convex combination of the old guess and the new estimate. Set $i = i + 1$ and return to (a).

I specify a grid of dimension $N_z \times N_s \times N_a \times N_X \times N_M = 5 \times 60 \times 2 \times 2 \times 10$, and I simulate the model for 5,500 periods and discard the first 500 periods to reduce the influence of initial conditions. I specify a tolerance of 10^{-3} for the forecasting rules to converge.

Solution and Accuracy Statistics

The converged forecast rules are reported in Table D2. Table D3 reports accuracy statistics for each model based on static and dynamic forecasts for aggregate materials and prices. The first row of the table reports the mean percentage difference between realized simulated data and its corresponding dynamic forecast based on the forecasting rules (Den Haan, 2010). The bottom two rows report the R^2 of each forecast rule based on static forecasts.

Table D2: Forecasting Rules

	Baseline	Counterfactual
$\beta_0^M(x = x_1)$	-2.661	-3.176
$\beta_0^M(x = x_2)$	-0.054	-0.247
$\beta_1^M(x = x_1)$	-2.339	-2.577
$\beta_1^M(x = x_2)$	0.089	1.5e-10
$\beta_0^p(x = x_1)$	1.449	1.836
$\beta_0^p(x = x_2)$	-0.204	-0.056
$\beta_1^p(x = x_1)$	1.324	1.678
$\beta_1^p(x = x_2)$	-0.262	-0.129

Note: The table reports the forecasting rules.

Table D3: Accuracy Statistics

	Baseline		Counterfactual	
	M	p	M	p
Mean percentage difference	0.286	0.195	0.277	0.029
Forecast regression R^2				
$x = x_1$	1.000	1.000	1.000	1.000
$x = x_2$	1.000	1.000	1.000	1.000

Note: The table reports accuracy statistics.

Impulse Response to a Fixed Order Cost Shock

To compute the impulse response to a shock to x_t, I follow the Koop et al. (1996) approach and compute the generalized impulse response. More specifically, I simulate 5,000 economies for 50 periods. I do this twice: once for a set of simulations that does not impose the shock, and again for simulations which do impose the shock. I then compute the impulse response for GDP as,[23]

$$\text{GDP}_t^{\text{IRF}} = \frac{1}{5000}\sum_{j=1}^{5000}\left[\log\left(\frac{\text{GDP}_{jt}^{\text{shock}}}{\text{GDP}_{jt}^{\text{no shock}}}\right)\times 100\right].$$

Figure D2 plots the impulse response and cumulative impulse response in the baseline and counterfactual economies. The left panel shows that the counterfactual no JIT economy contracts more than the baseline JIT economy on impact. This is in contrast to the unanticipated shock exercise and has to do with the fact, under aggregate uncertainty, ordering probabilities are less sensitive to the supply disruption. The reason that ordering probabilities become less sensitive when we introduce aggregate uncertainty about fixed order costs is because firms are able to build precautionary stocks of inventories. Indeed, aggregate inventories are about 11% higher in an unconditional simulation of the baseline model with aggregate uncertainty relative to the no aggregate uncertainty steady state baseline model.

However, as the shock unwinds the baseline economy indeed contracts more than the counterfactual economy based on the cumulative impulse responses plotted in the right panel of Figure D2. This is because, while orders are less sensitive to the supply disruption under aggregate uncertainty, orders still contract. Since the baseline economy still features lower inventory stocks than the counterfactual economy, firms in the baseline economy still have fewer inventories to draw from when producing and therefore still experience a stronger and more protracted contraction in material input usage relative to the counterfactual economy.

Overall, the findings documented in this section are consistent those obtained from the unanticipated shock exercise and indicate that the vulnerability of JIT producers to supply disruptions is

[23] See Koop et al. (1996) for more details. Terry (2017) also provides a detailed summary of this approach.

Figure D2: Output Response to Fixed Order Cost Shock

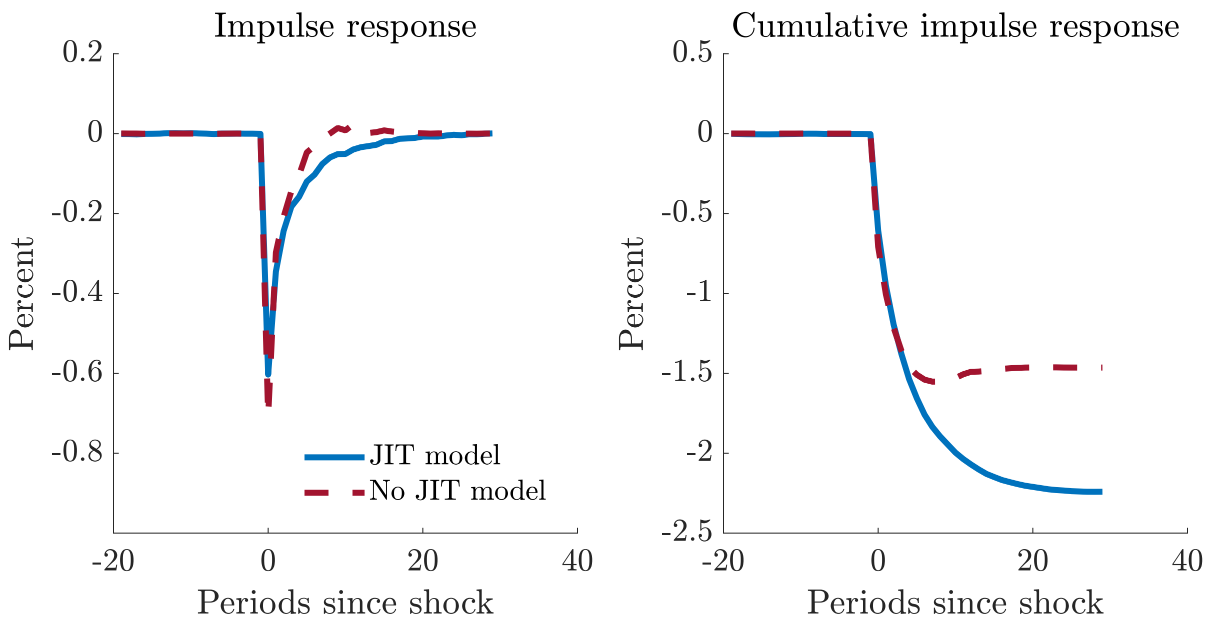

Note: The figure depicts the response of output to a positive shock to fixed order costs, averaged over 5,000 simulated economies.

also present in an environment that features aggregate uncertainty.

Figure D3 plots the cumulative impulse response of other endogenous aggregates in the model, mirroring the endogenous responses plotted in Figure 7.

Figure D3: Cumulative Impulse Response of Other Aggregates to Fixed Order Cost Shock

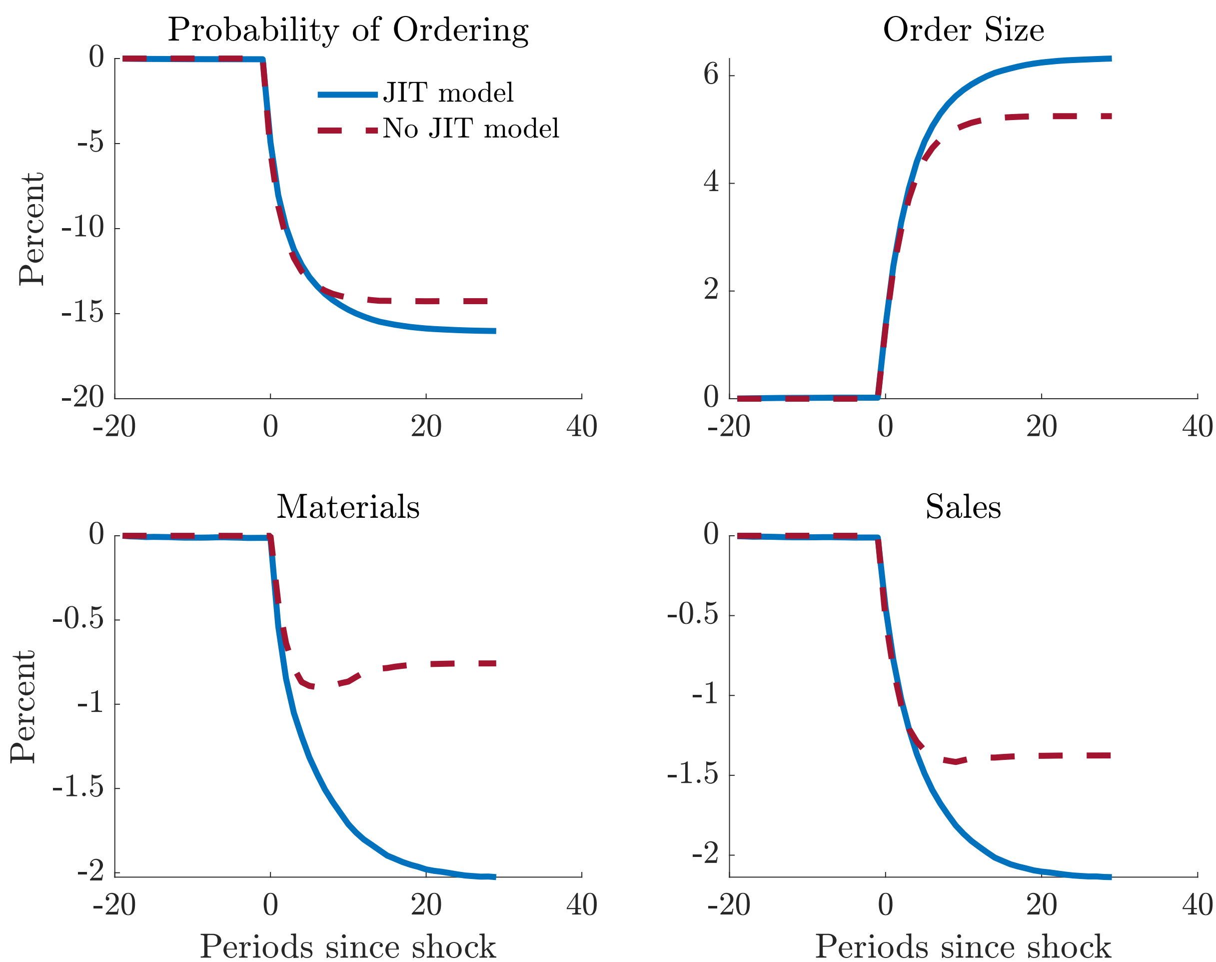

Note: The figure depicts the response of the probability of placing an order, order size, materials, and sales to a positive shock to fixed order costs, averaged over 5,000 simulated economies.