SUPPLEMENTARY MATERIALS FOR “SPREAD TOO THIN: THE IMPACT OF LEAN INVENTORIES”

Appendix A Empirics

This section includes further details on the JIT adopters sample, an alternative measure of JIT, and estimates of JIT firms' sensitivity to weather events faced by their suppliers.

A.1 Sample Construction

My data come from three sources. First, I use annual Compustat data to obtain information on firm-level inventory holdings, sales, and other outcomes. Second, I use the GSCPI in the regressions that estimate the JIT sensitivity to aggregate supply conditions. Lastly, for the weather regressions reported in Appendix A.3, I collect county-level weather event data from NOAA and map weather events to firm headquarter zip codes.

My compustat sample covers the years 1971-2019. I keep only manufacturers (two-digit NAICS codes 31, 32, and 33). In addition, I drop firm-years in which acquisitions exceed 5% of total assets to avoid the influence of large mergers. To guard against 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.

As discussed in the text, I define input inventories as the sum of raw material and work-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.

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

In this section, I explore the robustness of my measure of JIT to alternative assumptions. I then explore an alternative measure of JIT that identifies the adoption of lean production based on fi-

Table A1: Compustat Summary Statistics

Variable name	Compustat code definition	Mean	Median	Standard deviation	25%	75%
Earnings growth	$\frac{\Delta \mathtt{ib}}{\mathtt{ib}_{-1}}$	-0.285	-0.069	6.083	-0.805	0.449
Employment growth	$\frac{\Delta \mathtt{emp}_t}{\mathtt{emp}_{t-1}}$	-0.001	0.000	0.294	-0.083	0.100
IHS earnings	$ihs(\mathtt{ib}_t)$	0.839	0.850	3.231	-1.578	3.201
Log sales	$\log(\mathtt{sale}_t)$	4.381	4.373	2.338	2.824	5.977
Sales growth	$\frac{\Delta \mathtt{sale}_t}{\mathtt{sale}_{t-1}}$	0.067	0.055	0.413	-0.062	0.178
Log sales per worker	$\frac{\mathtt{sale}_t}{\mathtt{emp}_t}$	4.978	4.982	0.868	4.454	5.516
JIT adoption		0.676	1.000	0.468	0.000	1.000

Note: The table reports summary statistics for the relevant variables in Section 2 of the text. The sample is constructed from Compustat Fundamentals Annual files for 1980-2019. Sample consists of 5,912 unique firms.

nancial 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.

A.2.1 Measuring JIT Under Alternative Assumptions

In this section, I consider alternative assumptions in the construction of my JIT measure. First, I define JIT based on a comparison of inventory-to-sales ratios relative to pre-1980s mean inventory-to-sales ratios within sectors rather than median inventory-to-sales ratios. Second, I measure JIT by comparing inventory-to-sales to historical inventory-to-sales within a sector from 1971-1985 rather than 1971-1980. Third, I define JIT by comparing inventory-to-sales to historical inventory-to-sales within a narrower 4-digit NAICS industry rather than 2-digit NAICS sector. Fourth, I define JIT based on total inventories (raw material, work-in-process, and finished goods) rather than only input inventories (raw material and work-in-process). Fifth, I define JIT based on real inventory-to-sales ratios rather than nominal inventory-to-sales ratios, using manufacturing inventory and sales deflators from the National Income and Product Accounts (NIPA) tables. Figure A1 plots the evolution of the share of measured JIT adopters in my sample according to the baseline definition as well as these alternative definitions. Overall, these various measures of JIT evolve similarly over time, with my baseline definition generally residing in the middle of the range of these various measures.

Figure A1: Alternative Measures of JIT Adoption

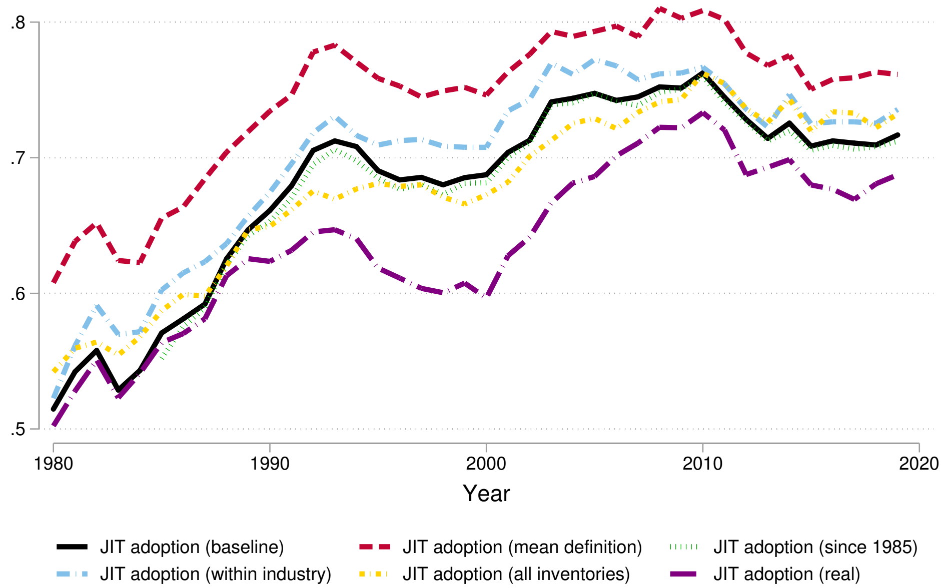

Note: The figure plots alternative measures of the share of JIT adopters in the sample based on variations to the definition in equation (1) of the main text.

A.2.2 A Text-Based Measure of JIT

I next examine a measure of JIT that identifies adoption of lean production based on financial news reports, press releases, and Form 10-K filings for a narrower set of manufacturing firms.

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

Figure A2: Validation of JIT Indicator

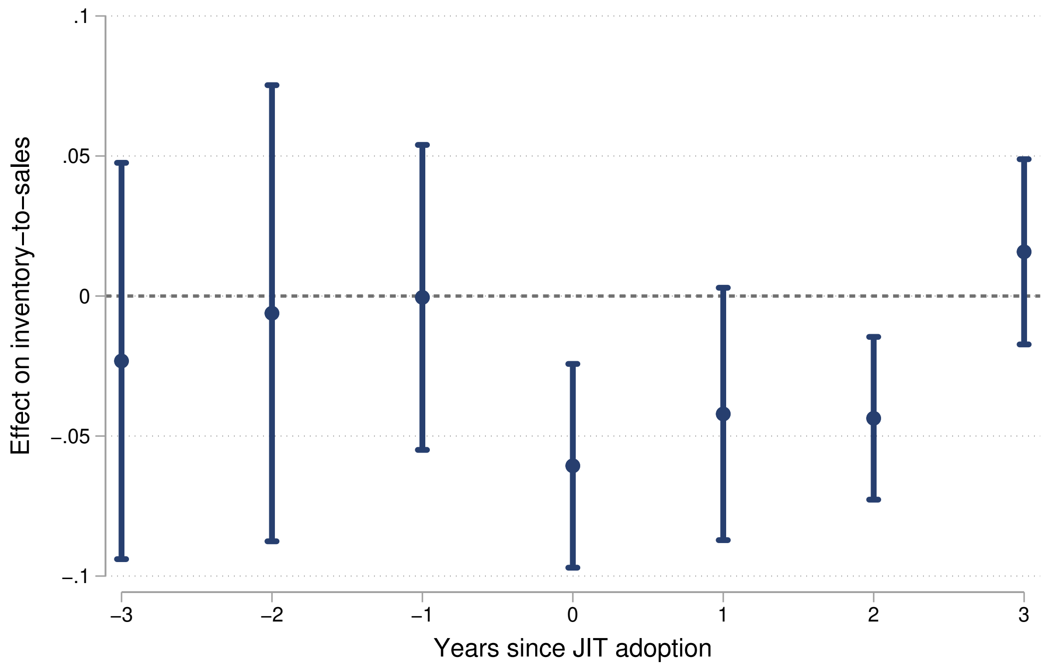

Note: The figure plots the estimated effect of JIT adoption on the level of inventory-to-sales. 95% confidence bands are displayed alongside point estimates.

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.

Next, I compare the two measures of JIT adoption. Table A2 reports the share of JIT firms in a given 3-digit NAICS industry (and identified through the text-based approach). Based on these reported frequencies, we see that most of the 177 JIT firms identified through the text-based approach reside in NAICS sector 33 and, specifically, within fabricated metal manufacturing, machinery manufacturing, computer and electronic product manufacturing, electrical equipment, appliance, and component manufacturing, and transportation equipment manufacturing. The final column of Table A2 reports each industry's average inventory-to-sales ratio from 1980 to 2019 relative to its pre-1980 median. Across all industries, we see a decline in the inventory-to-sales ratio relative to pre-1980, albeit to varying degrees.

Table A2: Text-Based JIT Adopters by Industry

Description	3-digit NAICS	Percent of text-based adopters	Inventory-to-sales relative to pre-1980s medians (percent)
Food manufacturing	311	1.96	-6.76
Textile mills	313	1.31	-0.05
Textile product mills	314	0.65	-0.78
Leather and allied product manufacturing	316	1.31	-4.93
Wood product manufacturing	321	0.65	-0.74
Paper manufacturing	322	1.31	-2.97
Printing and related support activities	323	0.65	-5.61
Chemical manufacturing	325	2.61	-1.74
Plastics and rubber manufacturing	326	1.96	-4.05
Nonmetallic mineral product manufacturing	327	0.65	-2.73
Primary metal manufacturing	331	3.27	-0.03
Fabricated metal product manufacturing	332	7.19	-7.15
Machinery manufacturing	333	15.03	-6.42
Computer and electronic product manufacturing	334	35.29	-5.36
Electrical equipment, appliance, and component manufacturing	335	7.19	-5.19
Transportation equipment manufacturing	336	11.76	-5.78
Furniture and related product manufacturing	337	3.27	-9.39
Miscellaneous manufacturing	339	3.92	-3.34

Note: The table lists 3-digit NAICS industry codes and descriptions as well as the share of JIT of firms identified through the text-based approach. The final column reports the average inventory-to-sales ratio in each industry relative to the pre-1980 median.

Figure A3 pools across the text-based JIT adopters and plots their average inventory-to-sales ratio, in each year, relative to their respective pre-1980s sector medians. The figure indicates that on average inventory-to-sales ratios among text-based JIT firms declined relative to their respective pre-1980 median inventory-to-sales ratios.

Figure A3: Inventory-to-sales Among Text-Based JIT Adopters

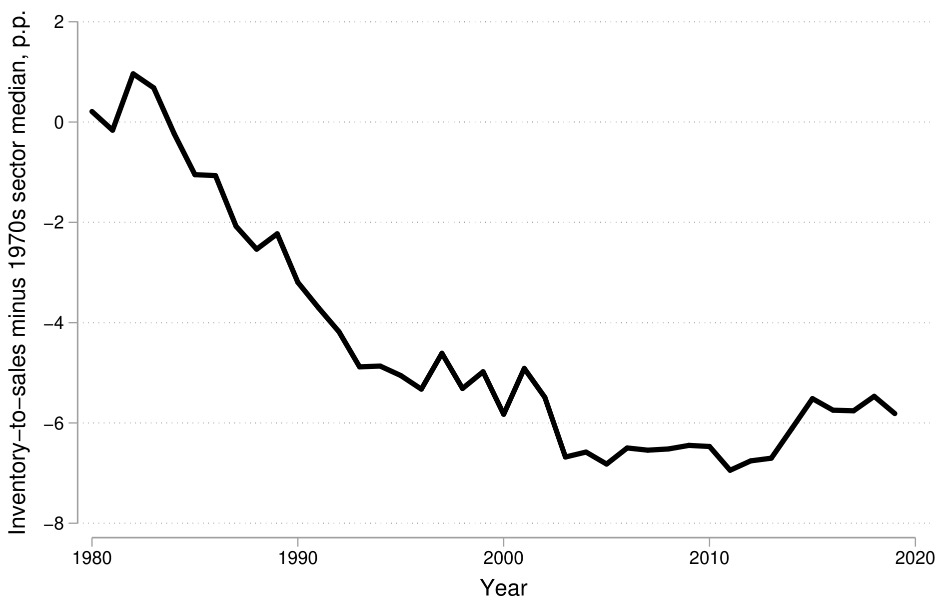

Note: The figure plots inventory-to-sales ratio of text-based JIT adopters relative to their respective 1970s sector median inventory-to-sales ratios.

A.3 JIT and Weather Events

In addition to being more sensitive to aggregate supply conditions, I find that JIT adopters appear to be 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:

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

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. 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 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 A3 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

Table A3: 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 (1). 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.

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.[1] Nonetheless, in my preferred specifications, I find that, on average, supplier weather events in my sample predict an additional 19.5% decline in JIT firm sales. Furthermore, 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.

[1] 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 status.

Appendix B Model

In this section, I first provide details on how the final goods firm problem is reduced from a choice over inventory, s' and labor, n, to only a choice over inventory. Second, I report the expression for the price of the intermediate good, q. Finally, I detail the model's numerical solution.

B.1 Reducing the Final Goods Firm Problem to a Choice Over Inventories

Final goods firms choose s' and n to maximize their value in the production stage. We can optimize out the static labor decision. The firm's payoff in the current period solves,

$$\max_{s' \in [0,s], n} z n^{\theta_n} (s - s')^{\theta_m} - c_m s'.$$

Solving for labor, 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, equivalently,

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

B.2 Intermediate Goods Price

The intermediate goods firm problem is:

$$\max_{K,L} q K^{\alpha} L^{1-\alpha} - RK - wL.$$

The solution to this problem yields the following closed-form expression for the price of the intermediate good,

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

B.3 Numerical Solution

The model is solved using methods that are standard in the heterogeneous firms literature. The exogenous productivity process is discretized according to 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 state. Considering the binary adoption state, this implies that the discretized model has 4,400 grid points.

I can simplify the model by defining $p \equiv U_1(C, H) = \frac{1}{C}$ and reformulating the final goods firms' problem as the maximization of dividends weighted by marginal utility price, p. In doing so, firms weight current profits by p and discount expected future earnings by β (eliminating the time-varying discount factor, Λ_{t+1}). Defining the value function $v^A = pV^A$, the firm's reformulated problem is,

Stage 1: Adoption Decision

$$v^A(z,s,a) = \max\left\{ -pwc(a) + \int v^O(z,s,1,\xi)dF(\xi_A), \int v^O(z,s,0,\xi)dF(\xi_{NA})\right\}, \tag{2}$$

where

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

Stage 2: Order Decision

$$v^O(z,s,a,\xi) = \max\left\{ -pw\xi + v^*(z,s,a), v^P(z,s,a)\right\}, \tag{3}$$

where the value of placing an order is

$$v^*(z,s,a) = \max_{s^* \geqslant s} \Big[-pq(s^*-s) + v^P(z,s^*,a) \Big], \tag{4}$$

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

$$\xi^*(z,s,a) = \min\Big(\max\big(\underline{\xi}, \widetilde{\xi}(z,s,a)\big), \overline{\xi}\Big), \tag{5}$$

and

$$\widetilde{\xi}(z,s,a) = \frac{v^*(z,s,a) - v^P(z,s,a)}{pw}. \tag{6}$$

Stage 3: Production Decision

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

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:

$$v^P(z,\widetilde{s},a) = \max_{s' \in [0,\widetilde{s}]} p\Big[zn(z,\widetilde{s},s',a)^{\theta_n}(\widetilde{s}-s')^{\theta_m} - wn(z,\widetilde{s},s',a) - c_m s' \Big] + \beta\mathbb{E}\big[v^A(z',s',a')\big]. \tag{7}$$

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 that produce a better update of the value function. The ergodic distribution of firms is obtained via nonstochastic simulation as in Young (2010).

I solve the model by initiating a guess for p, p_0. Using the household and intermediate goods producer's optimality conditions, I obtain the implied wage and intermediate goods 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 via nonstochastic simulation. 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 using 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 optimal decisions, I push the distribution of final goods firms forward beginning in the steady state. 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, $\frac{1}{p_t} = C_t$.

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 $\theta = \begin{pmatrix}\rho_z & \sigma_z & \overline{\xi}_{NA} & \overline{\xi}_A & c_s & c_f & c_m\end{pmatrix}'$. 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 vector $\widehat{\theta}$ 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 as the inverse of the covariance matrix of the moments. I compute the covariance matrix, clustering by firm following Hansen and Lee (2019). I estimate the parameter vector via particle swarm optimization, 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.[2] I obtain the standard errors of the parameters from Σ, where I compute the Jacobian, $\frac{\partial m(\theta)}{\partial \theta}$, via numerical differentiation.

[2] S is set to be approximately 6.5.

Given the lumpiness in ordering and inventories induced by the fixed ordering costs, the heterogeneous producers in my model are best characterized as establishments. Accordingly, when matching the model moments to the data, a simulated firm is assumed to consist of ten establishments. I classify a simulated firm as JIT according to my empirical definition of JIT, using the pre-1980 average inventory-to-sales ratio across all manufacturing sectors (approximately 0.13).

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. Similarly, an increase in the fixed order cost upper bound for JIT firms leads to an increase in the inventory-to-sales ratio for JIT firms.

An increase in the sunk cost of adoption leads to a decrease in the share of JIT adopters. Moreover, an increase in the continuation cost of adoption, causes the share of firms switching out of JIT to rise. 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.

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

Figure C1: Monotonic Relationships

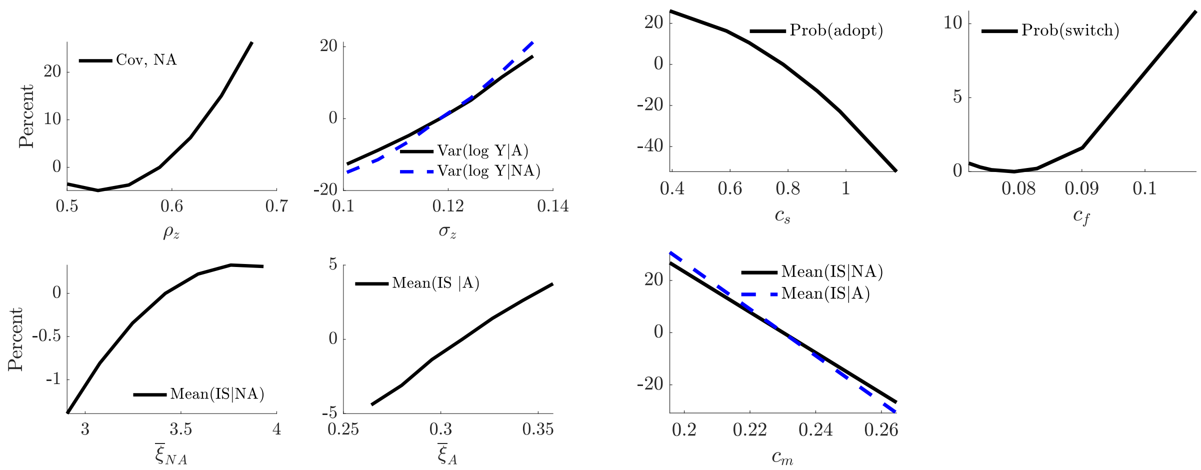

Note: The figure plots the changes in selected moments to changes in the parameters, in percent relative to the moment at the estimated parameter values.

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

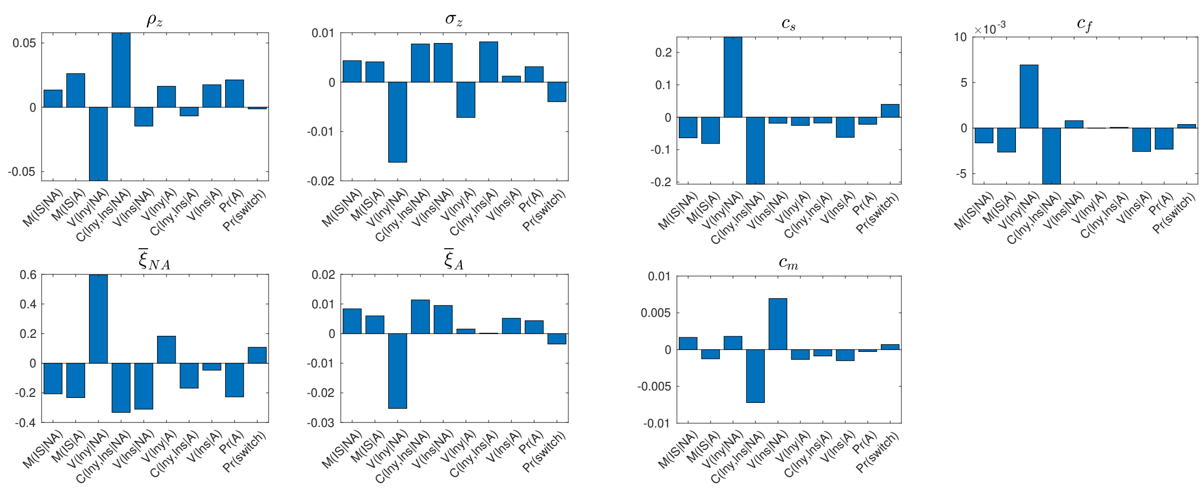

Note: The figure plots sensitivity estimates as in Andrews et al. (2017). These estimates describe the changes in each of the seven parameters to a one standard deviation increase in each moment.

Appendix D Robustness

In this section, I explore the sensitivity of the baseline economy to a fixed order cost shock based on alternative parameterizations, shock sizes, assumptions on the fixed order cost distribution, and under aggregate uncertainty.

D.1 Sensitivity Analysis

I start by analyzing the robustness of the unanticipated shock exercise in the main text to different parameterizations. I separately vary each estimated parameter by 5% in either direction, solve for the steady state of the JIT and no JIT economies, and solve for the transition path amid an unanticipated increase in fixed order costs. Figure D1 plots the total output contraction along the transition in each 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

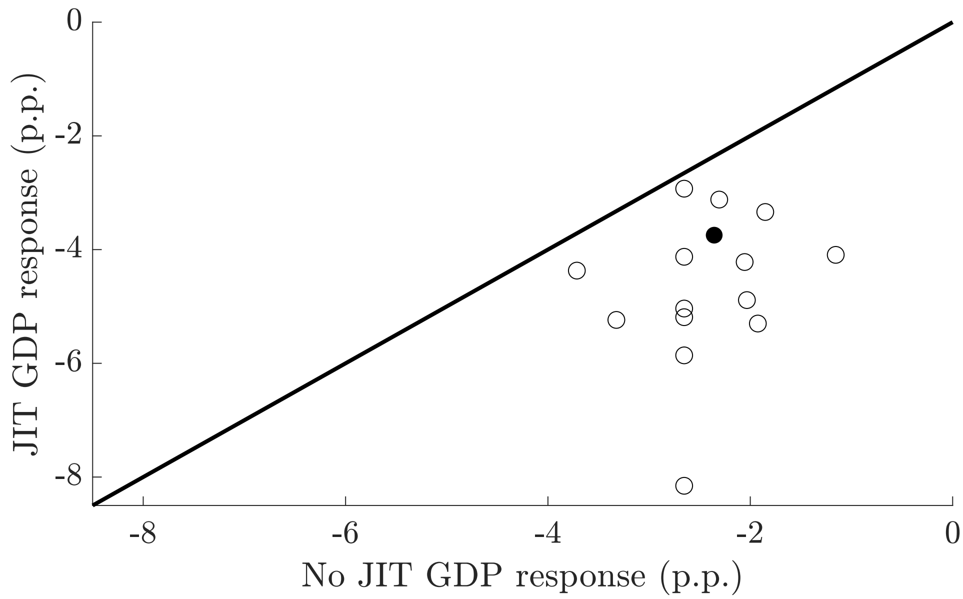

Note: The figure plots the GDP contraction in response to an unanticipated fixed order cost shock in the baseline economy (vertical axis) against the GDP contraction in the counterfactual model (horizontal axis) for a variety of different parameter specifications. The black dot denotes the baseline parameterization estimated in the main text, and the hollow dots represent the alternative parameterizations which reflect 5% deviations in each estimated parameter value in either direction.

D.2 Shock Size

In this section, I consider alternative shock sizes for the unanticipated increase in fixed ordering costs. Figure D2 plots the output dynamics when the shock is one half of the size specified in the main text. Figure D3 plots the output dynamics when the shock is twice the size specified in the main text.

Figure D2: Output Response to Smaller Fixed Order Cost Shock

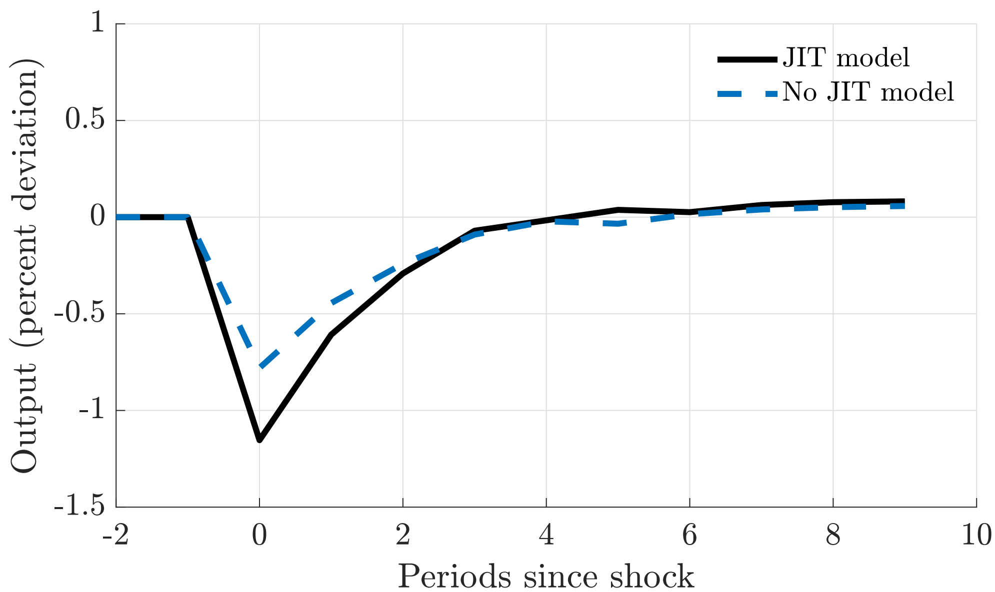

Note: The figure plots the output response to a fixed order cost shock that is half the size of the shock in Figure 4 of the text. The persistence of the shock is set to 0.50.

Figure D3: Output Response to Larger Fixed Order Cost Shock

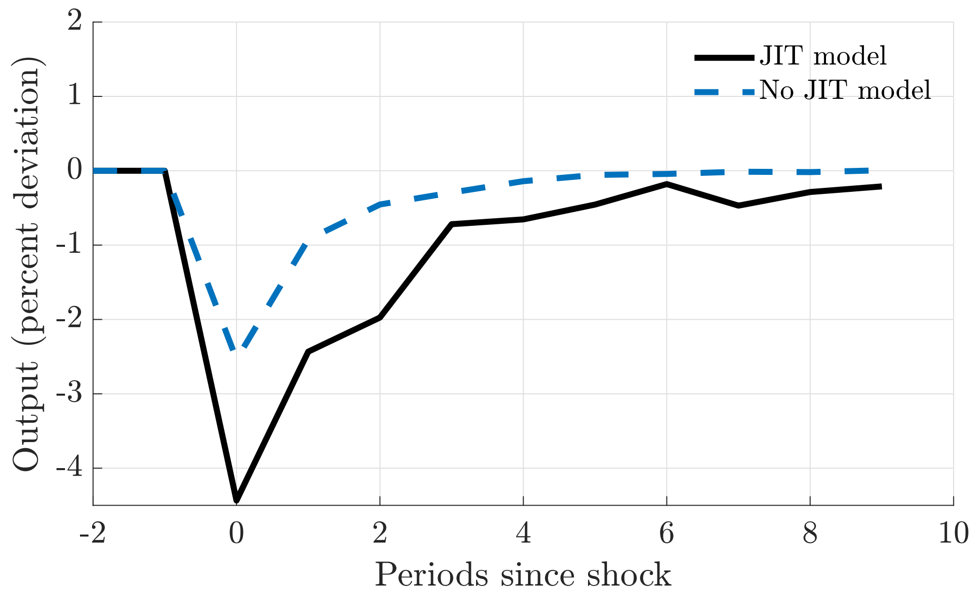

Note: The figure plots the output response to a fixed order cost shock that is twice the size of the shock in Figure 4 of the text. The persistence of the shock is set to 0.50.

D.3 Alternative Order Cost Distribution

In this section, I repeat my analysis under an alternative fixed ordering cost distribution. I assume that fixed order costs follow a four-parameter Beta distribution, $\mathcal{B}(\alpha, \beta, \underline{\xi}, \overline{\xi})$, where α and β are the shape parameters of the Beta distribution, and $\underline{\xi}$ and $\overline{\xi}$ are the lower and upper support. The probability density function is:

$$f(\xi) = \frac{(\xi - \underline{\xi})^{\alpha-1}(\overline{\xi} - \xi)^{\beta-1}}{(\overline{\xi} - \underline{\xi})^{\alpha+\beta-1}B(\alpha, \beta)},$$

where $B(\alpha, \beta)$ is the Beta function, $B(\alpha, \beta) = \int_0^1 t^{\alpha-1}(1-t)^{\beta-1}dt$.

I consider a version of the model with a Beta(3,3,$\underline{\xi}, \overline{\xi}$) order cost distribution, meaning that order costs are symmetrically distributed. Note that the support, $[\underline{\xi}, \overline{\xi}]$ differs for JIT and non-JIT firms, as in the uniform order cost case. Furthermore, as in the case of uniform order costs, I fix $\underline{\xi} = 0$ in the steady state.

First, I find that the economy-wide equilibrium stock of inventories is lower under Beta(3,3) distributed order costs, consistent with the notion that uniformly distributed order costs generate a relatively stronger precautionary inventory holding motive. The first row of Table D1 reports the aggregate inventory stock in Beta order cost model relative to the baseline model. The second row reports the relative output in the JIT model compared to the counterfactual model prior to the shock.

Table D1: Aggregate Outcomes Under Beta-Distributed Order Cost

	Uniform	Beta(3,3)
Steady state inventory stock relative to uniform (%)	–	-23.88
Output relative to no JIT counterfactual before shock (%)	15.28	22.17

Note: The first row of the table reports the steady state inventory stock relative to the baseline assumption of uniformly distributed fixed order costs. The second row of the table reports the output prior to the shock in the JIT economy relative to output prior to the shock in the analogous no-JIT economy.

Figures D4 and D5 plot the transition dynamics under uniform and Beta(3,3) distributed fixed order costs. Overall, the results are quite similar across order cost specifications. The JIT economy under Beta(3,3) experiences a slightly larger output contraction amid a shock to fixed order costs relative to the version with uniformly distributed fixed order costs.

Figure D4: Output Response with Uniformly Distributed Order Costs

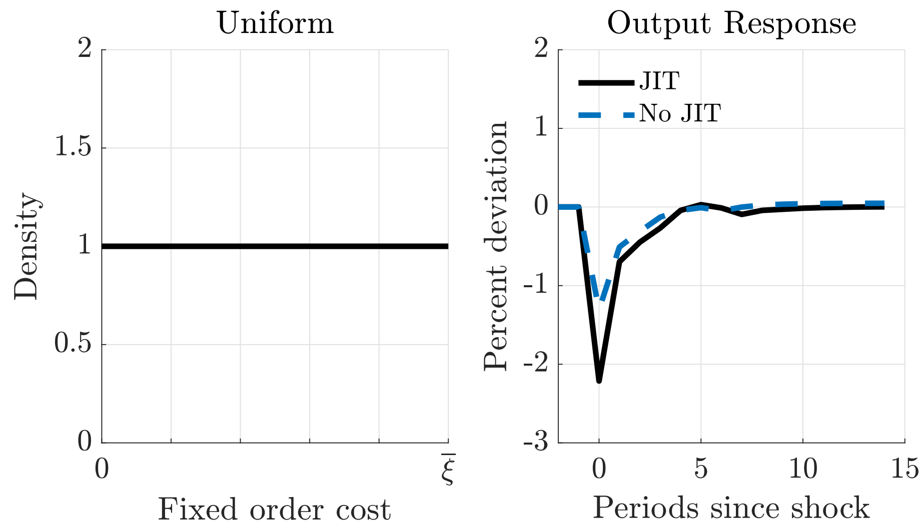

Note: The left panel plots the assumed pdf of fixed ordering cost in the model. The right panel plots the output response to a fixed order cost shock, as in Figure 4 of the main text.

Figure D5: Output Response with Beta(3,3) Distributed Order Costs

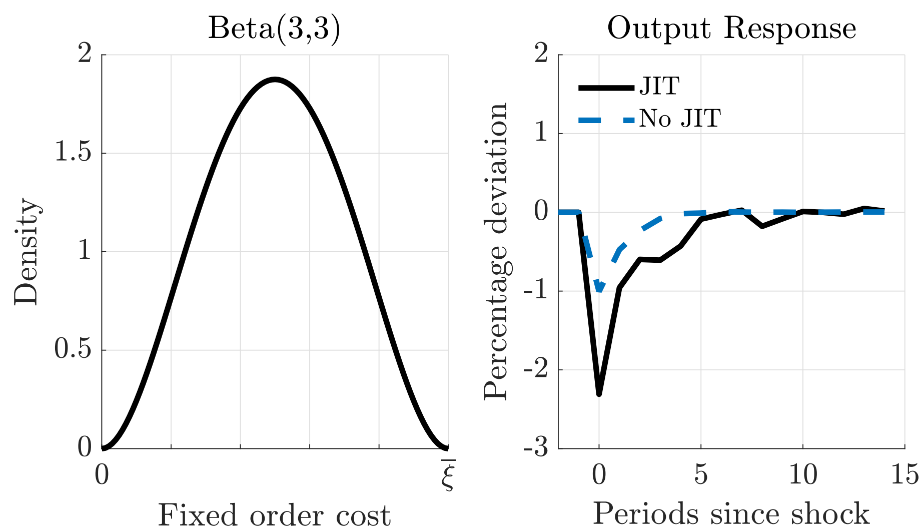

Note: The left panel plots the assumed pdf of fixed ordering costs in the model. The right panel plots the output response to a fixed order cost shock. The shock is equal in size and persistence to the one used in Figure 4 of the main text.

D.4 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 generally 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, $x_t = [0.0 \;\; 0.054]$, where $x_t = 0.0$ reflects normal times and $x_t = 0.054$ reflects a supply disruption, which is calibrated to be the same as in the unanticipated shock exercise. 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 input 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.

   (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 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 $N_z \times N_s \times N_a \times N_X \times N_M = 5 \times 120 \times 2 \times 2 \times 10$, and I simulate the model for 5,500 periods, discarding the first 500 periods to reduce the influence of initial conditions. I specify a tolerance of 10^{-3} for convergence of the forecasting rules.

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
$\beta_0^M(x = x_1)$	-1.536	-2.759
$\beta_0^M(x = x_2)$	-1.806	-2.695
$\beta_1^M(x = x_1)$	0.397	-0.041
$\beta_1^M(x = x_2)$	0.297	-0.010
$\beta_0^p(x = x_1)$	1.219	2.165
$\beta_0^p(x = x_2)$	1.578	2.210
$\beta_1^p(x = x_1)$	-0.373	-0.044
$\beta_1^p(x = x_2)$	-0.240	-0.031

Note: The table reports the forecasting rules.

Table D3: Accuracy Statistics

	Baseline		Counterfactual	
	M	p	M	p
Mean percentage difference	0.633	0.378	0.437	0.009
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,[3]

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

Figure D6 plots the impulse response in the baseline and counterfactual economies. The top left panel shows that output initially increases in the baseline economy relative to the counterfactual economy, but subsequently contracts further and recovers more gradually. The top right panel plots sales, which contracts more in the baseline model relative to the counterfactual without an initial relative increase.

The relative increase in output on impact is therefore due to a relative increase in inventory investment on impact which is ultimately attributed to the behavior of aggregate orders in the model. Under aggregate uncertainty, firms recognize the likelihood of transiting to the high fixed order cost state and optimally hold precautionary stocks of inventory in response. This, in turn, renders ordering probabilities less sensitive on impact to a fixed order cost shock.[4] The aggregate ordering probability declines by about 1.7% on impact but then falls further as inventories are depleted, reaching a trough of -2.6% one period after the shock. Meanwhile, order sizes increase on impact. Taken together, this yields a relative increase in aggregate orders in the period in which the aggregate shock is realized.

After the period in which the shock occurs, however, output declines by more in the baseline economy than in the counterfactual economy as inventory buffers are rundown and aggregate orders decline. Therefore, the findings documented in this section are broadly consistent those obtained

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

[4] One way to see this by noting that the slope of ordering probabilities in panel (a) of Figure 3 is flatter at higher levels of inventory. Though ordering probabilities still decline more in the baseline economy, they do so less than in the unanticipated shock exercise.

Figure D6: Output Response to Fixed Order Cost Shock

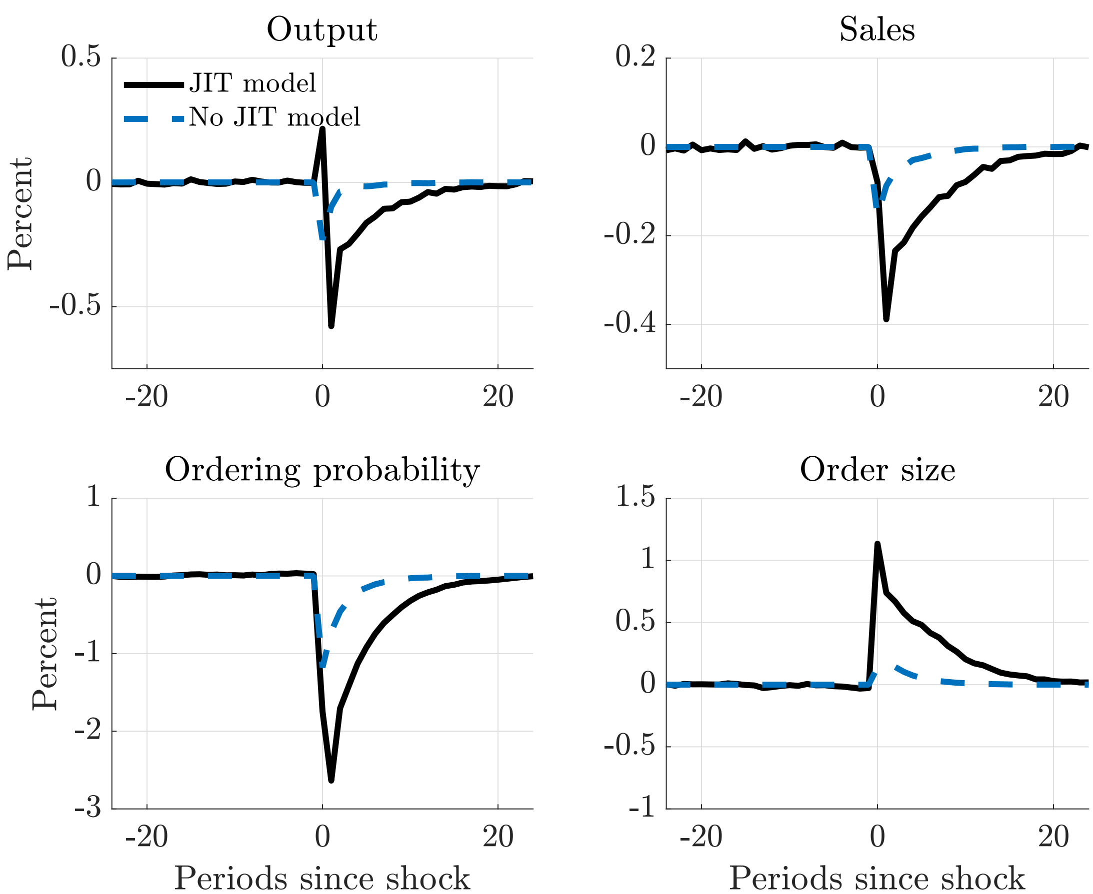

Note: The figure depicts the response of output to a positive shock to fixed order costs, averaged over 5,000 simulated economies.

from the unanticipated shock exercise and indicate that the vulnerability of JIT producers to supply disruptions is also present in an environment that features aggregate uncertainty, with the caveat that the extent of this vulnerability is shaped in part by the size of the shock and the degree to which firms hold precautionary stocks of inventories.[5]

[5] The initial increase in aggregate orders may be tempered, or potentially even disappear, in an environment with sufficiently large capital adjustment costs, as capital is an input in the production of orders and such a friction would dampen the investment response to the shock. I abstract away from capital adjustment costs in my model for simplicity.

References

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.

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.

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.

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

Hansen, Bruce E. and Seojeong Lee (2019), "Asymptotic theory for clustered samples." *Journal of Econometrics*, 210.

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.

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

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.

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.