Online Appendix (Not Intended for Publication)

"Rethinking Inventories and Markups over the Business Cycle"

by Lukasz A. Drozd and Marina M. Tavares

Contents:

Section A provides omitted proofs from the paper.

Section B details the calibration of model parameters and the solution algorithm.

Sections C and D describe the replication package and list data sources.

(Supporting code and data files can be found in the online replication package. Instructions are provided at the end of this document.)

A Omitted Proofs

This section contains omitted proofs from text.

Proof of Lemma 1:

(It is instructive to read the section describing shoppers' search technology laid out after this lemma. We omit time subscripts to simplify notation.)

In general, the policy function of the distributor is a set-valued indicator function on the space of quoted prices and the taste shock, i.e., feasible tuples $(\tilde{p}, \eta)$. As assumed in the text, the distributor only cares about the implied surplus. Expression in (11) shows that, in that case, it is without loss to restrict attention to policy functions that are identical for any tuple $(\tilde{p}, \eta)$ that maps onto the same effective price p. (The definition of the effective price can be found underneath equation (11). The relevant distribution of effective prices is described by the probability measure Pr(.) generated by $\tilde{p}^*$ and the shock η. When we say probability measure (function) or write Pr(.), we mean this measure.)

Consider the optimal policy (a policy that maximizes distributor's profits), which by the above can be represented as measurable set $\mathcal{A} \subset \mathbb{R}_+$ of effective prices with measure $\pi = 1 - \text{Pr}(\mathcal{A}) > 0$ that the shopper is allowed to accept. This policy requires that the contracted shopper delivers a good within an instance of time t of duration dt at an effective price in set $\mathcal{A}$. As long as there is a positive measure of prices in the market that are consistent with this policy—which is implied by the fact that the contract has been accepted by the shopper—the shopper can deliver a good because she can draw an infinite number of quotes in the limit as $dt \to 0$. Next, we show that the closure of the policy set (i.e., $cl\mathcal{A}$) can be restricted to $I_0(\bar{p}) := \{p : 0 \leq p \leq \bar{p}\}$, or else the implied total search costs can be strictly lowered, which would contradict optimality. We refer to such policies as *reservation effective price policies* (REPP).

By contradiction, suppose the above claim is false. If so, we can pick a reservation policy $\mathbf{I_0} := I_0(\bar{p}_0)$ such that $\pi = 1 - \Pr(\mathbf{I_0})$—which, note, is uniquely defined by $\pi > 0$ because density g has full support, and hence $\Pr[0, \bar{p}]$ is a continuous and strictly increasing function in $\bar{p}$ (intermediate value theorem).[38] Since shoppers *randomly* pull prices from the population, as described below the statement of the lemma in text, the expected number of searches to draw a member from a set of the same measure is the same. Accordingly, the implied total search cost associated with the policy $\mathbf{I_0}$ is the same as under the original policy because the probability measure of $\mathbf{I_0}$ and $\mathcal{A}$ is the same.

We next show that the expected effective price under $\mathbf{I_0}$ is strictly lower than under the original policy $\mathcal{A}$, which give the contradiction.

Note that the set $E_0 = \mathcal{A}/\mathbf{I_0}$ (prices in $\mathcal{A}$ that exceed $\bar{p}_0$ and hence are not in $\mathbf{I_0}$) must be of a positive probability measure ($\Pr(E_0) > 0$). If this is not the case, $\mathcal{A}$ and $\mathbf{I_0}$ are identical up to a measure zero set, and so there is nothing to prove (which trivially contradicts the hypothesis and we are done). If this is not the case, define the sets $E_1 = \mathcal{A} \cap \mathbf{I_0}$ (prices in $\mathcal{A}$ that are in $\mathbf{I_0}$) and $E_2 = \mathbf{I_0}/E_1$ (prices in $\mathbf{I_0}$ that are not in $\mathcal{A}$), and note that: i) E_0 and E_1 form a partition of $\mathcal{A}$, and ii) E_1 and E_2 form a partition of $\mathbf{I_0}$. Since $\Pr(\mathcal{A}) = \Pr(\mathbf{I_0})$ and $\Pr(E_0 \cup E_1) = \Pr(E_1 \cup E_2)$, it thus must be that $0 < \Pr(E_0) = \Pr(E_2)$. This is a contradiction by the inequality below because, by construction, all prices in the positive measure set E_0 are above $\bar{p}_0$ and all prices in the equal measure set E_2 are below $\bar{p}_0$:

$$\int_{\mathcal{A}=E_0\cup E_1} p\,\Pr(dp) - \int_{\mathbf{I_0}=E_1\cup E_2} p\,\Pr(dp) = \int_{E_0} p\,\Pr(dp) - \int_{E_2} p\,\Pr(dp) > 0.$$

To see this, note that the left-hand side of this inequality is the difference in the mean price under policy $\mathcal{A}$ and under policy $\mathbf{I_0}$, multiplied by $\Pr(\mathcal{A}) = \Pr(\mathbf{I_0})$. Accordingly, this inequality shows that the expected price conditional on being in set $\mathbf{I_0}$ is *strictly* lower than the expected price conditional on being in set $\mathcal{A}$—a contradiction.

Finally, the fact that $\pi = \Pr(p \geq \bar{p})$ constitutes an equivalent representation of distributor's policy follows from the fact that $\Pr(p \geq \bar{p})$ is monotone in $\bar{p}$ (not necessarily strictly monotone, as is the case in our model). Define the reservation price as the lowest price on the "flat portions" of this function. This is without loss because distributor's surplus remains unchanged. This ensures a bijective mapping. Q.E.D.

Proof of Lemma 2:

The proof considers a discretization of continuous time setup by assuming a fixed period length $dt > 0$ and taking the limit $dt \to 0^+$. We omit time subscripts to simplify notation.

Part 1. By Lemma 1, the relation $1 - \pi = \Pr(p \leq \bar{p}(\pi))$ is well defined and it implies that the event $p \leq \bar{p}(\pi)$ occurs with probability $1 - \pi$ after a single random draw of an effective price by the shopper, after two draws with probability $(1 - \pi)\pi$, and so on and so forth. This defines a geometric process with a fixed success (stopping) probability $(1 - \pi)$. Accordingly, the number of searches

[38] This result trivially generalizes to distributions that are weakly increasing or do not have a full support.

until a match is formed corresponds to the mean of the geometric distribution with parameter $(1-\pi)$.

To calculate the implied search cost, we must take into account discounting with the instance of time of length dt—a tedious technical condition. Since search costs are borne at different points in time on that interval, the discounted value may be different. Define the highest discount factor on the interval $[t, t+dt)$ as $\gamma_{\text{sup}} \equiv e^{-\left(\sup_{l\in[t,t+dt)}\rho_l\right)d\tau}$, and the lowest discount factor as $\gamma_{\text{inf}} \equiv e^{-\left(\inf_{l\in[t,t+dt)}\rho_l\right)d\tau}$. Note that $\gamma_{\text{sup}}, \gamma_{\text{inf}} \rightarrow_{dt\downarrow 0} 1$. (We assume an MIT shock and we assume that the path for ρ_t is left continuous. Therefore, in a discretization of continuous time considered here there are no "jump" discontinuities on the time interval considered). Note that the mean discounted cost with discount $\gamma = \gamma_{\text{sup}}$ or $\gamma = \gamma_{\text{inf}}$ corresponds to the infinite summation defined by the recursion:

$$Sum_{\gamma} := c_0(1-\pi)\gamma + \pi c_0 \gamma + \pi\gamma \left(\underbrace{c_0(1-\pi)\gamma + \pi c_0\gamma + \pi\gamma\left(c_0(1-\pi)\gamma + ...\right)}_{=Sum_{\gamma}} \right). \tag{43}$$

This is similar to the standard proof of the mean value of a geometric distribution but it additionally takes discounting into account. As noted under the expression, the recursion boils down to solving $c_0(1-\pi)\gamma + \pi\gamma c_0 + \pi\gamma Sum_{\gamma} = Sum_{\gamma}$, and hence $Sum_{\gamma} = c_0\gamma(1-\pi\gamma)^{-1} \rightarrow_{\gamma\uparrow 1} c_0(1-\pi)^{-1}$. Since the actual search cost is bounded by $Sum_{\gamma_{\text{sup}}} \leq c(\pi) \leq Sum_{\gamma_{\text{inf}}}$, the result follows with the discount factor.

Part 2. To calculate the integral in (14), we change the variable of integration from p to η, as implied by the definition of the effective price $p := \tilde{p}^* - \eta P$. Let $\bar{p}(\pi)$ be the one-to-one correspondence implied by (12) and proven in Lemma 1. Before we take the integral, we note the following properties: i) $dp = -Pd\eta$ by $p := \tilde{p}^* - \eta P$, ii) the implied bounds of integration for η by the integral for $s(\pi)$ are $p := \tilde{p}^* - \eta P$ are

$$\bar{p}(\pi) = \tilde{p}^* - \bar{\eta}(\pi)P \Rightarrow \bar{\eta}(\pi) := \frac{\tilde{p}^* - \bar{p}(\pi)}{P} \tag{44}$$

$$-\infty = \tilde{p}^* - \eta P \Rightarrow \eta = +\infty, \tag{45}$$

iii) the memoryless property of the exponential distribution implies

$$\mathbb{E}\left[\eta | \eta > \bar{\eta}(\pi)\right] = \max\left\{\bar{\eta}(\pi) + \eta_0, 0\right\}, \tag{46}$$

iv) equation (12) gives $\bar{p}(\pi) = \tilde{p}^* - PG^{-1}(\pi)$ and hence $\bar{p}(\pi) = \tilde{p}^* + \eta_0 P\log(1-\pi)$, and v) the shopper accepts the equilibrium price $\tilde{p}^*$ whenever

$$\eta \geq \bar{\eta}(\pi) := \frac{\tilde{p}^* - \bar{p}(\pi)}{P}. \tag{47}$$

Using the change of variables and the above properties, we now calculate that

$$
\begin{aligned}
s(\pi) &= \frac{P}{1-\pi}\int_{(-\infty,\bar{p}(\pi)]}\left(\frac{P-p}{P}\right)g\left(\frac{\tilde{p}^{*}-p}{P}\right)dp \\
&=_{[p=\tilde{p}^{*}-\eta P,\, dp=-Pd\eta]}\ \frac{P}{1-\pi}\int_{[\bar{\eta}(\pi),\infty)}\left(\frac{P-\tilde{p}^{*}}{P}+\eta\right)g(\eta)\,d\eta \\
&= P-\tilde{p}^{*}+P\mathbb{E}\left[\eta|\eta>\bar{\eta}(\pi)\right] \\
&= P-\tilde{p}^{*}+PG^{-1}(\pi)+\eta_{0}P \\
&= P(1+\eta_{0})-\tilde{p}^{*}-\eta_{0}P\log(1-\pi)\,. \qquad (48)
\end{aligned}
$$

The calculation assumes $\bar{\eta}(\pi)+\eta_0>0$, which is equivalent to imposing $\pi\geq 0$ by the above expression. This constraint is imposed in text as a feasibility restriction in the distributor problem.

Remark 1. *The expected preference shock follows from the above result and it is* $\mathbb{E}_{\pi}\left[\eta|\eta>\bar{\eta}(\pi)\right]=PG^{-1}(\pi)+\eta_0 P=\eta_0\left(P+\log(1-\pi)^{-1}\right)$.

Q.E.D.

Proof of Lemma 3:

We suppress the 'ss' notation. Any variable without an explicit time subscript refers to the steady-state value of that variable. For example, when we write X instead of X_t, we mean X^{ss}, not X_t as used in the text.

(*) In equilibrium, $\lambda(\tilde{p}^{*},\tilde{p}^{*})=\Lambda$. Accordingly, (27) in steady state requires that the steady-state values of M and N satisfy:

$$0=\tau N-\Lambda M-\delta M, \qquad (49)$$

$$0=\Lambda M-\tau N-\delta N+a(M+N), \qquad (50)$$

which necessitates $a=\delta$. From the producer problem in (26), we infer $V_0=\phi v$ and hence $d=0$. This follows from the fact that $a\geq\delta$ and $d>0$ cannot be both true under the firm value maximization in (26).

This system establishes a unique steady state ratio M/N, which can be satisfied for any value $M>0$. Since $\Lambda=Q/M$ and $0<D(P)=Q\left(1+\mathbb{E}_{\pi}\{\eta\}\right)$, where D is exogenous and time-invariant demand in the steady state. A unique value of $\Lambda>0, P>0, \pi$ determines a unique value of $M>0$, in which case N follows. We need $M>0$ or else retail market cannot clear for $D(P)>0$ because there is no production in the steady state.

The values $V_0=\phi v$ and Λ, P, π are the links between the remaining equilibrium conditions and the conditions (variables) above. Thus, from now on, we focus on the remaining equilibrium conditions, while assuming $V_0=\phi v$ and seeking $\Lambda>0$.

We proceed in two steps. First, we show that if the steady state exists, it is unique (part 1). Then, we demonstrate that the steady state exists under the conditions stated in the lemma (part 2).

Part 1 (uniqueness). Consider the HJB equation for V_0 in (24). In the steady state, production must be positive to meet positive retail demand. Therefore, we know $\hat{\tau} = \tau$ and hence $X = V_1 - V_0 \geq v$ must be true. This follows from steady state existence, which here we assume.

The entry condition $V_0 = \phi v$ and (24) evaluated in the steady state implies

$$\rho\phi v = -\zeta_0 v + \tau\left(-v + X\right) - \delta\phi v, \tag{51}$$

and hence

$$X = \tau^{-1}\left(\left(\rho + \delta\right)\phi + \zeta_0 + \tau\right)v. \tag{52}$$

Remark 2. *(For later use) Note that the condition $X = \tau^{-1}\left(\left(\rho + \delta\right)\phi + \zeta_0 + \tau\right)v \geq v$ is satisfied for the default ranges of model parameters.*

Next, consider the distributor's zero profit condition in (16) rewritten as follows:

$$P = \tilde{p}^* + \chi v + \eta_0 P \log\left(\frac{c_0 v}{\eta_0 P}\right). \tag{53}$$

Given the first order condition for the wholesale price in (30), which yields

$$\tilde{p}^* = X + \eta_0 P, \tag{54}$$

given the above formula for X, we obtain the following fixed point for the retail price P:

$$(1 - \eta_0)P + \eta_0 \log\left(\frac{\eta_0}{c_0 v}P\right)^P = X + \chi v = \tau^{-1}\left(\left(\rho + \delta\right)\phi + \zeta_0 + \tau\right)v + \chi v. \tag{55}$$

(Note: In stating this condition, we used the fact that $P\log\left(\frac{c_0 v}{\eta_0 P}\right) = -\log\left(\frac{\eta_0}{c_0 v}P\right)^P$.)

If the steady state exists, as assumed, $\frac{\eta_0}{c_0 v}P = (1 - \pi)^{-1} > 1$, since this term comes from the distributor's profit maximization in (16) and feasibility requires $0 \leq \pi < 1$. Accordingly, for the range of values that P may take in the steady state, we know that the left-hand side of (55) is strictly increasing in P. Thus, if a solution exists, it must be unique because the right-hand side is invariant with respect to P. The steady-state value of Λ can be derived from (29) by substituting the steady-state value of X given by (52), setting $\dot{X} = 0$, and substituting $\tilde{p}^*$ using (54). After some manipulations, we derive:

$$\Lambda = \frac{v}{\eta_0 P}\left(\tau^{-1}\left(\delta + \rho + \tau\right)\left(\phi\left(\delta + \rho\right) + \zeta_0 + \tau\right) + \zeta - \tau\right). \tag{56}$$

We have now shown that there is at most a single set of equilibrium values that satisfy the steady state equilibrium conditions implied by Definition 1: $a, d, M, N, Q, \Lambda, X, V_0, V_1, P, \tilde{p}^*, \pi, \bar{p}$. (The unique value for $\bar{p}$ is stated in Lemma 2 and recall that, by definition, $V_1 \equiv V_0 + X$.)

Part 2 (existence). As discussed in the beginning (see *), the values of Q, M, N, a, d are uniquely pined down by $\Lambda > 0, P > 0, 0 \leq \pi < 1$, which requires $V_0 = \phi v$ as an equilibrium condition. The following equilibrium conditions are then satisfied: the law of motion (27), producer value maximization in (26), and retail market clearing.

To obtain the remaining variables and ensure the remaining equilibrium conditions are satisfied, we start from the requirement that $X \geq v$ for production to be positive—and hence for the retail market to clear for $D > 0$. The steady state value of X is given by (52) (see part 1)—which, recall, embeds $V_0 = \phi v$ and the HJB equation for V_0 in (24). As noted in Remark 2 (see part 1), this value satisfies $X > v$ for the default ranges of model parameters. What needs to be shown is that there exists $\Lambda > 0$ such that, for $\dot{X} = 0$ and X given by (52), the HJB equation in (29) is satisfied (note that this automatically ensures HJB equation for V_1 in (22) is also satisfied). We have derived the implied by that steady value relation for Λ in equation (56) (see part 1). By that equation, if $P > 0$, for all positive parameter values $\Lambda > 0$ because the expression numerator, $\tau^{-1}(\delta + \rho + \tau)(\phi(\delta + \rho) + \zeta_0 + \tau) + \zeta - \tau$, is positive for the default ranges of model parameters. (Note that τ^2 cancels out after bringing this expression to a common denominator τ. Since we are only left with a summation of positive parameters, the expression is positive.)

Next, we show that we can can find $P > 0$ that solves (55) (see part 1) and satisfies the feasibility restriction implied by (17): $0 \leq \pi = 1 - \frac{c_0}{\eta_0}\frac{v}{P} < 1$. Given the steady state value of X in (52), $\tilde{p}^*$ from (54), and assuming we solve for $P > 0$, we have satisfied all remaining equilibrium conditions (remaining relative to *): the entry condition $V_0 = \phi v$, steady state relations implied by the HJB equations in (24) and (22), the first order condition for $\tilde{p}$, positive production requirement $\hat{\tau} = \tau$, we are sure that $\Lambda > 0$, and distributor's zero profit condition is ensured by (55) (which we need to show has a solution).

To show that we can solve for the retail P, note that we can rewrite the above requirement $\frac{\eta_0}{c_0 v} P = (1 - \pi)^{-1} > 1$ for P that solves (55), and hence we need to show $P > \frac{c_0 v}{\eta_0}$.

Define $\underline{P} = \frac{c_0 v}{\eta_0}$ so that $\frac{\eta_0}{c_0 v}\underline{P} = 1$, and note that the following properties apply to the second term on the left-hand side of (55): 1) $\log\left(\frac{\eta_0}{c_0 v}\underline{P}\right)^{\underline{P}} = 0$, and 2) $\log\left(\frac{\eta_0}{c_0 v}\underline{P}\right)^{\underline{P}}$ is strictly increasing in P on the restricted domain $[\underline{P}, \infty)$. Accordingly, if at $P = \underline{P}$ we can be sure that the left-hand side of (55) is strictly lower than the right-hand side, the result follows: a unique solution of (55) exists because the left-hand side of (55) is strictly increasing and unbounded in P, lower than the right hand side at $\underline{P}$, and the right-hand side is independent of P. Conversely, if this is not the case, there is no solution, and so this is an 'if and only' relation (recall here that $\underline{P}$ is a tight lowest value, or else $\pi < 0$). Plugging in $P = \underline{P}$ to (55), we obtain the inequality in Assumption (2):

$$(1 - \eta_0)\underbrace{(c_0/\eta_0)}_{\underline{P}/v} + \eta_0 0 < \underbrace{\tau^{-1}((\delta + \rho)\phi + \zeta_0 + \tau)}_{X/v} + \chi. \tag{57}$$

Remark 3. *To gain some intuition, rewrite the above condition as*

$$\underline{P} = \frac{c_0 v}{\eta_0} < v + \tau^{-1}((\delta + \rho)\phi v + \zeta_0 v) + \chi v + c_0 v.$$

The terms on the right-hand side represent all the costs associated with producing and selling a good when the search precision π is zero. These include the production cost, operational costs, and sunk costs (recall that producing a good takes τ^{-1} units of time). This scenario represents the worst case for the distributor's surplus, since any equilibrium retail price that higher yields a greater surplus. Consequently, the distributor's zero-profit condition cannot be satisfied if this inequality is violated (for reasons discussed in Step 1).

We have now shown that (55) has a solution on the admissible domain $P \in (\underline{P}, \infty)$, and as we noted above, all equilibrium conditions and domain restriction are satisfied when X is given by (52), $\tilde{p}^*$ is given by (54) (given X), Λ is given by (56), $V_0 = \phi v$, $V_1 \equiv \phi v + X$, π is given by (17), and a, d, M, N, Q are pinned down as discussed in the beginning ($\bar{p}$ can be found using Lemma 2). We have also shown that the restriction of the lemma (A2) is a necessary condition, or else the candidate steady state equilibrium implies $\pi < 0$. Q.E.D.

Proof of Corollary 1:

We use the formula for P^{ss} (steady state value) given by the linearized equation in (35). We substitute the steady state value X^{ss} from (52) (proof of Lemma 3) and plug the obtained expression into the steady state formula for Λ^{ss} in (56) (proof of Lemma 3). This gives:

$$\Lambda^{ss} = \frac{\tau\left(\tau^{-1}(\delta+\rho+\tau)\left(\phi(\delta+\rho)+\zeta_0+\tau\right)+\zeta-\tau\right)}{\Theta\left(\Gamma\tau+\phi(\delta+\rho)+\zeta_0+\tau\right)}. \tag{58}$$

Substituting the above steady state value of Λ into the expressions for a_1 and a_0 in the statement of the corollary, after basic algebraic manipulations to put all terms under common denominator, we obtain:

$$a_0 = \frac{\left(\Gamma(\delta+\rho+\tau)-\zeta+\tau\right)\left(\phi(\delta+\rho)+\zeta_0+\tau\right)}{\Gamma\tau+\phi(\delta+\rho)+\zeta_0+\tau} > 0, \tag{59}$$

$$a_1 = \frac{\tau(\Gamma(\delta+\rho+\tau)-\zeta+\tau)}{\Gamma\tau+\phi(\delta+\rho)+\zeta_0+\tau} > 0. \tag{60}$$

(Note that Θ cancels out.) The inequalities above follow from the fact that $\Gamma = \eta_0 P^{ss}/v + \chi$ in (35), and we know from the proof of Lemma 3 (step 2) that in the steady state $P^{ss} > (c_0 v)/(\eta_0)$; hence, $\Gamma > c_0 + \chi$, and therefore $\tau - \zeta + (c_0 + \chi)(\delta + \rho + \tau) > 0$ guarantees that all terms in the formulas for a_1 and a_0 are strictly positive. Q.E.D.

Proof of Lemma 4:

The expressions stated in the lemma follow by substituting the linearized term in text and solving for prices using equations (54) and (53) from the proof of Lemma 3. The fact that $\Gamma > 0$ is trivial and the fact $\Theta > 0$ follows from the proof of Lemma 3. The lemma derives the lower bound on the steady state retail price P^{ss} (step 2). This bound ensures the logarithmic term in the expression for $\Theta > 0$ is positive. (Online Appendix in Section B.1 provides an alternative derivation of these expressions.)

B Model Numerical Solution and Calibration

This appendix discusses the numerical solution of the model and the calibration of its parameters/shocks. Supporting *Mathematica* notebook can be found in the replication package (solve_Model_vX.nb.)

B.1 Equilibrium System and Steady State

Consider zero profit condition of the distributor in (16); that is

$$P - \tilde{p}^* - \eta_0 P \log\left(\frac{c_0}{\eta_0}\frac{v}{P}\right) = \chi v, \tag{61}$$

where we drop time subscripts since we focus here on calculating the steady state.

The above condition cannot be solved in closed form because of the last term. To solve the model numerically, we resort a linear approximation of the last term:

$$\log\left(\frac{c_0}{\eta_0}\frac{v}{P_t}\right) \approx \frac{P^{ss}}{P} - 1 + \log\left(\frac{c_0 v}{\eta_0 P^{ss}}\right), \tag{62}$$

implying

$$\chi v = P(1-\eta_0) - \tilde{p}^* - \eta_0 P^{ss} - \eta_0 P \log\left(\frac{c_0}{\eta_0}\frac{v}{P^{ss}}\right). \tag{63}$$

We introduce two composite parameters to soak up the unwanted terms (P^{ss} and $\log\left(\frac{c_0 v}{\eta_0 P^{ss}}\right)$):

$$\Gamma = \frac{P^{ss}\eta_0}{v} + \chi,\ \Theta = \frac{\eta_0}{1 - \eta_0 \log\left(\frac{c_0 v}{\eta_0 P^{ss}}\right)}. \tag{64}$$

Substituting into the above expression, we obtain

$$P = \frac{\Theta}{\eta_0}\frac{\tilde{p}^* + \Gamma v}{1+\Theta}. \tag{65}$$

The wholesale quoted price is given by (30), and hence

$$\tilde{p}^* = V_1 - V_0 + \eta_0 P. \tag{66}$$

Accordingly,

$$\tilde{p}^* = (V_1 - V_0)(1+\Theta) + v\Gamma\Theta,\ P = \frac{\Theta}{\eta_0}(V_1 - V_0 + v\Gamma). \tag{67}$$

Adding the entry condition $V_0 = \phi v$ and the Bellman equations (22) and (24) evaluated in the steady state ($\dot{V}_0 = 0, \dot{V}_1 = 0$), we obtain the following steady state system augmented by the auxiliary parameters Γ, Θ:

$$V_0 = \phi v, \tag{68}$$
$$(\rho+\delta)V_1 = -(\zeta+\zeta_0)v + \Lambda(\tilde{p}^* + V_0 - V_1), \tag{69}$$
$$(\rho+\delta)V_0 = -\zeta_0 v + \tau(-v + V_1 - V_0), \tag{70}$$
$$X = V_1 - V_0, \tag{71}$$
$$\tilde{p}^* = (V_1 - V_0)(1+\Theta) + v\Gamma\Theta, \tag{72}$$
$$P = \frac{\Theta}{\eta_0}(V_1 - V_0 + v\Gamma). \tag{73}$$

The analytic solution gives

$$V_0^{ss} = v\phi, \tag{74}$$
$$V_1^{ss} = \frac{v(\phi(\delta+\rho+\tau) + \zeta_0 + \tau)}{\tau}, \tag{75}$$
$$X^{ss} = \frac{v(\phi(\delta+\rho) + \zeta_0 + \tau)}{\tau}, \tag{76}$$
$$\tilde{p}^{*ss} = \frac{v(\Gamma\Theta\tau + (\Theta+1)\phi(\delta+\rho) + \zeta_0\Theta + \zeta_0 + \Theta\tau + \tau)}{\tau}, \tag{77}$$
$$P^{ss} = \frac{\Theta v(\Gamma\tau + \phi(\delta+\rho) + \zeta_0 + \tau)}{\eta_0\tau}, \tag{78}$$
$$\Lambda^{ss} = \frac{\delta\zeta_0 + \phi(\delta+\rho)(\delta+\rho+\tau) + \delta\tau + \zeta_0\rho + \zeta\tau + \zeta_0\tau + \rho\tau}{\Theta(\Gamma\tau + \phi(\delta+\rho) + \zeta_0 + \tau)}, \tag{79}$$
$$\pi^{ss} = 1 - \frac{c_0\tau}{\Theta(\Gamma\tau + \phi(\delta+\rho) + \zeta_0 + \tau)} \tag{80}$$

(These formulas a were automatically generated by Mathematica.)

Lemma 3 shows that the steady exists and is unique.

B.2 Calibration (Extended Version)

This section describes the calibration of the model. The baseline period of the model ($t = 1, 2, 3$) corresponds to one month. To map model onto the data, we assume (real) sales is Q, real inventory stock is M, the CPI is P, and nominal value added is $P_t Q_t(1 + \mathbb{E}\{\eta\})$. Unlike in the analytic section, gross margin/markup is measured in logs as $\log(p/v)$. Parameter values are listed in the table below.

B.2.1 Static Steady State-Based Targets

Consider first the targets that map directly onto model parameters. These include the weighted-average cost of capital (WACC) of 10 percent (annual rate) and the delivery delay of 60 days in U.S. manufacturing.[39] WACC pins down the steady state value of ρ according to the formula that first converts it to a monthly rate and then to a continuously compounded rate:

$$\rho^{ss} = \ln(1 + \text{WACC}/100)^{12}. \tag{81}$$

The average delivery time in the model is τ^{-1}, which for the given target yields $\tau = 0.5$. In the baseline calibration we set inventory holding cost $\zeta = 0$ and assume it is σ that prevents firms from holding more than one unit—so as to satisfy Assumption 1.

We next discuss how we calibrate the values of $\chi, \eta_0, c_0, \phi, \zeta_0$, as well as the auxiliary parameters Θ and Γ. We first describe data targets for these parameters.

Since manufacturing sector went through a major transformation in 2000s and 2010s, and much of this period lies outside of our sample period, we use the 1997 input-output (IO) tables published by Bureau of Economic Analysis (BEA) to set targets for margins and distribution costs. Using later tables would increase the targeted moment for by about 10-20 percent depending on the exact date (before of after GFC). (The raw tables can be found in the replication package, folder IO-Calibration.)

Our first data target is the distribution margin,

$$\mathscr{M}_1 = \frac{P - \tilde{p}^*}{\tilde{p}^*} = .38, \tag{82}$$

which we associate with the total trade and transportation margins borne on a unit of manufactured good and indirect (net) taxes borne by the final users. To obtain this number, we proceed as follows. We use the 1997 domestic-*supply*-of commodities IO table (for 15 industries) and focus on manufacturing commodities (in row of the IO table). We calculate the following ratios using this table (as shown in Table 2 below):

$$\tilde{\mathscr{M}}_1^I = \frac{\text{Total trade and transport margin in manufacturing}}{\text{Total supply of manufacturing commodities in basic prices}}, \tag{83}$$

$$\tilde{\mathscr{M}}_1^F = \frac{\text{Total supply of manufacturing commodities in purchaser prices}}{\text{Total supply of manufacturing commodities in basic prices}}. \tag{84}$$

The first ratio is the share of trade and transportation margins relative to the basic price of

[39]For an overview of cost of capital estimates for various industries, see the listing compiled by Aswath Adamodar at https://pages.stern.nyu.edu/~adamodar/New_Home_Page/datafile/wacc.html. Delivery delays in manufacturing come from Deloitte Research Center for Energy & Industrials (2024)—based on the source data from Institute for Supply Management (ISM)—and the report can be found at https://www2.deloitte.com/us/en/insights/industry/manufacturing/manufacturing-industry-outlook.html (see Figure 4). This number has remained steady between 2015 and 2019 according to the earlier reports.

manufacturing commodities that are sold as either intermediate good or the final goods. The second ratio includes (net) taxes and tariffs, where we associate this particular margin with purchases by the final users, as we shall see. In the second step—shown in Table 3—we calculate the following two ratios: i) the fraction of manufacturing output sold to final users

$$f = \frac{\text{Final use of manufacturing commodities}}{\text{Total use of manufacturing commodities}} \tag{85}$$

and ii) the share of manufacturing commodities sold to manufacturing sector as an intermediate input used by that sector:

$$x = \frac{\text{Use of manufacturing commodities as intermediate inputs by manufacturing sector}}{\text{Total use of intermediate inputs by manufacturing sector}}. \tag{86}$$

Table 2: Trade and transportation margins on manufacturing goods, US 1997 (millions of 1997 \$).

Commodity view	C1	C2	C3	$\tilde{\mathscr{M}}_1^I$	$\tilde{\mathscr{M}}_1^F$
(row in supply IO table)	Product supply (basic prices)	Product supply (purchaser prices)	Total trade & transportation margins	$\frac{C3}{C1}$	$\frac{C2-C1}{C1}$
Manufacturing	\$4,507,147	\$5,914,729	\$1,074,854	**24%**	**31%**

Notes: This extract comes from BEA's 1997 domestic-supply-of-commodities IO table.

Table 3: Share of manufacturing goods in production and final demand (millions of 1997 \$).

Commodity view	C1	C2	C3	**f**	**x**
(row in supply IO table)	Total use of products	Total use as intermediate inputs	Use as intermediates in manufacturing sector	$\frac{C1-C2}{C1}$	$\frac{C3}{C2}$
Manufacturing	\$5,914,730	\$2,794,901	\$1,656,638	**53%**	**59%**

Notes: This extract comes from BEA's 1997 use-of-commodities IO table.

Using these ratios, we calculate the total distribution margin by cumulating the cycles a manufacturing commodity is used as an intermediate input to produce manufacturing goods, where the outflows from the cycle are purchases of manufacturing commodities by final users (final users or other sectors) that accrue trade and transportation margins. Specifically, the outflows accrues trade and transportation margin $\tilde{\mathscr{M}}_1^F$ in proportion to purchases of manufacturing commodities as intermediate goods and $\tilde{\mathscr{M}}_1^I$ in proportion to purchases as final goods. This procedure gives the following geometric series:

$$\mathscr{M}_1 = \tilde{\mathscr{M}}_1^F f + \tilde{\mathscr{M}}_1^I (1-f) + x(1-f)(\tilde{\mathscr{M}}_1^F f + \tilde{\mathscr{M}}_1^I (1-f) + x(1-f)(\tilde{\mathscr{M}}_1^F f + \tilde{\mathscr{M}}_1^I (1-f) + \ldots). \tag{87}$$

Plugging in the numbers from Tables 3 and 2, we obtain

$$\mathscr{M}_1 = \frac{\tilde{\mathscr{M}}_1^F f + \tilde{\mathscr{M}}_1^I (1-f)}{1 - x(1-f)} = .38. \tag{88}$$

The key assumption is that similar margins apply along the production chain. If the bulk of trade and transportation margins accrue closer to the end of the supply chain, or the beginning, our procedure may bias the results upwards or downwards, respectively. The so-called PCE Bridge published by BEA sheds light on this issue by reporting separate transportation and distribution margins on final consumption that go into PCE.[40] The PCE bridge data suggests that there is no significant discrepancy of this sort in the data. Given the distribution margins used in the related literature, our number is on the conservative side (Burstein et al., 2000).

Our procedure intentionally ignores trade margins on intermediate goods that are implicitly embedded in the value of intermediate commodities purchased by the manufacturing sector from other sectors, since the goal of this calculation is to get to the margins paid on a unit of manufacturing output.

Our second target is the gross profit margin in manufacturing and trade industries:

$$\mathscr{M}_2 = \frac{\tilde{p}^* - v}{v} = .37. \tag{89}$$

To set this target, we use use-of-commodities IO table and take the ratio of gross surplus to the value added in producer prices for the three sectors (manufacturing/retail/wholesale). The extract of relevant data is in Table 4.

Table 4: Gross surplus in manufacturing and trade industries, US 1997 (millions of $).

		Sector view (column in use IO table)			
		Manufacturing	Wholesale	Retail	**All three sectors**
R1	Total intermediate inputs	\$2,514,885	\$219,597	\$277,373	\$3,011,855
R2	Compensation of employees	\$788,978	\$266,144	\$338,975	\$1,394,097
R3	Net taxes on production	\$26,559	\$6,940	\$2,994	\$36,493
R4	Gross operating surplus	\$552,767	\$138,413	\$134,696	\$825,876
R5	**Value added (basic prices) R2+R3+R4**	\$1,368,304	\$411,497	\$476,665	\$2,256,466
	Ratio: Gross margin R4/R5	40%	34%	28%	**37%**

Notes: This extract comes from BEA's 1997 use-of-commodities IO table.

[40] Available at https://www.bea.gov/industry/industry-underlying-estimates.

Our third target is the inventory-to-sales ratio during our sample period of:

$$\mathscr{M}_4 = \frac{M^{ss}}{Q^{ss}} = \frac{M^{ss}}{\tau N^{ss}} = 1.5. \tag{90}$$

Our fourth target is the share of expenses on sales infrastructure borne by producers, which we set equal to 25 percent:

$$\mathscr{M}_5 = \frac{\phi\delta + \zeta_0(M^{ss} + N^{ss})}{\tau N^{ss}} = 0.3. \tag{91}$$

To set this target, we use the expenses on SG&A relative to sales in the Compustat data for the manufacturing sector, which are about 17 percent. However, these costs are paid multiple times as manufactured goods are processed within the manufacturing sector, our target is 25 percent to reflect the same input-output multiplier of $(1 - x(1 - f))^{-1}$ =1.38 that we used above to cumulate trade and transportation margins in the calculation of distribution margin above (first target).[41]

Our last, fifth target from this group, pertains to search process and assumes that, on average, shoppers seek three quotes from suppliers before making an offer—which is the standard used by purchasing departments:

$$\mathscr{M}_3 = (1 - \pi)^{-1} = 3. \tag{92}$$

Finally, we set $\zeta = 0$ and relate our model to Assumption 1 by considering a hypothetical firm that faces equilibrium conditions but can accumulate inventory in the marketing state. To prevent that firm from doing so, we find $\sigma = 0.64$. Generally, if we set a different value for σ, we would need to impose $\zeta = 0.0638324 - 0.112234\sigma$ on that firm, which quantifies the friction that is needed to support our model's structure. Recall that $\sigma = 0.57$ operates on the cost of maintaining outposts, which is set at $\zeta_0 v = 0.094v$ in the calibration. Accordingly, 2.25 percent of value added in manufacturing and trade industries would be required to double inventory holding capacity. After accounting for the cost of all structures, land, and other resources that contribute to inventory holding costs, it is a large number.

The mapping between parameters and targets is analytic and invertible. The complete formulas are

[41] For an overview of this ratio across industries and countries, see the listing compiled by Aswath Adamodar at https://pages.stern.nyu.edu/~adamodar/New_Home_Page/datafile/margin.html. These numbers (relative to sales) range between 15 and 25 percent across sectors. The estimated share of value added in gross output comes from KLEMS for manufacturing and trade industries (value added weighted) can be found at https://www.bea.gov/data/special-topics/integrated-industry-level-production-account-klems (production account tables).

too long to state and we state them here for $\zeta = 0$, which is what we assume in calibration:

$$
\begin{aligned}
c_0 &= \frac{\mathscr{M}_3((\mathscr{M}_2+1)(\delta+\rho)+\mathscr{M}_2\tau)}{\mathscr{M}_3\mathscr{M}_4(\delta+\rho+\tau)+\mathscr{M}_4}, \\
\eta_0 &= \frac{\mathscr{M}_3((\mathscr{M}_2+1)(\delta+\rho)+\mathscr{M}_2\tau)}{(\mathscr{M}_1+1)(\mathscr{M}_2+1)(\mathscr{M}_3(\delta+\rho+\tau)+1)}, \\
\phi &= \frac{\tau(\mathscr{M}_2-\mathscr{M}_3(\delta+\rho))-\zeta_0(\mathscr{M}_3(\delta+\rho+\tau)+1)}{(\delta+\rho)(\mathscr{M}_3(\delta+\rho+\tau)+1)}, \\
c_0 &= \frac{\mathscr{M}_3((\mathscr{M}_2+1)(\delta+\rho)+\mathscr{M}_2\tau)}{\mathscr{M}_3\mathscr{M}_4(\delta+\rho+\tau)+\mathscr{M}_4}, \\
\chi &= \mathscr{M}_1(\mathscr{M}_2+1)+\frac{\mathscr{M}_3\log(\mathscr{M}_4)\left((\mathscr{M}_2+1)(\delta+\rho)+\mathscr{M}_2\tau\right)}{\mathscr{M}_3(\delta+\rho+\tau)+1}, \\
\zeta_0 &= \frac{\tau(\delta+\rho)\left(\mathscr{M}_5\left(\delta\mathscr{M}_3^2\tau-\mathscr{M}_3(\rho+\tau)-1\right)-\delta\mathscr{M}_3(\mathscr{M}_3\tau+1)\right)-\delta\mathscr{M}_2(\mathscr{M}_5-1)\tau(\mathscr{M}_3\tau+1)}{(\mathscr{M}_5-1)\rho(\mathscr{M}_3\tau+1)(\mathscr{M}_3(\delta+\rho+\tau)+1)}, \\
\Theta &= \frac{\mathscr{M}_3((\mathscr{M}_2+1)(\delta+\rho)+\mathscr{M}_2\tau)}{(\mathscr{M}_2+\chi+1)(\mathscr{M}_3(\delta+\rho+\tau)+1)}, \\
\Gamma &= \frac{\mathscr{M}_3((\mathscr{M}_2+1)(\delta+\rho)+\mathscr{M}_2\tau)}{\mathscr{M}_3(\delta+\rho+\tau)+1}+\chi.
\end{aligned}
\tag{93}
$$

(These formulas were automatically generated by Mathematica.)

Table 5 provides the numeric value of the assumed parameters.

Table 5: Parameter Values.

η_0	c_0	δ	ϕ	χ	τ	(ζ_0, ζ, σ)	ρ^{ss}	Γ	Θ
0.088	0.056	4.0×10^{-4}	0.78	0.70	0.5	$(0.095, 0.0, 0.57)$	7.94×10^{-3}	0.871	0.081

B.2.2 Dynamic Impulse Response-Based Targets

We set the remaining parameters and calibrate the monetary policy shock path to match the impulse responses generated by our SVAR.

Calibration of δ and κ_a.— We use the impulse responses to set the values of δ and κ_a. The parameter δ dictates the rate at which sales and production decline toward the trough, while κ_a controls the rate of recovery after exiting the inaction zone and when the monetary authority shifts policy. Our analysis does not assume an "active" recovery; instead, the recovery is driven by the impulse responses we incorporate as shocks. We target these rates over a 12-month window following the trough in sales as implied by the SVAR.

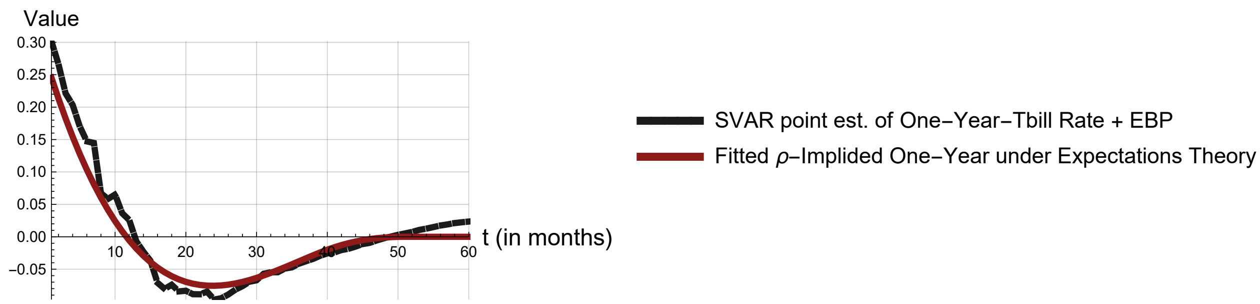

Figure 7: Fitted path of ρ using expectations theory and smooth polynomial approximation.

Notes: The figure shows the implied one-year rate by the fitted path of ρ_t that we assume in the model.

Calibration of the shock.— To calibrate the path of ρ_t, we target the impulse response of the sum of the one-year T-bill rate and the EBP as implied by the SVAR (point estimate). We assume that this rate converges to the steady state after 50 months. Using the expectations theory, we calculate the monthly rates implied by the annual series under perfect foresight. These are then converted to continuously compounded rates by taking the natural logarithm.

Specifically, let i_t^1 represent the monthly rate and i_t^{12} the one-year rate (the sum of the one-year T-bill rate and EBP), where $t = 1, 2, \ldots, 60$ denotes the months of data. To derive the monthly rate, we use the expectations theory under perfect foresight, which implies

$$1 + i_t^{12} = \prod_{j=1}^{12} \left(1 + i_{t+j-1}^1\right), \tag{94}$$

and hence

$$\frac{1 + i_{t+1}^{12}}{1 + i_t^{12}} = \frac{1 + i_{t+12}^1}{1 + i_t^1}, \tag{95}$$

and

$$i_t^1 = \frac{1 + i_t^{12}}{1 + i_{t+1}^{12}} \left(1 + i_{t+12}^1\right) - 1, \tag{96}$$

where, for $t \geq 50$, we assume steady state values: $i_t^1 = i^{ss} := \exp\left(\rho^{ss}\right) - 1$ and $\frac{1+i_t^{12}}{1+i_{t+1}^{12}} = 1$. We fit a third degree polynomial to a natural logarithm of the obtained monthly series $\{i_t^1\}$ to obtain a continuous function for ρ_t that enters the differential equation. Figure 7 shows the result after converting the interpolated function ρ_t back to a one-year forward rate using expectations theory.

To set the path of Q_t, we target the SVAR-implied trajectory of the inventory-to-sales ratio. We achieve this by fitting a third-degree polynomial to the positive portion of the impulse response and a connected polynomial to the negative portion. These two segments are joined where the data's impulse response transitions from positive to negative, ensuring continuity. The estimated polynomials provide a smooth target for the inventory-to-sales ratio, which our model is designed to match precisely by adjusting the path of Q_t. The next section will explain how this is implemented

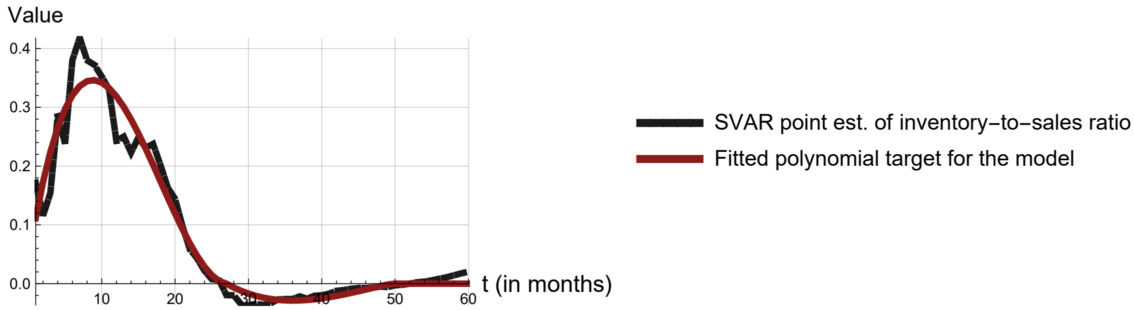

Figure 8: Fitted target for inventory-to-sales ratio.

Notes: The figure shows the fitted smooth target for the inventory-to-sales ratio that our calibrated model must hit via a particular path for Q_t.

within our model's solution framework. Figure 8 illustrates the fit of the two-part polynomial.

B.3 Numerical Procedure to Solve the Model

This section outlines the numerical approach we use to solve the model and implement our calibration strategy.

First, we solve the system of Hamilton-Jacobi-Bellman (HJB) equations in (22) and (24), substituting in the equilibrium wholesale price from (35) and the calibrated polynomial paths for ρ_t and Λ_t. We solve this system forward in time, assuming the model converges to the steady state at a distant horizon. The dependence of the HJB equations on ρ_t and Λ_t simplifies the problem, making the calibration of the inventory-to-sales ratio straightforward.

As discussed in the text, convergence to the steady state requires that ρ_t and Λ_t stabilize, but M_t and Q_t do not need to return to their steady-state values—in fact, they do not in our simulation. This is due to the model's assumption of wage rigidity, which allows for permanent changes in labor supply. We do not model how the economy recovers, other than it is implied by the impulse responses we feed (which does bring partial recovery). Importantly, this rigidity does not hinder the convergence of V_{0t} and V_{1t}.

Solving the HJB equations yields the paths for V_{0t}, V_{1t}, and X_t. Using V_{0t}, we identify the time $T > 0$ when the economy exits the inaction region, defined by $0 < V_{0t} < \phi v$. We confirm that V_{0t} remains within these bounds before T and does not reach the lower boundary $V_0 = 0$, which would trigger liquidations—which is never the case. This behavior is evident in the impulse responses shown in Figure 4, which shows how much larger than shock would need to be to trigger liquidations. At time T, $V_{0T} = \phi v$, and for $t > T$, $V_{0t} > \phi v$, indicating that $a_t > 0$ as given by (26). Specifically, for $t > T$, $a_t = \delta + (1/2)(V_{0t} - \phi v)/\kappa_a$, while $a_t = 0$ within the inaction zone (and $d_t = 0$, since the lower boundary is never reached).

Given the path for a_t, and the calibrated path for Λ_t, we next solve the differential equations in (27) forward, using the steady-state values M_0 and N_0 as initial conditions (the boundary value is the initial value). From the resulting paths for M_t and N_t, we back out $Q_t = \Lambda_t M_t$.

Although we do not explicitly recover the original shock, it can be reconstructed using the relation $D_{0t} P_t^{-\varepsilon} = (1 + \mathbb{E}\eta) Q_t$. Detailed implementation can be found in the annotated Mathematica notebook (`solve_Model.nb`) included in the replication package.

C SVAR Robustness and Extensions

In this appendix, we provide extended results for the SVAR analysis in Section 2. We provide alternative interpolation of markups, consider markups for all firms, and provide extended results that are presented in the paper.

C.1 Proxy SVAR Setup

As noted in the paper, we follow Gertler and Karadi (2015) (GK15, hereafter).[42] The SVAR includes 12 lags, a constant, and it takes the following form:

$$\mathbf{Y}_t = \mathbf{C} + \mathbf{A}_1 \mathbf{Y}_{t-1} + \mathbf{A}_2 \mathbf{Y}_{t-2} + \cdots + \mathbf{A}_{12} \mathbf{Y}_{t-12} + \mathbf{u}_t \tag{97}$$

where $\mathbf{Y}_t$ is an $n \times 1$ vector of endogenous variables in month t between 1979:m7 and 2012:m6, $\mathbf{C}$ is an $n \times 1$ set of constants, and $\mathbf{A}_i$ is an $n \times n$ matrix of coefficients for lag i (where $i = 1, 2, \ldots, 12$). We include the following variables:

$$\mathbf{Y}_t = \begin{bmatrix} \text{1yr T-Bill Rate (GS1) in percent} \\ \text{Excess Bond Premium (EBP) in percent} \\ \log\left(\text{Consumer Price Index (CPI)}\right) \times 100 \\ \log\left(\text{Real Sales}\right) \times 100 \\ \log\left(\text{Real Inventory Stock}\right) \times 100 \\ \text{Compustat-derived Margin in (9)} \times 100 \end{bmatrix} \tag{98}$$

The first four variables are the same as in GK15, but we replace industrial production by real sales for manufacturing and trade industries given our focus on the inventory-to-sales ratio. This makes little different and the two behave similarly. Omission of mining and utilities is reasonable in the given context.

As in GK15, the residuals $\mathbf{u}_t$ are used in conjunction with their instrument (ff4_tc in GK15) to identify monetary policy shocks. The time series for ff4_tc and other variables can be found in the dataM.csv file in the replication package (see last section of this document). The variable 1yr T-Bill Rate serves as the policy indicator and the residuals associated with this variable (u_{1t}) are regressed

[42] Our Python implementation of GK15 estimation is based on the Matlab VAR Toolbox version of Gertler and Karadi (2015) due to Ambrogio Cesa-Bianchi (`https://github.com/ambropo/VAR-Toolbox`).

A. Original GK15 specification with industrial production:

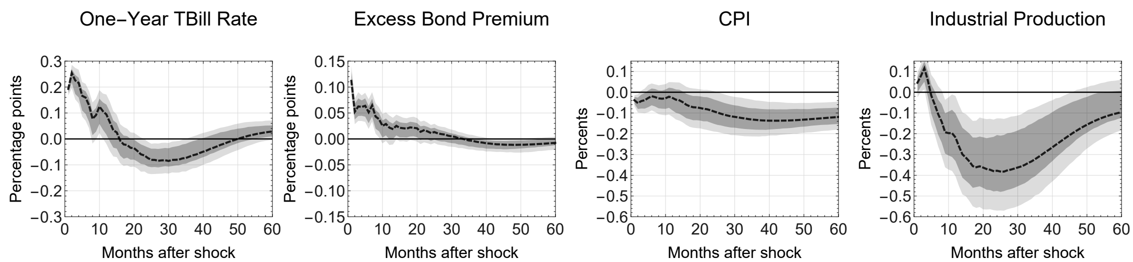

B. GK15 specification with real sales replacing industrial production:

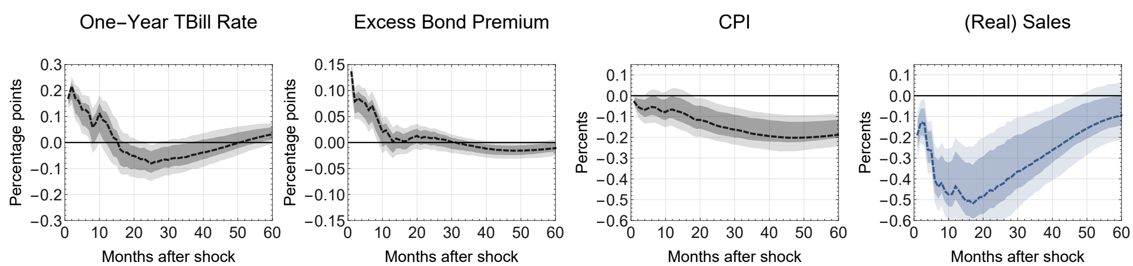

C. Baseline SVAR specification with no markup (first 4 variables reported):

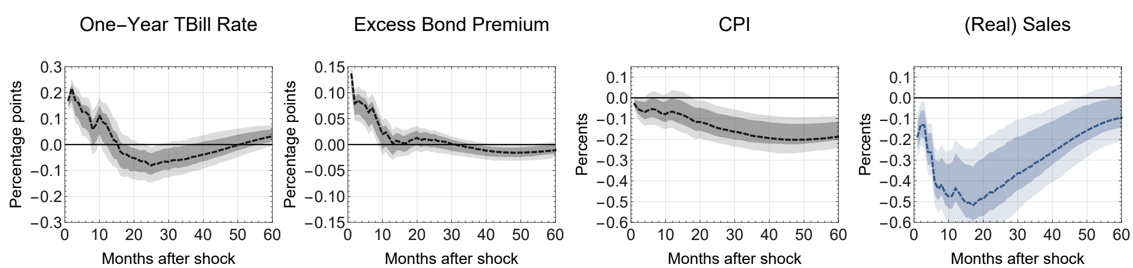

D. Baseline SVAR as in the paper (first 4 variables reported):

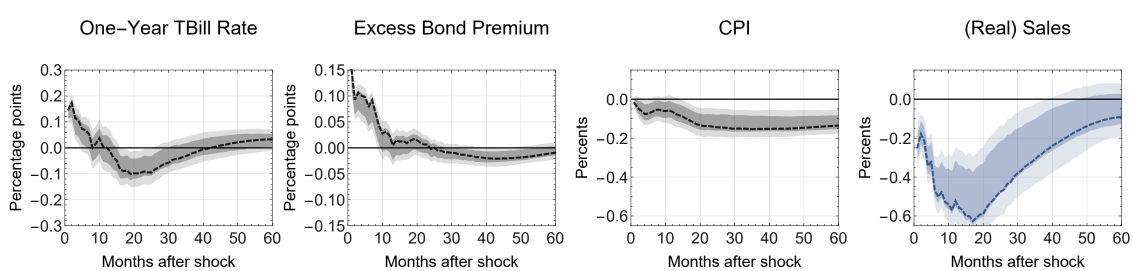

Figure 9: Comparison of Baseline SVAR to GK15 for the original variables.

on the policy instrument ff4_tc to obtain the predicted policy shocks. The data range for the instrument ff4_tc is limited and narrower and ranges from 1990:m1 to 2012:m6. Data for overlapping variables (GS1, EBP and the CPI) are taken from GK15's replication file.

C.1.1 Comparison to GK15

In this section, we show how the results from our SVAR compare to the original GK15 specification. To that end, we report results from three different SVARs and focus on the original set of variables in GK15:

A. The original GK15 specification with industrial production (in place of real sales); specifically, variables Y_1, Y_2, Y_3, and Y_4 replaced by industrial production as in GK15.

B. The original GK15 specification with real sales replacing industrial production; specifically, this SVAR features a limited set of variables: Y_1, Y_2, Y_3, and Y_4 as stated above.

C. The baseline SVAR specification that includes Y_1, Y_2, Y_3, Y_4, Y_5, Y_6 but no gross margin (Y_7) (inventory impulse response for Y_5 and Y_6 is note reported).

D. The baseline SVAR as in the paper. Figure 2 in the paper shows the other impulse responses.

The results of this exercise are shown in Figure 9. As we can see, sales are slightly more volatile, as they exclude the less responsive mining and utilities sectors, but the dynamics are similar across all cases. Figure 9 (panel D) reports individual impulse responses for GS1 and EBP (in the paper we reported 'GS1+EBP', given this is what we used in calibration)

C.1.2 SVAR with Real Wages

This extension additional includes real wages in the SVAR as laid out above. The results are shown in Figure 10.

C.1.3 SVAR with Margin Series Interpolated using ChowLin Method

This extensions replaces gross margin series interpolated linearly in the paper by series interpolated using the ChowLin method. We use payroll series for manufacturing and trade industries to interpolate within quarter monthly values. This does not fix the delayed timing of earnings releases, and to remedy this issue we use a moving average of the obtained series over a 3 month period (average of the current month and the next two months). The results are similar except for size of the response, which is smaller. This is shown in Figure 11.[43]

[43] Replication code for ChowLin interpolation is in the folder TDReplicate (execute the file markup_to_monthly.m). It uses the Matlab replication package for Quilis (2018).

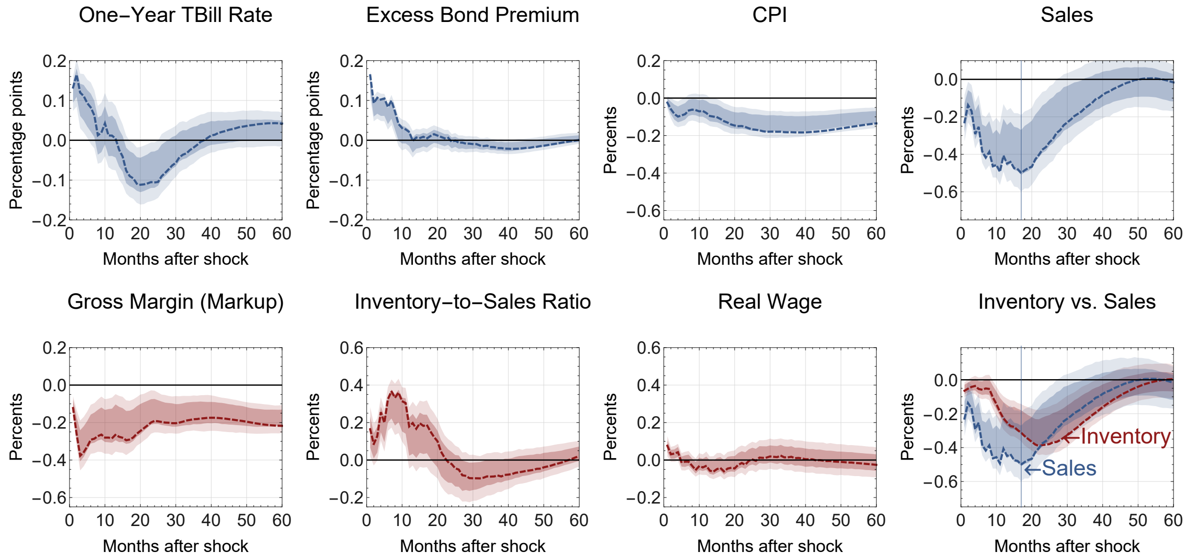

Figure 10: Extended baseline SVAR with real wage.

Notes: Notes to paper's Figure 2 apply.

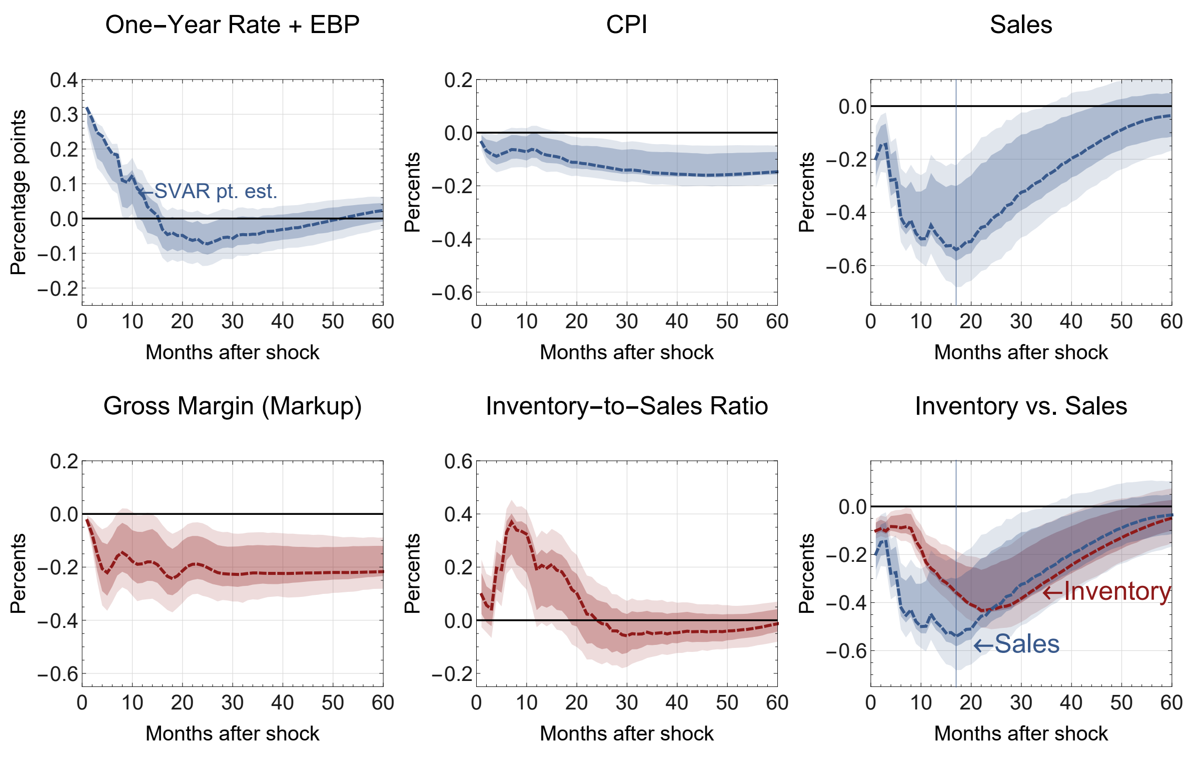

Figure 11: Baseline SVAR with gross margin interpolated using ChowLin method.

Notes: Notes to paper's Figure 2 apply.

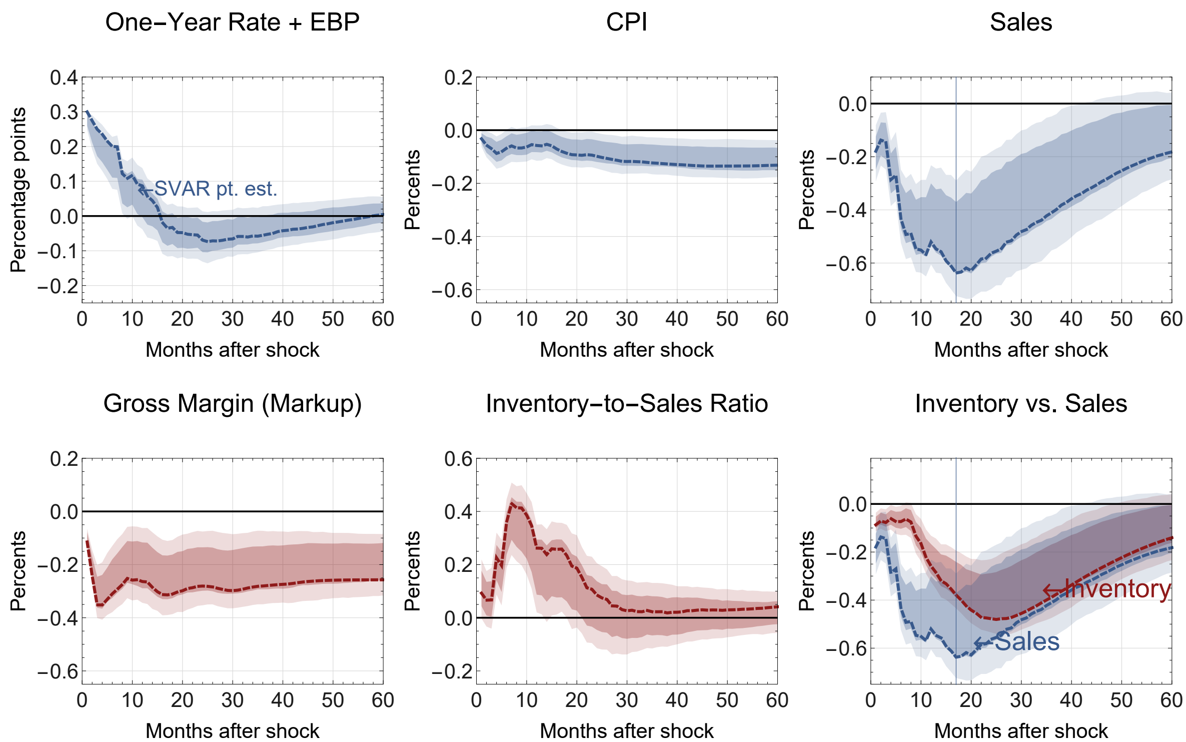

Figure 12: Baseline specification of SVAR with margin series for all firms (ex. FIRE).

Notes: Notes to paper's Figure 12 apply.

C.1.4 SVAR and Figure 1 with Margins for All Firms

This extensions constructs gross margins using data for all firms, except for firms in finance, insurance and real estate sectors (FIRE). The results are shown in Figure 12 and they are almost identical. Figure 13 shows unconditional evidence as in Figure 1 in the paper for all firms (ex. FIRE).

D Model Extension: Low Search Cost Calibration

This section explores the relationship between calibration targets and overall search costs, which amount to 7 percent of the value of produced retail goods in our baseline calibration. To do so, we consider an alternative calibration that excludes the inventory-to-sales ratio target and instead targets search costs at half the value assumed in the paper: 3.5 percent of the final retail price of a good. The assumed parameter values are detailed in Table 6, and the replication code for this extension is available in the *Mathematica* notebook `solve_Model_LS.nb`.

As shown in Figure 14, the results are almost identical. However, since this calibration omits the inventory-to-sales ratio as a target, that ratio drops to 0.5—significantly at odds with the data. This alternative calibration thus highlights a conflict between these two targets. In other words, our model,

A. Great Recession:

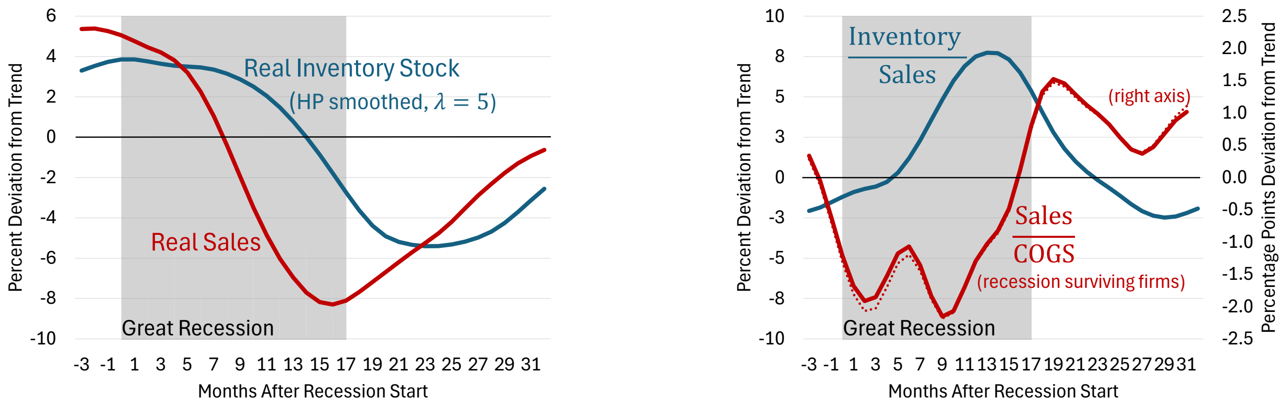

B. Average for U.S. recessions, 1979–2007:

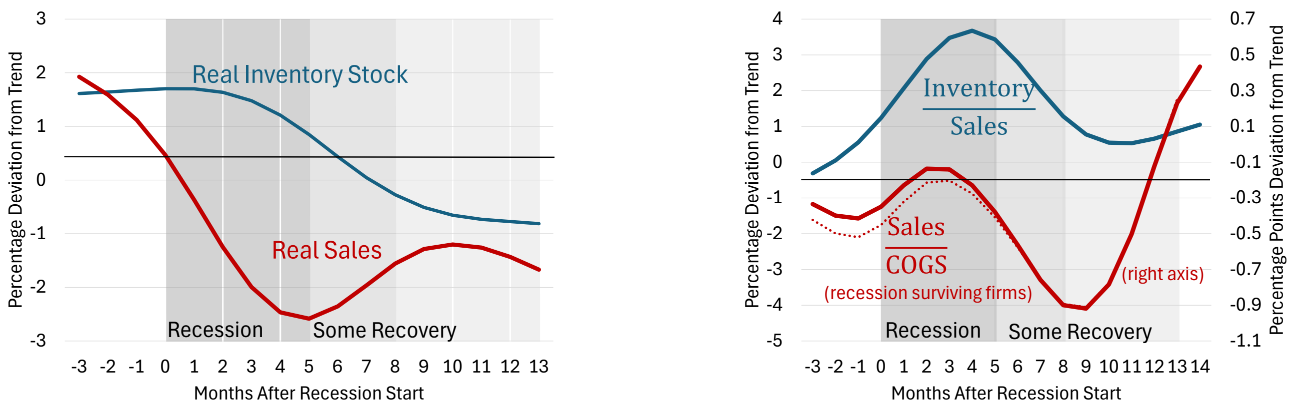

Figure 13: U.S. inventory, sales, and gross profit margins for all firms, 1979–2011.

Notes: The figure shows deviations from HP trend using a smoothing parameter of $\lambda = 10,000$—with additional smoothing applied to deviation (HP filter with $\lambda = 5$). Real inventory and sales data are plotted for the manufacturing and trade industries. The sales-to-COGS ratio is derived from Compustat Quarterly Fundamentals (North America) as described in Sesction 2 and here includes all firms (ex. FIRE) with positive sales right before each recession and one year after each recession (recession-surviving firms). Dotted lines incorporate output elasticities to map margins onto implied markups accordingly to the methodology developed by De Loecker et al. (2020). Shaded areas represent the duration of each recession. A detailed list of data sources is in the last section of this Online Appendix.

as specified, lacks the flexibility to simultaneously target this ratio and deliver lower search costs. We conjecture that introducing a form of congestion that reduces the productivity of search could provide the model with enough flexibility to match both targets.[44] The results shown here and our understanding of the model's mechanism is that such a modification would not affect inventory dynamics.

Are search costs in the range of 7 percent excessive for the entire supply chain? It is difficult to say

[44]This conjecture is based on the following reasoning: search costs are necessary to prevent excessive customer searches, ensuring that outposts spend less time in the production state relative to the marketing state. This balance is required to achieve the inventory-to-sales ratio target of 1.5. A reduction in the marginal productivity of search would produce a similar effect, lowering the search costs implied by the model.

Table 6: Parameter Values.

η_0	c_0	δ	ϕ	χ	τ	(ζ_0, ζ, σ)	ρ^{ss}	Γ	Θ
0.040	0.018	4.30×10^{-4}	1.12	0.69	0.5	$(0.137, 0.0, 0.64)$	7.940×10^{-3}	0.77	0.04

because search costs cannot be directly measured. However, as a reference, consider the calibration in Kaplan and Menzio (2016), who report that consumer search costs at the retail level alone amount to about two percent of the value of the good sold. A number within this range is consistent with other studies that focus on consumer search. Input-output tables indicate that well over two-thirds of transactions in the manufacturing sector are B2B transactions, and the process and risks involved in supplier selection suggest that search costs are substantial across the entire value chain.[45] Scaling the estimate from Kaplan and Menzio (2016) to account for the full supply chain brings search costs close to the level implied by the model. Therefore, while on the high side, we conclude that the 7 percent estimate falls within the range of model-based estimates of search frictions in the literature.

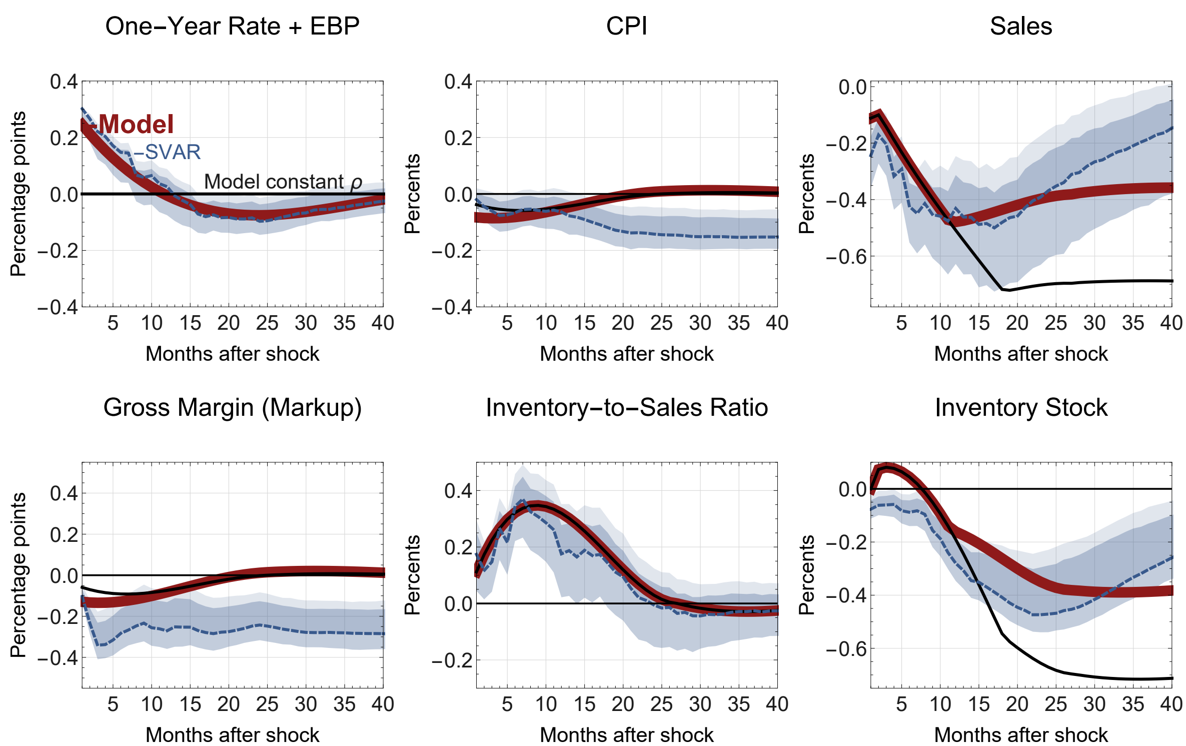

Figure 14: Results from Model with Low Search Costs (3.5 percent).

Notes: Notes to paper's Figure 6 apply.

[45] Beil (2011) provides an overview and discusses best practices in supplier selection.

E Replication Files and Procedures

Supporting codes and data files can be found in the online replication package 'DTreplication.zip'. The replication package uses the following software: *Mathematica*, Stata, Excel, and Python in the Jupyter Notebook environment (version 3.5). Some of the Excel files require HP filter toolbox written by von Kurt Annen and available at `https://econpapers.repec.org/software/dgeqmrbcd/165.htm`. To run replication files, the user must manually adjust the path to source files as described below. The contents of this package can be used under the standard GNU license GNU license agreement.

E.1 Replication of Main Results

The replication package contains four folders: 'Intro_Figure1', 'Model_Replication', 'IO_Calibration', and 'CL_Replicate'. Below we describe their content and the replication procedures:

Folder: `ModelReplication`

This folder contains codes that replicate all results reported in the paper's **Section 2** (Data), **Section 4** (Quantitative Analysis), and the results in the Online Appendix (Sections C and D). The folder contains the following files with codes/data:

- ***Mathematica* notebook `solve_Model_main.nb`**:

  This file generates the results reported in Sections 2 (Figure 2) and 4 (Figure 4 and 6). The notebook also contains supporting calculations for Sections 3.2, and Appendix Sections A and B. The notebook uses input CSV files found in the `/graphics/` subfolder and places all output files (PDF files) in that folder.

  To run the code, adjust the path at the top of the code so that it points to the `/graphics/` subfolder of the replication package. Input CSV files contain SVAR impulse responses and confidence bands and can be replicated by running `est_SVAR.py` (as described next). **Table 1** (Parameters) is also generated by the notebook (last section of the code).

- **Jupyter notebook `est_SVAR.ipy`**:

  This file estimates the baseline SVAR and generates CSV files found in the `graphics` folder for `solve_Model_main.nb`. Adjust the path in the first input cell so that it points to the `ModelReplication` folder of the replication package. The main data file, `dataM.csv`, contains time series for SVAR estimation except for the gross margin series, which is generated from raw data found in `collapsed_output_inv_csv.csv` (and placed in subfolder `./input/`). This file has been created by Stata do files `prep_data.do` and `proc_data_inv.do` described below.

- **Stata do files `prep_data.do` and `proc_data_inv.do`** (in /stata_replication/):

  These files construct the gross margin series as described in Section 2 from the raw S&P Compustat Quarterly Fundamentals dataset (AllDataQrtly7824.dta). The raw Compustat dataset file is not provided here due to copyright restrictions. Instructions for obtaining the data from WRDS are detailed in the next section. To replicate results, obtain the source file, rename it to AllDataQrtly7824.dta and adjust the paths accordingly.

Folder: Intro_Figure1

This folder contains Excel and supporting data files required to generate **Figure 1** in **Section 1**. To regenerate the source data, follow the steps:

1. **Pull Data from Compustat by Running Stata Scripts:**
   - Run proc_data_inv_intro.do (must run prep_data.do once prior to running proc files), located in the Model_Replication folder. Output will be saved in the ./input/ folder within Intro_Figure1 as collapsed_output_inv_fig_csv.
2. **Seasonally Adjust Constructed Data:**
   - Run Jupyter notetebook SeasonalAdjFigureIntro.ipy (in ./intro_Figure1/ folder) to apply seasonal adjustments (ensure the path is correctly set). Output file will be placed in the ./output/ subfolder.
3. **Update Excel file FigureIntro.xls:**
   - Open the Excel file in Intro_Figure1. On the first sheet (MarkupsFromCompustatRawData), adjust the path in cell **C1**.
   - Press the **Refresh** button or manually run the macro LoadDataMarkups in Module1 (press **Alt+F11** for Visual Basic, or go to the **Developer** tab and select **Macros**)

Folder: IO_Calibration

This folder contains miscellaneous files used in calibration (Steady State Targets), including input-output tables for **Appendix Tables 3** and **2**, and BEA's inventory and sales series (Inventories_Sales.xls).

E.2 Replication of Supplementary Results

Next, we describe additional files and the replication of results in this Online Appendix. These results are generated by derivatives of the main codes described above:

- ***Mathematica* notebook `solve_Model_LS.nb`**: This generates results in Appendix Section D (Appendix Figure 14). Run analogously to the main file.

- **Jupyter notebooks `est_SVAR_XXX.ipy`**: These generate raw data for Appendix Section C. Output files are placed in `./graphicsXXX/` subfolders. Run analogously to the main file. To replicate Appendix figures, run `gen_SVAR_figs.nb` in each `./graphicsXXX/` subfolder. The subfolders contain the following results:

  - `./graphicsGK/`: Results for the original GK15 SVAR specification with industrial production (Appendix Figure 9—panel A).
  - `./graphicsGKs/`: Results for the GK15 specification with industrial production replaced by real sales (Appendix Figure 9—panel B).
  - `./graphicsE1NoMup/`: Results for the baseline SVAR excluding the margin variable (Appendix Figure 9—panel C). (Panel D is the baseline model and can be replicated by running `gen_SVAR_figs.nb` in the `/graphics/` subfolder.)
  - `./graphicsE1/`: Results for the baseline SVAR with gross margin series interpolated using the Chow-Lin method (Appendix Figure 11).

- **Replicating the Baseline SVAR with Real Wage Series (Appendix Figure 10)**: Run `ext_SVAR` and set `solve_SVAR_version.ipy` to 3 in input cell `In[25]`. Adjust path information in the first input cell. Output files are in `graphics/SVAR_with_real_wage/`. Run `gen_SVAR_figs.nb` in that folder to generate the figure.

- **Replicating the Baseline SVAR with Margin Series for All Firms (Appendix Figure 12)**: Run the Jupyter Notebook `ext_SVAR_all`. Adjust path information as noted. Output files are in the `./graphicsAll/` subfolder. Run `gen_SVAR_figs.nb` in that folder. Source files are in `./inputAll/` and were generated by the Stata do file `proc_data_all.do` in `./stata_replication/`.

To replicate the monthly interpolation using the ChowLin package, download the replication package for Quilis (2018) from the Journal of Monetary Economics data service. Replace the file `run_mainfile.m` and the `./csvfiles/` subfolder of the downloaded package in `./replication_files_jme_public/data` using the contents of the `CL_replication` subfolder in `ModelReplication` folder. Run the file `run_mainfile.m` after adjusting the path to ensure it points to the `./csvfiles/` folder. Output file series are in `/csvfiles/output/monthly_Mup.csv`. The generated series needs to be manually copy-pasted to `ModelReplicate/input/data_M.csv` file (in column `logMup100chowlin`) containing data for the SVAR estimation. The notebook `example_Sec6.nb` details the calculations for the back-of-the envelope discussed in Section 5 ("Mismeasurement of Variable Costs", footnote 37).

F List of Raw Data Sources

Overlapping series are sourced from the replication files of Gertler and Karadi (2015) for consistency. This includes: CPI, Gilchrist-Zakrajsek Excess Bond Premium (EBP), One-Year T-bill rate (GS1), and the monetary policy shock instrument FF4 (ff4_tc). These files were downloaded from the AEA website replication file for Gertler and Karadi (2015). Data on sales (saleq) and COGS (cogsq) come from S&P Compustat Quarterly Fundamentals (North America), accessed via Wharton Research Data Services. Other series were downloaded from the FRED II service of the Federal Reserve Bank of St. Louis in December 2023. These include: (1) real inventories and in manufacturing and trade industries (Real Manufacturing and Trade Inventories, Millions of Chained 2017 Dollars, Monthly, Seasonally Adjusted, and Real Manufacturing and Trade Industries Sales, Millions of Chained 2017 Dollars, Monthly, Seasonally Adjusted), originally sourced from various series published by Bureau of Economic Analysis and compiled by the FRED service of the Federal Reserve Bank of St. Louis to build long time series;[46] (2) payroll employment in manufacturing (All Employees, Manufacturing, Thousands of Persons, Monthly, Seasonally Adjusted), originally sourced from the Bureau of Labor Statistics. Most raw files—except for Compustat-derived series—are included in the replication package (Excel files contain metadata). Processed output pulled from the Compustat dataset is provided, but raw data cannot be shared. To replicate these results, download the Compustat Quarterly Fundamentals (NA) with all main variables for 1970-2015 from Wharton Research Data Services (WRDS) (or similar data outlet) and run the Stata replication code as described in the previous section. Output elasticities come from the online replication package of De Loecker et al. (2020) (downloaded from the personal website of Jan Eckout).

[46]Federal Reserve Bank of St. Louis, Real Manufacturing and Trade Industries Inventories [INVCMRMT] and Sales [CMRMTSPL], retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/CMRMTSPL. The source data are U.S. Bureau of Economic Analysis Real Manufacturing and Trade Industries Sales for various time periods.