Order Backlog*

Bjarne Horst Matthias Meier Jan Schymik

September 2026

Abstract

Many firms receive orders months before they deliver the corresponding output, creating a backlog of unfilled orders. Using German administrative manufacturing microdata, we document widespread backlog and an average *time to fill*, from order to sale, of six months. Backlog renders sales prices predetermined by past order agreements, challenging conventional sales-based identification of firm demand. We construct theoretically coherent order-book prices and use them to identify firm-level demand and supply shocks. We show, empirically and theoretically, that backlog shapes the transmission of such shocks: it rises after expansionary demand shocks and falls after expansionary supply shocks, thereby dampening the new order price response to both shocks on impact.

Keywords: Order books, backlog, time to fill, prices, supply shock, demand shock.
JEL Codes: D21, D22, D25, E23, E31.

*We thank Klaus Adam and Christian Bayer, as well as participants at various seminars and conferences, for many helpful comments. We are grateful for excellent cooperation with the Federal Statistical Office of Germany that facilitated this project, in particular Director Daniel Vorgrimler for establishing the cooperation, and Hanna Brenzel, Benedikt Gerst, and Bayram Oruc for their vigorous technical support. Bjarne Horst, Matthias Meier, and Jan Schymik acknowledge support from the German Research Foundation (DFG) through CRC TR 224 (projects C02, C05, B07). Bjarne Horst acknowledges support from Stiftung Geld & Währung. Matthias Meier acknowledges funding from DFG Projects 550469881/550470072 and from the UniCredit & Universities Foundation. Jan Schymik acknowledges funding from DFG Project 557666971.
Bjarne Horst (University of Mannheim): bjarne.horst@uni-mannheim.de. Matthias Meier (University of Mannheim): m.meier@uni-mannheim.de. Jan Schymik (University of Mannheim): jan.schymik@uni-mannheim.de.

1 Introduction

Order backlog is the stock of accepted but unfilled delivery agreements. Backlog links current sales to past orders, and current sales prices to past order prices. It creates a time lag between new orders and sales, which we refer to as *time to fill*. Surprisingly, order backlogs have received little attention in economics, despite their prominence in the policy and public debate.[1] This paper addresses the gap and shows that order backlog is large, creates challenges for demand identification, and shapes the transmission of shocks.

This paper makes three contributions. First, we introduce novel administrative manufacturing microdata on order books in Germany. We document that backlog is widespread and sizable with an average time to fill of half a year. We further document that orders commonly constitute commitments to future sales price and quantity. Second, we show that predetermined sales prices fundamentally challenge a conventional class of sales-based strategies to identify firm demand. Instead, we construct theoretically coherent order prices that permit valid identification of firm demand and supply shocks. Third, we show that backlog matters for the propagation of shocks. We provide empirical evidence and a theoretical model to show that backlog rises after expansionary demand shocks but falls after expansionary supply shocks. Relative to a counterfactual without backlog adjustment, backlog dampens the sales response to demand shocks, amplifies the sales response to supply shocks, but dampens price responses to both shocks.

The empirical backbone of the paper is novel data, the Monthly Reports of Manufacturing Establishments, which are highly granular and well suited to study order backlog.[2] The data are collected by the Federal Statistical Office of Germany (Destatis).[3] Reporting is legally mandatory for all German manufacturing establishments with at least 50 employees. The reports track sales, new orders, and order backlog (all nominal) at the level of four-digits business segments within establishments, that we denote by establishment^{4d}. The dataset is a panel of several million establishment^{4d}-month observations from 2014 through 2022. Backlog data are collected in sectors for which backlog is deemed important. These sectors

[1] In many countries, the press regularly reports new readings of order backlog based on data from statistical agencies or surveys among purchasing managers. In addition, reports by economic advisory bodies, such as the SVR in Germany and the CEA in the US, frequently refer to order backlogs in their analyses.

[2] The US Census M3 collects similar order book data. However, while the M3 is a relatively small survey of firms or business segments (e.g., Schuh, 1996), our dataset covers nearly the universe of manufacturing establishments. US Compustat reports annual backlog of public firms (e.g., Rajgopal et al., 2003; Chang et al., 2018; Meier, 2020), but backlog is frequently missing, even in industries where backlog is known to be important. For Germany, the ifo survey also collects information on new orders, sales, and backlog, but these are mainly qualitative appraisals and the sample size is limited.

[3] Destatis uses the reports to construct various aggregate economic series but the data has been unexplored in research. Destatis does provide some manufacturing microdata to researchers (AfiD), which builds on the reports, but omits information on new orders and order backlog, and aggregates to establishments and years.

represent 69–76% of sales, employment, and value-added in aggregate manufacturing.[4]

A key implication of backlog is time to fill: the time it takes for new orders to pass through the backlog and turn into sales. In a sense, time to fill measures backlog in units of time. We construct two complementary measures of time to fill: an ex-ante measure that can be computed in real time and an ex-post measure capturing realized time to fill. The two measures are statistically quite similar. Across establishment[4d]-months, the average time to fill varies between 5 and 7 months, depending on measurement and weighting. The distribution has a pronounced right tail: time to fill exceeds one year for a tenth of our sample. While time to fill differs across sectors, most variation is within narrowly defined sectors, both between establishments[4d] and within establishments[4d] over time.

The order book data further indicate that new orders typically commit firms to future sales quantities and prices. Given the data design, order backlog, new orders, and sales satisfy a stock-flow equation up to a residual that captures cancellations, ex-post price changes, and measurement error. This residual has a mode of zero and is small for most observations, consistent with cancellations and ex-post price adjustments being the exception rather than the rule. The residual also does not meaningfully comove with price growth, further suggesting a limited role for ex-post price changes. Note that we do not argue that new orders are commitments to a specific delivery period. To the contrary, we provide evidence showing that the speed of order fulfillment is a rather flexible margin of adjustment.

Given commitment, the prices agreed in new order contracts (for short, new order prices) become sales prices in the future. Hence, for any establishment[4d]-month, the price of new orders can differ from the price of sales. That creates a measurement challenge because we only observe sales prices, from linked product-level production data, but new-order prices are not directly observed. While the challenge has long been recognized (e.g., Foss et al., 1981; Haltiwanger and Maccini, 1989), we propose a novel, theoretically coherent procedure that recovers all order book prices under a *first in, first out* assumption on the backlog, allowing for heterogeneity in time to fill across and within establishments[4d].[5] We document that sales prices differ widely from the computed new order prices, highlighting the importance of accounting for these price differences.

The price dynamics challenge the identification of demand. A conventional class of identification strategies infers the demand curve from the co-movement of current sales quantity and price, addressing endogeneity with suitable instruments.[6] In the presence of backlog,

[4] Backlog is prevalent well beyond manufacturing. In 2015–2022 US Compustat data, non-missing backlog averages more than six months of sales in several sectors: mining, utilities, construction, manufacturing, transportation, information, professional, administrative services. These sectors account for 53% of US GDP.

[5] Our procedure is related to perpetual inventory methods used to construct real capital stocks.

[6] Prominent examples for this class of identification strategies include Bresnahan (1981), Berry (1994),

however, as sales prices inherit past order prices, they need not reflect contemporaneous demand. Instead, current demand is defined by contemporaneous new order price and quantity. We characterize the bias that arises when inferring demand based on current sales quantity and price and show that endogenous time to fill makes bias particularly acute. To circumvent this problem, we identify establishment[4d]-level demand shocks in the data by regressing real new orders on new order prices, addressing endogeneity with a Bartik instrument based on imported input prices. We further identify supply shocks from sales surprises orthogonal to demand shocks. Estimating an analogous sales-price regression yields substantially flatter demand curves and the estimated responses to the associated shocks differ quantitatively and qualitatively.

We use the identified demand and supply shocks to empirically characterize the role of order books for their propagation. We find that expansionary demand shocks raise new order prices and new orders, sales rise, but less strongly on impact, and backlog increases. Hence, firms absorb part of the demand shock through higher backlog, and thus spread higher demand over time. Expansionary supply shocks also raise new orders, but lower new order prices and lower backlog. Lowering backlog allows the firm to raise sales more strongly in response to the supply expansion. The responses of backlog are economically significant and show that backlog is an important margin of adjustment. The responses further underline the importance of understanding the drivers of backlog, since high backlog can reflect expansionary demand or adverse supply conditions.

To rationalize the evidence and explore the role of backlog, we develop a tractable model in which backlog is a dynamic margin of adjustment, alongside the new order price and sales. The model trades off a production smoothing motive with the cost of backlog adjustment. We show analytically that the model qualitatively explains our empirical responses to both shocks. Despite being highly stylized, the model can also quantitatively match the empirical responses to demand and supply shocks well.

To analyze the role of backlog for shock propagation, we compare the baseline model with a counterfactual in which the costs of backlog adjustment are excessively high. After a demand shock, the increase in backlog allows the firm to accept new orders while letting production rise by less, dampening the increases in production costs. Absent backlog adjustment, the firm hikes order prices to suppress part of the increase in demand in order to moderate the increase in production costs. Backlog adjustment therefore dampens the sales and order price responses but amplifies new orders in response to demand shocks. After a supply shock, the firm draws down backlog to exploit lower production costs. Without this margin, it must cut order prices more to attract new demand. Thus, the backlog decrease dampens the response

Berry et al. (1995), Nevo (2001), Foster et al. (2008), and Hottman et al. (2016).

of new order price and quantity to supply shocks but amplifies the sales response.

Our paper contributes to a small, existing literature studying order backlog. Classical papers (e.g., Zarnowitz, 1961, 1962; Childs, 1967; Belsley, 1969; Zarnowitz, 1973) empirically study the co-movement in aggregate order book data. Zarnowitz (1962) further hypothesizes that operating backlog is favorable if output is costly to store and demand is specific and uncertain. We provide new micro evidence, including evidence in support of the Zarnowitz (1962) hypothesis, and study identification and propagation. Classical work modeling order backlog (e.g., Belsley, 1969; West, 1989; Haltiwanger and Maccini, 1989) treats backlog analogously to inventories. We instead emphasize that backlog is a collection of past price and quantity commitments with implications for price measurement and shock identification. We also contribute to a more recent literature that studies delays in production (Antràs and Tubdenov, 2025), in the supply of capital goods and construction (e.g., Meier, 2020; Oh and Yoon, 2020; Fernandes and Rigato, 2025), in transportation industries (e.g., Benkard, 2004; Kalouptsidi, 2014), and their propagation through production networks (Schaal and Taschereau-Dumouchel, 2025; Antràs and Kulesza, 2026). Our key contribution to this literature is to provide new stylized facts and findings on the identification and propagation of demand and supply shocks.

Our paper further relates to a literature on inventories (e.g., Cooper and Haltiwanger, 1990; Ramey and West, 1999; Khan and Thomas, 2007; Kryvtsov and Midrigan, 2012; Görtz et al., 2024). Backlog is related but distinct from finished goods inventory. Firms that produce to stock use inventory to store past supply and fill future demand. In contrast, firms that produce to order use backlog to store past demand and fill it with future supply. In other words, inventories are goods waiting for buyers; backlogs are buyers waiting for goods. This distinction matters because backlog, unlike finished-goods inventories, renders sales prices predetermined by past orders. In the manufacturing sectors reporting backlog data, backlog is several times larger than finished goods inventories. Hence, production typically starts after an order, consistent with production to order. Finally, we show that finished goods inventory and backlog have qualitatively different implications for shock propagation. While backlog dampens the sales response to demand shocks, but amplifies the sales response to supply shocks, finished goods inventory has the opposite effect on sales.

The remainder of the paper is organized as follows. Section 2 describes our data and documents the prevalence of order backlogs. Section 3 introduces theoretically coherent order book prices. Section 4 presents empirical evidence on time to fill. Section 5 discusses the empirical identification and propagation of demand and supply shocks. In Section 6, we introduce a theory of a firm with order backlog, quantify it and evaluate counterfactual shock propagation absent adjustments to order backlog. Finally, Section 7 concludes.

2 Microdata on order books

This section introduces novel microdata on order books in Germany. We document that backlog is prevalent across manufacturing sectors. The data show that establishments with backlog commonly produce to order and that backlog represents commitments to future sales quantities and prices.

2.1 Monthly Reports of Manufacturing Establishments

We introduce the *Monthly Reports of Manufacturing Establishments*, a novel dataset that has previously not been available to researchers. The reports are administrative, high-frequency, highly granular, and contain order book information.

The reports are collected by the Federal Statistical Office of Germany (Destatis). Submitting monthly reports is a legal obligation for all manufacturing establishments in Germany with at least 50 employees in the reporting year. Establishments have to submit separate monthly reports for each four-digit manufacturing product market in which they are active. We denote these within-establishment business segments as establishments^{4d}. We have access to the (anonymized) reports from January 2014 through December 2022. The reports go much further back in time, but backlog was added to the reports starting in January 2014. Since order backlog data was often missing or implausible in 2014 (Linz et al., 2016), we use an effective sample from January 2015 through December 2022.[7]

The reports contain the monthly nominal value of sales, new orders, and order backlog at the establishment^{4d}-level.[8] New orders are the nominal value of all firmly accepted new orders for future delivery in the reporting month. Order backlog is the nominal value of all firmly accepted orders at the end of the reporting month that have not yet resulted in sales and have not been canceled. Sales are invoiced revenues in the reporting month.[9] For details on the legislative background and coverage of the reports, see Appendix A.1.

2.2 Sector-level prevalence of backlog

Order backlog is reported only in manufacturing sectors in which backlog is deemed relevant by Destatis. These *backlog sectors* are textiles, cloths, paper, chemicals, pharmaceuticals, primary metals, fabricated metals, electronics, electrics, machinery, automotives, and other vehicles, see Appendix A.2 for details.

[7]We use 2014 observations only to construct lags of non-backlog variables.

[8]Destatis uses the reports in the production of aggregate indices of sales, new orders, and order backlog, all closely monitored business cycle indicators, and in the National Income and Product Accounts.

[9]Sales are recorded in the month they are invoiced, not the month payments are received.

Table 1: Sectoral prevalence of order backlog

	Sales (% of total)	Employment (% of total)	Value added (% of total)	Backlog /Inventory
Backlog sectors	72.5	69.0	76.3	7.0
– *Textiles, Cloths, Paper*	2.9	3.6	3.0	3.4
– *Chemicals, Pharma*	10.7	7.2	10.9	2.8
– *Primary, Fabricated Metals*	10.5	14.4	11.8	6.0
– *Electronics, Electrics*	9.7	12.1	13.4	7.0
– *Machinery*	13.0	16.5	16.9	6.4
– *Automobiles*	23.3	13.1	17.7	18.4
– *Other Vehicles*	2.3	2.0	2.6	21.7
Non-Backlog sectors	27.5	31.0	23.7	–

Note: The sample spans 2015M1–2022M12. Sales (% of total) provides the sector-specific share of total manufacturing sales, and analogously for employment and value added. Backlog/Inventory is the average annual ratio of sectoral backlog divided by sectoral inventory of finished goods.

Overall, backlog is a highly prevalent feature of manufacturing. Table 1 shows that backlog sectors account for 72% of total sales, 69% of employment, and 76% of value added in the manufacturing sector. The two backlog subsectors with the largest value added are the machinery and automobile sectors. In addition, backlog sectors represent the more volatile and more cyclical segment of manufacturing.

Backlog arises naturally when firms produce to order: new orders first enter the backlog and are fulfilled through future production. Firms that produce to stock, in contrast, produce before the order, store output as finished-goods inventory, and fill new orders from inventory (Belsley, 1969; Zarnowitz, 1962). In reality, however, many establishments[4d] may operate with both backlog and finished-goods inventory. Table 1 shows that order backlog is on average 7 times larger than finished-goods inventory across backlog sectors.[10] Moreover, backlog exceeds inventory in all backlog subsectors, and finished-goods inventory amounts

[10] We obtain the sum of finished and semi-finished goods inventory for backlog sectors from an annual Destatis survey of manufacturing firms, the Cost Structure Survey. To disentangle finished and semi-finished goods inventory, we compute the share of finished goods inventory for corresponding two-digit manufacturing sectors in the U.S. Census M3 Survey. Across backlog sectors, this share is 0.5 on average. In addition, the survey only contains a subset of the establishments[4d] present in the reports. Naively comparing aggregate inventory in the survey with aggregate backlog in the reports would be misleading. Instead, we compute the ratio of inventory to sales in the survey, the ratio of backlog to sales in the reports, and compute the last column in Table 1 as a ratio of ratios, adjusted for the finished-goods inventory share.

to half a month of sales on average. These facts suggest that firms in backlog sectors predominantly produce to order. The large size of backlog relative to inventories is also worth highlighting given the much larger literature on inventories than on backlog.

2.3 Micro-level backlog and commitment

Having established the prevalence of backlog across manufacturing sectors, we now zoom into the micro data, within the backlog sectors. We first describe the establishment$^{\text{4d}}$-month sample used in the remainder of the paper and document stylized facts about sales, new orders, and backlog. We argue that the micro data is consistent with backlog commonly representing commitments to future sales quantity and price.

The order backlog follows a stock-flow equation. Denoting by B_{it} the end-of-month t nominal value of backlog of establishment$^{\text{4d}}$ i, by O_{it} the nominal value of new orders, and by S_{it} the nominal value of sales, the following (nominal) stock-flow equation must hold,

$$B_{it} = B_{it-1} + O_{it} - S_{it} - R_{it}, \tag{2.1}$$

where R_{it} is not reported but implicitly defined by changes in B_{it} that are not explained by O_{it} and S_{it}. We will discuss R_{it} in detail further below in this section.

For some establishments$^{\text{4d}}$, we observe for most months that backlog is negligible relative to sales, amounting to less than two weeks of sales, sales and new orders are near-identical, or sales exceed new orders by a factor of five or higher. The former two cases may indicate that backlog is irrelevant in magnitude or constant across time. All three cases may indicate incomplete reporting. We exclude these observations from our analysis. The excluded observations account for 27% of observations and are similar in sales and employment as the included observations. The resulting sample, which we refer to as *final backlog sample*, still accounts for more than half of manufacturing sales and employment. For further details, see Appendix A.2.

Table 2 provides descriptive statistics for the final backlog sample, which contains 1.3 million establishment$^{\text{4d}}$-month observations from 2015 through 2022. Median monthly sales and new orders are about 1 million Euro, median order backlog is 3.7 million Euro, all expressed in 2019 prices. The averages mask substantial heterogeneity across establishments$^{\text{4d}}$. For example, monthly sales range from 0.1 million Euro at the 10th percentile to 7.4 million Euro at the 90th percentile. Table A.2 further shows that most firms and establishments have precisely one establishment$^{\text{4d}}$.

Importantly, the descriptive statistics on the residual are informative about the extent to which new orders represent commitments to quantity and price. Deviations from commit-

Table 2: Descriptive statistics for the final backlog sample

	Mean	Sd	P10	P25	P50	P75	P90	Obs
Sales (S_{it})	3.31	7.16	0.10	0.40	1.07	2.92	7.38	1,315,354
New orders (O_{it})	3.46	7.63	0.07	0.36	1.06	3.02	7.74	1,315,354
Order backlog (B_{it})	19.96	55.01	0.16	0.93	3.74	12.95	42.19	1,315,354
Order book residual (R_{it})	0.05	1.00	-0.26	0.00	0.00	0.02	0.36	1,315,354
Relative residual ($\|R_{it}\|/B_{it-1}$)	0.09	0.25	0.00	0.00	0.00	0.05	0.20	1,256,847

Note: The table provides descriptive statistics at the establishment^{4d}-month level for 2015M1–2022M12. All variables are deflated by the PPI normalized to the year 2019, expressed in EUR million, and winsorized at the 1st and 99th percentile. The last row refers to the absolute order book residual normalized by the order backlog in the previous month.

ment may show up as order cancellations or ex-post price changes. Given the definition of backlog, new orders, and sales (Section 2.1), cancellations and ex-post price changes will generate a non-zero residual R_{it}. Cancellations and price reductions lead to a positive R_{it}, price increases a negative R_{it}.[11] An additional potential source of non-zero residuals is measurement error, which can naturally create positive or negative residuals.[12]

The last two rows of Table 2 show that the residual R_{it} is typically zero or small relative to backlog. In 76% of observations the residual is either zero or corresponds to less than 5% of lagged backlog, limiting the scope for ex-post adjustments in quantity and price.[13] If cancellations were common and dominant, R_{it} should be mostly positive. Instead, it is frequently zero, and almost as often negative as positive, meaning at most half of the non-zero observations for R_{it} represent cancellations. If ex-post price revisions were frequent and dominant, the near-symmetric distribution of R_{it} around zero would be less surprising, but the large mass at and near zero indicates that such price changes are typically small at best. Furthermore, if the residual was dominated by ex-post price changes, it would be systematically linked to price changes. We find this link to be extremely weak, see Section 3.4 for details. Overall, we consider it plausible that accepted orders typically constitute commitments to future sales quantity and price. Note that we do not argue that new orders are commitments to a specific delivery period. To the contrary, we provide evidence showing that the speed

[11] A cancellation in t lowers S_{it}, or B_{it}, or both, and therefore appears as $R_{it} > 0$. New orders in t account for price discounts in t, so only ex-post price changes affect R_{it}: e.g., discounts in the sales period or price escalation clauses indexing the sales price to input prices.

[12] Examples of measurement error are orders not recorded in backlog but resulting in sales, or delayed adjustments of backlog to new orders and sales.

[13] We have 4% fewer observations for $|R_{it}|/B_{it-1}$ than for R_{it} because B_{it-1} is missing for the first observation of an establishment^{4d} and because B_{it-1} occasionally equals zero.

of order fulfillment is a rather flexible margin of adjustment.

3 Order book prices

The previous section documented nominal order book variables. This section shows that the respective prices of sales, new orders, and backlog differ substantially. Sales, new orders, and backlog reflect orders placed at different points in time. In particular, new order prices lead sales prices by several months. The underlying prices are not directly collected. Instead, we tackle a non-trivial measurement challenge by developing a theoretically coherent procedure to construct all order book prices when only sales prices are observable.

3.1 Measurement challenge

A simple example helps to explain the challenge of measuring the respective prices of sales, new orders, and backlog. Suppose orders remain in the backlog for two periods before becoming sales. For an order received in period t and delivered in $t+2$, the sales price in $t+2$ equals the new order price in t, provided orders are commitments. If the new order price changes between t and $t+2$, period $t+2$ prices of sales and new orders differ. Moreover, the price of backlog in $t+2$ can differ from both because it contains new orders placed in $t+1$ and $t+2$.

Formally, the prices of backlog, new orders, sales, and order book residuals are dynamically linked to each other through a *real* stock-flow equation, given by

$$\underbrace{b_{it}}_{B_{it}/p^B_{it}} = \underbrace{b_{it-1}}_{B_{it-1}/p^B_{it-1}} + \underbrace{o_{it}}_{O_{it}/p^O_{it}} - \underbrace{s_{it}}_{S_{it}/p^S_{it}} - \underbrace{r_{it}}_{R_{it}/p^R_{it}}, \tag{3.1}$$

where lowercase letters denote real variables, and uppercase letters the nominal values in (2.1). We refer to the collection $\{p^B_{it}, p^O_{it}, p^S_{it}, p^R_{it}\}$ as order book prices.

In general, order book prices differ even within an establishment[4d]-month, except for two special cases. If new order prices are constant over time, all order book prices coincide. If the firm fulfills new orders immediately, the new order price will equal the sales price. Deflating all variables with an aggregate PPI, as in the descriptive statistics of Section 2, improves the comparability across time but ignores heterogeneous price changes across establishments[4d] and across order book variables.

The ideal dataset would report all four price indices. The manufacturing reports do not, and we are not aware of any dataset that systematically records them.

3.2 A scheme for constructing order book prices

We next provide a scheme that allows constructing all order book prices in a theoretically coherent way based on three assumptions:

> *Assumption 1: The sales price (p^S_{it}) is observable.*

The first assumption is commonly satisfied as statistical agencies do collect sales prices.

> *Assumption 2: Backlog (b_{it}) follows a FIFO (First In, First Out) principle.*

The assumption means that the next order filled, and therefore turned into sales, is always the oldest unfilled order in the backlog. Cancellations captured by the order book residual also apply to the oldest unfilled orders, and we allocate between sales and cancellations in proportion to their relative size. Information on the order in which unfilled orders are filled is necessary to construct the prices of new orders from sales prices. However, we do not observe the filling order directly. That would require even more granular data at the level of individual orders. Instead, we impose the FIFO assumption on the filling order, which we consider the most plausible ordering assumption.

> *Assumption 3: The price of a new order (p^O_{it}) does not change subsequently and remains the price at order fulfillment. New orders (o_{it}) are not fully canceled.*

Assumption 3 puts additional structure on the order book residual. We rule out ex-post price changes, so nominal residuals reflect real flows rather than price revisions. Given that the residual is typically zero or small (Table 2), we view the assumption as quantitatively innocuous.[14] While a positive residual represents cancellations, a negative residual may represent orders that quickly translate into sales without entering backlog, representing a type of measurement error. We further rule out full cancellation of orders. This guarantees the existence of a sales price that includes information on the original new order price. In case the assumption is violated, we will impute the new order price using sales prices associated to new orders in subsequent periods.

Assumptions 1–3 allow us to construct theoretically coherent prices p^O_{it}, p^R_{it}, and p^B_{it} as we describe next:

[14] Alternatively, we could construct order book prices assuming the order book residual reflects only price adjustments or any convex combination between price and real adjustments.

Step 1: Price of new orders (p^O_{it})

$$
\begin{aligned}
p^O_{it} &= \sum_{j=\underline{\tau}_{it}}^{\overline{\tau}_{it}} w_{it+j} p^S_{it+j} \\
\underline{\tau}_{it} &= \min\left\{\tau \in \mathbb{N}^0 \,\middle|\, \sum_{k=0}^{\tau} (S_{it+k} + R_{it+k}) > B_{it-1}\right\} \\
\overline{\tau}_{it} &= \min\left\{\tau \in \mathbb{N}^0 \,\middle|\, \sum_{k=0}^{\tau} (S_{it+k} + R_{it+k}) \geq B_{it-1} + O_{it}\right\}
\end{aligned}
$$

Under Assumptions 2–3, $t + \underline{\tau}_{it}$ is the first period in which i starts depleting period t orders O_{it}. In between t and $t + \underline{\tau}_{it}$, the firm is still depleting the preceding backlog (B_{it-1}). Similarly, $t + \overline{\tau}_{it}$ is the last period in which i still fills period t orders. If an order O_{it} is fully filled in a single period, then $\overline{\tau}_{it} = \underline{\tau}_{it}$. In that case, the new order price is simply $p^O_{it} = p^S_{it+\underline{\tau}_{it}}$. If the order is filled across several periods, we compute the new order price p^O_{it} as a sales-weighted average of sales prices between $t + \underline{\tau}_{it}$ and $t + \overline{\tau}_{it}$, with sales weights

$$
w_{it+j} = \frac{\tilde{S}_{it+j}}{\sum_{k=\underline{\tau}_{it}}^{\overline{\tau}_{it}} \tilde{S}_{it+k}},
$$

where $\tilde{S}_{it+k}$ takes into account the possibility that only a fraction of sales in $t+\underline{\tau}_{it}$ and $t+\overline{\tau}_{it}$ is attributed to the period t order.[15] Appendix B.1 details the construction of $\tilde{S}_{it+k}$. Of course, $\underline{\tau}_{it}$ and $\overline{\tau}_{it}$ are interesting quantities *per se*. They measure the time it takes for an order to pass through the order backlog, as we discuss in more detail in Section 4.

Step 2: Price of order book residual (p^R_{it})

$$
p^R_{it} = p^S_{it}
$$

Under Assumptions 2–3, the order book residuals are real flows that apply to the oldest orders in the backlog. Hence, the price of the order book residual equals the sales price.

Step 3: Price of order backlog (p^B_{it})

$$
p^B_{it} = \frac{B_{it}}{b_{it}}, \qquad b_{it} = b_{it-1} + o_{it} - s_{it} - r_{it}
$$

Given Assumption 1 and Steps 1–2, prices p^S_{it}, p^O_{it}, and p^R_{it} are known. Hence, we can construct real flow variables s_{it}, o_{it}, and r_{it}. The real stock-flow equation (3.1) implies that we can

[15] In case $w_{it+j} \in (0, 1)$, we assign the same sales price p^S_{it+j} partially to two different, adjacent new order prices, potentially creating measurement error in new order prices.

compute b_{it} for all periods given an initial level of the real backlog. To obtain the initial real backlog, we exploit that the real and the nominal stock-flow equation both have to hold simultaneously. Thus, the time to deplete the initial real backlog and the initial nominal backlog has to be the same. That condition uniquely pins down the initial real backlog, as we describe in more detail in Appendix B.2.

3.3 Empirical implementation

We next describe the empirical implementation of the above scheme. We refer to the resulting sample for which order book prices are non-missing as *final price sample*.

The price construction scheme requires as input an index for the sales price p_{it}^S, which we construct using the *Production Census* of manufacturing in Germany (briefly: census). It records the quantity of production and the associated sales value at the level of nine-digit products and per establishment^{4d}-month. The coverage of establishment^{4d}-months almost perfectly overlaps with the manufacturing reports. We compute unit prices as value divided by quantity and aggregate to the establishment^{4d}-month level via Törnqvist indices with sales values acting as weights (cf., Gagliardone et al., 2023). This procedure allows us to construct a sales price index that covers 77.8% of establishment^{4d}-month observations. We cannot construct sales prices for some establishments^{4d} because the census does not elicit production quantities for some products. Appendix B.4 provides further detail on the census and the construction of the sales price index.

Given the sales price index, Steps 1–3 yield price indices for the remaining order book variables. Two practical issues require additional treatments. First, Step 1 requires future order book variables, which become unavailable near the end of the sample. Without adjustment, the final price sample would under-represent observations toward the end of the sample and for establishments^{4d} with long time to fill. To avoid such selection bias, we forecast new order prices whenever step 1 cannot be implemented because it requires observations beyond the end of our sample. For simplicity and transparency, we assume new order prices remain constant after the last period in which they can be computed. The data offer some justification for this assumption; sales price inflation averages close to zero between 2014 and 2019, and producer price inflation had nearly returned to zero by the end of 2022 and stayed close to zero through 2025.[16]

A second practical issue is that Step 3 occasionally yields negative real backlog. Given that nominal backlog is non-negative, this can only arise from our assumptions being occasionally violated. We address this issue by imposing that the real backlog b_{it} cannot be smaller than

[16] The respective end-of-year yoy PPI growth rates in 2022–2025 were 1.6%, -0.4%, 0.1%, and -0.2%.

the nominal backlog B_{it} divided by the largest sales prices observed for establishment[4d] i. Whenever this bound is first violated, we adjust the real order book residual to ensure that the real backlog satisfies the bound with equality. Then we continue iterating the real stock-flow equation, repeating the adjustment if needed. Adjustments in r_{it} that exceed 5% of real backlog in absolute terms represent 2.1% of observations in our final price sample. The resulting real backlog series is weakly positive and equals zero if and only if nominal backlog equals zero.[17] In that case, we consider the price of backlog as missing.

3.4 Evidence on sales and new order prices

The sample for which we have order book prices, and thus real order book variables, comprises nearly one million observations. We refer to this subset of the final backlog sample as the final price sample. Panel (d) of Table A.1 presents descriptive statistics for the final price sample. On average, establishments[4d] in this sample are slightly larger than those in the final backlog sample–by 7% for sales and new orders–but have slightly smaller order backlog. Overall, the final price sample does not appear to be a strongly selected subsample.

Figure 1 highlights the quantitative importance of taking into account that prices of sales and new orders generally differ. Panel (a) shows that the price of sales p^S_{it} and the price of new orders p^O_{it} differ substantially at the micro-level. For 10% of establishment[4d]-months, the order price exceeds the sales price by at least 20%; for another 10%, the sales price exceeds the order price by at least 30%. Using sales prices to deflate new orders would therefore generate large errors in real quantities. Panel (b) shows that these micro price differences also show up in the aggregate. Because sales prices reflect past new-order prices, new-order price growth should lead sales-price growth. Consistent with this pattern, aggregate new order prices rise ahead of aggregate sales prices during the recent inflation surge.[18]

Furthermore, the order book prices allow us to revisit from another angle the question to what extent the residual captures ex-post price changes. Suppose the price of an order is fixed until the sales period, in which it becomes flexible. Further suppose the firm finds it optimal to adjust the sales price p^S_{it} to equal the new order price p^O_{it}. In that case, $R_{it} = -(p^O_{it} - p^O_{it-\tau})o_{it-\tau}$, where $t - \tau$ denotes the period in which the order was made that

[17] Appendix B.3 proposes an alternative Step 3 that ensures $b_{it} \geq 0$, but without imposing that the real stock-flow equation (3.1) holds every period. The alternative is equivalent under our assumptions but may differ when the assumptions are violated in the data. The resulting quantity and price series for backlog are highly similar. We prefer Step 3 as described in Section 3.2 because it requires fewer and smaller adjustments.

[18] The aggregate price of backlog typically falls between the aggregate price of new orders and sales.

Figure 1: Sales prices and new order prices

(a) Micro differences

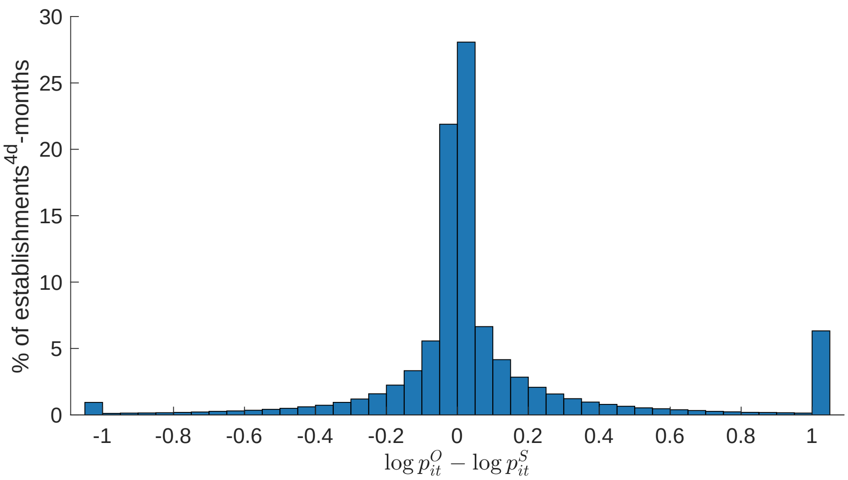

(b) Aggregate dynamics

Note: Panel (a) shows the distribution of the log difference between the new order price p_{it}^{O} and the sales price p_{it}^{S} across all observations in the final price sample. The distribution is winsorized at ± 1. Panel (b) shows aggregate year-on-year price growth for sales and new orders, computed by sales-weighted Törnqvist aggregation. The sample spans 2015–2022.

is sold in t. Substituting new orders by sales and similarly for prices, yields equivalently

$$\frac{R_{it}}{S_{it}} = -\frac{p_{it+\tau}^{S} - p_{it}^{S}}{p_{it}^{S}}. \tag{3.2}$$

If prices are commonly adjusted in the delivery period, we should expect sales price growth to correlate strongly with the residual-sales ratio. Importantly, the equation provides a test

about the importance of ex-post price changes for the residual that only requires sales prices, but not prices that were constructed under assumptions on the residual.
In the data, we compute τ-period sales price growth, where τ is the establishment^{4d}-specific rounded median ex-ante time to fill. we find that sales price growth correlates with the residual-sales ratio extremely weakly. The unconditional correlation is -0.00001. Adding establishment^{4d} and sector-time fixed effects still yields a tiny correlation of 0.0004, that additionally has the wrong sign. Thus, our test suggests ex-post price changes are of limited importance.

4 Time to fill

The previous section showed that prices of new orders and sales differ substantially. The difference arises because backlog creates a time lag between new orders and sales, corresponding to the time it takes for a new order to be filled. We refer to this lag as *time to fill* and document its empirical properties in this section.

4.1 Measuring time to fill

We use two complementary measures of time to fill. The first is an ex-ante (EA) measure of time to fill, defined as

$$\tau_{it}^{EA} = \frac{b_{it}}{s_{it}^{12}}, \qquad s_{it}^{12} = \frac{1}{12}\sum_{j=0}^{11} s_{it-j}. \tag{4.1}$$

It measures the coverage of real order backlog expressed in months of average recent real sales. In addition, τ_{it}^{EA} can be interpreted as the (ex ante) expected time until period-t new orders are filled. That interpretation holds if expected future sales (net of the residual) equal s_{it}^{12}, and orders are filled according to FIFO.[19] Its construction only requires real sales and backlog through month t.
The second measure is ex-post (EP) time to fill, defined as

$$\tau_{it}^{EP} = \min\left\{\tau \in \mathbb{N}^0 \,\middle|\, \sum_{k=0}^{\tau}(S_{it+k} + R_{it+k}) \geq B_{it-1} + O_{it}\right\}. \tag{4.2}$$

This measure coincides with $\overline{\tau}_{it}$ from Step 1 of the price construction scheme in Section 3.2.

[19]If an establishment^{4d} sells multiple products that differ in their time to fill, τ_{it}^{EA} measures the sales-weighted average time to fill across products. Formally, denote a product by $j \in \{1, \ldots, J\}$ and the associated ex-ante time to fill by $\tau_{it}^{j,EA} = b_{it}^{j}/s_{it}^{j,12}$. Then we obtain $\tau_{it}^{EA} = \sum_{j=1}^{J} w_{it}^{j}\tau_{it}^{j,EA}$ with $w_{it}^{j} = s_{it}^{j,12}/s_{it}^{12}$.

Under Assumptions 2–3, it is the (ex post) realized number of months until period-t new orders are completely filled.

4.2 Distribution of time to fill

We present robust evidence that time to fill averages half a year across establishments[4d] and months in our sample and features a pronounced right tail.

Table 3 presents descriptive statistics for τ_{it}^{EA} and τ_{it}^{EP}. The arithmetic average ranges from 5.4 to 5.7 months across measures, so roughly half a year. The finding that time to fill averages about half a year is a robust feature of the data as the sensitivity analysis in Table C.1 reveals. In particular, weighting either time to fill measure by nominal sales (deflated by the aggregate PPI) lowers the average τ_{it}^{EA} from 5.7 to 5.4 months and raises the average τ_{it}^{EP} from 5.4 to 6.8. Table 3 presents descriptive statistics for the largest possible samples. The sample size for the ex-post measure is larger in Table 3, because it does not require prices. If we compute the mean of τ_{it}^{EP} in the smaller sample for which the ex-ante measure is available, it slightly decreases by 0.1 months.

The ex-ante measure depends on the imputation of backlog prices. An alternative ex-ante measure is the nominal backlog-sales ratio B_{it}/S_{it}^{12}, which is common in previous work on order backlogs (see, e.g., Zarnowitz, 1973) and in official statistics. However, this measure can be distorted in the presence of price movements. E.g., increasing new order prices distort upwards the nominal backlog-sales ratio relative to its real counterpart. In our sample, we find that the nominal backlog-sales ratio is 5.9, similar to its real counterpart, suggesting a limited impact of price distortions on the average nominal backlog-sales ratio.

Table C.1 also shows that average time to fill is robust to accounting for the residual in the construction of τ_{it}^{EA}. As we discussed above, τ_{it}^{EA} measures expected time to fill if the firm expects future sales (net of the residual) to equal s_{it}^{12}. When the residual is typically non-zero, the net sales expectation may be better represented by $s_{it}^{12} + r_{it}^{12}$. Computing ex-

Table 3: Descriptive statistics on time to fill

		Cumulative distribution function				
	Mean	3 months	6 months	12 months	24 months	Obs
Time to fill (ex ante)	5.73	50.39%	71.99%	89.03%	96.37%	949,734
Time to fill (ex post)	5.41	48.94%	72.31%	91.71%	98.46%	1,311,391

Note: The table reports descriptive statistics for τ_{it}^{EA} and τ_{it}^{EP}, defined in equation (4.1) and (4.2), across establishment[4d]-months for 2015M1–2022M12, respectively in the final price sample and the final backlog sample, and winsorized at 60 months.

ante time to fill as $b_{it}/(s_{it}^{12}+r_{it}^{12})$ changes the average time to fill by very little, from 5.73 to 5.68. Finally, average time to fill for both measures rises only modestly when lifting the winsorization threshold from 5 to 10 years. In fact, there are barely any observations with τ_{it}^{EP} above 5 years. Hence, the rise is not even visible up to the presented digits.
Table 3 further shows that both time to fill measures feature a pronounced right tail: 8-11% of observations have time to fill above one year, and 1-4% above two years. The right-skew is similarly salient in the histograms of the two time to fill measures in Figure 2. Many establishment^{4d}-month observations have time to fill of three months or less, but a non-trivial right tail extends beyond one year.

Figure 2: Histogram of time to fill

(a) Ex-ante time to fill (τ_{it}^{EA})

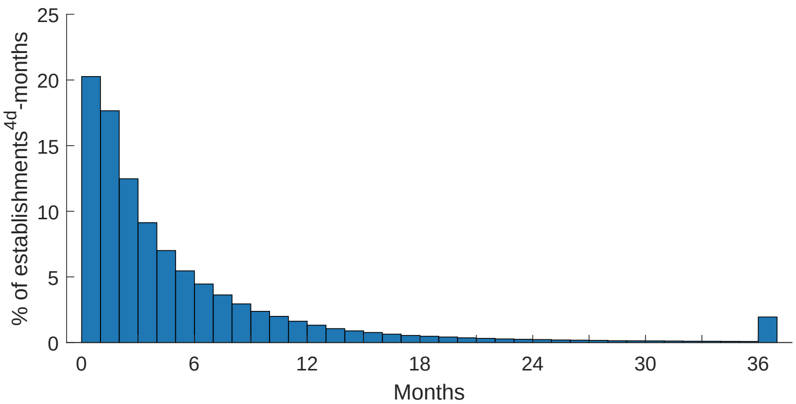

(b) Ex-post time to fill (τ_{it}^{EP})

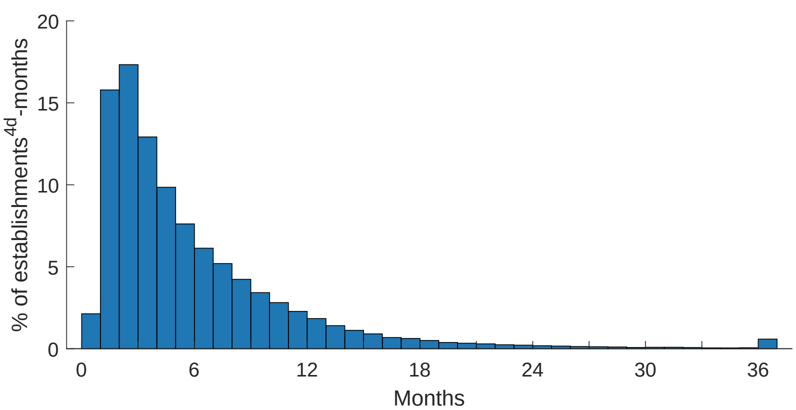

Note: The figure shows the histogram of τ_{it}^{EA}, defined in equation (4.1), and τ_{it}^{EP}, defined in equation (4.2), across establishment^{4d}-months for 2015M1–2022M12, respectively in the final price sample and the final backlog sample, and winsorized at 36 months.

4.3 Covariates of time to fill

We next show that most variation in time to fill is within sectors. In addition, time to fill is longer where demand appears more specific and uncertain, in line with Zarnowitz (1962), but is by and large unrelated to size, markups, or productivity.

Time to fill differs notably across backlog sectors. It is above the sample average in machinery and other vehicles, and below average in textiles, cloths, paper, chemicals, and pharmaceuticals, see Table C.2.[20] To systematically investigate the sources of variation, we consider the variance decomposition of time to fill τ_{it}^k for $k \in \{EA, EP\}$,

$$Var\left(\tau_{it}^k\right) = \underbrace{Var\left(\tau_{it}^k - \overline{\tau}_i^k\right)}_{\text{within establishment}^{4d}} + \underbrace{Var\left(\overline{\tau}_i^k - \overline{\tau}_s^k\right)}_{\substack{\text{between establishments}^{4d} \\ \text{(within sector)}}} + \underbrace{Var\left(\overline{\tau}_s^k - \overline{\tau}^k\right)}_{\text{between sectors}}, \tag{4.3}$$

where $\overline{\tau}_i^k$ denotes the establishment^{4d}-specific average τ_{it}^k across time, $\overline{\tau}_s^k$ the sector-specific average $\overline{\tau}_i^k$ across establishments^{4d} within sector s, $\overline{\tau}^k$ the average $\overline{\tau}_s^k$ across sectors s.

Table 4 provides the results when defining s as a four-digit sector. Even with such narrowly-defined sectors, differences across sectors explain only 7–12% of the variance. Most variation is either variation across establishments^{4d} within the same sector or variation within establishments^{4d} over time. For τ_{it}^{EA}, these two components are of broadly similar importance. For τ_{it}^{EP}, differences between establishments^{4d} dominate, but differences within establishments^{4d} are still more important than differences between sectors. In contrast, variation in time to fill across time is rather limited as Figure C.1 shows.

We next examine establishment-level covariates of time to fill in Table 5. Motivated by Zarnowitz (1962), we ask whether time to fill is longer where demand is more specific or uncertain. We find some evidence in support of these two hypotheses. To measure

Table 4: Variance decomposition of time to fill

	τ_{it}^{EA}	τ_{it}^{EP}
Within establishment^{4d}	36.66%	13.25%
Between establishments^{4d}	56.31%	74.76%
Between sectors	7.02%	11.99%

Note: The table reports the percentage share of the variance of time to fill accounted for by the components in equation (4.3). We use the final price sample for τ_{it}^{EA} and the final backlog sample for τ_{it}^{EP}, 2015M1–2022M12. Both τ_{it}^{EA} and τ_{it}^{EP} are winsorized at 60 months.

[20] For other vehicles, our statement relies on panel (b) of Table C.2. In this sector, the final price sample is highly selected because sales prices are missing for the vast majority of products.

Table 5: Covariates of establishment^{4d}-level time to fill

	$Corr(\cdot, \tau_i^{EA})$	$Corr(\cdot, \tau_i^{EP})$
Differentiated good	0.093	0.081
Avg. period of production	0.432	0.431
New order variance	0.345	0.340
Sales variance	0.336	0.264
Demand shock variance	0.230	0.250
Supply shock variance	0.455	0.472
Sales	0.015	0.097
Markup	0.002	0.032
TFP	-0.035	-0.053
Labor productivity	-0.059	-0.061

Note: The table reports correlations between establishment^{4d}-specific median time to fill and selected establishment^{4d} characteristics. The differentiated-good variable is based on a Rauch-type classification of nine-digit product codes. Avg. period of production is based on Antràs and Tubdenov (2025). Markup, TFP, and labor productivity are in logs.

demand specificity, we adapt the differentiated-goods classification of Rauch (1999) to nine-digit product categories in Germany. We construct an index that equals one if a good is classified as differentiated (i.e., more specific), and zero else. The variation is somewhat limited as more than 80% of products are classified as differentiated goods.[21] Nonetheless, we find a positive correlation with establishment^{4d}-level time to fill. To capture demand uncertainty, we consider the establishment^{4d}-level variance in log real new orders, log real sales, and of demand and supply shocks (the latter are constructed in the next section). Across all measures of establishment^{4d}-level uncertainty, we find a positive correlation with establishment^{4d}-level time to fill.

In contrast, time to fill is only weakly related to sales, markups, TFP, and labor productivity. Long time to fill therefore appears more closely associated with demand specificity and uncertainty than with establishment size or productivity.

5 Shock identification and propagation

This section shows how backlog affects the identification and propagation of micro-level demand and supply shocks. We formally highlight an identification challenge for a conven-

[21] Appendix C.4 provides more details on the index construction. Our Rauch classification of product categories can be retrieved from: `https://drive.google.com/file/d/1HuCr1RJLIhACEkoUnEjmN9YsbK0KunJm/view?usp=drive_link`

tional approach arising from the presence of backlog. We then propose an identification strategy that is valid in the presence of backlog and apply it to our sample. Finally, we estimate the dynamic responses of backlog, sales, new orders, and order prices to the identified shocks. The results show that backlog is an important adjustment margin: it absorbs expansionary demand shocks and is drawn down after expansionary supply shocks.

5.1 Identification challenge

We next analyze demand identification in the presence of backlog. A conventional class of identification strategies infers the demand curve from the co-movement of current sales quantities and sales prices, using supply-side variation to address price endogeneity. This logic is valid when sales reflect contemporaneous demand. With backlog, however, current sales are the fulfillment of backlogged orders, and current sales prices inherit past new order prices. Current demand instead depends on contemporaneous new order prices and quantities.

We formally analyze the implications of backlog for demand identification through a stylized example. Suppose an establishment[4d] faces the iso-elastic demand curve for new orders

$$\log o_{it} = \beta \log p_{it}^{O} + \log d_{it}, \qquad \beta < 0, \tag{5.1}$$

where $\log d_{it}$ denotes an exogenous demand shifter and we assume all variables are mean zero. In general, $\log d_{it}$ affects the order price p_{it}^{O}.[22] The econometrician observes (o_{it}, p_{it}^{O}) and a mean-zero instrument v_{it} that satisfies $\mathbb{E}[v_{it} \log p_{it}^{O}] \neq 0$ and $\mathbb{E}[v_{it} \log d_{it}] = 0$. The IV estimate $\widehat{\beta}^{IV}$ of (5.1) consistently identifies β. Instead consider the conventional sales-based identification strategy, where the econometrician observes $(s_{it}, p_{it}^{S}, v_{it})$ and estimates

$$\log s_{it} = \gamma \log p_{it}^{S} + u_{it}, \tag{5.2}$$

using instrument v_{it}. The IV estimate of γ generally differs from β:

$$\text{plim } \widehat{\gamma}^{IV} - \beta = \frac{\text{Cov}\left(v_{it}, \log s_{it} - \log o_{it}\right) + \beta \text{Cov}\left(v_{it}, \log p_{it}^{O} - \log p_{it}^{S}\right)}{\text{Cov}\left(v_{it}, \log p_{it}^{S}\right)}. \tag{5.3}$$

Absent time to fill ($\tau_{it} = 0$), it holds trivially that plim $\widehat{\gamma}^{IV} = \beta$. That is, the sales-based regression identifies the demand elasticity β. Any discrepancy between $\widehat{\beta}^{IV}$ and $\widehat{\gamma}^{IV}$ may reflect the endogeneity bias arising from correlation between the instrument and the sales-order wedge, in quantity and prices. If time to fill was constant, such that $s_{it} = o_{it-\tau}$ and $p_{it}^{S} = p_{it-\tau}^{O}$, the numerator of the right-hand side of (5.3) simplifies to Cov $(v_{it}, \log d_{it-\tau})$. The

[22]The model developed in Section 6 provides a theory of the endogeneity of p_{it}^{O} with respect to d_{it}.

covariance is zero if v_{it} is also orthogonal to the lagged demand shifter $\log d_{it-\tau}$.

In general, however, time to fill is endogenous to shocks, making this a harder identification problem. To fix ideas, suppose v_{it} is a cost-increasing supply shifter, implying that $\mathrm{Cov}\left(v_{it}, \log p^O_{it}\right) > 0$. The signs of both covariances in the numerator of (5.3) are ambiguous. As the empirical responses in Section 5.4 show, sales react more strongly to a contemporaneous supply shock than new orders, but new orders react more strongly at later horizons. Similarly, the response of order prices to a contemporaneous supply shock exceeds that of sales prices, since the latter are predetermined. Mechanically, that ordering flips at later horizons. The signs of both covariances therefore depend on whether the short-run or the long-run effects of supply side disturbances dominate in the data, which hinges on the autocorrelation structure of the supply shifter v_{it}. If the supply shifter is only weakly correlated over time, the bias is additionally inflated by the fact that the instrument induces little variation in current sales prices, even if it is a strong instrument for the current order price. In the limit as the supply shifter becomes uncorrelated over time, v_{it} is not a relevant instrument for the sales price at all.

5.2 Demand curve identification in practice

As we argued above, the demand curve is defined by new orders and associated prices, rather than sales and sales prices. Based on this understanding, we next implement a valid empirical identification strategy for the demand curve.

We consider the following regression

$$\log o_{it} = \alpha_i + \alpha_t + \beta \log p^O_{it} + \Gamma C_{it} + \varepsilon^d_{it}, \tag{5.4}$$

which generalizes (5.1) by controlling for establishment[4d] and month fixed effects. The vector C_{it} includes $\log o_{it-1}$ and $\log o_{it-12}$ which accounts for persistent order fluctuations and means the residual ε^d_{it} is rather a shock to the demand shifter $\log d_{it}$.

To address the endogeneity of p^O_{it} with respect to the demand shock ε^d_{it}, we instrument p^O_{it} with a Bartik shift-share measure of foreign input prices

$$v_{it} = \sum_{s \in S \setminus k(i)} \phi_{is} p^{imp}_{st}, \tag{5.5}$$

where s denotes a sector, ϕ_{is} is an exposure measure of establishment[4d] i to input costs from sector s, and p^{imp}_{st} is an import price index of sector s in month t. We construct the instrument v_{it} excluding the exposure of establishment[4d] i to the sector $k(i)$ to which it belongs. That avoids the potential endogeneity of v_{it} with respect to ε^d_{it} that arises if the latter directly affected prices in sector $k(i)$. Similarly, it should help that we use import prices, rather than domestic prices. The instrument is hence the average import price across other sectors weighted by the input cost exposure of establishment[4d] i across sectors.

We next provide more details on the implementation. S denotes the set of two-digit material- and energy-supplying sectors, of which there are 22 in the data. The price p_{st}^{imp} is the Destatis import price index of a two-digit sector s.[23] The exposure measure is constructed as

$$\phi_{is} = \begin{cases} \alpha_i^M \, \omega_{k(i),s}, & \text{if } s \in S_M, \\ \alpha_i^E \, \omega_{k(i),s}, & \text{if } s \in S_E, \end{cases} \tag{5.6}$$

where α_i^M is for establishment[4d] i the share of expenses for materials divided by the sum of expenses for materials, energy, and labor, and α_i^E is the corresponding share of energy expenses. $S_M \subset S$ is the subset of two-digit sectors supplying materials, $S_E \subset S$ the subset of two-digit sectors supplying energy. Finally, $\omega_{k(i),s}$ is the share of industry $k(i)$'s expenditure on imported inputs from sector s divided by all expenditures on imported inputs across sectors $s \in S$.[24]

We construct the instrument to limit its endogeneity to the demand shock ε_{it}^d by using only foreign input-price movements, excluding the establishment[4d] 's own sector, and holding exposure shares fixed over time. Additionally, establishment[4d] and month fixed effects absorb permanent differences across establishments[4d] and aggregate demand shocks. To address remaining concerns about endogeneity of ϕ_{is}, conditional on controls, we follow Borusyak et al. (2022) and include in the control vector C_{it} interactions between the time fixed effect and average exposure defined as $\overline{\phi}_i = \sum_{s \in S \setminus k(i)} \phi_{is}$. Then identification does not require the exposure shares ϕ_{is} to be randomly assigned across establishments[4d]. Finally, even if some establishments[4d] are granular, so their idiosyncratic demand shocks affect import prices, that does not necessarily invalidate the instrument. For example, a demand shock to an establishment[4d] i may raise prices in sectors from which i buys inputs, but lower prices in sectors that compete with i on output markets. That leaves endogeneity of v_{it} with respect to ε_{it}^d undetermined in sign and possibly negligible in magnitude.

Table 6 presents estimates of equation (5.4). Column (1) presents the baseline results. The estimated price elasticity of new orders is -1.11 and significantly differs from zero at the 1% level. The corresponding first-stage coefficient of the instrument is 0.13, estimated with

[23] To be precise, we compute the year-over-year growth rate of the Destatis price index, and then cumulate the growth rates to construct p_{st}^{imp}, normalizing the price for all sectors to 100 in 2015M1. This procedure mitigates seasonal and other short-lived fluctuations. Instead using the month-over-month growth rate of prices has little impact: e.g., the estimated $\widehat{\beta}^{IV}$ in column (1) of Table 6 changes from -1.1120 to -1.077.

[24] Appendix D provides further details on the construction of the instrument.

high precision, and the Kleibergen–Paap F-statistic is 196. Column (2) presents the results when not controlling for lagged new orders. The estimated elasticity is somewhat larger but broadly similar and remains significant. In column (3), we exclude the top one percent establishments^{4d} with the highest average real sales.[25] That addresses the concern that some establishments^{4d} are granular and generate endogeneity between the instrument and the demand shock which biases our estimates. We find that the estimated elasticity barely changes (-1.15) and is still significant at the 1% level. The finding suggests a limited role for potential endogeneity bias from granular establishments^{4d}.[26] Column (4) presents the OLS estimate of (5.4). The OLS estimate is substantially flatter (-0.69). That is consistent with the upward endogeneity bias expected in demand estimation arising from positive correlation between demand shocks and new order prices.

Finally, we estimate the demand elasticity using sales and sales prices instead of orders and order prices in (5.4) with the same instrument v_{it}. Table D.1 shows the estimates. Across

Table 6: Estimates of the demand curve

	TSLS			OLS
	(1)	(2)	(3)	(4)
	Baseline	*No controls*	*Trimmed*	—
$\log(p^O_{it})$	-1.1120	-1.2716	-1.1483	-0.6910
	(0.1316)	(0.1100)	(0.1351)	(0.0034)
Order controls	Yes	No	Yes	Yes
Establishment^{4d} FE	Yes	Yes	Yes	Yes
Month $\times$ $\overline{\phi}_i$ FE	Yes	Yes	Yes	Yes
R^2	0.1634	0.1177	0.1565	0.9053
Kleibergen–Paap F	196.13	319.26	198.00	—
First stage: v_{it}	0.1292	0.1687	0.1273	—
	(0.0092)	(0.0094)	(0.0090)	
N	842,803	899,454	833,529	842,803

Note: This table shows the estimated slopes of the demand curves for new orders. Columns (1)–(3) use TSLS with instrument v_{it}; column (4) reports OLS on the sample used in column (1). All regressions contain establishment^{4d} fixed effects and interactions of monthly time fixed effects with the sum of instrument exposure shares $\overline{\phi}_i \equiv \sum_{s \in S \setminus k(i)} \phi_{is}$. Columns (1), (3) and (4) control for $\log(o_{it-1})$ and $\log(o_{it-12})$. Column (3) excludes establishments above the 99^{th} percentile of average real sales in the estimation sample. The table reports the Kleibergen–Paap F-statistic and the first-stage coefficient of v_{it}. Standard errors in parentheses are clustered at the four-digit industry-month level.

[25] Real sales s_{it} are based on sales price indexes and thus not comparable across firms, only across time. Instead, we identify the largest firms using nominal sales S_{it} deflated by the PPI.

[26] Also excluding the top ten percent of establishments^{4d} from the estimation sample changes the demand slope only mildly to -1.31 (significant at the 1% level).

all specifications, the estimated demand slope is about half as steep as in Table 6. We decompose the sources of this bias into the residualized covariance terms according to (5.3) in Appendix Table D.2. This decomposition suggests that the positive covariance between v_{it} and the sales-order quantity wedge is a key factor for the misspecification bias.[27] Through the lens of our discussion in Section 5.1, this finding points toward a non-trivial bias in the sales-based demand regression.

5.3 Shock identification in practice

We identify establishment^{4d}-level demand shocks as the estimated residual ε_{it}^{d} in (5.4) based on the baseline TSLS estimate and specification in column (1) of Table 6. We denote the estimated shocks by $\widehat{\varepsilon}_{it}^{d}$.

We identify establishment^{4d}-level supply shocks based on the following regression

$$\log s_{it} = \alpha_i + \alpha_t + \eta \widehat{\varepsilon}_{it}^{d} + \boldsymbol{\gamma}' \mathbf{X}_{it} + \varepsilon_{it}^{s}, \tag{5.7}$$

where we control for establishment^{4d} and month fixed effects, the estimated demand shock, and a vector of lagged order-book variables $\mathbf{X}_{it} = (\log o_{it-1},\ \log o_{it-12},\ \log p_{it-1}^{O},\ \log s_{it-1},\ \log s_{it-12},\ \log b_{it-1})'$. We estimate (5.7) by OLS and denote the estimated supply shock by $\widehat{\varepsilon}_{it}^{s}$. The estimated η coefficient is positive (0.144) and highly significant, well below the 1% level. That is reassuring as we should expect positive demand shocks to raise sales.

The supply shocks are surprises in sales that are orthogonal to demand shocks. Put differently, they capture all movements in sales surprises that cannot be explained by demand shocks. The identification strategy is consistent with the firm's optimal sales policy function derived in the model in the next Section.

The estimated demand innovations $\widehat{\varepsilon}_{it}^{d}$ display modest serial correlation: their first-order autocorrelation is 0.117 and their twelve-month autocorrelation is 0.063. The supply innovations $\widehat{\varepsilon}_{it}^{s}$ are essentially serially uncorrelated, with the corresponding correlations -0.005 and 0.006. Since our estimation approach absorbs time and establishment^{4d} fixed effects, these identified shock series reflect idiosyncratic demand and supply shocks, rather than aggregate shocks.

[27] This positive covariance term is also consistent with the local projections of an expansionary supply shock on log sales and on log orders that we will estimate in the remainder of this Section. The local projections suggest a positive co-movement between a positive supply shock and $\log s_{it} - \log o_{it}$ on impact, which turns negative for longer horizons.

5.4 Evidence on shock propagation

Equipped with identified shocks, we estimate dynamic responses of order book variables using panel local projections. Formally, we estimate the following panel local projections for each outcome y_{it} and shock type $k \in \{d, s\}$:

$$\log y_{it+h} - \log y_{it-1} = \alpha_i^h + \alpha_{st}^h + \beta_{y,k}^h \hat{\varepsilon}_{it}^k + \boldsymbol{\gamma}_{y,k}^{h\prime} \mathbf{Z}_{i,t-1}^{y,k} + v_{it,y,k}^h, \tag{5.8}$$

where $y_{it} \in \{b_{it}, s_{it}, o_{it}, p_{it}^O\}$ denotes the order-book outcome of interest. The terms α_i^h and α_{st}^h denote establishment[4d] and two-digit-sector-time fixed effects, vector $\mathbf{Z}_{i,t-1}^{y,k}$ includes a set of controls $(\log y_{it-1}, \log b_{it-1}, \hat{\varepsilon}_{it-1}^k)'$.

We estimate (5.8) for horizons h from zero to twelve months,[28] separately for each outcome and shock type. We standardize the estimated shocks to have a unit standard deviation. The coefficients $\{\beta_{y,k}^h\}_{h=0}^{12}$ therefore trace the cumulative response of outcome y to a one-standard-deviation shock of type k, expressed in percentage log points relative to the month before the shock.

Figure 3 shows the estimated cumulative responses to one-standard-deviation expansionary demand and supply shocks. The responses are expressed in percent deviations relative to the month before the shock occurs.

An expansionary demand shock raises all four order-book variables on impact. New orders increase by 52%, substantially more than sales, which increase by 4.7%. Backlog rises by 6.8%, while order prices increase only slightly by 0.5%. The new-order response falls to 20% after one month and declines further over the subsequent months. Sales adjust more gradually and remain above their pre-shock level throughout the horizon. The backlog response also declines gradually but remains economically and statistically significant, reaching 5.4% after twelve months.

An expansionary supply shock raises sales by 37% on impact, while new orders increase by only 1.3%. At the same time, backlog falls by 6.0% and order prices decline by 0.03%. The sales response falls to 7.0% after one month and becomes considerably smaller at subsequent horizons. By contrast, the new-order response rises to 3.3% after one month and remains positive thereafter. The backlog response gradually moves toward zero but remains negative and statistically significant throughout the horizon, reaching 3.0% after twelve months. The decline in order prices reaches 0.24% after six months and subsequently moves toward zero.[29]

Overall, the estimates show that backlog is an economically important adjustment margin.

[28]The reason we do not extend the horizon beyond twelve months is the relatively short time series dimension of our sample.

[29]Figure D.2 in Appendix D provides the estimated response of the real residual.

Figure 3: Responses to demand and supply shocks

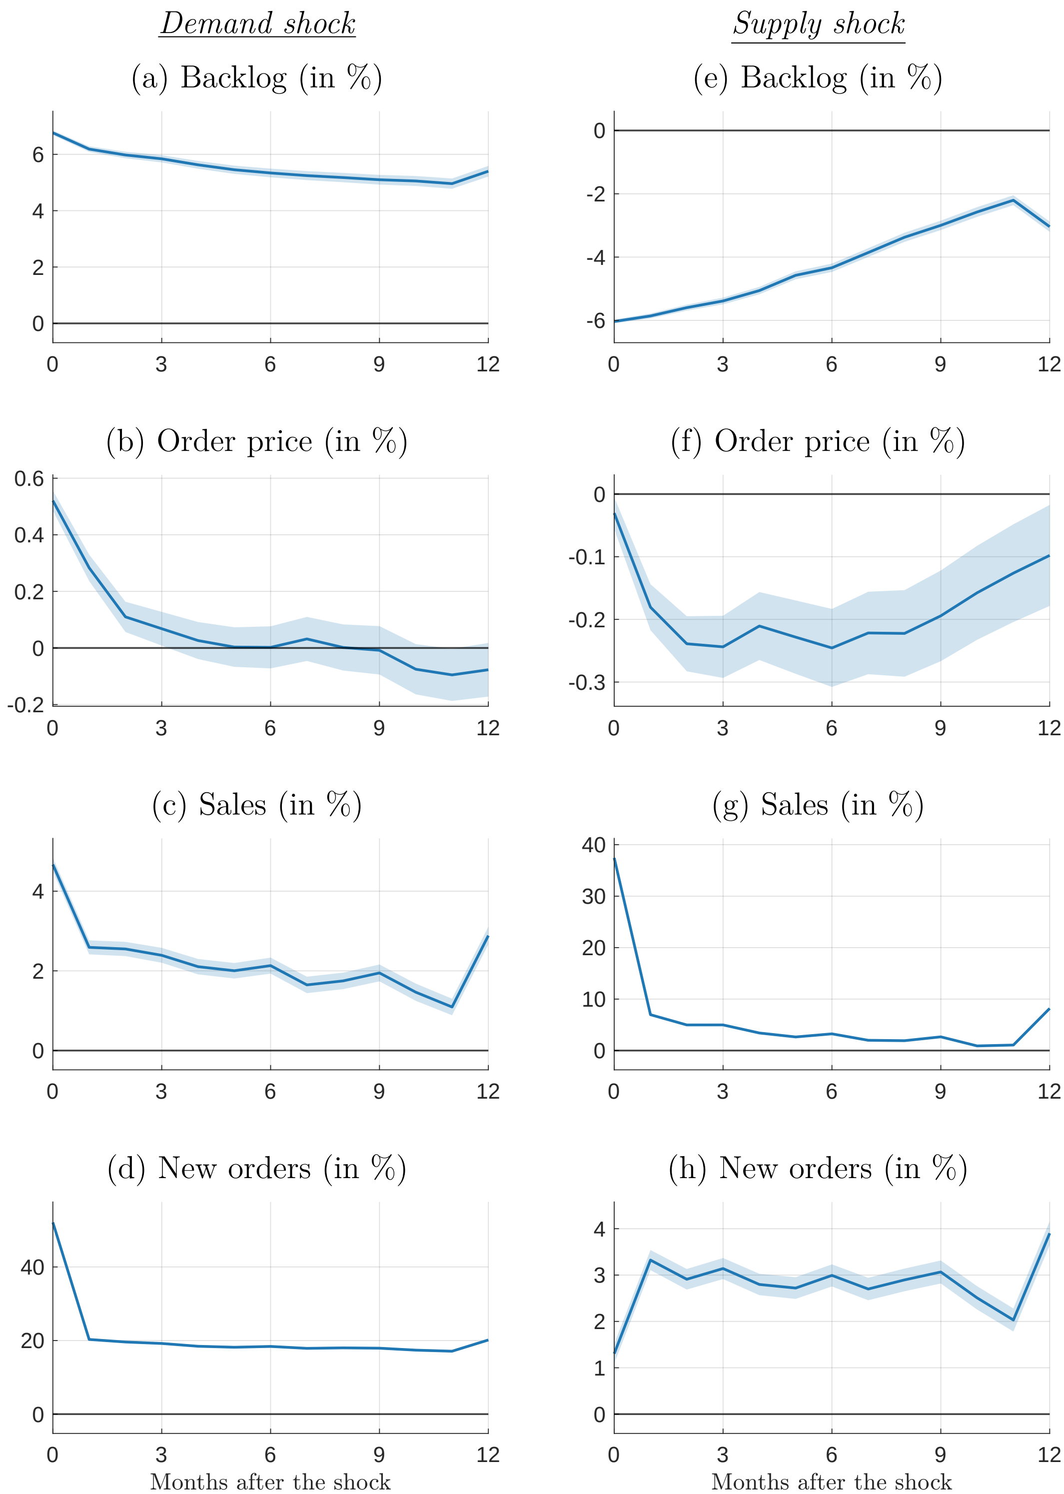

Note: This figure shows the IRFs to expansionary demand and supply shocks of one standard deviation expressed in percent. The 95% confidence bands are computed using robust standard errors. The observations were trimmed at 1% and 99% for each horizon.

Firms absorb parts of a demand expansion into the order book rather than translating it immediately into sales. When production conditions improve, firms draw down outstanding backlog. These responses match the central mechanism emphasized in the introduction: backlog dampens the sales response to demand shocks, amplifies the sales response to supply shocks, and dampens overall price responses.

Finally, we compare these results with shocks identified from the conventional sales-price and sales-quantity regression. Figure D.1 in the Appendix shows that the estimated responses differ substantially. Two results are most striking. First, order backlog falls in response to an expansionary sales-based demand shock, whereas backlog increases in response to the baseline, order-based, demand shock in Figure 3. Second, new orders fall in response to an expansionary sales-based supply shock, again in contrast to our finding for the baseline shock. In fact, the next section shows that a general model of backlog adjustment rationalizes the responses in Figure 3. The model hence will contradict the findings in Figure D.1. The comparison between order-based and sales-based identification suggests that the identification challenge posed by order backlog is not merely a theoretical curiosity, but has empirical bite for the estimated propagation of supply and demand shocks.

6 A theory of backlog adjustment

In this section, we introduce a theory of backlog adjustment. Backlog is a dynamic control variable that constitutes a third margin of firm adjustment to shocks, next to the price and production. We characterize the responses to demand and supply shocks analytically and show that they are qualitatively consistent with the estimated responses in Section 5. Next, we show that the model also matches the estimated responses quantitatively and use the calibrated model to quantify the role of backlog adjustment in shaping shock propagation.

6.1 Problem of the firm

We consider the problem of a firm that seeks to maximize its expected discounted profits. Time is discrete, indexed by t, and runs forever. The firm's choice variables are real new orders o_t, the price of new orders p_t, real sales s_t, and the real backlog b_t.[30] The firm's period profit is revenue from new order contracts less production costs and a backlog penalty function. Production costs depend on sales and an exogenous supply shifter z_t. The backlog penalty depends on the real backlog and steady state sales s. The firm's choices are subject to two constraints: backlog follows a real stock-flow equation and new orders reflect an

[30] To lighten the notation, we denote the price of new orders as p_t instead of p_t^O as in previous sections.

isoelastic demand function, subject to the exogenous demand shifter d_t.

$$\max_{\{o_t,p_t,s_t,b_t\}_{t=0}^{\infty}} \quad E_0 \sum_{t=0}^{\infty} \beta^t \left[p_t o_t - \frac{C_s(s_t)}{z_t} - C_\tau \left(\frac{b_t}{s} \right) \right] \tag{6.1}$$

$$s.t. \quad o_t = d_t \left(p_t\right)^{-\sigma} \tag{6.2}$$

$$b_t = b_{t-1} + o_t - s_t \tag{6.3}$$

$$\log(d_t) = \rho_d \log(d_{t-1}) + \varepsilon_t^d, \qquad \varepsilon_t^d \overset{iid}{\sim} (0, \sigma_d^2) \tag{6.4}$$

$$\log(z_t) = \rho_z \log(z_{t-1}) + \varepsilon_t^z, \qquad \varepsilon_t^z \overset{iid}{\sim} (0, \sigma_z^2) \tag{6.5}$$

$$\sigma > 1, \quad \beta,\ \rho_d,\ \rho_z \in (0,1)$$

The demand and supply shocks, ε_t^d and ε_t^z, are mutually uncorrelated. Variables without time subscripts, e.g., s, denote their deterministic steady-state values. We next state two assumptions and discuss them subsequently.

Assumption 1 (Steady state)**.** The model has an interior deterministic steady state with $z = d = 1$, $s = o > 0$, $p = s^{-\frac{1}{\sigma}}$, $s^{-\frac{1}{\sigma}} = \frac{\sigma}{\sigma-1} C_s'(s)$, $C_\tau'(\tau) = (1-\beta) s C_s'(s)$, and $\tau = \frac{b}{s} > 0$.

Assumption 2 (Local curvature)**.** The functions C_s and C_τ are twice continuously differentiable in neighborhoods of s and τ, respectively, with $C_s''(s) > 0$ and $C_\tau''(\tau) > 0$.

Assumption 1 is a high-level assumption on the existence of an interior steady state, defining local restrictions on the first derivatives of C_s and C_τ. In Assumption 2, twice continuous differentiability permits sharp analytical characterization up to first order. Strictly positive second derivatives, given all other assumptions, are sufficient for the firm problem to admit a stable optimal policy. Neither assumption restricts the cost functions globally.[31]

Assumptions 1 and 2 are not only mathematically convenient but also economically meaningful. Positive curvature of C_s generates a production smoothing motive which may reflect decreasing returns to scale of the underlying production function. In the short run, expanding production requires increasingly expensive overtime or more intense capacity use, while reducing production generates diminishing cost savings as the most expensive production margins are eliminated first.

We view C_τ as representing, in reduced-form, the various technological, organizational, financial, strategic, and customer-side considerations associated with backlog. Modeling these mechanisms separately would require a substantially richer model, whereas C_τ keeps the

[31] A set of global assumptions sufficient for a unique positive steady state and for the local curvature conditions is: (i) C_s is twice continuously differentiable on $(0,\infty)$, with $C_s'(x) > 0$ and $C_s''(x) > 0$ for every $x > 0$; and (ii) C_τ is twice continuously differentiable on $(0,\infty)$, with $C_\tau''(x) > 0$ for every $x > 0$, $\lim_{x\downarrow 0} C_\tau'(x) \leq 0$, and $\lim_{x\to\infty} C_\tau'(x) = \infty$.

model parsimonious and tractable. The argument of C_τ equals time to fill τ in steady state. Several non-exclusive mechanisms can explain positive (local) curvature of C_τ, which generates a backlog smoothing motive. When backlog is low, new orders quickly turn into sales, leaving little time for planning, procurement, and scheduling production, thus creating expediting costs.[32] Low backlog also means future capacity utilization depends more strongly on yet unknown future orders. That raises planning uncertainty, renders investment riskier, in turn making external finance costlier. Conversely, a large backlog increases expected time to fill, raising the producer's exposure to nominal input prices and increasing profitability risk given precommitted sales prices. Long time to fill can further reduce customer demand, shift orders to competitors, and may trigger contractual penalties.[33]

6.2 Optimal firm behavior

We next analyze the firm's optimal choices. We analytically characterize the propagation of supply and demand shocks, and the role of backlog adjustment.

The following proposition analytically characterizes optimal firm behavior in response to shocks up to first order, where we denote log deviations of backlog from its steady state by $\hat{b}_t = \log b_t - \log b$ and analogously for p_t, s_t, and o_t.

Proposition 1 (Shock Propagation). *Under Assumptions 1-2, the first-order conditions of* (6.1)*, log-linearized around the steady state, admit the unique bounded solution*

$$\widehat{b}_t = \theta_b^b\, \widehat{b}_{t-1} + \theta_b^d\, \widehat{d}_t + \theta_b^z\, \widehat{z}_t, \qquad \widehat{p}_t = \theta_p^b\, \widehat{b}_{t-1} + \theta_p^d\, \widehat{d}_t + \theta_p^z\, \widehat{z}_t,$$

$$\widehat{s}_t = \theta_s^b\, \widehat{b}_{t-1} + \theta_s^d\, \widehat{d}_t + \theta_s^z\, \widehat{z}_t, \qquad \widehat{o}_t = \theta_o^b\, \widehat{b}_{t-1} + \theta_o^d\, \widehat{d}_t + \theta_o^z\, \widehat{z}_t,$$

where $\theta_b^b \in (0,1)$*, hence* $\widehat{b}_t$*,* $\widehat{p}_t$*,* $\widehat{s}_t$*,* $\widehat{o}_t$ *follow stable processes. Moreover,* $\theta_b^d, \theta_p^d, \theta_s^d, \theta_o^d > 0$*, i.e.,*

- *expansionary demand shocks raise backlog, prices, sales, and new orders,*

and $\theta_b^z, \theta_p^z < 0$*,* $\theta_s^z, \theta_o^z > 0$*, i.e.,*

- *expansionary supply shocks lower backlog and prices, but raise sales and new orders.*

Proof: See Appendix E.1.

Importantly, for both the demand and supply shock, all predictions of the theory regarding the impact responses of the two shocks match the estimated impact responses in Figure 3.

[32]Note that these forces are distinct from the costs of operating under high utilization of labor and capital, which are captured by the cost function C_s.

[33]For related discussions of the mechanisms behind backlog, see Zarnowitz (1973) and Belsley (1969).

The model hence rationalizes the empirical evidence. In both data and theory, the demand and supply shock behave like textbook demand and supply shocks. The supply shock moves the economy along the demand curve (6.2), leading to negative co-movement of price and quantity of new orders, while the demand shock generates positive co-movement. Interestingly, the co-movement between backlog and sales is ambiguous and shock-dependent. An increase in backlog may reflect an expansionary demand shock or a contractionary supply shock.

We next explore the role of backlog for shock propagation. That will also clarify why backlog increases after demand shocks but falls after supply shocks. In particular, we contrast our baseline firm problem to a counterfactual in which the costs of backlog adjustment are excessively high such that backlog adjustment is not optimal. Formally, we consider the firm problem (6.1) when replacing the backlog penalty function C_τ by an alternative backlog penalty function $\widetilde{C}_\tau$ that satisfies

$$\widetilde{C}_\tau'(\tau) = C_\tau'(\tau), \qquad C_\tau''(\tau) < \widetilde{C}_\tau''(\tau) \to \infty. \tag{6.6}$$

A simple example is $\widetilde{C}_\tau(x) = C_\tau(x) + \eta(x-\tau)^2$ when $\eta \to \infty$. Under $\widetilde{C}_\tau$, it is optimal never to adjust backlog, hence $\widehat{b}_t = 0$. A direct implication is that sales equal new orders in every period, $\widehat{s}_t = \widehat{o}_t$. The following proposition characterizes how the policy coefficients change in this case.

Proposition 2 (Backlog counterfactual). *Consider the θ coefficients in Proposition 1 and compare them to the coefficients $\widetilde{\theta}$ obtained when replacing C_τ by $\widetilde{C}_\tau$ as defined in* (6.6)*. In the limit of $\widetilde{C}_\tau''(\tau) \to \infty$, backlog adjustment is not optimal, $\widetilde{\theta}_b^d = \widetilde{\theta}_b^z = 0$. It further holds that $\widetilde{\theta}_p^d > \theta_p^d > 0$, $\widetilde{\theta}_s^d > \theta_s^d > 0$, and $\theta_o^d > \widetilde{\theta}_o^d > 0$, i.e.,*

- *absent backlog adjustment: expansionary demand shocks raise prices and sales more strongly, while new orders increase less strongly,*

and $\widetilde{\theta}_p^z < \theta_p^z < 0$, $\theta_s^z > \widetilde{\theta}_s^z > 0$, and $\widetilde{\theta}_o^z > \theta_o^z > 0$, i.e.,

- *absent backlog adjustment: expansionary supply shocks lower prices more strongly, raise sales less strongly, and raise new orders more strongly.*

Proof: See Appendix E.2.

We first discuss the firm response to an (expansionary) demand shock. If the firm can adjust its backlog, it shifts some of the increase in new orders coming from the demand shock toward future sales, raising its backlog in the meantime. That is optimal because it allows the firm to dampen the increase in production costs (C_s). If the firm cannot adjust its backlog, any increase in new orders immediately raises sales and thus production costs. To dampen the increase in production costs, the firm raises its new order price more strongly, leading to a weaker expansion of new orders and thus sales. As a consequence, backlog adjustment dampens the price and sales response, but amplifies the new order response to an expansionary demand shock.

Next, consider an (expansionary) supply shock. If the firm cannot adjust its backlog, it can only exploit low production costs by lowering the new order price, which attracts more new orders, and thus sales. If the firm can adjust its backlog, it can also raise sales by lowering its order backlog. Given the backlog margin, it is optimal for the firm to lower its new order price by less. As a consequence, backlog adjustment dampens the price and order response, but amplifies the sales response to an expansionary supply shock.

So far we have discussed the impact response. But backlog is an endogenous state variable and thus also shapes the dynamic propagation. Whether or not backlog amplifies or dampens the responses to demand and supply shocks beyond the impact period depends on the interaction of the direct effects discussed in Proposition 2 and the indirect effects due to $\widehat{b}_t$. Which of the channels dominates is ultimately a quantitative question, which we address in the subsequent subsection.

6.3 Model calibration

We next show that despite its stylized nature, the model can quantitatively explain the dynamic response estimates shown in Figure 3.

In line with the evidence, we calibrate the model at monthly frequency. We calibrate the discount factor β to match a yearly interest rate of 3%. Since the empirical estimates in Figure 3 are average responses across establishments[4d] with vastly heterogeneous time to fill, we allow for ex-ante heterogeneity in steady state time to fill τ. We split the empirical time to fill distribution into ten groups, each containing an equal share of observations, and compute the group-specific median ex-ante time to fill. Those numbers are shown in Table 7 and define ten types of firms in the model. We assume that the firms are otherwise identical. We compute the average shock responses, corresponding to Figure 3, as the arithmetic average of the type-specific shock responses.

All remaining model parameters are calibrated targeting the estimates in Figure 3. In total, these estimates provide 104 calibration targets (across 8 impulse responses with each 13 horizons). We calibrate to minimize the squared distance between these empirical targets and their model counterparts. The goal is to see how closely the model can match the evidence. The proof of Proposition 1 provides closed-form solutions for all policy coefficients, which shows that those coefficients only depend on β and τ (calibrated above), the demand elasticity σ, the shock processes, as well as the following two parameters

$$\varphi_s = \frac{sC''_s(s)}{C'_s(s)}, \qquad \varphi_\tau = \frac{\tau C''_\tau(\tau)}{sC'_s(s)}, \tag{6.7}$$

which are proportional to the steady state curvature of C_s and C_τ, respectively. They fully summarize how C_s and C_τ shape optimal firm policy.

We slightly generalize the processes for d_t and z_t in (6.4) and (6.5). The motivation is that the estimated response of new orders to the demand shock and of sales to the supply shock in Figure 3 exhibits large impact responses which drop substantially at a one-month horizon but remains persistently different from zero for longer horizons. To better allow the model to match this pattern, we allow for both a transitory and a persistent shock, according to

$$\log d_t = u^d_t + \eta^d_t, \qquad u^d_t = \rho_d u^d_{t-1} + \varepsilon^d_t, \qquad \eta^d_t \overset{iid}{\sim} (0, \sigma^2_{\eta d}), \qquad \varepsilon^d_t \overset{iid}{\sim} (0, \sigma^2_{\varepsilon d}), \tag{6.8}$$

$$\log z_t = u^z_t + \eta^z_t, \qquad u^z_t = \rho_z u^z_{t-1} + \varepsilon^z_t, \qquad \eta^z_t \overset{iid}{\sim} (0, \sigma^2_{\eta z}), \qquad \varepsilon^z_t \overset{iid}{\sim} (0, \sigma^2_{\varepsilon z}), \tag{6.9}$$

where η^d_t, ε^d_t, η^z_t, ε^z_t are mutually uncorrelated. We define a one-standard deviation shock as the sum of a one-standard deviation transitory shock and a one-standard deviation persistent shock, respectively for demand and supply shocks.

Finally, we address an inconsistency between model and evidence. The stock-flow equation in (6.3) is an identity that must hold. By construction, the shock responses in the model satisfy this identity up to its log-linear approximation

$$\tau(\widehat{b}_t - \widehat{b}_{t-1}) = \widehat{o}_t - \widehat{s}_t. \tag{6.10}$$

In contrast, the estimates in Figure 3 are not restricted to satisfy (6.10). Deviations can arise from the log-linear approximation of (6.10), from implicit variance weighting in the OLS estimates, and because of non-zero responses of the order book residual. To address the inconsistency between model and data, we have re-estimated the empirical responses when imposing (6.10) as a cross-equation restriction that binds together the backlog, new orders, and sales responses.

Table 7 summarizes the calibrated parameters. The demand elasticity takes a value of 15.76, implying a moderate steady state markup of around 7%. The model requires relatively little curvature in the backlog penalty function relative to the production cost function. The persistent shocks dominate fluctuations in the total demand and supply shocks: 83% of the unconditional variance of the demand shock $\log d_t$ is due to the persistent shock, for the supply shock it is even 99%. The autocorrelation of both persistent shocks exceeds 0.95.

Table 7: Monthly calibration

Externally calibrated		
β	Discount factor	0.9975
τ	Time to fill	0.24, 0.73, 1.26, 1.84, 2.54, 3.46, 4.76, 6.65, 9.95, 24.56
Internally calibrated		
σ	Demand elasticity	15.762
φ_s	Proportional to curvature of $C_s(s)$	0.0213
φ_τ	Proportional to curvature of $C_\tau(\tau)$	0.0004
ρ_d	Autocorrelation of persistent demand shock	0.9614
ρ_z	Autocorrelation of persistent supply shock	0.9950
$\sigma_{\varepsilon d}$	Standard deviation of persistent demand shock	0.0554
$\sigma_{\eta d}$	Standard deviation of transitory demand shock	0.0913
$\sigma_{\varepsilon z}$	Standard deviation of persistent supply shock	0.0025
$\sigma_{\eta z}$	Standard deviation of transitory supply shock	0.0019

Note: This table reports the calibrated model parameters.

Figure 4 shows that the model matches the empirical estimates remarkably well. The figure compares the estimated responses when imposing the stock-flow constraint, the blue dotted line, with the corresponding model responses, the red solid line. For all four outcomes and both shocks, the model generates the correct sign and the right order of magnitude. Furthermore, for most responses, it closely traces the impact and dynamic response. The model does not fit well the impact response of sales to the demand shock, and the dynamic shape of the order price response to both shocks. The price responses are matched well in terms of their average magnitudes, but the model overestimates their persistence. Overall, our stylized model, with 20 parameters, matches the 104 moments describing the empirical responses quite well.

6.4 Backlog counterfactuals

We use the calibrated model to quantify the role of backlog adjustment, following the backlog counterfactual in Proposition 2. We increase the curvature of the backlog penalty function (φ_τ) to infinity, while keeping all other parameters unchanged. Figure 5 compares the counterfactual responses (blue, dashed lines) to the baseline model responses (solid, red lines). Consistently with Proposition 2, the price responses, the sales response to the demand shock,

Figure 4: Model fit

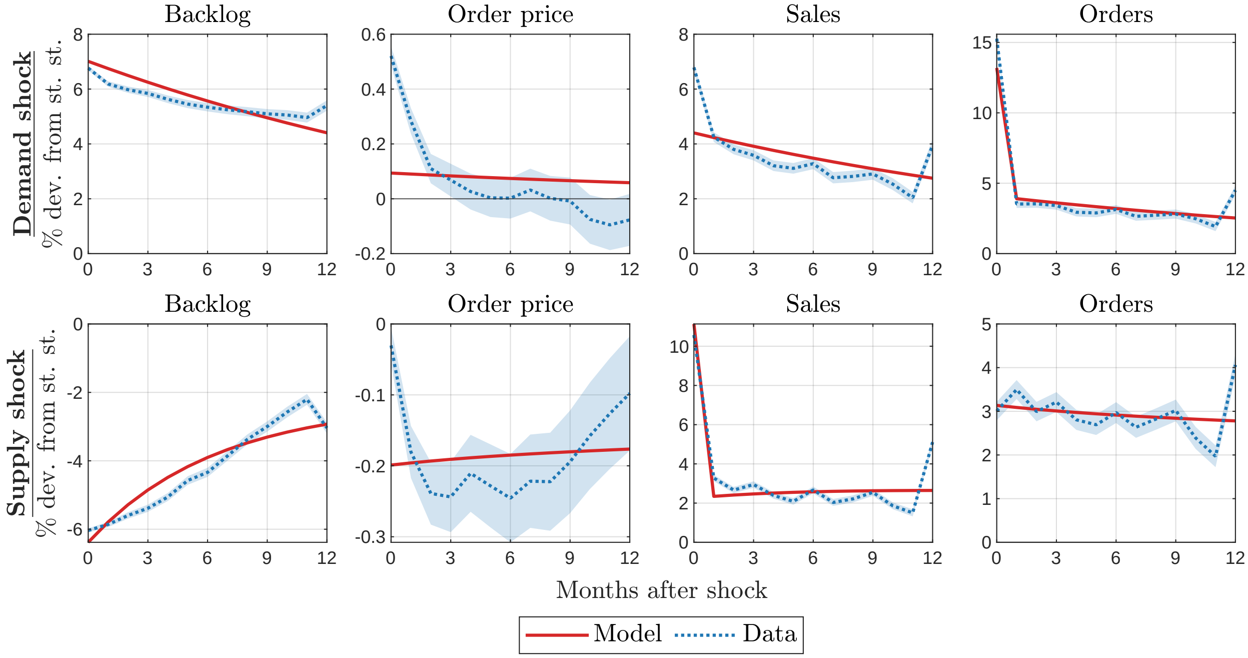

Note: This figure shows the estimated IRFs together with 95% confidence bands and the aggregate model IRFs for the parameters reported in Table 7.

and the order response to the supply shock are dampened by backlog adjustment, whereas the order response to the demand shock and the sales response to the supply shock are amplified. We find quantitatively large effects for the impact responses. The price responses are dampened by 40-60% and the sales response by 60-120%.

Figure 5 also shows that the quantitative effects of backlog adjustment beyond the impact period are rather small and typically reverse signs relative to the impact period. Making backlog adjustment excessively costly exerts two effects on the model dynamics. In all periods, the policy coefficients associated with demand and supply shocks change in the way discussed in Proposition 2. At the same time, because backlog remains at its steady state value, the path of the endogenous state variable changes to $\widehat{b}_{t-1} = 0$, inducing a second indirect effect as of horizon 1, which we investigate through a second counterfactual.

The blue dotted line in Figure 5 shows the responses in a naive counterfactual, where the policy coefficients are as in the baseline, but the backlog is set (forced) to its steady state value in all periods.[34] On impact, lagged backlog is at its steady state so the naive counterfactual

[34]Note that by construction of the naive counterfactual, the responses do not satisfy the stock-flow equation anymore. Rather, it serves as an accounting exercise, allowing us to separate the direct and indirect effects of backlog adjustments.

Figure 5: Backlog counterfactual

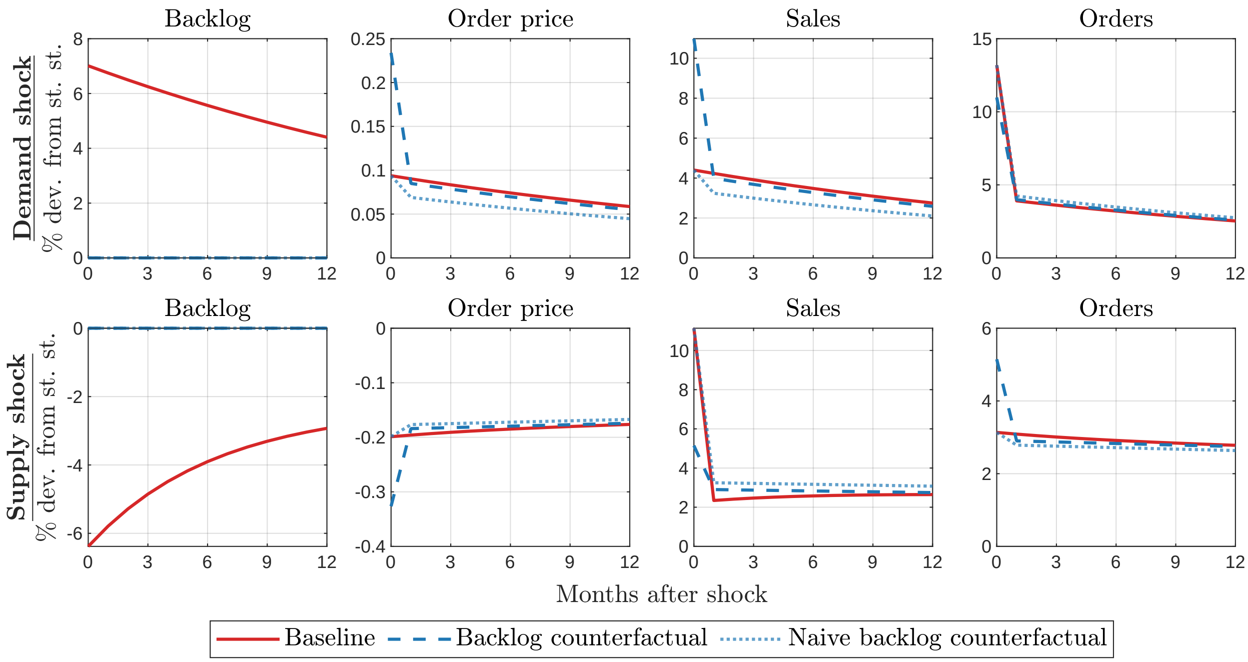

Note: This figure shows the baseline responses (red, solid lines) from Figure 4 as well as two counterfactuals. The backlog counterfactual (blue, dashed lines) represents the counterfactual responses when $\varphi_\tau \to \infty$. The naive backlog counterfactual (blue, dotted line) represents a counterfactual where all parameters are chosen as in the baseline, but $\widehat{b}_t$ is set to 0 in all periods.

mechanically coincides with the baseline. Beyond the impact period, the difference between baseline and naive counterfactual reflects the indirect effects working through the alternative backlog path dominate, generating a sign reversal. The naive counterfactual shows that backlog adjustment amplifies the price responses, the sales response to demand shocks, and the new order response to supply shocks whereas it dampens the order response to demand shocks and the sales response to supply shocks. Hence, the indirect effects are the exact opposite of the direct effects associated with the changing policy coefficients discussed in Proposition 2.

We next provide some intuition. Consider the demand shock. At horizon 1 in the baseline, the firm starts into the period with a backlog exceeding its steady state value. Due to the penalty cost, the firm wants to decrease the backlog, all else equal, so it needs to increase sales and decrease new orders, which it achieves by charging a higher price. Following a supply shock, backlog falls short of its steady state value, so the firms decreases its price and sales and increases its new orders, relative to a scenario where backlog is at the steady state in the beginning of the period. Put differently, because the model is stationary and backlog eventually returns to steady state, an amplified response at some horizon must be dampened at other horizons.

Summing up, both in response to demand and supply shocks, the firm finds it optimal to allow for a wedge between sales and new orders. This is because periods in which production is cheap do not necessarily coincide with periods of exceptionally high demand. In response to demand shocks, the firm collects new orders while demand is high and only gradually works them down to avoid increasing the marginal production cost too much. In response to supply shocks, the firm produces while it is cheap to produce, but only restocks it order backlog gradually to avoid large price cuts in any given period. As a consequence, the possibility to adjust backlog tends to be dampening the responses to shocks, except for the order response to demand shocks and the sales response to supply shocks. Quantitatively, those amplifying and dampening effects exist predominantly on impact and can even reverse signs at later horizons.

7 Conclusion

This paper highlights the importance of order books in the manufacturing sector. We introduce a novel administrative micro dataset on order books in Germany, which includes order book data for two thirds of the manufacturing sector. We observe the order backlog, the inflow of new orders, and the order book outflow, sales. The presence of order backlog induces a time wedge between new orders and sales—time to fill—which averages between 5 and 7 months across establishments in our sample, and which features a pronounced right tail. Our analysis requires prices for all order book variables, which is an informational requirement not met in the data. Instead, we propose and apply a scheme to construct theoretically coherent establishment-level prices for all order book variables for which it suffices to observe establishment-level sales prices. The scheme exploits the fact that all prices are dynamically linked through the order backlog. We show that sales prices are predetermined by past prices of new orders, creating a novel source of nominal rigidity. The rigidity of sales prices invalidates conventional strategies to identify demand and supply shocks. Instead, variation in quantity and price of new orders can identify such shocks and we propose and apply a valid identification strategy. Finally, we provide a theory of the firm to analyze the propagation of demand and supply shocks when backlog constitutes an adjustment margin. Using the identified shocks, we test the theory in the data. We empirically characterize shock propagation and find responses that are qualitatively in line with the model.

References

Adão, R., M. Kolesár, and E. Morales (2019): "Shift-Share Designs: Theory and Inference," *Quarterly Journal of Economics*, 134, 1949–2010.

Amiti, M., O. Itskhoki, and J. Konings (2019): "International Shocks, Variable Markups, and Domestic Prices," *The Review of Economic Studies*, 86, 2356–2402.

Antràs, P. and A. Kulesza (2026): "An 'Austrian' Model of International Specialization," Working Paper 34983, National Bureau of Economic Research.

Antràs, P. and V. Tubdenov (2025): "Measuring the Average Period of Production," mimeo.

Belsley, D. A. (1969): *Industry production behavior: The order-stock distinction*, vol. 62, North-Holland.

Benkard, C. L. (2004): "A Dynamic Analysis of the Market for Wide-Bodied Commercial Aircraft," *The Review of Economic Studies*, 71, 581–611.

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63, 841–890.

Berry, S. T. (1994): "Estimating Discrete-Choice Models of Product Differentiation," *The RAND Journal of Economics*, 25, 242–262.

Bighelli, T., F. di Mauro, M. J. Melitz, and M. Mertens (2022): "European Firm Concentration and Aggregate Productivity," *Journal of the European Economic Association*, 21, 455–483.

Borusyak, K., P. Hull, and X. Jaravel (2022): "Quasi-Experimental Shift-Share Research Designs," *Review of Economic Studies*, 89, 181–213.

Bresnahan, T. F. (1981): "Departures from Marginal-Cost Pricing in the American Automobile Industry," *Journal of Econometrics*, 17, 201–227.

Chang, H., J. Chen, S.-W. Hsu, and R. Mashruwala (2018): "The Impact of the Bullwhip Effect on Sales and Earnings Prediction Using Order Backlog," *Contemporary Accounting Research*, 35, 1140–1165.

Childs, G. L. (1967): *Unfilled orders and inventories a structural analysis*, Contributions to economic analysis; 49, Amsterdam: North-Holland.

Cooper, R. and J. Haltiwanger (1990): "Inventories and the Propagation of Sectoral Shocks," *American Economic Review*, 80, 170–190.

Eslava, M., J. Haltiwanger, A. Kugler, and M. Kugler (2004): "The effects of structural reforms on productivity and profitability enhancing reallocation: evidence from Colombia," *Journal of Development Economics*, 75, 333–371.

Fernandes, A. and R. Rigato (2025): "K Wasn't Built in a Day," *IMF Working Papers*, 2025, 1.

Foss, M. F., G. Fromm, and I. Rottenberg (1981): *Measurement of Business Inventories*, National Bureau of Economic Research, Inc.

Foster, L., J. Haltiwanger, and C. Syverson (2008): "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?" *American Economic Review*, 98, 394–425.

Gagliardone, L., M. Gertler, S. Lenzu, and J. Tielens (2023): "Anatomy of the Phillips Curve: Micro Evidence and Macro Implications," *NBER working paper*, 31382.

Görtz, C., C. Gunn, and T. A. Lubik (2024): "What Drives Inventory Accumulation? News on Rates of Return and Marginal Costs," *Journal of Money, Credit and Banking.*

Haltiwanger, J. C. and L. J. Maccini (1989): "Inventories, Orders, Temporary and Permanent Layoffs: An Econometric Analysis," *Carnegie-Rochester Conference Series on Public Policy*, 30, 301–366.

Hottman, C. J., S. J. Redding, and D. E. Weinstein (2016): "Quantifying the Sources of Firm Heterogeneity," *The Quarterly Journal of Economics*, 131, 1291–1364.

Kalouptsidi, M. (2014): "Time to Build and Fluctuations in Bulk Shipping," *American Economic Review*, 104, 564–608.

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

Kryvtsov, O. and V. Midrigan (2012): "Inventories, Markups, and Real Rigidities in Menu Cost Models," *The Review of Economic Studies*, 80, 249–276.

Linz, S., P. Mehlhorn, and S. Wolf-Göbel (2016): "Neue Statistik zum Auftragsbestand in der Industrie," *WISTA*, 1, 33–46.

Meier, M. (2020): "Supply Chain Disruptions, Time to Build, and the Business Cycle," mimeo.

Mertens, M. (2022): "Micro-mechanisms behind declining labor shares: Rising market power and changing modes of production," *International Journal of Industrial Organization*, 81, 102808.

Nevo, A. (2001): "Measuring Market Power in the Ready-to-Eat Cereal Industry," *Econometrica*, 69, 307–342.

Oh, H. and C. Yoon (2020): "Time to build and the real-options channel of residential investment," *Journal of Financial Economics*, 135, 255–269.

Rajgopal, S., T. Shevlin, and M. Venkatachalam (2003): "Does the Stock Market Fully Appreciate the Implications of Leading Indicators for Future Earnings? Evidence from Order Backlog," *Review of Accounting Studies*, 8, 461–492.

Ramey, V. A. and K. D. West (1999): "Inventories," Elsevier, vol. 1 of *Handbook of*

Macroeconomics, 863–923.

Rauch, J. E. (1999): "Networks versus markets in international trade," *Journal of International Economics*, 48, 7–35.

Schaal, E. and M. Taschereau-Dumouchel (2025): "Echoes and Delays: Time-to-Build in Production Networks," Working Paper 1482, Barcelona School of Economics.

Schuh, S. (1996): "Evidence on the Link between Firm-Level and Aggregate Inventory Behavior," Finance and economics discussion series.

West, K. D. (1989): "Order Backlogs and Production Smoothing," in *The Rational Expectations Equilibrium Inventory Model: Theory and Applications*, ed. by T. Kollintzas, New York, NY: Springer US, 246–269.

Zarnowitz, V. (1961): "The Timing of Manufacturers' Orders During Business Cycles," in *Business Cycle Indicators, Volume 1*, ed. by G. H. Moore, Princeton, NJ: Princeton University Press / National Bureau of Economic Research, chap. 14, 420–484.

——— (1962): "Unfilled Orders, Price Changes, and Business Fluctuations," *The Review of Economics and Statistics*, 44, 367–394.

——— (1973): *Orders, Production, and Investment: A Cyclical and Structural Analysis*, National Bureau of Economic Research, Inc.

Appendix A Manufacturing Reports

A.1 Data details

Legislative details. The *Monthly Reports of Manufacturing Establishments* (official name: *Monatsbericht für Betriebe des Verarbeitenden Gewerbes*) are collected by the (State) Statistical Offices of Germany and compiled by the Federal Statistical Office (Destatis). The reporting obligation is based on federal law, in particular the Federal Statistics Act (*BStatG*) and the Law on Statistics in the Manufacturing Sector (*ProdGewStatG*). Order backlog was added to the reports through a legal change in 2012, effective from 2014M1. The change was motivated by the 2008/09 recession, during which the lack of backlog data was considered an impediment to assessing the extent to which plummeting new orders would depress manufacturing activity.

Coverage. The monthly reports cover all manufacturing establishments with at least 50 employees in the reporting year. Destatis also publishes annual establishment-level statistics that include all establishments with at least 20 employees. Relative to the 20+ sample, the 50+ sample covers 49% of establishments, 89% of employment, and 94% of sales in manufacturing.[35] To gauge the importance of establishments with fewer than 20 employees, we compare the 20+ sample to the economy-wide enterprise statistics.[36] Aggregate manufacturing sales among establishments in the 20+ sample amount to 81% of aggregate manufacturing sales, and 78% of employment. However, the comparison is imperfect because the sample of manufacturing firms may include non-manufacturing establishments and exclude manufacturing establishments that are part of non-manufacturing firms. Further details on data collection, data quality, and the questionnaire are provided in Destatis' quality report.[37]

A.2 Sample construction

We next describe the *final backlog sample* used in our analysis. Starting from the full sample, we first restrict to the full backlog sample, which contains establishments[4d] assigned to backlog sectors, and then consecutively apply the cleaning steps described subsequently.

[35] `https://www.destatis.de/DE/Themen/Branchen-Unternehmen/Industrie-Verarbeitendes-Gewerbe/Publikationen/Downloads-Struktur/betriebe-taetige-personen-2040412207004`

[36] For details on the enterprise statistic, see: `https://www.destatis.de/DE/Methoden/Qualitaet/Qualitaetsberichte/Unternehmen/bereichsuebergreifende-unternehmensstatistik.html`

[37] `https://www.destatis.de/DE/Methoden/Qualitaet/Qualitaetsberichte/Industrie-Verarbeitendes-Gewerbe/einfuehrung.html`

Backlog sectors. The backlog sectors, for which backlog is collected, are originally based on Eurostat guidelines based on the understanding that backlog is non-negligible in these sectors (Linz et al., 2016).[38] Throughout our sample, the selection of backlog sectors has remained unchanged. The backlog sectors are defined by a set of two-digit manufacturing sectors, classified according to the German classification system *WZ 2008* that is comparable to the international NACE Rev. 2 standard. The backlog sectors are textiles (*WZ 2008* code 13), cloths (14), paper (17), chemicals (20), pharmaceuticals (21), primary metals (24), fabricated metals (25), electronics (26), electrics (27), machinery (28), automotives (29), and other vehicles (30). Non-backlog manufacturing sectors are the following two-digit sectors: food (10), beverages (11), tobacco (12), leather (15), lumber products (16), print (18), petroleum and coal (19), rubber and plastics (22), other non-metallic minerals (23), furniture (31), other manufacturing (32), repair and installation (33).

Observations with negligible or implausible order book data. We exclude observations for which order books appear either negligible or implausible. For an est[4d] i observed in periods $t_{1i}, \dots, T_i$, we consider all subsamples $\mathbb{T}_{it} = \{t_{1i}, \dots, t\}$, for $t \in (t_{1i}, T_i]$, and determine whether one of the following three conditions holds:

$$(1) \quad B_{it} \leq 0.5 \cdot S_{it}, \quad \text{in at least 90\% of the periods in } \mathbb{T}_{it}$$
$$(2) \quad |S_{it} - O_{it}| < 0.005 \cdot S_{it}, \quad \text{in at least 90\% of the periods in } \mathbb{T}_{it}$$
$$(3) \quad O_{it} < 0.2 \cdot S_{it}, \quad \text{in at least 90\% of the periods in } \mathbb{T}_{it}$$

Criterion (1) captures establishments[4d] for which reported backlog is usually less than half of monthly sales, suggesting that backlog is either negligible or underreported. Distinguishing between these, (2) holds if sales closely track new orders in most periods, leaving little meaningful backlog dynamics and again suggesting that backlog is either of negligible importance or misreported. (3) captures establishments[4d] that usually report far fewer new orders than sales.[39] Rationalizing this pattern would require implausibly large negative growth in new orders, so it likely reflects incomplete new order records. We apply these criteria to each establishment[4d] and each subsample $\mathbb{T}_{it}$ and drop all observations in any subsample that satisfies at least one criterion.

Transitory noise in residual. The residual R_{it} occasionally displays a transitory zigzag pattern: a large residual in one month followed by a residual of similar magnitude but

[38]For comparison, the U.S. Census M3 Survey also collects order backlog only for some manufacturing sectors. The selection overlaps well with the set of backlog sectors in our data.

[39]The inverse case, $S_{it} < 0.2 \cdot O_{it}$ in at least 90% of observations, is empirically irrelevant.

opposite sign in the next month. We interpret such patterns as likely reporting mistakes in the timing of order backlog.[40] We smooth out such transitory zigzags and adjust backlog so that the stock-flow equation remains satisfied. Formally, we define a zigzag as two adjacent residuals that (1) have opposite signs, $R_{it} \cdot R_{it+1} < 0$, (2) are large relative to other order book variables, $0.5(|R_{it}| + |R_{it+1}|) > \min\{B_{it}, O_{it}, S_{it}\}$, and (3) approximately offset each other, $|R_{it} + R_{it+1}| < 0.05(|R_{it}| + |R_{it+1}|)$. We define

$$X_{it} = \min(|R_{it}|, |R_{it+1}|), \tag{A.1}$$

and adjust the residuals and backlog for observations that satisfy the three criteria by

$$R_{it}^{corrected} = R_{it} - sign(R_{it})X_{it}, \tag{A.2}$$
$$R_{it+1}^{corrected} = R_{it+1} - sign(R_{it})X_{it}, \tag{A.3}$$
$$B_{it}^{corrected} = B_{it} + sign(R_{it})X_{it}. \tag{A.4}$$

This correction preserves the stock-flow equation and requires adjusting backlog only in month t.[41] Such two-month zigzags describe 0.7% of observations in the raw sample.

A.3 Descriptive statistics

Table A.1 presents descriptive statistics on the sample. Panel (a) reports the full sample of est[4d]-months 2015M1–2022M12, including all manufacturing sectors and before dropping any data. The sample contains 3.0 million observations and the median est[4d] has 76 employees and monthly sales of 0.9 million EUR in 2019 prices. Panel (b) restricts the sample to the full backlog sample, defined by the backlog sectors above. This subsample contains 1.8 million observations and the establishments[4d] tend to be larger in terms of employees and sales. Panel (c) further restricts the sample (in b) to the final backlog sample. That is, observations with negligible or implausible order book data are excluded, criteria (1)–(3) above, and transitory noise in the residual is smoothed. Sample size drops by 27% to 1.3 million observations. Criterion (3) is quantitatively minor and applies to only 0.4% of observations. Criteria (1) and (2) apply to 22% and 25% of observations in backlog sectors, respectively, and frequently overlap. The final backlog sample covers 98% of aggregate order backlog. Hence, the criteria select on high backlog observations. But the excluded

[40] For example, suppose an order of value Z is recorded in new orders O_{it}, but mistakenly not in B_{it} but only in B_{it+1}. This generates a positive residual $R_{it} = +Z$ followed by a negative one $R_{it+1} = -Z$.

[41] To ensure a non-negative corrected backlog, we define a lower bound for the corrected backlog as $\underline{B}_{it} = \min\{B_{it}, 0.1 \cdot \frac{1}{12}\sum_{j=0}^{11} S_{it-j}\}$. If $sign(R_{it}) < 0$ and $B_{it} + sign(R_{it})X_{it} < \underline{B}_{it}$, we redefine X_{it} as $X_{it} = B_{it} - \underline{B}_{it}$ and compute $R_{it}^{corrected}$ and $B_{it}^{corrected}$ as described above.

observations are otherwise similar, in terms of sales or employment.

Table A.1: Sample construction

	Mean	Sd	P10	P25	P50	P75	P90	Obs
				(a) Full sample				
Employment	139	206	6	38	76	150	313	3,022,069
Sales (S_{it})	2.81	5.83	0.07	0.31	0.91	2.52	6.63	3,022,069
				(b) Full backlog sample				
Employment	162	258	9	45	81	167	359	1,805,063
Sales (S_{it})	3.35	7.33	0.09	0.37	1.03	2.90	7.59	1,805,063
New orders (O_{it})	3.41	7.48	0.06	0.33	1.01	2.95	7.79	1,805,063
Order backlog (B_{it})	14.13	42.34	0.00	0.02	1.61	8.19	29.66	1,805,063
Order book residual (R_{it})	0.02	0.82	-0.18	0.00	0.00	0.00	0.23	1,805,063
Relative residual ($\|R_{it}\|/B_{it-1}$)	0.13	0.43	0.00	0.00	0.00	0.06	0.25	1,362,767
				(c) *Final backlog sample*				
Employment	165	259	11	48	84	173	362	1,315,354
Sales (S_{it})	3.31	7.16	0.10	0.40	1.07	2.92	7.38	1,315,354
New orders (O_{it})	3.46	7.63	0.07	0.36	1.06	3.02	7.74	1,315,354
Order backlog (B_{it})	19.96	55.01	0.16	0.93	3.74	12.95	42.19	1,315,354
Order book residual (R_{it})	0.05	1.00	-0.26	0.00	0.00	0.02	0.36	1,315,354
Relative residual ($\|R_{it}\|/B_{it-1}$)	0.09	0.25	0.00	0.00	0.00	0.05	0.20	1,256,847
				(d) *Final price sample*				
Employment	173	261	15	52	88	182	381	996,903
Sales (S_{it})	3.52	7.26	0.14	0.47	1.20	3.21	7.92	996,903
New orders (O_{it})	3.72	7.87	0.11	0.43	1.21	3.34	8.41	996,903
Order backlog (B_{it})	18.92	47.11	0.23	1.12	4.20	13.79	42.85	996,903
Order book residual (R_{it})	0.06	1.02	-0.27	0.00	0.00	0.02	0.40	996,903
Relative residual ($\|R_{it}\|/B_{it-1}$)	0.08	0.23	0.00	0.00	0.00	0.05	0.19	958,835

Note: The table provides descriptive statistics at the establishment[4d]-month level for 2015M1–2022M12. All variables, except employment, are deflated by the PPI normalized to the year 2019, expressed in EUR million, and winsorized at the 1st and 99th percentile. The last row refers to the absolute order book residual relative to backlog in the previous month.

Table A.2: Distribution of establishments^{4d} across establishments and firms

	Mean	Sd	P10	P50	P90	P99	Obs
Establishments^{4d}/Firm	1.40	1.31	1.00	1	2	5	937,112
Establishments^{4d}/Establishment	1.30	0.79	1.00	1	2	4	1,013,805

Note: The table provides descriptive statistics on the composition of establishment^{4d}-month observations in the final backlog sample for 2015M1–2022M12.

Table A.2 documents the distribution of establishments^{4d} across establishments and firms. The median establishment and the median firm have precisely one establishment^{4d}. The arithmetic average is respectively 1.3 and 1.4. Even at the 99th percentile, the number of establishments^{4d} within establishments and firms is limited, 4 and 5, respectively.

Appendix B Prices

B.1 Details on Step 1

Step 1 of the price construction scheme in Section 3.2 constructs the new order price as a weighted average of the future sales prices associated with the same order cohort

$$p^O_{it} = \sum_{j=\underline{\tau}_{it}}^{\overline{\tau}_{it}} w_{it+j} p^S_{it+j}, \qquad w_{it+j} = \frac{\tilde{S}_{it+j}}{\sum_{k=\underline{\tau}_{it}}^{\overline{\tau}_{it}} \tilde{S}_{it+k}},$$

where $\tilde{S}_{it+k}$ denotes sales attributed to period-t new orders. Formally, we define $\tilde{S}_{it+k}$ as

$$\begin{aligned} \tilde{S}_{it+k} &= S_{it+k}, \quad \text{if } \underline{\tau}_{it} < k < \overline{\tau}_{it}, \\ \tilde{S}_{it+\underline{\tau}_{it}} &= \underline{\zeta}_{it+\underline{\tau}_{it}} S_{it+\underline{\tau}_{it}}, \\ \tilde{S}_{it+\overline{\tau}_{it}} &= \overline{\zeta}_{it+\overline{\tau}_{it}} S_{it+\overline{\tau}_{it}}. \end{aligned}$$

$\underline{\zeta}_{it+\underline{\tau}_{it}}$ is the fraction of net outflows in period $t + \underline{\tau}_{it}$ attributed to period-t new orders, implicitly defined by

$$B_{it-1} = \sum_{k=0}^{\underline{\tau}_{it}-1} (S_{it+k} + R_{it+k}) + (1 - \underline{\zeta}_{it+\underline{\tau}_{it}})(S_{it+\underline{\tau}_{it}} + R_{it+\underline{\tau}_{it}}).$$

$\overline{\zeta}_{it+\overline{\tau}_{it}}$ is the fraction of net outflows in period $t+\overline{\tau}_{it}$ attributed to period-t new orders, implicitly defined by

$$B_{it-1}+O_{it}=\sum_{k=0}^{\overline{\tau}_{it}-1}(S_{it+k}+R_{it+k})+\overline{\zeta}_{it+\overline{\tau}_{it}}(S_{it+\overline{\tau}_{it}}+R_{it+\overline{\tau}_{it}}).$$

The two ζ shares ensure that sales prices in partially overlapping delivery periods receive weights proportional to their information about period-t orders.

B.2 Details on Step 3

Step 3 of the price construction scheme in Section 3.2 describes how real backlog can be computed given an initial real backlog. We uniquely determine the initial real backlog by requiring that the initial nominal and initial real backlog are filled at equal speed.
Let t_{i0} be the first period in which an establishment[4d] is observed, then $b_{it_{i0}-1}$ denotes the initial real backlog and $B_{it_{i0}-1}$ the initial nominal backlog. We first define the time span to fully deplete the initial nominal backlog $B_{it_{i0}-1}$ by

$$\tau_{i0}=\min\left\{\tau\in\mathbb{N}^0\,\middle|\,\sum_{k=0}^{\tau}(S_{it_{i0}+k}+R_{it_{i0}+k})\geq B_{it_{i0}-1}\right\}. \tag{B.1}$$

A part of net outflows in period $t_{i0}+\tau_{i0}$ may fill orders not contained in $B_{it_{i0}-1}$. Let $\zeta_{it_{i0}+\tau_{i0}}$ be the share of net outflows $(S_{it_{i0}+\tau_{i0}}+R_{it_{i0}+\tau_{i0}})$ attributed to the initial backlog. $\zeta_{it_{i0}+\tau_{i0}}$ is implicitly defined by

$$B_{it_{i0}-1}=\sum_{k=0}^{\tau_{i0}-1}(S_{it_{i0}+k}+R_{it_{i0}+k})+\zeta_{it_{i0}+\tau_{i0}}(S_{it_{i0}+\tau_{i0}}+R_{it_{i0}+\tau_{i0}}). \tag{B.2}$$

If the initial real backlog $b_{it_{i0}-1}$ satisfies the same τ_{i0} and $\zeta_{it_{i0}+\tau_{i0}}$, the following equation uniquely pins down the initial real backlog

$$b_{it_{i0}-1}=\sum_{k=0}^{\tau_{i0}-1}(s_{it_{i0}+k}+r_{it_{i0}+k})+\zeta_{it_{i0}+\tau_{i0}}(s_{it_{i0}+\tau_{i0}}+r_{it_{i0}+\tau_{i0}}). \tag{B.3}$$

B.3 Alternative Step 3

The scheme for constructing order book prices proposed in Section 3.2 occasionally yields $b_{it}<0$ (before further adjustment). We propose an alternative Step 3 that ensures $b_{it}\geq 0$.

The alternative backlog price index departs from equation (B.2) above, which we rewrite as

$$p_{it}^{B} b_{it-1} = \sum_{k=0}^{\tau_{it}-1} p_{it}^{S} (s_{it+k} + r_{it+k}) + \zeta_{it+\tau_{it}} p_{it}^{S} (s_{it+\tau_{it}} + r_{it+\tau_{it}}), \tag{B.4}$$

where we used $p_{it}^{S} = p_{it}^{R}$ from Step 2 in Section 3.2. Since the nominal and the real stock-flow equation have to hold simultaneously, (B.4) also has to hold in real terms

$$b_{it-1} = \sum_{k=0}^{\tau_{it}-1} (s_{it+k} + r_{it+k}) + \zeta_{it+\tau_{it}} (s_{it+\tau_{it}} + r_{it+\tau_{it}}). \tag{B.5}$$

Dividing (B.4) by b_{it-1} and using (B.5) yields the order backlog price

$$\begin{aligned} p_{it}^{B} = \sum_{k=0}^{\tau_{it}-1} p_{it}^{S} & \frac{(s_{it+k} + r_{it+k})}{\sum_{k=0}^{\tau_{it}-1} (s_{it+k} + r_{it+k}) + \zeta_{it+\tau_{it}} (s_{it+\tau_{it}} + r_{it+\tau_{it}})} \\ & + p_{it}^{S} \frac{\zeta_{it+\tau_{it}} (s_{it+\tau_{it}} + r_{it+\tau_{it}})}{\sum_{k=0}^{\tau_{it}-1} (s_{it+k} + r_{it+k}) + \zeta_{it+\tau_{it}} (s_{it+\tau_{it}} + r_{it+\tau_{it}})}. \end{aligned} \tag{B.6}$$

In other words, p_{it}^{B} is a weighted average of different sales prices.
Reassuringly, we find the resulting real backlog series from Step 3 and Alternative Step 3 to be highly similar with a correlation of 0.97. Even for the backlog price index, the correlation is still quite high at 0.86.

B.4 Sales prices

We construct establishment[4d]-level sales prices, p_{it}^{S}, by merging our primary dataset, the reports, with the *Monthly Production Census* of the manufacturing sector in Germany.[42] The census records the sales value V_{ijt} and quantity Q_{ijt} of each nine-digit product j produced by establishment[4d] i in month t.
The nine-digit product codes in the census are highly detailed.[43] E.g., code `289317703` includes machines and devices for the beverage industry (excluding presses, mills, etc.) and `245310100` includes light metal foundry products used in the production of road vehicles.

[42]The official name of the dataset is *Monatliche Produktionserhebung*. The census is used to construct an index of industrial production following the EU PRODCOM standard. It is comparable to similar data in other EU countries (e.g., those used in Amiti et al., 2019; Gagliardone et al., 2023). Annual aggregates of the census have been used in prior research (e.g., Mertens, 2022; Bighelli et al., 2022).

[43]The underlying product classification is the GP (*Güterverzeichnis für Produktionsstatistiken*). During our sample, the classification changes from GP 2009 to GP 2019. We use official concordance tables to map 2009 codes to the 2019 classification. For the few unmatched codes, we use establishment[4d]-level production patterns to identify additional matches. If no clear match exists, we treat the unmatched code as a new product introduced in 2019. A similar procedure is applied in Amiti et al. (2019) and Mertens (2022).

The quantity measures of production are time-consistent: pieces for `289317703` and tons for `245310100`. Establishments with at least 50 employees are legally obliged to report in the census, so the census and the reports have nearly identical establishment^{4d}-month coverage. We compute product-level unit prices as

$$p_{ijt}^{S} = \frac{V_{ijt}}{Q_{ijt}}. \tag{B.7}$$

Quantity Q_{ijt} is missing for 24.9% of products (in all periods), so unit prices are unavailable for these observations. We find a particularly high share of missing Q_{ijt} in the two-digit sectors pharmaceuticals and other vehicles. In these sectors, more than 80% of products do not report Q_{ijt}.[44] We further clean implausible observations. V_{ijt} is 0 or 1 Euro in 1.3% of observations and Q_{ijt} is zero in 1.2% of observations. We treat these entries as missing and linearly interpolate single-month gaps. We also smooth over isolated unit-price spikes, which occur in 0.9% of observations, via linear interpolation.[45]

Finally, we aggregate product-level unit prices p_{ijt}^{S} to the establishment^{4d}-level and merge them to the manufacturing reports. Let F_{it} denote the set of products j produced by establishment^{4d} i in month t with non-missing unit price p_{ijt}^{S}. We construct establishment^{4d}-level sales price growth factors using Törnqvist aggregation (as in Eslava et al., 2004):

$$\frac{p_{it}^{S}}{p_{it-1}^{S}} = \prod_{j \in F_{it}} \left(\frac{p_{ijt}^{S}}{p_{ijt-1}^{S}} \right)^{0.5 s_{ijt} + 0.5 s_{ijt-1}}, \quad s_{ijt} = \frac{V_{ijt}}{\sum_{j \in F_{it}} V_{ijt}}. \tag{B.8}$$

The weights s_{ijt} are product sales shares. Following Gagliardone et al. (2023), we set product-level price growth to missing if the absolute log growth factor exceeds one. We create a sales price index by normalizing $p_{i1}^{S} = 1$, where $t = 1$ is the first month in the sample. Missing establishment^{4d}-level growth factors are imputed using the value-weighted average growth factor among observed establishments^{4d} in the same four-digit industry (again following Eslava et al., 2004). We then set p_{it}^{S} to missing whenever its growth factor was imputed.

[44] In addition, observations with missing Q_{ijt} tend to belong to establishments^{4d} with higher backlog, relative to sales, rendering our analysis on the importance of backlog conservative.

[45] We define isolated spikes as months t in which the unit price: 1) grows by a factor of more than two or shrinks by a factor of less than 1/2 between $t-1$ and t, and between t and $t+1$, 2) the product of the growth rate between $t-1$ and t and the growth rate between t and $t+1$ is negative, and 3) the two-month growth factor between $t-1$ and $t+1$ is less than half the average one-month growth factors in t and $t+1$. The latter criterion means the spike is large relative to trend.

Appendix C Time to fill

C.1 Sensitivity analysis

Table C.1: Sensitivity analysis of descriptive statistics on time to fill

	Mean	Cumulative distribution function				Obs
		3 months	6 months	12 months	24 months	
(a) Sales-weighted distributions						
Time to fill (ex ante)	5.45	47.51%	72.77%	91.12%	97.24%	949,734
Time to fill (ex post)	6.83	45.25%	69.17%	89.03%	95.52%	1,238,945
(b) Restriction to final price sample						
Time to fill (ex post)	5.34	48.57%	72.17%	92.07%	98.79%	949,734
(c) Nominal ex-ante time to fill						
Time to fill (ex ante)	5.87	47.27%	69.47%	88.31%	96.62%	1,238,945
(d) Ex-ante time to fill incl. residual						
Time to fill (ex ante)	5.68	51.06%	72.66%	89.29%	96.32%	943,432
(e) Winsorize at 120 months						
Time to fill (ex ante)	6.04	50.39%	71.99%	89.03%	96.37%	949,734
Time to fill (ex post)	5.41	48.94%	72.31%	91.71%	98.46%	1,311,391

Note: The table reports alternative distributions of τ_{it}^{EA} and τ_{it}^{EP}, across establishment^{4d}-months for 2015M1–2022M12. Unless stated differently, the measures are computed according to equation (4.1) and (4.2). Panel (a) presents sales-weighted moments (using S_{it}^{12}, the average nominal sales in the preceding 12 months). Panel (b) presents moments of τ_{it}^{EP} for the final price sample. Panel (c) presents moments of τ_{it}^{EA} computed as in equation (4.1), but using nominal backlog B_{it} and sales S_{it}^{12}. Panel (d) presents moments of τ_{it}^{EA} computed as in equation (4.1), but replacing s_{it} by $s_{it} + r_{it}$. Panel (e) presents moments of τ_{it}^{EA} and τ_{it}^{EP} when winsorizing at 120 months, instead of the baseline 60 months.

C.2 Industry-level evidence

Table C.2: Time to fill across industries

Sector	Mean	Cumulative distribution function 3 months	6 months	12 months	24 months	Obs
		a) Ex-ante time to fill (τ_{it}^{EA})				
Textiles, Cloths, Paper	2.74	74.39%	91.42%	97.64%	99.12%	67,910
Chemical, Pharma	2.47	79.85%	91.31%	96.93%	99.33%	67,255
Primary and Fabricated Metals	5.25	49.56%	72.25%	90.78%	97.57%	314,606
Electronics, Electrics	5.50	52.93%	74.40%	89.54%	96.45%	186,917
Machinery	7.85	38.01%	61.71%	82.86%	93.57%	281,467
Automobiles	6.56	36.15%	61.55%	87.57%	96.56%	28,643
Other Vehicles	3.26	72.58%	87.87%	95.50%	98.13%	2,936
		b) Ex-post time to fill (τ_{it}^{EP})				
Textiles, Cloths, Paper	2.91	75.11%	92.73%	98.35%	99.73%	86,270
Chemical, Pharma	2.89	76.46%	91.46%	97.75%	99.67%	95,186
Primary and Fabricated Metals	5.01	50.72%	74.48%	93.33%	99.05%	387,754
Electronics, Electrics	5.21	50.21%	74.52%	92.48%	98.55%	259,871
Machinery	6.65	36.46%	62.28%	88.03%	97.87%	400,558
Automobiles	5.71	41.36%	65.15%	92.01%	99.17%	60,743
Other Vehicles	12.19	27.70%	45.92%	66.91%	84.77%	21,009

Note: The table reports descriptive statistics at the level of backlog sectors for τ_{it}^{EA} and τ_{it}^{EP}, defined in equation (4.1) and (4.2), across establishment^{4d}-months for 2015M1–2022M12, respectively in the final price sample and the final backlog sample, and winsorized at 60 months.

C.3 Time series evidence

Figure C.1: Distribution of ex-ante time to fill over time

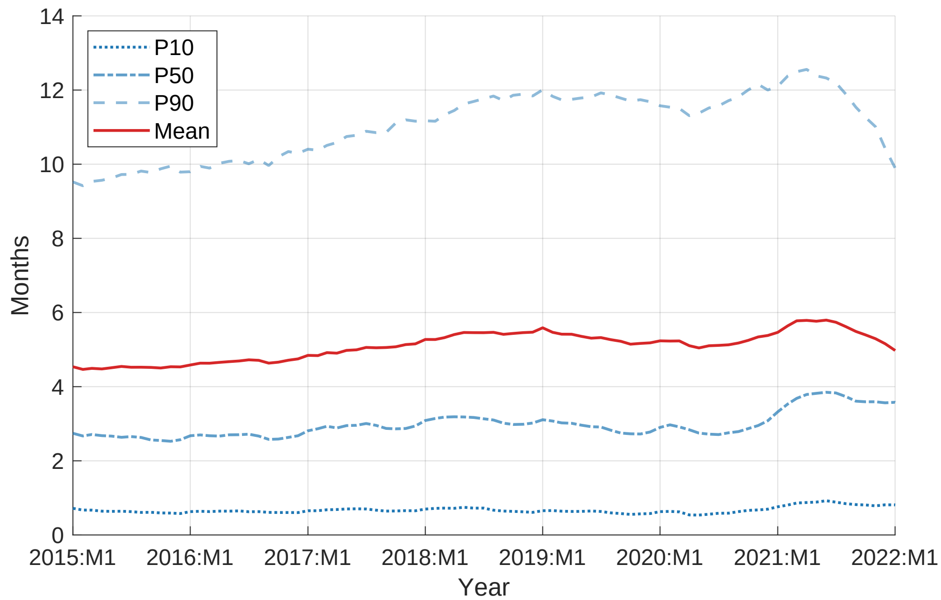

Note: This figure shows the cross-sectional mean, and the 10th, 50th, and 90th percentile of ex-ante time to fill (τ^{EA}) for the final price sample over time. Time to fill is winsorized at 60 months.

C.4 Rauch Classification for German Product Categories

Rauch (1999) classifies internationally traded goods into three categories according to the way they are traded: *(i)* goods traded on organized exchanges, *(ii)* goods with reference prices quoted in trade publications, and *(iii)* differentiated goods whose prices are determined in bilateral relationships. The original classification defines a good as a four-digit level item in the SITC Rev. 2 classification. We adapt the classification principle to the nine-digit product codes in the GP classification in Germany.

In particular, we use GP Rev. 2019 (GP 2019) and the associated product descriptions published by Destatis. In total, there are 7,153 distinct product descriptions in the GP 2019 classification. Based on these descriptions and the classification principle of Rauch (1999) we let an LLM (OpenAI GPT-4) classify each product code according to *(i)*-*(iii)*. We ensure that product codes are classified independently from other product codes, only on the basis of their own description.

At the nine-digit level, the majority of product codes are classified. Of those, 83.7% are classified as differentiated products, 9.5% as reference priced, and 6.8% as goods traded on organized exchanges. The predominance of differentiated products mirrors the original classification of Rauch (1999), in which also most manufactured goods are differentiated. Fewer differentiated products are concentrated where one would expect them: in coal, crude petroleum and natural gas, metal ores and other mining products (GP divisions 05–08), certain food products (10), refined petroleum products (19), basic chemicals (20), and basic metals (24).

Appendix D Shock identification and propagation

D.1 Instrument

Instrument construction. The instrument v_{it} in (5.5) depends on an input cost exposure ϕ_{is} defined in (5.6). We compute α_i^M and α_i^E, the establishment^{4d}-specific share of expenses for materials and energy, using the *Cost Structure Survey* of German manufacturing firms. The survey is an administrative dataset maintained by Destatis and documents the annual firm-level expenses for material inputs, energy, and labor. We can match firms in the survey to establishments^{4d} in our primary dataset, the manufacturing reports. Different from the reports, the survey is annual instead of monthly and covers a random subsample of firms instead of all establishments^{4d} above a size threshold. For our sample from 2015 through 2022, most firms appear only once in the survey. We compute the establishment^{4d}-specific expense shares α_i^M and α_i^E as the respective cost shares of the firm owning the establishment^{4d}. If a firm is surveyed more than once, we compute α_i^M and α_i^E based on the average firm-level expense shares across surveyed years. Finally, we compute $\omega_{k(i),s}$, the share of industry $k(i)$'s expenditure on imported inputs from sector s divided by all expenditures on imported inputs across sectors $s \in S$, using the direct requirements input-output table from Destatis.

Shock-level independence. Following Borusyak et al. (2022), we assess two features of the shocks underlying our shift-share instrument v_{it}: their dependence across sectors, as well as the concentration of exposure across shocks. Sufficiently dispersed exposure to weakly dependent shocks allows chance correlations with unobserved determinants of orders to average out. Common movements in import prices and concentrated exposure can limit this averaging.

We first examine the dependence of import prices across the input sectors S. The estimated within-sector intra-class correlation accounts for 6.8% of the total variation in import prices, which suggests that price movements in foreign input sectors are sufficiently independent between sectors.

We next measure exposure concentration of establishments[4d] to individual input shocks. In total, we have $(22 - 1)$ input industries and 96 months, hence 2,016 possible input-sector-month shifters. We can calculate the inverse Herfindahl–Hirschman index as

$$N_{\text{eff}} = \frac{\left(\sum_t \sum_i \overline{\phi}_i\right)^2}{\sum_t \sum_{s \in S} \left(\sum_{\substack{i \\ k(i) \neq s}} \phi_{is}\right)^2}, \tag{D.1}$$

with $\overline{\phi}_i \equiv \sum_{s \in S \setminus k(i)} \phi_{is}$. The resulting value of 607.7 measures the number of equally weighted sector–month shocks, corresponding to an effective sample size in our TSLS demand estimation.

D.2 Identification

Table D.1: Estimates of the sales-based demand curve

	TSLS		
	(1)	(2)	(3)
	Baseline	*No controls*	*Trimmed*
$\log(p^S_{it})$	-0.5307	-0.8165	-0.5535
	(0.1247)	(0.1230)	(0.1315)
Sales controls	Yes	No	Yes
Establishment[4d] FE	Yes	Yes	Yes
Month $\times$ $\overline{\phi}_i$ FE	Yes	Yes	Yes
R^2	0.3390	0.2415	0.3368
Kleibergen–Paap F	170.70	300.05	161.10
First stage: v_{it}	0.1362	0.1722	0.1301
	(0.0104)	(0.0099)	(0.0103)
N	833,478	869,395	824,294

Note: This table shows the estimated slopes of the demand curves for sales. All columns (1)–(3) use TSLS with instrument v_{it}. All regressions contain establishment[4d] fixed effects and interactions of month fixed effects with the sum of instrument exposure shares $\overline{\phi}_i \equiv \sum_{s \in S \setminus k(i)} \phi_{is}$. Columns (1) and (3) control for $\log(s_{it-1})$ and $\log(s_{it-12})$; column (1) is the main specification. Column (3) excludes establishments above the 99^{th} percentile of average real sales in the estimation sample. The table reports the Kleibergen–Paap F-statistic and the first-stage coefficient of v_{it}. Standard errors in parentheses are clustered at the (4-digit) industry-time level.

Table D.2: Bias of the sales-based demand curve

IV slopes and implied bias			Residualized covariances		
Order demand $\widehat{\beta}^{IV}$	*Sales demand* $\widehat{\gamma}^{IV}$	*Implied bias* $\widehat{\gamma}^{IV} - \widehat{\beta}^{IV}$	$\log s_{it} - \log o_{it}$	$\log p^O_{it} - \log p^S_{it}$	$\log p^S_{it}$
-1.1166	-0.4902	0.6263	0.000297	-0.000002	0.000479

Note: This table decomposes the difference between sales-based and order-based IV demand slopes, using instrument v_{it}. Both slopes are re-estimated on a common sample with identical controls: $\log(o_{it-1})$, $\log(o_{it-12})$, establishment^{4d} fixed effects and interactions of month fixed effects with $\sum_j \phi_{ij}$. Covariances are calculated after residualizing all variables against controls and fixed effects.

D.3 Additional evidence

Figure D.1: Incorrectly identified shocks

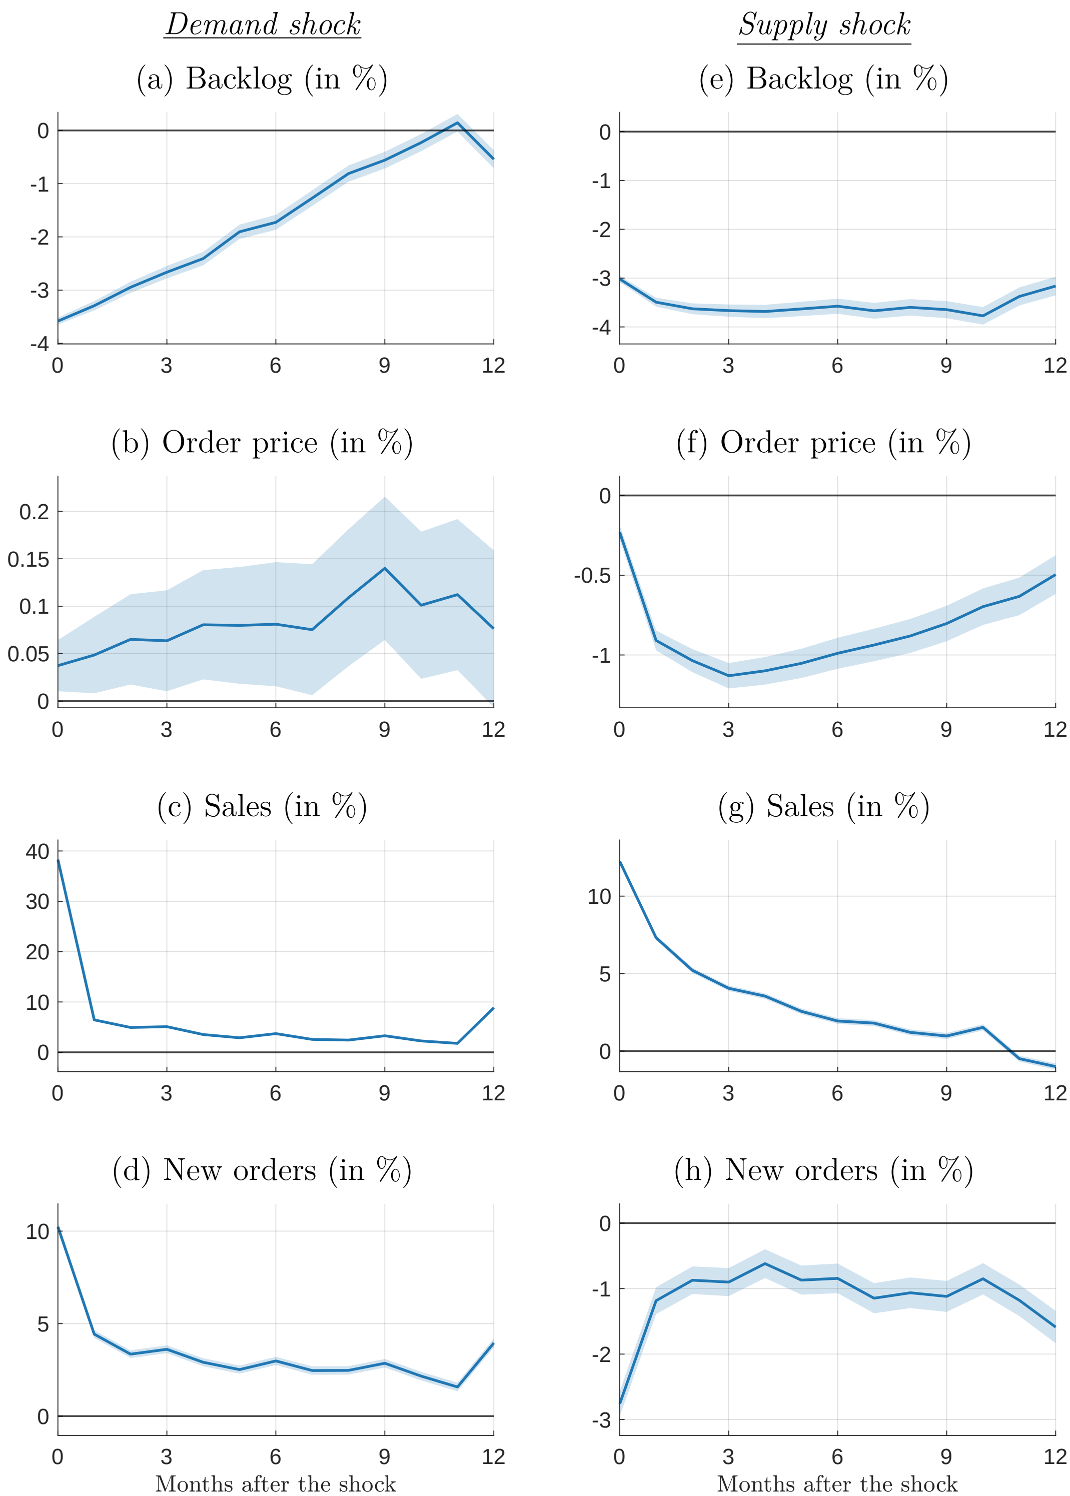

Note: This figure shows the IRFs to expansionary, sales-based demand and supply shocks of one standard deviation expressed in percent. The 95% confidence bands are computed using robust standard errors. The observations were trimmed at 1% and 99% for each horizon.

Figure D.2: Response of order book residual

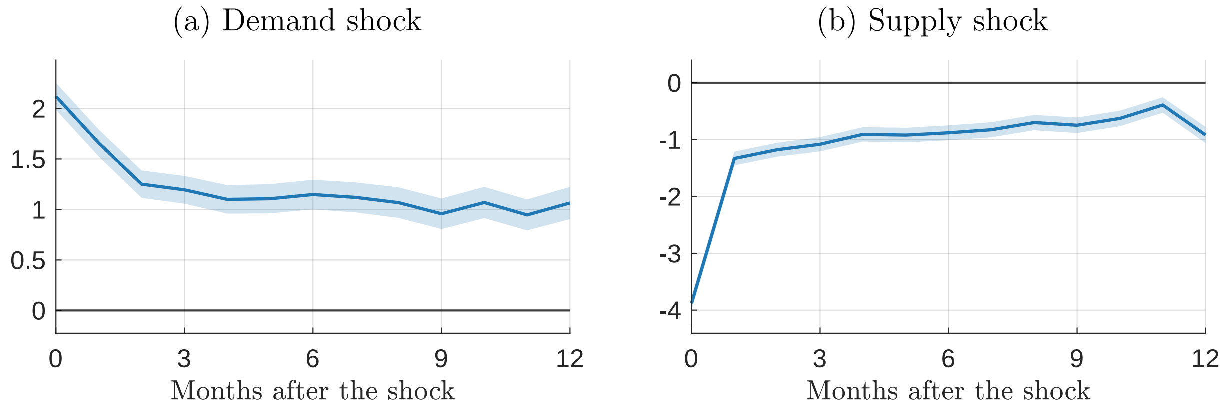

Note: This figure shows the IRFs to a one standard deviation expansionary demand and supply shock. The 95% confidence bands are computed using robust standard errors. The observations were trimmed at 1% and 99% for each horizon. As opposed to the other local projections, the left hand side of the regression is $(r_{it+h} - r_{it-1})/s_{it}^{12}$ to account for the fact that the residual can also be negative.

Appendix E Theory

E.1 Proof of Proposition 1

We organize the proof of Proposition 1 in four steps.

Step 1: Nonlinear optimality conditions. The first step is to substitute both constraints into the objective. Demand implies $o_t = d_t p_t^{-\sigma}$, and the stock-flow constraint then implies $s_t = b_{t-1} + d_t p_t^{-\sigma} - b_t$. Hence the problem can be written with p_t and b_t as the only remaining choice variables:

$$\max_{\{p_t, b_t\}_{t=0}^{\infty}} \quad \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[d_t p_t^{1-\sigma} - \frac{1}{z_t} C_s\Big(b_{t-1} + d_t p_t^{-\sigma} - b_t\Big) - C_\tau\left(\frac{b_t}{s}\right)\right]. \tag{E.1}$$

For notational compactness, we retain s_t below. The FOCs w.r.t. p_t and b_t give

$$0 = (1-\sigma) d_t p_t^{-\sigma} + \sigma d_t p_t^{-\sigma-1} \frac{C_s'(s_t)}{z_t}. \tag{E.2}$$

$$0 = \frac{C_s'(s_t)}{z_t} - \frac{1}{s} C_\tau'\left(\frac{b_t}{s}\right) - \beta \mathbb{E}_t \left[\frac{C_s'(s_{t+1})}{z_{t+1}}\right]. \tag{E.3}$$

Denoting marginal costs by

$$c_t \equiv \frac{C_s'(s_t)}{z_t}, \tag{E.4}$$

we can simplify the FOCs as

$$c_t = \frac{\sigma - 1}{\sigma} p_t, \qquad c_t = \beta \mathbb{E}_t c_{t+1} + \frac{1}{s} C'_\tau \left(\frac{b_t}{s} \right). \tag{E.5}$$

Step 2: Log-linearize optimality conditions. For every positive variable x_t, let $\widehat{x}_t \equiv \log x_t - \log x$, where x denotes its deterministic steady-state value. Define the elasticity of marginal production cost and the normalized local curvature of the backlog penalty as

$$\varphi_s \equiv \frac{s C''_s(s)}{C'_s(s)} > 0, \qquad \varphi_\tau \equiv \frac{\tau C''_\tau(\tau)}{s C'_s(s)} > 0. \tag{E.6}$$

Log-linearizing the demand curve, the stock-flow equation, the definition of marginal cost, the price FOC, and the backlog FOC gives

$$\widehat{o}_t = \widehat{d}_t - \sigma \widehat{p}_t, \tag{E.7}$$

$$\tau(\widehat{b}_t - \widehat{b}_{t-1}) = \widehat{o}_t - \widehat{s}_t, \tag{E.8}$$

$$\widehat{c}_t = \varphi_s \widehat{s}_t - \widehat{z}_t, \tag{E.9}$$

$$\widehat{p}_t = \widehat{c}_t, \tag{E.10}$$

$$\widehat{c}_t = \beta \mathbb{E}_t \widehat{c}_{t+1} + \varphi_\tau \widehat{b}_t. \tag{E.11}$$

Combining (E.7)–(E.10) and replacing $\widehat{c}_t$ by $\widehat{p}_t$ in (E.11) yields two equilibrium conditions:

$$\widehat{p}_t = \alpha_1(\widehat{b}_{t-1} - \widehat{b}_t) + \alpha_2 \widehat{d}_t - \alpha_3 \widehat{z}_t, \tag{E.12}$$

$$\widehat{p}_t = \beta \mathbb{E}_t \widehat{p}_{t+1} + \varphi_\tau \widehat{b}_t, \tag{E.13}$$

where

$$\alpha_1 \equiv \frac{\tau \varphi_s}{1 + \sigma \varphi_s} > 0, \qquad \alpha_2 \equiv \frac{\varphi_s}{1 + \sigma \varphi_s} > 0, \qquad \alpha_3 \equiv \frac{1}{1 + \sigma \varphi_s} > 0. \tag{E.14}$$

Equation (E.12) derives from the price FOC, while (E.13) derives from the backlog FOC. These two equations define the optimality conditions for $\widehat{p}_t$ and $\widehat{b}_t$. Fluctuations of all other variables follow from these two.

Step 3: Solution for backlog. We solve the two optimality conditions by conjecturing a linear policy rule for backlog:

$$\widehat{b}_t = \theta_b^b \widehat{b}_{t-1} + \theta_b^d \widehat{d}_t + \theta_b^z \widehat{z}_t. \tag{E.15}$$

Substituting (E.12) and (E.15) into (E.13), and using $\mathbb{E}_t\widehat{d}_{t+1} = \rho_d\widehat{d}_t$ and $\mathbb{E}_t\widehat{z}_{t+1} = \rho_z\widehat{z}_t$, yields

$$\begin{aligned}
&\alpha_1\left[(1-\theta_b^b)\widehat{b}_{t-1} - \theta_b^d\widehat{d}_t - \theta_b^z\widehat{z}_t\right] + \alpha_2\widehat{d}_t - \alpha_3\widehat{z}_t \\
=\ &\beta\alpha_1\left[\theta_b^b(1-\theta_b^b)\widehat{b}_{t-1} + \theta_b^d(1-\theta_b^b-\rho_d)\widehat{d}_t + \theta_b^z(1-\theta_b^b-\rho_z)\widehat{z}_t\right] \\
&+\beta\alpha_2\rho_d\widehat{d}_t - \beta\alpha_3\rho_z\widehat{z}_t + \varphi_\tau\left(\theta_b^b\widehat{b}_{t-1} + \theta_b^d\widehat{d}_t + \theta_b^z\widehat{z}_t\right). \qquad \text{(E.16)}
\end{aligned}$$

Since the equation must hold for any $\widehat{b}_{t-1}$, it must also hold that

$$\beta\alpha_1(\theta_b^b)^2 - \left[\alpha_1(1+\beta)+\varphi_\tau\right]\theta_b^b + \alpha_1 = 0. \qquad \text{(E.17)}$$

Let $F(x) \equiv \beta\alpha_1x^2 - [\alpha_1(1+\beta)+\varphi_\tau]x + \alpha_1$. Since $\alpha_1 > 0$ and $\varphi_\tau > 0$, $F(0) = \alpha_1 > 0$ and $F(1) = -\varphi_\tau < 0$, so one root lies in $(0,1)$ by the intermediate value theorem. The product of the two roots is $1/\beta > 1$ by Vieta's formula. Hence the other root equals $1/(\beta\theta_b^b) > 1$. The boundedness condition selects the unique stable root

$$\theta_b^b = \frac{\alpha_1(1+\beta)+\varphi_\tau - \sqrt{[\alpha_1(1+\beta)+\varphi_\tau]^2 - 4\beta\alpha_1^2}}{2\beta\alpha_1} \in (0,1). \qquad \text{(E.18)}$$

Matching the coefficients on the demand and supply shifters gives

$$\theta_b^d = \frac{\alpha_2(1-\beta\rho_d)}{\alpha_1[1+\beta(1-\theta_b^b-\rho_d)]+\varphi_\tau} = \frac{\alpha_2(1-\beta\rho_d)}{\alpha_1[(\theta_b^b)^{-1}-\beta\rho_d]} > 0, \qquad \text{(E.19)}$$

$$\theta_b^z = -\frac{\alpha_3(1-\beta\rho_z)}{\alpha_1[1+\beta(1-\theta_b^b-\rho_z)]+\varphi_\tau} = -\frac{\alpha_3(1-\beta\rho_z)}{\alpha_1[(\theta_b^b)^{-1}-\beta\rho_z]} < 0, \qquad \text{(E.20)}$$

where we have used that (E.17) can be rewritten as

$$\alpha_1(1+\beta-\beta\theta_b^b)+\varphi_\tau = \alpha_1(\theta_b^b)^{-1}. \qquad \text{(E.21)}$$

The denominators are positive because $(\theta_b^b)^{-1} > 1 > \beta\rho_d, \beta\rho_z$. Hence the coefficients in the backlog policy rule of Proposition 1 satisfy

$$\theta_b^b \in (0,1), \qquad \theta_b^d > 0, \qquad \theta_b^z < 0. \qquad \text{(E.22)}$$

Step 4: Solutions for prices, sales, and new orders. For $x \in \{d, z\}$, define

$$R_x \equiv \frac{(\theta_b^b)^{-1} - 1}{(\theta_b^b)^{-1} - \beta\rho_x}. \qquad \text{(E.23)}$$

Since $0 < \beta\rho_x < 1 < (\theta_b^b)^{-1}$, both the numerator and denominator are positive. Moreover, the denominator exceeds the numerator by $1 - \beta\rho_x > 0$, so $R_x \in (0,1)$. Substituting the backlog policy (E.15) into the price optimality condition (E.12) gives

$$\widehat{p}_t = \alpha_1(1-\theta_b^b)\widehat{b}_{t-1} + (\alpha_2 - \alpha_1\theta_b^d)\widehat{d}_t - (\alpha_3 + \alpha_1\theta_b^z)\widehat{z}_t. \tag{E.24}$$

Using (E.19)–(E.20) and the definition of R_x,

$$\alpha_2 - \alpha_1\theta_b^d = \alpha_2 R_d, \qquad -\alpha_3 - \alpha_1\theta_b^z = -\alpha_3 R_z. \tag{E.25}$$

Hence the price policy rule is

$$\widehat{p}_t = \underbrace{\alpha_1(1-\theta_b^b)}_{\theta_p^b}\widehat{b}_{t-1} + \underbrace{\alpha_2 R_d}_{\theta_p^d}\widehat{d}_t + \underbrace{-\alpha_3 R_z}_{\theta_p^z}\widehat{z}_t. \tag{E.26}$$

Because $\alpha_2, \alpha_3 > 0$ and $R_d, R_z \in (0,1)$, $\theta_p^d > 0$ and $\theta_p^z < 0$. Combining (E.10) and (E.9) gives $\widehat{s}_t = (\widehat{p}_t + \widehat{z}_t)/\varphi_s$. Substituting the price policy rule yields the sales policy rule

$$\widehat{s}_t = \underbrace{\frac{\alpha_1(1-\theta_b^b)}{\varphi_s}}_{\theta_s^b}\widehat{b}_{t-1} + \underbrace{\frac{\alpha_2 R_d}{\varphi_s}}_{\theta_s^d}\widehat{d}_t + \underbrace{\frac{1-\alpha_3 R_z}{\varphi_s}}_{\theta_s^z}\widehat{z}_t. \tag{E.27}$$

The demand coefficient is positive. The supply coefficient is also positive because $\alpha_3 \in (0,1)$ and $R_z \in (0,1)$ imply $1 - \alpha_3 R_z > 0$. Thus $\theta_s^d > 0$ and $\theta_s^z > 0$. Substituting the price policy rule into the demand curve (E.7) gives the new-orders policy rule

$$\widehat{o}_t = \underbrace{-\sigma\alpha_1(1-\theta_b^b)}_{\theta_o^b}\widehat{b}_{t-1} + \underbrace{(1-\sigma\alpha_2 R_d)}_{\theta_o^d}\widehat{d}_t + \underbrace{\sigma\alpha_3 R_z}_{\theta_o^z}\widehat{z}_t. \tag{E.28}$$

The supply coefficient is positive. The demand coefficient is also positive because $R_d < 1$ and $\sigma\alpha_2 = \frac{\sigma\varphi_s}{1+\sigma\varphi_s} \in (0,1)$. Thus $\theta_o^d > 0$ and $\theta_o^z > 0$.

Collecting (E.15), (E.26), (E.27), and (E.28) proves every sign claimed in Proposition 1. Since $\theta_b^b \in (0,1)$ and the exogenous processes have $\rho_d, \rho_z \in (0,1)$, backlog is stable. Prices, sales, and new orders are linear functions of the stable state vector and are therefore stable as well. □

E.2 Proof of Proposition 2

Given (6.6), the steady-state values of s, o, p, b, and τ remain unchanged. Hence, in the counterfactual the parameters $\varphi_s, \alpha_1, \alpha_2, \alpha_3$ are also unchanged. The parameter that changes from the baseline is φ_τ. We define

$$\widetilde{\varphi}_\tau \equiv \frac{\tau \widetilde{C}''_\tau(\tau)}{s C'_s(s)}. \tag{E.29}$$

By construction, $\widetilde{\varphi}_\tau \to \infty$. The log-linearized backlog FOC in (E.13) can be rewritten as

$$\widehat{b}_t = \frac{\widehat{p}_t - \beta \mathbb{E}_t \widehat{p}_{t+1}}{\widetilde{\varphi}_\tau}. \tag{E.30}$$

Hence, $\widehat{b}_t \to 0$. The remaining choice is $\widehat{p}_t$ with the price FOC in (E.12) simplifying to

$$\widehat{p}_t = \underbrace{\alpha_2}_{\widetilde{\theta}_p^d} \widehat{d}_t + \underbrace{(-\alpha_3)}_{\widetilde{\theta}_p^z} \widehat{z}_t. \tag{E.31}$$

Combining (E.31) with $\widehat{s}_t = (\widehat{p}_t + \widehat{z}_t)/\varphi_s$ gives

$$\widehat{s}_t = \underbrace{\frac{\alpha_2}{\varphi_s}}_{\widetilde{\theta}_s^d} \widehat{d}_t + \underbrace{\frac{1-\alpha_3}{\varphi_s}}_{\widetilde{\theta}_s^z} \widehat{z}_t. \tag{E.32}$$

Finally, substituting (E.31) into the demand curve (E.7) yields

$$\widehat{o}_t = \underbrace{(1 - \sigma\alpha_2)}_{\widetilde{\theta}_o^d} \widehat{d}_t + \underbrace{\sigma\alpha_3}_{\widetilde{\theta}_o^z} \widehat{z}_t. \tag{E.33}$$

Using the baseline coefficients in (E.26), (E.27), and (E.28), together with $R_d, R_z \in (0, 1)$, the differences between the counterfactual and baseline coefficients are

$$\widetilde{\theta}_p^d - \theta_p^d = \alpha_2(1 - R_d) > 0, \qquad \widetilde{\theta}_p^z - \theta_p^z = -\alpha_3(1 - R_z) < 0, \tag{E.34}$$

$$\widetilde{\theta}_s^d - \theta_s^d = \frac{\alpha_2}{\varphi_s}(1 - R_d) > 0, \qquad \widetilde{\theta}_s^z - \theta_s^z = -\frac{\alpha_3}{\varphi_s}(1 - R_z) < 0, \tag{E.35}$$

$$\widetilde{\theta}_o^d - \theta_o^d = -\sigma\alpha_2(1 - R_d) < 0, \qquad \widetilde{\theta}_o^z - \theta_o^z = \sigma\alpha_3(1 - R_z) > 0. \tag{E.36}$$

Finally, $1 - \sigma\alpha_2 = \alpha_3 > 0$ and $(1 - \alpha_3)/\varphi_s = \sigma\alpha_3 > 0$. Hence the counterfactual coefficients retain the signs stated in Proposition 1, and the strict comparisons above establish all claims in Proposition 2. □