

Inventories matter for the transmission of monetary policy: uncovering the cost-of-carry channel*

Diego Rodrigues and Tim Willems[†]

October 2025

[Click here for the most recent version](#)

Abstract

By setting interest rates, monetary policy affects the cost of carrying inventories – giving rise to a “cost-of-carry channel” of monetary policy transmission. Via a simple model, we show that higher inventory carrying costs drive firms, especially those holding larger inventories, to cut their prices. We test this hypothesis using data from the U.S. goods, housing, and oil markets – finding robust evidence supporting the cost-of-carry channel. We then introduce this channel into a New Keynesian setup and show that it makes optimal policy more focused on inflation stabilization when inventories are more plentiful – the reason being that the central bank faces a more favorable sacrifice ratio in such an environment.

JEL-classification: E30, E31, E32, E52, E58.

Key words: inventories, monetary policy, monetary transmission mechanism, inflation.

*First draft: September 2024. We thank Wouter den Haan, Xavier Ragot, Felipe E. Saffie, Pedro Teles, and many seminar audiences for useful comments and discussions. The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England or its committees.

[†]Rodrigues: Université du Québec à Montréal (UQAM) (de_sousa_rodrigues.diego@uqam.ca); Willems: Bank of England and Centre for Macroeconomics (tim.willems@bankofengland.co.uk).

1 Introduction

This paper offers empirical evidence in support of a “cost-of-carry channel” of monetary policy transmission and analyzes implications for optimal policy. The cost-of-carry channel captures the notion that, by setting interest rates, the central bank also shapes the costs of carrying inventories – and thereby firms’ incentives to keep (or shed) them. Since changing prices is one way to manage inventory levels, there may well be important links from the cost of carrying inventories, to inflation dynamics.

At a narrative level, commentators and firms often refer to inventory levels when discussing pricing decisions. As for example noted by [Robinson \(2022\)](#), in an *Investors Chronicle* article titled “Interest rates could spell trouble for inventories, liquidity, and IPOs”:

Carrying costs can rise appreciably as interest rates climb, a worrisome prospect given that they can represent an estimated 25-30 per cent of overall inventory value.

*This has become a more acute issue since the pandemic, as companies accelerated the trend towards building resilience into supply chains by increasing inventory levels. Nike is merely the latest big retailer to warn that, along with unfavourable currency movements, its earnings have come under pressure through the inventory-build it undertook following the pandemic and the **subsequent discounts aimed at alleviating the situation**. Part of the sportswear giant’s profit shortfall will be linked to increased carrying costs. [Emphasis added]*

While the above quote mentions Nike, other companies have also alluded to inventory levels in relation to their pricing decisions. [Trudell \(2024\)](#) for example mentions how “Tesla is slashing prices (...) in a bid to clear its biggest-ever stockpile. (...) Tesla is offering the deals after producing 46,561 more vehicles than it delivered in the first quarter, adding more cars to inventory than ever before”. Similar considerations have been raised in relation to retailers,¹ consumer goods producer Unilever ([Dominguez 2023](#)), while [Williamson \(2023\)](#) notes how “destocking policies are meanwhile adding to the downturn in pricing power”.

It is interesting to note the timing of these articles, in particular how they post-date the recent rise in interest rates (with some of them explicitly referring to this development). Consulting firms specializing in inventory management often point out that “lean” inventory management becomes more consequential when rates are higher. SAFIO Solutions, for example, notes on their website:² “As interest rates increase (...) the costs of carrying excess

¹See [CBS \(2022\)](#), which is titled “Target’s profit craters after it cut prices to clear inventory”, and [Reuters \(2022\)](#): “U.S. retailers’ ballooning inventories set stage for deep discounts”.

²See <https://safiosolutions.com/increasing-interest-rates-carrying-excess-inventory-can-have-an-even-greater-effect-on-your-cash-flow/>. Very similar points are made by Rackbeat (a provider of warehouse management systems), in a post titled “How an Inventory System Helps You Counteract the Red Hot

inventory will be increasing as well, impacting your company’s bottom line.” It then goes on to mention that “capital costs are typically the largest portion of total carrying costs. Capital costs represent the cash that is being tied up in the inventory. These costs include the money spent on the inventory, interest paid on the purchase, and the opportunity cost of the money invested in the inventory rather than other investments like mutual funds.”³ It is furthermore striking how many Western companies only adopted “just-in-time” inventory management when interest rates started soaring in the early 1980s (“post-Volcker”), even though the idea had been around for decades (Petersen 2002).

Next to goods markets, carrying costs may also be relevant in markets for housing services (more on which below) and commodities. The latter is an insight dating back to (at least) Deaton & Laroque (1992, 1995, 1996). Frankel (2008a,b, 2014), in particular, makes the case that the rate of interest is an important driver of oil prices (with higher interest rates providing greater incentives to economize on oil inventories, which raises available supply, thus depressing the price). In line with this notion, Miranda-Pinto et al. (2023) document how commodity prices tend to fall in response to a US monetary policy tightening – to a degree that is increasing in the storability of the commodity. This supports the cost-of-carry channel, over a parallel general equilibrium channel that operates by slowing down aggregate demand via more conventional transmission channels (Miranda-Pinto et al. 2024).

The essence of our argument can be captured through a simple model, developed in Section 3. There, we show that as inventory carrying costs rise (e.g., due to a monetary tightening), firms are incentivized to lower their prices to economize on inventory holdings.⁴ Firms carrying more inventory when a shock hits are more exposed to this dynamic, and thus have a stronger incentive to cut prices.

Informed by our simple model, we proceed (in Section 4) by testing whether these forces are at play in U.S. data. There, we find that a contractionary monetary policy shock does more to lower goods prices when retailers are sitting on more inventories. A potential worry is that our analysis relies on data which are aggregated at too high a level (summing across many different goods, in a way that biases results). To address this concern, we also examine two specific markets in isolation – namely those for oil and housing. Such analyses enable us to measure the key concepts (prices and inventory levels) with greater precision, while housing and oil are also major components of the consumption basket – making their prices of direct interest.⁵ For the housing market, we show that monetary policy has a stronger

Interest Rate” (<https://rackbeat.com/en/how-an-inventory-system-helps-you-counteract-the-increased-interest-rates/>).

³Also see Copeland et al. (2019) who refer to the example of a car dealership whose interest rate carrying costs (associated with their inventory of cars) averaged about 7% of gross profits over 2002-2011.

⁴One can also view this by noting that the accumulation of inventories is a form of investment – with the incentives to invest typically being negatively related to the interest rate.

⁵In the U.S., the shelter component accounts for over 30% of the CPI basket. For PCE, the housing-related share stands at over 15%. While the *direct* share of oil prices is lower, they are an important driver of price dynamics via their prevalence throughout the supply chain (Baqae & Rubbo 2023).

impact on the cost of housing services when housing inventory (the fraction of unoccupied homes) stands higher. High vacancy rates can be thought of as reducing landlords’ market power, and combined with the opportunity cost of higher interest rates, this may incentivize landlords to lower prices to fill properties more quickly. Looking at the oil market, we confirm that oil prices are more sensitive to U.S. monetary policy shocks when oil inventories are higher.

All exercises thus confirm the key prediction of our model: inventories matter to the transmission of monetary policy, with higher inventory levels making sellers more responsive (in the conventional direction) to changes in the monetary policy stance. This suggests that the cost-of-carry channel may be worthy of more attention than it has hitherto received in the monetary policy literature (where it is not mentioned in standard treatments of the monetary transmission mechanism; see, e.g., [Boivin et al. \(2010\)](#)). In our attempt to lessen this shortfall, we augment the standard New Keynesian model with inventories – showing how the associated cost-of-carry channel modifies optimal policy prescriptions. In particular, we show that the central bank faces a more favorable sacrifice ratio when inventories are plentiful – making it optimal to focus more on inflation stabilization in such an environment.

2 Related literature

Inventory dynamics have a rich history in business cycle models, as inventories are typically thought to account for a significant share of fluctuations in GDP ([Blinder & Maccini 1991](#), [Fitzgerald 1997](#), [Ramey & West 1999](#)). The idea central to this paper, that higher interest rates give firms a stronger incentive to economize on their inventory holdings, has been alluded to before (see, e.g., [Lieberman \(1980\)](#), [Irvine \(1981\)](#), [Blinder \(1981\)](#), [Akhtar \(1983\)](#), [Maccini et al. \(2004\)](#), [Alessandria et al. \(2010\)](#), [Kim \(2021\)](#)). While some of the aforementioned papers have offered empirical support for this hypothesis,⁶ other papers (such as [Maccini & Rossana \(1984\)](#), [Ramey \(1989\)](#), [Kryvtsov & Midrigan \(2013\)](#) and [Benati & Lubik \(2014\)](#)) have failed to do so, which has contributed to the theory’s declining popularity over time.

Armed with recent progress in monetary policy shock identification (in particular: the availability of high-frequency monetary policy shocks), we revisit this debate. When doing so, we deviate from the earlier literature along two dimensions. First, informed by our simple model (presented in Section 3) and aided by improved data availability, we broaden our focus to look at the response of prices. This contrasts with the aforementioned earlier literature (notable exceptions being [Alessandria et al. \(2010\)](#) and [Kim \(2021\)](#)), which solely looked at

⁶Also see [Gürkaynak et al. \(2022\)](#). Additional support is presented in papers focusing on the credit channel of monetary policy. [Gertler & Gilchrist \(1994\)](#) find that smaller firms’ inventory holdings are particularly sensitive to borrowing costs, while [Kashyap et al. \(1993\)](#) and [Kashyap et al. \(1994\)](#) make the same point for firms that are bank-dependent.

outcomes in firms’ inventory holdings (which are challenging to measure – especially at a high frequency, whereas analyzing outcomes at a lower frequency might bias results; [Jacobson et al. \(2023\)](#)). Our model, however, suggests that also *the price response* to monetary policy shocks should vary with inventory levels, which is an hypothesis that is arguably cleaner to test than looking at observed changes in inventory holdings.⁷ Second, we take our analysis beyond aggregate data (as those may give rise to worries about aggregating across different goods). In particular, we also test our main hypothesis in two specific markets: those for housing and oil. Across outcome variables, we find broad-based evidence supporting the cost-of-carry channel.

More generally, the wider inventory literature has mostly evolved around three stylized facts:

1. Production is more volatile than sales;
2. Inventories are procyclical;
3. The ratio of inventories-to-sales is countercyclical (implying that sales display stronger procyclicality than inventories).

While an early literature treated inventories as a way to smooth production over the cycle, this approach has fallen out of favor as it is inconsistent with stylized fact #1 ([Blinder 1986](#), [Eichenbaum 1989](#)). Instead, scholars have tried to reconcile (some of) the above stylized facts by modeling inventories as a factor that directly boosts sales ([Kahn 1987](#), [Bils & Kahn 2000](#)), whereas others have approached the issue by focusing on non-convex production costs (leading to “production bunching”; [Ramey \(1991\)](#)) or by taking an (S,s)-type approach ([Khan & Thomas 2007](#)). [Wen \(2011\)](#) and [McMahon \(2012\)](#) show how the stylized facts can be matched by introducing lags between a good being produced and a good becoming available for sale (leading to a stockout avoidance motive – the modeling approach we will adopt below). [Kryvtsov & Midrigan \(2013\)](#) do so via a model combining strongly procyclical marginal costs with countercyclical markups. [Den Haan & Sun \(2024\)](#) augment a standard New Keynesian model with a “sell friction”, which enables their model to replicate key stylized facts whilst also highlighting the importance of inventories for business cycle fluctuations. Of note, their model also offers a fully microfounded environment in which the cost-of-carry channel arises.

⁷With respect to inventory levels, standard logic predicts that higher interest rates should lower inventory holdings (as maintaining them becomes costlier). However, when all firms attempt to shed their inventories by cutting prices (leaving their relative prices unchanged), inventory holdings might show relatively little movement in aggregate. That is: the process of price adjustment might limit the response of quantities in general equilibrium. This concern might be particularly acute when the intertemporal substitution elasticity on the consumer side is low and/or when they become less willing to hold inventories too as interest rates rise ([Copeland et al. 2019](#)). In such cases, one would still see the cost-of-carry channel operate on prices though, which motivates our focus.

Relative to this last literature, our objective is more focused: rather than wishing to replicate all stylized facts (which are unconditional in nature, averaging over all shocks driving the business cycle – where this average might not be dominated by monetary shocks; [Angeletos et al. \(2020\)](#)), we mainly wish to understand *how monetary policy interacts with inventory levels* when it comes to shaping inflation dynamics.

3 Simple Model

The core of our argument can be captured through a very simple model. Consider a profit-maximizing firm that enters the period with an inventory of X_0 units, which are carried over from the past. The firm controls its production level (Y) and the price it charges (P), which implies a level of inventories (X) to carry into the next period. Both production and the carrying of inventories involve quadratic costs. Goods produced will only be available for sale in the future (not explicitly modeled here but addressed in further detail below), meaning that inventories arise due to the presence of demand uncertainty and a precautionary stockout-avoidance motive ([Kahn 1987](#), [Wen 2011](#)). The firm’s problem can be represented as:

$$\begin{aligned} \max_{p,y} \quad & PS(P) - \psi_y \frac{Y^2}{2} - \psi_x \frac{X^2}{2} + \vartheta Q, \\ \text{s.t.} \quad & X = X_0 - S(P), \\ & Q = X + Y, \\ & X \geq 0, \end{aligned} \tag{1}$$

where $S(P)$ is the demand function, assumed to be continuous, three-times differentiable, and with $S'(P) < 0$; the demand function also hosts a stochastic term (which we suppress for notational convenience), making the firm uncertain as to what level of demand it can expect to realize. The cost of producing Y units is given by $\psi_y \frac{Y^2}{2}$, and the cost of carrying X units in inventory by $\psi_x \frac{X^2}{2}$. The final term in the objective function, $+\vartheta Q$, serves as a shorthand to represent the positive value of carrying goods over into the future, where $\vartheta > 0$, and Q , the sum of end-of-period inventories X and production Y , represents the number of goods available for sale in the future.

As the firm begins with an inventory of X_0 , it will have $X = X_0 - S(P)$ units left in inventory after selling $S(P)$ units. When setting its price P , the firm considers the relationship between the price it charges, the quantity of goods it will sell, and, consequently, the amount of inventory it will carry forward – subject to a quadratic cost governed by ψ_x . The latter can be thought of as the opportunity cost of the funds being “locked up” in the inventory or, in case the firm is borrowing, the cost of it having to borrow additional working capital. We think of this cost as being an increasing function of the central bank’s policy

rate r , i.e. $\psi_x = \psi_x(r)$ with $\psi'_x(r) > 0$.

Solving the problem described by (1) leads a profit-maximizing firm to set its optimal production and price as follows:

$$Y = \frac{\vartheta}{\psi_y}, \quad (2)$$

$$0 = [P + \psi_x(X_0 - S(P)) - \vartheta] S'(P) + S(P). \quad (3)$$

The production component of the model (2) is intentionally simplified, allowing us to focus on price setting as governed by the implicit function in (3). In particular, we are interested in understanding how a firm's "exposure" to inventory carrying costs – reflected by its initial inventory level, X_0 – influences its pricing strategy when faced with changes in inventory carrying costs, ψ_x . This brings us to the following proposition:

Proposition 1. (*Price setting*) *At any interior optimum where the non-negativity constraint on inventories does not bind, as the cost of carrying inventories rises, profit-maximizing behavior induces the firm to lower its price, i.e.:*

$$\frac{\partial P}{\partial \psi_x} < 0.$$

The strength of this effect is increasing in the firm's pre-existing inventory level X_0 .

Proof. Applying the Implicit Function Theorem to (3), we obtain:

$$\frac{\partial P}{\partial \psi_x} = \frac{X_0 - S(P)}{\psi_x S'(P) + \frac{S''(P)}{S'(P)^2} S(P) - 2}.$$

If the non-negativity constraint binds ($X_0 - S(P) = 0$), then locally P is pinned down by $S(P) = X_0$ and is independent of ψ_x ; hence $\frac{\partial P}{\partial \psi_x} = 0$ at the corner. At an interior optimum, $S(P) < X_0$, we have $\frac{\partial P}{\partial \psi_x} < 0 \Leftrightarrow \psi_x S'(P) + \frac{S''(P)}{S'(P)^2} S(P) - 2 < 0$. Given that $S'(P) < 0$, this condition holds when $\frac{S''(P)}{S'(P)^2} S(P) < 2$. From the chain rule, it follows that $S'(P) = \frac{1}{P'(S)}$ and $S''(P) = -S'(P)^3 P''(S) = -\frac{P''(S)}{P'(S)^3}$. Using these relationships, we can rewrite $\frac{S''(P)}{S'(P)^2} S(P) = -\frac{P''(S)}{P'(S)} S(P)$, which reflects the convexity of the demand curve. With respect to this object, Mrázová & Neary (2017) show that profit-maximizing behavior guarantees that $-\frac{P''(S)}{P'(S)} S(P) < 2$, which implies $\frac{S''(P)}{S'(P)^2} S(P) - 2 < 0$, thereby proving that $\frac{\partial P}{\partial \psi_x} < 0$. The second part of the proposition follows trivially from the observation that the steepness of this derivative increases with X_0 . $\square$

Proposition 1 conveys the logic, frequently alluded to by many firms (recall the quotes featured in the Introduction) that as inventory carrying costs rise, for example due to an

interest rate hike, firms gain an incentive to lower their prices in an attempt to economize on inventory holdings. Crucially, our model illustrates that firms carrying more inventory when the shock hits (i.e., firms with higher X_0) are more exposed to this channel and thus have the strongest incentive to cut their prices.

This implies that the cost-of-carry channel of monetary policy transmission can also be tested by looking at the response *of prices*, as opposed to that of inventory levels themselves. The latter has been the traditional focus of the literature (recall the references in Section 2), but inference there is complicated by price adjustments potentially dampening the response of quantities (recall footnote 7).

4 Empirical evidence

In Section 4.1 we will test the model’s prediction that – when the cost-of-carry channel is at play – prices should be more responsive (in the conventional direction) to the stance of monetary policy when inventories stand at a higher level. Section 4.2 will subsequently revisit the traditional focus of the literature: the dynamics of inventory levels themselves. In both cases, we end up finding broad support for the cost-of-carry channel.

4.1 Empirical evidence: prices are more responsive to monetary policy when inventories are plentiful

To test our model’s prediction and, thereby, the cost-of-carry channel of monetary policy transmission, we run Local Projections (LPs) of the following form on monthly data:

$$\Delta^h \ln P_{t+h} = \alpha_h + \beta_h MPS_t + \gamma_h (MPS_t \times INV_t) + \delta_h Z_t + \epsilon_{t,h}, \quad (4)$$

where $\Delta^h \ln P_{t+h} \equiv \ln P_{t+h} - \ln P_{t-1}$ is the cumulative change in the natural log of prices over h months. The variable “ INV ” serves as a proxy for inventory levels. Finally, “ MPS ” is the monetary policy shock, which we draw from the series provided by [Bauer & Swanson \(2023\)](#).⁸ We run the LPs by controlling (in Z_t) for lags of (i) the shocks “ MPS_t ”, (ii) the lagged, logged price level “ $\ln P_{t-1}$ ”, (iii) the inventory metric “ INV_t ”, and (iv) the interactions between shocks and inventory levels “ $MPS_t \times INV_t$ ” (but, as shown in Appendix A, results are robust to including further controls). For all of these controls, we include 12 lags.

The remainder of this section displays the IRFs resulting from estimated versions of (4). First, Section 4.1.1 uses aggregate data to examine how the response of the goods component

⁸We have found this shock series to consistently produce intuitive responses in core variables (output, consumer prices, unemployment, equity indices) building confidence that the series captures true monetary policy shocks.

of the PCE price index varies with retailers’ inventory holdings. However, one might be concerned about aggregation issues here, as it is not obvious how price and inventory data should be summed across various types of goods for the purposes of our exercise.⁹

Given such concerns, we proceed (in Sections 4.1.2 and 4.1.3) by extending our analysis to specific markets – namely, those for housing services and oil. In these markets, the core concepts underpinning our theory (prices and inventory levels) are relatively easy to define and measure, while both are also major components of the consumption basket (recall footnote 5). In both cases, we continue to find support for the cost-of-carry channel of monetary policy.

4.1.1 Goods

We first test our theory in the market for goods (i.e., neglecting services, where the concept of “inventories” is generally more nebulous, but see Section 4.1.2 for the exception that is shelter). To do so, we use the goods-component of the PCE price index as our measure of “ P ” in (4). For inventories “ INV ”, we use the ratio of inventories to sales held by retailers (FRED code: RETAILIRSA).¹⁰ The sample runs from 1992m1 (when the inventory-sales ratio becomes available) to 2019m12 (to exclude Covid-driven dynamics).

As Figure 1 shows, the PCE goods-price index does not show a clear response to the monetary policy shock when the inventory-sales ratio stands at its historical average level. This can be consistent with the notion that monetary policy has – generally speaking – greater leverage over prices of services, which are more closely linked to wages.

Figure 2 displays our key result. It traces out the estimates for γ_h in equation (4), which captures how the responsiveness of prices varies with retailers’ inventory holdings. In this case, the negative coefficients indicate that a contractionary monetary policy shock ($MPS_t > 0$) *does more* to lower goods prices when retailers are sitting on more inventories (higher INV_t). This confirms the key prediction from our model, that inventories matter for the transmission of monetary policy – with higher inventory levels making prices more responsive (in the conventional direction) to changes in the stance of monetary policy (remember Proposition (1)). Note that this result is also consistent with Kim (2021), who finds – using the failure of Lehman Brothers as a quasi-experiment – that firms hit by a negative credit supply shock are more inclined to cut their prices to liquidate inventory (and generate cash, which is now more valuable to them); earlier work by Alessandria et al. (2010) reports

⁹Ideally, we would have access to detailed, high-frequency data on inventories *at the goods level* and the ability to match these with corresponding goods-level prices. The closest to this ideal is Kim (2021), who built a dataset of prices and inventory holdings *at the firm level*. This proved a major challenge, however: even after linking several proprietary datasets, Kim’s data only cover grocery items from 2004 to 2011 at the quarterly frequency.

¹⁰Here, looking at inventory holdings *relative to sales* is a convenient way to control for the level of demand (see Mehrotra et al. (2025) for a related approach).

similar findings by looking at inflation dynamics following large devaluations accompanied by interest rate hikes.




Figure 1: Response of PCE-goods price index to a 25-bp contractionary monetary policy shock, estimated via equation (4), when the inventory-sales ratio “ INV_t ” stands at its historical average (INV_{avg}). The figure plots $\hat{\beta}_h + \hat{\gamma}_h \cdot INV_{avg}$. Shaded area represents the 90% confidence interval, calculated via Newey-West standard errors.




Figure 2: Additional response of PCE-goods price index to a 25-bp contractionary monetary policy shock, due to a unit increase in the inventory-sales ratio, estimated via equation (4). The figure plots $\hat{\gamma}_h$. Shaded area represents the 90% confidence interval, calculated via Newey-West standard errors.

As shown in Appendix A, our core result is robust to adding further controls, such as the

5-year Treasury yield and the rate of unemployment, which suggests that we are not picking up variations driven by business cycle fluctuations.

One may, however, worry that we are conducting our exercise at a level of aggregation that is too high to be useful. Consequently, the next two subsections repeat our analysis in the markets for housing services (Section 4.1.2) and oil (Section 4.1.3), where we can analyze the issue at a more “micro” level (linking prices and inventory levels more directly).

4.1.2 Housing

To measure the price of housing services, we use the CPI owner’s equivalent rent (OER) series (but results are robust to using the housing-component in the PCE series instead, or to using the “shelter” component of the CPI, which is slightly broader than OER).¹¹ In this application, the variable “ INV ” represents the fraction of homes being vacant, which represents the inventory of homes looking to be utilized (see Appendix B for details on how this variable is constructed). The lag structure and set of controls for our regression is analogous to the one used in Section 4.1.1 (tying our hands in that regard). The sample runs from 1988m2 (when the monetary policy shock series becomes available) to 2019m12.

Results following from estimating (4) are plotted in the next two figures. Figure 3 shows that, when the housing vacancy rate stands at its sample average, a monetary tightening has virtually no effect on the cost of housing services. Looking at the coefficient “ γ ” on the interaction term (in Figure 4) however demonstrates that monetary policy has greater leverage when more properties are vacant – i.e., when there is a greater inventory of housing looking for an occupant.¹² As shown in Appendix A, this result is robust to adding further controls. This again supports the prediction from our theoretical model, that inventories matter for the transmission of monetary policy – with higher inventory levels making prices more responsive (in the conventional direction) to changes in the stance of monetary policy. This is consistent with the notion that a tighter housing market (lower INV_t) enables landlords to pass on any increases in their borrowing costs (e.g., following a monetary policy tightening). However, when there are more vacant properties (higher INV_t), landlords have less market power and they become more inclined to lower their price in response to a monetary contraction – reflecting the higher opportunity cost of not having the property occupied. As Figure 14 in Appendix A shows, our main result is also visible when using the *overall* CPI as dependent variable in (4) – which is perhaps to be expected, given that OER accounts for about one-third of the total CPI basket. A similar picture emerges when

¹¹Relative to OER, the CPI-shelter series also includes “lodging away from home” and insurance costs, among other items.

¹²The lagged nature of the response in Figure 4 is to be expected given the construction of the OER series (which not only looks at rentals offered on the market contemporaneously, but takes into account that rents only tend to change when leases expire; see Conner et al. (2024) and Cotton (2024) for more details on the calculation of OER).

looking at the overall PCE index.




Figure 3: Response of CPI OER to a 25-bp contractionary monetary policy shock, estimated via equation (4), when the home vacancy rate “ INV_t ” stands at its historical average (INV_{avg}). The figure plots $\hat{\beta}_h + \hat{\gamma}_h \cdot INV_{avg}$. Shaded area represents the 90% confidence interval, calculated via Newey-West standard errors.




Figure 4: Additional response of CPI OER to a 25-bp contractionary monetary policy shock, due to a 1-pp increase in the home vacancy rate, estimated via equation (4). The figure plots $\hat{\gamma}_h$. Shaded area represents the 90% confidence interval, calculated via Newey-West standard errors.

4.1.3 Oil

Alongside housing, the oil market is another natural candidate to test our theory. In this case, we again have data on both prices and inventories. For the latter, we rely on OECD

petroleum stocks data (also used in Kilian & Murphy (2014) and Känzig (2021)), while the former is proxied by the WTI price (sourced from FRED).¹³ Since the inventory data are trending, we first detrend this series using the HP filter with a smoothing parameter of 129,600, as recommended by Ravn & Uhlig (2002) for monthly data (but, as documented in Appendix A, similar results are obtained after linear detrending). The sample runs from 1988m2 (when the monetary policy shock series becomes available) to 2019m12.

First, Figure 5 captures the response of oil prices when detrended oil inventories stand at their sample average. It does not point to a strong, direct impact – with the medium-run estimate taking on the unintuitive sign, if anything.

Next, Figure 6 presents our key result: U.S. monetary policy shocks have a stronger effect (in the conventional direction) when oil inventories stand at a higher level. As with our earlier findings, Appendix A demonstrates that this result – which is again in line with cost-of-carry logic – is robust to the inclusion of further controls.




Figure 5: Response of oil prices to a 25-bp contractionary monetary policy shock, estimated via equation (4), when the oil “ INV_t ” stands at its historical average (INV_{avg}). The figure plots $\hat{\beta}_h + \hat{\gamma}_h \cdot INV_{avg}$. Shaded area represents the 90% confidence interval, calculated via Newey-West standard errors.

¹³The inventory data may not be all-encompassing, but – as noted by Frankel (2014) – even data with partial coverage can be sufficient, as all players in the underlying market, which is global in nature, are subject to the same forces, making them likely to move in similar directions over time.




Figure 6: Additional response of oil prices to a 25-bp contractionary monetary policy shock, due to a 1-pp increase in oil inventory, estimated via equation (4). The figure plots $\hat{\gamma}_h$. Shaded area represents the 90% confidence interval, calculated via Newey-West standard errors.

4.2 Empirical evidence: firms economize on inventory holdings when monetary policy tightens

As mentioned in Section 2, the traditional approach to test the cost-of-carry channel has been to look at the behavior of inventories more directly. If the channel is relevant, sellers should have a stronger incentive to economize on inventory holdings when interest rates are higher. Previous studies (like [Maccini & Rossana \(1984\)](#), [Ramey \(1989\)](#), and [Benati & Lubik \(2014\)](#)) have often struggled to find evidence that aligns with the cost-of-carry logic. This is the context in which [Blinder & Maccini \(1991\)](#) came to state that:

The financial press and business people often state that higher interest rates induce firms to reduce inventories (...). Yet little influence of real interest rates on inventory investment can be found empirically. Why? It is not clear whether the trouble is with the theory or with the empirical tests (...). Whatever the reason, the question of why inventory investment seems to be insensitive to changes in real interest rates remains open, important, and troublesome.

Building on advances in monetary policy shock identification (to the best of our knowledge, we are the first to study this issue with high-frequency shocks), we revisit this question. In particular, we repeat our analysis in Section 4.1 and estimate LPs of the form:

$$\Delta^h INV_{t+h} = \alpha_h + \beta_h MPS_t + \delta_h Z_t + \epsilon_{t,h}, \quad (5)$$

where $\Delta^h INV_{t+h} \equiv INV_{t+h} - INV_{t-1}$ is the cumulative change in the ratio of inventories to sales. In line with our specification in Section 4.1, Z_t controls for 12 lags of the monetary policy shock as well as for 12 lags of INV_{t-1} .

The resulting IRFs suggest that, following a contractionary monetary policy shock, firms *do* economize on their inventory holdings (see Figure 7). This is in line with cost-of-carry logic.¹⁴ In light of the equilibrium forces mentioned in footnote 7 (which are suggesting a bias towards zero), this is a strong finding which further strengthens the case for the relevance of the cost-of-carry channel in monetary policy transmission.




Figure 7: Response of goods inventory-sales ratio to a 25-bp contractionary monetary policy shock, estimated via equation (5). The figure plots $\hat{\beta}_h$. Shaded area represents the 90% confidence interval, calculated via Newey-West standard errors.

5 A New Keynesian Model with Inventories (“NK-inv”)

Motivated by the empirical results presented in Section 4, we now augment the standard New Keynesian model with physical inventories (“NK-inv”) and the cost-of-carry channel of monetary policy (along the lines of the approach taken in Section 3). This extension alters the Phillips curve in a fundamental way as it introduces two inventory-related wedges –

¹⁴This finding may seem at odds with the well-documented fact that the inventory-sales ratio is countercyclical. This stylized fact is however “unconditional” in nature (or conditional on “the average shock” driving the business cycle – which need not be monetary in nature; Angeletos et al. (2020)). Our theory, instead, only speaks to the correlation conditional on monetary policy shocks, which is what our IRFs illustrate.

breaking the Divine Coincidence [Blanchard & Galí \(2007\)](#). One wedge term is related to the stock of inventories, reflecting that increased demand can now be met by drawing down from stock; the other wedge term reflects a direct channel from the real interest rate to inflation.¹⁵ The strength of this direct channel increases with the steady-state level of inventories and gives rise to a sacrifice ratio that varies with inventory levels. In particular, holding the activity gap constant, a monetary contraction in NK-inv lowers inflation by more when inventories are more plentiful. This has important implications for policy, with optimality prescribing that the central bank's focus on inflation stabilization should be increasing in inventory levels.

We now set out the NK-inv model and present the resulting log-linearized equilibrium conditions.

5.1 Firms and inventories

Market structure and demand. Time is discrete, $t = 0, 1, 2, \dots$. A unit mass of monopolistically competitive producers $j \in [0, 1]$ sell directly to a Dixit-Stiglitz final-good aggregator. Final demand for (or sales of) variety j is:

$$S_{j,t} = \left(\frac{P_{j,t}}{P_t} \right)^{-\varepsilon} S_t, \quad \varepsilon > 1, \quad (6)$$

where S_t represents the final aggregator and ε the elasticity of substitution (where we impose $\varepsilon > 1$ to ensure positive marginal revenue). Associated with this problem we have the following price index:

$$P_t = \left(\int_0^1 P_{j,t}^{1-\varepsilon} dj \right)^{\frac{1}{1-\varepsilon}}. \quad (7)$$

Technology and inventory timing. Each firm produces with labor only, according to the technology:

$$Y_{j,t} = A_t L_{j,t}^\alpha, \quad 0 < \alpha < 1, \quad (8)$$

where A_t is the economy-wide level of productivity and $\alpha \in (0, 1)$ is the labor share. Goods produced become available for sale with a one-period delay. Combined with the presence of demand uncertainty, this introduces a role for inventories (as a means to prevent stockouts). Its presence allows sales S_t to draw on prior output. Let $X_{j,t-1}$ be the stock of inventories at the start of t , i.e., inventories carried from the previous period ($t - 1$). The quantity that

¹⁵The first wedge term (related to the stock of inventories) is also obtained in contemporaneous work by [Mehrotra et al. \(2025\)](#). Their model does not include the cost-of-carry channel though and does not explore interactions with monetary policy; instead, they show how the stock wedge is able to improve the NKPC's empirical fit.

firm j has available for sale in period t is:

$$Q_{j,t} = Y_{j,t-1} + X_{j,t-1}. \quad (9)$$

The inventory evolution equation (for the end-of-period t inventory stock) is:

$$X_{j,t} = Q_{j,t} - S_{j,t} = Y_{j,t-1} + X_{j,t-1} - S_{j,t}. \quad (10)$$

At this stage, we anticipate our imminent log-linearization and choose to abstract from the non-negativity constraint on inventories (similar to how the standard New Keynesian model abstracts from the lower bound on the policy rate, which we also do). In a symmetric equilibrium ($P_{j,t} = P_t$, $Y_{j,t} = Y_t$, etc.) equation (10) implies the goods identity, which says that output produced is either sold or added to the inventory stock:

$$Y_{t-1} = S_t + (X_t - X_{t-1}). \quad (11)$$

Profits and costs. Each firm j chooses its price $P_{j,t}$ taking S_t , P_t , and the stochastic discount factor (SDF) M_t as given. Since households own firms, M_t is proportional to the marginal utility of consumption. Let mc_t denote the real marginal cost of production. Assume that firms face a quadratic price adjustment cost à la Rotemberg (1982). Following the literature, this adjustment cost is proportional to the squared deviation of the firm's relative price change and is given by $\frac{\phi}{2}(P_{j,t}/P_{j,t-1} - 1)^2 S_t$. In addition, each firm j incurs a quadratic cost of holding inventories, $X_{j,t}$.¹⁶ The cost of carrying inventories depends on the real rate of interest r_t via the function $\psi(r_t)$. This function can be thought of as describing financial intermediaries, who pass on changes in the policy rate to the relevant interest rates faced by firms; since this sector is not our focus, we do not model it explicitly to keep the model concise. We assume $\psi(\cdot)$ is differentiable with elasticity $\eta_r \equiv \bar{r}\psi'(\bar{r})/\psi(\bar{r})$ at the steady state $(\bar{r}, \bar{X})$.

Firm j 's real profits at date t are therefore given by:

$$\Psi_{j,t} = \frac{P_{j,t}}{P_t} S_{j,t} - mc_t S_{j,t} - \frac{\phi}{2} \left(\frac{P_{j,t}}{P_{j,t-1}} - 1 \right)^2 S_t - \frac{\psi(r_t)}{2} X_{j,t}^2. \quad (12)$$

Pricing condition and Phillips curve in marginal-cost form. Using (6) and $X_{j,t} = Q_{j,t} - S_{j,t}$, the symmetric first-order condition for price setting reads:

$$(\varepsilon - 1) = \varepsilon mc_t - \phi \pi_t (1 + \pi_t) + \beta \phi \mathbb{E}_t \left[\pi_{t+1} (1 + \pi_{t+1}) \frac{S_{t+1}}{S_t} \frac{M_{t+1}}{M_t} \right] - \varepsilon \psi(r_t) X_t, \quad (13)$$

¹⁶Like with prices, the quadratic formulation provides analytical convenience, but can also be thought of as proxying the idea that carrying ever more inventories comes with increasing challenges: one might need to find extra space, or manage available space more carefully. When inventory levels are high, pulling a good from stock can also become more logistically complex and hence time-consuming.

where the main modification is the introduction of the final term, capturing the cost-of-carry channel. After log-linearization around $\bar{\pi} = 0$ and $(\bar{r}, \bar{X})$, we obtain the New Keynesian Phillips Curve modified for the presence of inventories (“NKPC-inv”, see Appendix C for details):

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + (\kappa + \eta_x) \widehat{mc}_t - \eta_x (\hat{x}_t + \eta_r \hat{r}_t), \quad (14)$$

where $\pi_t \equiv P_t/P_{t-1} - 1$ is the rate of inflation, $\widehat{mc}_t$ is the log-deviation of real marginal cost from its steady state, $\hat{x}_t$ is the log-deviation in the end-of-period inventory stock, and $\hat{r}_t$ is the log-deviation of the real-rate. The coefficients in (14) are given by:

$$\kappa \equiv \frac{\varepsilon - 1}{\phi}, \quad \eta_x \equiv \frac{\varepsilon \psi(\bar{r}) \bar{X}}{\phi}, \quad \eta_r \equiv \frac{\bar{r} \psi'(\bar{r})}{\psi(\bar{r})}. \quad (15)$$

Relative to the textbook NKPC, (14) features (i) a steeper slope (by $\eta_x > 0$)¹⁷ and (ii) two inventory-related wedges: a state-related term $-\eta_x \hat{x}_t$ and a direct real-rate channel $-\eta_x \eta_r \hat{r}_t$ (the absence of which is seen as problematic to the standard New Keynesian model; [Rupert & Šustek \(2019\)](#)). Reassuringly, [Den Haan & Sun \(2024\)](#), who introduce a role for inventories via a microfounded “sell friction”, and [Mehrotra et al. \(2025\)](#), who model inventories as directly boosting sales, arrive at a modified NKPC that has similar features. The latter paper also estimates their NKPC with inventories, finding that inclusion of the inventory-sales ratio leads to a much better fit. They find that both the “missing disinflation” of 2009-2011 as well as the Covid-era inflation surge can be understood better when taking note of the low inventory-sales ratio over these periods.

Special cases and interpretation. If $\psi \equiv 0$ or $\bar{X} = 0$, then $\eta_x = 0$ and (14) collapses to the standard NKPC. Otherwise, a higher real rate raises carrying costs, incentivizes stock drawdowns, and reduces current inflation via the $-\eta_x \eta_r \hat{r}_t$ term. A larger buffer stock ($\hat{x}_t > 0$) dampens price pressures through $-\eta_x \hat{x}_t$, which captures the idea that – with access to a greater stock of inventories – higher future demand can be met by drawing down inventories (lowering the need to immediately crank up production, which would raise marginal costs); since the NKPC-inv is forward looking, this lowers inflationary pressures in the present.

Resource constraint and rebate. Rotemberg adjustment costs (AC) and inventory carrying costs (IC) are uses of the final good:

$$AC_t \equiv \frac{\phi}{2} \left(\frac{P_t}{P_{t-1}} - 1 \right)^2 S_t = \frac{\phi}{2} \pi_t^2 S_t, \quad IC_t \equiv \frac{\psi(r_t)}{2} X_t^2. \quad (16)$$

¹⁷This is a relative statement, not necessarily inconsistent with the alleged flatness of the Phillips curve (as for example argued in [Hazell et al. \(2022\)](#)). Within the [Rotemberg \(1982\)](#) setup, a flat NKPC is typically thought of as being driven by high price adjustment costs, ϕ . Since η_x is decreasing in ϕ , a high ϕ could still give rise to a flat slope $(\kappa + \eta_x)$.

We follow a standard rebate scheme with respect to IC : linearize IC_t around $(\bar{r}, \bar{X})$ and define the rebate V_t as the constant-plus-linear Taylor part; consequently, the remainder $\check{IC}_t \equiv IC_t - V_t$ is second order (see Appendix D for the algebra). Goods allocation and the inventory identity are then given by:

$$S_t = C_t + AC_t + \check{IC}_t, \quad Y_{t-1} = S_t + (X_t - X_{t-1}). \quad (17)$$

Around the zero-inflation steady state, $AC_t = \frac{\phi}{2}\pi_t^2 S_t$ is second order and, by construction, so is $\check{IC}_t$. Hence the linearized feasibility conditions are:

$$\hat{s}_t \approx \hat{c}_t, \quad \hat{y}_{t-1} = \hat{c}_t + (\hat{x}_t - \hat{x}_{t-1}), \quad (18)$$

which will justify using the consumption gap in the IS block discussed below. The rebate V_t is financed lump-sum and rebated lump-sum to households, so it does not enter marginal conditions (the Euler equation, the first-order condition for labor supply, or that for price setting).¹⁸

5.2 Households

The household side of the model is standard. Households are expected utility maximizers with a time-separable utility function and a constant discount factor $\beta \in (0, 1)$. Their period utility is defined over aggregate consumption C_t and labor supply L_t . In each period t , the household earns wage income $W_t L_t$ and dividends from firms. The household chooses sequences $\{C_t, L_t\}_{t \geq 0}$ and saves in one-period nominal bonds to maximize:

$$\max_{\{C_t, L_t\}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\frac{C_t^{1-\sigma}}{1-\sigma} - \chi \frac{L_t^{1+\nu}}{1+\nu} \right). \quad (19)$$

Here $\sigma > 0$ is the coefficient of relative risk aversion (with $1/\sigma$ being the elasticity of intertemporal substitution), $\nu \geq 0$ is the inverse Frisch elasticity of labor supply, and $\chi > 0$ scales the disutility of labor.

The first-order condition for a one-period nominal bond gives the usual Euler equation:

$$1 = (1 + i_t) \mathbb{E}_t \left[\beta \left(\frac{C_t}{C_{t+1}} \right)^\sigma \frac{P_t}{P_{t+1}} \right]. \quad (20)$$

¹⁸One can think of this scheme as an approximation to the following. Start by viewing inventory-carrying costs as interest payments flowing to an (unmodelled) financial sector. Since households would own these intermediaries in aggregate, payments are returned to them – netting out at first order. This is exactly what the rebate implements in the linearized resource constraint. For full accounting or welfare calculations, one should include IC_t and V_t explicitly in the government and goods resource accounts, to track second-order losses; for linearized dynamics, V_t and $\check{IC}_t$ drop out and need not be included in the equilibrium conditions. See Appendix D for details.

The intratemporal first-order condition for labor supply is:

$$\frac{W_t}{P_t} = \chi L_t^\nu C_t^\sigma. \quad (21)$$

Log-linearizing (20) around the zero-inflation steady state yields the familiar IS equation:

$$\tilde{c}_t = \mathbb{E}_t \tilde{c}_{t+1} - \frac{1}{\sigma} (i_t - \mathbb{E}_t \pi_{t+1} - r_t^n), \quad (22)$$

where $\tilde{c}_t$ is the consumption gap (or sales gap) in deviations from the natural allocation, i.e., $\tilde{c}_t \equiv \hat{c}_t - \hat{c}_t^n \approx \hat{s}_t - \hat{s}_t^n$ (the approximation holds at first order under the rebate scheme explained above). Natural allocations (i.e., those arising under flexible prices) are denoted by a superscript n , and r_t^n is the natural real rate. Define the real-rate gap as:

$$\tilde{r}_t \equiv (i_t - \mathbb{E}_t \pi_{t+1}) - r_t^n = \sigma (\mathbb{E}_t \tilde{c}_{t+1} - \tilde{c}_t), \quad (23)$$

and the natural real rate by:

$$r_t^n = \bar{r} + \sigma \mathbb{E}_t (\hat{c}_{t+1}^n - \hat{c}_t^n), \quad (24)$$

where $\bar{r} = \frac{1}{\beta} - 1$. These results are standard and we do not derive them here for the sake of brevity.

5.3 Equilibrium conditions

Production. Now, before defining the equilibrium of our NK-inv model, we first need to derive the production Euler equation, which is modified by the presence of inventories. To do so, observe that production decided at $(t-1)$ becomes available for sale at t . This implies that the production decision at $(t-1)$ affects marginal cost at t (via inventories), but the price is set at t . This means that the firm's problem at $(t-1)$ is to choose Y_{t-1} (or, equivalently, the amount of goods for sale, Q_t) taking S_t , P_t , and the SDF M_t as given. Since households own firms, the real SDF between $(t-1)$ and t is $\beta M_t / M_{t-1}$, with $M_t \propto u'(C_t)$. By the envelope theorem (since P_t is chosen optimally at t), we have that $\partial \Pi_t / \partial X_t = -\psi(r_t) X_t$. The intertemporal first-order condition for production is therefore given by:

$$mc_{t-1} = \mathbb{E}_{t-1} \left[\beta \frac{M_t}{M_{t-1}} \psi(r_t) X_t \right]. \quad (25)$$

By log-linearizing the SDF and the carrying-cost around $(\bar{M}, \bar{r}, \bar{X})$, and mapping real marginal cost to gaps via intratemporal optimality and technology (see Appendix F), we obtain:

$$\sigma \tilde{c}_{t-1} + \frac{\nu + 1 - \alpha}{\alpha} \tilde{y}_{t-1} = \psi(\bar{r}) \bar{X} \mathbb{E}_{t-1} [\tilde{x}_t + \eta_r \tilde{r}_t] + \omega_{t-1}, \quad (26)$$

where ω_{t-1} collects SDF wedges (see Appendix F for the expression and discussion; since ω_{t-1} vanishes under certainty equivalence and the standard normalization, we simply drop this term going forward).

NKPC-inv in terms of natural gaps. Next, we rewrite the NKPC-inv. We can use the first-order condition for labor supply and the production function to map the marginal-cost gap to the sales/consumption- and production gaps:

$$\widehat{mc}_t - \widehat{mc}_t^n = \sigma \tilde{c}_t + \frac{\nu + 1 - \alpha}{\alpha} \tilde{y}_t. \quad (27)$$

Substituting this into (14) yields the NKPC-inv featuring the gap terms $\tilde{c}_t$ and $\tilde{y}_t$:

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + (\kappa + \eta_x) \left[\sigma \tilde{c}_t + \frac{\nu + 1 - \alpha}{\alpha} \tilde{y}_t \right] - \eta_x \hat{x}_t - \eta_x \eta_r \hat{r}_t + (\kappa + \eta_x) \widehat{mc}_t^n, \quad (28)$$

where $(\kappa + \eta_x) \widehat{mc}_t^n$ is a standard cost-push shock that depends on exogenous movements in the natural allocation (such as productivity shocks). See Appendix E for details.

Note that the formulation in (28) mixes gaps relative to steady state (indicated by hats above the variable) with gaps relative to the natural allocation (indicated by tildes above the variable, i.e., $\tilde{z}_t \equiv \hat{z}_t - \hat{z}_t^n$). We can rewrite purely in terms of the latter as:

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + (\kappa + \eta_x) \left[\sigma \tilde{c}_t + \frac{\nu + 1 - \alpha}{\alpha} \tilde{y}_t \right] - \eta_x \tilde{x}_t - \eta_x \eta_r \tilde{r}_t + \mu_t, \quad (29)$$

where the term μ_t collects all terms that arise from movements in natural allocations. In the specified model environment, its definition is:

$$\mu_t \equiv (\kappa + \eta_x) \widehat{mc}_t^n - \eta_x \hat{x}_t^n - \eta_x \eta_r \hat{r}_t^n. \quad (30)$$

Given the one period time to sale, end-of-period t inventories are related to the consumption surprise via (see Appendix G for the derivation):

$$\tilde{x}_t - \mathbb{E}_{t-1} \tilde{x}_t = -\frac{\bar{S}}{\bar{X}} (\tilde{c}_t - \mathbb{E}_{t-1} \tilde{c}_t). \quad (31)$$

Let the production Euler equation (26) deliver the predetermined inventory plan:

$$\mathbb{E}_{t-1} \tilde{x}_t = \frac{1}{\psi(\bar{r}) \bar{X}} \left[\sigma \tilde{c}_{t-1} + \frac{\nu + 1 - \alpha}{\alpha} \tilde{y}_{t-1} \right] - \eta_r \mathbb{E}_{t-1} \tilde{r}_t. \quad (32)$$

In Appendix G we show how this can be used to manipulate the NKPC-inv further into:

$$\begin{aligned}\pi_t = & \beta \mathbb{E}_t \pi_{t+1} + (\kappa + \eta_x) \left[\sigma \tilde{c}_t + \frac{\nu + 1 - \alpha}{\alpha} \tilde{y}_t \right] - \eta_x \eta_r (\tilde{r}_t - \mathbb{E}_{t-1} \tilde{r}_t) \\ & - \frac{\eta_x}{\psi(\bar{r}) \bar{X}} \left[\sigma \tilde{c}_{t-1} + \frac{\nu + 1 - \alpha}{\alpha} \tilde{y}_{t-1} \right] + \eta_x \frac{\bar{S}}{\bar{X}} (\tilde{c}_t - \mathbb{E}_{t-1} \tilde{c}_t) + \mu_t.\end{aligned}\quad (33)$$

A key implication is that the “direct” impact of an unanticipated change in the real rate is to lower inflation:

$$\frac{\partial \pi_t}{\partial (\tilde{r}_t - \mathbb{E}_{t-1} \tilde{r}_t)} = -\eta_x \eta_r < 0.$$

This confirms the presence of the cost-of-carry channel in the NKPC-inv.

Equilibrium. At this stage, we can define the equilibrium for the NK-inv model as follows:

Definition 1 (Rational expectations equilibrium for NK-inv). *Fix parameters $\beta, \sigma, \nu, \alpha, \varepsilon, \phi$, the carry-cost function $\psi(\cdot)$ evaluated at $\bar{r}$, and the steady-state inventory level $\bar{X}$. Define:*

$$\kappa \equiv \frac{\varepsilon - 1}{\phi}, \quad \eta_x \equiv \frac{\varepsilon \psi(\bar{r}) \bar{X}}{\phi}, \quad \eta_r \equiv \frac{\bar{r} \psi'(\bar{r})}{\psi(\bar{r})}.$$

Given exogenous sequences $\{r_t^n\}_{t \geq 0}$ and $\{\mu_t\}_{t \geq 0}$, an initial condition $(\tilde{x}_{-1}, \tilde{y}_{-1})$, a (linear) rational expectations equilibrium is a collection $\{\pi_t, \tilde{c}_t, \tilde{y}_t, \tilde{x}_t, i_t\}_{t \geq 0}$ such that, for all t , the equilibrium satisfies the following conditions:

(i) IS equation:

$$\tilde{c}_t = \mathbb{E}_t \tilde{c}_{t+1} - \frac{1}{\sigma} (i_t - \mathbb{E}_t \pi_{t+1} - r_t^n).$$

(ii) NKPC-inv:

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + (\kappa + \eta_x) \left[\sigma \tilde{c}_t + \frac{\nu + 1 - \alpha}{\alpha} \tilde{y}_t \right] - \eta_x \tilde{x}_t - \eta_x \eta_r \tilde{r}_t + \mu_t.$$

(iii) Inventory accounting:

$$\tilde{x}_t = \tilde{x}_{t-1} + \tilde{y}_{t-1} - \tilde{c}_t.$$

(iv) Production Euler equation:

$$\sigma \tilde{c}_{t-1} + \frac{\nu + 1 - \alpha}{\alpha} \tilde{y}_{t-1} = \psi(\bar{r}) \bar{X} \mathbb{E}_{t-1} [\tilde{x}_t + \eta_r \tilde{r}_t].$$

(v) Monetary policy rule:

$$i_t = \phi_\pi \pi_t + \phi_c \tilde{c}_t.$$

The real-rate gap is given by:

$$\tilde{r}_t \equiv i_t - \mathbb{E}_t \pi_{t+1} - r_t^n,$$

and expectations $\mathbb{E}_t[\cdot]$ are taken conditional on information available at time t .

5.4 Optimal monetary policy

We now study optimal monetary policy in the NK-inv model – looking at the optimal simple rule, as well as optimal policy under both commitment and discretion. In the baseline New Keynesian model, stabilizing the output gap also stabilizes inflation (the Divine Coincidence; [Blanchard & Galí \(2007\)](#)), unless an ad hoc cost-push term is appended to the Phillips curve (also see [Ravenna & Walsh \(2006\)](#)). Our NK-inv breaks the Divine Coincidence endogenously, with the NKPC-inv featuring a wedge related to the inventory carrying cost, $-n_x \eta_r \tilde{r}_t$, and an inventory wedge, $-n_x \tilde{x}_t$. Consequently, strict inflation targeting will not succeed in stabilizing real activity – making it optimal to pay more direct attention to the latter.

In our analysis of optimal policy, we start from the following “dual mandate” period loss function:

$$L_t = \pi_t^2 + \lambda_c \tilde{c}_t^2, \quad (34)$$

where $\lambda_c \geq 0$ captures the relative weight the central bank places on stabilization of the consumption gap. We study optimal policy in response to innovations in the cost-push wedge term μ_t , which we assume to evolve according to $\mu_t = \rho_\mu \mu_{t-1} + \epsilon_{\mu,t}$ (where we set $\rho_\mu = 0.9$ and have $\epsilon_{\mu,t} \stackrel{iid}{\sim} N(0, 0.01)$).

As shown in [Figure 8](#), for a given value of λ_c (this section uses $\lambda_c = 0.10$ but its precise value is unimportant to the broader point), optimal policy is more strongly focused on inflation stabilization when the cost-of-carry channel is at play.¹⁹ Remember that, in our model, the strength of this channel is governed by the steady-state level of inventories, $\bar{X}$. [Figure 8](#) therefore considers a low value of $\bar{X} = 0.01$ (approximating the standard NK model, which has $\bar{X} = 0$), an intermediate value of $\bar{X} = 1.4$ (our calibration has $\bar{S} = 1$, meaning that $\bar{X}$ also represents the steady-state ratio of inventories-to-sales; this ratio has historically averaged around 1.4 in US data), and a high value of $\bar{X} = 2.5$. Both under discretion (left panel) and under commitment (right panel) we see that the optimal policy works more heavily to keep inflation closer to target when $\bar{X}$ is higher (i.e., when the cost-of-carry channel is

¹⁹In this exercise, we use a relatively standard calibration for the traditional parameters. We set the discount factor to $\beta = 0.99$ and the curvature of the utility function to $\sigma = 1$. The labor share is normalized to $\alpha = 1$, and the inverse of the Frisch elasticity of labor supply is set to $\nu = 2$, as recommended by [Chetty et al. \(2011\)](#) for the intensive margin. Following the literature, we choose an elasticity of substitution across varieties of $\varepsilon = 6$ and a price adjustment cost of $\phi = 50$ (see [Bilbiie & Ragot \(2021\)](#) for a discussion and references). For the inventory-related parameters, we set $\bar{X} = 1.4$, corresponding to a steady-state inventory-to-sales ratio consistent with the historical U.S. average. Finally, we set $\bar{\psi} = 0.25$ and $\eta_r = 0.50$, where η_r governs the elasticity of the cost of carrying inventories with respect to the real interest rate.

stronger).




Figure 8: Optimal policy results for a shock to μ_t : Discretion (left) vs. Commitment (right).

Figure 9 shows that the same insight follows from an optimal simple rule exercise, where we search for the coefficients (ϕ_π, ϕ_c) that minimize the present discounted value of the central bank's loss (34): the greater the importance of the cost-of-carry channel (as governed by $\bar{X}$), the more the central bank should focus on inflation stabilization (relative to stabilization of the consumption gap).

The intuition for this result lies in the fact that the introduction of the cost-of-carry channel gives rise to a more favorable sacrifice ratio for the central bank. Recall that the introduction of inventories modifies the standard NKPC by (i) steepening its slope (by $\eta_x > 0$) and (ii) introducing a direct real rate channel (cf. the $-\eta_x \eta_r \tilde{r}_t$ -term in equation (29)). Both make it “cheaper” to lower inflation – in the sense of requiring less slack (instead, it is the price-cost markup desired by firms which compresses following a monetary contraction).²⁰

²⁰See Nekarda & Ramey (2013) for empirical evidence pointing to price-cost markups being procyclical in

This makes it optimal for the central bank to put more weight on inflation stabilization when faced with trade-off inducing shocks.²¹




Figure 9: Optimal Simple Rule as a Function of Steady-State Inventories.

6 Conclusion

Theory suggests that interest rates might affect firms’ inventory management and pricing strategies. According to this logic, higher interest rates give firms – in particular those carrying more inventories – a stronger incentive to cut prices in an attempt to cut back on their inventory holdings (as carrying inventories is more expensive when interest rates are higher). Testing this hypothesis on data from the U.S. goods, housing, and oil markets, we indeed find that higher inventory levels amplify the disinflationary effects of tighter monetary policy – suggesting that there is a “cost-of-carry channel” of monetary policy transmission at play. As a result, monetary policy may have more leverage over inflation in high-inventory environments (where a given interest rate increase can be expected to lower inflation by more, without this requiring slack to open up). Extending a standard New Keynesian model

response to monetary policy shocks. They go on to note how this is inconsistent with the prediction of the standard New Keynesian model. Our analysis suggests that part of the solution might lie in accounting for firms’ inventory holdings and associated dynamics. Also see [Van Der Ploeg & Willems \(2025\)](#), who provide a general analysis of optimal monetary policy when markups are cyclically sensitive.

²¹The intuition can also be understood by making a reduced-form change to the standard NKPC (making it more “NKPC-inv like”). In particular, consider a steeper slope ($\check{\kappa} > \kappa$) and introduce a direct real rate channel, with its strength governed by $\vartheta > 0$: $\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \check{\kappa} \tilde{c}_t - \vartheta r_t$. By following the standard steps to solve for optimal policy under discretion, one then finds that the optimal targeting criterion is $\pi_t = -\frac{\lambda_c}{\check{\kappa} + \sigma \vartheta} \tilde{c}_t$. Since $\check{\kappa} + \sigma \vartheta > \kappa$ (the latter being the denominator for the standard NK model), it follows that optimality calls for greater inflation stabilization under the modified NKPC.

with inventories and the cost-of-carry channel, we show that this modification makes optimal policy more focused on inflation stabilization when inventories are more plentiful – the reason being that the central bank faces a more favorable sacrifice ratio in such an environment.

Finally, this paper leaves several issues for future work. If one accepts the empirical findings of this paper, further steps on the modeling front might be desirable. Given our focus on the implications for optimal policy, we made the conscious choice to keep the model simple and abstract from formally introducing financial intermediaries (who pass on changes in the policy rate to the interest rates faced by inventory carrying firms). However, incorporation of such a sector could be an interesting extension. In addition, our analysis has followed the literature by focusing on the log-linearized equilibrium, but given the volatility in inventories – along with its non-negativity constraint – analyzing the original non-linear problem might have considerable value too. We hope that future work will be able to advance along these, and other, dimensions.

References

- Akhtar, M. A. (1983), ‘Effects of interest rates and inflation on aggregate inventory investment in the united states’, *American Economic Review* **73**, 319–328.
- Alessandria, G., Kaboski, J. P. & Midrigan, V. (2010), ‘Inventories, lumpy trade, and large devaluations’, *American Economic Review* **100**(5), 2304–2339.
- Angeletos, G.-M., Collard, F. & Dellas, H. (2020), ‘Business-cycle anatomy’, *American Economic Review* **110**(10), 3030–3070.
- Baqae, D. & Rubbo, E. (2023), ‘Micro propagation and macro aggregation’, *Annual Review of Economics* **15**(1), 91–123.
- Bauer, M. D. & Swanson, E. T. (2023), ‘A reassessment of monetary policy surprises and high-frequency identification’, *NBER Macroeconomics Annual* **37**.
- Benati, L. & Lubik, T. A. (2014), ‘Sales, inventories and real interest rates: a century of stylized facts’, *Journal of Applied Econometrics* **29**(7), 1210–1222.
- Bilbiie, F. O. & Ragot, X. (2021), ‘Optimal monetary policy and liquidity with heterogeneous households’, *Review of Economic Dynamics* **41**, 71–95.
- Bils, M. & Kahn, J. A. (2000), ‘What inventory behavior tells us about business cycles’, *American Economic Review* **90**(3), 458–481.
- Blanchard, O. & Galí, J. (2007), ‘Real wage rigidities and the new keynesian model’, *Journal of money, credit and banking* **39**, 35–65.
- Blinder, A. S. (1981), ‘Retail inventory behavior and business fluctuations’, *Brookings Papers on Economic Activity* **2**, 443–505.
- Blinder, A. S. (1986), ‘Can the production smoothing model of inventory behavior be saved?’, *Quarterly Journal of Economics* **101**, 431–453.
- Blinder, A. S. & Maccini, L. J. (1991), ‘Taking stock: A critical assessment of recent research on inventories’, *Journal of Economic Perspectives* **5**, 73–96.
- Boivin, J., Kiley, M. T. & Mishkin, F. S. (2010), How has the monetary transmission mechanism evolved over time?, in ‘Handbook of Macroeconomics’, Vol. 3, pp. 369–422.
- CBS (2022), ‘Target’s profit craters after it cut prices to clear inventory’, <https://www.cbsnews.com/news/targets-profit-sinks-after-it-cut-prices-to-clear-inventory>.

- Chetty, R., Guren, A., Manoli, D. & Weber, A. (2011), ‘Are micro and macro labor supply elasticities consistent? a review of evidence on the intensive and extensive margins’, *American Economic Review* **101**(3), 471–475.
- Conner, A., Campbell, S., Sheiner, L. & Wessel, D. (2024), ‘How does the consumer price index account for the cost of housing?’, *Brookings Commentary*.
- Copeland, A., Hall, G. & Maccini, L. J. (2019), ‘Interest rates and the market for new light vehicles’, *Journal of Money, Credit and Banking* **51**(5), 1137–1168.
- Cotton, C. D. (2024), The pass-through of gaps between market rent and the price of shelter, Technical Report 24-6, Federal Reserve Bank of Boston Research Department.
- Deaton, A. & Laroque, G. (1992), ‘On the behaviour of commodity prices’, *Review of Economic Studies* **59**, 1–23.
- Deaton, A. & Laroque, G. (1995), ‘Estimating a nonlinear rational expectations commodity price model with unobservable state variables’, *Journal of Applied Econometrics* **10**(S1), S9–S40.
- Deaton, A. & Laroque, G. (1996), ‘Competitive storage and commodity price dynamics’, *Journal of Political Economy* **104**, 896–923.
- Den Haan, W. J. & Sun, T. (2024), The role of sell frictions for inventories and business cycles, Technical Report 26, CFM Discussion Paper.
- Dominguez, L. (2023), ‘How unilever is tackling excess inventory with pricing intelligence’, <https://consumergoods.com/how-unilever-tackling-excess-inventory-pricing-intelligence>.
- Eichenbaum, M. S. (1989), ‘Some empirical evidence on the production level and production cost smoothing models of inventory investment’, *American Economic Review* **79**, 853–864.
- Fitzgerald, T. J. (1997), ‘Inventories and the business cycle: an overview’, *Federal Reserve Bank of Cleveland Economic Review* **33**(3).
- Frankel, J. A. (2008a), ‘An explanation for soaring commodity prices’, <https://voxeu.org/article/explanation-soaring-commodity-prices>.
- Frankel, J. A. (2008b), ‘Monetary policy and commodity prices’, <https://voxeu.org/article/monetary-policy-and-commodity-prices>.
- Frankel, J. A. (2014), ‘Effects of speculation and interest rates in a carry trade model of commodity prices’, *Journal of International Money and Finance* **42**, 88–112.
- Gertler, M. & Gilchrist, S. (1994), ‘Monetary policy, business cycles, and the behavior of small manufacturing firms’, *Quarterly Journal of Economics* **109**, 309–340.

- Gürkaynak, R., Karasoy-Can, H. G. & Lee, S. S. (2022), ‘Stock market’s assessment of monetary policy transmission: The cash flow effect’, *The Journal of Finance* **77**(4), 2375–2421.
- Hazell, J., Herreno, J., Nakamura, E. & Steinsson, J. (2022), ‘The slope of the phillips curve: evidence from us states’, *The Quarterly Journal of Economics* **137**(3), 1299–1344.
- Irvine, F. O. (1981), ‘Retail inventory investment and the cost of capital’, *American Economic Review* **71**, 633–648.
- Jacobson, M. M., Matthes, C. & Walker, T. B. (2023), ‘Temporal aggregation bias and monetary policy transmission’.
- Kahn, J. A. (1987), ‘Inventories and the volatility of production’, *American Economic Review* **77**, 667–679.
- Känzig, D. R. (2021), ‘The macroeconomic effects of oil supply news: Evidence from opec announcements’, *American Economic Review* **111**(4), 1092–1125.
- Kashyap, A. K., Lamont, O. A. & Stein, J. C. (1994), ‘Credit conditions and the cyclical behavior of inventories’, *Quarterly Journal of Economics* **109**, 565–592.
- Kashyap, A. K., Stein, J. C. & Wilcox, D. W. (1993), ‘Monetary policy and credit conditions: evidence from the composition of external finance’, *American Economic Review* **83**, 78–98.
- Khan, A. & Thomas, J. K. (2007), ‘Inventories and the business cycle: An equilibrium analysis of (s, s) policies’, *American Economic Review* **97**(4), 1165–1188.
- Kilian, L. & Murphy, D. P. (2014), ‘The role of inventories and speculative trading in the global market for crude oil’, *Journal of Applied econometrics* **29**(3), 454–478.
- Kim, R. (2021), ‘The effect of the credit crunch on output price dynamics: The corporate inventory and liquidity management channel’, *The Quarterly Journal of Economics* **136**(1), 563–619.
- Kryvtsov, O. & Midrigan, V. (2013), ‘Inventories, markups, and real rigidities in menu cost models’, *Review of Economic Studies* **80**(1), 249–276.
- Lieberman, C. (1980), ‘Inventory demand and cost of capital effects’, *Review of Economics and Statistics* **62**, 348–356.
- Maccini, L. J., Moore, B. J. & Schaller, H. (2004), ‘The interest rate, learning, and inventory investment’, *American Economic Review* **94**, 1303–1327.
- Maccini, L. J. & Rossana, R. J. (1984), ‘Joint production, quasi-fixed factors of production, and investment in finished goods inventories’, *Journal of Money, Credit and Banking* **16**(2), 218–236.

- McMahon, M. F. (2012), ‘Inventories in motion: A new approach to inventories over the business cycle’.
- Mehrotra, N., Oh, H. & Ortiz, J. L. (2025), ‘Retail inventories and inflation dynamics: The price margin channel’.
- Miranda-Pinto, J., Pescatori, A., Prifti, E. & Verduzco-Bustos, G. (2023), Monetary policy transmission through commodity prices, Technical Report 2023/215, IMF Working Paper.
- Miranda-Pinto, J., Pescatori, A., Prifti, E. & Verduzco-Bustos, G. (2024), ‘The commodity transmission channel of monetary policy and inflation dynamics’, <https://cepr.org/voxeu/columns/commodity-transmission-channel-monetary-policy-and-inflation-dynamics>.
- Mrázová, M. & Neary, J. P. (2017), ‘Not so demanding: Demand structure and firm behavior’, *American Economic Review* **107**(12), 3835–3874.
- Nekarda, C. J. & Ramey, V. A. (2013), The cyclical behavior of the price-cost markup, Technical report, National Bureau of Economic Research.
- Petersen, P. B. (2002), ‘The misplaced origin of just-in-time production methods’, *Management Decision* **40**, 82–88.
- Ramey, V. A. (1989), ‘Inventories as factors of production and economic fluctuations’, *American Economic Review* **79**, 338–354.
- Ramey, V. A. (1991), ‘Nonconvex costs and the behavior of inventories’, *Journal of Political Economy* **99**, 306–334.
- Ramey, V. A. & West, K. D. (1999), Inventories, in ‘Handbook of Macroeconomics’, Vol. 1, pp. 863–923.
- Ravenna, F. & Walsh, C. E. (2006), ‘Optimal monetary policy with the cost channel’, *Journal of Monetary Economics* **53**(2), 199–216.
- Ravn, M. O. & Uhlig, H. (2002), ‘On adjusting the hodrick-prescott filter for the frequency of observations’, *Review of economics and statistics* **84**(2), 371–376.
- Reuters (2022), ‘U.s. retailers’ ballooning inventories set stage for deep discounts’, <https://www.reuters.com/markets/us/us-retailers-ballooning-inventories-set-stage-deep-discounts-2022-05-27/>.
- Robinson, M. (2022), ‘Interest rates could spell trouble for inventories, liquidity, and ipos’, <https://www.investorschronicle.co.uk/news/2022/10/06/interest-rates-could-spell-trouble-for-inventories-liquidity-and-ipos>.

- Rotemberg, J. J. (1982), ‘Sticky prices in the united states’, *Journal of political economy* **90**(6), 1187–1211.
- Rupert, P. & Šustek, R. (2019), ‘On the mechanics of new-keynesian models’, *Journal of Monetary Economics* **102**, 53–69.
- Trudell, C. (2024), ‘Tesla offers steep discounts on suvs piling up in inventory’, <https://www.bloomberg.com/news/articles/2024-04-05/tesla-aims-discounts-at-unprecedented-number-of-evs-in-inventory>.
- Van Der Ploeg, F. & Willems, T. (2025), ‘Battle of the markups: conflict inflation and the aspirational channel of monetary policy transmission’.
- Wen, Y. (2011), ‘Input and output inventory dynamics’, *American Economic Journal: Macroeconomics* **3**(4), 181–212.
- Williamson, C. (2023), ‘Price pressures alleviated by falling demand, fewer supply delays and inventory reduction policies’, *S&P Global Economics Commentary* .

Appendix

A Robustness

Our main finding, that prices are more sensitive (in the conventional direction) to monetary policy shocks when inventory levels are higher, is very robust. Here, we document some of the robustness exercises we have conducted.

First, our finding on aggregate goods prices (shown in Figure 2) is robust to adding the rate of unemployment alongside the 5-year Treasury yield to the controls vector Z_t in equation (4); see Figure 10. This suggests that the effect we are picking up is not driven by variations in the state of the business cycle. The same holds when adding the same controls to the regressions for the price of housing services, as well as that for the oil price; see Figures 11 and 12, respectively.




Figure 10: Additional response of PCE-goods price index to a 25-bp contractionary monetary policy shock, due to a unit increase in the inventory-sales ratio, estimated via equation (4), with the rate of unemployment and the 5-year Treasury yield added to the vector of controls (Z_t). The figure plots $\hat{\gamma}_h$. Shaded area represents the 90% confidence interval, calculated via Newey-West standard errors.

Further robustness checks are warranted with respect to our result for housing services, where different price series are available. Replicating our baseline analysis when the dependent variable is the housing-related component of the PCE index produces an even stronger result (in the sense of being larger in magnitude and more persistent; see Figure 13). At the same time, one can even see a significant effect coming from using the home vacancy rate as inventory proxy “*INV*” when having *overall* CPI as dependent variable (Figure 14), which is perhaps not that surprising given OER accounts for about one-quarter of the overall CPI.




Figure 11: Additional response of CPI OER to a 25-bp contractionary monetary policy shock, due to a 1-pp increase in the home vacancy rate, estimated via equation (4), with the rate of unemployment and the 5-year Treasury yield added to the vector of controls (Z_t). The figure plots $\hat{\gamma}_h$. Shaded area represents the 90% confidence interval, calculated via Newey-West standard errors.




Figure 12: Additional response of oil prices to a 25-bp contractionary monetary policy shock, due to a 1-pp increase in oil inventory, estimated via equation (4), with the rate of unemployment and the 5-year Treasury yield added to the vector of controls (Z_t). The figure plots $\hat{\gamma}_h$. Shaded area represents the 90% confidence interval, calculated via Newey-West standard errors.




Figure 13: Additional response of the housing component of the PCE index to a 25-bp contractionary monetary policy shock, due to a 1-pp increase in the home vacancy rate, estimated via equation (4). The figure plots $\hat{\gamma}_h$. Shaded area represents the 90% confidence interval, calculated via Newey-West standard errors.




Figure 14: Additional response of overall CPI to a 25-bp contractionary monetary policy shock, due to a 1-pp increase in the home vacancy rate, estimated via equation (4). The figure plots $\hat{\gamma}_h$. Shaded area represents the 90% confidence interval, calculated via Newey-West standard errors.

Finally, with respect to the core of our result for oil prices, Figure 15 shows that performing linear detrending produces a result which is similar to that displayed in the main text (which was obtained after applying the HP-filter to the oil inventory series).




Figure 15: Additional response of oil prices to a 25-bp contractionary monetary policy shock, due to a 1-pp increase in oil inventory, estimated via equation (4), when applying linear detrending. The figure plots $\hat{\gamma}_h$. Shaded area represents the 90% confidence interval, calculated via Newey-West standard errors.

B Constructing a monthly housing inventory metric

Our empirical exercise in Section 4 investigates whether the response of the cost of housing (to monetary policy shocks) differs depending on the size of the housing inventory “ INV_t ”, which is the fraction of homes that is not being occupied. The US Census Bureau produces two series which are getting at this concept: one for rental properties (FRED code: RRVUSQ156N) and one for owner-occupied properties (FRED code: RHVRUSQ156N). These two series are highly correlated (with a correlation coefficient of 0.74), which is intuitive. We proceed by combining these two series into a single “home vacancy rate”, which is constructed as a weighted-average between the two – with the weight determined by the homeownership rate (FRED code: RSAHORUSQ156S). Finally, since the original series are only available at the quarterly frequency, we use linear interpolation to obtain a monthly series. Given the high degree of persistence in the quarterly series (an autocorrelation coefficient of 0.96 at the quarterly frequency), this is unlikely to be a major issue.

C Price setting and NKPC-inv derivation

A firm j chooses $P_{j,t}$ taking S_t, P_t, M_t as given (households own firms, so $M_t \propto u'(C_t)$). Real profits (in units of the final good) are:

$$\Psi_{j,t} = \frac{P_{j,t}}{P_t} S_{j,t} - mc_t S_{j,t} - \frac{\phi}{2} \left(\frac{P_{j,t}}{P_{j,t-1}} - 1 \right)^2 S_t - \frac{\psi(r_t)}{2} X_{j,t}^2, \quad (35)$$

Using (6) and $X_{j,t} = Q_{j,t} - S_{j,t}$, and letting $p_t \equiv P_{j,t}/P_t$, the static derivative of $\Psi_{j,t}$ with respect to p_t (ignoring Rotemberg terms temporarily) yields, in a symmetric equilibrium where $p_t = 1$,

$$(\varepsilon - 1) = \varepsilon mc_t - \varepsilon \psi(r_t) X_t. \quad (36)$$

Including Rotemberg costs—whose current-period derivative contributes $-\phi\pi_t(1 + \pi_t)$ and whose discounted next-period derivative contributes $+\beta\phi\mathbb{E}_t \left[\pi_{t+1}(1 + \pi_{t+1}) \frac{S_{t+1}}{S_t} \frac{M_{t+1}}{M_t} \right]$ —we obtain the nonlinear symmetric pricing first-order condition:

$$(\varepsilon - 1) = \varepsilon mc_t - \phi\pi_t(1 + \pi_t) + \beta\phi\mathbb{E}_t \left[\pi_{t+1}(1 + \pi_{t+1}) \frac{S_{t+1}}{S_t} \frac{M_{t+1}}{M_t} \right] - \varepsilon \psi(r_t) X_t. \quad (37)$$

Steady state. At $\bar{\pi} = 0$, $S_{t+1}/S_t = 1$, and $M_{t+1}/M_t = 1/\beta$, (37) implies:

$$\bar{mc} = \frac{\varepsilon - 1}{\varepsilon} + \psi(\bar{r}) \bar{X}. \quad (38)$$

Define log-deviations from steady state $\hat{z}_t \equiv \ln Z_t - \ln \bar{Z}$ and the elasticity $\eta_r \equiv \bar{r}\psi'(\bar{r})/\psi(\bar{r})$. Linear approximations around (38) are given by:

$$-\phi\pi_t(1 + \pi_t) \approx -\phi\pi_t, \quad \beta\phi\mathbb{E}_t[\pi_{t+1}(1 + \pi_{t+1}) \cdots] \approx \beta\phi\mathbb{E}_t\pi_{t+1}. \quad (39)$$

$$\psi(r_t) X_t \approx \psi(\bar{r}) \bar{X} [1 + \hat{x}_t + \eta_r \hat{r}_t]. \quad (40)$$

Collect terms, subtract the steady-state identity implied by (38) and define:

$$\kappa \equiv \frac{\varepsilon - 1}{\phi}, \quad \eta_x \equiv \frac{\varepsilon \psi(\bar{r}) \bar{X}}{\phi}, \quad (41)$$

to obtain the log-linearized NKPC with inventories (“NKPC-inv”):

$$\pi_t = \beta\mathbb{E}_t\pi_{t+1} + (\kappa + \eta_x) \widehat{mc}_t - \eta_x (\hat{x}_t + \eta_r \hat{r}_t). \quad (42)$$

D Resource constraints and rebate scheme

Both Rotemberg adjustment costs (AC) and inventory carrying costs (IC) are uses of the final good:

$$AC_t = \frac{\phi}{2} \left(\frac{P_t}{P_{t-1}} - 1 \right)^2 S_t = \frac{\phi}{2} \pi_t^2 S_t, \quad (43)$$

$$IC_t = \frac{\psi(r_t)}{2} X_t^2. \quad (44)$$

We will now linearize IC_t around the steady state $(\bar{r}, \bar{X})$. Define the carrying cost elasticity and log deviations:

$$\eta_r \equiv \bar{r} \frac{\psi'(\bar{r})}{\psi(\bar{r})}, \quad \hat{r}_t \equiv \ln r_t - \ln \bar{r}, \quad \hat{x}_t \equiv \ln X_t - \ln \bar{X}. \quad (45)$$

The following steps lead to the rebate scheme, which is a modeling trick to ensure that inventory carrying costs do not create first-order wedges in the resource constraint after linearization. The steps are:

Step 1. First order expansions around $(\bar{r}, \bar{X})$. Carrying cost schedule:

$$\psi(r_t) \approx \psi(\bar{r}) (1 + \eta_r \hat{r}_t). \quad (46)$$

Square of inventories:

$$X_t^2 = (\bar{X} e^{\hat{x}_t})^2 \approx \bar{X}^2 (1 + 2 \hat{x}_t). \quad (47)$$

Step 2. First order approximation of IC_t . Multiply (46) and (47) and drop the cross term $\eta_r \hat{r}_t \hat{x}_t$ which is second order:

$$IC_t \approx \frac{\psi(\bar{r})}{2} \bar{X}^2 (1 + 2 \hat{x}_t + \eta_r \hat{r}_t) + \text{second order terms}. \quad (48)$$

Step 3. Rebate definition and net waste. Define the rebate as the constant plus linear Taylor part of IC_t :

$$V_t \equiv \frac{\psi(\bar{r})}{2} \bar{X}^2 (1 + 2 \hat{x}_t + \eta_r \hat{r}_t). \quad (49)$$

$$\check{IC}_t \equiv IC_t - V_t = \text{second order}. \quad (50)$$

Step 4. Final uses and physical identity. Goods allocation and the inventory identity in levels are:

$$S_t = C_t + AC_t + \check{I}C_t, \quad Y_{t-1} = S_t + (X_t - X_{t-1}). \quad (51)$$

Step 5. Linearized feasibility. Around a zero inflation steady state we have $\pi = 0$ so $AC_t = \frac{\phi}{2}\pi_t^2 S_t$ is second order and contributes zero to first order; $\check{I}C_t$ is second order by construction. Therefore, the linearized feasibility conditions are:

$$\hat{s}_t \approx \hat{c}_t, \quad \hat{y}_{t-1} = \hat{c}_t + (\hat{x}_t - \hat{x}_{t-1}). \quad (52)$$

The rebate V_t is financed with lump sum taxes and rebated lump sum to households. It cancels the constant and linear parts of IC_t so it does not enter marginal conditions. The Euler equation and the labor supply condition are unchanged. Firms take V_t as given, so the pricing first order condition keeps the inventory term $\psi(r_t)X_t$.

E Mapping marginal cost to the gap

From (21) and the technology (8), real marginal cost is:

$$mc_t = \frac{W_t/P_t}{\partial Y_t/\partial L_t} = \frac{\chi L_t^\nu C_t^\sigma}{\alpha A_t L_t^{\alpha-1}} = \frac{\chi}{\alpha} C_t^\sigma A_t^{-1} L_t^{\nu-\alpha+1}. \quad (53)$$

Using $Y_t = A_t L_t^\alpha \Rightarrow \ln L_t = (\ln Y_t - \ln A_t)/\alpha$, log-linearization gives:

$$\widehat{mc}_t = \sigma \hat{c}_t + \frac{\nu + 1 - \alpha}{\alpha} \hat{y}_t - \frac{\nu + 1}{\alpha} \hat{a}_t. \quad (54)$$

Let the natural allocation (flexible prices, same A_t) be denoted by a superscript n . Define the *sales/consumption gap* and the *output gap*:

$$\tilde{c}_t \equiv \hat{c}_t - \hat{c}_t^n \approx \hat{s}_t - \hat{s}_t^n, \quad \tilde{y}_t \equiv \hat{y}_t - \hat{y}_t^n. \quad (55)$$

Subtracting the flex-price counterpart of (54) (same A_t) yields the exact mapping in two gaps, $\tilde{c}_t$ and $\tilde{y}_t$:

$$\widehat{mc}_t - \widehat{mc}_t^n = \sigma \tilde{c}_t + \frac{\nu + 1 - \alpha}{\alpha} \tilde{y}_t. \quad (56)$$

Therefore, substituting into (14) yields:

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + (\kappa + \eta_x) \left[\sigma \tilde{c}_t + \frac{\nu + 1 - \alpha}{\alpha} \tilde{y}_t \right] - \eta_x \hat{x}_t - \eta_x \eta_r \hat{r}_t + (\kappa + \eta_x) \widehat{mc}_t^n, \quad (57)$$

where $(\kappa + \eta_x) \widehat{mc}_t^n$ is a standard cost-push shock that depends on exogenous natural-

allocation movements (e.g., productivity).

F Production Euler equation with inventories

Timing and inventory accounting. Production chosen at $t - 1$ becomes available for sale at t :

$$Q_t = Y_{t-1} + X_{t-1}, \quad X_t = Q_t - S_t = Y_{t-1} + X_{t-1} - S_t. \quad (58)$$

FOC for Y_{t-1} . At $t - 1$, producing one more unit raises X_t one-for-one at t as we can see in (58). By knowing the households own the firms; the real SDF between $t - 1$ and t is $\beta M_t/M_{t-1}$ with $M_t \propto u'(C_t)$. Using the envelope theorem (price P_t is chosen optimally at t), the marginal value of an extra unit of X_t at t is $\partial \Pi_t / \partial X_t = -\psi(r_t) X_t$. Thus the intertemporal first-order condition for production reads:

$$mc_{t-1} = \mathbb{E}_{t-1} \left[\beta \frac{M_t}{M_{t-1}} \psi(r_t) X_t \right]. \quad (59)$$

Linearization and mapping to gaps. Standard log-linearization around the steady state $(\bar{M}, \bar{r}, \bar{X})$ for the SDF block yields:

$$\beta \frac{M_t}{M_{t-1}} \approx 1 + m_t, \quad m_t \equiv -\sigma(\hat{c}_t - \hat{c}_{t-1}). \quad (60)$$

For the carrying-cost block we get:

$$\psi(r_t) \approx \psi(\bar{r})(1 + \eta_r \hat{r}_t), \quad X_t \approx \bar{X}(1 + \hat{x}_t), \quad \eta_r \equiv \frac{\bar{r} \psi'(\bar{r})}{\psi(\bar{r})}. \quad (61)$$

Multiplying (60) and (61) and dropping cross products gives:

$$\mathbb{E}_{t-1} \left[\beta \frac{M_t}{M_{t-1}} \psi(r_t) X_t \right] \approx \psi(\bar{r}) \bar{X} \mathbb{E}_{t-1} [1 + \hat{x}_t + \eta_r \hat{r}_t + m_t]. \quad (62)$$

The left-hand side maps to gaps via the labor-supply/technology block:

$$\widehat{mc}_{t-1} - \widehat{mc}_{t-1}^n = \sigma \tilde{c}_{t-1} + \frac{\nu + 1 - \alpha}{\alpha} \tilde{y}_{t-1}. \quad (63)$$

Subtracting the natural counterpart of (59) and using (62)-(63) yields the *log-linearized production Euler equation*:

$$\sigma \tilde{c}_{t-1} + \frac{\nu + 1 - \alpha}{\alpha} \tilde{y}_{t-1} = \psi(\bar{r}) \bar{X} \mathbb{E}_{t-1} [\hat{x}_t + \eta_r \hat{r}_t] + \omega_{t-1}, \quad (64)$$

where $\tilde{z}_t \equiv \hat{z}_t - \hat{z}_t^n$ and $\tilde{r}_t \equiv (i_t - \mathbb{E}_t \pi_{t+1}) - r_t^n$ is the real-rate gap. The residual ω_{t-1} collects the SDF wedge $\psi(\bar{r})\bar{X}\mathbb{E}_{t-1}(m_t - m_t^n)$; using feasibility $\hat{s}_t \simeq \hat{c}_t$,

$$\omega_{t-1} \approx -\psi(\bar{r})\bar{X}\sigma\mathbb{E}_{t-1}(\tilde{c}_t - \tilde{c}_{t-1}). \quad (65)$$

G Reduced form of the NKPC-inv

Derivation of the inventory-sales surprise relation. Start from the stock-flow identity with one-period time-to-sale:

$$X_t = Y_{t-1} + X_{t-1} - S_t. \quad (66)$$

Conditioning on information at $t-1$, note that Y_{t-1} and X_{t-1} are known at $t-1$. Hence:

$$\mathbb{E}_{t-1}X_t = Y_{t-1} + X_{t-1} - \mathbb{E}_{t-1}S_t. \quad (67)$$

Subtract (67) from (66):

$$X_t - \mathbb{E}_{t-1}X_t = -(S_t - \mathbb{E}_{t-1}S_t). \quad (68)$$

Linearize around the steady state $(\bar{X}, \bar{S})$ using $X_t = \bar{X}(1 + \hat{x}_t)$ and $S_t = \bar{S}(1 + \hat{s}_t)$:

$$X_t - \mathbb{E}_{t-1}X_t = \bar{X}(\hat{x}_t - \mathbb{E}_{t-1}\hat{x}_t) + \mathcal{O}(2), \quad (69)$$

$$S_t - \mathbb{E}_{t-1}S_t = \bar{S}(\hat{s}_t - \mathbb{E}_{t-1}\hat{s}_t) + \mathcal{O}(2). \quad (70)$$

Plugging these into (68) and dropping second-order terms yields:

$$\hat{x}_t - \mathbb{E}_{t-1}\hat{x}_t = -\frac{\bar{S}}{\bar{X}}(\hat{s}_t - \mathbb{E}_{t-1}\hat{s}_t) \equiv -\frac{\bar{S}}{\bar{X}}(\hat{s}_t - \mathbb{E}_{t-1}\hat{s}_t). \quad (71)$$

With the rebate convention, linear feasibility implies $\hat{s}_t \approx \hat{c}_t$ at first order, and our gap is $\tilde{s}_t \equiv \hat{s}_t - \hat{s}_t^n$. Because the natural allocation is taken as exogenous to firms' within-period sales choice, $\mathbb{E}_{t-1}\hat{s}_t^n$ is known at $t-1$ (or, under rational expectations for the natural process, $\hat{s}_t^n - \mathbb{E}_{t-1}\hat{s}_t^n$ is an exogenous shock). Hence the same structure applies to gaps:

$$\tilde{x}_t = \mathbb{E}_{t-1}\tilde{x}_t - \frac{\bar{S}}{\bar{X}}(\tilde{c}_t - \mathbb{E}_{t-1}\tilde{c}_t). \quad (72)$$

Final reduced form NKPC-inv. Substitute (31) and (32) into (29). The terms in $\mathbb{E}_{t-1}\tilde{r}_t$ cancel and the Phillips curve becomes:

$$\begin{aligned}\pi_t = & \beta\mathbb{E}_t\pi_{t+1} + (\kappa + \eta_x) \left[\sigma \tilde{c}_t + \frac{\nu + 1 - \alpha}{\alpha} \tilde{y}_t \right] - \eta_x\eta_r (\tilde{r}_t - \mathbb{E}_{t-1}\tilde{r}_t) \\ & - \frac{\eta_x}{\psi(\bar{r})\bar{X}} \left[\sigma \tilde{c}_{t-1} + \frac{\nu + 1 - \alpha}{\alpha} \tilde{y}_{t-1} \right] + \eta_x \frac{\bar{S}}{\bar{X}} (\tilde{c}_t - \mathbb{E}_{t-1}\tilde{c}_t) + \mu_t.\end{aligned}\tag{73}$$