Bank of England

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

Staff Working Paper No. 1,153
April 2026

Diego Rodrigues and Tim Willems

This is an updated version of the Staff Working Paper originally published on 14 November 2025

Bank of England

Staff Working Paper No. 1,153

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

Diego Rodrigues[(1)] and Tim Willems[(2)]

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 US 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 stabilisation when inventories are more plentiful – the reason being that the central bank faces a more favourable sacrifice ratio in such an environment.

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

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

---

(1) Université du Québec à Montréal (UQAM). Email: de_sousa_rodrigues.diego@uqam.ca
(2) Bank of England and Centre for Macroeconomics. Email: tim.willems@bankofengland.co.uk

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.

The Bank's working paper series can be found at www.bankofengland.co.uk/working-paper/staff-working-papers

Bank of England, Threadneedle Street, London, EC2R 8AH
Email: enquiries@bankofengland.co.uk

ISSN 1749-9135 (on-line)

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 inventory will be increasing as well, impacting your company's bottom line." It then goes

[1] See 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".

[2] See 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 Interest Rate" (https://rackbeat.com/en/how-an-inventory-system-helps-you-counteract-the-increased-interest-rates/).

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 (2008*a*,*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 get an incentive to lower their prices to economize on inventory holdings.[4] Firms carrying more inventory when the 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 uses 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 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

[3]Also 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.

[4]One 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.

[5]In 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 (Baqaee & Rubbo 2023).

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 address 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] the broader literature provides mixed or limited evidence. Empirically, Maccini & Rossana (1984) and Ramey (1989) find at best weak support for a negative effect of real interest rates on inventories, and Benati & Lubik (2014) reach similarly skeptical conclusions in a structural VAR framework. Together, these contributions have contributed to the hypothesis'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 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

[6] Also 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.

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 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.

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.

[7] With 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.

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:

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

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 starts out with an inventory level of X_0, it will have $X = X_0 - S(P)$ 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:

$$Y = \frac{\vartheta}{\psi_y}, \tag{2}$$

$$0 = [P + \psi_x(X_0 - S(P)) - \vartheta]\, S'(P) + S(P). \tag{3}$$

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. □

As we show in Appendix A, the same result obtains if we instead start from a specification where inventories *directly* boost sales, which is a standard way to tractably introduce a role for inventories in DSGE models (see, e.g., Bils & Kahn (2000) and Mehrotra et al. (2025)) and also the modelling approach we adopt in Section 5 below.

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 carrying-cost inspired logic that 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, our findings support 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:

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

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" represents the inventory level. 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 monetary policy shocks, (ii) the logged price level, (iii) the inventory metric, and (iv) the interactions between shocks and inventory levels "$MPS_t \times INV_t$" (but, as shown in Appendix B, 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 PCE goods price index varies with retailers' inventory holdings. However, given possible concerns regarding aggregation issues (it is not obvious how price and inventory data should be summed across various types of goods for the purposes of our exercise[9]), Sections 4.1.2 and 4.1.3 extend our baseline 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

[8] We 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.

[9] Ideally, 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.

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.

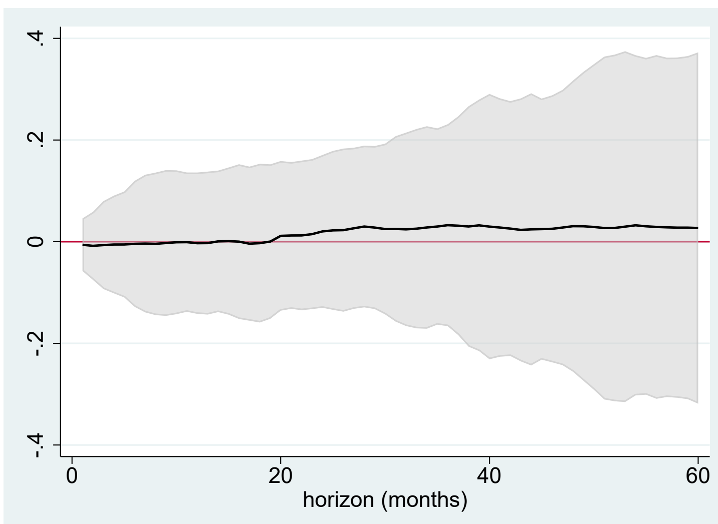

Figure 1: Response of PCE-goods price index to a 25-bp contractionary monetary policy shock, estimated via equation (4), when the inventory-sales ratio "INV_t" stands at its historical average (INV_{avg}). The figure plots $\hat{\beta}_h + \hat{\gamma}_h \cdot INV_{avg}$. Shaded area represents the 90%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

Figure 2 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

[10] Here, 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).

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

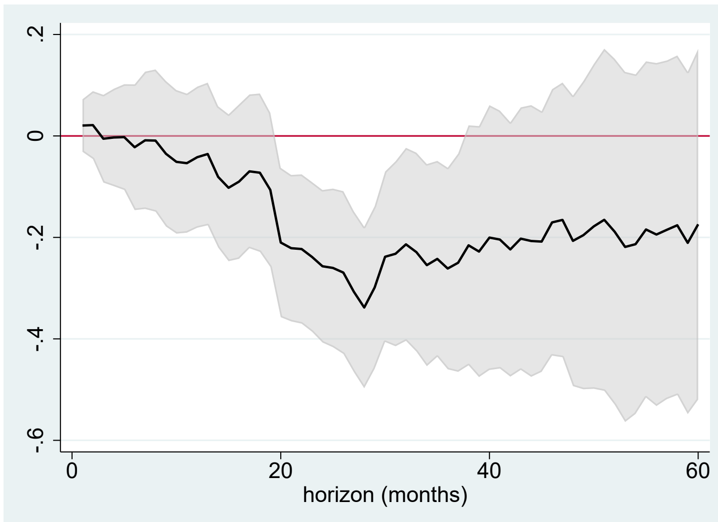

Figure 2: Additional response of PCE-goods price index to a 25-bp contractionary monetary policy shock, due to a unit increase in the inventory-sales ratio, estimated via equation (4). The figure plots $\hat{\gamma}_h$. Shaded area represents the 90% confidence interval, calculated via Newey-West standard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

As shown in Appendix B, 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 too high a level of aggregation. 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

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

represents the inventory of homes looking to be utilized (see Appendix C 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 that follow 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.

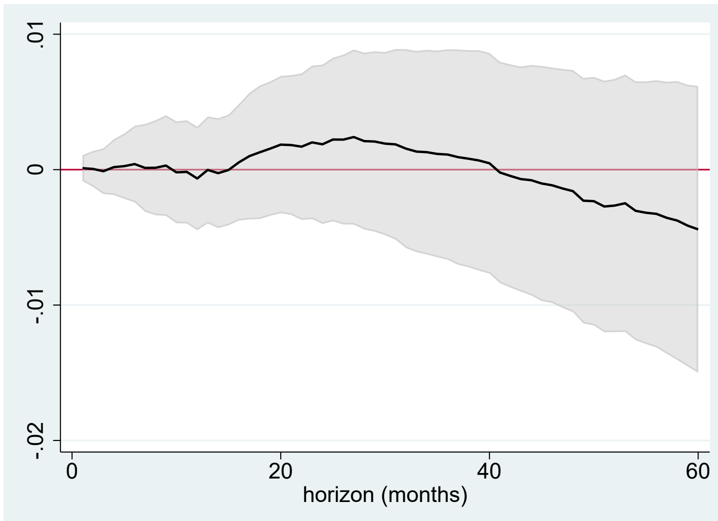

Figure 3: Response of CPI OER to a 25-bp contractionary monetary policy shock, estimated via equation (4), when the home vacancy rate "INV_t" stands at its historical average (INV_{avg}). The figure plots $\hat{\beta}_h + \hat{\gamma}_h \cdot INV_{avg}$. Shaded area represents the 90% confidence interval,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

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 to be occupied.[12] As shown in Appendix B, 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 15 in Appendix B shows, our main result is also visible when using the *overall* CPI as dependent variable in (4) – which is perhaps

[12]The 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).

to be expected, given that OER accounts for about one-third of the total CPI basket. A similar picture emerges when looking at the overall PCE index.

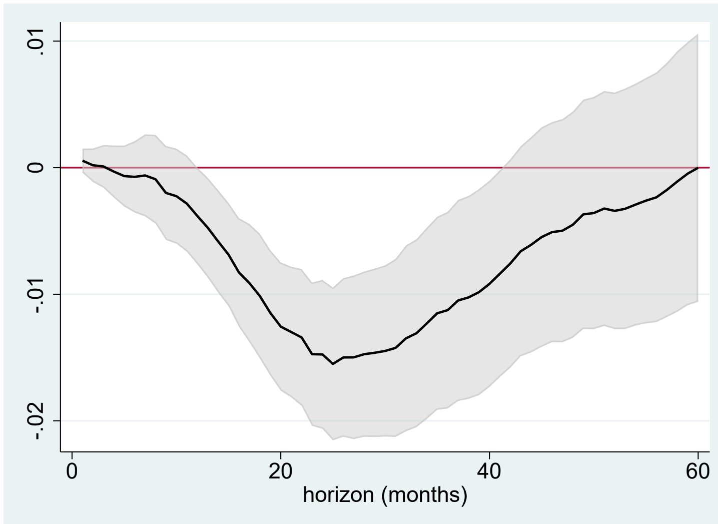

Figure 4: Additional response of CPI OER to a 25-bp contractionary monetary policy shock, due to a 1-pp increase in the home vacancy rate, estimated via equation (4). The figure plots $\hat{\gamma}_h$. Shaded area represents the 90% confidence interval, calculated via Newey-West standard errors.

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 B, similar results follow from 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 a positive 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 B demonstrates that this result – which is again in line with cost-of-carry logic – is robust to the inclusion of further controls.

[13]The 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.

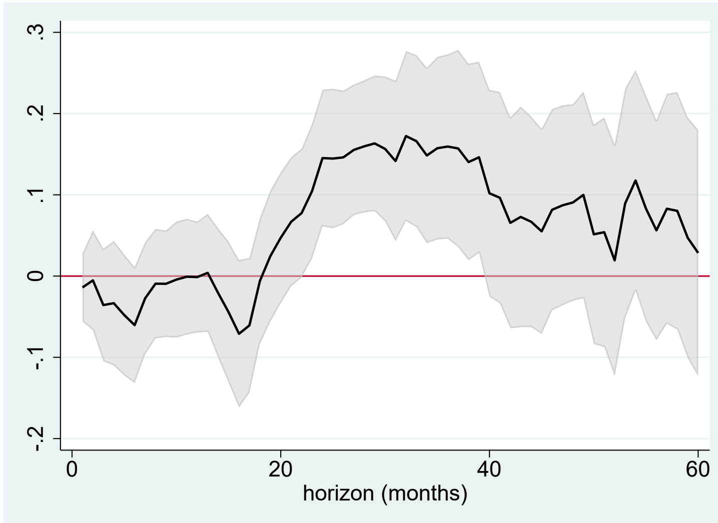

Figure 5: Response of oil prices to a 25-bp contractionary monetary policy shock, estimated via equation (4), when the oil "INV_t" stands at its historical average (INV_{avg}). The figure plots $\hat{\beta}_h + \hat{\gamma}_h \cdot INV_{avg}$. Shaded area represents the 90% confidence interval, calculated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

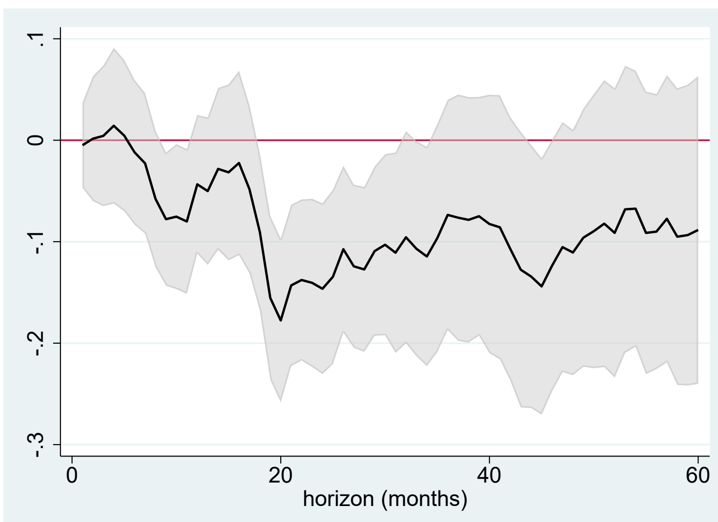

Figure 6: Additional response of oil prices to a 25-bp contractionary monetary policy shock, due to a 1-pp increase in oil inventory, estimated via equation (4). The figure plots $\hat{\gamma}_h$. Shaded area represents the 90% confidence interval, calculated via Newey-West standard errors.

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}.

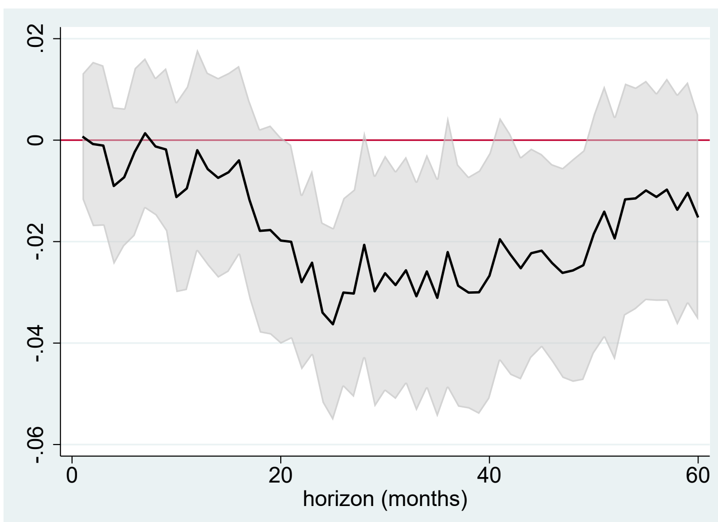

Figure 7: Response of goods inventory-sales ratio to a 25-bp contractionary monetary policy shock, estimated via equation (5). The figure plots $\hat{\beta}_h$. Shaded area represents the 90% confidence interval, calculated via Newey-West standard errors.

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

[14]This 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"

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.

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 and the cost-of-carry channel of monetary policy (along the lines of the approach taken in Section 3). This extension ("NK-inv") alters the Phillips curve in a fundamental way, as it introduces two inventory-related wedges – 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, that is, the cost-of-carry channel.[15] Together, these two wedges strengthen with the steady-state level of inventories and give rise to a sacrifice ratio that varies with inventory holdings. In particular, holding slack 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 our NK-inv model and present the resulting log-linearized equilibrium conditions. Full derivations are provided in Appendices D–H.

5.1 Firms and inventories

Market structure and demand. Time is discrete, $t = 0, 1, 2, \ldots$. A unit mass of monopolistically competitive producers $j \in [0, 1]$ sell directly to a Dixit-Stiglitz final-good aggregator. Final demand for variety j depends on the firm's relative price $(P_{j,t}/P_t)$, as governed by the elasticity of substitution across varieties $(\varepsilon > 1)$, and on its available stock on the shelf $(Q_{j,t}/Q_t)$:

$$S_{j,t} = \left(\frac{Q_{j,t}}{Q_t}\right)^{\zeta} \left(\frac{P_{j,t}}{P_t}\right)^{-\varepsilon} S_t, \qquad \varepsilon > 1,\ \zeta > 0, \tag{6}$$

where Q_t is the economy-wide "availability index" (taken as given by each firm) and ζ governs the *stock-on-shelf elasticity*: when $\zeta > 0$, a firm holding a relatively larger stock

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.

[15] The 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.

ends up selling more (e.g., because of lower stock-out risk). Such a "stock-elastic demand specification" is a standard way to tractably introduce a role for inventories in DSGE models, see e.g. Bils & Kahn (2000) and Mehrotra et al. (2025). Associated with this problem we have the following price index:

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

Technology and inventory timing. Each firm j produces according to the technology:

$$Y_{j,t} = A_t L_{j,t}^{\alpha}, \qquad 0 < \alpha \leq 1, \tag{8}$$

where A_t is the economy-wide level of productivity and α is the output elasticity of labor, $L_{j,t}$. The quantity available for sale is:

$$Q_{j,t} = Y_{j,t} + X_{j,t-1}, \tag{9}$$

where $X_{j,t-1}$ is the inventory stock carried from the previous period. Available stock is allocated between sales and end-of-period inventories:

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

We abstract from the non-negativity constraint on inventories (analogous to how the standard New Keynesian model abstracts from the zero lower bound on the policy rate). In a symmetric equilibrium ($P_{j,t} = P_t$, $Y_{j,t} = Y_t$, etc.), the goods identity states that output is either sold or added to inventory:

$$Y_t = S_t + \Big(X_t - X_{t-1} \Big). \tag{11}$$

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 $\mathcal{C}_t(Y_{j,t})$ denote the real production cost with marginal cost $mc_{j,t}^p \equiv \mathcal{C}_t'(Y_{j,t})$. 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

[16] Like with prices, the quadratic formulation provides analytical convenience, but can also be thought of as proxying the idea that as firms decide to hold more in stock, they have less liquid working capital (or end up borrowing more). This makes them riskier to lend to from a creditor's perspective, which translates into higher borrowing costs through a risk premium. That could be captured more generally by modelling total carrying cost as $\psi(r_t, X_{j,t})X_{j,t}$, with $\partial\psi(r_t, X_{j,t})/\partial X_{j,t} > 0$ capturing the risk premium channel, but note that (12) obtains from working with $\psi(r_t, X_{j,t}) = \psi(r_t)X_{j,t}$.

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:

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

Two marginal costs and stocking conditions. Because firms simultaneously set prices, produce, and manage inventories, the model features two distinct marginal costs. The *production marginal cost* equals $mc_t^p \equiv \mathcal{C}_t'(Y_{j,t})$. The *sales/delivery marginal cost* mc_t^s is the shadow cost of releasing a unit from available stock for sale (the Lagrange multiplier on the constraint $Q_{j,t} = S_{j,t} + X_{j,t}$). The full Lagrangian derivation is in Appendix D. These two costs are linked by a static and a dynamic stocking condition, along the following lines:

Static stocking condition. Define the stock-on-shelf demand weight $\theta_t \equiv \zeta\, S_t/(S_t + X_t)$, where $S_t/Q_t = S_t/(S_t + X_t)$ is the sales share of available stock and ζ is the stock-on-shelf elasticity from (6). Then:

$$mc_t^p = \theta_t + (1 - \theta_t)\, mc_t^s. \tag{13}$$

Dynamic stocking condition. The intertemporal first-order condition for end-of-period inventories yields:

$$mc_t^s = \mathbb{E}_t[m_{t,t+1}\, mc_{t+1}^p] - \psi(r_t)\, X_t. \tag{14}$$

This equates the shadow value of selling a unit today to the discounted production cost saved tomorrow, net of the associated carrying cost. A rise in the real rate r_t raises $\psi(r_t)$, which *directly reduces* mc_t^s: the cost-of-carry channel operates through this dynamic stocking condition, lowering the effective marginal cost relevant for pricing and thus inflation.

Pricing condition and Phillips curve in marginal-cost form. Pricing depends on the cost of supplying units for sale, so the Rotemberg first-order condition is expressed in terms of mc_t^s. The symmetric nonlinear pricing condition reads:

$$(\varepsilon - 1) = \varepsilon\, mc_t^s - \phi\pi_t(1 + \pi_t) + \beta\phi\mathbb{E}_t\Big[\pi_{t+1}(1 + \pi_{t+1})\frac{S_{t+1}}{S_t}\frac{M_{t+1}}{M_t}\Big]. \tag{15}$$

After log-linearization around $\bar{\pi} = 0$ and $(\bar{r}, \bar{X})$, and substituting the log-linearized dynamic stocking condition (see Appendix E for the full derivation), we obtain the inventory-augmented Phillips Curve ("NKPC-inv") expressed in log-deviations from the deterministic steady state:

$$\pi_t = \beta\mathbb{E}_t\pi_{t+1} + (\kappa + \eta_x)\,\mathbb{E}_t[\hat{m}_{t,t+1} + \widehat{mc}_{t+1}^p] - \eta_x\Big(\hat{x}_t + \eta_r\hat{r}_t\Big), \tag{16}$$

where hats denote log-deviations from steady state (i.e., $\hat{z}_t \equiv \ln Z_t - \ln \bar{Z}$). Here, $m_{t,t+1} \equiv \beta M_{t+1}/M_t$ is the one-step stochastic discount factor (SDF), with $M_t \propto u'(C_t)$ being the marginal utility of the representative household; in steady state $m_{t,t+1} = \beta$, and $\hat{m}_{t,t+1} \equiv \ln m_{t,t+1} - \ln \beta$ measures its log-deviation from β. The coefficients are given by:

$$\kappa \equiv \frac{\varepsilon - 1}{\phi}, \qquad \eta_x \equiv \frac{\varepsilon \psi(\bar{r})\bar{X}}{\phi}, \qquad \eta_r \equiv \frac{\bar{r}\psi'(\bar{r})}{\psi(\bar{r})}. \tag{17}$$

Relative to the textbook NKPC, (16) 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 for the standard New Keynesian model; Rupert & Šustek (2019)). Both Den Haan & Sun (2024) and Mehrotra et al. (2025) also obtain a modified NKPC with similar features (except for the cost-of-carry channel, which is novel to our paper). Mehrotra et al. (2025) estimate their NKPC with inventories and find that the inventory-sales ratio substantially improves the NKPC's empirical fit – explaining both the "missing disinflation" of 2009-2011 and the Covid-era inflation surge.

Special cases and interpretation. If $\psi \equiv 0$ or $\bar{X} = 0$, then $\eta_x = 0$ and (16) collapses to the standard NKPC. Otherwise, a higher real rate raises carrying costs, incentivizes stock drawdowns (lowering mc_t^s via (14)), and reduces current inflation via $-\eta_x \eta_r \hat{r}_t$. A larger inventory buffer ($\hat{x}_t > 0$) dampens price pressures through $-\eta_x \hat{x}_t$: with more stock available, higher future demand can be met from inventory rather than by cranking up production and raising marginal costs. And since the NKPC-inv is forward-looking, this lowers inflationary pressures in the present.

Steady state and calibration of ζ. A positive steady-state inventory level ($\bar{X} > 0$) exists if and only if the stock-on-shelf demand elasticity satisfies:

$$\beta \zeta \frac{\bar{S}}{\bar{S} + \bar{X}} > (\varepsilon - 1)(1 - \beta). \tag{18}$$

Given a target inventory-to-sales ratio $\bar{X}/\bar{S}$ (calibrated to the U.S. historical average of 1.4), the parameter ζ is internally determined by:

$$\zeta = \frac{\bar{S} + \bar{X}}{\beta \bar{S}} \Big[\varepsilon \, \psi(\bar{r}) \bar{X} + (\varepsilon - 1)(1 - \beta) \Big], \tag{19}$$

which pins down ζ from steady-state targets. The derivation of the steady-state conditions can be found in Appendix F.

[17] This is a relative statement, not necessarily inconsistent with the alleged flatness of the Phillips curve (Hazell et al. (2022)). Within the Rotemberg (1982) setup, a flat NKPC is driven by a high degree of price stickiness (high ϕ); since η_x is decreasing in ϕ, a high ϕ could still give rise to a flat slope ($\kappa + \eta_x$).

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

$$AC_t \equiv \frac{\phi}{2}\pi_t^2 S_t, \qquad IC_t \equiv \frac{\psi(r_t)}{2}X_t^2. \tag{20}$$

We follow a standard rebate scheme: the carrying cost IC_t is decomposed into its first-order Taylor approximation V_t around $(\bar{r}, \bar{X})$ (rebated lump-sum to households) and a second-order remainder $\widetilde{IC}_t \equiv IC_t - V_t$ (see Appendix G for details). Since AC_t and $\widetilde{IC}_t$ are both second order, the linearized feasibility conditions are:

$$\hat{s}_t \approx \hat{c}_t, \qquad \hat{y}_t = \hat{c}_t + \frac{\bar{X}}{\bar{Y}}\big(\hat{x}_t - \hat{x}_{t-1}\big). \tag{21}$$

The rebate V_t is financed and returned lump-sum; it does not enter the Euler equation or the labor supply condition.[18] Importantly, the price-setting first-order condition (15) retains the inventory carrying term $\psi(r_t)X_t$ through mc_t^s; this is the channel through which interest rate changes have a direct effect on inflation.

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:

$$\max_{\{C_t, L_t\}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\frac{C_t^{1-\sigma}}{1-\sigma} - \chi \frac{L_t^{1+\nu}}{1+\nu} \right). \tag{22}$$

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:

$$1 = (1 + i_t)\mathbb{E}_t \left[\beta \left(\frac{C_t}{C_{t+1}} \right)^{\sigma} \frac{P_t}{P_{t+1}} \right]. \tag{23}$$

[18] One 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 $\widetilde{IC}_t$ drop out and need not be included in the equilibrium conditions. See Appendix G.

The intratemporal first-order condition for labor supply is:

$$\frac{W_t}{P_t} = \chi L_t^{\nu} C_t^{\sigma}. \tag{24}$$

Log-linearizing (23) around the zero-inflation steady state yields the familiar IS equation:

$$\tilde{c}_t = \mathbb{E}_t \tilde{c}_{t+1} - \frac{1}{\sigma}\Big(i_t - \mathbb{E}_t \pi_{t+1} - r_t^n\Big), \tag{25}$$

where $\tilde{c}_t$ is the consumption gap (or sales gap) in deviations from the natural allocation, i.e., $\tilde{c}_t \equiv \widehat{c}_t - \widehat{c}_t^n \approx \widehat{s}_t - \widehat{s}_t^n$ (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:

$$\tilde{r}_t \equiv \Big(i_t - \mathbb{E}_t \pi_{t+1}\Big) - r_t^n = \sigma\Big(\mathbb{E}_t \tilde{c}_{t+1} - \tilde{c}_t\Big), \tag{26}$$

and the natural real rate by:

$$r_t^n = \bar{r} + \sigma \mathbb{E}_t\Big(\widehat{c}_{t+1}^n - \widehat{c}_t^n\Big), \tag{27}$$

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

Mapping marginal cost to natural gaps. Before defining the equilibrium of our NK-inv model, we rewrite (16) in terms of gaps relative to the natural allocation. We can use the first-order condition for labor supply (24) and the production function (8) to map the production marginal cost gap to the sales/consumption and output gaps (see Appendix H for details):

$$\widehat{mc}_t^p - \widehat{mc}_t^{p,n} = \sigma \tilde{c}_t + \frac{\nu + 1 - \alpha}{\alpha} \tilde{y}_t, \tag{28}$$

where natural allocations (arising under flexible prices) are denoted by a superscript n, and we define the production marginal cost gap as $\widetilde{mc}_t^p \equiv \widehat{mc}_t^p - \widehat{mc}_t^{p,n}$. The log-linearized stock-on-shelf share in natural gaps is:

$$\tilde{\theta}_t = \omega_X\Big(\tilde{c}_t - \tilde{x}_t\Big), \qquad \omega_X \equiv \frac{\bar{X}}{\bar{S} + \bar{X}}. \tag{29}$$

NKPC-inv in terms of natural gaps. Decompose each hatted variable in (16) as $\hat{z}_t = \tilde{z}_t + \hat{z}_t^n$, where $\tilde{z}_t \equiv \hat{z}_t - \hat{z}_t^n$ denotes the gap relative to the natural (flexible-price) allocation and r_t^n is the natural real rate. Substituting into (16) and collecting all natural-allocation terms into the composite residual gives:

$$\mu_t \equiv (\kappa + \eta_x)\, \mathbb{E}_t\Big[\hat{m}_{t,t+1}^n + \widehat{mc}_{t+1}^{p,n}\Big] - \eta_x(\hat{x}_t^n + \eta_r \hat{r}_t^n)\,, \tag{30}$$

and allows us to express the NKPC-inv in gaps relative to natural levels:

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + (\kappa + \eta_x)\, \mathbb{E}_t[\tilde{m}_{t,t+1} + \widetilde{mc}^p_{t+1}] - \eta_x \tilde{x}_t - \eta_x \eta_r \tilde{r}_t + \mu_t. \tag{31}$$

The term μ_t is not an additional structural shock; it is the bookkeeping residual that appears when the NKPC is rewritten purely in natural gaps (see Appendix H for its full derivation and the definition of natural allocations). Under CRRA preferences, $\tilde{m}_{t,t+1}$ and $\widetilde{mc}^p_{t+1}$ can be further reduced using $\tilde{m}_{t,t+1} = -\sigma(\tilde{c}_{t+1} - \tilde{c}_t)$ and (28); see Appendix H. The cost-of-carry channel is confirmed from (31) by noting that, holding producer marginal costs fixed:

$$\left.\frac{\partial \pi_t}{\partial \tilde{r}_t}\right|_{mc^p_t} = -\eta_x \eta_r < 0.$$

A rise in the real rate therefore lowers inflation through this direct cost-of-carry channel, independently of any indirect effect operating through slack and the marginal cost of production.

Inventory accounting. From (10), the linearized inventory law of motion (in natural gaps) is given by:

$$\tilde{x}_t = \tilde{x}_{t-1} + \frac{\bar{S}}{\bar{X}}\big(\tilde{y}_t - \tilde{c}_t\big). \tag{32}$$

Equilibrium.

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

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

Let $b_{sm} \equiv \beta \bar{mc}^p / \bar{mc}^s$, $b_{sx} \equiv \psi(\bar{r})\bar{X}/\bar{mc}^s$, and a_{ps}, $a_{p\theta}$ as defined in Appendix H. Given exogenous sequences $\{r^n_t\}_{t\geq 0}$ and $\{\mu_t\}_{t\geq 0}$, and an initial condition $\tilde{x}_{-1}$, a (linear) rational expectations equilibrium is a collection $\{\pi_t, \tilde{c}_t, \tilde{y}_t, \tilde{x}_t, i_t, \widetilde{mc}^p_t, \widetilde{mc}^s_t, \tilde{\theta}_t, \tilde{m}_{t,t+1}\}_{t\geq 0}$ such that, for all $t \geq 0$:

(i) IS equation:

$$\tilde{c}_t = \mathbb{E}_t \tilde{c}_{t+1} - \frac{1}{\sigma}\tilde{r}_t.$$

(ii) NKPC-inv:

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + (\kappa + \eta_x)\, \mathbb{E}_t[\tilde{m}_{t,t+1} + \widetilde{mc}^p_{t+1}] - \eta_x \tilde{x}_t - \eta_x \eta_r \tilde{r}_t + \mu_t.$$

(iii) Inventory accounting:

$$\tilde{x}_t = \tilde{x}_{t-1} + \frac{\bar{S}}{\bar{X}}\big(\tilde{y}_t - \tilde{c}_t\big).$$

(iv) Production marginal cost:

$$\widetilde{mc}_t^p = \sigma \tilde{c}_t + \frac{\nu + 1 - \alpha}{\alpha} \tilde{y}_t.$$

(v) Stock-on-shelf share:

$$\tilde{\theta}_t = \omega_X (\tilde{c}_t - \tilde{x}_t).$$

(vi) Static stocking condition:

$$\widetilde{mc}_t^p = a_{ps} \widetilde{mc}_t^s + a_{p\theta} \tilde{\theta}_t.$$

(vii) Dynamic stocking condition:

$$\widetilde{mc}_t^s = b_{sm} \, \mathbb{E}_t[\tilde{m}_{t,t+1} + \widetilde{mc}_{t+1}^p] - b_{sx} \Big(\tilde{x}_t + \eta_r \tilde{r}_t \Big).$$

(viii) SDF gap:

$$\tilde{m}_{t,t+1} \equiv \ln m_{t,t+1} - \ln m_{t,t+1}^n.$$

(ix) Monetary policy rule:

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

The real-rate gap is $\tilde{r}_t \equiv i_t - \mathbb{E}_t \pi_{t+1} - r_t^n$, and expectations $\mathbb{E}_t[\cdot]$ are conditional on time-t information.

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 studying 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, $-\eta_x \eta_r \tilde{r}_t$, and an inventory wedge, $-\eta_x \tilde{x}_t$. 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, which is thought to be a good description of central bank practices (Carney (2013)):[19]

$$\mathcal{L}_t = \pi_t^2 + \lambda_c \tilde{c}_t^2, \tag{33}$$

[19] The model-implied loss function, given by $\mathcal{L}_t = \frac{1}{2} \sum_{t=0}^{\infty} \beta^t \left[\pi_t^2 + \frac{\sigma}{\phi} \tilde{c}_t^2 + \frac{\nu}{\phi} \tilde{y}_t^2 + \frac{\lambda_x}{\phi} \tilde{x}_t^2 \right]$, also calls on the central bank to stabilize inventories. Since this is difficult to square with observed central bank mandates, we instead start from (33), which is more practically relevant. However, our core insight – that optimal policy is more inflation-focused when $\bar{X}$ is higher – equally follows when starting from the model-implied loss function. Basing our analysis on the standard loss function, given by (33), makes clear that this insight

where $\lambda_c \geq 0$ captures the relative weight the central bank places on stabilization of the consumption gap $\tilde{c}_t$ (which coincides with the sales gap at first order; see Appendix G). 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 $\epsilon_{\mu,t} \overset{iid}{\sim} N(0, 0.01)$).

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.[20] Remember that, in our model, the strength of this channel is governed by the steady-state level of inventories, $\bar{X}$.

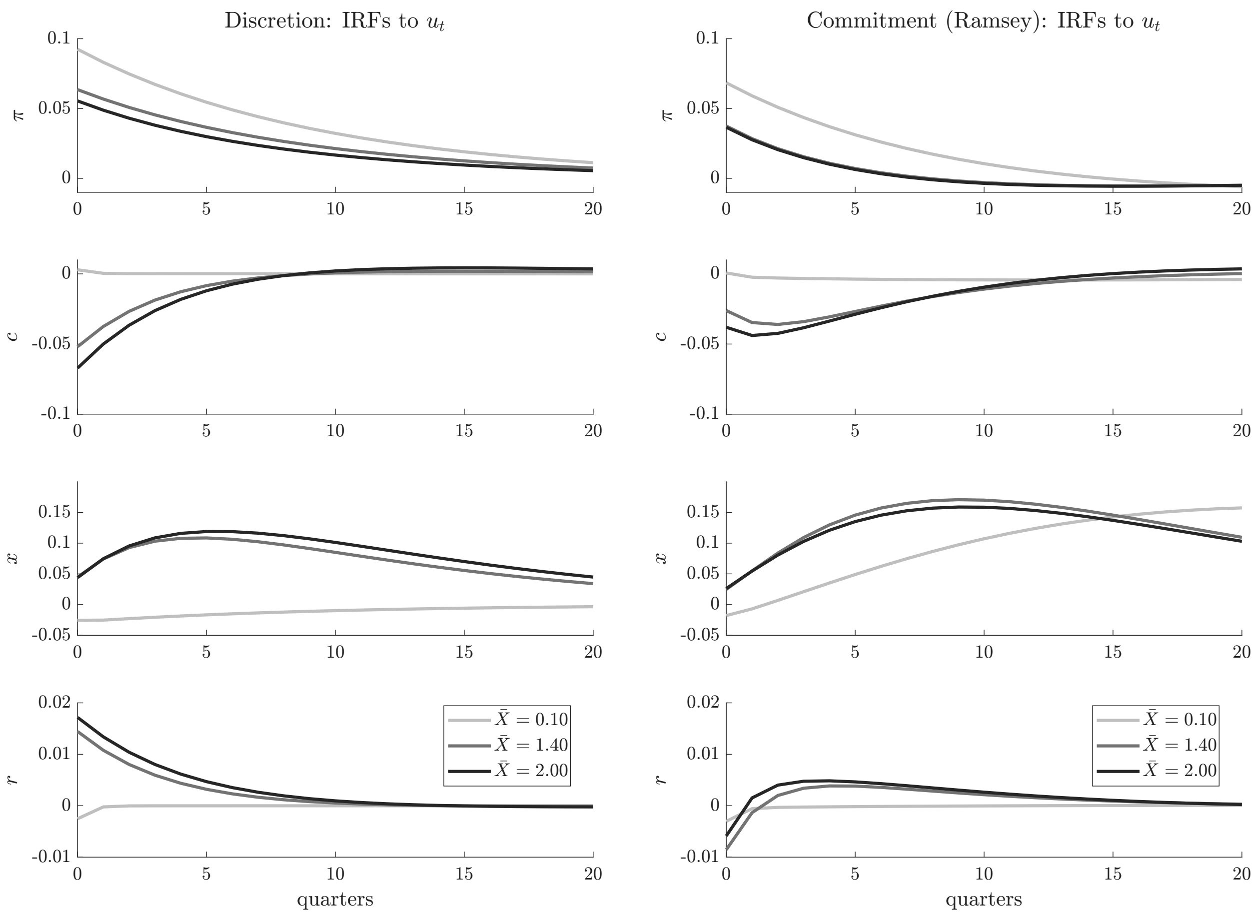

(a) Discretion (b) Commitment

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

Figure 8 therefore considers a low value of $\bar{X} = 0.10$ (which lies in the neighborhood of the

comes from inventory-driven changes to the monetary transmission mechanism; not from changes to the loss function.

[20] We use a relatively standard quarterly calibration. Household parameters: $\beta = 0.99$, $\sigma = 1$, $\alpha = 1$, $\nu = 2$ (Chetty et al. (2011)). Firm parameters: $\varepsilon = 6$, $\phi = 100$ (Bilbiie & Ragot (2021)). Inventory parameters: $\bar{X}/\bar{S} = 1.4$ (U.S. historical average inventory-to-sales ratio, with $\bar{S} = 1$), $\psi(\bar{r}) = 0.05$, $\eta_r = 1.50$. The stock-on-shelf elasticity ζ is calibrated from (19) given these targets, satisfying the existence condition (18).

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 $\bar{X} = 2.0$. 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 stronger).

In Appendix I we show that the same insight obtains if the central bank's loss function were specified in terms of the output gap (as opposed to the consumption gap).

Figure 9 presents the results of an optimal simple rule exercise in which the central bank follows a pure inflation-targeting rule, $i_t = \phi_\pi \pi_t$, and we search for the coefficient ϕ_π^* that minimizes the present discounted value of the loss (33). Two features stand out. First, for low values of $\bar{X}$ — in the neighborhood of the standard New Keynesian model — the Blanchard–Kahn conditions require $\phi_\pi \geq 1$ (the Taylor principle) for equilibrium determinacy, and this constraint binds: the optimum is pinned at $\phi_\pi^* = 1$. This is the familiar result that, absent any additional propagation mechanism, an inflation-only rule already sits at the determinacy boundary, leaving no scope to respond less aggressively to inflation. Second, once $\bar{X}$ is large enough (approximately $\bar{X} > 0.5$ in Figure 9), the constraint ceases to bind: the unconstrained optimum rises strictly above unity and increases monotonically with $\bar{X}$. This confirms our central finding — that the cost-of-carry channel makes inflation stabilization less costly, making the optimal simple rule call for a stronger focus on inflation as inventories accumulate. Appendix J shows that this conclusion is robust when considering a more general interest rate rule, also responding to the consumption gap.

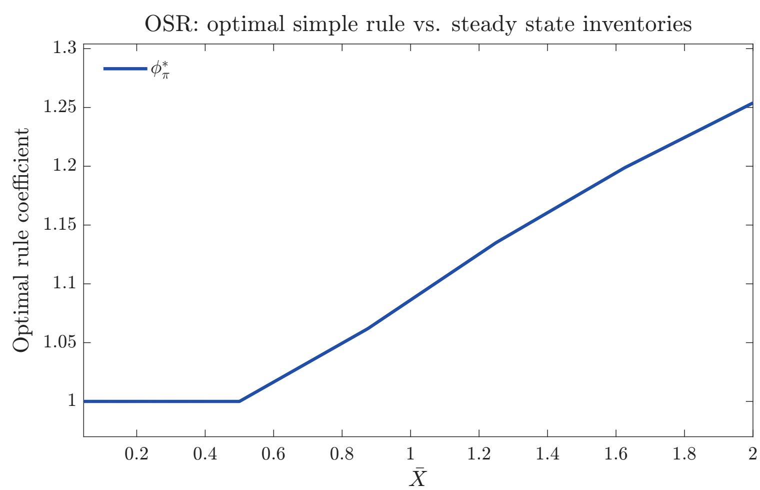

Figure 9: Optimal inflation-targeting coefficient ϕ_π^* as a function of steady-state inventories $\bar{X}$. Rule: $i_t = \phi_\pi \pi_t$.

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 (31)). 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).[21] This makes it optimal for the central bank to put more weight on inflation stabilization when faced with trade-off inducing shocks.[22]

5.5 Generalization to a non-linear environment

So far, we have been analyzing optimal policy in the first-order, log-linearized version of our model. There, we find that the optimal degree of focus on inflation stabilization is increasing in *the steady state level* of inventories, $\bar{X}$. While this insight is mathematically accurate (up to a first-order approximation), the volatility of inventories implies that the (mathematically) higher-order terms may well be *economically* relevant. In this section we therefore demonstrate that the nature of our key insight is more general – with a non-linear analysis showing that optimal policy in our model calls for a greater focus on inflation stabilization whenever inventories stand at a higher level (going beyond a sole reliance on the steady-state object $\bar{X}$).

We demonstrate this by solving the *nonlinear* transition path of the economy under perfect foresight following a one-time cost-push shock (which hits when inventories stand at their "initial" level X_0) and optimizing the policy coefficient over the resulting welfare loss.

Nonlinear equilibrium. Let $m_t \equiv \beta(C_{t+1}/C_t)^{-\sigma}$ denote the gross discount factor and let $\psi_t \equiv \bar{\psi}(1+\eta_r(r_t - r_{ss})/r_{ss})$ represent the carrying cost.

Definition 2 (Perfect-foresight equilibrium, NK-inv)**.** *Given an initial inventory level X_0, a one-time cost-push shock $u_1 > 0$, and a monetary policy rule, a* perfect-foresight equilibrium *is a sequence $\{C_t, S_t, Y_t, \pi_t, i_t, r_t, mc_t^p, mc_t^s, \theta_t, X_t\}_{t\geq 1}$ converging to the deterministic steady state such that, for all $t \geq 1$:*

(i) Resource constraint:

$$S_t\Big(1 - \tfrac{\phi}{2}\pi_t^2\Big) = C_t + \tfrac{\bar{\psi}}{2}(X_t - \bar{X})^2.$$

(ii) Stock-on-shelf share:

$$\theta_t = \zeta \, \frac{S_t}{S_t + X_t}.$$

[21] See Nekarda & Ramey (2013) for empirical evidence pointing to price-cost markups being procyclical in 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.

[22] The 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 ($\breve{\kappa} > \kappa$) and introduce a direct real rate channel, with its strength governed by $\vartheta > 0$: $\pi_t = \beta\mathbb{E}_t\pi_{t+1} + \breve{\kappa}\widetilde{c}_t - \vartheta r_t$. By following the standard steps to solve for optimal policy under discretion, one then finds that the optimal targeting criterion is $\pi_t = -\frac{\lambda_c}{\breve{\kappa}+\sigma\vartheta}\widetilde{c}_t$. Since $\breve{\kappa} + \sigma\vartheta > \kappa$ (the latter being the denominator for the standard NK model), it follows that optimality calls for greater inflation stabilization under the modified NKPC.

(iii) Static stocking condition:

$$mc_t^p = \theta_t + (1-\theta_t)\, mc_t^s.$$

(iv) Dynamic stocking condition:

$$mc_t^s + \psi_t\, X_t = m_t\, mc_{t+1}^p.$$

(v) NKPC:

$$(\varepsilon - 1) - \varepsilon\, mc_t^s + \phi\, \pi_t(1+\pi_t) - \phi\, m_t \tfrac{S_{t+1}}{S_t} \pi_{t+1}(1+\pi_{t+1}) = u_t.$$

(vi) Euler equation:

$$m_t\, \frac{1+i_t}{1+\pi_{t+1}} = 1.$$

(vii) Fisher relation:

$$(1+r_t)(1+\pi_{t+1}) = 1 + i_t.$$

(viii) Monetary policy rule:

$$i_t = i_{ss} + \phi_\pi\, \pi_t.$$

(ix) Production marginal cost:

$$mc_t^p = \chi\, Y_t^{\frac{\nu+1-\alpha}{\alpha}}\, C_t^\sigma,$$

where $\chi > 0$ collects the remaining parameters.

(x) Inventory accumulation:

$$X_t = X_{t-1} + Y_t - S_t.$$

Solution method and optimal simple rule. The system of $10 \times T$ nonlinear equations (with $T = 30$ quarters) is solved by Newton-based continuation.[23] We vary the initial inventory level across $X_0 \in \{1.4, 1.5, \ldots, 2.0\}$ and, for each X_0, we search for the optimal inflation-response coefficient ϕ_π^* by minimizing:

$$\mathcal{L}(\phi_\pi;\, X_0) = \sum_{t=1}^{T} \beta^{t-1} \Big[\pi_t^2 + \lambda_c \Big(\ln \tfrac{C_t}{\bar{C}} \Big)^2 + \lambda_i\, (i_t - \bar{i})^2 \Big] \tag{34}$$

In this case, we include λ_i (set to 0.3) which is a penalty term on interest rate deviations $(i_t - i_{ss})^2$ to ensure an interior optimum for all values of X_0 considered (without it, the

[23]The $10 \times T$ equations are stacked into a single nonlinear system and solved by a trust-region Newton algorithm. To ensure convergence from a non-trivial starting point we employ a homotopy approach: the cost-push shock is ramped up in ten equal steps from zero to its target value, with the solution at each step serving as the warm start for the next. See Juillard (1996) for an early treatment of Newton-based methods for solving forward-looking models by stacking the full transition path.

central bank obtains an incentive to send $\phi_\pi \to \infty$). Otherwise, the calibration follows the linearized model.[24]

Main result: a stronger focus on inflation in high-inventory states. Figure 10 reports ϕ^*_π as a function of X_0. It generalizes the core insight from the linearized model: the optimal response coefficient on inflation rises with initial inventories, from $\phi^*_\pi = 1.09$ when $X_0 = 1.4$ to $\phi^*_\pi = 2.09$ when $X_0 = 2.0$. The intuition mirrors the linearized analysis, but goes beyond a sole dependence on *the steady-state level* of inventories. Instead, when *the actual level* of inventories (X_0) stands at a higher level when the shock hits, the central bank faces a more favorable sacrifice ratio, which makes a more aggressive inflation response optimal. This channel seems quantitatively important: moving the inventory-sales ratio from its steady-state value (around 1.4) by two standard deviations (towards 1.7) raises the optimal response coefficient on inflation by almost 30 percent (from $\phi^*_\pi = 1.09$ to around 1.41).

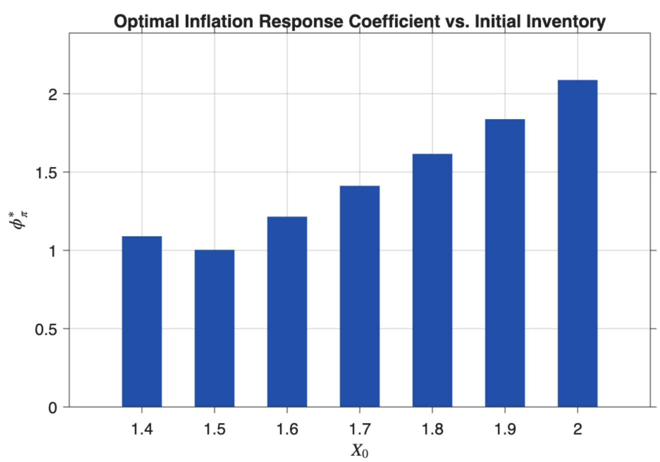

Figure 10: Optimal inflation-response coefficient ϕ^*_π in the nonlinear model as a function of the initial inventory X_0. Each bar is the welfare-minimising value of ϕ_π for that initial condition.

Appendix K shows that the same qualitative insight follows when the interest rate rule is extended to include the output gap as an additional argument; using the consumption gap instead yields similar results.

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

[24]In particular, we use: $\beta = 0.99$, $\sigma = 1$, $\nu = 2$, $\varepsilon = 6$, $\phi = 100$, $\bar{\psi} = 0.05$, $\eta_r = 1.5$, $\bar{X} = 1.40, \lambda_c = 0.1$. The cost-push shock $u_1 = 2.0$ is active only at $t = 1$.

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 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.

This paper also 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 and important extension.

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.

Baqaee, 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.

Carney, M. (2013), 'Monetary policy after the fall', *Remarks at Eric J. Hanson Memorial Lecture, Üniversity of Alberta, Edmonton, Alberta* **1**.

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. (2008*a*), 'An explanation for soaring commodity prices', https://voxeu.org/article/explanation-soaring-commodity-prices.

Frankel, J. A. (2008*b*), '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'.

Juillard, M. (1996), Dynare: A program for the resolution and simulation of dynamic models with forward variables through the use of a relaxation algorithm, Couverture Orange 9602, CEPREMAP.

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 Inventories directly boosting sales in the simple model

The core of our argument can also be captured through a version of the simple static model in which inventories directly stimulate sales. Consider a profit-maximizing firm that enters the period with X_0 units of inventory and produces Y additional units, both immediately available for sale. Having goods on the shelf is assumed to stimulate sales directly (for example because it lowers incidence of stockouts), so the sales function depends on both price and total shelf stock: $S(P, Q)$, assumed to be continuous, three-times differentiable in P at given Q, with $S_P < 0$ and $S_Q > 0$. The total stock of goods available is $Q = X_0 + Y$. After selling $S(P, Q)$ units, unsold inventory $X = Q - S(P, Q)$ is carried into the next period. The firm's static problem is:

$$\begin{aligned} \max_{P,Y} \quad & PS(P,Q) - \psi_y \frac{Y^2}{2} - \psi_x \frac{X^2}{2}, \\ \text{s.t. } & Q = X_0 + Y, \\ & X = Q - S(P,Q), \\ & X \geq 0. \end{aligned} \tag{35}$$

The cost of producing Y units is $\psi_y \frac{Y^2}{2}$, and the cost of carrying X unsold units is $\psi_x \frac{X^2}{2}$. This specification is consistent with the one we adopt in our NK-inv model of Section 5.

The first-order condition for price-setting is:

$$[P + \psi_x X]\, S_P(P,Q) + S(P,Q) = 0. \tag{36}$$

The first-order condition for production determines Q jointly with (36):

$$[P + \psi_x X]\, S_Q(P,Q) - \psi_x X = \psi_y Y. \tag{37}$$

The price-setting condition (36) has the same structure as in the baseline simple model: the firm trades off the revenue from a higher price against the carrying cost of any unsold inventory. As before, 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 ψ_x. As the following proposition shows, the model with stock-on-shelf demand has the exact same implication in this regard (provided that $S_Q < 1$, meaning that adding one more unit to the shelf raises demand by less than one unit, which we view as a very mild condition):

Proposition 2. ***(Price setting with stock-on-shelf demand)*** *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.:*

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

Moreover, under $S_Q < 1$, the strength of this effect is increasing in the firm's pre-existing inventory level X_0.

Proof. Applying the Implicit Function Theorem to (36) at given Q:

$$\left.\frac{\partial P}{\partial \psi_x}\right|_Q = -\frac{X \cdot S_P(P,Q)}{\frac{\partial}{\partial P}[(P+\psi_x X)S_P + S]_Q}.$$

At an interior optimum ($X > 0$) with $S_P < 0$, the numerator satisfies $-XS_P > 0$. For the denominator, differentiating at fixed Q (so $\frac{\partial X}{\partial P} = -S_P$):

$$\frac{\partial}{\partial P}[(P+\psi_x X)S_P + S]_Q = 2S_P + (P+\psi_x X)S_{PP} - \psi_x S_P^2.$$

Substituting $P + \psi_x X = -S/S_P$ from (36):

$$\frac{\partial}{\partial P}[(P+\psi_x X)S_P + S]_Q = 2S_P - \frac{S}{S_P}S_{PP} - \psi_x S_P^2 = S_P\left(2 - \frac{S \cdot S_{PP}}{S_P^2} - \psi_x S_P\right).$$

Hence,

$$\left.\frac{\partial P}{\partial \psi_x}\right|_Q = -\frac{X \cdot S_P(P,Q)}{S_P\left(2 - \frac{S \cdot S_{PP}}{S_P^2} - \psi_x S_P\right)}.$$

Since $X > 0$ and $S_P < 0$, it follows that $\left.\frac{\partial P}{\partial \psi_x}\right|_Q < 0$ provided $2 - \frac{S \cdot S_{PP}}{S_P^2} - \psi_x S_P > 0$. Because $S_P < 0$, the last term satisfies $-\psi_x S_P > 0$, so a sufficient condition is $\frac{S \cdot S_{PP}}{S_P^2} < 2$. As in the proof of Proposition 1, this condition is ensured by the regularity restriction emphasized by Mrázová & Neary (2017). Under that condition $\left.\frac{\partial P}{\partial \psi_x}\right|_Q < 0$.

For the second part of the proposition, note that at fixed P and Y a unit increase in X_0 raises Q by one, which increases S by S_Q and hence raises unsold inventory by $\frac{\partial X}{\partial X_0} = 1 - S_Q$. Therefore, provided $S_Q < 1$, a higher initial inventory level increases the firm's inventory exposure. Since the numerator in $\partial P/\partial \psi_x$ is proportional to X, this amplifies the magnitude of the price response to a rise in carrying costs. □

B 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 11. 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 12 and 13, respectively.

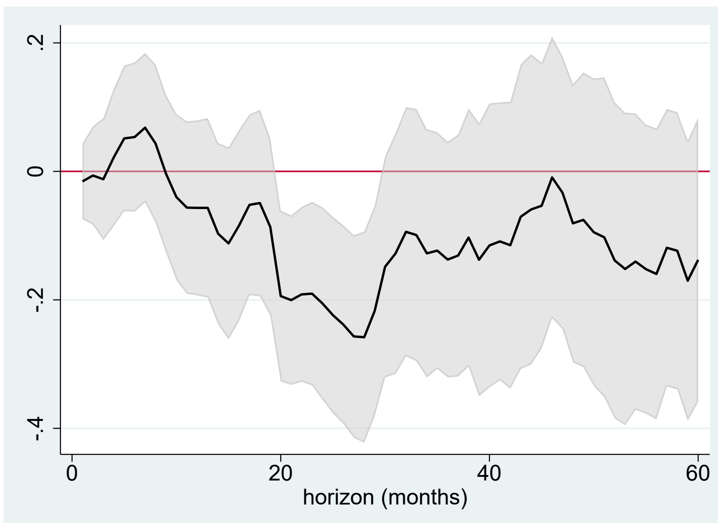

Figure 11: Additional response of PCE-goods price index to a 25-bp contractionary monetary policy shock, due to a unit increase in the inventory-sales ratio, estimated via equation (4), with the rate of unemployment and the 5-year Treasury yield added to the vector of controls (Z_t). The figure plots $\hat{\gamma}_h$. 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 14). 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 15), which is perhaps not that surprising given OER accounts for about one-third of the overall CPI.

Finally, with respect to the core of our result for oil prices, Figure 16 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).

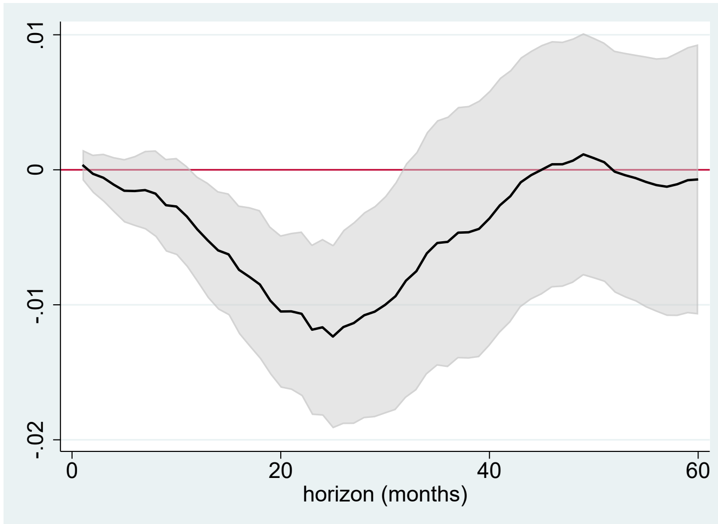

Figure 12: Additional response of CPI OER to a 25-bp contractionary monetary policy shock, due to a 1-pp increase in the home vacancy rate, estimated via equation (4), with the rate of unemployment and the 5-year Treasury yield added to the vector of controls (Z_t). The figure plots $\hat{\gamma}_h$.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

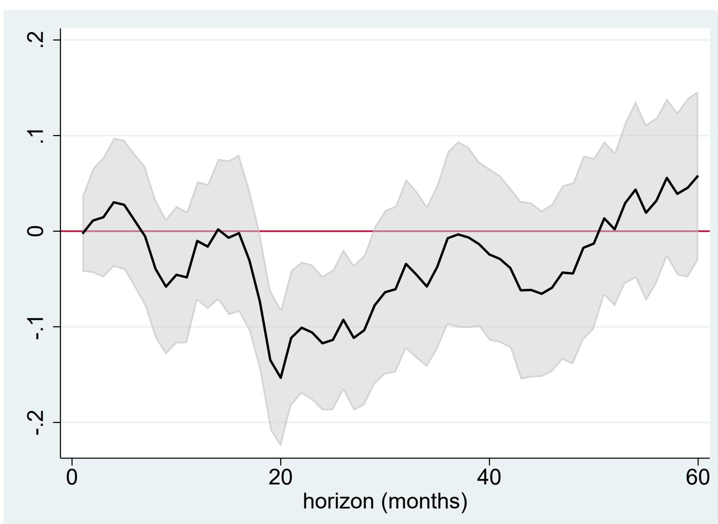

Figure 13: Additional response of oil prices to a 25-bp contractionary monetary policy shock, due to a 1-pp increase in oil inventory, estimated via equation (4), with the rate of unemployment and the 5-year Treasury yield added to the vector of controls (Z_t). The figure plots $\hat{\gamma}_h$. Shaded area represents the 90% confidence interval, calculated via Newey-West standard errors.

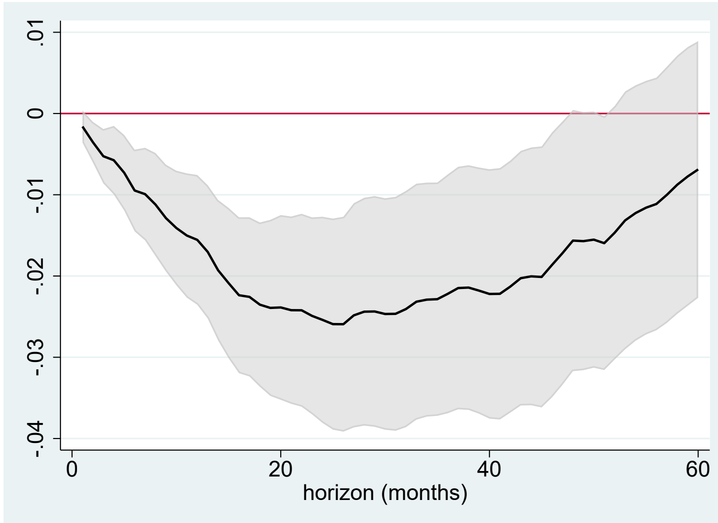

Figure 14: Additional response of the housing component of the PCE index to a 25-bp contractionary monetary policy shock, due to a 1-pp increase in the home vacancy rate, estimated via equation (4). The figure plots $\hat{\gamma}_h$. Shaded area represents the 90% confidence interval, calculated via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

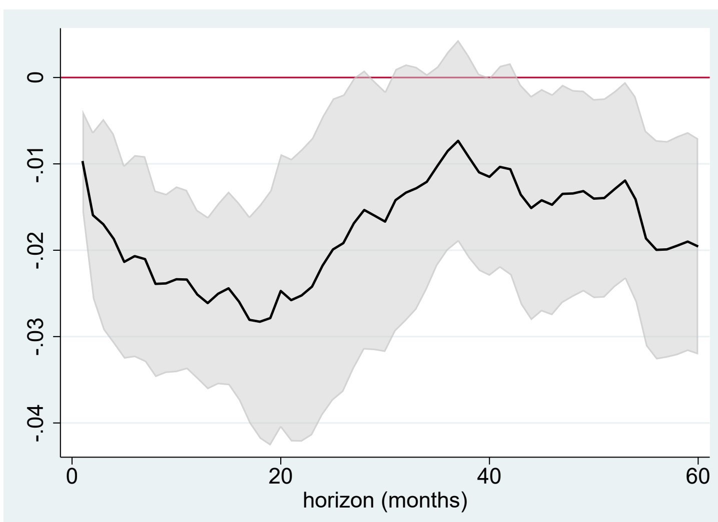

Figure 15: Additional response of overall CPI to a 25-bp contractionary monetary policy shock, due to a 1-pp increase in the home vacancy rate, estimated via equation (4). The figure plots $\hat{\gamma}_h$. Shaded area represents the 90% confidence interval, calculated via Newey-West standard errors.

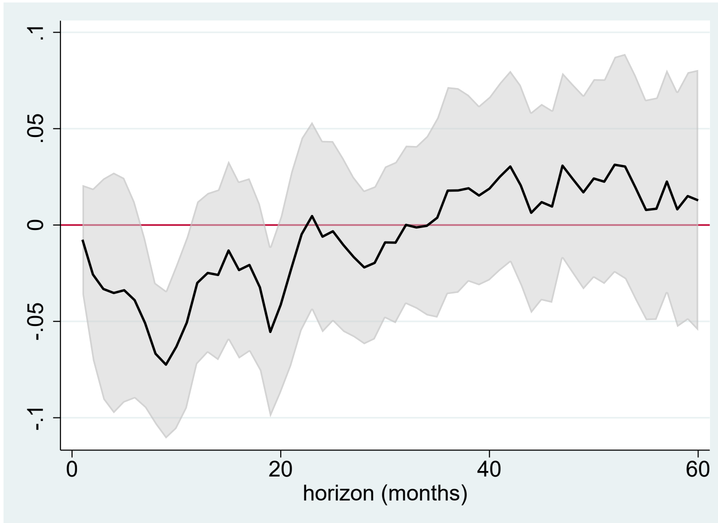

Figure 16: Additional response of oil prices to a 25-bp contractionary monetary policy shock, due to a 1-pp increase in oil inventory, estimated via equation (4), when applying linear detrending. The figure plots $\hat{\gamma}_h$. Shaded area represents the 90% confidence interval, calculated via Newey-West standard errors.

C 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: RRVRUSQ156N) 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.

D Price setting, stocking conditions, and NKPC-inv

This appendix derives the firm's optimal conditions under stock-on-shelf demand (6), contemporaneous production (9), and Rotemberg adjustment costs. We work with the full intertemporal Lagrangian.

D.1 Firm problem and intertemporal Lagrangian

A firm chooses sequences $\{P_{j,t}, Y_{j,t}, Q_{j,t}, S_{j,t}, X_{j,t}\}_{t\geq 0}$ to maximise $\mathbb{E}_0 \sum_{t=0}^{\infty} \Lambda_t \Psi_{j,t}$ subject to three period-t constraints. Introducing multipliers $\kappa_{j,t}$, $\mu_{j,t}$, $\lambda_{j,t}$, the constraints are:

$$g^D_{j,t}: \quad S_{j,t} - \left(\frac{Q_{j,t}}{Q_t}\right)^{\zeta} \left(\frac{P_{j,t}}{P_t}\right)^{-\varepsilon} S_t = 0, \tag{38}$$

$$g^Q_{j,t}: \quad Q_{j,t} - Y_{j,t} - X_{j,t-1} = 0, \tag{39}$$

$$g^X_{j,t}: \quad Q_{j,t} - S_{j,t} - X_{j,t} = 0. \tag{40}$$

The intertemporal Lagrangian is:

$$\mathcal{L} = \mathbb{E}_0 \sum_{t=0}^{\infty} \Lambda_t \Big[\Psi_{j,t} + \kappa_{j,t} g^D_{j,t} + \mu_{j,t} g^Q_{j,t} + \lambda_{j,t} g^X_{j,t}\Big]. \tag{41}$$

D.2 First-order conditions

We differentiate (41) with respect to each choice variable and then impose symmetric equilibrium ($P_{j,t} = P_t$, $Q_{j,t} = Q_t$, $S_{j,t} = S_t$, $X_{j,t} = X_t$).

FOC w.r.t. sales $S_{j,t}$. Revenue $\frac{P_{j,t}}{P_t} S_{j,t}$, the demand constraint $g^D_{j,t}$, and the stock split $g^X_{j,t}$ all involve $S_{j,t}$:

$$\frac{P_{j,t}}{P_t} + \kappa_{j,t} - \lambda_{j,t} = 0 \quad \Rightarrow \quad \lambda_{j,t} = \frac{P_{j,t}}{P_t} + \kappa_{j,t}. \tag{42}$$

Under symmetry $P_{j,t}/P_t = 1$, so $\lambda_t = 1 + \kappa_t$.

FOC w.r.t. production $Y_{j,t}$. Only $-\mathcal{C}_t(Y_{j,t})$ and $g^Q_{j,t}$ involve $Y_{j,t}$:

$$-\mathcal{C}'_t(Y_{j,t}) - \mu_{j,t} = 0 \quad \Rightarrow \quad \mu_{j,t} = -mc^p_{j,t}. \tag{43}$$

FOC w.r.t. availability $Q_{j,t}$. Availability enters linearly through $g^Q_{j,t}$ and $g^X_{j,t}$, and also through the demand constraint via the $(Q_{j,t}/Q_t)^{\zeta}$ term. Computing the demand derivative:

$$\frac{\partial}{\partial Q_{j,t}} \left[\left(\frac{Q_{j,t}}{Q_t}\right)^{\zeta} \left(\frac{P_{j,t}}{P_t}\right)^{-\varepsilon} S_t\right] = \zeta \, \frac{S_{j,t}}{Q_{j,t}}, \tag{44}$$

the first-order condition is $\mu_{j,t} + \lambda_{j,t} - \kappa_{j,t} \, \zeta \, S_{j,t}/Q_{j,t} = 0$. Imposing symmetry, using $\mu_t = -mc^p_t$ and $\kappa_t = mc^s_t - 1$ (from (42)), and defining $\theta_t \equiv \zeta S_t/Q_t$ yields the *static stocking condition*:

$$mc^p_t = \theta_t + (1 - \theta_t) \, mc^s_t, \qquad \theta_t \equiv \zeta \frac{S_t}{Q_t} = \zeta \frac{S_t}{S_t + X_t}. \tag{45}$$

FOC w.r.t. end-of-period inventories $X_{j,t}$. Inventories affect (i) carrying costs, (ii) the stock split $g^X_{j,t}$ at t, and (iii) next-period availability $g^Q_{j,t+1}$ because $Q_{j,t+1} = Y_{j,t+1} + X_{j,t}$:

$$0 = \Lambda_t\Big[- \psi(r_t)X_{j,t} - \lambda_{j,t}\Big] + \mathbb{E}_t\Big[\Lambda_{t+1}\Big(- \mu_{j,t+1}\Big)\Big]. \tag{46}$$

Dividing by Λ_t, using $m_{t,t+1} = \Lambda_{t+1}/\Lambda_t$, and substituting $-\mu_{j,t+1} = mc^p_{j,t+1}$ from (43) gives the *dynamic stocking condition* (with $mc^s_t \equiv \lambda_t$):

$$mc^s_t = \mathbb{E}_t[m_{t,t+1}\, mc^p_{t+1}] - \psi(r_t)X_t. \tag{47}$$

D.3 Price-setting FOC and nonlinear NKPC

Price $P_{j,t}$ affects (i) current real revenue, (ii) the demand constraint via the price-elasticity term ε, (iii) the current Rotemberg cost through $\pi_{j,t} = P_{j,t}/P_{j,t-1}-1$, and (iv) the next-period Rotemberg cost because $P_{j,t}$ appears as the lagged price in $\pi_{j,t+1} = P_{j,t+1}/P_{j,t} - 1$.

Step 1: current-period derivatives. The period-t terms in (41) that depend on $P_{j,t}$ contribute:

$$\frac{\partial}{\partial P_{j,t}}\Big(\frac{P_{j,t}}{P_t}S_{j,t}\Big) = \frac{S_{j,t}}{P_t}, \tag{48}$$

$$\frac{\partial}{\partial P_{j,t}}\left[\Big(\frac{Q_{j,t}}{Q_t}\Big)^{\zeta}\Big(\frac{P_{j,t}}{P_t}\Big)^{-\varepsilon}S_t\right] = -\varepsilon\,\frac{S_{j,t}}{P_{j,t}}, \tag{49}$$

$$\frac{\partial}{\partial P_{j,t}}\Big(-\tfrac{\phi}{2}\pi^2_{j,t}S_t\Big) = -\phi\,\pi_{j,t}S_t\frac{1}{P_{j,t-1}}, \tag{50}$$

where (49) uses $S_{j,t} = \Big(\frac{Q_{j,t}}{Q_t}\Big)^{\zeta}\Big(\frac{P_{j,t}}{P_t}\Big)^{-\varepsilon}S_t$ from the demand constraint. The period-t contribution to $\partial\mathcal{L}/\partial P_{j,t}$ is therefore:

$$\Lambda_t\left[\frac{S_{j,t}}{P_t} + \kappa_{j,t}\varepsilon\,\frac{S_{j,t}}{P_{j,t}} - \phi\,\pi_{j,t}S_t\frac{1}{P_{j,t-1}}\right]. \tag{51}$$

Step 2: next-period Rotemberg derivative. Because $P_{j,t}$ is the lagged price entering $\pi_{j,t+1}$:

$$\frac{\partial \pi_{j,t+1}}{\partial P_{j,t}} = -\frac{P_{j,t+1}}{P^2_{j,t}} = -\frac{1+\pi_{j,t+1}}{P_{j,t}}, \tag{52}$$

so the future Rotemberg contribution (discounted back to t) is:

$$\mathbb{E}_t\left[\Lambda_{t+1}\,\phi\,\frac{S_{t+1}}{P_{j,t}}\,\pi_{j,t+1}(1+\pi_{j,t+1})\right]. \tag{53}$$

Step 3: combine and impose symmetry. Setting $\partial\mathcal{L}/\partial P_{j,t} = 0$, adding (51) and (53), dividing by Λ_t, imposing $P_{j,t} = P_t$, $S_{j,t} = S_t$, $P_t/P_{t-1} = 1+\pi_t$, and multiplying by P_t/S_t

yields:

$$0 = 1 + \varepsilon\,\kappa_t - \phi\,\pi_t(1+\pi_t) + \phi\,\mathbb{E}_t\Big[m_{t,t+1}\,\frac{S_{t+1}}{S_t}\,\pi_{t+1}(1+\pi_{t+1})\Big]. \tag{54}$$

Step 4: substitute $\kappa_t = mc^s_t - 1$. From (42) under symmetry, $\kappa_t = \lambda_t - 1 = mc^s_t - 1$. Substituting into (54) gives the *nonlinear Rotemberg pricing condition*:

$$(\varepsilon - 1) = \varepsilon\,mc^s_t - \phi\,\pi_t(1+\pi_t) + \phi\,\mathbb{E}_t\Big[m_{t,t+1}\frac{S_{t+1}}{S_t}\pi_{t+1}(1+\pi_{t+1})\Big]. \tag{55}$$

D.4 Steady state

Consider a deterministic steady state with $\bar{\pi} = 0$, $S_{t+1}/S_t = 1$, and $M_{t+1}/M_t = 1$, so that $m_{t,t+1} = \beta$.

Pricing block. Setting $\bar{\pi} = 0$ in (55) causes all Rotemberg adjustment terms to vanish, leaving:

$$\bar{mc}^s = \frac{\varepsilon - 1}{\varepsilon}. \tag{56}$$

Stocking block. Steady-state availability is $\bar{Q} = \bar{S} + \bar{X}$, so $\bar{\theta} = \zeta\bar{S}/(\bar{S} + \bar{X})$. The static stocking condition (45) at steady state gives:

$$\bar{mc}^p = \bar{\theta} + (1 - \bar{\theta})\bar{mc}^s = \bar{mc}^s + \bar{\theta}(1 - \bar{mc}^s). \tag{57}$$

The dynamic stocking condition (47) at steady state gives:

$$\bar{mc}^s = \beta\bar{mc}^p - \psi(\bar{r})\bar{X} \quad \Rightarrow \quad \beta\bar{mc}^p = \bar{mc}^s + \psi(\bar{r})\bar{X}. \tag{58}$$

Substituting (57) into (58):

$$\psi(\bar{r})\bar{X} = \beta\bar{\theta}(1 - \bar{mc}^s) - \bar{mc}^s(1 - \beta). \tag{59}$$

Existence of $\bar{X} > 0$. Since $\psi(\bar{r}) > 0$, the right-hand side of (59) must be positive. Using $1 - \bar{mc}^s = 1/\varepsilon$ and $\bar{mc}^s = (\varepsilon - 1)/\varepsilon$, this is equivalent to:

$$\beta\bar{\theta} > (\varepsilon - 1)(1 - \beta). \tag{60}$$

The log-linearization of (55) and the substitution of the dynamic stocking condition are carried out in Appendix E, yielding the NKPC-inv (16).

E Log-linearization of the NKPC-inv

E.1 From the Rotemberg condition to the NKPC in mc_t^s

Starting from the nonlinear Rotemberg condition (55), expand each term around $\bar{\pi} = 0$. For the current Rotemberg term, since π_t^2 is second order:

$$-\phi\,\pi_t(1+\pi_t) = -\phi\pi_t - \phi\pi_t^2 \approx -\phi\pi_t. \tag{61}$$

For the forward Rotemberg term, deviations of $m_{t,t+1}$ and S_{t+1}/S_t from their steady-state values multiply π_{t+1} and are second order, so:

$$\phi\,\mathbb{E}_t\Big[m_{t,t+1}\frac{S_{t+1}}{S_t}\pi_{t+1}(1+\pi_{t+1})\Big] \approx \beta\phi\mathbb{E}_t[\pi_{t+1}]. \tag{62}$$

Substitute (61)–(62) into (55) and subtract the steady-state identity $(\varepsilon - 1) = \varepsilon\bar{mc}^s$:

$$0 = \varepsilon(mc_t^s - \bar{mc}^s) - \phi\pi_t + \beta\phi\mathbb{E}_t[\pi_{t+1}]. \tag{63}$$

Using $mc_t^s - \bar{mc}^s \approx \bar{mc}^s\,\widehat{mc}_t^s$ and $\varepsilon\bar{mc}^s = \varepsilon - 1$ yields:

$$\pi_t = \beta\mathbb{E}_t[\pi_{t+1}] + \kappa\,\widehat{mc}_t^s, \qquad \kappa \equiv \frac{\varepsilon - 1}{\phi}. \tag{64}$$

E.2 Log-linearization of θ_t, mc_t^p, mc_t^s

Define steady-state shares $\omega_S \equiv \bar{S}/(\bar{S}+\bar{X})$ and $\omega_X \equiv \bar{X}/(\bar{S}+\bar{X})$, so $\omega_S + \omega_X = 1$.

Log-linearization of θ_t. From $\theta_t = \zeta S_t/Q_t$ with $Q_t = S_t + X_t$, take logs: $\ln\theta_t = \ln\zeta + \ln S_t - \ln Q_t$. Differentiating around the steady state ($\hat{z}_t \equiv \ln Z_t - \ln\bar{Z}$):

$$\hat{\theta}_t = \hat{s}_t - \hat{q}_t.$$

To linearize $\hat{q}_t = \ln Q_t - \ln\bar{Q}$ note that $Q_t = S_t + X_t$ implies $Q_t - \bar{Q} = \omega_S\bar{Q}\,\hat{s}_t + \omega_X\bar{Q}\,\hat{x}_t$ (dividing by $\bar{Q}$), so:

$$\hat{q}_t = \omega_S\hat{s}_t + \omega_X\hat{x}_t.$$

Therefore:

$$\hat{\theta}_t = \hat{s}_t - (\omega_S\hat{s}_t + \omega_X\hat{x}_t) = \omega_X(\hat{s}_t - \hat{x}_t). \tag{65}$$

Static stocking condition. Log-linearize $mc_t^p = \theta_t + (1-\theta_t)mc_t^s$ around $(\bar{\theta}, \bar{mc}^s, \bar{mc}^p)$. Converting to level deviations $mc_t^p - \bar{mc}^p \approx \bar{mc}^p\,\widehat{mc}_t^p$ and likewise for mc_t^s and θ_t, a first-order expansion gives:

$$\bar{mc}^p\,\widehat{mc}_t^p = (1-\bar{\theta})\bar{mc}^s\,\widehat{mc}_t^s + (1-\bar{mc}^s)\bar{\theta}\,\hat{\theta}_t, \tag{66}$$

where the cross-term $(\theta_t - \bar{\theta})(mc_t^s - \bar{mc}^s)$ is second order. Dividing by $\bar{mc}^p$ and defining $a_{ps} \equiv (1-\bar{\theta})\bar{mc}^s/\bar{mc}^p$ and $a_{p\theta} \equiv (1-\bar{mc}^s)\bar{\theta}/\bar{mc}^p$:

$$\widehat{mc}_t^p = \underbrace{\frac{(1-\bar{\theta})\bar{mc}^s}{\bar{mc}^p}}_{a_{ps}} \widehat{mc}_t^s + \underbrace{\frac{(1-\bar{mc}^s)\bar{\theta}}{\bar{mc}^p}}_{a_{p\theta}} \hat{\theta}_t. \tag{67}$$

Note $a_{ps} + a_{p\theta} = 1$ follows from $\bar{mc}^p = (1-\bar{\theta})\bar{mc}^s + \bar{\theta}$.

Dynamic stocking condition. Start from the dynamic stocking condition:

$$mc_t^s = \mathbb{E}_t[m_{t,t+1}\, mc_{t+1}^p] - \psi(r_t)\, X_t, \tag{68}$$

where $m_{t,t+1} \equiv \beta M_{t+1}/M_t$ is the one-step stochastic discount factor (SDF), with $M_t \propto u'(C_t)$ the household marginal utility; the household discounts future profits at $m_{t,t+1}$ because it owns the firms. In steady state $m_{t,t+1} = \beta$. Define the log-deviation of the SDF from its steady-state value:

$$\hat{m}_{t,t+1} \equiv \ln m_{t,t+1} - \ln \beta. \tag{69}$$

Step 1: linearize the SDF-marginal-cost product. Use $m_{t,t+1} \approx \beta(1+\hat{m}_{t,t+1})$ and $mc_{t+1}^p \approx \bar{mc}^p(1+\widehat{mc}_{t+1}^p)$; dropping second-order cross-products gives:

$$\mathbb{E}_t[m_{t,t+1}\, mc_{t+1}^p] \approx \beta\bar{mc}^p\left(1 + \mathbb{E}_t[\hat{m}_{t,t+1} + \widehat{mc}_{t+1}^p]\right).$$

Step 2: linearize the carrying-cost term. With $\eta_r \equiv \bar{r}\psi'(\bar{r})/\psi(\bar{r})$:

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

Step 3: subtract the steady state. The steady-state version of (68) is $\bar{mc}^s = \beta\bar{mc}^p - \psi(\bar{r})\bar{X}$. Subtracting and using $mc_t^s - \bar{mc}^s \approx \bar{mc}^s\,\widehat{mc}_t^s$ yields:

$$\bar{mc}^s\,\widehat{mc}_t^s = \beta\bar{mc}^p\,\mathbb{E}_t[\hat{m}_{t,t+1} + \widehat{mc}_{t+1}^p] - \psi(\bar{r})\bar{X}\Big(\hat{x}_t + \eta_r\hat{r}_t\Big). \tag{70}$$

Dividing through by $\bar{mc}^s$ and defining $b_{sm} \equiv \beta\bar{mc}^p/\bar{mc}^s$ and $b_{sx} \equiv \psi(\bar{r})\bar{X}/\bar{mc}^s$ gives the compact form:

$$\widehat{mc}_t^s = \underbrace{\frac{\beta\bar{mc}^p}{\bar{mc}^s}}_{b_{sm}}\mathbb{E}_t[\hat{m}_{t,t+1} + \widehat{mc}_{t+1}^p] - \underbrace{\frac{\psi(\bar{r})\bar{X}}{\bar{mc}^s}}_{b_{sx}}(\hat{x}_t + \eta_r\hat{r}_t). \tag{71}$$

E.3 Substituting the dynamic stocking condition

The log-linearized dynamic stocking condition (71) gives $\bar{mc}^s\,\widehat{mc}^s_t$ explicitly. Substituting into (64):

$$\pi_t = \beta\mathbb{E}_t[\pi_{t+1}] + \frac{\kappa}{\bar{mc}^s}\Big\{\beta\bar{mc}^p\,\mathbb{E}_t[\hat{m}_{t,t+1} + \widehat{mc}^p_{t+1}] - \psi(\bar{r})\bar{X}(\hat{x}_t + \eta_r\hat{r}_t)\Big\}.$$

Two key identities (verified from the steady-state link $\beta\bar{mc}^p = \bar{mc}^s + \psi(\bar{r})\bar{X}$):

$$\frac{\kappa\beta\bar{mc}^p}{\bar{mc}^s} = \kappa + \underbrace{\frac{\kappa\psi(\bar{r})\bar{X}}{\bar{mc}^s}}_{\eta_x}, \qquad \frac{\kappa\psi(\bar{r})\bar{X}}{\bar{mc}^s} = \eta_x, \qquad \eta_x \equiv \frac{\varepsilon\psi(\bar{r})\bar{X}}{\phi}.$$

Substituting:

$$\pi_t = \beta\mathbb{E}_t[\pi_{t+1}] + (\kappa+\eta_x)\mathbb{E}_t[\hat{m}_{t,t+1} + \widehat{mc}^p_{t+1}] - \eta_x(\hat{x}_t + \eta_r\hat{r}_t). \tag{72}$$

This is the NKPC-inv in deviation-from-steady-state form, equation (16) of the main text.

The natural-gap form of (72) is derived in Appendix H.

F Steady state and calibration of ζ

We work with a deterministic steady state satisfying $\bar{\pi} = 0$ and $m_{t,t+1} = \beta$.

F.1 Pricing block: $\bar{mc}^s$ pinned by the markup

From (56), $\bar{mc}^s = (\varepsilon - 1)/\varepsilon$.

F.2 Stocking block: solving for ζ

Steady-state availability is $\bar{Q} = \bar{S} + \bar{X}$, so:

$$\bar{\theta} = \zeta\frac{\bar{S}}{\bar{S}+\bar{X}}. \tag{73}$$

The static stocking condition at steady state gives:

$$\bar{mc}^p = \bar{\theta} + (1-\bar{\theta})\bar{mc}^s = \bar{mc}^s + \bar{\theta}(1-\bar{mc}^s). \tag{74}$$

The dynamic stocking condition at steady state gives:

$$\bar{mc}^s = \beta\bar{mc}^p - \psi(\bar{r})\bar{X} \quad\Rightarrow\quad \beta\bar{mc}^p = \bar{mc}^s + \psi(\bar{r})\bar{X}. \tag{75}$$

Multiply (74) by β and use (75):

$$\bar{mc}^s + \psi(\bar{r})\bar{X} = \beta\bar{mc}^s + \beta\bar{\theta}(1 - \bar{mc}^s). \tag{76}$$

Rearranging to isolate the carrying-cost term:

$$\psi(\bar{r})\bar{X} = \beta\bar{\theta}(1 - \bar{mc}^s) - \bar{mc}^s(1 - \beta). \tag{77}$$

Solve for $\bar{\theta}$ (using $1 - \bar{mc}^s = 1/\varepsilon$):

$$\begin{aligned} \frac{\beta\bar{\theta}}{\varepsilon} &= \psi(\bar{r})\bar{X} + \frac{\varepsilon - 1}{\varepsilon}(1 - \beta) \\ \bar{\theta} &= \frac{1}{\beta}\Big[\varepsilon\,\psi(\bar{r})\bar{X} + (\varepsilon - 1)(1 - \beta)\Big]. \end{aligned} \tag{78}$$

Substituting $\bar{\theta} = \zeta\bar{S}/(\bar{S} + \bar{X})$ from (73) and solving for ζ:

$$\zeta = \frac{\bar{S} + \bar{X}}{\beta\,\bar{S}}\Big[\varepsilon\,\psi(\bar{r})\bar{X} + (\varepsilon - 1)(1 - \beta)\Big]. \tag{79}$$

F.3 Existence condition and calibration

Since $\psi(\bar{r}) > 0$, the right-hand side of (77) must be positive. Using $1 - \bar{mc}^s = 1/\varepsilon$ and $\bar{mc}^s = (\varepsilon - 1)/\varepsilon$, this requires $\beta\bar{\theta} > (\varepsilon - 1)(1 - \beta)$, which in terms of ζ reads:

$$\beta\,\zeta\,\frac{\bar{S}}{\bar{S} + \bar{X}} > (\varepsilon - 1)(1 - \beta). \tag{80}$$

Given calibration targets $(\bar{X}/\bar{S}, \psi(\bar{r}), \beta, \varepsilon)$, equation (79) pins down ζ uniquely, and (80) is verified after calibration.

G Rebate scheme and linearized feasibility

G.1 Rebate construction

Define the (raw) carrying cost as $IC_t \equiv \frac{1}{2}\psi(r_t)X_t^2$. Compute the three objects needed for a first-order Taylor expansion around $(\bar{r}, \bar{X})$:

$$IC(\bar{r}, \bar{X}) = \tfrac{1}{2}\psi(\bar{r})\bar{X}^2, \tag{81}$$

$$IC_X(\bar{r}, \bar{X}) = \frac{\partial}{\partial X}\tfrac{1}{2}\psi(r)X^2\bigg|_{(\bar{r},\bar{X})} = \psi(\bar{r})\bar{X}, \tag{82}$$

$$IC_r(\bar{r}, \bar{X}) = \frac{\partial}{\partial r}\tfrac{1}{2}\psi(r)X^2\bigg|_{(\bar{r},\bar{X})} = \tfrac{1}{2}\psi'(\bar{r})\bar{X}^2. \tag{83}$$

Define the *rebate* V_t as the first-order Taylor part:

$$V_t \equiv IC(\bar{r}, \bar{X}) + IC_X(\bar{r}, \bar{X})(X_t - \bar{X}) + IC_r(\bar{r}, \bar{X})(r_t - \bar{r}). \tag{84}$$

The rebated remainder is $\widetilde{IC}_t \equiv IC_t - V_t$. Substituting (81)–(83):

$$\widetilde{IC}_t = \tfrac{1}{2}\psi(r_t)X_t^2 - \tfrac{1}{2}\psi(\bar{r})\bar{X}^2 - \psi(\bar{r})\bar{X}(X_t - \bar{X}) - \tfrac{1}{2}\psi'(\bar{r})\bar{X}^2(r_t - \bar{r}). \tag{85}$$

By construction, $\widetilde{IC}_t$ is the Taylor remainder: its value and both first partial derivatives vanish at $(\bar{r}, \bar{X})$,

$$\widetilde{IC}(\bar{r}, \bar{X}) = 0, \quad \left.\frac{\partial \widetilde{IC}}{\partial X}\right|_{(\bar{r},\bar{X})} = 0, \quad \left.\frac{\partial \widetilde{IC}}{\partial r}\right|_{(\bar{r},\bar{X})} = 0,$$

so $\widetilde{IC}_t = O(2)$ around the steady state. The rebate V_t is financed and returned to households lump-sum; it does not enter marginal pricing or stocking conditions.

G.2 Linearized feasibility

Goods allocation in terms of sales is:

$$S_t = C_t + AC_t + \widetilde{IC}_t, \qquad AC_t = \frac{\phi}{2}\pi_t^2 S_t. \tag{86}$$

Since $AC_t = O(\pi_t^2) = O(2)$ around $\bar{\pi} = 0$, and $\widetilde{IC}_t = O(2)$ by construction, the first-order approximation of (86) is simply $S_t \approx C_t$, so:

$$\hat{s}_t \approx \hat{c}_t. \tag{87}$$

Linearized inventory law. The inventory accumulation identity (10) is $X_t = X_{t-1} + Y_t - S_t$. In steady state $\bar{X} = \bar{X} + \bar{Y} - \bar{S}$, so $\bar{Y} = \bar{S}$. Linearizing in level deviations:

$$(X_t - \bar{X}) = (X_{t-1} - \bar{X}) + (Y_t - \bar{Y}) - (S_t - \bar{S}).$$

Converting to log deviations via $Z_t - \bar{Z} \approx \bar{Z}\,\hat{z}_t$:

$$\bar{X}\,\hat{x}_t = \bar{X}\,\hat{x}_{t-1} + \bar{Y}\,\hat{y}_t - \bar{S}\,\hat{s}_t.$$

Dividing by $\bar{X}$ and using $\bar{Y} = \bar{S}$ and (87):

$$\hat{x}_t = \hat{x}_{t-1} + \frac{\bar{S}}{\bar{X}}(\hat{y}_t - \hat{c}_t). \tag{88}$$

In natural gaps: $\tilde{x}_t = \tilde{x}_{t-1} + (\bar{S}/\bar{X})(\tilde{y}_t - \tilde{c}_t)$.

H Natural-gap form and the composite term μ_t

H.1 Natural-gap forms of θ_t, mc_t^p, mc_t^s

Mapping the log-linearized results of Appendix E.2 to natural gaps ($\tilde{z}_t \equiv \hat{z}_t - \hat{z}_t^n$):

Stock-on-shelf share. Applying (65) with $\hat{s}_t \approx \hat{c}_t$ from (87):

$$\tilde{\theta}_t = \omega_X(\tilde{c}_t - \tilde{x}_t).$$

Static stocking condition. With $a_{ps} \equiv (1-\bar{\theta})\bar{mc}^s/\bar{mc}^p$ and $a_{p\theta} \equiv (1-\bar{mc}^s)\bar{\theta}/\bar{mc}^p$ (so $a_{ps} + a_{p\theta} = 1$):

$$\widetilde{mc}_t^p = a_{ps}\widetilde{mc}_t^s + a_{p\theta}\tilde{\theta}_t.$$

Dynamic stocking condition. With $b_{sm} \equiv \beta\bar{mc}^p/\bar{mc}^s$ and $b_{sx} \equiv \psi(\bar{r})\bar{X}/\bar{mc}^s$:

$$\widetilde{mc}_t^s = b_{sm}\mathbb{E}_t[\tilde{m}_{t,t+1} + \widetilde{mc}_{t+1}^p] - b_{sx}(\tilde{x}_t + \eta_r\tilde{r}_t).$$

Production marginal cost. From the household labor supply condition (24) and technology (8):

$$mc_t^p = \frac{\chi}{\alpha}C_t^\sigma A_t^{-1}L_t^{\nu-\alpha+1}. \tag{89}$$

Log-linearizing and mapping to natural gaps via $\tilde{z}_t \equiv \hat{z}_t - \hat{z}_t^n$:

$$\widetilde{mc}_t^p = \sigma\tilde{c}_t + \frac{\nu+1-\alpha}{\alpha}\tilde{y}_t. \tag{90}$$

H.2 Derivation of the natural-gap NKPC-inv

Start from (72) in deviation-from-steady-state form and decompose every hat variable as $\hat{z}_t = \tilde{z}_t + \hat{z}_t^n$, where $\tilde{z}_t \equiv \hat{z}_t - \hat{z}_t^n$ is the natural gap and $\hat{z}_t^n$ denotes the natural allocation. Substituting into each term:

$$\begin{aligned}\pi_t &= \beta\mathbb{E}_t[\pi_{t+1}] \\ &\quad + (\kappa+\eta_x)\mathbb{E}_t\Big[(\tilde{m}_{t,t+1} + \hat{m}_{t,t+1}^n) + (\widetilde{mc}_{t+1}^p + \widehat{mc}_{t+1}^{p,n})\Big] \\ &\quad - \eta_x\left[(\tilde{x}_t + \hat{x}_t^n) + \eta_r(\tilde{r}_t + \hat{r}_t^n)\right].\end{aligned} \tag{91}$$

Collecting all natural-allocation terms ($\hat{z}^n$) into a single composite and defining

$$\mu_t \equiv (\kappa+\eta_x)\mathbb{E}_t[\hat{m}_{t,t+1}^n + \widehat{mc}_{t+1}^{p,n}] - \eta_x(\hat{x}_t^n + \eta_r\hat{r}_t^n), \tag{92}$$

(91) reduces to:

$$\pi_t = \beta\mathbb{E}_t[\pi_{t+1}] + (\kappa+\eta_x)\mathbb{E}_t[\tilde{m}_{t,t+1} + \widetilde{mc}_{t+1}^p] - \eta_x\tilde{x}_t - \eta_x\eta_r\tilde{r}_t + \mu_t, \tag{93}$$

which coincides with (31) of the main text.

H.3 CRRA reduction of the leading term

Under CRRA preferences, the household Euler equation implies $\tilde{m}_{t,t+1} = -\sigma(\tilde{c}_{t+1} - \tilde{c}_t)$. Substituting into $(\kappa + \eta_x)\mathbb{E}_t[\tilde{m}_{t,t+1} + \widetilde{mc}^p_{t+1}]$ and using (90):

$$\begin{aligned}(\kappa + \eta_x)\mathbb{E}_t[\tilde{m}_{t,t+1} + \widetilde{mc}^p_{t+1}] &= (\kappa + \eta_x)\mathbb{E}_t\left[-\sigma\tilde{c}_{t+1} + \sigma\tilde{c}_t + \sigma\tilde{c}_{t+1} + \frac{\nu + 1 - \alpha}{\alpha}\tilde{y}_{t+1}\right] \\ &= (\kappa + \eta_x)\mathbb{E}_t\left[\sigma\tilde{c}_t + \frac{\nu + 1 - \alpha}{\alpha}\tilde{y}_{t+1}\right], \end{aligned} \tag{94}$$

so that the SDF and marginal-cost leads partially cancel, leaving only $\sigma\tilde{c}_t$ and an expectation of $\tilde{y}_{t+1}$. This confirms that, under CRRA, the leading term further reduces to a combination of the contemporaneous consumption gap and the expected future output gap.

I Optimal policy: robustness

The main optimal-policy results of Section 5.4 use a loss function penalizing deviations in inflation and the consumption gap ($L_t = \pi_t^2 + \lambda_c\tilde{c}_t^2$). Here we verify that the key qualitative finding – more inventories make inflation stabilization more attractive – survives when the loss is defined over the output gap:

$$L_t^y = \pi_t^2 + \lambda_y\,\tilde{y}_t^2, \qquad \lambda_y = 0.10. \tag{95}$$

In line with this change of objective, the monetary policy rule used throughout this appendix replaces the consumption gap with the output gap:

$$i_t = \phi_\pi\pi_t + \phi_y\tilde{y}_t. \tag{96}$$

This modifies condition (ix) of the equilibrium definition in Section 5.3; all other equilibrium conditions remain unchanged.

Figure 17 repeats the analysis of Figure 8 under the loss (95). The core result is unchanged: both under commitment and under discretion, optimal policy keeps inflation closer to target when $\bar{X}$ is higher, because the cost-of-carry channel makes inflation stabilization less costly in terms of the output gap.

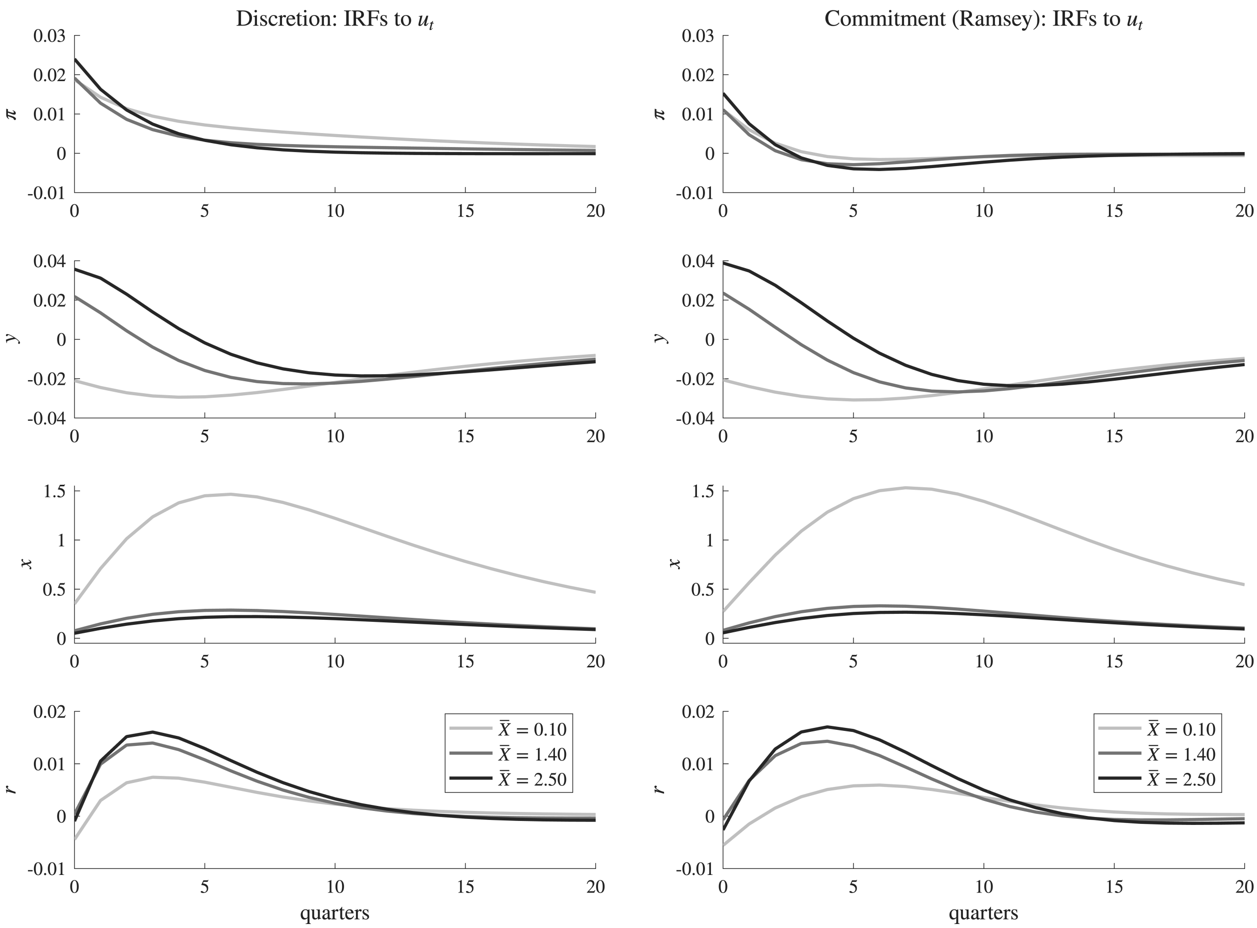

(a) Discretion (b) Commitment

Figure 17: Optimal policy under commitment and discretion for a shock to μ_t, for three levels of steady-state inventories $\bar{X}$. Loss function: $L_t^y = \pi_t^2 + \lambda_y \tilde{y}_t^2$ with $\lambda_y = 0.10$.

J Linearized OSR: robustness

Section 5.4 establishes that the welfare-optimal pure inflation-targeting coefficient is increasing in $\bar{X}$. Here we verify that the same qualitative conclusion holds when the interest rate rule is extended to include the consumption gap:

$$i_t = \phi_\pi \pi_t + \phi_c \tilde{c}_t, \tag{97}$$

and both coefficients (ϕ_π, ϕ_c) are jointly optimized. Figure 18 reports the results. 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).

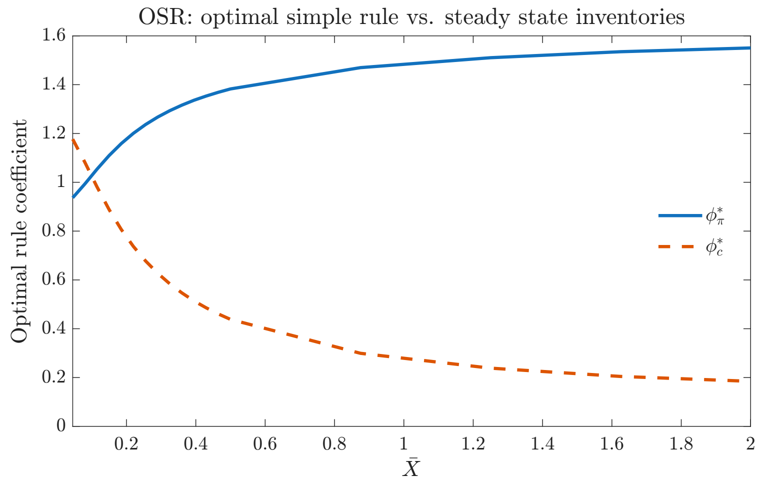

Figure 18: Optimal simple rule coefficients (ϕ_π^*, ϕ_c^*) as a function of steady-state inventories $\bar{X}$. Rule: $i_t = \phi_\pi \pi_t + \phi_c \tilde{c}_t$.

K Non-linear OSR: robustness

Section 5.5 established that the optimal inflation-response coefficient is increasing in the level of inventories under a pure inflation-targeting rule. This appendix shows that the same qualitative conclusion holds when the interest rate rule is extended to include the output gap as an additional argument. In particular, consider the augmented rule:

$$i_t = i_{ss} + \phi_\pi \, \pi_t + \phi_y \left(\frac{Y_t}{Y_{ss}} - 1 \right), \tag{98}$$

where both ϕ_π and ϕ_y are jointly optimized using sequential quadratic programming (SQP). Figure 19 presents the result – again confirming the notion that optimality calls for a greater focus on inflation stabilization when inventories are plentiful.

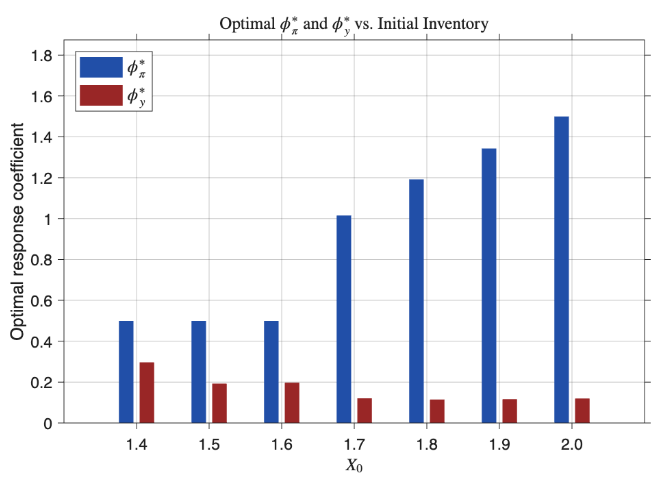

Figure 19: Jointly optimal coefficients under the output-gap Taylor rule (98). Blue bars: ϕ_π^*. Orange bars: ϕ_y^*.