BANK OF CANADA
BANQUE DU CANADA

Staff Working Paper/Document de travail du personnel—2025-30

Last updated: November 3, 2025

Non-homothetic Preferences and the Demand Channel of Inflation

Stephen Murchison
Executive Leadership
Bank of Canada
smurchison@bankofcanada.ca

DOI: https://doi.org/10.34989/swp-2025-30 | ISSN 1701-9397

Acknowledgements

I would like to thank Tom Carter, Don Coletti, José Dorich, Oleksiy Kryvstov and conference participants of the conference Pandemic, Labour Markets and Inflation, held June 30–July 1, 2023, at the University of Siena, Italy, for helpful discussions and feedback. I would also like to extend a special thanks to Alison Collins and Clémence Plourde for excellent research assistance. Any remaining errors are my own. The views expressed herein are my own and do not reflect those of the Bank of Canada or its Governing Council.

Abstract

The post-pandemic rise in global inflation has renewed interest in the relative roles played by demand and supply factors in determining prices. Many central banks have stressed the joint role that persistent increases in input costs and excess demand played in boosting inflation in 2021 and 2022. Yet the latter influence plays no independent role in the workhorse New Keynesian models used by many central banks. Under the typical assumption of constant elasticity of substitution (CES) preferences, variations in consumption shift the firm's profit function up and down, but do not influence its curvature. As a result, the optimal markup is not a function of demand. This assumption is contradicted by both evidence about household shopping behaviour and survey evidence about how firms set prices. This paper proposes an alternative structure based on non-homothetic household preferences over varieties of consumption goods. Specifically, the elasticity of substitution between goods is state dependent, declining during periods of strong per-capita consumption and vice versa. This captures the stylized fact that individual consumers are less price sensitive during economic booms and more price sensitive during downturns. These substitution effects in turn give the firm an incentive to adjust its markup in response to consumption demand. In aggregate, this generates desired markups that increase nonlinearly in consumption demand. When strategic complementarities in pricing are present, these preferences also give rise to state-dependent pass-through of cost shocks.

A New Keynesian sticky price model featuring non-homothetic preferences is estimated on Canadian data, and strong evidence favouring a direct role for per-capita consumption demand is presented. This model is better able to capture the increase in core inflation that occurred in 2021–22, particularly when simulated in its nonlinear form.

Topics: Economic models; Market structure and pricing; Inflation and prices
JEL codes: D1, D4, E3

Résumé

La hausse de l'inflation dans le monde après la pandémie a ravivé l'intérêt pour le rôle relatif des facteurs de demande et d'offre dans la détermination des prix.
De nombreuses banques centrales ont souligné la façon dont les augmentations persistantes du coût des intrants et la demande excédentaire ont toutes les deux contribué à la hausse de l'inflation en 2021 et 2022. Or, l'influence de la demande n'a pas de rôle indépendant dans les modèles de référence de type nouveau keynésien utilisés par beaucoup de banques centrales. Dans l'hypothèse courante des préférences de type CES (à élasticité de substitution constante), les variations de la consommation font

osciller la fonction de profit de l’entreprise de haut en bas, mais n’ont aucun effet sur sa courbure. Par conséquent, la marge bénéficiaire optimale n’est pas fonction de la demande. Cette hypothèse est contredite à la fois par les données sur les habitudes d’achat des ménages et par les résultats d’enquêtes sur la façon dont les entreprises établissent leurs prix. Cette étude propose une structure alternative fondée sur les préférences non homothétiques des ménages entre des variétés de biens de consommation. Plus précisément, l’élasticité de substitution dépend de la conjoncture : elle diminue lorsque la consommation par habitant est forte et augmente dans le cas contraire. Cela correspond à l’exemple stylisé selon lequel les consommateurs sont moins sensibles aux prix en période d’expansion économique, et le sont plus en période de ralentissement économique. Ces effets de substitution incitent à leur tour les entreprises à ajuster leurs marges en fonction de la demande de consommation. Globalement, il en résulte des marges désirées qui augmentent de façon non linéaire par rapport à la demande de consommation. En présence de complémentarités stratégiques dans l’établissement des prix, ces préférences donnent aussi lieu à une transmission des chocs de coûts qui dépend de la conjoncture.

L’auteur estime, sur la base de données canadiennes, un modèle néokeynésien à prix rigides intégrant des préférences non homothétiques et présente de solides données soutenant une approche selon laquelle de la demande de consommation par habitant joue un rôle direct. Ce modèle permet de mieux rendre compte de la hausse de l’inflation fondamentale qui s’est produite en 2021-2022, surtout lorsqu’il est simulé sous sa forme non linéaire.

Sujets : Modèles économiques; Structure de marché et établissement des prix; Inflation et prix
Codes JEL : D1, D4, E3

1 Introduction

Understanding firm-level price-setting behaviour is central to the conduct of monetary policy for inflation-targeting central banks. The post-pandemic rise in global inflation has renewed interest in the relative roles played by demand and supply factors in determining prices. This episode has highlighted the importance of shifts in both the supply and demand curves as inflation sources. Central bank communications have generally highlighted the interaction between pent-up demand following the acute phase of the pandemic and heavily constrained supply as the main factors behind the large and rapid increase in inflation. Based on recent survey evidence for Canada, the Euro Area and the US,[1] *cost pressures, competitors' prices* and the *strength of demand* rank among the top considerations for firms. In addition, the pandemic experience revealed that there can be important interactions between these factors. Specifically, the degree of short-run pass-through of cost increases, though typically low, may increase during periods of strong consumer demand. For instance, markups for customer-facing firms did not decline in Canada during the pandemic (Bilyk, Grieder and Khan (2023)) despite very large increases in costs, suggesting an unusually high degree of short-run pass-through that coincided with strong demand for consumer goods.[2] Using survey evidence from the Bank of Canada's Business Outlook Survey, Asghar, Fudurich and Voll (2023) indicate that the incidence of pass-through of cost increases increased notably in 2021 and 2022, which coincided with strong demand for many consumer goods. Firms noted more favourable conditions to pass through cost increases, including strong demand relative to supply and greater customer acceptance.

Profit maximization implies that a firm's desired markup of price over marginal cost depends on the sensitivity of consumer demand to price—that is, the slope of the demand curve it faces. The more price elastic demand is, the lower the desired markup. This elasticity, in turn, depends in part on the relative availability of close substitutes to the firm's product. Other things being equal, in a marketplace featuring a wide variety of close substitutes and low switching costs, customers will be more price sensitive. Available evidence suggests that other factors can influence the degree of price sensitivity, including the consumer's overall real spending (or income). Specifically, consumers become more price conscious during periods of low *real* spending (high prices or low nominal income), switching to similar but cheaper products. Thus, the elasticity of demand is not necessarily fixed for a given degree of product selection, but instead can vary according to aggregate and market-specific conditions. This helps to rationalize survey evidence indicating that the strength of market demand matters for price setting, separate from its effects on costs.

In contrast, New Keynesian DSGE models used by central banks accord no independent role for demand as a determinant of inflation. Demand forces matter only to the extent that they raise factor prices or reduce productivity, or both.[3] Absent a rise in marginal costs,

[1] See Appendix for additional survey details. Survey results for Canada are taken from Amirault, Kwan and Wilkinson (2006); for the Euro Area, see ECB (2019); and the United States, see Dogra et al. (2023).

[2] During the pandemic, this likely stemmed from a combination of consumers' increased willingness to accept higher prices and low inventories of many products.

[3] In the benchmark New Keynesian model, it is straightforward to show that the deviation of real marginal cost from steady state is proportional to the level of output relative to the flex-price level. However, the key point is that this is simply an equivalent representation of costs—it is not a separate channel for output to

firms have no incentive to raise prices in a high-demand environment or vice versa. This is not to suggest that these factors don't affect the firm's profits. A higher relative price or lower market demand both make the firm less profitable. The key insight is that the firm cannot mitigate this profit loss in the standard model by lowering its price. This stems from the assumption that households' intratemporal substitution elasticity across products is constant. While convenient from a modeller's perspective, this assumption is contradicted by both evidence about household shopping behaviour and survey evidence about how firms set prices.

This paper suggests a simple, tractable modification to household *intratemporal* preferences to allow for non-homotheticity, which simply means that the elasticity of substitution between similar goods is not fixed but rather depends on the household's overall spending. Intuitively, homothetic preferences imply that the shares of different goods purchased by households are independent of household income or overall consumption. Homotheticity implies, for instance, that the share of food purchased is independent of income, contradicting Engel's law and the widespread observation that the share varies inversely with income across countries. In our application, breaking homotheticity means that indifference curves have relatively less (more) curvature at low (high) income/consumption, implying that the relative price between the two goods becomes less important as indifference curves shift out from the origin. Countercyclical substitution elasticities that arise from our specification of preferences give firms an explicit incentive to vary their desired markup in a procyclical manner. When sources of nominal rigidity are introduced, the behaviour of the observed markup will depend on the relative strength of the cost and demand channels.

Non-homothetic preferences (NHP) have mainly been used in partial equilibrium models, since they are useful for explaining the evolution of consumption shares during the process of economic development. Their use in business cycle analysis is much more limited. Cavallari and Etro (2020) (RBC model) and Cavallari (2018) (Rotemburg (1982) costly price adjustment) specify the instantaneous utility function $u(c_{it}) = \gamma c_{it} + \frac{\epsilon}{\epsilon - 1} c_{it}^{\frac{\epsilon-1}{\epsilon}}$, which nests constant elasticity of substitution (CES) preferences but is non-homothetic when $\gamma \neq 0$. Castillo and Montoro (2008) demonstrate that non-homothetic *intertemporal* preferences can help explain the observed asymmetric impact of positive and negative monetary policy shocks.

While many classes of preference functions break homotheticity, so-called implicit CES specifications are particularly flexible as they give rise to demand functions that are easy to work with and do not require the calculation of any additional variables like Lagrange multipliers. Implicit CES preferences can also give rise to interesting interactions between other model features, such as strategic complementarities in pricing behaviour. In this paper, we present a simple model of implicit CES preferences and combine it with the assumption of sticky prices and firm-specific marginal cost. Section 3 provides the details of this setup and presents estimation results of the key model parameters of the linearized system using Canadian data. Section 3.3 explores two key nonlinearities emerging from the model, a Phillips curve that is convex in per-capita consumption, and state-dependent pass-through of marginal cost shocks. Section 4 concludes.

affect inflation.

2 Motivation

2.1 Consumer behaviour: Harrod's Law of Diminishing Elasticity of Demand

In *The Trade Cycle,* Roy Harrod proposed the law of diminishing elasticity of demand, which states that with greater affluence comes a decline in the sensitivity to price changes, and by consequence:

> "The commodity price fluctuation has greater amplitude than that of the money rewards to prime factors" - Harrod (1936: 84)

Since then, many authors have documented a robust relationship between reductions to employment (or hours worked), increased time spent on shopping and lower prices paid. Using data from the *American Time Use Survey,* Aguiar, Hurst and Karabarbounis (2013) document significant increases in time spent shopping during economic downturns in which work hours fall, suggesting that price may play a more significant role in purchase decisions when demand is weak. The authors interpret shopping as a form of non-market work, suggesting that it is not a leisure activity from which individuals derive utility, but rather it is an activity that produces some other benefit, such as lower prices. Such an increased sensitivity to price differences for similar goods could come as the result of lower income, increased future income insecurity or both. Similarly, Aguiar and Hurst (2007) document that older households spend more time shopping and pay lower prices, suggesting that the search for lower prices may be a key motivation for increased shopping time. They estimate that a doubling in shopping frequency lowers a good's price by 7 to 10 percent. Krueger and Muller (2010) find that unemployed people spend between 15 and 30 percent more time shopping than employed people and they pay lower prices. Kaplan and Menzio (2015) find that households with fewer employed members pay lower prices, and they do so by visiting a larger number of stores.

Using a large US dataset of prices and quantities of individual goods, Coibion, Gorodnichenko and Hong (2015) explore differences between posted and "effective" prices, where the latter reflect the prices actually paid by consumers as opposed to posted prices. They find that effective price inflation is much more cyclical than official inflation, primarily because of large substitution effects by consumers toward cheaper products during downturns. Specifically, within each metropolitan area and UPC, the share of items at the bottom of the price distribution rises with the unemployment rate. They also demonstrate a substitution from retailers that charge higher prices to those selling at lower prices for identical goods, and that this substitution effect occurs at the individual household level.

Finally, Auer et al. (2023) present evidence for Switzerland indicating that elasticity of substitution between goods in their dataset is substantially lower for higher-income households.

2.2 Firm behaviour

While the evidence presented in the previous section draws an important link between the business cycle and the price sensitivity of demand, it is also necessary to establish that firms take this into consideration in their pricing behaviour. Survey evidence is certainly suggestive that consumer demand is an important factor in price-setting behaviour, even after controlling for other factors like costs. Based on recent survey evidence for Canada, the Euro Area and the United States, *cost pressures, competitors' prices* and the *strength of demand* rank among the top considerations for firms. For Canada, the strength of demand was found to be particularly important in the services and construction sectors (Amirault, Kwan and Wilkinson (2006)).

Kryvtsov and Vincent (2021) document using micro data for the US and UK that the frequency of "sales" is strongly countercyclical, which they interpret as evidence favouring time-use models in which bargain hunting increases during periods of higher unemployment because the opportunity cost of shopping declines.

Standard models of monopolistic competition with nominal rigidities predict countercyclical observed markups of price over marginal cost, and sources of real rigidity, such as strategic complementarities in pricing, augment this behaviour. Using data on the retail sector for Canada and the US, Anderson, Rebelo and Wong (2018) find that margins are instead largely acyclical or mildly procyclical.

Finally, Cavallo and Krystov (2024) find strong evidence of *cheapflation* for 10 countries, including Canada, during the pandemic. Cheapflation refers to a situation in which the prices of cheaper goods rise more than more expensive goods during a period of high inflation. There are both supply- and demand-side explanations for this observation, including that relative demand for cheaper goods increased as higher prices reduced real income and consumption in certain product segments, and that sellers raised the prices of cheaper goods by more in response. For Canada, they present evidence that expenditure shares did indeed shift from more expensive to cheaper goods, but that the savings were partly offset by higher relative price growth.

3 The Role of Household Preferences in Firm Pricing

With flexible prices, profit maximization by a given firm i implies a possibly time-varying markup of price over marginal cost:

$$P_{i,t} = \frac{\epsilon_t}{\epsilon_t - 1}\lambda_{i,t},$$

where $\lambda_{i,t}$ captures firm-level nominal marginal cost and ϵ_t is the elasticity of demand for the firm's product. Under the standard assumption of CES preferences over varieties, the demand elasticity and hence the flexible-price markup are constant, $\epsilon_t = \epsilon$. In addition, much of the literature makes the simplifying assumption that firm-level marginal cost is independent of its output, meaning all firms face the marginal cost, $\lambda_{i,t} = \lambda_t$. In this case, firm $i's$ profit function can be expressed in terms of aggregate marginal cost, consumption and the price level:

$$\Pi_{i,t} = (P_{i,t} - \lambda_t)\, C_{i,t} = \left(P_{i,t}^{1-\epsilon} - P_{i,t}^{-\epsilon}\lambda_t\right) \times C_t P_t^{\epsilon},$$

where C_t and P_t are, respectively, the sector-wide demand and price level. The curvature of the profit function comes from the terms in $P_{i,t}$. C_t and P_t scale the profit function, but they do not affect the location of its maximum, μ^*. In the example presented in Figure 1 with $\epsilon = 11$, a decline in competitors' prices or lower market demand each make the firm less profitable (a shift from the red to the blue line), but the optimal gross markup remains unchanged at $\mu^* = 1.1$. The key insight is that the firm cannot reduce this profit loss by lowering its price under CES preferences. Lower demand or a decline in competitors' price shift down the entire profit schedule; they don't affect the price elasticity of demand and therefore the curvature of the schedule. CES preferences yield indifference curves (between two differentiated goods) that are homothetic to the origin, meaning the optimal consumption bundle is independent of income or aggregate consumption. As a result, marginal cost is the only variable determinant of prices.

Figure 1: Shifts in demand or competitors' prices don't affect μ^*

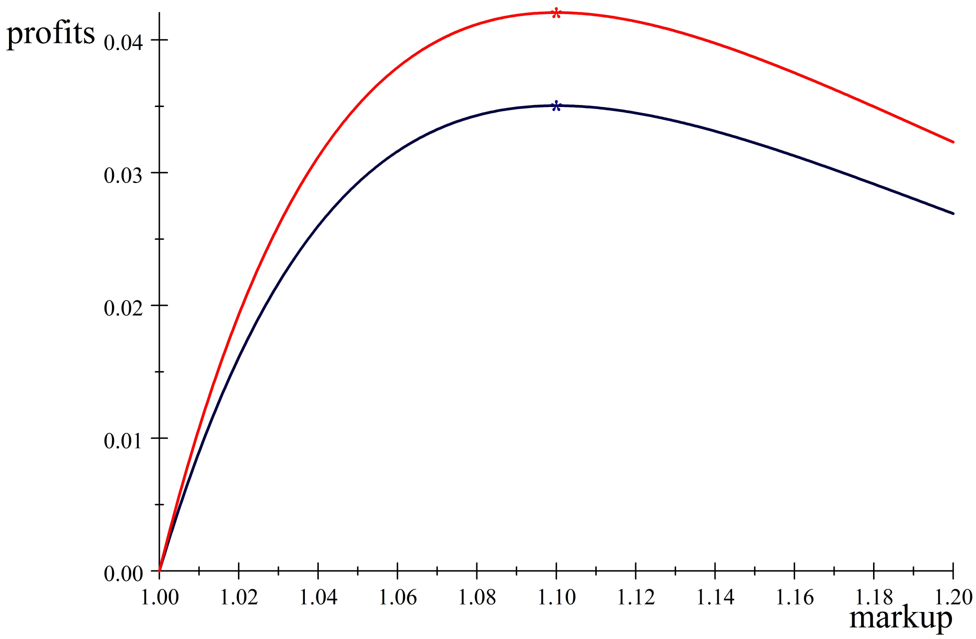

3.1 A Model with Implicit CES Preferences

Consider the non-homothetic utility function for the representative consumer given by the implicit function

$$\int_I \left(\frac{c_{i,t}}{\mathbf{U}_t}\right)^{\frac{\epsilon(\mathbf{U}_t)-1}{\epsilon(\mathbf{U}_t)}} di = 1, \tag{1}$$

where $c_{i,t}$ is consumption of the i^{th} variety of good (produced by firm i) by the representative consumer, $\mathbf{U}_t$ is utility and $\epsilon(\mathbf{U}_t)$ is the elasticity of substitution across goods, which depends on utility. $\epsilon'(\mathbf{U}_t) < 0$ by assumption in our application. According to (1), the elasticity of substitution is common across goods but varies across indifference curves (see Bertoletti and Etro (2020) and Matsuyama (2023) for a detailed discussion). A particularly convenient feature of this form of preferences is that it gives rise to Marshallian demand functions that are identical to that under standard CES except that the demand elasticity depends on the level of consumption.[4] Similar to CES, and in contrast to other forms of preferences (notably the popular aggregator detailed in Kimball (1995)), the demand elasticity does not depend on the firm's relative price, so there are no strategic considerations across firms associated with pricing. Unlike many other non-homothetic specifications, relative demand does not depend on the distribution of prices, but only the aggregate price level, given by:

$$P_t = \left[\int_I P_{i,t}^{1-\epsilon(U_t)} di\right]^{\frac{1}{1-\epsilon(U_t)}}$$

Maximizing (1) subject to the standard budget constraint $\int_I P_{i,t}c_{i,t} - P_t c_t = 0$ yields the demand function for good i:

$$c_{i,t} = \left(\frac{P_{i,t}}{P_t}\right)^{-\epsilon(c_t)} c_t,$$

where c_t is the utility-maximizing basket of consumption goods across all i. In effect, as household spending increases (decreases), the importance of the relative price term $\frac{P_{i,t}}{P_t}$ for the spending share $\frac{c_{i,t}}{c_t}$ diminishes (increases). In a fully articulated model, household consumption would then depend on factors such as wealth, income, real interest rates and the overall price level (affordability).

To gauge the impact of departing from the assumption of homotheticity, a functional form for $\epsilon(c_t)$ is specified such that $\epsilon'(c_t) < 0$. Importantly, it is also required that $\epsilon(c_t) > 1\ \forall c_t$ to ensure that the markup, $\epsilon(c_t)/(\epsilon(c_t) - 1)$, remains finite. Finally, in order to estimate the model (discussed in Section 3.2), the issue of trend growth in per-capita consumption due to productivity gains must be addressed, since it implies a positive trend in markups. At the cost of additional complication, endogenous firm entry could be incorporated, such that higher markups associated with growth in technology encourage the formation of new firms, thereby increasing the variety of goods. Greater variety makes goods closer substitutes, raising ϵ for a given level of consumption. It would be reasonably straightforward to modify the model described here to allow for fixed costs of production, which would eliminate monopoly profits associated with the steady-state markup $\epsilon/(\epsilon-1)$, and then to specify barriers to entry/exit that would make new firm creation/destruction unprofitable for business-cycle variations in markups but profitable for permanent shocks to productivity. With such a setup, markups would remain stationary even in the presence of productivity growth. For the purposes of this paper, the simplifying assumption is made that firm entry is cointegrated with technology, a trend level of consumption is specified, $\bar{c}_t \sim A_t$, such that $c_t/\bar{c}_t$ and hence $\epsilon(c_t)$ is mean stationary.

The generalized logistical function below has the property that $\epsilon(c_t) > 1\ \forall c_t$:

$$\epsilon(c_t) \equiv \frac{2(\epsilon - 1)}{\left(1 + e^{\eta\left(\frac{c_t}{\bar{c}_t} - 1\right)}\right)} + \left(u_t^\epsilon\right)^{-1},\ \eta \geq 0,\ u_t^\epsilon \sim N(1, \sigma_\epsilon^2) \tag{2}$$

where u_t^ϵ is a random shock to the desired markup.[5] In the special case where $c_t = \bar{c}_t$,

[4] Following the usual convention, the utility-based aggregate consumption index is equal to U_t .

[5] Note that the shock appears in inverse form so that a positive shock reduces ϵ_t and therefore increases the desired (flexible-price) markup and inflation.

$\epsilon(c_t) = \epsilon$. The function also nests explicit CES preferences ($\eta = 0$) and possesses a well-defined upper and lower bound, $\epsilon(c) \in (1, 2\epsilon - 1)$. The relationship between the demand elasticity and consumption is plotted in Figure 2.[6]

Figure 2: Demand elasticities are state dependent

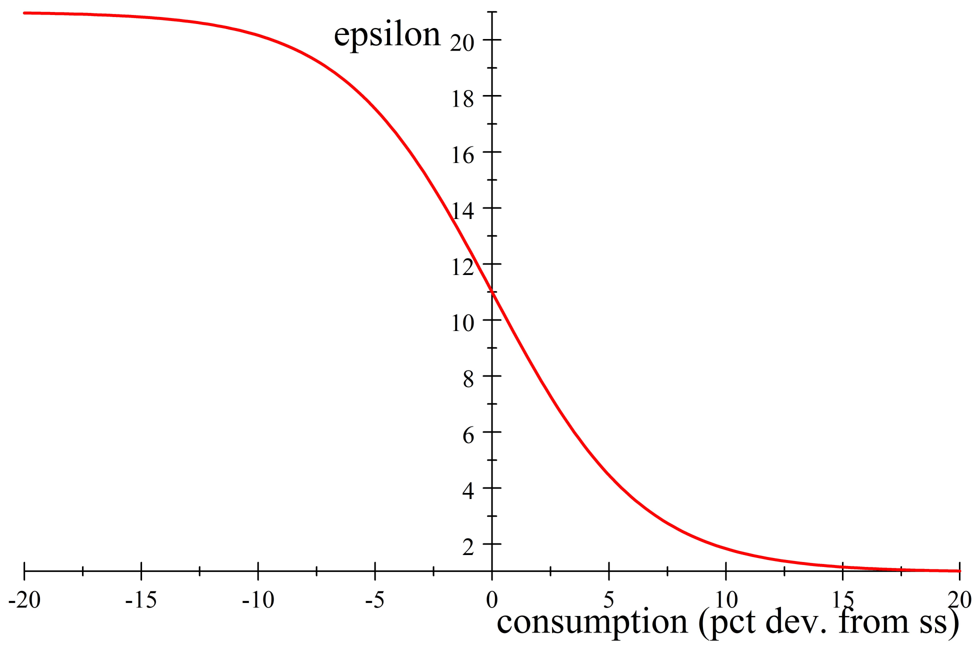

Figure 3 plots the relationship between the share of nominal income spent on the more expensive of two hypothetical goods (vertical axis) and overall consumption spending. In this example, the price of the expensive good is assumed to be 12% higher than the cheaper good. With CES preferences, the household allocates 25% of their nominal income to the more expensive good and the remaining 75% to the cheaper good. Under NHP, the share is also 25% when consumption is equal to its steady-state level but varies considerably depending on overall spending relative to the steady state (horizontal axis). As spending rises, households become less concerned about the price difference (indifference curves become more bow shaped) and the nominal spending share of the expensive good increases. In the limit, households would spend an equal share of their income on each good regardless of any price differential, since the goods increasingly become complements. Likewise, during periods of weak spending, the share spent on the expensive good falls as households become more price conscious. The relationship is not symmetric, however, as strong demand reduces households' sensitivity to relative price differences by more than weak demand increases it. The nominal spending share of any good i is given by $(p_{it}/P_t)^{1-\epsilon(c_t)}$, meaning that as overall consumption increases, the price sensitivity of spending on that good diminishes at an increasing rate in the vicinity of the steady state.

[6]Simulation values are $\epsilon = 11$ and $\eta = 31$, which are taken from Section 3.2.

Figure 3: Relative spending on expensive good rises with overall spending

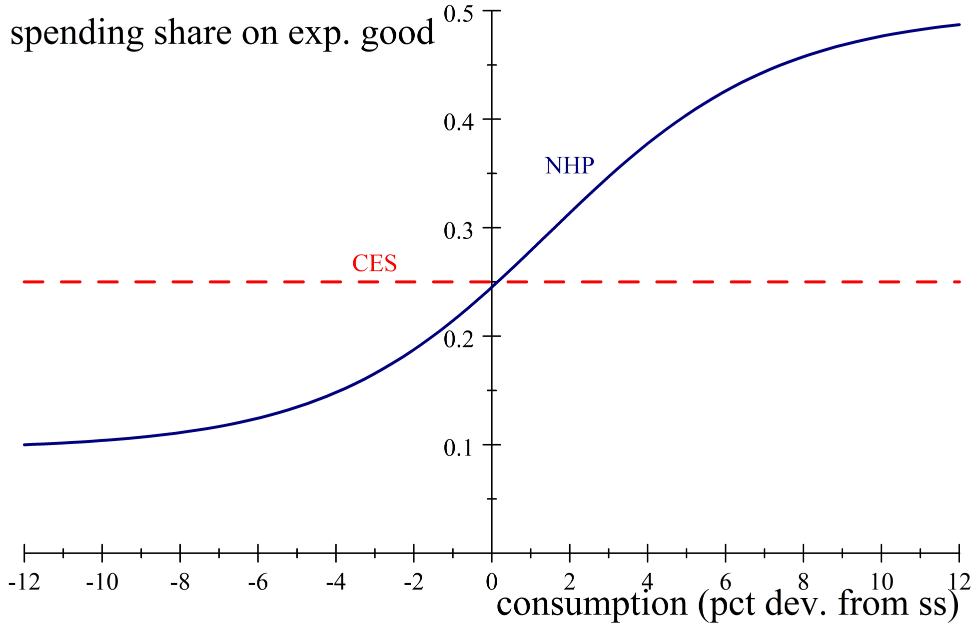

3.1.1 Flexible Prices

Under standard assumptions, the flexible price markup of price over marginal cost in a symmetric equilibrium is given as:

$$\mu_t \equiv \frac{P_t}{\lambda_t} = \frac{\epsilon\left(c_t\right)}{\epsilon\left(c_t\right) - 1},$$

and the elasticity of the markup with respect to consumption is $\frac{\epsilon'(c_t)c_t}{\epsilon(c_t)}\left(1 - \mu_t\right)$. Since $\epsilon'(c_t)$ and $(1 - \mu_t)$ are both negative, the elasticity is positive. Figure 4 shows the relationship between household consumption and the desired markup of a given firm, assuming a steady-state net markup of 10% ($\epsilon = 11$) and using the estimated value $\eta = 31$ from Section 3.2. The estimated historical minimum and maximum of the consumption gap are, respectively, about -6% and 4% over the 1997–2024 sample, suggesting a range for the desired net markup of 6–22%. The desired markup is highly convex in the consumption gap, suggesting that excess demand is significantly more inflationary than excess supply is disinflationary. It's important to note that this convexity does not stem from the particular choice of function form for $\epsilon(c_t)$, which is essentially linear in consumption around the steady state. Rather, it stems from the optimal markup, $\epsilon\left(c_t\right)/\left(\epsilon\left(c_t\right) - 1\right)$, increasing at an increasing rate as ϵ declines.

For a given level of household consumption, pass-through of shocks to economy-wide marginal cost remains complete; only the assumption of nominal rigidity can reduce short-run pass-through. Micro-evidence on price change frequency for Canada suggests that nominal rigidity alone is insufficient to produce empirically plausible short-run pass-through. Estimates of price-change frequency from Bilyk et al. (2024) imply an average nominal contract length of just 4–6 months, though multi-stage value chains could be an additional source of

Figure 4: Desired markup is increasing and convex in consumption

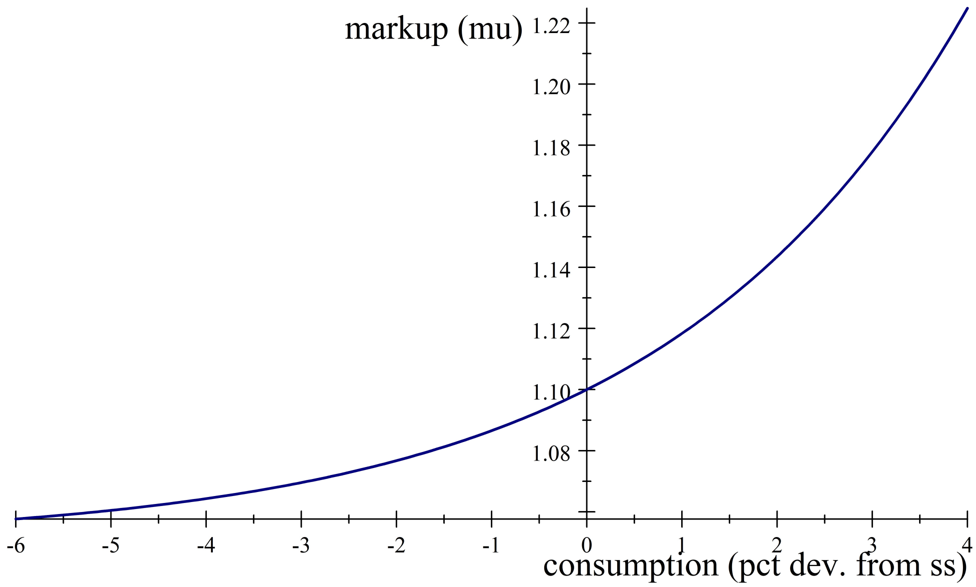

nominal rigidity. Two broad approaches have been suggested to address this shortcoming in sticky-price models. The first allows the demand elasticity to increase with the firm's relative price (firm-specific demand elasticity, FSDE), while the second makes the firm's marginal cost a decreasing function of its relative price (firm-specific marginal cost, FSMC). Both of these mechanisms give firms an incentive to choose a price that is close to that of their competitors. In the context of a model with staggered nominal price contracts, it flattens the Phillips curve and reduces relative price dispersion.

FSDE and FSMC will also generally have higher-order effects and will interact with other aspects of market structure. In our context, these mechanisms will interact in potentially important ways with NHP. Specifically, FSMC will cause the desired level of pass-through of marginal cost to increase when household consumption is strong and desired markups are high, whereas FSDE will have the opposite effect. Since we wish to explore the former possibility, we allow for FSMC by assuming that the level of capital allocated to each firm is equal, and therefore doesn't vary with the firm's relative price or its market share. This has the effect of reducing the returns to scale in the remaining factors, thereby making the marginal product of each factor a declining function of overall production. A firm considering a low relative price, for instance, will take into consideration that doing so will increase its marginal cost, thereby partially offsetting any rise in profits associated with greater market share. In practice, this gives firms an incentive to absorb a portion of variations in aggregate marginal cost, leading to incomplete cost pass-through.

Since the sensitivity of market share, and therefore marginal cost, to the firm's relative price is determined by ϵ, so too is desired pass-through. Since $\epsilon'(c_t) < 0$, desired pass-through of cost shocks to prices will be higher when demand is strong and vice versa. To see this, consider the simple model of firm-specific marginal cost in the spirit Levin et al. (2007), in which the aggregate capital stock is chosen based on aggregate conditions and then equally allocated among individual firms. Specifically, production of the good produced by firm i (c_{it}) is a Cobb-Douglas aggregate of variable inputs (V_{it}) and capital (K_t), with shares α and $1-\alpha$:

$$c_{it} = (V_{it})^{\alpha} (K_t)^{1-\alpha}.$$

The variable input is itself a constant returns-to-scale CES aggregate of labour, raw materials and imported inputs, with marginal cost denoted λ_t^v. Given the returns-to-scale assumption (along with the assumption that factor prices are taken as given), λ_t^v is independent of the level of production. Firm-specific marginal cost can then be expressed as

$$\lambda_{it} = \lambda_t^v \frac{1}{\alpha} \left(\frac{c_{it}}{K_t}\right)^{\frac{1-\alpha}{\alpha}} = \lambda_t \left(\frac{c_{i,t}}{c_t}\right)^{\rho} = \lambda_t \left(\frac{P_{i,t}}{P_t}\right)^{-\rho\epsilon(c_t)} \qquad \rho = \frac{1-\alpha}{\alpha} \geq 0$$

since average marginal cost is simply $\lambda_t = P_t^v \frac{1}{\alpha} \left(\frac{c_t}{K_t}\right)^{\frac{1-\alpha}{\alpha}}$ and $\rho \equiv \frac{1-\alpha}{\alpha} = 0$ corresponds to the standard assumption of full flexibility for all factor inputs. Substituting out firm-specific marginal cost from the firm's FOC yields the following simple expression:

$$P_{it} = \left(\frac{\epsilon(c_t)}{\epsilon(c_t)-1}\lambda_t\right)^{\frac{1}{1+\rho\epsilon(c_t)}} (P_t)^{\frac{\rho\epsilon(c_t)}{1+\rho\epsilon(c_t)}}. \tag{3}$$

The firm's optimal price now represents a weighted geometric average of sector-wide marginal cost and the average price chosen among its competitors, where the weights vary with consumption demand. During periods of strong demand, $\rho\epsilon(c_t)$ declines and cost pass-through increases. This interaction is not present if marginal cost is independent of production/demand, but for $\rho > 0$, it will depend positively on ρ and the sensitivity of $\epsilon(c_t)$ to consumption, $\epsilon'(c_t)$.

When firms' marginal cost is independent of the level of demand they face, desired pass-through remains complete even when preferences are specified as non-homothetic $(\eta > 0)$. This can be seen from equation (3), where the elasticity of the desired price with respect to marginal cost is $\xi(c_t) \equiv (1+\rho\epsilon(c_t))^{-1}$, which reduces to 1 for $\rho = 0$. Allowing for a positive ρ has two effects, one direct and one that results from an interaction with preferences. The direct effect is a reduction in desired pass-through of marginal cost increases to below one, where the reduction is increasing in ρ. Even for a modest degree of diminishing returns to scale at the firm level, pass-through declines significantly when the demand elasticity is high (see Figure 5). The interaction effect is created when the elasticity becomes state dependent. When consumption demand increases, $\epsilon(c_t)$ declines, as does the impact of ρ, since the former multiplies the latter. In effect, firms become less concerned about losing business from a price increase when per-capita demand is strong. For example, when $\rho = 0.25$, pass-through is about 40% higher with a consumption gap of 3% compared to when the gap is zero. Weak demand reduces desired pass-through by less than strong demand increases it, since $\xi''(\bar{c}_t) > 0$.

Figure 5: State-dependent cost pass-through

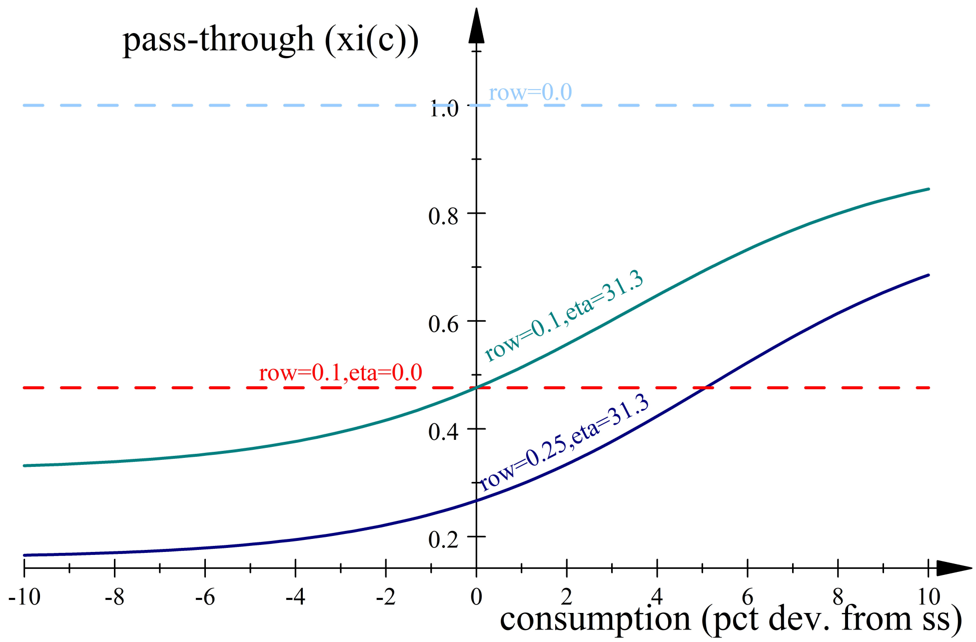

3.1.2 Sticky Prices

In order to study the impact of NHP in a more realistic environment of both nominal and real rigidities, we adopt a variant of the Gali and Gertler (1999) setup in which a fraction $(1-\theta)$ of firms are chosen to reset their price, of which ω follow a simple rule of thumb (ROT) and $(1-\omega)$ choose their price optimally. Unlike in Gali and Gertler's setup, ROT firms can *partially* index to lagged inflation. Optimizing firms do not index to past inflation. The equation governing the average across updated prices in a given period is the geometric average of the prices chosen by ROT $(\widetilde{P}_t^*)$ and optimizing (P_t^*) firms:

$$\overline{P}_t^* = (P_t^*)^{1-\omega} \left(\widetilde{P}_t^*\right)^{\omega}. \tag{4}$$

ROT price-setters base their price on the following:

$$\widetilde{P}_t^* = P_{t-1} \left(\frac{P_{t-1}}{P_{t-2}}\right)^{\gamma}. \tag{5}$$

Finally, the aggregate price level in the sector is given by:

$$P_t = \left[\theta \left(P_{t-1}\right)^{1-\epsilon(c_t)} + (1-\theta)(\overline{P}_t^*)^{1-\epsilon(c_t)}\right]^{\frac{1}{1-\epsilon(c_t)}}. \tag{6}$$

Combining these three expressions and defining $\pi_t \equiv P_t/P_{t-1}$ and $p_t^* \equiv P_t^*/P_t$ allows us to express the relative contract price for optimizing firms in period t as:

$$p_t^* = \left(\frac{1-\theta\left(\pi_t\right)^{\epsilon(c_t)-1}}{(1-\theta)}\right)^{\frac{1}{(1-\epsilon(c_t))(1-\omega)}} \left(\pi_t\right)^{\frac{\omega}{1-\omega}} \left(\pi_{t-1}\right)^{\frac{-\omega\gamma}{1-\omega}}. \tag{7}$$

The first-order condition for optimizing firms is given as:

$$\mathbf{E}_t \sum_{s=t}^{\infty} (\beta\theta)^{s-t} c_s \left((1-\epsilon_s) \left(\frac{p_t^*}{\pi_{s,t}} \right)^{-\epsilon_s} + \tilde{\lambda}_s \epsilon_s \left(\frac{p_t^*}{\pi_{s,t}} \right)^{-(1+\epsilon_s(1+\rho))} \right) = 0, \tag{8}$$

where c_s is per-capita consumption, $\tilde{\lambda}_s$ is average real marginal cost, $\pi_{s,t} \equiv P_{t+s}/P_t$ and β is the discount factor. Unlike in the standard model where ϵ is constant through time, the contract price, p_t^*, cannot be isolated since it is raised to the power $\epsilon_s \equiv \epsilon(c_s)$, which is indexed to s. This also means that the equation cannot be written recursively in its nonlinear form. However, we can approximate (8) around the model's zero-inflation steady state and then write that expression recursively. The resulting linearized NKPC can be expressed compactly as:

$$\hat{\pi}_t = (1-\theta)\gamma\omega\phi^{-1}\hat{\pi}_{t-1} + \beta\theta\phi^{-1}E_t(\hat{\pi}_{t+1}) + \vartheta(\hat{\lambda}_t + \frac{\eta}{2\epsilon}\hat{c}_t + \frac{\mu}{\epsilon^2}\hat{u}_t^{\epsilon}), \tag{9}$$

with

$$\begin{aligned} \phi &\equiv \theta + \omega(1-\theta)(1+\beta\theta\gamma) \\ \vartheta &\equiv \frac{(1-\omega)(1-\theta)(1-\beta\theta)}{\phi(1+\epsilon\rho)}, \end{aligned}$$

where $\mu > 1$ corresponds to the steady-state gross markup of price over marginal cost. As expected, variations in aggregate per-capita consumption spending relative to steady state, $\hat{c}_t$, now enters as a separate driver of inflation in the NKPC with a coefficient $\eta/2\epsilon$ in addition to variations in real marginal cost, $\hat{\lambda}_t$.

Identifying marginal cost requires a number of assumptions about production technology and market structure since it is not directly observable. The most widely used proxy is unit labour costs (ULC), or wages divided by labour productivity. Proportionality between ULC and marginal costs requires an elasticity of substitution of 1 (Cobb-Douglas production technology). When $\sigma < 1$, marginal cost will be more sensitive to changes in factor prices than implied by ULC precisely because firms find it more difficult to substitute away from more expensive inputs.

For some countries, a unit elasticity is a reasonable long-run assumption when nominal factor shares are mean stationary. Estimates of σ for Canada typically range from 0.3 to 0.6 using aggregate data (Perrier (2003)). Dissou, Karnizova and Sun (2014) use industry-level data and find that elasticities for most industries are significantly less than 1, and in many cases below 0.5. As a result, empirical support for a long-run level relationship between consumer prices and nominal ULC is very weak.[7] ULC relative to total and core CPI are non-stationary, with deviations from the mean lasting up to a decade.[8]

We consider a simple production structure featuring four factor inputs that are combined using a constant returns-to-scale CES production technology for the consumption sector:

[7] Another issue is that ULC measures are typically for the overall business sector, not the consumption sector, since there are no readily available cost or producivitity measures for the expenditure components of GDP, such as consumption.

[8] The null hypothesis of a unit root cannot be rejected at conventional levels for either measure of real ULC using the ADF test.

labour (augmented by technology), $A_t L_{it}$; capital; non-energy raw materials, COM_{it}; and imported inputs, M_{it}:

$$C_{it} = \left(\delta_1 (A_t L_{it})^{\frac{\sigma-1}{\sigma}} + \delta_2 COM_{it}{}^{\frac{\sigma-1}{\sigma}} + \delta_3 M_{it}{}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\alpha\sigma}{\sigma-1}} (K_t)^{1-\alpha},$$

where σ is substitution elasticity between variables inputs. For convenience, we assume unitary elasticity between capital, whose level is the same across all forms, and factor inputs chosen at the firm level.[9] Given our assumptions about production technology, combined with the assumption of profit maximization and exogenous factor prices (at the firm level), the nonlinear and the log-linearized (relative to steady state) expressions for real marginal cost are:

$$\tilde{\lambda}_{it} = \alpha^{-\sigma} s_{it} \left[1 - (c_{k,it})^{\rho(1-\sigma)} \left((\delta_2)^\sigma \left(\frac{p_t^{com}}{\tilde{\lambda}_{it}} \right)^{1-\sigma} + (\delta_3)^\sigma \left(\frac{p_t^m}{\tilde{\lambda}_{it}} \right)^{1-\sigma} \right) \right]^{-1} \tag{10}$$

$$\hat{\lambda}_{it} = \frac{1}{1+\varpi_{com}+\varpi_m} \left(\hat{s}_{it} + \rho \left(\varpi_{com} + \varpi_m \right) \hat{c}_{k,it} + \varpi_{com} \hat{p}_t^{com} + \varpi_m \hat{p}_t^m \right), \tag{11}$$

where the ϖ weights depend on nominal factor shares, steady-state capital productivity and the substitution elasticity, σ. $\hat{c}_{k,t}$, $\hat{w}_t$, $\hat{p}_t^{com}$ and $\hat{p}_t^m$ capture percent deviations in the output-to-capital ratio, labour's share of income, the relative price of commodities and imported consumption goods. In the special case where $\sigma = 1$, deviations in real marginal cost are proportional to deviation in labour's share of income, $\hat{s}_t$ ($s_t = W_t L_t / P_t C_t$), since the weights ϖ_{com} and ϖ_m equate to zero.

3.2 Data and Model Estimation

The measure of (log) prices used is CPI-trim (p_t^{cpi}), which excludes CPI components whose rates of change in a given month are located in the tails of the distribution of price changes. This measure helps filter out extreme price movements that might be caused by factors specific to certain components. In particular, CPI-trim excludes 20 percent of the weighted monthly price variations at both the bottom and top of the distribution of price changes, and thus it always removes 40 percent of the total CPI basket. The measure of real consumption is taken from the quarterly National Accounts and divided by population. The import and commodity price series used are, respectively, the monthly fixed-weight measure of the price of consumer imports from the trade accounts and the Bank of Canada commodity price index excluding energy (BCNE). The capital stock measure corresponds to the annual constant dollar fixed non-residential capital stock, which is converted to a quarterly frequency. Nominal unit labour costs are for the total economy and taken from the productivity accounts.

In order to estimate the parameters of (9), we require an estimate of real marginal cost, which is not observable, and therefore we require estimates of the parameters ϖ_{com} and ϖ_m, as well as the capital share parameter ρ. Under the assumption that the markup of price over

[9] Allowing for a Cobb-Douglas structure between variable inputs and capital yields a simple expression relating the ratio of firm-specific marginal to sector-wide average and the ratio of firm-specific to average output.

marginal cost is stationary, the price level and nominal marginal cost should be cointegrated, or equivalently, $\upsilon_t \sim I(0)$ in:

$$p_t^{cpi} = \xi_s s_t + \xi_k c_{k,t} + \xi_{com} p_{com,t} + \xi_m p_{m,t} + \upsilon_t,$$

where the ξ's are composite parameters of the ϖ and ρ parameters in (11). The Johansen rank test and the Engle-Granger approach of testing the stationarity of the residuals of (11) both strongly reject the null of no cointegration (see Table 1). The estimated parameter values are given in Table 2, based on Fully Modified Least Squares estimation of the above equation.

A significant complicating factor associated with this or any other type of empirical analysis involving the pandemic period is that some of the data exhibit very high volatility and compositional changes in 2020. For instance, wages and productivity both increase significantly, but the increases do not fully cancel out, so unit labour costs rise and then decline rapidly. The same is true for the output-to-capital ratio. Such high volatility can distort the estimated parameters of (9) in small samples. The second difficulty is that many price series appear to be correlated with CPI-trim over this short period of time, since just about every price was rising at this time. This increases the risk of a spurious correlation where no causal relationship exists. For these reasons, a stronger test of cointegration is to estimate (11) on pre-pandemic data only. Given that the imported consumption goods price series begins in 1997, we elect to estimate the model using data from 1997Q1 to 2023Q4. Unfortunately, ULC and capital productivity data don't exist for the consumption sector, so we opt to use their economy-wide counterpart as a proxy.

The second requirement for estimation is a model for the level of marginal cost and per-capita consumption deviations in order to generate the expectations that appear in (9). For this, we specify a simple bivariate VAR(1) model. The augmented system therefore consists of (9) and the VAR model, where real marginal cost is defined according to (11) and trend per-capita consumption is treated as an unobserved component subject to growth-rate shocks and estimated using the Kalman filter as part of the system (see Figure 12 in the Appendix). This simple system of difference equations is solved using Sims' (2002) method, which allows for unit roots in the transition matrix, and the model's parameters are estimated by maximum likelihood over the sample 1997Q1–2023Q4. Since θ, μ and ρ cannot be separately identified in estimation using the linearized NKPC, we elect to use the estimate of $\rho = 0.24$ taken from the FMLS estimation of (11) and further impose the steady-state net markup of 10%, $\mu = 1.1$. β is calibrated to be 0.994 using quarterly data. This leaves θ, η, ω and γ to be freely estimated. The point estimates for the parameters can be found in Table 2. The share of ROT price-setters is 0.65 and $\gamma \approx 1$, indicating they fully index to lagged inflation. This relatively high share (and high degree of indexation) is partly due to the choice of inflation series, which displays a high degree of persistence due to the omission of volatile price sub-indices. The average duration of price contracts $(1/(1-\theta))$ is about two quarters $(\theta = 0.53)$ and roughly in line with micro data on price change frequency for Canada.[10]

[10] Bilyk et al. (2024), who use micro CPI data for Canada, estimate that prior to the pandemic, about 25% of prices change each month ($\theta_m = 0.75$). In order to convert this to a quarterly estimate, we make the simplifying assumption of a constant hazard rate at the monthly frequency. Under this assumption, the quarterly hazard rate will be $\theta_m^3 = 0.422$.

Finally, the parameter of greatest interest, η, is estimated to be 31.3 ($\eta/(2\epsilon) = 1.4$), and a likelihood ratio test strongly rejects the restricted model with $\eta = 0$ ($p < 0.0001$). Based on this result, per-capita consumption appears to play an important role in explaining historical inflation dynamics in Canada, even after controlling for its indirect effect on inflation through real marginal cost. Indeed, while our estimate of real marginal cost and the consumption gap are positively correlated over the sample period (maximum correlation is 0.38), there is a substantial amount of variation that is uncorrelated, indicating that there are shocks that cause independent variation in each series. The maximum correlation occurs between current real marginal cost and consumption lagged one year, suggesting that movements in the former tend to significantly lag the latter. During periods of strong demand, inflation initially rises due to a rise in the desired markup and it is subsequently sustained by higher input costs. Overall, the estimated model does a good job of explaining the high degree of persistence in core inflation. The largest root in quarterly CPI-trim inflation is 0.77. The coefficient on lagged inflation in the reduced-form representation of our estimated NKPC is just 0.38, suggesting most of the intrinsic persistence is captured by the two forcing variables.

In order to decompose the relative contribution of each forcing variable, we first note that the reduced-form Phillips curve consistent with (9), after substituting the VAR(1) used for expectations, is given by:

$$\hat{\pi}_t = \varkappa_1 \hat{\pi}_{t-1} + \varkappa_2 \hat{\lambda}_t + \varkappa_3 \hat{c}_t + \varkappa_4 \hat{u}^{\epsilon}_t \approx \left((1-\varkappa_1) \sum_{i=0}^{k} \varkappa_1^i \right)^{-1} \sum_{i=0}^{k} \varkappa_1^i \left(\varkappa_2 \hat{\lambda}_{t-i} + \varkappa_3 \hat{c}_{t-i} + \varkappa_4 \hat{u}^{\epsilon}_{t-i} \right).$$

Fundamental inflation is defined as $\left((1-\varkappa_1) \sum_{i=0}^{k} \varkappa_1^i \right)^{-1} \sum_{i=0}^{k} \varkappa_1^i \left(\varkappa_2 \hat{\lambda}_{t-i} + \varkappa_3 \hat{c}_{t-i} \right)$ for a sufficiently large k. Figure 6 provides the decomposition based on the filtered (not smoothed) estimate of trend per-capita consumption. From this, it is clear that real marginal cost and the consumption gap have each played an important role at different points in history.[11] For instance, most of the variation in fundamental inflation between 2000 and 2006 is explained by real marginal cost. Real marginal cost also played an important role in delaying the decline in inflation following the Great Financial Crisis, partly because wage pressures did not abate in line with the fall in aggregate demand. In contrast, most of the weakness in inflation in the 2000s reflects weak demand, with cost pressures actually providing a small offset. The oil price shock in early 2015 is mainly captured through marginal cost, though consumption demand also declined during this period. The decline in inflation prior to the pandemic is entirely explained by a decline in consumption demand. Finally, the pandemic itself can be broken into two distinct phases. Initially, inflation rose almost exclusively due to a rise in cost pressures, notably higher import and commodity prices. Beginning in the second half of 2021, the rebound in consumption demand contributed to a further increase in inflation. Since then, the decline in inflation is mostly explained by demand factors, as the widening gap between wages and productivity have largely offset declines in the prices of commodities and imported goods.

[11] Figures 12 and 13 in the Appendix provide plots of the real marginal cost and consumption gaps and a comparison of actual versus fundamental inflation.

Figure 6: Decomposition of fundamental inflation

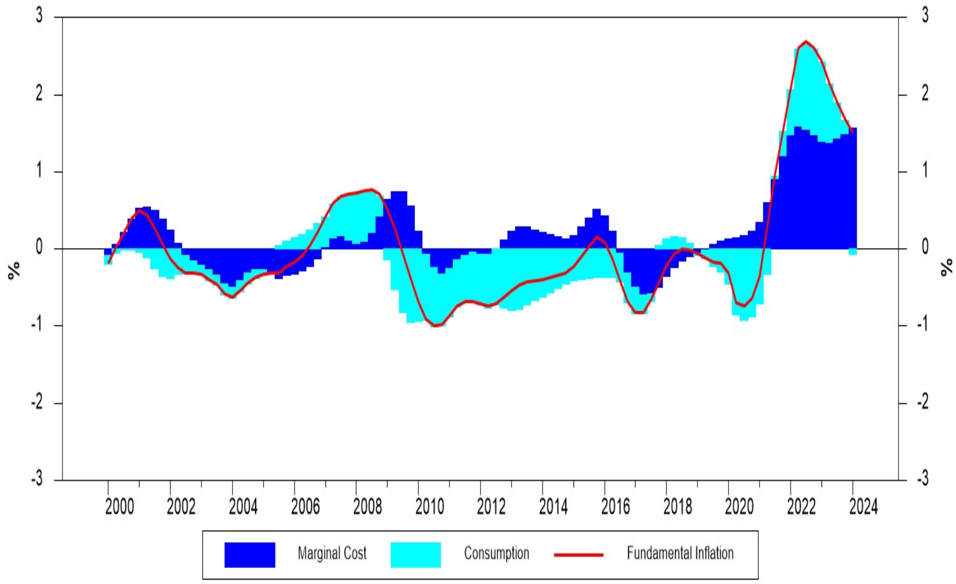

3.3 Nonlinear dynamics

Two potentially important sources of nonlinearity arise from the model presented here. First, inflation is a convex function of the state of per-capita consumption demand, meaning large deviations from steady state will generate a disproportionately larger inflation response. The second source comes about through the interaction between NHP and an upward-sloping firm-level marginal cost function, which gives rise to state-dependent cost pass-through. Neither the convex response of inflation nor the interaction between the state of demand and pass-through can be captured in the linearized NKPC given by (9). In order to explore these nonlinearities, we must first replace the constant hazard rate model of Calvo (1983) with a model featuring a finite maximum nominal contract length, since the former cannot be written recursively when combined with NHP. We follow Dotsey, King and Wolman (1999)[12] and Wolman (1999) and allow for the possibility that firms fix their prices for up to j ($j > 1$) periods (hereafter, this pricing model is referred to as the DKW model). This differs from the original Taylor (1980) specification in that Taylor assumed all firms set price contracts for exactly j periods. It also differs from the Calvo (1983) constant hazard rate specification in that conditional price change probability increases in the time since the last price change and that this probability goes to 1 for finite j. The details of the model can be found in the Appendix, but the additional parameters are a vector of conditional price-change probabilities, α, and the choice of maximum contract length, which in this application is $j = 6$ quarters. Since estimating the parameters of the NKPC using nonlinear methods is beyond the scope of this paper, we instead calibrate the conditional price-change probabilities to match the inflation responses to marginal cost and consumption shocks using the estimated version of (9). This gives rise to the set of conditional probabilities of price

[12]For this version of the model, we exclude the state-dependent component discussed in Dotsey, King and Wolman (1999). Thus, our price-change probabilities are invariant to the state of the economy.

change $\alpha = \begin{bmatrix} 0.13 & 0.18 & 0.25 & 0.36 & 0.5 & 0.71 \end{bmatrix}'$, which implies unconditional probabilities of a given price remaining fixed $\mathbf{\Lambda} = \begin{bmatrix} 0.87 & 0.72 & 0.53 & 0.34 & 0.17 & 0.05 \end{bmatrix}'$, $\Lambda_0 = 1$. The inflation responses of the linearized version of (12) match very closely those from (9), but the hazard rates in the DKW required to generate this equivalence are greater than in the Calvo specification. Using this modified pricing model, the first-order condition for resetting firms is identical to that under Calvo pricing except that the unconditional probability of a price reset need not follow a Poisson process and goes to one at a finite horizon:

$$\mathbf{E}_t \sum_{s=t}^{t+j} \Lambda_{s-t} \beta^{s-t} C_s \left((1-\epsilon_s) \left(\frac{p_t^*}{\pi_{s,t}} \right)^{-\epsilon_s} + \tilde{\lambda}_s \epsilon_s \left(\frac{p_t^*}{\pi_{s,t}} \right)^{-(1+\epsilon_s(1+\rho))} \right) = 0, \tag{12}$$

and (6) becomes:

$$P_t = \left[\sum_{i=0}^{j-1} \varphi_i \left(\overline{P}_{t-i}^* \right)^{1-\epsilon_t} \right]^{\frac{1}{1-\epsilon_t}}$$

The nonlinear dynamics implied by (12) are now analyzed using perfect foresight simulations. Figure 7 plots the peak year-over-year inflation response as a function of the consumption gap for both the linearized and nonlinear DKW models under the assumption of NHP. Large positive shocks to consumption generate substantially more inflation in the nonlinear version relative to the linearized model. The difference is somewhat smaller for large negative shocks. This nonlinearity of the inflation response in consumption arises because the desired markup, $\mu = \epsilon/(\epsilon - 1)$, is nonlinear in ϵ. The relationship between consumption and ϵ is pretty linear for the size of shocks to consumption considered here. The size of the nonlinearity in the vicinity of the model's steady state will also be greater in markets already characterized by low competition and high markups, since $\mu'''(\epsilon) < 0$.

Figure 7: NHP yields a nonlinear NKPC

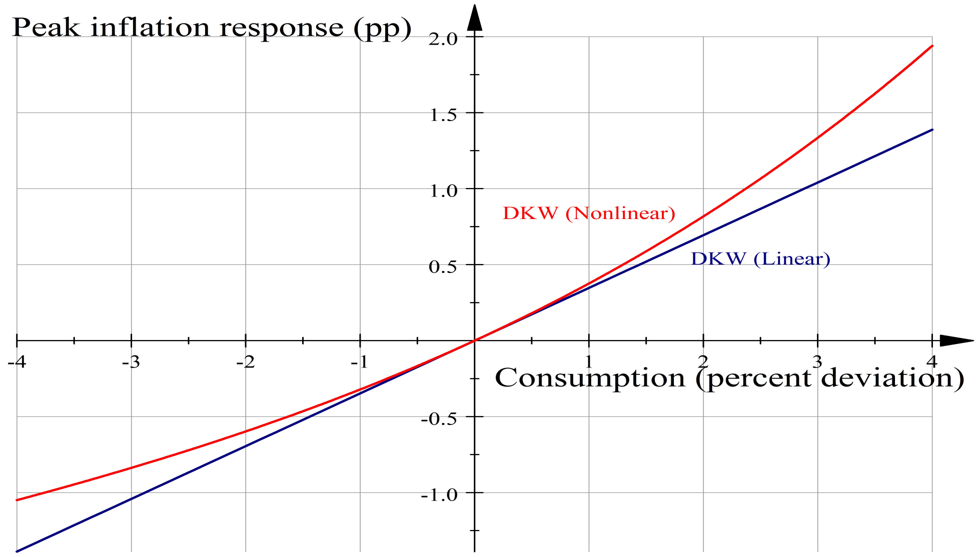

Figure 8: Short-run cost pass-through is higher when demand is strong

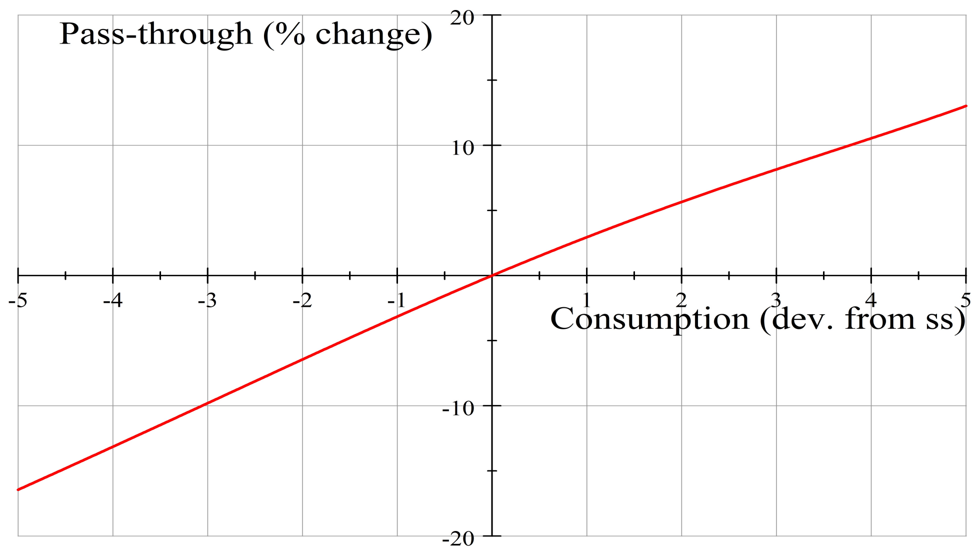

As discussed in the previous section, there is also a positive interaction effect between real marginal cost and per-capita consumption such that short-run pass-through of cost shocks will be greater when $\hat{c}_t > 0$. To gauge the interaction effect, Figure 8 plots the change in the pass-through to prices from a marginal cost shock as a function of consumption, expressed relative to the linearized DKW model with constant pass-through. For instance, when the consumption gap is 3.5%, the optimal short-run pass-through of a cost shock by firms is about 10% higher as a result of ϵ being state dependent. The lower value of ϵ when demand is strong means that the loss in sales to a firm associated with a given price increase is smaller, so the optimal price increase is larger.

4 Conclusion

This paper introduces a new channel through which economic shocks can have first-order effects on price inflation by relaxing the assumption that the elasticity of substitution between goods, and by extension firms' optimal markup, is constant through time. To our knowledge, this paper is the first to introduce non-homothetic preferences into a sticky, staggered-price contracts model. Non-homothetic preferences are motivated by empirical evidence that (1) consumers shop around more during economic downturns and consequently pay lower prices and (2) consumer demand is an important independent factor considered by firms when setting prices. Strong empirical support is found for per-capita consumption as a separate determinant of Canadian core inflation in the context of a Calvo-style NKPC that permits rule-of-thumb price-setters. In addition, when strategic complementarities are introduced via firm-specific marginal cost, it is demonstrated that short-run pass-through of cost shocks becomes state dependent, increasing during periods of strong consumer spending and vice versa. Finally, non-homothetic preferences give rise to a convex Phillips curve, which arises as a result of the nonlinear relationship between consumers' demand elasticity and firms' optimal markup. Future work will focus on the implications of non-homothetic preferences for optimal monetary policy in a fully specified DSGE model.

References

1. Aguiar, M., and Hurst, E. (2007). *Life-Cycle Prices and Production.* American Economic Review, 97(5), 1533–1559.

2. Aguiar, M., Hurst, E., and Karabarbounis, L. (2013). *Time Use During the Great Recession.* American Economic Review, 103(5), 1664–96.

3. Amirault, D., Kwan, C., and Wilkinson, G. (2006). *A Survey of the Price-Setting Behaviour of Canadian Companies.* Bank of Canada Review.

4. Anderson, E., Rebelo, S. T., and Wong, A. (2018). *Markups Across Space and Time.* NBER Working Paper No. 24434.

5. Auer, R., Burstein, A., Lein, S., and Vogel, J. (2023). *Unequal Expenditure Switching: Evidence from Switzerland.* BIS Working Paper No. 1001.

6. Bertoletti, P. and Etro, F. (2020). *Monopolistic competition with generalized additively separable preferences.* Oxford Economic Papers, Volume 73, Issue 2, April 2021, Pages 927–952.

7. Bilyk, O., Khan, M., and Kostyshyna, O. (2024). *Pricing Behaviour and Inflation During the COVID-19 Pandemic.* Bank of Canada Staff Analytical Note 2024-6.

8. Calvo, G. A. (1983). *Staggered Prices in a Utility-Maximizing Framework.* Journal of Monetary Economics, 12(3), 383–398.

9. Castillo, P., and Montoro, C. (2008). *The Asymmetric Effects of Monetary Policy in General Equilibrium.* Journal of CENTRUM Cathedra, 1(2), 28–46.

10. Cavallari, L. (2018). *Monetary Policy with Non-homothetic Preferences.* MPRA Paper No. 85147.

11. Cavallari, L., and Etro, F. (2020). *Demand, Markups and the Business Cycle.* European Economic Review, 127, 103471.

12. Cavallo, A., and Kryvtsov, O. (2024). *Price Discounts and Cheapflation During the Post-Pandemic Inflation Surge.* NBER Working Paper No. 32626.

13. Coibion, O., Gorodnichenko, Y., and Hong, G. H. (2015). *The Cyclicality of Sales, Regular and Effective Prices.* American Economic Review, 105(3), 993–1029.

14. Dogra, K., et al. (2023). *Estimates of Cost-Price Passthrough from Business Survey Data.* Federal Reserve Bank of New York Staff Reports, No. 1062.

15. Dissou, Y., Karnizova, L., and Sun, Q. (2015). *Industry-Level Econometric Estimates with a Nested CES Production Function.* Atlantic Economic Journal, 43(1), 107–121.

16. Dotsey, M., King, R. G., and Wolman, A. L. (1999). *State-Dependent Pricing and General Equilibrium Dynamics.* Quarterly Journal of Economics, 114(2), 655–90.

17. ECB Economic Bulletin. (2019). Issue 7, Box 4.

18. Gali, J., and Gertler, M. (1999). *Inflation Dynamics: A Structural Econometric Analysis.* Journal of Monetary Economics, 44(2), 195–222.

19. Harrod, R. (1936). *The Trade Cycle: An Essay.*

20. Kaplan, G., and Menzio, G. (2015). *The Morphology of Price Dispersion.* International Economic Review, 56, 1165–1206.

21. Kimball, M. S. (1995). *The Quantitative Analytics of the Basic Neomonetarist Model.* Journal of Money, Credit and Banking, 27(4), 1241–1277.

22. Kryvtsov, O., and Vincent, N. (2021). *The Cyclicality of Sales and Aggregate Price Flexibility.* Review of Economic Studies, 88(1), 334–377.

23. Krueger, A. B., and Mueller, A. (2010). *Job Search and Unemployment Insurance: New Evidence from Time Use Data.* Journal of Public Economics, 94(3–4), 298–307.

24. Levin, A., López-Salido, J. D., and Yun, T. (2007). *Strategic Complementarities and Optimal Monetary Policy.* CEPR Discussion Paper No. 6423.

25. Matsuyama, K. (2023). *Non-CES Aggregators: A Guided Tour.* Annual Review of Economics, 15, 235–265.

26. Perrier, P. (2005). *La Fonction de Production et les Données Canadiennes.* Bank of Canada Working Paper No. 2005-20.

27. Rotemberg, J. J. (1982). *Sticky Prices in the United States.* Journal of Political Economy, 90(6), 1187–1211.

28. Sims, C. A. (2002). *Solving Linear Rational Expectations Models.* Computational Economics, 20, 1–20.

29. Taylor, J. B. (1980). *Aggregate Dynamics and Staggered Contracts.* Journal of Political Economy, 88(1), 1–23.

30. Wolman, A. L. (1999). *Sticky Prices, Marginal Cost, and the Behavior of Inflation.* Federal Reserve Bank of Richmond Economic Quarterly, 85(4).

Appendix: Additional Materials

Table 1: Cointegration test statistics

Test	H_0 : No Coint P-value
ADF Test	-4.4 (-3.5)

Table 2: Estimated cointegration equation $p_t^{cpi} = \xi_s s_t + \xi_k c_{k,t} + \xi_{com} p_{com,t} + \xi_m p_{m,t} + u_t$

Method	Structural param.					
	ξ_s	ξ_k	ξ_m	ξ_{com}	σ	$\alpha \rightarrow (\rho)$
FMOLS	0.84 (0.00)	0.21 (0.00)	0.14 (0.00)	0.03 (0.04)	0.4	0.2 $\rightarrow$ (0.25)

Table 3: Parameter values for (9)

Method	Structural param.				Fit	
	θ	ω	γ	η	R^2	$s.e.$
ML	0.53	0.65	1.0	31.3	0.76	0.00124

Details of the DKW Pricing Model

Let α be a j-dimensional vector in which the ith row, α_i, represents the probability that a firm adjusts its price, conditional on having last adjusted i periods ago by assumption $\alpha_j = 1$. The fraction of firms, φ_i, in a given period that charge prices that were set i periods ago is therefore given by:

$$\varphi_i = (1 - \alpha_i) \cdot \varphi_{i-1} \quad i = 1, 2, ..., j-1, \tag{13}$$

or

$$\Lambda_i \equiv \frac{\varphi_i}{\varphi_0} = \prod_{q=0}^{i} (1 - \alpha_q) \qquad \alpha_0 = 0, \tag{14}$$

which states that the probability, Λ_i, of a contract price remaining in effect i periods in the future is equal to the product of the probabilities of not changing prices in each of the preceding periods up to the ith. φ_0 represents the (constant) proportion of firms that adjust their price each period. Since each firm must fall into one of the categories (in terms of the number of periods since their last price change), we have:

$$\varphi_0 = 1 - \sum_{i=1}^{j-1} \varphi_i. \tag{15}$$

The average contract length is then given by φ_0^{-1}.

Figure 9: Main factors affecting firm pricing (Canada)

(1) Triggers / Causes [a]	Total sample		Cons.	Mfg.	Whole-sale retail trade	Info & trans[c]	FIRE[d]	Services[e]
	Mean score[b]	Rankings based on mean score						
Price changes by competitors	3.16*	1	4	1	1	2	1	1
Change in domestic inputs costs (non-labour)	2.90	2	1	2	2	5	3	5
Change in demand for product/ service	2.89*	3	2	3	3	1	2	3
Change in wage costs	2.53*	4	3	5	7	3	6	2
We routinely change prices	2.18	5	7	7	4	4	8	4
Change in taxes, fees, and other charges	2.09	6	6	6	8	8	5	6
Change in economic/inflation forecast	2.01	7	5	9	9	6	4	7
Change in exchange rates	1.87	8	9	4	5	9	9	8
Sales campaigns	1.84	9	8	8	6	7	7	9

Figure 10: Main factors affecting firm pricing (US)

Figure 11: Main factors affecting firm pricing (Euro Area)

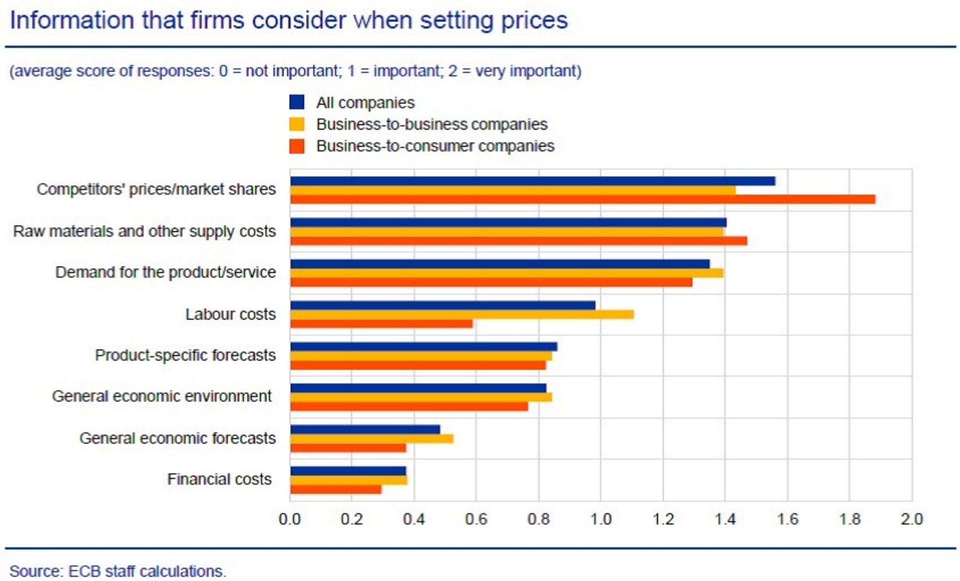

Figure 12: Estimated forcing variables

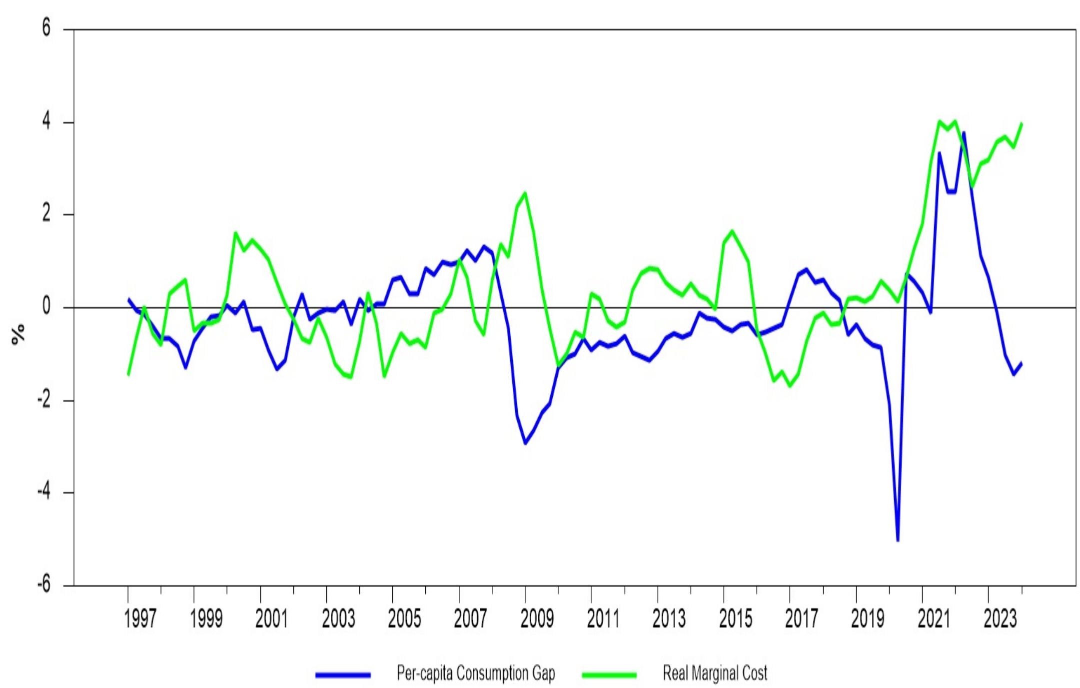

Figure 13: Observed vs. fundamental core inflation