Wage-Price Spirals: The Contract Accelerator

Xilin Yuan *

November 9, 2026

Click here for the latest version

Abstract

I study how endogenous wage indexation generates self-reinforcing wage-price spirals. Leveraging a contract-level dataset, I document that explicit indexation is more prevalent when inflation is high and volatile, while the pass-through from prices to wages is typically incomplete. I explain these patterns through risk-sharing in a general equilibrium model with endogenous indexation. The insurance value of wage indexation rises with inflation volatility. However, greater indexation transmits price increases more quickly to wages, fueling further inflation. This steepens the Phillips curve and amplifies inflation volatility. Together, these forces generate the *contract accelerator*: higher inflation risk encourages indexation, which in turn amplifies inflation risk and further encourages indexation. Consequently, contract choice exhibits strategic complementarity, and the economy may admit multiple self-fulfilling indexation equilibria. A more hawkish policy rule lowers inflation risk and discourages indexation. It also weakens the strategic complementarity and can ensure uniqueness. I provide empirical evidence for key model mechanisms: uncertainty encourages indexation, indexation amplifies inflation, and indexation choices are strategic complements. Quantitatively, the contract accelerator amplifies inflation by 14 percent across the 1973-75 and 1979-81 inflation spikes, relative to a fixed-indexation counterfactual. It also doubles the contribution of the post-Volcker rule to inflation stabilization during the Great Moderation.

*Department of Economics, Yale University. Contact: xilin.yuan@yale.edu. I am deeply grateful to Giuseppe Moscarini, Joel Flynn, and Zhen Huo for their invaluable guidance and unwavering support throughout this project. I thank Andy Atkeson, Yan Bai, Mark Bils, Gideon Bornstein, Laura Castillo-Martinez, Ryan Chahrour, Mishel Ghassibe, Borys Grochulski, Alex Kohlhas, Moritz Kuhn, Ilse Lindenlaub, Jeremy Lise, Zheng Liu, Yang Lu, George Nikolakoudis, Guangyu Pei, Michael Peters, Pascual Restrepo, Martin Schneider, Sanjay Singh, Kjetil Storesletten, Yeji Sung, Nicolas Trachter, Aleh Tsyvinski, Alex Wolman, Chen Yeh, Fabrizio Zilibotti, and other participants at Yale Macro Lunch, HKUST brown bag seminar, CUHK brown bag seminar, FRB Richmond brown bag seminar, FRB San Francisco brown bag seminar for their helpful suggestions. All errors are my own.

1 Introduction

The indexation of nominal wages to prices is more prevalent during inflation spikes. Indexation can be either explicit, through formal cost-of-living adjustment (COLA) clauses in labor contracts (Hendricks and Kahn, 1985), or implicit, through contract reopenings or more frequent wage reviews (Board of Governors of the Federal Reserve System, 2022). These arrangements protect workers' purchasing power from inflation surprises, but they also transmit price increases more quickly to labor costs, adding to inflationary pressure. Hence, they have long raised concerns that indexation may generate a wage-price spiral.

A vast literature sheds light on this spiral, but almost always treats the *indexation regime*—the prevalence and strength of indexation—as fixed (Barro, 1976, Fischer, 1977, 1983). This omission can be consequential, especially during high-inflation episodes. Inflation dynamics shape private incentives to use and design wage indexation, while the resulting indexation regime in turn changes inflation transmission. This two-way interaction may generate an additional spiral, with implications for both equilibrium outcomes and monetary policy transmission.

In this paper, I develop a theory of this interaction that envisions the indexation regime as a general equilibrium outcome. My main contribution is to identify a new spiral between the wage indexation regime and inflation risk and a new role for monetary policy in weakening it. Higher inflation risk raises private demand for the insurance provided by wage indexation. Greater indexation then amplifies the very inflation risk it insures against. This feedback generates strategic complementarity: expectations of greater indexation by others raise perceived inflation risk, thereby encouraging private agents to use indexation themselves. A more hawkish policy rule lowers inflation risk, weakens the strategic complementarity, and can anchor the economy to a low-indexation regime.

I also contribute contract-level empirical measures of wage indexation and causal evidence on my key theoretical mechanism. Using a novel dataset at the labor-contract level, I observe whether a contract is indexed and, if so, how strongly wages respond to prices. I examine empirically three model implications: (1) aggregate uncertainty raises the adoption of indexation, (2) more indexation amplifies the inflation response to monetary policy shocks, and (3) adoption exhibits strategic complementarity. I also use the evidence to discipline the model and quantify the impact of endogenous indexation. I find that it amplifies inflation by about 14 percent during the two major Great Inflation spikes, relative to holding contract composition fixed, and almost doubles the contribution of monetary policy to the subsequent decline in inflation during the Great Moderation. Thus, inflation changes its own transmission mechanism through private insurance choices, while monetary policy can stabilize inflation not only directly, but also by anchoring the indexation regime that governs its amplification.

Data. I use explicit COLA provisions in collective bargaining agreements to measure wage indexation systematically, providing a window into the broader practice of inflation-dependent wage setting. The data come primarily from contract-level records that cover nearly all major US

collective bargaining agreements between 1974 and 1995. These records identify whether a contract includes COLA and, if so, the formula that governs how strongly wages respond to prices.

Two facts motivate the analysis. First, COLA coverage among workers represented in these agreements changed sharply across inflation regimes. It rose to 60 percent during the 1970s and 1980s, when inflation was high and volatile, and subsequently fell to 20 percent. A separate sample of recent agreements shows renewed adoption during the 2021–22 inflation episode. Second, the pass-through from prices to nominal wages within COLA contracts is systematically incomplete. The average pass-through varies between 80 and 90 percent. These facts motivate treating COLA adoption and pass-through as endogenous choices, rather than fixed and mechanical technological parameters.

Theory. A firm and its workers jointly design a nominal wage contract to maximize their total expected value. A non-contingent contract specifies a preset nominal wage path at signing. A COLA contract allows wages to respond to inflation over the life of the contract, at a fixed cost that captures the cognitive burden of planning around contingent nominal income (Martínez-Marquina et al., 2019, Gabaix and Graeber, 2024). The pair first optimizes the wage terms under each contract type—the nominal wage path under a non-contingent contract and the wage formula under a COLA contract (the intensive margin)—and then adopts COLA if it justifies its fixed cost (the extensive margin). Workers are risk-averse, while the firm faces financial constraints and is reluctant to make large wage payments when revenues are low. However, revenues are not contractible. Thus, inflation can serve as a signal of the firm's revenues, the "pie" to be shared.

The contracting model explains both the extensive margin and the intensive margin through risk sharing. COLA insures workers against inflation surprises and can raise the pair's joint surplus relative to preset wages. Greater inflation volatility increases the gain from COLA and encourages adoption. However, when inflation signals low firm revenues, full COLA insurance would require nominal wages to rise one-for-one with prices precisely when the firm can least afford it. Partial pass-through arises optimally from this tension.

Moving to general equilibrium, I embed the contracting problem in a New Keynesian environment similar to Erceg et al. (2000). However, endogenous contracting introduces three interrelated challenges. First, heterogeneous COLA adoption costs and staggered contracting generate a distribution of outstanding wage contracts, an endogenous aggregate state variable one must track. Second, optimal COLA formulas depend on what inflation reveals about firm revenues through their joint equilibrium process, which in turn depends on those COLA formulas themselves. Third, adoption changes the contract distribution and hence inflation risk, which in turn affects the private gain from COLA and thus adoption itself. These two fixed-point problems remain nonlinear even under a first-order approximation of aggregate dynamics.

Despite these challenges, the model retains a tractable structure: conditional on a path of the contract distribution, aggregate dynamics admit a familiar three-equation representation that incorporates the optimal wage terms, while a contract-choice condition governs the evolution of that distribution.

I fully characterize the core mechanisms in an analytical benchmark with one-period contracts and serially independent shocks. The analysis reveals two nested *spirals*—feedback loops that reinforce an initial change.[1] The first operates under a given contract distribution, which is summarized by the aggregate share of COLA contracts. For a given positive inflation surprise, real wages decline by less (if at all), the larger the COLA share, keeping production costs higher and adding to inflationary pressure. This feedback steepens the Phillips curve, makes inflation more responsive to aggregate shocks, and thus raises inflation volatility. I call this propagation mechanism the *conditional wage-price spiral*. This result formalizes, in a modern framework, the long-standing concern that wage indexation can amplify inflationary shocks.

My main contribution is to identify a second spiral, between the wage indexation regime and inflation risk, when contracts are endogenously chosen. Greater inflation risk encourages COLA adoption. Wider adoption strengthens the conditional wage-price spiral, raising inflation risk and further encouraging adoption. I call this spiral the *contract accelerator*. Because of that, contract choices are strategic complements: a firm and its workers have a stronger incentive to adopt COLA when they expect others to do so. In some cases, the economy can admit multiple self-fulfilling indexation regimes under the same fundamentals, even when equilibrium is unique conditional on a given contract distribution.

Monetary policy acts on both transmission mechanisms. A more hawkish Taylor rule lowers inflation risk through the standard aggregate demand channel, thereby reducing the private gain from COLA. This discourages COLA adoption and weakens the conditional wage-price spiral. A more hawkish rule also weakens the strategic complementarity and dampens the contract accelerator. If sufficiently hawkish, it ensures a unique low-indexation regime.

Model Validation. I provide causal evidence supporting three key model mechanisms. First, innovations in the Jurado et al. (2015) macroeconomic uncertainty index raise the share of newly signed contracts that include COLA, consistent with uncertainty increasing the private value of insurance. Second, state-dependent local projections show that expansionary monetary-policy shocks generate larger inflation responses when the economy begins with a larger COLA share, consistent with the prediction that a larger COLA share amplifies inflation propagation. Third, within the same bargaining pair, the probability that a new agreement adopts COLA increases with the inherited aggregate stock of COLA contracts, instrumented with either long-lagged uncertainty shocks or oil-supply shocks that affected earlier contract choices, but whose direct effects on current bargaining conditions should have dissipated. This reduced-form evidence is consistent with strategic complementarity, through which COLA adoption is self-reinforcing.

[1] A spiral in this sense need not imply explosive dynamics.

Quantification. I calibrate the full model separately to the Great Inflation and the Great Moderation, allowing only the monetary-policy response to inflation and the volatilities of primitive shocks to differ across the two eras. The Great Inflation provides a setting in which heightened inflation risk coincided with a large increase in COLA adoption, while the Great Moderation allows me to study the reversal.

The contract accelerator amplified inflation during the Great Inflation. To quantify the effect, I recover the sequence of structural shocks under which the model reproduces the observed paths of inflation and output. I then feed the same shocks into a counterfactual economy in which the contract composition is fixed at the pre-Great-Inflation level. Relative to this counterfactual, the contract accelerator raised average annualized inflation by 1.10 p.p. in 1973-1975 and 1.02 p.p. in 1979-1981, an amplification of about 14 percent across the two episodes.

The same fundamentals also support another rational expectations equilibrium. Holding fixed the policy rule, the joint process of primitive shocks, the initial contract distribution, and the realized shocks, I find an alternative path along which the COLA share converges to about 95 percent. Inflation at the two historical peaks is then about 40 percent higher than observed. Although this path did not occur, it reveals the economy's vulnerability to the self-reinforcing contract accelerator, especially when monetary policy responds weakly to inflation.

In contrast, the contract accelerator made monetary policy more effective during the Great Moderation. I decompose the decline in inflation's responses to markup and demand shocks, each scaled to one standard deviation in the respective period. On average across the two shocks, the more hawkish post-Volcker rule accounts for about 16 percent of the decline through its direct aggregate-demand effect, and another 16 percent through its effect on contract choice. Lower volatility of aggregate shocks accounts for the remainder. Thus, the contract accelerator gives policy more credit for stabilizing inflation, shedding new light on the policy-versus-luck debate (Bernanke, 2004, Stock and Watson, 2003, Primiceri, 2005).

Welfare and Policy. Private inflation insurance creates an aggregate externality. A bargaining pair values the protection provided by COLA, but does not internalize how its adoption raises inflation risk for others. Workers without COLA because of higher adoption costs remain exposed, and greater inflation fluctuations increase resource misallocation. Under the Great Inflation calibration, these externalities outweigh the private insurance gains: both the average worker and the representative firm owner are worse off in an equilibrium with COLA than in an economy where COLA is prohibited.

Finally, I conduct a policy experiment in which the Taylor-rule coefficient on inflation increases with the aggregate COLA share. Thus, policy responds more aggressively to inflation precisely when COLA amplifies it more. Under the Great Inflation calibration, the optimal policy responds very aggressively to inflation, reflecting the additional stabilization gains from discouraging COLA adoption.

Related Literature. The paper builds on the insurance view of labor contracts (Azariadis, 1978, Azariadis and Cooper, 1985), and on the use of observable prices as a substitute for richer state contingency (Blanchard, 1979). Empirical studies examine COLA adoption, adjustment formulas, and other contract terms across bargaining relationships through this lens (Ehrenberg et al., 1983, Hendricks and Kahn, 1985, Card, 1983, 1986). Recent work studies other channels through which wages can become more responsive to inflation, including more frequent renegotiation, shorter contract durations, and job-to-job transitions (Hurst et al., 2026, Afrouzi et al., 2024, Gödl et al., 2026, Moscarini and Postel-Vinay, 2023, Pilossoph and Ryngaert, 2024). I contribute to this literature in two ways. Empirically, I contribute measurements of adoption, pass-through, and COLA coverage in a broad panel of major US agreements. These measurements allow empirical tests of how inflation risk and others' contract choices affect a bargaining pair's own contract choice, and how the resulting choices affect inflation transmission. Theoretically, my model explains adoption and partial pass-through in the contracting problem and traces their aggregate consequences.

The second literature studies how wage rigidity and indexation shape inflation propagation. Classic analyses establish that indexation changes the transmission of real and nominal shocks (Barro, 1976, Gray, 1976, Fischer, 1977, 1983). Modern DSGE frameworks make the propagation channel operational through staggered wages, prices, and often exogenous indexation parameters (Erceg et al., 2000, Christiano et al., 2005, Smets and Wouters, 2007). Blanchard (1986) and its modern counterpart, Lorenzoni and Werning (2023), study how staggered wage and price setting turns workers' and firms' desired real-wage and markup adjustments into wage-price spirals, taking the wage-setting arrangements as given. My work endogenizes wage indexation, an important component of the wage-setting arrangements, and shows this changes equilibrium outcomes and monetary policy transmission substantially.

The closest papers also endogenize particular aspects of wage indexation. Blanchard (1979) and Alexanian (2025) characterize optimal price-contingent wage formulas and their interaction with aggregate outcomes. Carrillo et al. (2022) study workers' choice of full indexation to past inflation or the central bank's inflation target, while Ball (1988) studies the potential inefficiency of costly indexation adoption in a model with risk-neutral agents. Building on these studies, I endogenize both indexation adoption and the optimal wage formulas in a risk-sharing framework disciplined by contract-level evidence. As just discussed, the adjustment along all these margins that we observe in the data has profound implications for our understanding of business cycles and for optimal monetary policy.

More broadly, the paper connects research on uncertainty (Bloom, 2009, Bloom et al., 2018, Jurado et al., 2015, Flynn et al., 2026a,b) with research on coordination and strategic complementarity (Cooper and John, 1988, Benhabib and Farmer, 1994, Morris and Shin, 2003, Angeletos and La'O, 2010, 2013). Wage-contract choice provides a concrete link between them. The spiral between uncertainty and indexation generates strategic complementarity in contract choice, which shapes the equilibrium indexation regime through which aggregate shocks and monetary policy are transmitted.

Roadmap. The rest of the paper proceeds as follows. Section 2 describes the contract data and motivating evidence. Section 3 lays out the model environment and defines the contracting equilibrium. Section 4 fully characterizes an analytical benchmark and derives the core theoretical results. Section 5 empirically validates the model's key predictions. Section 6 specifies the quantitative environment and algorithm, and then presents the calibration, counterfactuals, and welfare and optimal policy exercises. Section 7 concludes.

2 Data and Motivating Evidence

Wage indexation can be either explicit through written COLA formulas in wage contracts, or implicit through more frequent renegotiations and wage reviews. The latter is difficult to measure systematically across different inflation episodes, and it is difficult to separate the part that responds to inflation from the total observed wage change. I use explicit COLA provisions in collective bargaining agreements as a measurement device: I can observe whether wages are indexed to inflation and if so, the formula that governs the adjustment. I provide some suggestive evidence of state-dependent wage setting in the broader economy at the end of this section.

2.1 Data and Institutional Setting

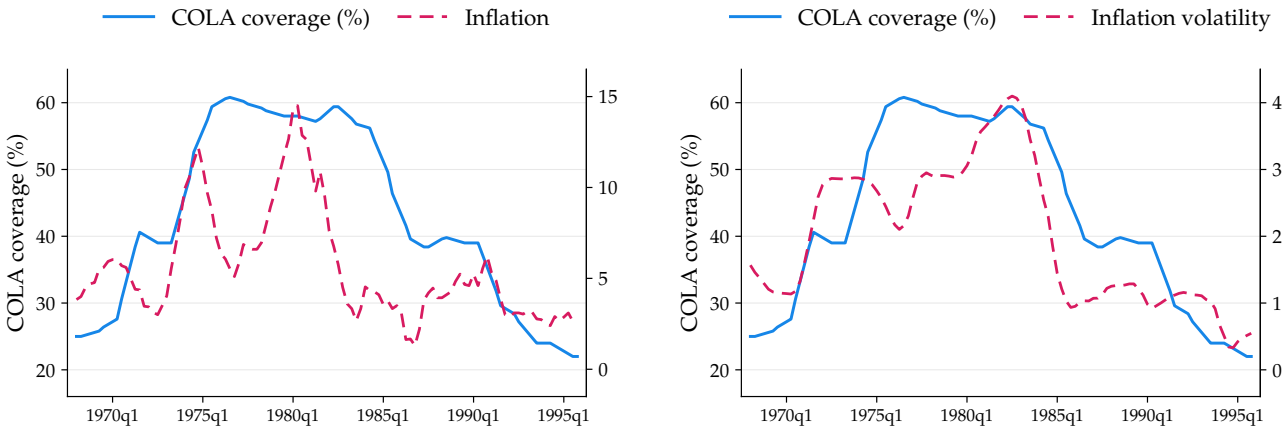
The main data source is the Major Collective Bargaining Settlements, Restricted File (1974-1995), compiled by the Bureau of Labor Statistics.[2] It covers nearly all major US collective bargaining agreements, defined as agreements covering at least 1,000 workers, between 1974 and 1995.[3] The resulting sample contains 21,543 contracts from 3,918 firm-union bargaining pairs in the private and public sectors. For each agreement, I observe its signing and expiration dates, the number of workers covered, scheduled wage changes, whether it contains a COLA clause, and the corresponding formula if it does. I also match firms to Compustat to obtain financial conditions for the matched subsample. Appendix Sections A.1-A.3 describe the archive, sample construction, coding procedure, and some summary statistics.

Collective bargaining agreements specify nominal wage rules over a known contractual horizon. A contract in this sample typically lasts two to three years, and the average duration is 28 months. Among those agreements, two wage-setting arrangements are the most common. The first one is often called scheduled wage increases, which determine a nominal wage path at the time of signing. They may incorporate expected inflation, productivity growth, tenure, and other anticipated developments. The other one is COLA, which allows wages to adjust automatically to prices (mostly CPI) over the life of the contract. Among all the agreements, about 20 percent ever adopted a COLA clause. A typical formula, for example, raises the hourly wage by one cent

[2]The data is now managed by the National Archives and Records Administration (NARA). To my knowledge, this paper is the first one that uses this dataset after it was transferred to NARA. NARA originally only provided a public-use file. I filed a FOIA request and got access to the restricted file. Now, the restricted file is also publicly available.

[3]The file also contains some contracts signed before 1974.

Figure 1: COLA coverage vs. inflation (left) and inflation volatility (right)



Notes: COLA coverage is the fraction of workers with a COLA provision among all workers covered by major collective bargaining agreements in the US. Inflation on the left panel is quarterly inflation, and inflation volatility on the right panel is the standard deviation of inflation from the 12-quarter rolling window.

for each 0.4-point increase in the CPI. Most COLA clauses review wages at a quarterly frequency. As a side note, a few contracts adopted profit-sharing rules between firms and workers, but they were limited to a few giant auto firms in the US (Katz, 1985).

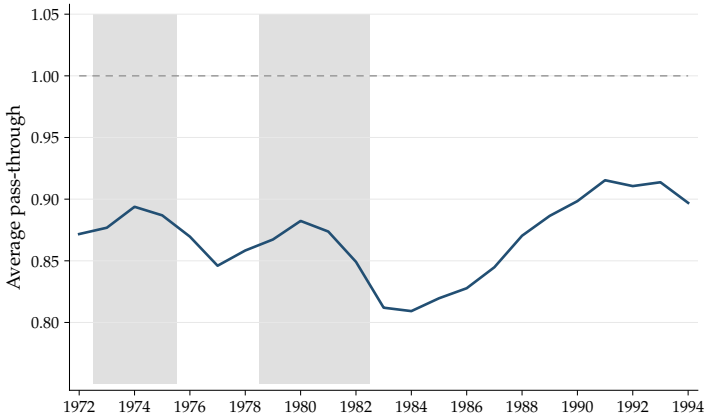
Scheduled wage increases appear in almost all contracts. COLA, if adopted, is layered on top of them. But the two substitute along the inflation-compensation margin. I show in Appendix A.4 that contracts with COLA specify smaller scheduled increases on average than otherwise comparable contracts without COLA, which is consistent with anticipated inflation compensation being shifted from the predetermined schedule to the contingent clause (LeRoy, 1982). This institutional structure motivates a discrete contract choice between a predetermined nominal wage path and inflation contingency along the inflation-compensation margin.

Given the relatively long contract duration, it is also useful to distinguish new COLA adoption from the existing COLA contracts. The share of newly signed contracts that adopt COLA is an inflow, whereas the share of active COLA contracts among all active contracts is a stock inherited from earlier choices.

2.2 COLA Coverage Across Inflation Regimes

I first examine the extensive margin. Figure 1 plots COLA coverage, measured by the share of workers covered by COLA among all workers covered by these agreements, against inflation and inflation volatility. COLA coverage varies sharply across inflation regimes. It rose to approximately 60 percent during the 1970s and early 1980s, when inflation was high and volatile, and then fell to around 20 percent in the 1990s. Because the coverage is calculated within those agreements, it cannot be mechanically explained by the secular decline in the unionization rate in the

Figure 2: Average pass-through of new COLA contracts each year



Notes: COLA pass-through is defined as the percentage change in nominal hourly wage in response to a one-percent increase in the price level implied by the COLA formula. The average is taken over all newly signed contracts with a COLA clause each year. The time series is smoothed using the moving average method on a three-year rolling window.

US. The magnitude of these movements indicates that wage indexation should not be treated as a fixed feature of the economy.

This pattern is not unique to the 1970s-1990s time period. Appendix A.6 shows that during the 2021-2022 inflation episode, explicit COLA adoption in collective bargaining agreements rose in the United States,[4] as well as in other countries like Canada and Spain. Because these agreements are long-lived, the adoption surge also produces a persistent increase in the stock of workers covered by COLA. For instance, the COLA coverage increased from 10 percent to 17 percent in the US during this inflation spike.

2.3 Incomplete COLA Pass-through

The second fact concerns the intensive margin. Figure 2 plots the average pass-through from prices to nominal wages among newly signed COLA contracts each year. COLA pass-through is relatively high: when the price level increases by 1 percent, nominal wages increase by about 0.8-0.9 percent. However, it remains systematically below one, meaning that even workers covered by COLA still retain some exposure to inflation risk.

COLA pass-through also exhibits some time variation. Its first two peaks coincide with the 1973-1975 and 1979-1982 inflation spikes, while the decline in the early 1980s coincides with the Volcker disinflation.[5] Appendix A.9 also shows that in the subsample matched to Compustat, both COLA adoption and COLA pass-through are lower among firms with weaker financial conditions, as measured by lower fixed assets and greater long-term leverage. These relationships

[4]The recent US COLA series comes from a separate source which is only a selected sample of collective bargaining agreements. So I leave it to the appendix.

[5]The increase during the 1990s partly reflects selection: as aggregate COLA coverage declined, the contracts that continued to use COLA were those that tended to specify relatively high pass-through.

Table 1: Inflation pass-through to nominal wages for job-stayers

$$g_t^{W,stayer} = \beta \pi_t^{12m} + \gamma \left(\pi_t^{12m} \times \text{Spike}_t \right) + \delta_{y(t)} + \varepsilon_t$$

Stayer wage growth	(1) OLS	(2) Positive WTI IV
12-month CPI inflation	-0.007	-0.000
	(0.038)	(0.080)
Inflation × Spike	0.176***	0.218**
	(0.061)	(0.094)
Observations	310	300
Year FE	Yes	Yes

Notes: The dependent variable is the Atlanta Fed Wage Growth Tracker measure of year-over-year median nominal wage growth for job stayers. The inflation variable is same-window twelve-month CPI inflation. The spike indicator equals one from January 2021 to December 2022, and the reported interaction coefficient measures the increase in inflation pass-through relative to the pre-2021 period. Year fixed effects are partialled out in all specifications. Standard errors are HAC/Newey-West with a Bartlett kernel and twelve monthly lags. In column (2), twelve-month CPI inflation and its interaction with the spike indicator are instrumented with the cumulative positive WTI price changes over months $t-1$ through $t-12$.

are descriptive, but they suggest that contract choices reflect risks on both sides of the bargaining relationship.

2.4 State-Dependent Wage-Setting Beyond the Unionized Sector

The 2021-2022 inflation episode also provides complementary evidence for the broader labor market. Table 1 reports the regression results where I regress the aggregate year-over-year median nominal wage growth of US job-stayers on the twelve-month CPI inflation of the same window, and its interaction with an indicator of the 2021-2022 inflation spike. The series spans from 1997m3 to 2022m12. The indicator equals one from 2021m1 to 2022m12, and zero prior to 2021. I also instrument inflation using an oil shock, constructed as the sum of positive WTI price changes over the preceding twelve months. I include year fixed effects to control for other macro dynamics. Accordingly, the first row estimates the pass-through of inflation that deviates from the yearly average to aggregate nominal wages, and the second row estimates the additional pass-through during the 2021-2022 inflation spike. For both specifications, the pass-through is almost zero prior to 2021. However, it surges to about 20 percent in the latest inflation episode.

The increase in the aggregate pass-through may come from the adoption of formal COLA clauses as well as more off-cycle wage raises through contract reopenings and renegotiations. The latter is consistent with Hurst et al. (2026), who find a surge of off-cycle wage raises in 2021 and 2022 using the ADP payroll data. Both explicit COLA and off-cycle wage raises increase wage change frequency and thus make nominal wages effectively more responsive to inflation surprises. This is the sense in which I view the latter as implicit indexation.[6]

[6]One caveat is that one cannot identify the underlying reasons for wage adjustments merely from the observed wage changes. Thus, I am not claiming a causal relationship between inflation spikes and more off-cycle wage raises. The

Taken together, these findings suggest that the wage responsiveness to inflation is neither fixed nor mechanically one-for-one. To study its interaction with inflation dynamics, we need a model with endogenous contract choice that can reconcile those facts, to which I now turn.

3 Model

I now develop a model of contract choice in a general equilibrium environment close to Erceg et al. (2000). This section describes the model setup and defines the rational expectations contracting general equilibrium.

3.1 Setup

The economy is a New Keynesian production economy with both wage rigidity and price rigidity. Upstream service firms hire labor from unions under staggered nominal wage contracts, and downstream goods firms transform the service bundle into differentiated varieties subject to Calvo price-setting frictions (Moscarini and Postel-Vinay, 2023). The departure from the standard environment is that service firms and unions jointly design a nominal wage contract in which payments can be potentially contingent on future states. I describe the blocks of the economy in turn.

Upstream Service Firms. There is a continuum of upstream service firms indexed by $i \in [0,1]$. Each of them produces a differentiated service using labor as the only input according to a linear technology:

$$X_{it} = \eta_{it}^{\frac{\varepsilon_X}{\varepsilon_X - 1}} L_{it}$$

where $\eta_{it} \sim \log N(-\sigma_\eta^2/2, \sigma_\eta^2)$ is an idiosyncratic productivity shock with mean one, and $\sigma_\eta > 0$. A competitive CES aggregator combines services into a service bundle X_t:

$$X_t = \left(\int_0^1 X_{it}^{\frac{\varepsilon_X - 1}{\varepsilon_X}} di \right)^{\frac{\varepsilon_X}{\varepsilon_X - 1}}$$

Labor Choice Given Wage. The CES aggregator implies the standard service demand curve

$$P_{it}^X = \left(\frac{X_{it}}{X_t} \right)^{-\frac{1}{\varepsilon_X}} P_t^X$$

latter may arise from the motive of worker retention when worker outside options rise with inflation, e.g., after a large demand shock. This is likely for the 2021-2022 inflation episode. However, during the stagflation period in the 1970s and 1980s, high inflation could coincide with weaker worker outside options, making worker retention less important. In any case, they make wages more responsive to the changing environment, including inflation.

where P_{it}^X is the price set by firm i and is flexible, and P_t^X is the price index for X_t. Denote by $w_{it} = W_{it}/P_t$ the real wage. Combining the production function and the demand for service gives the real operating profit

$$q_{it} = \zeta_t \eta_{it} L_{it}^{1-1/\varepsilon_x} - w_{it} L_{it}$$

where $\zeta_t := \frac{P_t^X}{P_t} X_t^{1/\varepsilon_x}$ is labeled as the "real revenue factor", which is exogenous to the firm and summarizes the aggregate component affecting firm profit. Given the revenue factor and the contractual wage to be specified shortly, firm i chooses labor demand facing uncertainty in its idiosyncratic TFP shock η_{it}. A very low η_{it} realization can lead to a negative operating profit because the firm committed to hire and pay its workers for this period. If that happens, firm i must pay its wage bill by issuing new equity to raise external funds, which incurs a real and convex financing cost $\phi_0 q_{it}^2$ with $\phi_0 > 0$. This is a stylized way to model firm financial frictions (Gilchrist et al., 2017). It makes firms effectively risk-averse, so that they dislike paying high wages especially in low-revenue states. Given (ζ_t, w_{it}), the firm's employment problem is

$$Q(\zeta_t, w_{it}) = \max_{L_{it}} \mathbb{E}^{\eta} \left[q_{it} - \phi_0 \min\{q_{it}, 0\}^2 \right]$$

and I denote by $L^*(\zeta_t, w_{it})$ the optimal labor choice.

Union. Each service firm i hires specialized workers from a firm-specific union i. Union members are immobile, and the size of each union L is sufficiently large that the union can supply the labor demanded by the firm at the contractual wage without labor rationing (Erceg et al., 2000). Given labor demand L_{it}, the union randomly allocates that many workers to production, and each employed worker supplies one unit of labor inelastically. Non-employed members receive home production b. Workers are hand-to-mouth and have concave utility u and discount factor β. Unions are utilitarian, so they care about the total utility received by their workers

$$\mathcal{U}_t(L_{it}, w_{it}) = L_{it} u(w_{it}) + (L - L_{it})\, u(b)$$

Taylor Staggered Contracts. Each service firm i signs a nominal wage contract with union i to maximize the weighted sum of their values, which can be interpreted as an efficient bargaining outcome.[7] Both parties commit to the contract. Contracts last for a fixed number $N \geq 1$ of periods.[8,9] Renewing cohorts are uniformly staggered, so a measure $1/N$ of firms renews at the end

[7] The union's outside option is that all workers get home production, and the firm's outside option is being idle with zero profits. The firm and the union are also subject to their respective participation constraints. The subsequent analysis assumes neither participation constraint binds. This is verified in the calibration ex post.

[8] I use Taylor staggered contracts rather than Calvo wage resetting for both tractability and realism: the fixed-duration assumption maps naturally to collective bargaining agreements, whose contractual durations are known at signing.

[9] In the data, the average duration of newly signed contracts falls modestly from 28 to 26 months during high-inflation periods. This shortening is consistent with the broader pattern of wage setting becoming more responsive to inflation, but the variation on this margin is small. Thus, I hold contract duration fixed and focus on the COLA margin.

of each period. A contract signed at the end of period t specifies a potentially state-contingent nominal wage sequence that applies in periods $t+1, \ldots, t+N$.

Contract Space. For the time-t renewing cohort, denote by $\mathcal{F}_t$ the common information set at signing, and by $\mathcal{I}_{t,s}$ the information set that is observed and contractible when the wage for period $s > t$ is determined. The contract space contains two classes of contracts. A non-contingent class, corresponding to the empirical scheduled-wage contracts, sets the future wage sequence conditional only on information available at contracting, resembling the sticky-information pricing in Mankiw and Reis (2002):

$$\mathcal{C}_t^{\mathcal{N}} = \left\{ \{W_s(.)\}_{s=t+1}^{t+N} : W_s \text{ is } \mathcal{F}_t\text{-measurable} \right\}$$

A contingent contract instead allows W_s to be contingent on future states:

$$\mathcal{C}_t^{\mathcal{C}} = \left\{ \{W_s(.)\}_{s=t+1}^{t+N} : W_s \text{ is } \mathcal{I}_{t,s}\text{-measurable} \right\}$$

Implementing a state-contingent contract incurs a fixed cost χ_i. This reduced-form cost captures the idea that most (if not all) of a worker's other spending is in nominal terms, so planning household spending around contingent nominal income is complex and mentally costly (Martínez-Marquina et al., 2019, Oprea, 2020, Gabaix and Graeber, 2024). The cost is drawn from an exogenous CDF F_χ and is independent across different pairs. Unless otherwise stated, F_χ is twice continuously differentiable with bounded density f_χ on its support $[0, +\infty)$.

The information structure determines which states wages can be contingent on. Inflation risk π_s and revenue risk ζ_s are payoff-relevant. Inflation is publicly observed and contractible during the life of the contract (e.g., published by the BLS). In contrast, unions only learn revenue factors with a delay at the next renewal. Thus, future revenue factors are not contractible for the *current* contract cycle.[10] Hence, for a contract signed at t, the enforceable information sets for the two contract classes are

$$\mathcal{F}_t = \sigma\left(\{\pi_{t-k}, \zeta_{t-k}\}_{k\geq 0}\right), \ \mathcal{I}_{t,s} = \mathcal{F}_t \vee \sigma(\pi_{t+1}, \ldots, \pi_s)$$

where $\sigma(.)$ means the sigma-field spanned by the relevant random variables. Thus, the feasible contingent contract specifies an inflation-contingent wage sequence, which I call a COLA contract.[11]

[10] Workers' delayed observation of firm revenue is one microfoundation for the scarcity of wage indexation to revenues in the real world. There are also other microfoundations; e.g., firms may manipulate their revenue reports once wages are indexed to revenue. The macro implications studied here do not hinge on the particular microfoundations.

[11] I rule out the case where the wage sequence can depend on labor demand. We almost never observe that type of contract in the real world. One possible reason is that, if we did, firms may adopt more labor-substituting methods like outsourcing, automation, etc., which may hurt workers eventually.

The Contracting Problem. The firm-union pair first chooses and evaluates the optimal wage terms within each feasible class, and then adopts COLA if the gain from contingency exceeds its fixed cost. Conditional on the COLA class and $\mathcal{F}_t$, the optimal COLA contract signed at the end of period t specifies N nominal wage functions of contemporaneous inflation history π^s that solve

$$V_t^{\mathcal{C}} = \max_{\{W_s(\pi^s)\}\in\mathcal{C}_t^{\mathcal{C}}} \sum_{s=t+1}^{t+N} \beta^s \mathbb{E}\left[\mathcal{U}_t\left(L^*\left(\zeta_s, \frac{W_s(\pi^s)}{P_s}\right), \frac{W_s(\pi^s)}{P_s}\right) + \lambda Q\left(\zeta_s, \frac{W_s(\pi^s)}{P_s}\right) \middle| \mathcal{F}_t\right]$$

where λ is the weight on firm value in the joint objective.[12] Conditional on the non-contingent class and $\mathcal{F}_t$, the optimal non-contingent contract specifies N nominal wage scalars that solve

$$V_t^{\mathcal{N}} = \max_{\{W_s\}\in\mathcal{C}_t^{\mathcal{N}}} \sum_{s=t+1}^{t+N} \beta^s \mathbb{E}\left[\mathcal{U}_t\left(L^*\left(\zeta_s, \frac{W_s}{P_s}\right), \frac{W_s}{P_s}\right) + \lambda Q\left(\zeta_s, \frac{W_s}{P_s}\right) \middle| \mathcal{F}_t\right]$$

This formulation implies that the firm has the discretion to adjust its labor demand based on the contractual wage each period, which is often called the right to manage (Oswald, 1985). But the contractual wage affects the firm's subsequent labor demand, captured by $L^*(\cdot)$ inside the union payoff.

Let

$$\mathcal{V}_t := V_t^{\mathcal{C}} - V_t^{\mathcal{N}}$$

be the gross gain from the optimal COLA contract relative to the optimal non-contingent contract in period t. For pair i that renews its contract at the end of period t, denote by $\mathbb{I}_{it}^{\mathcal{C}}$ the indicator that equals one if the pair adopts COLA, and zero otherwise. The choice is characterized by

$$\mathbb{I}_{it}^{\mathcal{C}} = \mathbb{I}\left(\mathcal{V}_t > \chi_i\right)$$

Downstream Goods Firms. A continuum of downstream firms, indexed by $f \in [0,1]$, uses the service bundle to produce differentiated varieties:

$$Y_{ft} = X_t$$

A competitive final-good producer aggregates varieties according to

$$Y_t = \left(\int_0^1 Y_{ft}^{\frac{\varepsilon_{yt}-1}{\varepsilon_{yt}}} df\right)^{\frac{\varepsilon_{yt}}{\varepsilon_{yt}-1}}$$

where ε_{yt} is a possibly time-varying elasticity of substitution with mean $\bar{\varepsilon}_y$. Downstream firms face Calvo price-setting frictions: in each period, a firm can reset its nominal price with probability $1-\theta$, while with probability θ its price remains unchanged.

[12] COLA therefore combines time- and state-dependent wage setting: contracting (contract renewal) is time-dependent, while the contract (wage functions) is state-dependent.

Consumers. There are two types of consumers. The hand-to-mouth workers have already been described earlier. A representative entrepreneur, who owns all upstream and downstream firms, can perfectly diversify idiosyncratic productivity risk and price-resetting risk. He trades a one-period nominal bond at gross nominal return $1 + i_t$. He has the same concave utility u as workers and has discount factor $\beta_t = \beta e^{\rho_t}$ in period t, where ρ_t is a demand shock. The entrepreneur's problem is

$$\max_{\{C_t^e, B_{t+1}\}} \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \left(\prod_{s=0}^{t} \beta_s \right) u\left(C_t^e\right) \right]$$
$$s.t. P_t C_t^e + B_{t+1} - (1 + i_t) B_t = \mathcal{R}_t$$

where C_t^e is the entrepreneur's consumption, and $\mathcal{R}_t$ denotes aggregate nominal firm profits.

Central Bank. The model is closed by a central bank that sets the nominal rate according to a standard Taylor rule

$$i_t = i_0 + \phi_\pi \pi_t + \phi_y \hat{y}_t$$

where $\hat{y}_t$ is the log deviation of output from the steady state.

Exogenous Shock Processes. The primitive shocks follow independent AR(1) processes:

$$\begin{aligned} u_t &= \varrho_u u_{t-1} + \varepsilon_{ut} \\ \rho_t &= \varrho_\rho \rho_{t-1} + \varepsilon_{\rho t} \end{aligned} \tag{1}$$

where $u_t = -\frac{1}{\bar{\varepsilon}_y - 1} \hat{\varepsilon}_{yt}$ is a markup shock from the time-varying elasticity ε_{yt}. Their innovations follow independent Gaussian processes:

$$\begin{pmatrix} \varepsilon_{ut} \\ \varepsilon_{\rho t} \end{pmatrix} \sim N\left(0, \begin{pmatrix} \sigma_{ut}^2 & 0 \\ 0 & \sigma_{\rho t}^2 \end{pmatrix}\right)$$

3.2 Equilibrium Definition

I now define a rational expectations contracting general equilibrium. Relative to a standard New Keynesian model, there is an additional aggregate state variable: the distribution of outstanding wage contracts signed by different cohorts. This distribution records not only how many firms are covered by COLA, but also when those contracts were signed, because different cohorts may have chosen different wage terms under different information sets.

Contract Distribution. Denote by $\mathcal{M}_{t,k}$ the measure of firm-union pairs at time t whose contracts were signed at the end of period $t - k$ and whose contract type is COLA, for $k = 1, ..., N$.

Since each renewing cohort has measure $1/N$, $1/N - \mathcal{M}_{t,k}$ is the corresponding measure of non-contingent contracts in the same cohort. By the law of large numbers, under the random COLA cost assumption, the mass of COLA inflow at $t-k$ is

$$\mathcal{M}_{t,k} = \frac{1}{N}\mathbb{P}_{t-k}\left[\mathbb{I}^{\mathcal{C}}_{i,t-k} = 1\right] = \frac{1}{N}F_{\chi}\left(\mathcal{V}_{t-k}\right)$$

The outstanding contract distribution is summarized by the vector

$$\mathcal{M}_t := (\mathcal{M}_{t,1}, \ldots, \mathcal{M}_{t,N})$$

This contract distribution is predetermined in period $t-1$ under the timing assumption.

Equilibrium Definition. Let $\mathcal{Z}_t$ collect the aggregate allocation, prices, interest rate, profits, and exogenous shocks in period t. A *rational expectations contracting general equilibrium* is a collection of stochastic processes

$$\left\{\left\{L_{it}, P^X_{it}, \mathbb{I}^{\mathcal{C}}_{it}\right\}_{i\in[0,1]}, \left\{P_{ft}, P^*_{ft}, Y_{ft}\right\}_{f\in[0,1]}, \left\{W^{\mathcal{C}}_{t,k} \in \mathcal{C}^{\mathcal{C}}_{t-k}, W^{\mathcal{N}}_{t,k} \in \mathcal{C}^{\mathcal{N}}_{t-k}\right\}_{k=1}^{N}, \mathcal{M}_t, \mathcal{Z}_t\right\}_{t\in\mathbb{Z}}$$

such that the following conditions hold.

1. For each renewing cohort, the firm-union pair chooses the contract type according to $\mathbb{I}^{\mathcal{C}}_{it} = \mathbb{I}(\mathcal{V}_t > \chi_i)$. Conditional on the chosen type, it chooses the wage term that solves the corresponding problem defining $V^{\mathcal{C}}_t$ or $V^{\mathcal{N}}_t$.

2. Given the contractual wage and the revenue factor, each service firm chooses labor to solve $Q(\zeta_t, w_{it})$, and service prices and quantities satisfy the CES demand system.

3. Downstream goods firms choose reset prices optimally subject to Calvo frictions. Entrepreneurs choose consumption and bond holdings optimally. The central bank follows the Taylor rule, and the exogenous shocks follow their stated laws of motion.

4. The contract distribution evolves according to the accounting law above, aggregate quantities and prices are consistent with the service and final-good aggregators, and the labor, goods, bond, and profit-allocation conditions clear.

5. Expectations are rational: all choices are optimal under the stochastic law of motion generated by the equilibrium itself, and everyone understands the equilibrium law of motion.

I will also use as an intermediate object a *conditional general equilibrium*, where I take as given the path of the contract distribution $\{\mathcal{M}_t\}_t$, while firms and unions still optimize nominal wage terms within their assigned contract type, and all other equilibrium conditions remain unchanged. This object isolates how a given path of contract distribution affects the aggregate dynamics, while the contracting general equilibrium characterizes how that path is then determined by the aggregate dynamics through individual contract choices.

4 An Analytical Benchmark

Four hurdles make it difficult to analyze the model. The first two are standard. First, under persistent shocks agents have to forecast the future path of aggregate variables, including the contract distribution, and wage terms set in the past can affect dynamics today. Second, with long-term staggered contracts, the aggregate state includes the whole distribution of outstanding contracts.

The other two are novel and appear as two fixed-point problems. Third, the optimal COLA formula depends on how inflation signals the payoff-relevant revenue factor. Solving for the formula requires their endogenous joint distribution, which depends on the optimal COLA formula itself. Finally, adopting COLA changes the contract distribution and thus the gain from COLA, which determines adoption in the first place.

In order to isolate the two fixed-point problems that are central to this paper, I first make some restrictive assumptions to shut down the two standard dynamic complications. This allows me to fully characterize core model mechanisms analytically. I will go back to the full model for the quantitative exercise in Section 6.

Assumption 1. *$\varrho_u = \varrho_\rho = 0$, and $\sigma_{ut} = \sigma_u, \sigma_{\rho t} = \sigma_\rho$ for all $t \in \mathbb{Z}$ where σ_u and σ_ρ are known constants that are small enough that a log linear approximation is accurate. Contract length $N = 1$.*

First, serially independent shocks eliminate the forward- and backward-looking complications.[13] Second, one-period contracts imply that the contract distribution is represented by the current aggregate COLA share, denoted by m_t. Thus, the model contains three aggregate state variables (u_t, ρ_t, m_t). I log-linearize the system around the no-shock steady state *except for* the dynamics of the contract distribution. I retain its nonlinearity, as we do observe large variations of the aggregate COLA share in the data.

4.1 Firm-Union Contracting Problem

I first solve the problem of an individual firm-union pair who takes the aggregate dynamics as given. I solve it backwards: after the real wage is realized, the firm chooses labor; before that, the pair chooses the optimal wage term, depending on the contract type; and at the beginning, the pair chooses the type of the contract.

Stage 3: Labor Choice. Given the realized real wage and the real revenue factor, the firm's labor decision is summarized by the following linear policy rule.

[13] Technically, a full analytical characterization does not hinge on i.i.d. shocks. I assume it here only to ease the exposition.

Lemma 1. *(**Firm labor input choice**) Given the realized real wage and the real revenue factor,*

$$\hat{l}_{it} = \iota_1 \hat{\zeta}_t - \iota_2 \hat{w}_{it}$$

where $\iota_1, \iota_2 > 0$ are constants determined by the steady state, and for any variable x, $\hat{x}$ denotes its log deviation from steady state.

Proof. See Appendix D.2.1. □

Firms hire more labor when the real revenue factor is high, less when the realized real wage is high. The proof follows from approximating the labor demand problem.

Stage 2: Wage Term. The contractual wage splits the pie between the firm and its workers, and the size of the pie hinges on the revenue factor $\hat{\zeta}_t$. I show in Appendix D.2.1 that if both $\hat{\zeta}_t$ and π_t were contractible, the optimal real wage would be proportional to $\hat{\zeta}_t$:

$$\hat{w}_t^* = \mu \hat{\zeta}_t$$

where μ is a constant from the steady state and measures the real wage exposure to the revenue risk. Under reasonable parameter ranges, $\mu \geq 0$, which I maintain throughout this section.[14] When there is no external financing cost ($\phi_0 = 0$), the firm is truly risk-neutral. In that case, we have $\mu = 0$, and the firm absorbs all the risk. For a given $\mu > 0$, when ζ_t is high (low), the firm faces a loose (tight) financial constraint, and is willing to pay a high (low) real wage. Conditional on $\hat{\zeta}_t$, the optimal nominal wage moves one-for-one with inflation π_t.

However, the information structure specified in Section 3.1 prevents writing $\hat{\zeta}_t$ directly into the contract. COLA allows only inflation contingency, and inflation can provide an imperfect signal about the revenue factor $\hat{\zeta}_t$. A non-contingent contract is even coarser: it infers current revenue conditional on information only at contract signing.

Therefore, solving for the exact wage term requires knowing the endogenous data-generating process of π_t and $\hat{\zeta}_t$. Under Assumption 1, I conjecture the following particular joint process in equilibrium: the real revenue factor $\hat{\zeta}_t$ and inflation π_t are linear in exogenous shocks u_t and ρ_t, but the coefficients can be arbitrary functions of the endogenous state m_t. Namely,

$$\begin{aligned} \hat{\zeta}_t &= \zeta_\rho(m_t)\rho_t + \zeta_u(m_t)u_t \\ \pi_t &= \pi_\rho(m_t)\rho_t + \pi_u(m_t)u_t \end{aligned} \tag{2}$$

where $\zeta_\rho(m_t), \pi_\rho(m_t), \pi_u(m_t) > 0$ and $\zeta_u(m_t) < 0$. Thus, a positive demand shock raises both revenue and inflation, while a positive markup shock raises inflation but reduces revenue. In the next subsection, I verify both the functional forms and signs of the coefficients in general

[14] For fixed employment, Appendix C shows $\mu \geq 0$. For variable labor demand, it is possible that a higher revenue factor leads to a lower optimal real wage to expand employment. Appendix D.2.1 discusses this point in more detail.

equilibrium. Individual firms and workers take the joint process as given when they choose their contracts.

Lemma 2. *(**Realized real wage**) Given Assumption 1 and the perceived joint process in (2), the realized real wages implied by the two contract types are*

$$\hat{w}_t^{\mathcal{C}} = \mu \mathbb{E}\left[\hat{\zeta}_t | \mathcal{F}_{t-1}\right] + \delta_t \left(\pi_t - \mathbb{E}\left[\pi_t | \mathcal{F}_{t-1}\right]\right) = \delta_t \pi_t$$
$$\hat{w}_t^{\mathcal{N}} = \mu \mathbb{E}\left[\hat{\zeta}_t | \mathcal{F}_{t-1}\right] - \left(\pi_t - \mathbb{E}\left[\pi_t | \mathcal{F}_{t-1}\right]\right) = -\pi_t$$

where

$$\delta_t = \mu \frac{\mathrm{Cov}\left(\hat{\zeta}_t, \pi_t\right)}{\mathrm{Var}(\pi_t)}$$

Proof. See Appendix E.1. □

Given the conjecture in (2), both optimal realized real wages come from regressing the target $\hat{w}_t^* = \mu \hat{\zeta}_t$ on the respective information sets. $\mu \mathbb{E}_{t-1}\left[\hat{\zeta}_t\right]$ in both formulas is the forecast conditional on the common information set $\mathcal{F}_{t-1}$ at signing. COLA additionally leverages information from the inflation surprise, and the regression coefficient $\mathrm{Cov}(\hat{\zeta}_t, \pi_t)/\mathrm{Var}(\pi_t)$ translates the inflation surprise to the predicted revenue surprise. Because a non-contingent contract specifies a preset nominal wage, the implied real wage adjusts one for one against the inflation surprise. Since structural shocks are i.i.d. over time, expected inflation and expected revenue conditional on $\mathcal{F}_{t-1}$ are both zero, leading to Lemma 2.

The coefficient δ_t has a meaningful interpretation. Notice that the COLA pass-through from prices to nominal wages is exactly $1 + \delta_t$.[15] If $1 + \delta_t = 1$, COLA keeps the real wage constant and gives full real wage insurance to workers. If $1 + \delta_t = 0$, workers bear the full inflation risk, exactly as under a non-contingent contract. When $\mathrm{Cov}(\hat{\zeta}_t, \pi_t) < 0$, we have $1 + \delta_t < 1$. Thus, the COLA nominal pass-through will be optimally incomplete if inflation signals low revenue. Intuitively, in that case, firm revenues tend to be low in high-inflation states, making full inflation insurance difficult to afford. In general equilibrium, this is likely if inflation is primarily driven by markup shocks. In contrast, when inflation signals high revenue, the optimal COLA pass-through can be greater than one.

The following condition on the optimal COLA pass-through is the relevant case that I consider.

Condition 1. (**Positive COLA pass-through**) The optimal COLA pass-through from prices to nominal wages is positive, i.e.,

$$1 + \delta_t > 0$$

I provide more primitive assumptions for this condition to hold in equilibrium in Appendix E.2.2. In general, this is a very loose condition and is empirically grounded. Given $\mu > 0$, if $\mathrm{Cov}(\hat{\zeta}_t, \pi_t) \geq 0$, then $1 + \delta_t \geq 1$, so the condition is trivially satisfied. This can happen if inflation is mostly driven by demand shocks in equilibrium. Condition 1 fails only when $\mathrm{Cov}(\hat{\zeta}_t, \pi_t) < 0$,

[15] To see this, notice $\log W_t = \log w_t + \pi_t + \log P_{t-1} = (1 + \delta_t)\pi_t + \log P_{t-1} + \log w_{ss}$.

and μ is sufficiently large. This can happen if the firm's financing cost is extremely high relative to workers' risk aversion, and inflation is primarily driven by markup shocks, so that high inflation signals low revenues. In that case, workers actually compensate the firm in high-inflation states when COLA is adopted.

Stage 1: Contract Type. The last step is the choice of contract type. The pair compares the gross value of the best COLA contract with that of the best non-contingent contract. Given the real wage expressions above, the gain from COLA can be summarized by a simple variance term.

Lemma 3. *(**Contract type choice**) Given Assumption 1 and the conjecture in (2), the gross gain from COLA for a pair renewing at the end of period $t-1$ is*

$$\mathcal{V}_{t-1} \propto (1+\delta_t)^2 \operatorname{Var}_{t-1}(\pi_t)$$

Proof. See Appendix E.1. □

The formula illustrates the main economic content of the contracting problem. The COLA nominal pass-through $1+\delta_t$ measures the amount of inflation risk absorbed by COLA relative to non-contingent wages. When COLA provides more insurance against inflation surprises, or when the perceived inflation volatility is higher, the value of contingency rises, which captures the risk-sharing effect from COLA contracts.

It is worth mentioning here that both the gain from COLA $\mathcal{V}_{t-1}$ and the inflation risk $\operatorname{Var}_{t-1}(\pi_t)$ are the perceived quantities from the individual pair's perspective. Define $m_t^e = \mathbb{E}_{t-1}[m_t]$ as the pair's expected COLA share in period t. As will become clear shortly, both quantities are functions of m_t^e, i.e.,

$$\mathcal{V}_{t-1}(m_t^e), \;\; \operatorname{Var}_{t-1}(\pi_t | m_t^e)$$

It enters the individual problem through the perceived aggregate process (2), with m_t replaced with m_t^e. Thus, the expected COLA share summarizes the general equilibrium feedback to the individual problem.

4.2 Aggregate Effects from Any Given COLA Share m_t

I now show how a given contract distribution affects an otherwise standard New Keynesian block using the familiar three-equation representation augmented by the COLA share m_t. I will go back to the determination of m_t in the contracting general equilibrium in the next subsection.

Downstream firm pricing yields the familiar New Keynesian Phillips curve in marginal-cost form,

$$\pi_t = \beta \mathbb{E}_t[\pi_{t+1}] + \kappa \hat{mc}_t + \kappa u_t \tag{3}$$

where $\kappa = \frac{(1-\theta)(1-\theta\beta)}{\theta}$. Hence, this price-setting block is standard, and the departure from the canonical model is inside the real marginal cost, which is the real price of the service bundle from

upstream. In log-linear form, it depends on output and on the aggregate real wage faced by upstream service firms:

$$\hat{mc}_t = \frac{\varepsilon_x - \iota_1}{\iota_1 \varepsilon_x} \hat{y}_t + \frac{\iota_2}{\iota_1} \hat{\omega}_t$$

where

$$\hat{\omega}_t := m_t \hat{w}_t^{\mathcal{C}}(\pi_t) + (1 - m_t) \hat{w}_t^{\mathcal{N}}(\pi_t) = \underbrace{[m_t \delta_t - (1 - m_t)]}_{:= K(m_t, \delta_t)} \pi_t \tag{4}$$

is an average of the realized real wages weighted by their contract shares.

The coefficient $K(m_t, \delta_t)$ defined in equation (4) measures the sensitivity of the aggregate real wage to inflation surprises. It is useful to keep the magnitude in mind. When all firms use non-contingent wages, $K(0, \delta_t) = -1$: the aggregate nominal wage is predetermined, so the real wage falls one-for-one with inflation. Under Condition 1, a higher COLA share and/or a higher COLA pass-through raise $K(m_t, \delta_t)$ and reduce the sensitivity. When all firms use fully indexed COLA, $K(1, 0) = 0$, and the aggregate real wage does not respond to inflation surprises.

The entrepreneur's consumption-saving problem leads to the standard dynamic IS curve

$$\hat{c}_t^e = \mathbb{E}_t[\hat{c}_{t+1}^e] - \frac{1}{\gamma}(i_t - \mathbb{E}_t[\pi_{t+1}] - \rho_t) \tag{5}$$

Since labor is the only primary input into production, up to first order the aggregate production function is

$$\hat{y}_t = m_t \hat{l}_t^{\mathcal{C}} + (1 - m_t) \hat{l}_t^{\mathcal{N}} = \iota_1 \hat{\zeta}_t - \iota_2 \hat{\omega}_t \tag{6}$$

where we plug in the firms' labor input choices under different types of contracts and use the definition of $\hat{\omega}_t$. Final output is used for consumption and to cover external financing costs. We can write the aggregate resource constraint in terms of consumption, the aggregate real wage, and the aggregate real revenue factor

$$\hat{y}_t = \frac{\bar{C}^e}{\bar{Y}} \hat{c}_t^e + \varphi_\omega \hat{\omega}_t + \varphi_\zeta \hat{\zeta}_t \tag{7}$$

where $\bar{C}^e$ and $\bar{Y}$ are entrepreneur consumption and total output in the steady state, and φ_ω and φ_ζ are constants from the steady state. Given (m_t, δ_t), equations (3)-(7) completely characterize the economy in period t, because the two expectation terms in (3) and (5) are zero under i.i.d. shocks. Recall that the COLA nominal pass-through $1 + \delta_t$ is itself endogenous to $\hat{\zeta}_t$ and π_t. Therefore, combining the above five equations, the conjecture in (2), and the wage expressions in Lemma 2, we can solve for the four coefficient functions $\zeta_\rho(m_t), \pi_\rho(m_t), \zeta_u(m_t), \pi_u(m_t)$ by the method of undetermined coefficients, which characterize the conditional general equilibrium. Appendix E.2.1 shows its existence and uniqueness for any given $m \in [0, 1]$ under some regularity conditions, which I maintain throughout this section.

In this conditional GE, δ_t is itself determined by the aggregate COLA share m_t. I therefore write

$$\mathcal{K}(m_t) := K(m_t, \delta(m_t))$$

which nests both the COLA share and the optimal COLA pass-through. After substituting $\hat{\omega}_t = \mathcal{K}(m_t)\pi_t$ back into the above equations and rearranging, we obtain the following:

Proposition 1. ***(Three-equation representation)*** *Given Assumption 1, for any predetermined m_t, the conditional general equilibrium has the representation*

$$\begin{aligned}
NKPC: \pi_t &= \frac{\kappa}{1-\kappa\frac{\iota_2}{\iota_1}\mathcal{K}(m_t)}\left[\frac{\varepsilon_x-\iota_1}{\iota_1\varepsilon_x}\hat{y}_t + u_t\right] \\
DIS: \hat{y}_t &= \frac{\iota_1}{\iota_1-\varphi_\zeta}\left[-\frac{\bar{C}^e}{\bar{Y}\gamma}(i_t-\rho_t) + \left(\varphi_\omega + \frac{\iota_2}{\iota_1}\varphi_\zeta\right)\mathcal{K}(m_t)\pi_t\right] \\
MP: i_t &= \phi_\pi\pi_t + \phi_y\hat{y}_t
\end{aligned}$$

Under Condition 1, $\mathcal{K}(m_t)$ increases with m_t.

Proof. See Appendix E.2.2. □

This proposition clarifies the connection of the current model with the canonical New Keynesian model.[16] Holding m_t fixed, the economy is still characterized by a Phillips curve, an aggregate-demand equation, and a monetary-policy rule. The main departure is that the Phillips curve contains the COLA share m_t, whose aggregate effect is summarized by $\mathcal{K}(m_t)$.[17] In particular, a higher COLA share m_t raises $\mathcal{K}(m_t)$. This makes the aggregate real wages more rigid and therefore steepens the Phillips curve. Consequently, the same output gap or cost-push shock produces a larger inflation response. [18]

The economic mechanism is intuitive. Consider a given inflationary shock. An economy with a higher COLA share will have less real wage erosion given Condition 1. Higher real wages then reduce labor demand by service firms, which in turn raises the service price and thus the marginal cost of downstream goods firms. Goods firms that reset prices then choose higher prices, further adding to inflationary pressure. The fixed point of this wage-price feedback is embodied in the state-dependent Phillips curve. I refer to this feedback mechanism as the "conditional wage-price spiral", which describes how a given stock of COLA contracts affects aggregate dynamics.

As a technical remark, Condition 1 does not mechanically imply that $\mathcal{K}(m_t)$ increases with m_t. Notice $\mathcal{K}_m(m_t) = 1 + \delta(m_t) + m_t\delta_m(m_t)$. Condition 1 makes sure a marginal switch from a non-contingent contract to COLA raises the aggregate real wage rigidity, holding the COLA

[16] This representation does not hinge on Assumption 1 and will be augmented with forward-looking terms like the usual inflation and output expectations, and backward-looking terms from wage formulas determined in the past.

[17] Note $\mathcal{K}(m_t)$ also enters the DIS curve, but this is a more mechanical effect: $\mathcal{K}(m_t)\pi_t = \hat{\omega}_t$ is the total demand from those hand-to-mouth workers.

[18] One can also derive the wage Phillips curve: $\pi_t^w := W_t^{agg} - W_{t-1}^{agg} = \omega_t - \omega_{t-1} + \pi_t = [1 + \mathcal{K}(m_t)]\,\pi_t - \omega_{t-1}$. A higher COLA share also steepens the wage Phillips curve.

pass-through fixed. But such a marginal change can also change the optimal COLA pass-through. The structure of the conditional GE, including the inference results from Lemma 2, ensures that Condition 1 is sufficient (and actually also necessary) for the comparative statics.

An immediate implication of the conditional wage-price spiral is that, for fixed primitive shock variances, inflation volatility is increasing in the COLA share.

Corollary 1. *(**Inflation volatility in equilibrium**) Given Assumption 1 and under Condition 1, inflation volatility* $\text{Var}(\pi_t|m_t)$ *increases with the COLA share* m_t*.*

The corollary follows because both shock loadings $\pi_u(m_t)$ and $\pi_\rho(m_t)$ in the conditional inflation process are positive and increase with m_t.

Finally, given the analytical complexity of the model, to highlight the core economic intuition and avoid unnecessary discussion about unrealistic scenarios, I will work with the next condition on the optimal COLA pass-through in the subsequent analysis.

Condition 2. (**Positive slope of the COLA pass-through**) The optimal COLA pass-through in the conditional GE weakly increases with the COLA share, i.e.,

$$\delta_m(m_t) \geq 0$$

Appendix E.2.2 states more primitive assumptions for Condition 2 to hold. This condition is consistent with the empirical pattern in Figure 2. In the conditional GE, a higher COLA share affects the optimal COLA pass-through in two different directions through the inference problem. As a higher COLA share amplifies inflation, the information about the revenue factor carried by the same observed inflation is diluted. For the part of inflation that is driven by markup shocks, this means the same observed inflation implies a smaller drop in revenue. This raises the optimal pass-through. For the part of inflation that is driven by demand shocks, the same observed inflation implies a smaller increase in revenue. This lowers the optimal pass-through. Condition 2 holds if the first force weakly dominates.

4.3 The Contract Accelerator

I now characterize the contracting general equilibrium in which the aggregate COLA share is determined by individual contract choices.

It follows immediately from Lemma 3 and Corollary 1 that under Condition 2, higher inflation risk triggers more COLA adoption, which in turn amplifies inflation risk by strengthening the conditional wage-price spiral and further triggers COLA adoption and higher COLA pass-through. This spiral between inflation risk and the indexation regime is what I call "the contract accelerator".[19]

[19]Note that Condition 2 is sufficient but not necessary for the self-reinforcing feedback mechanism. It remains effective as long as the increase in the inflation risk generated by more COLA adoption dominates the change in the gain from COLA.

What is the implication of the contract accelerator for the contracting general equilibrium? Recall that the gain from COLA in Lemma 3 is the perceived gain by an individual firm-union pair. This gain depends crucially on their expected aggregate COLA share:

Proposition 2. *(**Strategic complementarity in contract choice**) Given Assumption 1 and under Conditions 1-2, the expected gain from COLA increases with the expected COLA share:*

$$\frac{\partial \mathcal{V}\left(m_t^e\right)}{\partial m_t^e} > 0$$

where $m_t^e = \mathbb{E}_{t-1}[m_t]$ is the expected COLA share in t by an individual firm-union pair.

Proof. See Appendix E.3. □

That is, if a firm-union pair expects more other pairs to use COLA, it should also expect more volatile inflation, which in turn raises their expected gain from COLA. Therefore, through the contract accelerator, an individual pair has a stronger incentive to adopt COLA precisely when it expects others to do so: contract choices are strategic complements.

Because there is no private information across firm-union pairs, they have the same expected COLA share in a symmetric equilibrium. Denote by $\mathrm{BR}\left(m_t^e\right)$ the induced aggregate COLA share for any given belief m_t^e, which is given by

$$\mathrm{BR}\left(m_t^e\right) = F_\chi\left(\left(1+\delta(m_t^e)\right)^2\left[\pi_u(m_t^e)^2\sigma_u^2 + \pi_\rho(m_t^e)^2\sigma_\rho^2\right]\right) \tag{8}$$

The right-hand side is exactly the share of firm-union pairs whose adoption costs are below the expected gain $\mathcal{V}(m_t^e)$. Thus, it is the aggregate best response to that expectation.

Under rational expectations, the expected COLA share should coincide with the aggregate share generated by individual adoption decisions. This characterizes the equilibrium COLA share, as the following proposition shows.

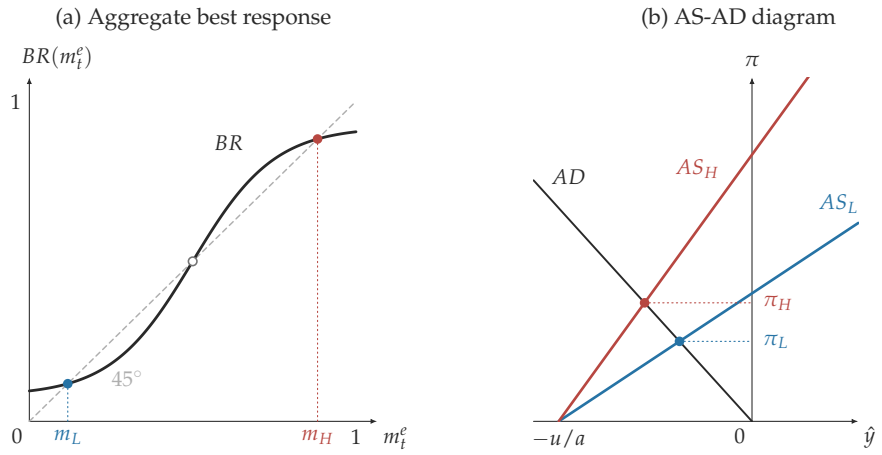
Proposition 3. *(**Contracting general equilibrium and multiple COLA equilibria**) Given Assumption 1 and under Conditions 1-2, there exists a contracting general equilibrium in which the COLA share m_t is deterministic and solves*

$$m_t = \mathrm{BR}(m_t)$$

Moreover, if the COLA adoption cost is common across firm-union pairs, and $\sigma_u = \varsigma, \sigma_\rho = k\varsigma$ for some constant $k > 0$, then there exist some $\underline{\varsigma} < \bar{\varsigma}$ such that the economy has a unique no-COLA equilibrium when $\varsigma < \underline{\varsigma}$, multiple COLA equilibria when $\varsigma \in (\underline{\varsigma}, \bar{\varsigma})$, and a unique all-COLA equilibrium when $\varsigma > \bar{\varsigma}$.

Proof. See Appendix E.3.1. □

Figure 3: Illustration of the two spirals



Notes: The left panel illustrates the aggregate best response with heterogeneous COLA adoption costs. The right panel plots the AS-AD diagram in response to a given positive markup shock. The case for a demand shock is similar. The AS curves are the Phillips curves that correspond to different COLA regimes (indicated by the same color), and the AD curve combines the DIS and MP equations in Proposition 1 and summarizes the demand block. For illustration, the AD curve is held fixed to highlight the supply-side amplification channel. Allowing the AD curve to change with m_t does not change the key message.

The equilibrium COLA share m_t can in principle be stochastic and respond to markup or demand shocks. The first half of Proposition 3 guarantees the existence of a non-stochastic m_t in equilibrium,[20] though it does not rule out the possibility of multiple such fixed points. The strategic complementarity in contract choice implies that the aggregate best response $BR(m_t)$ is increasing in m_t, and it can cross the 45° line more than once.

The second half of Proposition 3 gives a clean example of multiplicity when the random adoption cost degenerates. The volatility of primitive shocks ς measures the total amount of macroeconomic uncertainty in the economy, while the ratio σ_u/σ_ρ is fixed.[21] When uncertainty is low, inflation risk is limited, so the insurance value of COLA is too small to justify the cost. The economy settles in a no-COLA equilibrium. When uncertainty is high, inflation risk is large anyway, so COLA becomes privately attractive, and the economy settles in an all-COLA equilibrium. For intermediate uncertainty, beliefs matter: the economy settles in a no-COLA equilibrium if all pairs expect no other pairs to adopt COLA, and vice versa for the all-COLA equilibrium.

Figure 3 illustrates the economic content of the self-fulfilling regimes when they emerge. The figure assumes heterogeneous adoption costs, which smooth the best-response function on the left

[20] I will explain why this is the case and discuss the possibility of stochastic contract distribution in Section 6.1 when I characterize the full model.

[21] I show in Appendix E.5 that fixing σ_u/σ_ρ implies that the optimal COLA pass-through does not change with ς. This simplifies the analysis and also isolates the effect of the total amount of aggregate uncertainty from the composition of uncertainty.

panel and generate interior pure equilibria.[22] Other things equal, a firm-union pair that expects a higher COLA share also expects a steeper NKPC and stronger inflation response to any given aggregate shock. This is visualized in the AD-AS diagram on the right panel, which summarizes the conditional wage-price spiral from the conditional GE. Higher sensitivity of inflation to aggregate shocks then raises the anticipated inflation volatility at contracting, which feeds back to the left panel and supports the high-COLA regime. The reasoning is similar for other self-fulfilling equilibria.

4.4 Role of Monetary Policy

Given the two nested spirals, it is natural to ask whether monetary policy can weaken them. The policy response to inflation, captured by ϕ_π in the Taylor rule, affects the best-response curve (8), because it affects coefficients in the conditional inflation process for any given m_t through the conditional GE. To obtain sensible results, I first restrict attention to the following type of equilibrium point.

Definition 1. A fixed point m_t^* is locally stable if

$$\partial_m \, \mathrm{BR}(m_t^*) < 1.$$

When the best-response function crosses the 45° line from above, its slope there is below one, and the equilibrium is locally stable.[23] The next proposition gives the local comparative statics for monetary policy at locally stable fixed points.

Proposition 4. ***(Policy reduces COLA adoption)*** *Given Assumption 1 and under Conditions 1-2, a stronger policy response to inflation reduces the COLA share in a locally stable equilibrium:*

$$\frac{dm_t^*}{d\phi_\pi} < 0$$

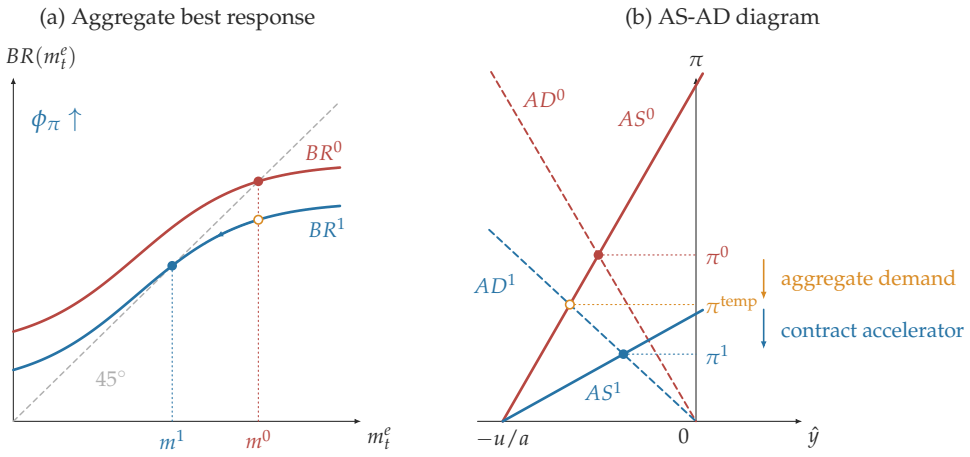
Proof. See Appendix E.4. □

Figure 4 illustrates the mechanism and the implication of this proposition. Suppose the economy is initially in a locally stable equilibrium with COLA share m^0. A more hawkish rule first stabilizes inflation through the conventional aggregate demand channel, which flattens the AD curve from AD^0 to AD^1. Inflation then becomes less responsive to the same aggregate shock (from π^0 to

[22] The common-cost limit resembles a complete-information coordination game, and the best response curve has a reverse "Z" shape. Heterogeneous adoption costs smooth the aggregate best-response curve and generate interior pure equilibria. Appendix E.3.1 shows that a higher degree of dispersion in the adoption costs can further flatten the best-response curve and can even restore uniqueness. This echoes the equilibrium-selection logic of global games. However, unlike standard global games, heterogeneity here lies in private adoption costs rather than noisy signals about a common fundamental (Morris and Shin, 2003).

[23] I select locally eductively stable equilibria in the sense of Guesnerie (1992). Starting from any sufficiently small common-knowledge neighborhood of conjectures around m_t^*, higher-order rational reasoning contracts the admissible set to m_t^* if $BR_m(m_t^*) < 1$. An equilibrium with $BR_m(m_t^*) > 1$ is a fixed point under rational expectations, but it requires knife-edge coordination because any local imprecision in conjectures is amplified through higher-order reasoning.

Figure 4: Illustration of a more hawkish Taylor rule



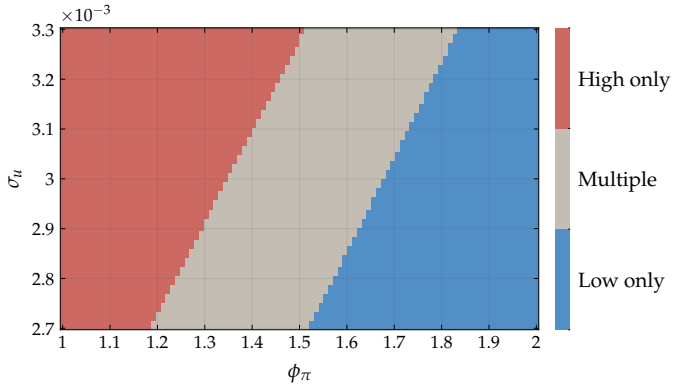
Notes: The left panel illustrates the aggregate best response, where BR^1 corresponds to a higher ϕ_π than BR^0. The corresponding AS curve, AD curve, and inflation response to a given markup shock are indicated by the same color. The case for a demand shock is similar and is reported in Appendix E.6. The AS curves are the Phillips curves corresponding to different COLA shares, and the AD curves combine the DIS and MP equations in Proposition 1, which summarize the demand block. For illustration, I omit the change in the AD curve when m_t changes, which does not affect the main message.

π^{temp}), and inflation volatility declines. This shifts the best-response curve downwards from BR^0 to BR^1, and the equilibrium COLA share starts to decline, as the insurance value is now smaller. Moreover, I show in Appendix E.4 that a higher ϕ_π also reduces COLA pass-through. Both margins of adjustment flatten the Phillips (or AS) curve from AS^0 towards AS^1, generating an even smaller inflation response and lower inflation volatility. The COLA share continues to move along BR^1, and the AS curve continues to flatten, until the COLA share reaches the new equilibrium m^1, and AS^0 reaches AS^1. Thus, beyond the conventional demand channel, a stronger policy response to inflation generates a second-round inflation stabilization (from π^{temp} to π^1) through the contract accelerator.

As a remark on the local stability requirement, I show in Appendix E.3.1 that if there is a unique equilibrium COLA share, it must be locally stable. Thus, a more aggressive policy rule will unambiguously reduce that share. This requirement matters when there are multiple equilibria. For example, when there are three fixed points as in the left panel of Figure 3, the low and high intersections are stable, and the middle intersection is unstable. In that case, a downward shift in the best-response curve moves the stable low- and high-COLA equilibria towards lower COLA shares.

While a more aggressive policy rule can reduce the equilibrium COLA share locally, can it weaken the strategic complementarity and restore global uniqueness? The next proposition says it can.

Figure 5: Policy-uncertainty frontier of the endogenous contract regimes



Notes: This figure plots the equilibrium properties under different policy-uncertainty combinations in the stationary contracting general equilibrium. "High/Low only" means there is a unique equilibrium with a COLA share above/below 50 percent, and "Multiple" means there are multiple COLA equilibria.

Proposition 5. *(**Policy and global uniqueness**) Given Assumption 1 and under Conditions 1-2, a stronger policy response to inflation weakens strategic complementarity for any expected COLA share, i.e., for all* $m \in [0,1]$,

$$\frac{\partial^2 \mathcal{V}_{t-1}}{\partial m \partial \phi_\pi} < 0$$

Moreover, there exists some finite $\phi_\pi^U > 0$ *such that for all admissible* $\phi_\pi > \phi_\pi^U$, *there is a unique contracting general equilibrium.*

Proof. See Appendix E.4. □

Therefore, a more aggressive Taylor rule not only reduces the level of the expected gain but also flattens it, weakening the strategic complementarity in contract choice. More importantly, when policy becomes sufficiently aggressive, the slope of the aggregate best response $\text{BR}(m)$ will be uniformly smaller than one, implying that it will only intersect the 45° line once. In other words, a sufficiently aggressive monetary policy can eliminate the potential multiplicity and anchor the COLA regime. Notice this is distinct from the standard determinacy issue (Clarida et al., 2000, Blanchard and Kahn, 1980). I will elaborate on this point shortly.

In Appendix E.5, I also discuss formally how this anchoring role of monetary policy interacts with aggregate uncertainty, measured by the volatility of aggregate shocks. I illustrate the core implications using Figure 5, which plots the equilibrium properties under different combinations of policy and uncertainty. The horizontal axis is the monetary policy response to inflation ϕ_π, and the vertical axis is the volatility of markup shocks σ_u. The gray band is the set of policy-uncertainty combinations that support multiple COLA equilibria. The band is upward-sloping, implying that when uncertainty rises, an economy in a low-COLA regime may drift to the multiplicity region or to the high-COLA regime, if the policy rule does not change. In turn, a more aggressive rule is needed to prevent those scenarios and anchor the economy in a low-COLA regime.

The figure also makes it clear why endogenous indexation is not a cosmetic extension of the New Keynesian model. For every point in the policy-uncertainty plane, the conditional general equilibrium is unique for any given contract distribution. That corresponds to a standard sticky-wage model augmented with exogenous indexation. However, once contracts are endogenously chosen, the indexation regime can be indeterminate, or the economy may directly settle in a high-COLA regime with strong inflation propagation. This distinction matters for policy. If policymakers overlook the contract accelerator, a stabilization policy may turn out to be insufficient, leaving the economy vulnerable to uncertainty shocks or belief shocks.

4.5 Taking Stock of the Analytical Benchmark

The model delivers a simple economic message: private risk-sharing can have sizable macroeconomic effects and policy implications. COLA allows more flexible risk-sharing between firms and workers. Private contract choices give rise to two nested spirals at the aggregate level: the conditional wage-price spiral for any given contract distribution, and the contract accelerator through the feedback loop between inflation risk and the indexation regime. A more aggressive policy rule discourages COLA adoption, weakens the contract accelerator, and can anchor the economy in a low-indexation regime.

5 Validating the Model

The following three mechanisms are the backbone of the contract accelerator in Section 4:

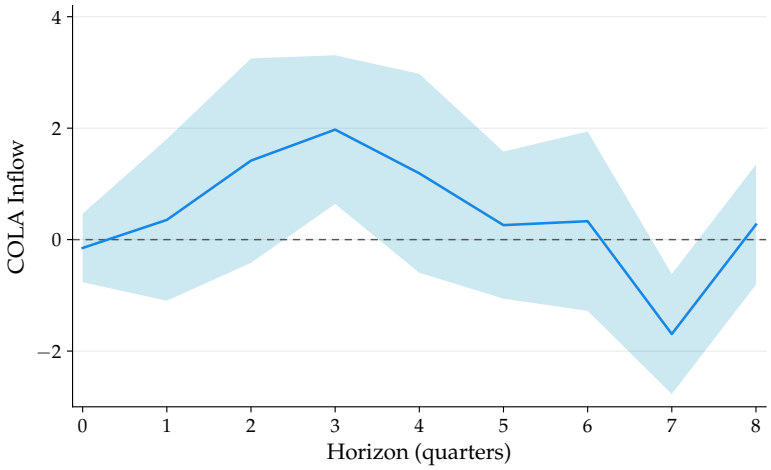
1. Higher macroeconomic uncertainty raises new COLA adoption.
2. A higher COLA share amplifies the inflation response to an inflationary shock.
3. COLA adoption exhibits strategic complementarity.

Quantitatively, these mechanisms remain operative in the full model with persistent shocks and long-term contracts, as I will show momentarily. Before going there, I empirically validate these mechanisms using the contract data in this section.

5.1 Higher Uncertainty Triggers More COLA Adoption

The model predicts that through the risk-sharing channel, higher primitive uncertainty makes COLA more valuable and encourages COLA adoption. To test this prediction, I first construct a sequence of uncertainty shocks. I measure uncertainty using the one-year-ahead macroeconomic uncertainty index of Jurado et al. (2015, hereafter JLN), which is available from the early 1970s and is constructed from a broad set of macroeconomic and financial variables. To obtain innovations in uncertainty, I regress the index on its own lags and on lagged macroeconomic controls, including output, inflation, and the federal funds rate. I then estimate local projections of the quarterly share of newly signed contracts that include COLA on the residual uncertainty innovation.

Figure 6: The response of the COLA adoption rate to an uncertainty shock



Notes: This figure plots the coefficient δ_h for $h = 0, 1, ..., 8$ in the regression $COLAIn_{t+h} = \delta_h u_t + \sum_{j=1}^{l} \Gamma_{h,j} Z_{t-j} + \varepsilon_{t+h}$. $COLAIn_t$ is the fraction of COLA contracts among all the newly signed contracts in quarter t. The control Z_t includes real GDP, inflation, and FFR. The shaded area shows the 90 percent confidence interval. Similar results can be found in Appendix B.1 where the adoption rate is weighted by the union size.

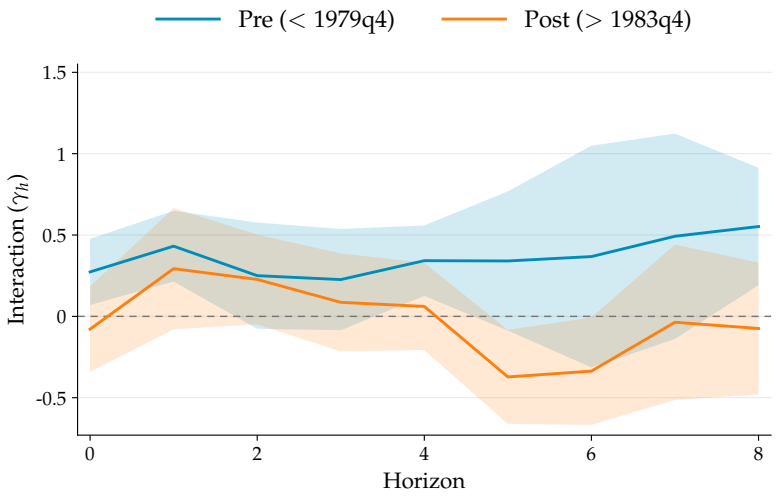
Figure 6 plots the results. A positive uncertainty innovation raises the COLA adoption rate, and the effect gradually fades away as uncertainty returns toward normal. The pattern is consistent with the risk-sharing channel in the model. I also show that higher subjective inflation uncertainty from the Michigan Consumer Survey of Expectations (instrumented by the JLN uncertainty innovation) raises the probability of adopting COLA at the individual firm-union level, providing a complementary check. However, the survey began in the late 1970s and missed the 1973-75 inflation spike, so I leave it to Appendix B.2.

5.2 More COLA Amplifies the Inflation Response to Policy Shocks

Next, I test the prediction that a higher stock of COLA amplifies the inflation response to a given inflationary shock. I use Romer-Romer monetary policy shocks (Romer and Romer, 2004), which are available for the period in which the aggregate COLA share can be measured. I estimate local projections of inflation on the monetary-policy shock and its interaction with a high-COLA indicator. A quarter is classified as high-COLA if the COLA share is above the median within the relevant subsample. I use the one-period lagged share to avoid the case that some changes in the share are already the response to the shock. Because the conduct of monetary policy is supposed to have changed substantially after the Volcker disinflation, I estimate the responses separately for the pre-Volcker and post-Volcker periods.

Figure 7 plots the coefficients of the interaction term, which captures the additional inflation response in high-COLA states. In the pre-Volcker period, a high COLA share amplifies the inflationary effect of an expansionary monetary shock, and the effect is persistent. In the post-Volcker

Figure 7: State-dependency of inflation IRF to monetary policy shocks



Notes: This figure shows the coefficient γ_h in $\pi_{t+h} = \beta_h RR_t + \gamma_h RR_t \times \mathbb{I}_{t-1}^{COLAHigh} + \sum_{j=1}^{l} \Gamma_{h,j} Z_{t-j} + \varepsilon_{t+h}$ for different horizons $h = 0, 1, ..., 8$. The blue line corresponds to the pre-Volcker period, while the orange line is the post-Volcker period. RR_t denotes the Romer–Romer monetary-policy shock, and $\mathbb{I}_t^{COLAHigh} = 1$ if the COLA share is higher than the sample median. The control Z_t includes real GDP, inflation, and FFR with lags up to 4 quarters. The shaded area shows the 90 percent confidence interval.

period, the response is positive on impact but reverses at longer horizons, which is consistent with a policy regime that reacts more aggressively to inflation.

5.3 Strategic Complementarity in Contract Choice

One of the key theoretical results of the model is that contract choices are strategic complements. I test it at the firm-union level. I regress an indicator for whether a newly signed contract includes COLA on the aggregate COLA stock, controlling for firm-union pair fixed effects and aggregate macroeconomic conditions. Pair fixed effects allow comparison within the same bargaining relationship across time. As in the previous exercise, I use the one-period lagged COLA share to avoid any new adoption this period that enters the stock mechanically.

A remaining concern is that the COLA stock may be correlated with persistent macroeconomic conditions that also affect current adoption. To address this concern as much as possible, I instrument the COLA stock with long-lagged (8-12 quarters) uncertainty shocks and oil-supply shocks separately. The identifying assumption is that, conditional on recent macro controls and pair fixed effects, these old and distant shocks affect current COLA adoption only through their effect on earlier contract choices. The contract composition induced by those earlier choices persists until today due to the long duration of these contracts.

Table 2 reports the results from both the OLS and the IV specifications. Across specifications, the coefficient on the COLA stock is positive and statistically significant. The OLS estimate implies that a one percentage point increase in the inherited COLA stock is associated with a 0.33

Table 2: Strategic Complementarity in COLA Adoption

COLA	(1) OLS	(2) JLN IV	(3) Oil IV
COLA share	0.330***	0.458***	0.511***
	(0.056)	(0.141)	(0.141)
Pair FE	✓	✓	✓
Observations	18642	15467	14248

Notes: The dependent variable is an indicator for COLA adoption at renegotiation. COLA share is the one-period lagged COLA share. Column 1 estimates the relationship by OLS. Column 2 instruments lagged COLA share with lags 8–12 of residualized JLN uncertainty shocks, and column 3 instruments lagged COLA share with lags 8–12 of Baumeister-Hamilton oil supply shocks. Macro controls include lags 1–4 of JLN uncertainty, the federal funds rate, and the detrended log real GDP, in addition to log covered workers and contract duration. All specifications include bargaining-pair fixed effects. Standard errors are two-way clustered by bargaining pair and negotiation quarter. *, **, and *** denote significance at the 10, 5, and 1 percent levels.

percentage point increase in the probability that a new contract includes COLA. The IV estimates are larger, between 0.46 and 0.51 percentage points.

Taken together, the reduced-form evidence supports the mechanisms developed in the model. The next section then assesses the quantitative importance of the contract accelerator.

6 Quantitative Analysis

I use the full model to quantify the relevance of the contract accelerator for two historical episodes: the Great Inflation (1965Q1:1983Q4) and the Great Moderation (1984Q1:2007Q4). The exercise asks two questions. First, how much did the contract accelerator amplify inflation during the 1970s and early 1980s? Second, how much more did the post-Volcker policy rule stabilize inflation once its effect on the contract accelerator is considered? Finally, I will also use the Great Inflation as a laboratory for welfare and optimal policy analysis.

6.1 The Quantitative Environment

I restore both persistent shocks and long-term staggered contracts in the full model, which will be important for quantitative purposes. Still, I will make one restriction on the volatility of primitive shocks, whose extensions will be discussed in Section 6.6.

Assumption 2. *The volatility of primitive shocks has a deterministic time path. Namely, at any t_0, agents know* $\left\{\Sigma_t : \Sigma_t = diag(\sigma_{ut}^2, \sigma_{\rho t}^2)\right\}_{t \geq t_0}$.

Notice that Assumption 2 does not make the economy deterministic. The economy is always in a stochastic equilibrium, as markup and demand shocks remain stochastic. Instead, it allows me to focus on a class of contracting equilibria in which the time path of the contract distribution is deterministic.

Formally, recall that $\mathcal{M}_t = (\mathcal{M}_{t,1}, \ldots, \mathcal{M}_{t,N})$ denotes the distribution of outstanding contracts at date t, where $\mathcal{M}_{t,k}$ is the mass of COLA inflows at the end of period $t-k$. Suppose the current renewing cohort at the end of period $t-1$ expects a deterministic path of the contract distribution $\{\mathcal{M}^e_{t+\tau}\}_{\tau\geq 0}$. Conditional on that expected path and on a deterministic path of primitive volatility $\{\Sigma_{t+\tau}\}_{\tau\geq 0}$, the conditional general equilibrium has a convenient finite-dimensional representation. Namely, if the vector of aggregate shocks $v_t = (u_t, \rho_t)'$ follows an AR(1) process, and the contract duration is N periods, then equilibrium inflation and the revenue factor can be written as finite moving averages:

$$\pi_t = \sum_{s=0}^{N} q^{\pi}_{t,s} v_{t-s}, \qquad \hat{\zeta}_t = \sum_{s=0}^{N} q^{\zeta}_{t,s} v_{t-s} \tag{9}$$

where these coefficients are functions of $\{\mathcal{M}^e_{t+\tau}\}_{\tau\geq 0}$ and $\{\Sigma_{t+\tau}\}_{\tau\geq 0}$ and are thus deterministic. Notice that the conjecture (2) in the analytical benchmark is a special case of equation (9).

This further implies that the time path of the conditional variance-covariance matrix of $\left(\pi_t, \hat{\zeta}_t\right)$ is deterministic. This variance-covariance structure pins down the expected gain from COLA, as shown in Appendix F.1. Hence, in reduced form, one can write the expected gain from COLA for this renewing cohort as

$$\mathcal{V}_{t-1}\left(\{\mathcal{M}^e_{t+\tau}\}_{\tau\geq 0}; \{\Sigma_{t+\tau}\}_{\tau\geq 0}\right) \tag{10}$$

Aggregating individual choices under heterogeneous adoption costs leads to the new COLA inflow from this cohort. The distribution of the actual outstanding contracts evolves according to

$$\mathcal{M}_{t,1} = \frac{1}{N} F_{\chi}\left(\mathcal{V}_{t-1}\left(\{\mathcal{M}^e_{t+\tau}\}_{\tau\geq 0}; \{\Sigma_{t+\tau}\}_{\tau\geq 0}\right)\right), \qquad \mathcal{M}_{t,k} = \mathcal{M}_{t-1,k-1}, \quad k = 2, \ldots, N \tag{11}$$

which is independent of first-order shocks. This supports the initial expectation that the path of the contract distribution is deterministic. Appendix F.1 provides a detailed characterization of the full model.

I deliberately separate first-order shocks from the evolution of the contract distribution. Realized first-order shocks affect inflation and real activity, but not the contract distribution. This separation makes the model transparent and tractable. It results from both Assumption 2 and the feature of the model that COLA gains depend only on second-order moments (recall inflation volatility in the analytical benchmark). In Section 6.6, I discuss extensions where first-order shocks affect the contract distribution, making it a genuinely stochastic state.

Finally, to better match the observed COLA pass-through, I further assume that service firms face a quadratic labor adjustment cost for any employment deviation from the steady state.[24] This is a quantitative device and will not change any qualitative feature of the model.

[24] This increases the firm's effective risk aversion while keeping the tractability of the contracting problem. This extension makes it easier for the model to match the observed COLA pass-through without pushing other parameters to unreasonable regions.

Solution Algorithm. The separation of first- and second-order moments in contract adoption significantly reduces the difficulty of solving the model. The equilibrium appears as two layers of fixed points. In the outer loop, I conjecture a path for the contract distribution. In the inner loop, taking that path as given, I solve the conditional equilibrium, including the associated filtering problem, using the finite moving-average representation in equation (9). That allows me to evaluate the expected gain from COLA as in (10). The initial conjecture of the contract distribution is updated according to the laws of motion (11). This process repeats until the conjectured and the induced paths coincide.

For stationary environments, the outer fixed point contains only a stationary distribution of outstanding contracts. For transitional dynamics, such as unexpected changes in the volatility of primitive shocks or in the policy rule, the outer fixed point is the entire time path of the contract distribution. Appendix F.2 provides the computational details.

6.2 Calibration

I calibrate the full model to the Great Inflation (GI) and the Great Moderation (GM), allowing only two things to change across the two eras: the monetary policy response to inflation and the volatility of primitive shocks. I discuss my calibration strategy next.

Panel A of Table 3 shows the externally calibrated parameters. The discount factor β, worker risk aversion γ, the elasticities of substitution across services and across varieties $\varepsilon_x, \bar{\varepsilon}_y$, and the price-resetting probability $1-\theta$ all take standard values. Contract length N is set to 8 quarters, the typical duration of a contract in the data. The standard deviation of the idiosyncratic firm shock σ_η is set to 0.25 to replicate the annual value of 0.5 as suggested by Bloom et al. (2018). I set the policy response to inflation $\phi_{\pi,GI} = 1.1$ during the Great Inflation and $\phi_{\pi,GM} = 1.7$ during the Great Moderation, implying a more hawkish policy regime after the Volcker disinflation. The response to output is set to 0.4 for both periods. The rule is active in both periods, which keeps the conditional general equilibrium determinate.[25] Finally, similar to the bargaining power in a standard search model with Nash bargaining, the firm's weight in the joint value λ is difficult to separately identify. I set λ to one in the baseline calibration. Appendix G.5.2 shows that the main quantitative results are robust to alternative values of λ.

The rest of the parameters are estimated internally by the simulated method of moments. They include the slope of the external financing cost ϕ_0, the slope of the labor adjustment cost ν, worker home production b, the AR(1) persistence and innovation volatilities of the markup shock and the demand shock, and the COLA cost distribution F_χ. I assume F_χ is the CDF of a logistic distribution with location parameter χ_0 and scale parameter χ_1.[26] Except for the volatility of shock

[25] All of them lie in the plausible ranges estimated by Coibion and Gorodnichenko (2011). They follow the idea of Orphanides (2001) who argues that during the Great Inflation period, the Federal Reserve actually did respond to inflation given its best available information, and $\phi_{\pi,GI} > 1$. I take this view here.

[26] A logistic shock can be obtained from the difference of two i.i.d. type-I extreme value shocks. Its support includes negative values, which I bunch at zero. I use this to capture COLA adoption that is independent of the macroeconomic environment and instead reflects firm or union culture or inertia. Bunching negative values to zero does not affect the

Table 3: Calibrated parameters

Parameter	Value	Meaning
Panel A: Externally calibrated parameters		
β	0.99	Discount factor (quarterly)
γ	3	Worker risk aversion
$\varepsilon_x = \bar{\varepsilon}_y$	5	Elasticity of substitution
θ	0.75	Probability of not resetting prices
N	8	Contract length
σ_η	0.25	Std. dev. of firm idiosyncratic shock
$\phi_{\pi,GI}$	1.1	MP response to inflation, GI
$\phi_{\pi,GM}$	1.7	MP response to inflation, GM
ϕ_y	0.4	MP response to output gap
Panel B: Internally calibrated parameters		
ϕ_0	0.3027	Slope of the external financing cost
ν	0.0093	Slope of the labor adjustment cost
b	0.5087	Worker home production
ϱ_u	0.8487	AR(1) persistence of markup shocks
ϱ_ρ	0.8358	AR(1) persistence of demand shocks
$\sigma_{u,GI}$	0.0013	Std. of innovation to markup shocks, GI
$\sigma_{\rho,GI}$	0.005	Std. of innovation to demand shocks, GI
$\sigma_{u,GM}$	0.0007	Std. of innovation to markup shocks, GM
$\sigma_{\rho,GM}$	0.0019	Std. of innovation to demand shocks, GM
χ_0	0.0054	COLA cost distribution location parameter
χ_1	0.0028	COLA cost distribution scale parameter

Notes: Panel A reports the externally calibrated parameters, while panel B reports internally estimated parameters by the simulated method of moments. I allow the policy response to inflation and the volatility of shock innovations to vary across the Great Inflation and the Great Moderation. Other parameters are assumed to be stable across the two eras.

innovations, all other internally estimated parameters are assumed to be stable across the two eras.

The estimation follows two steps. In the first step, given the externally calibrated parameters, I assume the economy reaches its stationary stochastic equilibrium in both GI and GM, and the COLA share is set to the observed average, namely, 50 percent and 15 percent for GI and GM, respectively. I then choose all the non-cost parameters (i.e., except for χ_0, χ_1) to match model moments generated from the conditional GE with their data counterparts, which I will discuss momentarily. Importantly, the conditional GE is unique and is independent of the COLA-cost parameters.[27] In the second step, I choose the remaining cost parameters that govern COLA adoption to

calibration nor the results in Sections 6.3 and 6.4. It only affects welfare calculations, though the effect is very small. I also tried distributions supported only on positive values, such as the log-normal and gamma distributions. The quantitative implications barely change.

[27]The uniqueness of the conditional GE in the full model is a numerical statement. For the restricted model, I show the conditional GE is unique. For the general model, many different initial guesses converge to the same equilibrium. Also notice that this is not a statement of identification. Equilibrium uniqueness guarantees a unique mapping from structural parameters to equilibrium outcomes. Identification means a unique mapping from selected moments to structural parameters. However, I do verify that different initial guesses of the non-cost parameters converge when targeting the selected moments.

Table 4: Targeted moments

Targeted moments	Data values	Model values
Value of cash	1.13	1.15
Labor share	0.60	0.66
COLA pass-through	0.87	0.88
ACF(1) of real GDP	0.85	0.84
ACF(1) of inflation	0.85	0.88
Std. dev. of real GDP during GI	0.020	0.020
Std. dev. of inflation during GI	0.008	0.009
Std. dev. of real GDP during GM	0.009	0.010
Std. dev. of inflation during GM	0.003	0.004
Average COLA stock during GI	0.50	0.51
Average COLA stock during GM	0.15	0.15

Notes: This table reports the empirical value and the simulated value of those targeted moments.

guarantee the consistency of the observed average COLA shares with their respective stationary contracting equilibria. Given all other parameters, the two observed average COLA shares in the GI and GM uniquely pin down the two COLA-cost parameters. This strategy avoids the potential multiplicity issue, which might arise if one tries to estimate all parameters together by matching moments generated directly from the contracting GE to their data counterparts.

The first two columns of Table 4 report the targeted moments and their empirical values. The ACF(1) and the standard deviation of real GDP and inflation should directly inform the shock process. Their values are constructed at quarterly frequency. Real GDP is the HP-filtered cyclical component, and inflation is constructed from CPI and is defined as the quarter-to-quarter inflation rate.[28] The value of cash, defined by the marginal expected value of one unit of internal funds to firms, reflects the firm-side risk exposure (Gilchrist et al., 2017), and the COLA pass-through reflects the risk-sharing arrangements between firms and workers. Along with the shock process, they are informative about the slope of the external financing cost and the labor adjustment cost. Still, this discussion only provides some intuition as all the parameters are jointly estimated.

Panel B of Table 3 reports internally estimated parameters, and the third column of Table 4 reports the model-generated moments under the calibration. This model can fit the data relatively well. Notice that there is a noticeable decline in the volatility of shock innovations from GI to GM, which is consistent with the previous literature (Primiceri, 2005). I also report impulse response functions of inflation and output to one-standard-deviation markup and demand shocks in both eras in Appendix G.4. Finally, the estimated average cost of COLA adoption is modest. In consumption-equivalent terms, it roughly requires all covered workers to give up two days of their wages over a two-year contract. For non-contingent contracts to remain prevalent in times of low inflation risk despite such a small adoption cost, the private insurance value of COLA must itself be limited. The flip side of this result is that contract choice can respond sharply when inflation risk rises: a modest shift in the value of risk-sharing can move many firm-union pairs across

[28] I tried the GDP deflator, headline CPI, and the PCE price index, and the results are similar.

the adoption margin. Therefore, a small contractual friction can have sizable effects on aggregate inflation dynamics (as one will see shortly), echoing the insight from the menu-cost literature that small private costs may generate large macroeconomic consequences (Mankiw, 1985).

6.3 The Contribution of the Contract Accelerator to the Great Inflation

During the Great Inflation period, inflation risk rose sharply, COLA coverage started to increase, and concerns about wage-price spirals were central to contemporary policy debates (Nixon, 1971, Jacoby, 1987). I investigate the contribution of the contract accelerator to inflation through the lens of the model. I model the onset of the episode as an unexpected and permanent increase in the volatility of primitive shocks to the calibrated level in 1965Q1.

Baseline Counterfactual. I first recover the unique sequence of markup and demand shocks that rationalizes the observed paths of output and inflation in the calibrated Great Inflation economy. I then feed the same recovered shocks into a counterfactual economy in which the contract composition is fixed at its pre-shock level. Wage terms are still chosen optimally conditional on contract type. Thus, the only difference is that COLA adoption cannot increase after the volatility shock in the counterfactual economy.[29]

Figure 8 compares the observed path of inflation with the counterfactual path. The gap is concentrated in the two major inflation spikes, during which high COLA shares interact with large inflationary shocks. Quantitatively, the contract accelerator raises average annualized inflation by 1.10 percentage points during 1973-75 and by 1.02 percentage points during 1979-81, an amplification of about 14 percent across the two episodes. The amplification exceeds 20 percent at inflation peaks, precisely when a high COLA share interacts with large inflationary shocks.

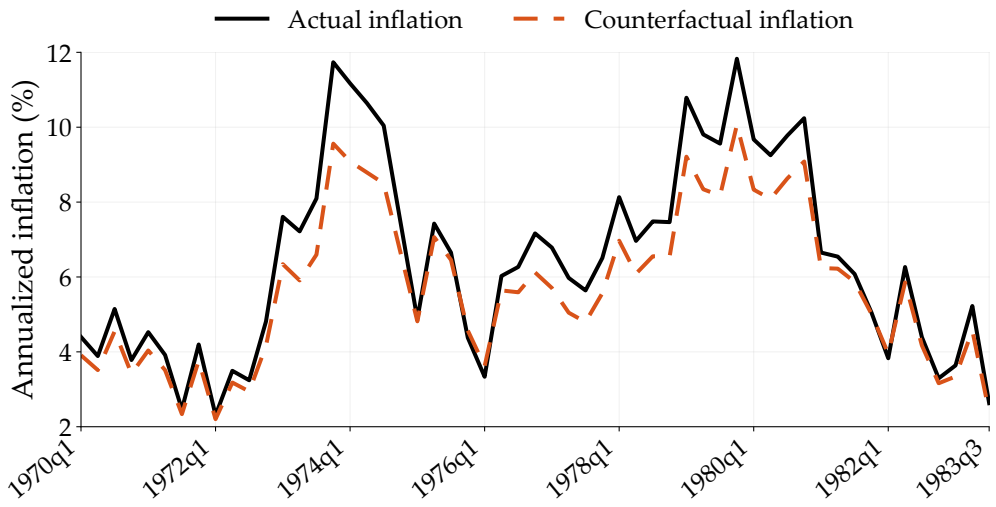
The baseline extrapolates the inflation-contingent wage setting measured in union contracts to the broader economy, assuming similar reactions through COLA or implicit indexation outside the union sector. Appendix G.5.1 considers two other assumptions for extrapolation. One is informed by the observed increase in both COLA coverage in unions and off-cycle wage raises in the broader economy during the post-pandemic inflation episode. I assume the ratio of their changes is stable over time, and use that to extrapolate for the Great Inflation period. The results are similar to the baseline extrapolation. The other is much more conservative and assumes the non-union sector only uses non-contingent wages. The amplification is undoubtedly smaller but still nontrivial.

Other Rational Expectations Equilibria. Although the calibrated policy rule ensures a unique conditional equilibrium given the observed COLA share during the Great Inflation, it does not necessarily ensure a unique contracting general equilibrium. I examine this issue here.

Holding fixed the structural parameters, the initial distribution of outstanding contracts, and the same shocks (including the volatility shock and the recovered markup and demand shocks), I

[29] Since the REE solution depends on past shocks up to N periods ago, I need to feed in some initial shocks. I set them to zero and discard the shocks from the first few years. The result is not sensitive to the initial values.

Figure 8: Path of inflation: actual and without COLA choice



Notes: The Great Inflation is modeled as a permanent and unexpected volatility shock that raises the volatility of primitive shocks to the calibrated level. The black solid line is the actual path of inflation in the US, and the red dashed line is the counterfactual path where COLA composition is fixed at the initial level, while conditional on the contract type, wage terms are still chosen optimally.

solve for alternative rational expectations equilibrium paths. Specifically, I initialize the solution algorithm described in Section 6.1 with two conjectures about future adoption (which determines the path of the contract distribution): one in which adoption remains at its pre-shock level, and another in which all newly signed contracts adopt COLA. These initial conjectures are updated iteratively, and in each resulting equilibrium path, the induced path from individual optimal choices is consistent with the conjectured path.[30]

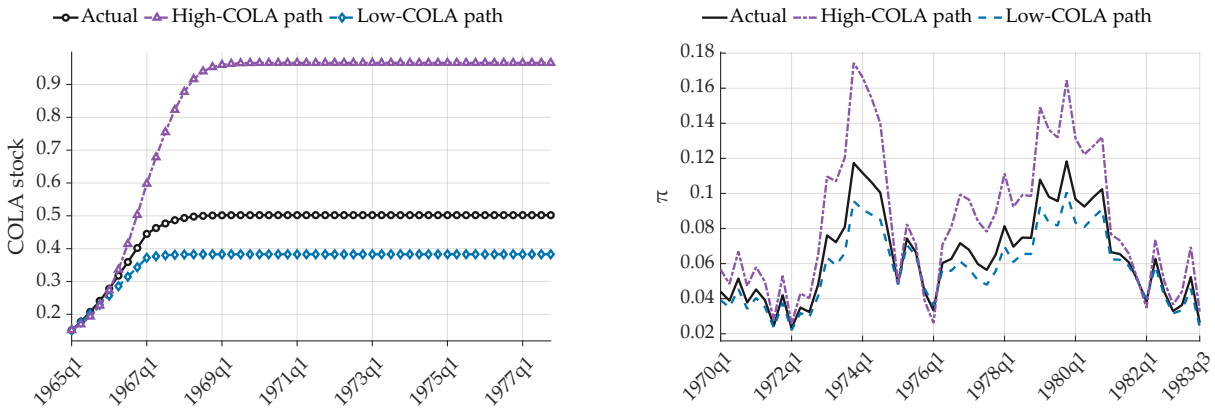
Figure 9 shows that the two initial conjectures lead to distinct equilibrium paths. Along the low-adoption path, the COLA share rises after the volatility shock, but remains below the observed path. Along the high-adoption path, it rises to about 95 percent, which in turn greatly amplifies the inflation response to the same markup and demand shocks.[31]

The high-COLA equilibrium illustrates the more extreme outcome that, fortunately, did not occur. It nevertheless reveals the fragility of the observed regime: had beliefs become sufficiently unanchored, the economy could have shifted to a self-fulfilling high-COLA path, and at the two inflation peaks, inflation would have been about 40 percent higher than observed. This highlights the importance of belief management on the margin of labor contract choice, especially in times of already high inflation risk.

[30] This algorithm that searches for a fixed point iteratively has a meaningful interpretation. The initial conjecture can be thought of as the initial belief of private agents. Given that belief, they choose their contract. The induced path serves as their new belief. This process repeats until their belief is consistent with the induced path, i.e., the belief is rational. This interpretation resembles the level-k thinking model (Nagel, 1995, Farhi and Werning, 2019).

[31] There are other equilibrium paths if one uses other initial conjectures. I numerically checked that the equilibrium path obtained from the all-COLA conjecture is the upper bound of all possible paths.

Figure 9: Paths of COLA (left) and inflation (right) after the GI shock in different equilibria



Notes: The left panel plots the equilibrium dynamic path of COLA shares after the GI volatility shock from different initial conjectures about the COLA adoption rate. The black line marked by circles is the actual path, the blue line marked by diamonds is the path where the initial conjecture remains at the pre-GI level in the stationary equilibrium, and the purple line marked by triangles is the path where the initial conjecture is that everyone else would always use COLA in the future. The right panel plots the paths of inflation implied by the three scenarios given the same recovered markup and demand shocks.

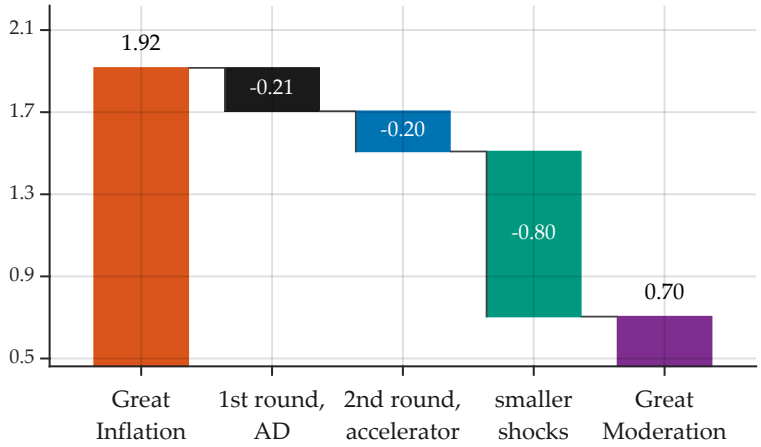
6.4 Monetary Policy Transmission through the Contract Accelerator in the Great Moderation

The Great Moderation, characterized by remarkable macroeconomic stability, provides the reverse experiment. Previous literature about the "policy vs. luck" debate suggests that luck—the exogenous decline in shock sizes—appears quantitatively more important in explaining the tranquility, compared to the more aggressive post-Volcker policy regime (Stock and Watson, 2003, Bernanke, 2004, Primiceri, 2005, Sims and Zha, 2006). However, conventional decomposition may underestimate the contribution from policy, as it weakens the contract accelerator and generates a second round of inflation stabilization. I revisit this debate in light of this new channel in this section.

In 1984Q1, I introduce into the model economy the more aggressive monetary policy rule and the lower volatility of primitive shocks from the calibration. Three things change in the model economy from the Great Inflation to the Great Moderation: the policy response to inflation becomes stronger, shock volatility falls, and the COLA share declines. The first two are exogenous to the economy, while COLA changes endogenously. To separate different channels, I decompose the *decline* in the on-impact inflation response to a typical shock as follows. I first change the policy rule, while holding the COLA share and the volatility of primitive shocks at their Great Inflation values. This isolates the effect of policy through the aggregate demand channel (AD). I then additionally allow the COLA share to respond to the more aggressive policy rule, which captures the effect of policy by dampening the contract accelerator. The residual goes to the exogenous decline in shock sizes.

Figure 10 reports the decomposition result for typical markup shocks. Through the AD channel alone, the more aggressive rule accounts for 17 percent of the decline in the inflation response

Figure 10: Decomposing the inflation responses to a one-std markup shock on impact



Notes: This figure decomposes the decline in the on-impact inflation response to a one-std markup shock from GI to GM. GM is modeled as a more aggressive policy rule and a decline in the volatility of primitive shocks. "AD" means the effect of policy through the aggregate demand channel, "accelerator" is the effect of policy by dampening the contract accelerator, and "smaller shocks" means the exogenous decline in shock sizes from GI to GM.

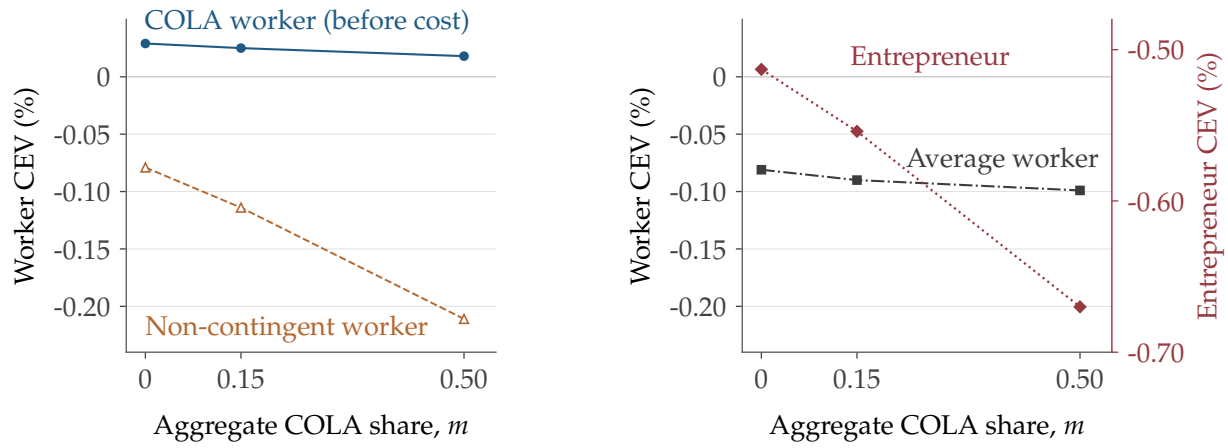
to a typical markup shock. By dampening the contract accelerator, policy explains an additional 17 percent. Thus, monetary policy accounts for one-third of the decline in the inflation response to a typical markup shock. The case for demand shocks is similar and is reported in Appendix G.6. Therefore, luck remains the dominant explanation for inflation stabilization during the Great Moderation, but the transmission channel through the contract accelerator gives policy more credit.

6.5 Welfare and Optimal Policy

The analysis so far focused on the positive implications of the contract accelerator. I close with a brief analysis of welfare and optimal policy. The central welfare tension follows directly from the key model mechanisms. Under incomplete insurance markets, COLA serves as a risk-sharing device and thus partially completes the missing markets. However, through nominal rigidities, COLA amplifies inflationary shocks and raises inflation risk. Higher inflation causes more misallocation through the standard channel, and more inflation risk hurts workers who are not protected by COLA. Firms and workers do not internalize those social costs when they choose their contracts. Hence, whether decentralized adoption decisions are socially desirable depends on the relative magnitude of the private gain and the aggregate externality. I use the Great Inflation calibration as a laboratory to compare the two forces, so all parameters and shocks remain those calibrated for that episode.

The left panel of Figure 11 plots the value expected by an individual worker under COLA (before the adoption cost) and under a non-contingent contract in three different scenarios for the aggregate COLA share, again in the same macro environment: COLA is completely banned (no COLA), COLA is fixed at the pre-GI level (pre-GI), and COLA responds endogenously to the equilibrium level where the gain from COLA offsets the adoption cost (GI eqm). The gap

Figure 11: Welfare comparison during the Great Inflation



Notes: The left panel plots the expected value of an individual worker under COLA (before the adoption cost) and under a non-contingent contract in different scenarios, including no COLA, pre-GI COLA, and COLA at its equilibrium level. The expected values are the values from the corresponding stationary equilibrium relative to the no-shock steady state, and are expressed in consumption-equivalent terms. The right panel plots the expected value for an average worker and for the entrepreneur.

between the two curves captures the private insurance value from COLA for a worker (before the adoption cost). The slope of the curves captures the macro externality from inflation risk: when more workers are covered by COLA, inflation becomes more volatile, which not only reduces workers' expected value even under COLA, but also greatly hurts those without COLA.

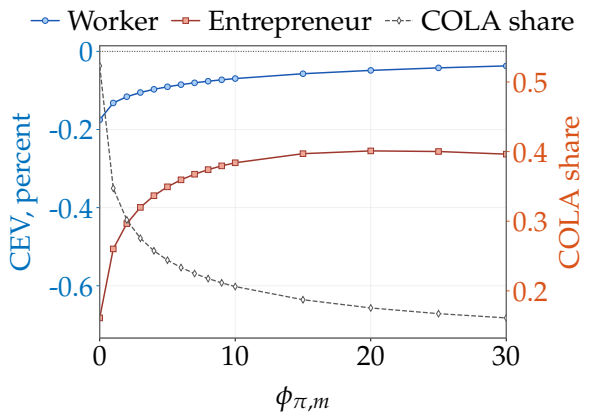
To assess which force dominates the welfare calculation, I plot the expected values of an average worker (before the adoption cost) and the representative entrepreneur in the right panel of Figure 11. During the Great Inflation, more COLA makes an average worker worse off, even before the COLA adoption cost is taken into account. This implies that the aggregate externality from inflation risk outweighs the private gain from inflation insurance. Moreover, higher inflation risk also raises resource misallocation and reduces output, so the entrepreneur is also worse off. Thus, under the Great-Inflation calibration, it turns out that the no-COLA case yields the highest value.[32]

A direct COLA cap can eliminate the wedge. Appendix H.3 reports the experiment. In practice, however, those policies may not be able to prevent implicit indexation. Therefore, I study the role of monetary policy next.

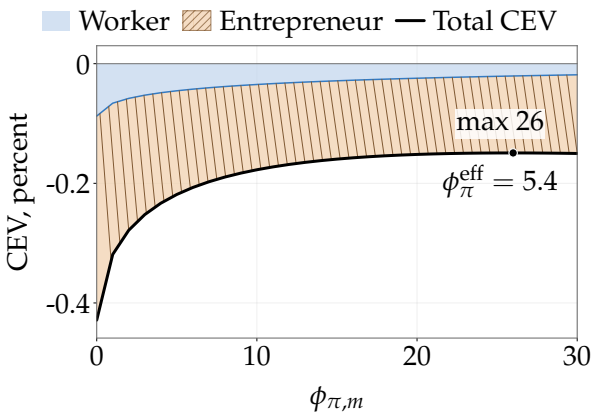
[32] In Appendix H.4, I also report how an average worker's expected value changes as I increase the COLA share up to one (i.e., compulsory COLA adoption). The gross value exhibits a U shape. It starts to increase when a sufficiently large share of workers is covered by COLA, because the aggregate externality mainly hurts non-adopters. However, the net value after accounting for the COLA cost decreases monotonically with the aggregate COLA share. Meanwhile, this aggregate externality always hurts the entrepreneur. Hence, the optimal COLA share is indeed zero under the Great-Inflation calibration.

Figure 12: Welfare under the COLA-augmented Taylor rule

(a) Welfare net of COLA adoption cost



(b) Average welfare



COLA-Augmented Taylor Rule. Theory suggests that the strength of the monetary policy response to inflation should depend on the indexation regime. To make this idea operational, I consider a simple COLA-augmented Taylor rule

$$i_t = (\phi_{\pi 0} + \phi_{\pi m} m_t)\, \pi_t + \phi_y \hat{y}_t \tag{12}$$

where $\phi_{\pi m} \geq 0$. Thus, the policy response to inflation becomes stronger as the COLA share rises. I set $\phi_{\pi 0} = 1.1$, as in the Great Inflation calibration, and vary $\phi_{\pi m}$.

I show in Appendix H.6 that the rule in (12) with a positive $\phi_{\pi m}$ can effectively change the equilibrium COLA path and even ensure uniqueness during the Great Inflation. Here, I evaluate welfare in different equilibria as $\phi_{\pi m}$ increases. The left panel of Figure 12 reports separately the welfare of an average worker net of COLA adoption costs and the welfare of the entrepreneur, both in consumption-equivalent units. An average worker benefits from a more aggressive rule that lowers inflation risk. The welfare of the entrepreneur initially rises as inflation is stabilized, but it eventually falls as a more aggressive response hurts output. Thus, the two groups value policy differently when it is sufficiently aggressive.

Given the two heterogeneous groups, comparing aggregate welfare across different economies requires taking a stand on the social welfare weights. As an illustration, I assign equal weight to a one-percent consumption-equivalent change in their respective welfare. The right panel of Figure 12 plots the two weighted welfare components and their sum. The optimal policy balances the worker's gain from lower inflation risk against the entrepreneur's loss from greater output risk. Appendix H.6 reports the results under different social welfare weights.

Under this criterion, the optimal augmented rule implies that the effective inflation coefficient $\phi_\pi^{eff} := \phi_{\pi 0} + \phi^*_{\pi m} m_t(\phi^*_{\pi m})$ is approximately 5.4. The large implied response to inflation should be interpreted with caution. The model abstracts from other concerns associated with a very aggressive policy rule, e.g., household and firm financial vulnerability due to very volatile nominal

interest rates. This exercise nevertheless illustrates the welfare impact of the externality from the contract accelerator, which is large during the Great Inflation through the lens of the model.

6.6 Extension: Endogenously Stochastic Inflation Volatility

As mentioned in Section 6.1, the solution of the baseline model relies on the separation of first- and second-order moments in determining the contract distribution. Removing this separation can have interesting implications. In Appendix I.2, I discuss one departure from the baseline quantitative model by allowing inflation levels to enter the expected gain from COLA for the renewing cohort. In that case, an inflationary shock not only raises inflation, but also raises COLA adoption, which in turn amplifies inflation volatility by triggering the contract accelerator. Thus, inflation volatility can be endogenously stochastic even if the volatility of primitive shocks is fixed, providing one structural explanation of why high inflation and high inflation volatility often appear together (Primiceri, 2005). The cost is that the contract distribution becomes a genuinely stochastic state, so the clean finite moving-average representation in (9) no longer applies. In Appendix I.2, I solve the model by the method of parametrized expectations (den Haan and Marcet, 1990) and show that the simulated economy can deliver the expected results. However, the computation for the case of staggered two-period contracts is already heavy and less transparent. Therefore, I leave the extension of this mechanism to the full quantitative environment for future work.

7 Concluding Remarks

Can the interaction between endogenous wage indexation and inflation dynamics lead to wage-inflation spirals? The answer is yes. This paper shows that private risk-sharing, optimally chosen, can lead to sizable macroeconomic effects. A more inflation-contingent wage setting steepens the Phillips curve and amplifies inflation risk, which in turn reinforces the adoption of inflation-contingent contracts. The contract accelerator makes contract choices strategic complements, and can potentially support multiple self-fulfilling indexation equilibria. Through the lens of the model, the contract accelerator magnifies inflation by 14 percent during the 1973-75 and 1979-81 inflation spikes. However, if beliefs about the aggregate adoption rate of indexed contracts had become completely unanchored during the Great Inflation, the inflation peaks could have been 40 percent higher than observed.

Can a more aggressive monetary policy stance prevent such spirals? The answer is also yes. A more hawkish policy rule stabilizes inflation twice: first through the conventional aggregate demand channel, and second by dampening the contract accelerator. Quantitatively, the size of the second-round stabilization is almost the same as that of the first-round stabilization for Volcker-type policy during the Great Moderation. A more hawkish policy also weakens strategic complementarity in contract choice and can anchor the economy in a low-indexation regime. The flip side of the coin is that if policymakers overlook the endogenous response of private contract choices, the economy can be vulnerable to belief and/or uncertainty shocks.

This paper also opens up many interesting directions for future research. First, this paper uses collective bargaining agreements as the main measurement device. I view other responses like off-cycle wage raises through renegotiation or shorter contract durations as implicit indexation, because they effectively make wages more responsive to inflation. I provide some suggestive evidence on them in the broader economy and conduct some conservative robustness checks, but more direct and systematic evidence about those responses in the broader labor market is desirable.

Second, the limits on contractibility and the COLA adoption cost still call for deeper microfoundations. Why are revenue- or employment-contingent contracts rare? I provide an information-based explanation for the scarcity of revenue-contingent contracts, but other explanations are also possible. As for the cost of adopting COLA, I take a shortcut by assuming a reduced-form cognitive calculation cost. This may also represent some deeper contractual frictions. Understanding the cost and its heterogeneity is important for welfare calculations and policy design. For example, while compulsory COLA adoption may insure all workers against inflation risk, some workers and firms with huge adoption costs can be hurt.

Third, this paper demonstrates the importance of beliefs about the COLA adoption rate in determining the equilibrium. Thus, it is important to model belief formation more explicitly, which, for example, may help select a unique equilibrium and yield unambiguous analysis. Moreover, I impose full-information rational expectations in the model economy. A departure from these strong assumptions can also be interesting.

Fourth, the baseline model separates first-order shocks from the evolution of the contract distribution. In an extension, I show that without this separation, the contract choice channel can generate endogenous stochastic inflation volatility, even if primitive volatility is constant. It also generates positive comovements between inflation and inflation volatility, a robust empirical feature in the data (Primiceri, 2005). However, due to its computational complexity, I only treat the exercise as an illustration. A formal quantitative assessment with more advanced computational tools seems promising.

Fifth, I shut down worker mobility in the labor market for tractability. Unions in the model do not have direct interactions (although they do have indirect interactions through contract adoption). Adding more labor market features would not only make the model more realistic but also allow them to interact with contract choice. For example, COLA may have a spillover effect on non-adopters through wage competition.

References

Hassan Afrouzi, Andrés Blanco, Andrés Drenik, and Erik Hurst. A theory of how workers keep up with inflation. Working Paper 33233, National Bureau of Economic Research, 2024.

Narek Alexanian. Endogenous wage indexation and aggregate supply. Job market paper, 2025.

George-Marios Angeletos and Jennifer La'O. Noisy business cycles. Working Paper 14982, National Bureau of Economic Research, 2010.

George-Marios Angeletos and Jennifer La'O. Sentiments. *Econometrica*, 81(2):739–779, 2013.

Costas Azariadis. Escalator clauses and the allocation of cyclical risks. *Journal of Economic Theory*, 18(1):119–155, 1978.

Costas Azariadis and Russell Cooper. Predetermined prices and the allocation of social risks. *The Quarterly Journal of Economics*, 100(2):495–518, 1985.

Laurence Ball. Is equilibrium indexation efficient? *The Quarterly Journal of Economics*, 103(2): 299–311, 1988.

Robert J. Barro. Indexation in a rational expectations model. *Journal of Economic Theory*, 13(2): 229–244, 1976.

Jess Benhabib and Roger E. A. Farmer. Indeterminacy and increasing returns. *Journal of Economic Theory*, 63(1):19–41, 1994.

Ben S. Bernanke. The Great Moderation. Remarks at the meetings of the Eastern Economic Association, Washington, DC, February 2004.

Olivier J. Blanchard. The wage price spiral. *The Quarterly Journal of Economics*, 101(3):543–566, 1986.

Olivier Jean Blanchard. Wage indexing rules and the behavior of the economy. *Journal of Political Economy*, 87(4):798–815, 1979.

Olivier Jean Blanchard and Charles M. Kahn. The solution of linear difference models under rational expectations. *Econometrica*, 48(5):1305–1311, 1980. doi: 10.2307/1912186.

Nicholas Bloom. The impact of uncertainty shocks. *Econometrica*, 77(3):623–685, 2009.

Nicholas Bloom, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J. Terry. Really uncertain business cycles. *Econometrica*, 86(3):1031–1065, 2018.

Board of Governors of the Federal Reserve System. The Beige Book: Current economic conditions by Federal Reserve district, 2022.

David Card. Cost-of-living escalators in major union contracts. *ILR Review*, 37(1):34–48, 1983.

David Card. An empirical model of wage indexation provisions in union contracts. *Journal of Political Economy*, 94(3, Part 2):S144–S175, 1986.

Julio A. Carrillo, Gert Peersman, and Joris Wauters. Endogenous wage indexation and aggregate shocks. *Journal of Macroeconomics*, 72:103417, 2022.

Lawrence J. Christiano, Martin Eichenbaum, and Charles L. Evans. Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy*, 113(1):1–45, 2005.

Richard Clarida, Jordi Galí, and Mark Gertler. Monetary policy rules and macroeconomic stability: Evidence and some theory. *The Quarterly Journal of Economics*, 115(1):147–180, 2000.

Olivier Coibion and Yuriy Gorodnichenko. Monetary policy, trend inflation, and the Great Moderation: An alternative interpretation. *American Economic Review*, 101(1):341–370, 2011.

Russell Cooper and Andrew John. Coordinating coordination failures in Keynesian models. *The Quarterly Journal of Economics*, 103(3):441–463, 1988.

Wouter J. den Haan and Albert Marcet. Solving the stochastic growth model by parameterizing expectations. *Journal of Business & Economic Statistics*, 8(1):31–34, 1990.

Ronald G. Ehrenberg, Leif Danziger, and Gee San. Cost-of-living adjustment clauses in union contracts: A summary of results. *Journal of Labor Economics*, 1(3):215–245, 1983.

Christopher J. Erceg, Dale W. Henderson, and Andrew T. Levin. Optimal monetary policy with staggered wage and price contracts. *Journal of Monetary Economics*, 46(2):281–313, 2000.

Emmanuel Farhi and Iván Werning. Monetary policy, bounded rationality, and incomplete markets. *American Economic Review*, 109(11):3887–3928, 2019. doi: 10.1257/aer.20171400.

Stanley Fischer. Wage indexation and macroeconomic stability. *Carnegie-Rochester Conference Series on Public Policy*, 5:107–147, 1977.

Stanley Fischer. Indexing and inflation. *Journal of Monetary Economics*, 12(4):519–541, 1983.

Joel P. Flynn, George Nikolakoudis, and Karthik A. Sastry. A theory of supply function choice and aggregate supply. *American Economic Review*, 116(2):710–748, 2026a. doi: 10.1257/aer.20240443.

Joel P. Flynn, George Nikolakoudis, and Karthik A. Sastry. Organizational targets in general equilibrium. Working Paper 35099, National Bureau of Economic Research, 2026b.

Xavier Gabaix and Thomas Graeber. The complexity of economic decisions. Working Paper 33109, National Bureau of Economic Research, 2024.

Simon Gilchrist, Raphael Schoenle, Jae Sim, and Egon Zakrajsek. Inflation dynamics during the financial crisis. *American Economic Review*, 107(3):785–823, 2017.

Maximilian Gödl, Isabel Gödl-Hanisch, and Simon Velten. Wage setting in times of high and low inflation. Working Paper 11319, CESifo, 2026.

Jo Anna Gray. Wage indexation: A macroeconomic approach. *Journal of Monetary Economics*, 2(2): 221–235, 1976.

Roger Guesnerie. An exploration of the eductive justifications of the rational-expectations hypothesis. *American Economic Review*, 82(5):1254–1278, 1992.

Wallace E. Hendricks and Lawrence M. Kahn. *Wage Indexation in the United States: COLA or Un-COLA?* Ballinger Publishing Company, Cambridge, MA, 1985.

Erik Hurst, Christina Patterson, Nela Thomas Richardson, and Ye Liv Wang. Sticky wage norms and the real wage cost of unexpected inflation. Working Paper 35624, National Bureau of Economic Research, 2026.

Sanford M. Jacoby. The development of cost-of-living escalators in the United States. *Labor History*, 28(4):515–533, 1987.

Kyle Jurado, Sydney C. Ludvigson, and Serena Ng. Measuring uncertainty. *American Economic Review*, 105(3):1177–1216, 2015.

Harry C. Katz. *Shifting Gears: Changing Labor Relations in the U.S. Automobile Industry*. MIT Press, Cambridge, MA, 1985.

Douglas R. LeRoy. Scheduled wage increases and cost-of-living provisions in 1982. *Monthly Labor Review*, 105(1):16–20, January 1982.

Guido Lorenzoni and Iván Werning. Wage price spirals. *Brookings Papers on Economic Activity*, (2), 2023.

N. Gregory Mankiw. Small menu costs and large business cycles: A macroeconomic model of monopoly. *The Quarterly Journal of Economics*, 100(2):529–537, 1985.

N. Gregory Mankiw and Ricardo Reis. Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips curve. *The Quarterly Journal of Economics*, 117(4):1295–1328, 2002.

Alejandro Martínez-Marquina, Muriel Niederle, and Emanuel Vespa. Failures in contingent reasoning: The role of uncertainty. *American Economic Review*, 109(10):3437–3474, 2019.

Stephen Morris and Hyun Song Shin. Global games: Theory and applications. In *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress*, pages 56–114. Cambridge University Press, 2003.

Giuseppe Moscarini and Fabien Postel-Vinay. The job ladder: Inflation vs. reallocation. Working Paper 31466, National Bureau of Economic Research, 2023.

Rosemarie Nagel. Unraveling in guessing games: An experimental study. *American Economic Review*, 85(5):1313–1326, 1995.

Richard Nixon. Address to the nation outlining a new economic policy: The challenge of peace. The American Presidency Project, August 1971.

Ryan Oprea. What makes a rule complex? *American Economic Review*, 110(12):3913–3951, 2020.

Athanasios Orphanides. Monetary policy rules based on real-time data. *American Economic Review*, 91(4):964–985, 2001.

Andrew J. Oswald. The economic theory of trade unions: An introductory survey. *The Scandinavian Journal of Economics*, 87(2):160–193, 1985.

Laura Pilossoph and Jane M. Ryngaert. Job search, wages, and inflation. Working Paper 33042, National Bureau of Economic Research, 2024.

Giorgio E. Primiceri. Time varying structural vector autoregressions and monetary policy. *The Review of Economic Studies*, 72(3):821–852, 2005.

Christina D. Romer and David H. Romer. A new measure of monetary shocks: Derivation and implications. *American Economic Review*, 94(4):1055–1084, 2004.

Christopher A. Sims and Tao Zha. Were there regime switches in U.S. monetary policy? *American Economic Review*, 96(1):54–81, 2006.

Frank Smets and Rafael Wouters. Shocks and frictions in U.S. business cycles: A Bayesian DSGE approach. *American Economic Review*, 97(3):586–606, 2007.

James H. Stock and Mark W. Watson. Has the business cycle changed? Evidence and explanations. In *Monetary Policy and Uncertainty: Adapting to a Changing Economy*, pages 9–56. Federal Reserve Bank of Kansas City, 2003.

Online Appendix for Wage-Price Spirals: The Contract Accelerator

Xilin Yuan

A Data and Measurement

This section describes in more detail the contract-level data and the measurement of wage indexation. The main dataset consists of the records of major collective bargaining agreements (CBAs) collected by the Bureau of Labor Statistics. These records are useful because they directly contain information about explicit cost-of-living adjustment clauses and, in many cases, report the formula that governs wage adjustments to prices.

A.1 Sources and Examples

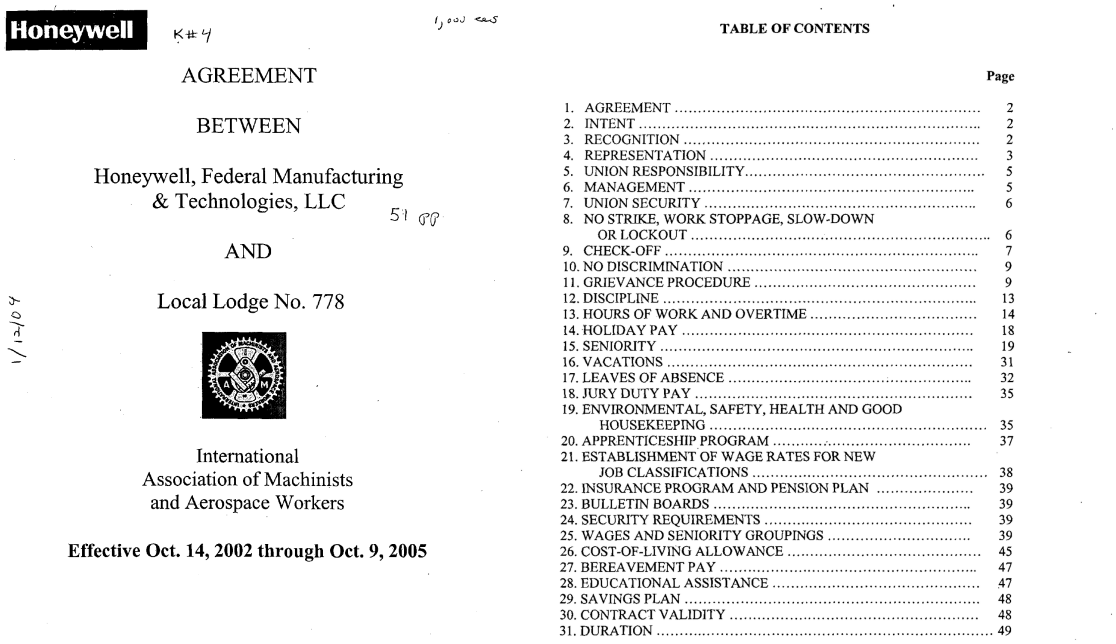
The baseline sample is called the Major Collective Bargaining Settlements Restricted File for 1974-1995. The file covers major CBAs, typically those covering at least 1,000 workers, and records the name of the bargaining pair, negotiation and effective dates, contract duration, employment covered, scheduled wage changes, and the presence and design of COLA clauses. I use the negotiation date as the relevant timing for contract choice, because it is the date at which the parties choose whether future wage adjustment will be predetermined or contingent on realized inflation.

Figure 13 illustrates the front page of one CBA, and all CBAs underlying the main records have a similar structure. The front page shows the bargaining parties and the contract duration, while the table of contents shows what they have bargained and the corresponding results, including wages, hours, seniority, holidays, and cost-of-living provisions.

Figure 14 shows examples of the two wage-setting arrangements emphasized in the paper. Scheduled-wage increases offer a preset nominal wage path at signing, often through annual across-the-board increases, while some contracts have detailed wage tables for different occupations. A COLA provision instead allows inflation-contingent adjustments over the life of the contract. For example, the COLA clause on the right panel has a quarterly adjustment frequency, and the formula is that every 0.4 point increase in CPI raises a worker's hourly wage by 1 cent. This is a typical COLA formula used by those CBAs. The two arrangements can appear in the same agreement. I will go back to this point in Section A.4 and argue that they substitute on the margin of inflation adjustment.

Finally, I also link the data to Compustat to see how COLA adoption and COLA formula may correlate with firm financial conditions. I do a fuzzy match based on employer names, industries,

Figure 13: A typical collective bargaining agreement



Honeywell

AGREEMENT

BETWEEN

Honeywell, Federal Manufacturing & Technologies, LLC

AND

Local Lodge No. 778

International Association of Machinists and Aerospace Workers

Effective Oct. 14, 2002 through Oct. 9, 2005

TABLE OF CONTENTS

	Page
1. AGREEMENT	2
2. INTENT	2
3. RECOGNITION	2
4. REPRESENTATION	3
5. UNION RESPONSIBILITY	5
6. MANAGEMENT	5
7. UNION SECURITY	6
8. NO STRIKE, WORK STOPPAGE, SLOW-DOWN OR LOCKOUT	6
9. CHECK-OFF	7
10. NO DISCRIMINATION	9
11. GRIEVANCE PROCEDURE	9
12. DISCIPLINE	13
13. HOURS OF WORK AND OVERTIME	14
14. HOLIDAY PAY	18
15. SENIORITY	19
16. VACATIONS	31
17. LEAVES OF ABSENCE	32
18. JURY DUTY PAY	35
19. ENVIRONMENTAL, SAFETY, HEALTH AND GOOD HOUSEKEEPING	35
20. APPRENTICESHIP PROGRAM	37
21. ESTABLISHMENT OF WAGE RATES FOR NEW JOB CLASSIFICATIONS	38
22. INSURANCE PROGRAM AND PENSION PLAN	39
23. BULLETIN BOARDS	39
24. SECURITY REQUIREMENTS	39
25. WAGES AND SENIORITY GROUPINGS	39
26. COST-OF-LIVING ALLOWANCE	45
27. BEREAVEMENT PAY	47
28. EDUCATIONAL ASSISTANCE	47
29. SAVINGS PLAN	48
30. CONTRACT VALIDITY	48
31. DURATION	49

and states. I then merge the matched bargaining-pair identifiers to annual Compustat variables and negotiation year.

A.2 Calculating COLA Pass-Through

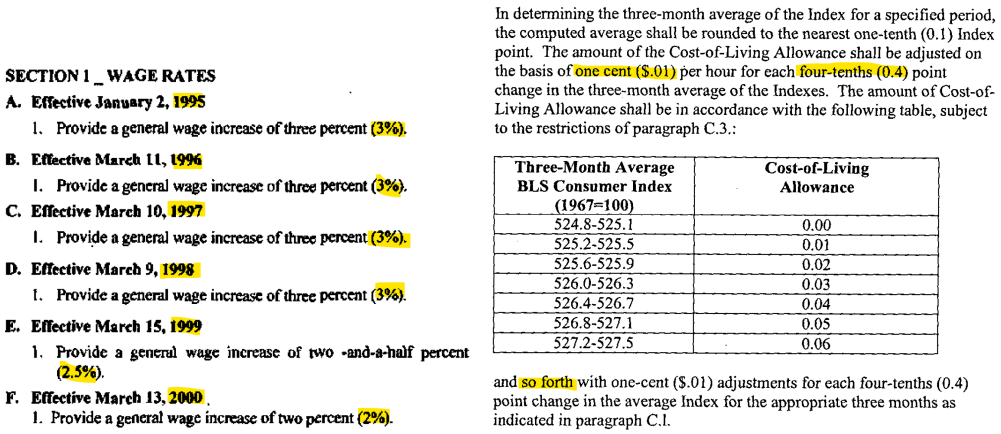
The data often reports the COLA formula if COLA is adopted, as the example in Section A.1. By itself, one cannot tell how generous the compensation is, because the inflation implied by a 0.4 point increase in CPI and the purchasing power of one cent can be very different across years. Table 5 reports the distribution of the main formula types. About 60 percent of COLA contracts contain a directly usable formula. For contracts with COLA but without a record of an explicit formula, the dataset does record the wage changes generated by COLA, which are separated from the scheduled wage increases.

To compare the generosity of different COLA contracts, I convert COLA formulas and recorded COLA wage changes into a standardized price-to-nominal-wage pass-through measure. That is, for a one percent increase in the price index (COLA contracts will specify a particular type of price index, mostly CPI-W), by what percentage does the nominal hourly wage increase? It is unit-free and thus comparable across contracts with different base wages and CPI base years.

Specifically, for contracts with an explicit formula, I compute the formula-implied pass-through at the time of negotiation.[33] Denote by W_0 the base hourly wage in cents and P_0 the price level with the base year used by the contract. If the clause pays b cents for each x point increase in the price index, the implied pass-through is

[33] Different timing choices yield almost no difference.

Figure 14: Examples of scheduled wage increases and COLA



SECTION 1 _ WAGE RATES

A. Effective January 2, 1995

1. Provide a general wage increase of three percent (3%).

B. Effective March 11, 1996

1. Provide a general wage increase of three percent (3%).

C. Effective March 10, 1997

1. Provide a general wage increase of three percent (3%).

D. Effective March 9, 1998

1. Provide a general wage increase of three percent (3%).

E. Effective March 15, 1999

1. Provide a general wage increase of two -and-a-half percent (2.5%).

F. Effective March 13, 2000

1. Provide a general wage increase of two percent (2%).

In determining the three-month average of the Index for a specified period, the computed average shall be rounded to the nearest one-tenth (0.1) Index point. The amount of the Cost-of-Living Allowance shall be adjusted on the basis of one cent ($.01) per hour for each four-tenths (0.4) point change in the three-month average of the Indexes. The amount of Cost-of-Living Allowance shall be in accordance with the following table, subject to the restrictions of paragraph C.3.:

Three-Month Average BLS Consumer Index (1967=100)	Cost-of-Living Allowance
524.8-525.1	0.00
525.2-525.5	0.01
525.6-525.9	0.02
526.0-526.3	0.03
526.4-526.7	0.04
526.8-527.1	0.05
527.2-527.5	0.06

and so forth with one-cent ($.01) adjustments for each four-tenths (0.4) point change in the average Index for the appropriate three months as indicated in paragraph C.1.

$$\tau^{formula} = \frac{b/W_0}{x/P_0}$$

For clauses written in percentage-point terms, one just replaces the denominator with the recorded percentage change. A few clauses directly specify percentage-to-percentage formulas, which are already pass-through measures.

For COLA contracts without a directly usable formula, I use recorded COLA wage changes. The data report COLA-generated wage adjustments in cents at several points during the contract. I sum these recorded adjustments to obtain the total COLA wage change over the life of the contract, ΔW^{COLA}. I then compute

$$\tau^{doc} = \frac{\Delta W^{COLA}/W_0}{\Delta P/P_0}$$

where ΔP is the cumulative growth in the price index over the same horizon.

Figure 15 plots the cross-sectional distribution of COLA pass-through. COLA pass-through is high but incomplete: the mean is about 0.89, the median is about 0.86. Thus COLA clauses typically provide substantial inflation compensation, but they are not mechanically equivalent to full real-wage indexation. The pattern is almost identical if one plots the same figure with the formula-based measure only.

A.3 Summary Statistics

Table 6 summarizes the baseline sample. The underlying contracts are often long term and last about 29 months on average. Bargaining units are typically large. The median contract covers 1,850 workers, and the mean is higher because a small number of agreements cover very large units. Scheduled wage changes are universal and appear in 93 percent of contracts, while 22 percent include a COLA clause. Conditional on COLA, the average size of the bargaining units

Table 5: CPI formula distribution

CPI Formula	Freq.	Percent	Cum.
1 cent/0.3 point rise	1,216	26.12	26.12
1 cent/0.35 point rise	26	0.56	26.68
1 cent/0.4 point rise	1,065	22.88	49.56
1 cent/0.45 point rise	3	0.06	49.62
1 cent/0.5 point rise	59	1.27	50.89
1 cent/0.3% rise	114	2.45	53.34
1 cent/0.4% rise	46	0.99	54.33
1 cent/0.5% rise	10	0.21	54.54
1%/1% rise	231	4.96	59.51
Others	1,885	40.49	100.00
Total	4,655	100.00	

Notes: The benchmark CPI series is CPI-W. "Others" include AT&T-style rules and unnamed wage-change methods for which the formula is not directly converted into cents per CPI point.

is even larger, suggesting that COLA is more likely in large firms and unions. The variation in COLA pass-through is limited. Most contracts have substantial but incomplete pass-through.

Table 6: Summary statistics

	N	Mean	SD	P25	P50	P75
Panel A: Full sample						
Contract duration (months)	21,543	28.8	11.0	24	36	36
Union size	21,543	4,190.6	14,109.0	1,240	1,850	3,500
Scheduled wage change (=1)	21,543	0.93	0.26	1	1	1
COLA (=1)	21,543	0.22	0.42	0	0	0
Panel B: Conditional on COLA						
COLA duration (months)	4,676	20.2	11.3	12	21	30
Union size	4,798	6,937.3	27,229.2	1,300	2,100	4,400
COLA pass-through	3,984	0.89	0.42	0.72	0.86	1

A.4 Substitution between COLA and Scheduled Wages in Inflation Adjustment

As shown in Section A.3, scheduled wage increases are almost universal and thus the default wage-setting arrangement in the sample. COLA, if adopted, is layered on top of it. Both of them allow nominal wage adjustments during the life of a contract. But they differ mainly in inflation contingency, and also in frequency. Scheduled increases are fixed at signing and only compensate for expected developments, including expected inflation, whereas COLA adjusts with inflation realizations. Thus, conceptually, if the two arrangements are used partly to compensate for future inflation, contracts with COLA should require smaller scheduled increases.

Figure 15: COLA pass-through distribution

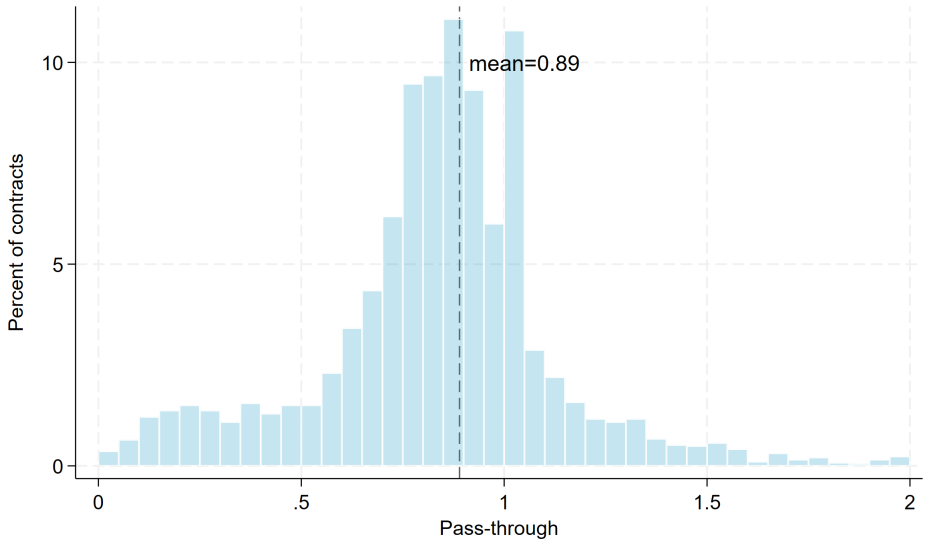


Table 7 provides evidence for this substitution. The dependent variable is annualized scheduled nominal wage growth, and the key regressor is an indicator for whether the contract contains COLA. The OLS specifications include bargaining-pair fixed effects. They show that, within the same bargaining pair, scheduled wage increases are about 0.8 to 1.1 percentage points lower in contracts with COLA compared to those without. Of course, COLA adoption is endogenous. So, the table also reports a 2SLS specification where COLA adoption is instrumented with the lagged Baumeister-Hamilton oil-supply shocks. It shifts the value of inflation contingency and hence COLA adoption. The IV estimate has the same direction but is significantly larger in magnitude. Taken together, the estimates motivate a discrete contract choice between scheduled increases and COLA on the margin of inflation compensation.

Table 7: COLA and scheduled wage increases

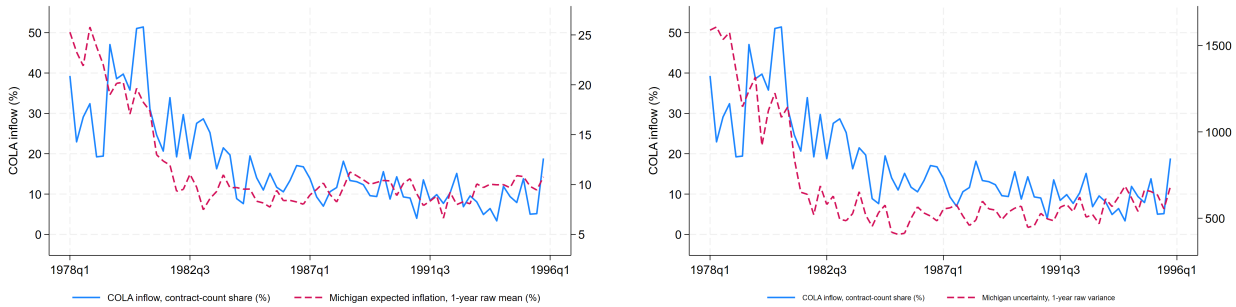
	OLS (1)	OLS (2)	2SLS (3)
COLA	-0.801	-1.121	-6.734
	(0.085)	(0.215)	(3.279)
Pair FE	✓	✓	✓
Financial controls		✓	✓
N	20,790	4,537	3,394

Notes: The dependent variable is annualized scheduled nominal wage growth. Columns (1) and (2) report OLS estimates in the full and Compustat-linked samples. Column (3) instruments COLA adoption with lagged Baumeister-Hamilton oil-supply shocks in the Compustat-linked sample. All specifications include negotiation-year and bargaining-pair fixed effects; the Compustat-linked specifications include firm financial controls. Standard errors are reported in parentheses.

A.5 Subjective Inflation Expectations and COLA Adoption

In the main text, I show the relationship between COLA and the realized inflation and a statistical measure of inflation volatility. Figure 16 provides a complementary view using subjective expectations from the Michigan Survey of Consumers. The left panel plots the COLA inflow rate against the mean one-year-ahead inflation forecast. The right panel plots the same COLA inflow rate against the cross-sectional variance of one-year-ahead forecasts. These measures provide more direct correspondences to the contract choice problem in which the perceived inflation dynamics at the individual level matter for their choices. This downside is that the data has a shorter horizon, as the survey started in the late 1970s. Still, the same pattern appears: COLA adoption is higher when inflation is expected to be high and when inflation forecasts are more dispersed.

Figure 16: COLA adoption and subjective inflation expectations



Notes: The left panel compares the COLA inflow rate with the mean one-year-ahead inflation forecast in the Michigan Survey of Consumers. The right panel compares the COLA inflow rate with the cross-sectional variance of one-year-ahead inflation forecasts. The figure shows that the adoption margin also comoves with subjective inflation expectations and uncertainty, complementing the realized-inflation and statistical-volatility evidence in the main text.

A.6 Recent Time Path of COLA Adoption

The main data source ends in 1995, as BLS stopped collecting those CBAs after that. Still, there are other sources to collect those CBAs, especially for recent years, though the coverage is not as universal as the main dataset. This subsection shows evidence from post-2000 US contracts and from other countries like Spain and Canada, which covers the post-pandemic inflation episode.

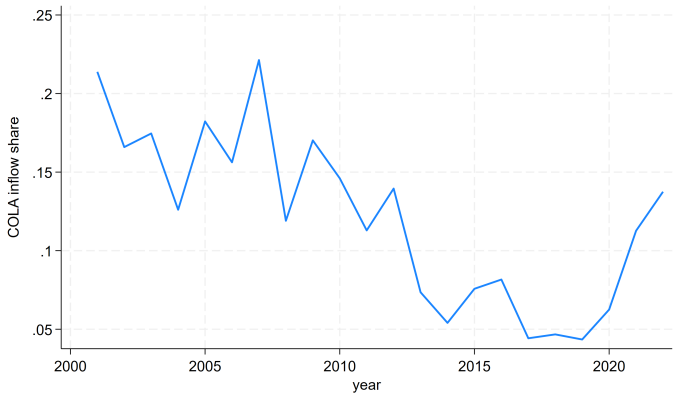
US Post-2000. US CBAs in recent years are collected by the Office of Labor-Management Standards. Part of these files were formerly collected by the BLS and were transferred to this agency. It also collects CBAs itself from firms and unions. These agreements, as scanned files, are publicly available. I digitize the files and classify a contract as containing COLA using a combination of text extraction and keyword search. I assign each contract to a year using the extracted starting

date. If that's missing, I impute from the listed expiration date from the metadata provided by OLMS.[34]

Figure 17 plots the annual share of newly recorded private agreements with COLA. The series is relatively stable before the Great Recession. It reaches very low levels in the late 2010s, possibly due to the recession. It rises again during the recent inflation episode.

Figure 17 shows the inflow of COLA contracts. I next convert it into the stock, making it more comparable to the series shown in the main text.[35] Figure 18 plots the time series of COLA coverage across different inflation episodes, including the coverage series from the main dataset (blue) and from the OLMS dataset (green), together against inflation and inflation volatility. Similar patterns repeat across those episodes: COLA coverage is high when inflation is high and volatile. As for the magnitude, COLA coverage increased from about 10 percent to 17 percent during the recent inflation episode.[36] The caveat is that submitting agreements to the OLMS archive is voluntary, so the post-2000 sample is not as universal as the main dataset.

Figure 17: Share of new contracts with COLA provisions in the US



Notes: The figure plots the share of newly recorded private-sector collective bargaining agreements with COLA provisions from 2001 to 2022. The data are digitized from agreements collected by the Office of Labor-Management Standards.

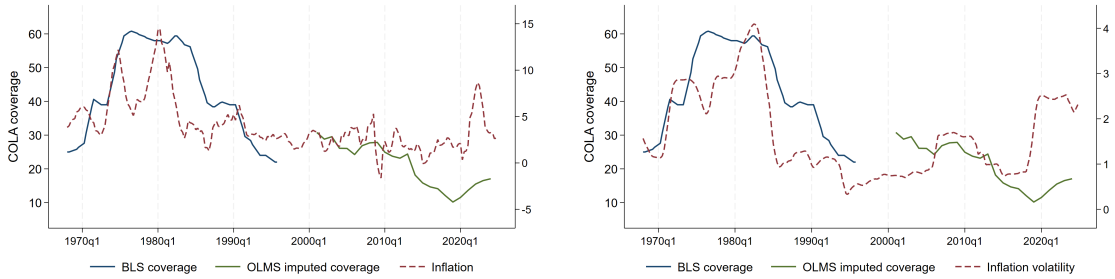
Spain. Figure 19 plots the fraction of new collective bargaining contracts with COLA clauses each year (the pink bar) against aggregate inflation (the red curve) in Spain from 2001 to 2023. This figure comes directly from Spain's central bank report (Economic Bulletin, 3/2022, Banco De

[34]OLMS provided me with metadata of the contract files which contains the expiration date. When it fails to extract the validity period of a contract, I impute the starting date by subtracting τ years from the expiration date. In practice, I tried $\tau = 2, 3$, which are the most common duration periods. The results are almost identical using the two values, and I reported the case for $\tau = 2$.

[35]The metadata from OLMS contains the size of the union that is covered by a CBA, but this variable is largely missing after 2010. I therefore calculated the ratio of the average numbers of workers covered for CBAs with and without COLA. I did this separately using the main 1974-1995 dataset and the OLMS dataset prior to 2010, and the ratio is surprisingly similar. I then used the ratio to impute COLA coverage for the post-2010 series.

[36]There is also a small surge in inflation uncertainty during the Great Recession, but the COLA coverage did not respond significantly. This is not inconsistent with the model: firm financial conditions deteriorated during the financial crisis, making it more difficult to sustain COLA.

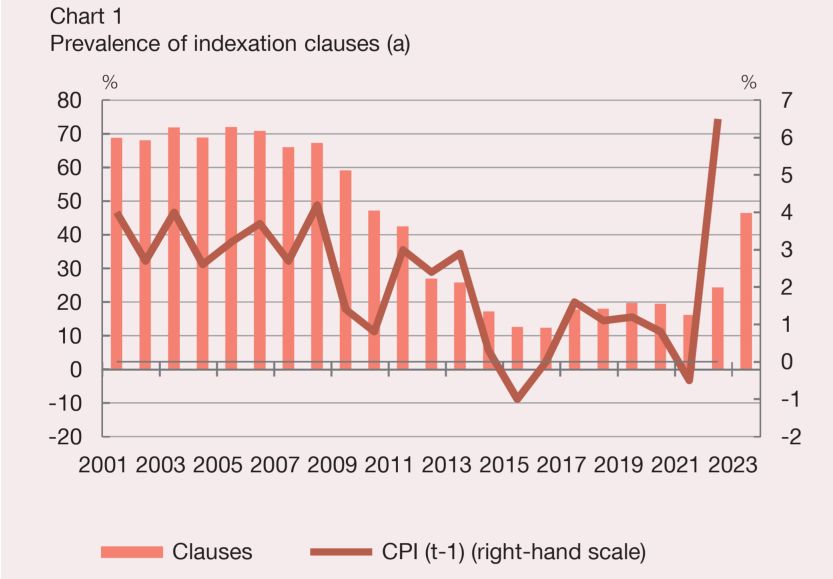
Figure 18: Extended COLA coverage in the US



Notes: The figure plots COLA coverage in the US against inflation and inflation volatility. BLS coverage means the series is from the main 1974-1995 dataset, and the OLMS imputed coverage is from the OLMS collected collective bargaining agreements from 2001 to 2022.

Espana), and it shows a very strong comovement between the COLA inflow and the aggregate inflation. In particular, the COLA inflow rate dropped to about 12 percent in 2015, but it bounced back to about 50 percent in 2023.

Figure 19: COLA adoption vs. inflation in Spain

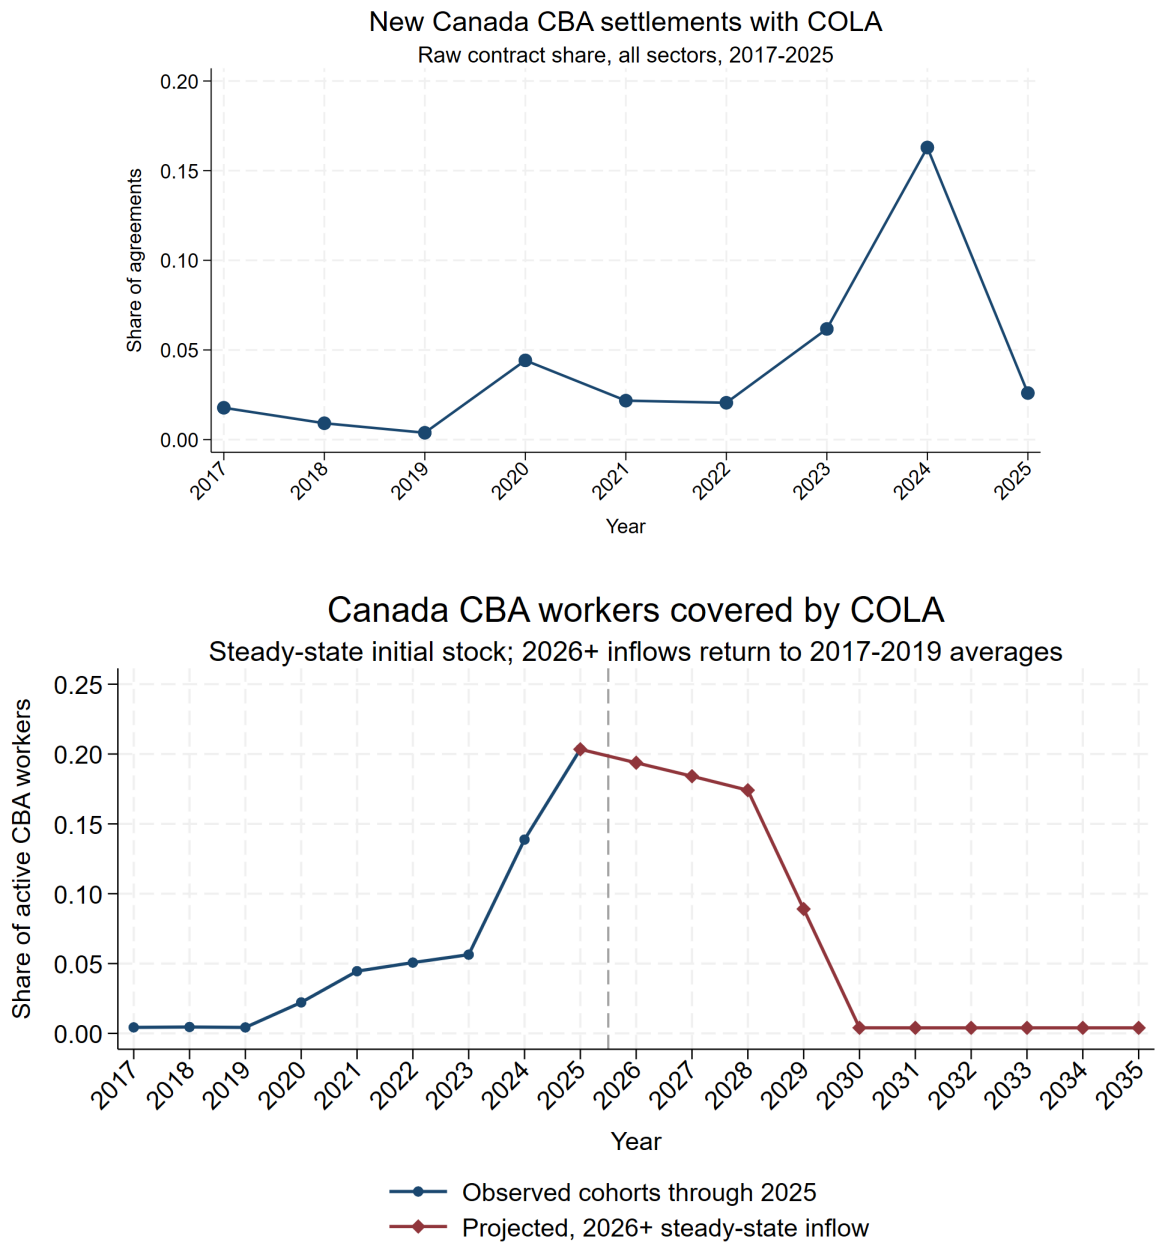


Canada. Canada provides another good example. The upper panel in Figure 20 plots the time path of the share of new collective bargaining contracts with COLA clauses in Canada from 2017 to 2025, while the lower panel plots the imputed worker coverage from the inflow sequence. The red part in the lower panel is the projection of COLA coverage from 2026 to 2035, assuming that COLA inflow goes back to the 2017-2019 average from 2026 on. The aggregate time series comes

directly from the Government of Canada.[37] The data also contains the average duration of the newly signed contracts, which allows me to impute the worker coverage series. Similar to the US and Spain, Canada also witnessed a surge of COLA contracts in 2023 and 2024, though the momentum receded in 2025, as inflation is relatively under control. However, if we look at the worker coverage (the stock), this COLA spike will persist into 2029, because the duration of those COLA contracts is quite long (about 3-4 years).

[37] See https://www.canada.ca/en/employment-social-development/services/collective-bargaining-data/wages/wages-year-cola.html.

Figure 20: COLA adoption in Canada



Overall, these empirical facts complement those in the main text, and they suggest that indexation is state-dependent. This state dependence is not specific to a particular inflation episode nor a particular country.

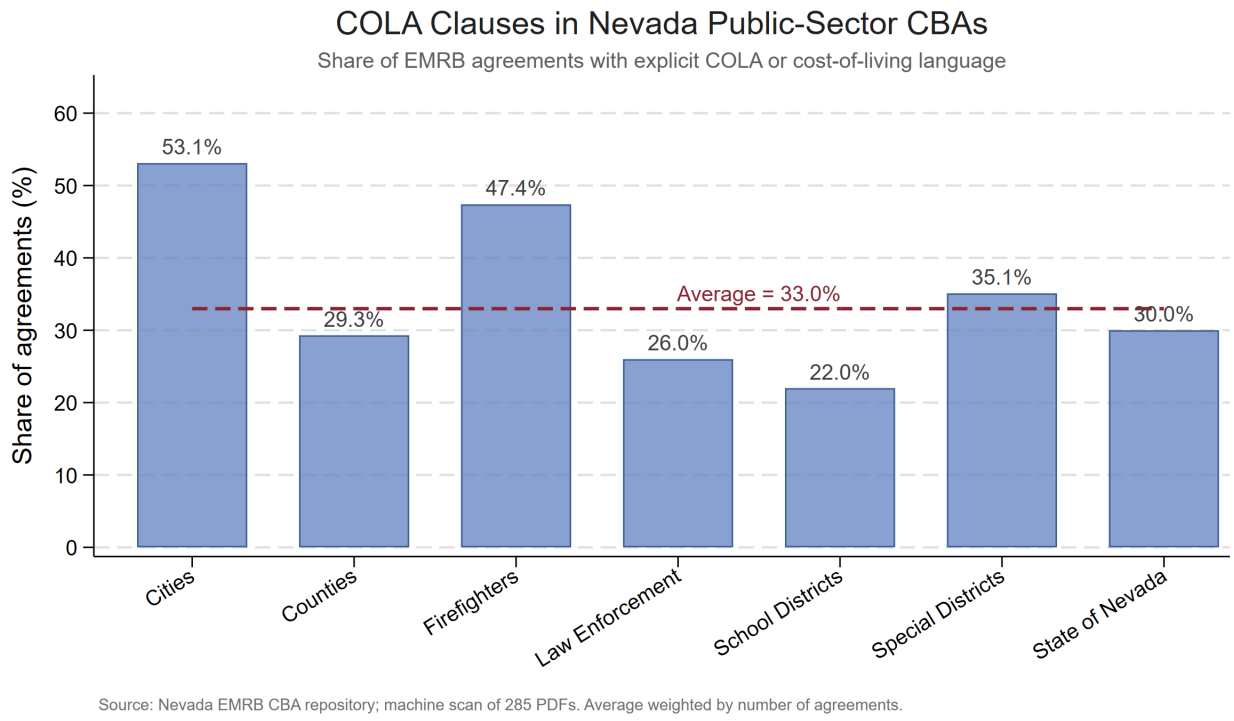
A.7 A Comprehensive Example of Recent COLA Adoption in the US

As another check on recent US COLA use, I collect collective bargaining agreements from the Nevada Government Employee-Management Relations Board. The EMRB states that employee

organizations are required to file a copy of their CBAs annually, and the agreements are publicly available.[38] Although this source is not nationally representative, it is systematic and more up to date.

Figure 21 summarizes a machine scan of 285 EMRB agreements. I classify an agreement as containing COLA if the extracted text contains explicit COLA or cost-of-living language. The average COLA share is 33 percent, with frequent use across cities, firefighters, special districts, state units, counties, law enforcement, and school districts. This suggests that formal COLA clauses are still an important component in wage setting, at least in the US public sector.

Figure 21: COLA clauses in Nevada public-sector CBAs



A.8 The Rise of Off-Cycle/Mid-Year Wage Adjustments in the Recent Inflation Episode

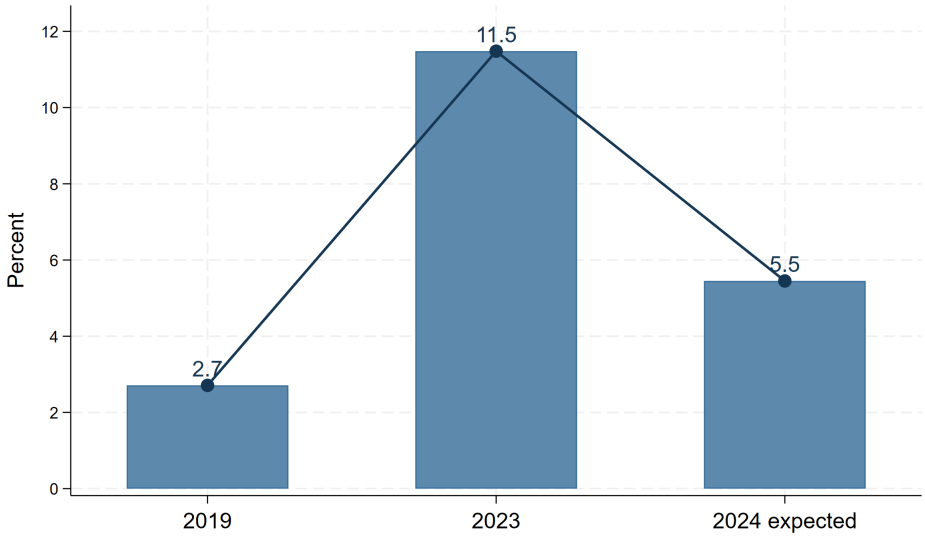
All the evidence that I showed before is about explicit indexation in union contracts. However, inflation-dependent wage setting is not limited to the union sector through formal COLA clauses. Even outside of the union sector, wages can be more responsive to inflation surprises through more frequent renegotiation and more mid-year adjustments. In that sense, I view them as implicit indexation.

[38] See https://emrb.nv.gov/Resources/Collective_Bargaining_Agreements/.

One piece of evidence of implicit indexation comes from the Bank of England's Decision Maker Panel (DMP). The DMP is a monthly survey of UK firms with at least ten employees, and the published data are properly weighted and thus representative of the economy. In a 2024 survey wave, firms were asked how often they changed, or expected to change, the wages of a typical employee in 2019, 2023, and 2024. I use the quarterly aggregate data to calculate the share of firms who changed or expected to change wages more than once a year.

Figure 22 reports the result. In 2019, only 2.7 percent of firms reported changing a typical employee's wage more than once per year. In 2023, this share rose to 11.5 percent. For 2024, when inflation pressure had started to normalize, the expected share declined to 5.5 percent. The pattern is consistent with state-dependent wage setting. One caveat is that, unlike COLA formulas that are specially designed for inflation compensation, one cannot tell whether the increased mid-year wage raises are for inflation compensation or for other purposes like worker retention.

Figure 22: Frequency of wage changes in the Decision Maker Panel



Survey evidence from compensation consultants and qualitative evidence from the Federal Reserve Beige Book point in the same direction. In WTW's 2022 North America Mid-Year Compensation Survey, which covers 884 organizations employing more than 15 million workers, 7 percent of employers had already made salary adjustments more frequently throughout the year and 31 percent were planning or considering doing so; among employers delivering or planning more frequent increases, many reported doing so as needed or more than once per year. Similarly, Mercer's 2023 US Compensation Planning Survey, which includes more than 1,200 US organizations, reports that ad-hoc off-cycle wage reviews or adjustments were a common employer response to inflation, with 38 percent of employers reporting this margin in the public summary. Taken together, they all point to the same pattern that inflation-dependent wage setting also appears in the broader labor market.

A.9 COLA Heterogeneity

The benchmark model abstracts from persistent ex-ante differences across firms. In the data, however, the adoption of COLA contracts varies with industry and other firm heterogeneities. The benchmark model also abstracts from heterogeneities across different COLA contracts. For example, COLA contracts may differ in contract duration, formula and pass-through, and firm financial conditions. This subsection documents these patterns.

Table 8 shows some summary statistics at the sector level. Manufacturing and transportation/utilities account for large worker shares and have high worker-weighted COLA inflow shares (share of new COLA among all the new agreements). In contrast, construction and services have relatively lower adoption rates. The table also shows that COLA contracts tend to be long, which is consistent with the argument that COLA guarantees "labor peace" and avoids the cost of frequent renegotiation.

Table 8: COLA moments by sector

Sector	N	Worker share	Inflow share	Wkr-wtd inflow	PT	COLA dur.	Non-COLA dur.
Mining	93	1.3	61.3	30.6	0.86	35.5	35.4
Construction	5,099	16.1	4.3	5.6	0.83	34.3	26.1
Manufacturing	6,928	32.1	45.8	65.4	0.85	35.5	29.9
Transp./Utilities	2,266	19.7	31.0	56.7	1.09	33.8	27.2
Wholesale	197	0.5	28.9	47.6	0.81	35.7	31.9
Retail	1,388	7.0	17.0	21.9	0.85	34.5	33.5
Finance	232	1.3	24.6	40.5	0.48	33.2	30.7
Services	5,289	22.1	5.0	5.0	0.92	33.8	23.9

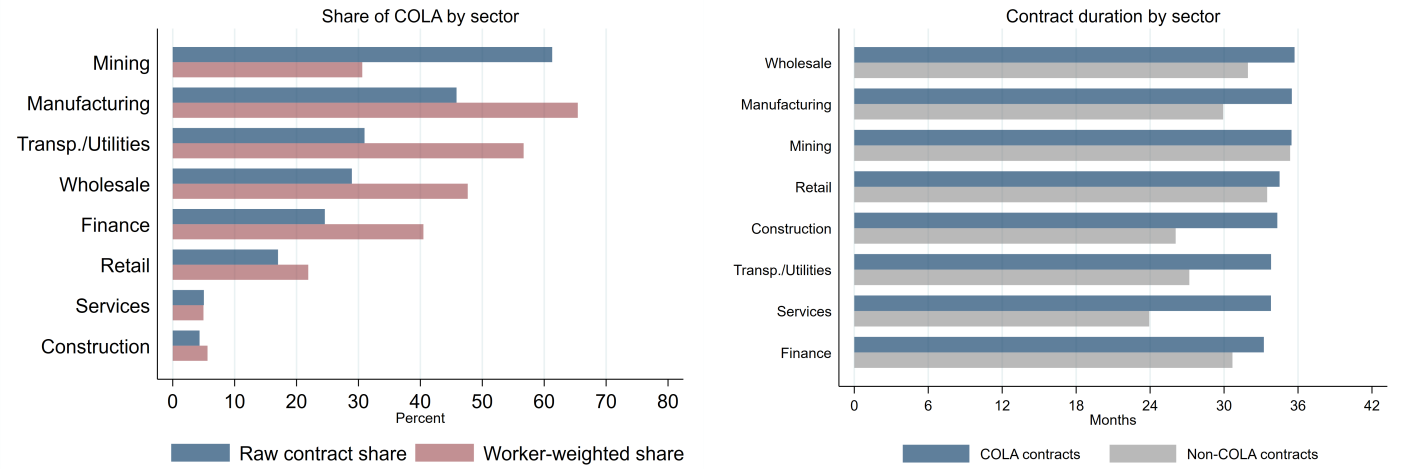
Notes: The inflow share is the fraction of newly negotiated contracts in the sector that include COLA. The worker-weighted inflow share weights contracts by covered workers. The construction of COLA pass-through is described in Appendix A.2. Durations are measured in months.

Figure 23 illustrates the adoption and duration patterns across different sectors. The left panel shows both the raw share of new COLA contracts among all CBAs in each sector and the worker-weighted share. While manufacturing has a high COLA adoption rate, COLA extends well beyond manufacturing. The right panel shows the duration by sector. Sectoral variation is small, and in almost all sectors, COLA contracts last longer than non-COLA contracts.

Figure 24 shows the average COLA pass-through across sectors. In general, the sectoral variation in their COLA pass-through is limited. Over pass-through (higher than one) in transportation and utilities is driven by some extreme values. In most of the contracts, COLA pass-through is above 50 percent but incomplete.

The frequency for wage adjustment to inflation within COLA contracts also shows some variation. Figure 25 reports the distribution of the COLA review frequency. Almost half of the contracts adjust wages to inflation at a quarterly frequency, which makes wages very responsive to inflation surprises. About one third of contracts review each year. If one zooms in and looks at the time

Figure 23: COLA adoption and contract duration by sector



Notes: The left panel reports the sector-level fraction of newly negotiated contracts with COLA clauses, both unweighted and weighted by covered workers. The right panel reports average contract duration by sector and whether COLA is provided.

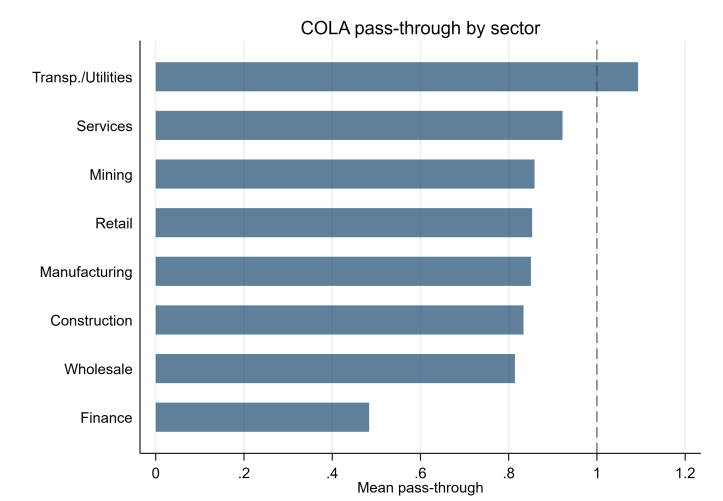
trend (not reported), COLA contracts are more likely to have annual reviews in the late 1980s and early 1990s, during which inflation has been much more stabilized.

Finally, Figure 26 shows the relationship between COLA adoption and simple proxies for firm financial conditions in the Compustat-linked subsample. Firms with more cash holdings are more likely to adopt COLA, while the adoption rate is lower among firms with higher long-term leverage. Of course, these figures only serve as suggestive evidence. Nevertheless, they are consistent with the view that COLA is not only related to inflation risk on the worker side, but also related to financing risk on the firm side.

B Additional Model Testing

The main text reports the three core empirical tests of the model: uncertainty raises COLA adoption, inherited COLA coverage predicts new adoption, and inflation responds more strongly to monetary-policy shocks when the COLA share is high. This appendix reports robustness checks and complementary specifications that help interpret the main evidence. The first subsection shows that the uncertainty response is similar when COLA inflow is weighted by covered workers. The second uses contract-level variation to relate adoption directly to subjective inflation risk. The third reports the aggregate time-series evidence for the strategic complementarity in contract choice, complementing the contract-level evidence in the main text.

Figure 24: COLA pass-through by sector



Notes: The figure reports sector-level averages of COLA pass-through among contracts.

B.1 Worker-Weighted COLA Adoption after Uncertainty Shocks

The baseline local projection in the main text uses the unweighted share of newly signed contracts with COLA. Figure 27 repeats the exercise using the worker-weighted inflow rate. This weighting gives more influence to large contracts. The response remains positive.

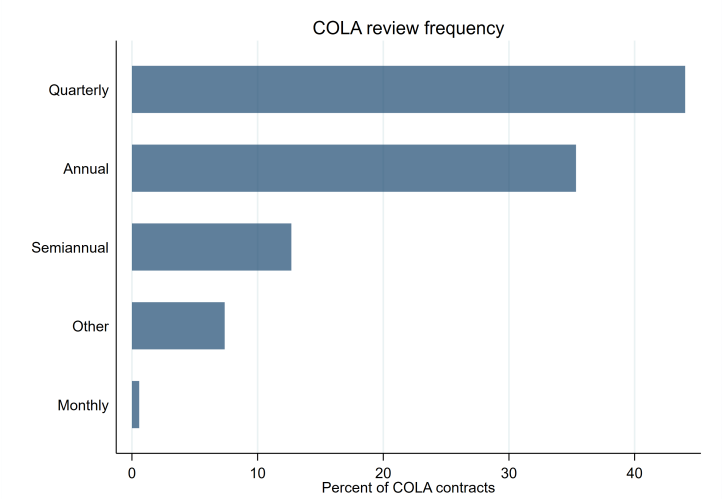
B.2 Subjective Inflation Risk and COLA Adoption

The local projection exercise using the aggregate time series of COLA adoption in the main text establishes that COLA adoption rises after macro uncertainty shocks. Table 9 provides complementary evidence at the firm-union pair level. The dependent variable is an indicator for whether a newly negotiated agreement adopts COLA. The key regressor is the subjective inflation uncertainty constructed by one-year-ahead inflation forecasts from the Michigan Household Survey. I include pair fixed effects to absorb persistent differences across different bargaining relationships. Moreover, I control for other macro variables, including inflation, to absorb contemporaneous aggregate conditions. I instrument the subjective inflation uncertainty with lagged uncertainty shocks constructed from the JLN uncertainty index in the main text.

The estimates show that given the same inflation realization and within a bargaining pair, contracts are more likely to have COLA when future inflation is more uncertain. This is consistent with the prediction from the contract choice model in partial equilibrium.

I also show explicitly the coefficient on inflation. The positive and significant coefficient is also informative (I instrument inflation with lagged Romer-Romer policy shocks). In the benchmark model, the adoption gain depends on conditional second moments but not the level of inflation. Empirically, even conditional on the same amount of inflation uncertainty, higher inflation realization makes COLA adoption more likely. This corroborates the model extension of endogenously

Figure 25: COLA review frequency



Notes: The figure reports the distribution of COLA review frequencies among contracts with COLA clauses.

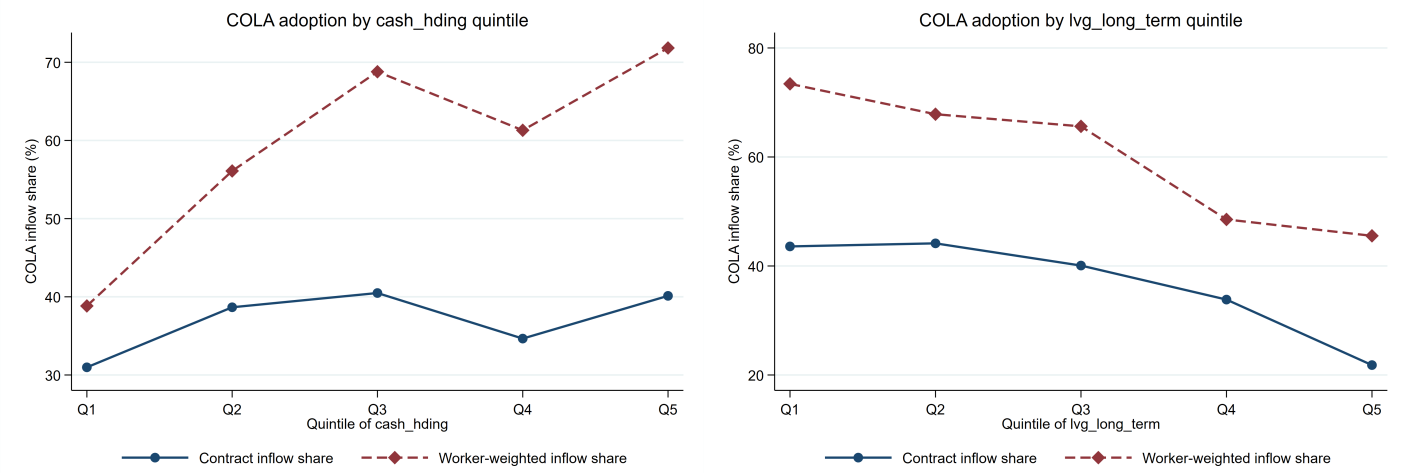
stochastic inflation volatility that I discuss at the end of the main text. More details can be found in Section I.2. The idea is, when high inflation itself changes the expected gain from inflation contingency, first-order inflation realizations can affect COLA adoption, through which it can then affect inflation volatility.

Table 9: Inflation risk and contract-level COLA adoption

	(1) IQR, 2SLS	(2) Var, 2SLS
Inflation	0.024***	0.028***
	(0.003)	(0.004)
Subjective inflation uncertainty, IQR	0.084**	
	(0.036)	
Subjective inflation uncertainty, variance		0.009**
		(0.004)
Pair FE	✓	✓
Macro controls	✓	✓
Observations	14248	14248

Notes: The dependent variable is an indicator for COLA adoption. The subjective inflation uncertainty is constructed by one-year-ahead inflation forecasts from the Michigan Household Survey. t corresponds to the quarter of renegotiation. The IQR specification uses the weighted monthly interquartile range of individual forecasts after normalizing each forecast by the monthly mean forecast, while the variance specification uses the corresponding weighted cross-sectional variance. Inflation and subjective uncertainty are instrumented with lags 1--4 of Romer-Romer monetary policy shocks and lags 1--4 of residualized JLN uncertainty shocks. Macro controls include lags 1--4 of the federal funds rate, lagged detrended log employment, and lags of detrended log real GDP, in addition to union size and contract duration. All specifications include bargaining-pair fixed effects. Standard errors are clustered at the bargaining-pair level. *, **, and *** denote significance at the 10, 5, and 1 percent levels.

Figure 26: COLA adoption by balance-sheet quintile



Notes: The figure uses the Compustat-linked subsample. The left panel groups contracts by cash holdings over assets. The right panel groups contracts by long-term debt over assets. Each panel reports contract-count and worker-weighted COLA inflow shares.

B.3 Aggregate Strategic Complementarity in COLA Adoption

The main text uses variation at the firm-union pair level to test whether a higher COLA stock raises the probability that a newly negotiated contract adopts COLA. Table 10 reports the counterpart using the aggregate time series of COLA adoption.

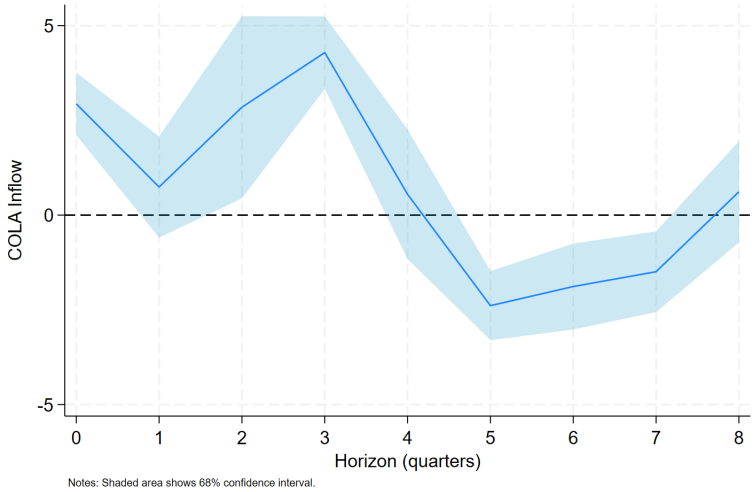
I instrument the current COLA stock with long-lagged Baumeister-Hamilton oil supply shocks. The identifying assumption is that these distant shocks affect new adoption today only through their effects on previously negotiated long-term agreements, after controlling for current and lagged macro conditions and for recent oil shocks. Similar to the OLS estimates, the IV estimates are positive in both unweighted and worker-weighted specifications.

Table 10: Aggregate strategic complementarity in COLA adoption

	OLS		2SLS	
	Unweighted	Wkr-wtd	Unweighted	Wkr-wtd
	(1)	(2)	(3)	(4)
COLA stock	0.892	1.237	1.133	2.787
	(0.188)	(0.232)	(0.504)	(0.831)
N	88	88	76	76
KP rk Wald F			15.25	15.25

Notes: The table reports regressions of the COLA inflow rate on the current COLA stock. Newey-West HAC standard errors are in parentheses. Macro controls include lagged inflation, detrended real output, the federal funds rate, and the JLN one-year-ahead macro uncertainty index. In the 2SLS specifications, the COLA stock is instrumented with the Baumeister-Hamilton oil-supply shock eight quarters earlier, and the controls include the recent lags of the same shock.

Figure 27: Worker-weighted COLA inflow after an uncertainty shock



Notes: The figure reports the local-projection response of the worker-weighted COLA inflow rate to an uncertainty shock. The shock is constructed as in the main text.

C Understanding the Contracting Problem: A One-Firm-One-Worker Case

For wages to have an allocative effect in the model presented in the main text, I turn on the employment choice margin. This is the key channel for COLA contracts to amplify inflationary shocks in the general equilibrium. However, this margin is not necessary for understanding the contracting problem per se, especially for the contract type choice and the optimal wage rules conditional on type. Therefore, to better understand the contracting problem, I study a one-firm-one-worker problem in great detail here. This problem is actually nested in the main model in the sense that if one assumes infinite labor adjustment costs or extreme decreasing returns to scale in the service firm's production function, one effectively turns off the employment choice margin.

I first set up a one-period contracting model, and then I characterize the first-best contract as well as the optimal inflation-contingent contract where firm profits are not contractible. The model can be easily extended to multiple periods. Finally, I study the non-contingent contracts and the gain from COLA relative to non-contingent wages.

C.1 Environment

Preferences and Shocks Consider a one-period model where a firm offers a wage contract to a worker. The worker has preferences $u(.)$ such that $u'' < 0$ and is hand-to-mouth, while the firm is risk neutral. They produce $z \geq 0$ units of product. The product is sold at nominal price P, which is also the aggregate price level. Let P_0 be the initial price level, which is normalized to one. Inflation

π is defined as $e^{\pi} = P/P_0 = P$. Assume π has CDF $F^{\pi}(\pi)$, while for each π, z has conditional CDF $F^{z}(z|\pi)$.[39] Both $F^{\pi}(\pi)$ and $F^{z}(z|\pi)$ are public information.

Firm financial constraint If the contractual wage W, which will be specified shortly, is higher than the firm's sales revenue Pz, the firm can issue new equity to pay the workers, which incurs an external financing cost ϕ.[40] That is, for each unit of equity raised, the firm has to pay an additional cost of ϕ. When $Pz > W$, the firm can pay the wage with internal funds only, and the nominal profit is $Pz - W$. When $Pz < W$, however, the firm must rely on external financing, and the nominal profit becomes $(1+\phi)(Pz - W)$, which is negative.

The contracting problem The contract is signed before z and π are realized, and the wage is paid in nominal terms. Both the firm and the worker commit to the contract. In offering the contract, the firm maximizes total value $u(c) + \lambda d$, where c is the worker's real consumption, d is the firm's real cash flow, and λ is the firm's relative weight in their total value.[41] Given any contract space $\mathcal{C}$, the problem is

$$\max_{W \in \mathcal{C}} \mathbb{E}_0 \left[u\left(We^{-\pi}\right) + \lambda \left[z - We^{-\pi} - \phi \max\left\{We^{-\pi} - z, 0\right\}\right]\right]$$

C.2 The Full-contingency Contract

Assume there are no restrictions on the contract space $\mathcal{C}$. Then, the contract can be contingent on all possible states (z, π) ex post. Hence, choosing this first-best contract reduces to choosing a wage function $W(z, \pi) : \mathbb{R}_+^2 \to \mathbb{R}_+$ from the state-by-state optimization problem:

$$\max_{W \in \mathbb{R}_+} u\left(We^{-\pi}\right) + \lambda \left[z - We^{-\pi} - \phi \max\left\{We^{-\pi} - z, 0\right\}\right]$$

Denote by $W^{FB}(z, \pi)$ the nominal wage from the first-best contract, and by $w^{FB}(z, \pi)$ the real wage implied by it. The next proposition characterizes the first-best contract.

Proposition 6. *(First-best contract) $w^{FB}(z, \pi)$ is independent of π and is given by*

$$w^{FB}(z) = \begin{cases} w_1 & z < w_1 \\ z & z \in [w_1, w_2] \\ w_2 & z > w_2 \end{cases}$$

[39] One may not interpret z as a pure firm-level idiosyncratic productivity shock. Rather, it may also contain an aggregate TFP shock, industry output price dispersion, and input cost dispersion, all of which can be correlated with the inflation shock. For example, if it is a supply-side shock that dominates the economy, we may expect a negative correlation between z and π; if, on the other hand, it is a demand-side shock that drives the economy, z and π may be positively correlated. As I will show later, this correlation turns out to be important for the contract choice problem.

[40] I assume a quadratic external financing cost function in the main model. Here, I assume linear costs for simplicity. This is not essential for the qualitative results.

[41] One can think of this as an efficient bargaining outcome.

where w_1, w_2 are defined by

$$u'(w_1) = \lambda + \lambda\phi$$
$$u'(w_2) = \lambda$$

and $W^{FB}(z,\pi) = w^{FB}(z)e^{\pi}$.

Proof. Let $w = We^{-\pi}$ be the real wage. First note that only the real wage w enters the objective function. The formal discussion is quite complicated because of the discrete choice introduced by the financial constraint. I use the following backward induction: First, for any given z, I maximize the objective function conditional on whether the firm uses external financing or not. Then, I compare the resulting value of using external financing and not using it. I use a graph to illustrate the problem in the end.

Case 1: The firm uses external financing. This means the firm will set $w \geq z$. In this case, the objective function becomes $u(w) + \lambda(1+\phi)(z-w)$, and we choose $w > z$ to maximize it. Notice this is a smooth and concave function, and its FOC gives the unique maximizer if there is no constraint. Denote that by w_1, i.e.,

$$u'(w_1) - \lambda(1+\phi) = 0$$

Then, if $z < w_1$, w_1 is indeed the maximizer. On the other hand, if $z \geq w_1$, they cannot choose w_1 on the interval $[z, \infty)$. But we know the objective function is decreasing on $[z, \infty)$. Hence, the firm will choose the boundary solution $w = z$ in this case with value $u(z)$.

Case 2: The firm does not use external financing. This means the firm will set $w \leq z$. The objective function is $u(w) + \lambda(z-w)$, and we choose $w \leq z$ to maximize it. Again this is a smooth and concave function, and its FOC yields the unique maximizer if there is no constraint. Denote that by w_2, namely

$$u'(w_2) - \lambda = 0$$

Then, if $z > w_2$, w_2 is the maximizer. But if $z \leq w_2$, w_2 cannot be chosen on the interval $(0, z]$. However, the objective function is increasing on $(0, z]$. Therefore, the firm chooses the boundary solution $w = z$, and the value is $u(z)$.

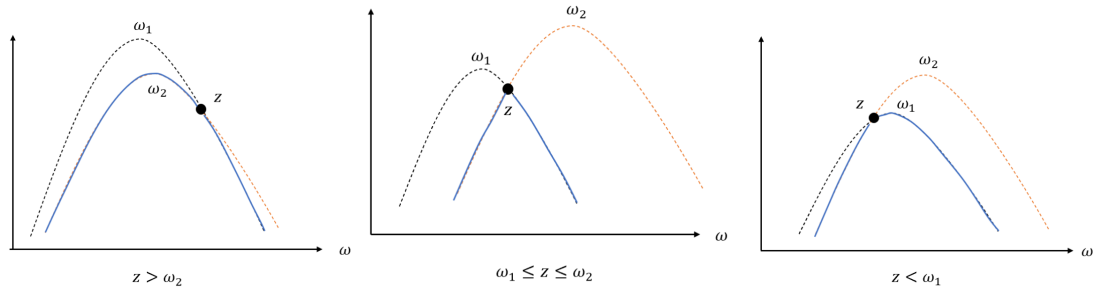
Compare the values from the two cases: We compare the weighted total values from having or not having external financing under different values of z. Note by concavity of u, $w_1 \leq w_2$. There are three scenarios based on the relative position of z, w_1, and w_2.

Scenario 1 $z > w_2$: If the firm uses external financing, because $z > w_2 \geq w_1$, the firm chooses $w = z$ with value $u(z)$. On the other hand, if the firm does not use external financing, since $z > w_2$, the firm sets $w = w_2$. Also notice that the resulting value is higher than choosing z. Hence, $w = w_2$ is indeed the maximizer in this scenario.

Scenario 2 $z \in [w_1, w_2]$**:** If the firm uses external financing, because $z \geq w_1$, again, the firm chooses $w = z$ with value $u(z)$. If the firm does not use external financing, since $z \leq w_2$, the firm still sets $w = z$, yielding a value of $u(z)$. Therefore, $w = z$ is the maximizer in this scenario.

Scenario 3 $z < w_1$**:** If the firm uses external financing, because $z < w_1$, the firm sets $w = w_1$. Notice that the value function is increasing in $(z, w_1]$, and it approaches $u(z)$ when w goes to z. On the other hand, if the firm does not use external financing, since $z < w_2$, the firm sets $w = z$, yielding a value of $u(z)$. Hence, the value from using external financing is higher, and $w = w_1$ is the maximizer in this scenario.

The following figures illustrate the value functions for different z's. The black dashed line corresponds to using external financing, with value function $u(w) + \lambda(1+\phi)(z-w)$. The orange dashed line corresponds to not using it, with value function $u(w) + \lambda(z-w)$. Note the two curves intersect at $w = z$. The blue solid line is the actual value function in each case: when $w < z$, the firm does not use external financing, but it does when $w > z$.



Conclusion. Combining all pieces together, the complete solution is

$$w^{FB}(z) = \begin{cases} w_1 & z < w_1 \\ z & z \in [w_1, w_2] \\ w_2 & z > w_2 \end{cases}$$

Note also that if there is no financial constraint, i.e., $\phi = 0$, $w_1 = w_2$, and the firm always pays a constant real wage no matter what the states are. □

This proposition shows that the optimal real wage depends only on the realized z because it completely determines the size of the pie to be shared. The firm dislikes paying high wages in a low revenue state because of its financial frictions. In addition, holding the real revenue fixed, the nominal wage moves one-for-one with the price level.

C.3 The inflation-contingent (COLA) Contract

Restricted contract space Assume the realization of z is the firm's private information, and there is no credible way to reveal it. In contrast, the inflation realization π is observed by both. Hence,

the contract can only depend on inflation π but not real profit z.[42] The contract space contains all nominal wage functions $W(\pi) : \mathbb{R}_+ \to \mathbb{R}_+$. Each $W(\pi)$ also implies a unique real wage function $w(\pi) : \mathbb{R} \to \mathbb{R}_+$ where $w(\pi) = W(\pi)e^{-\pi}$.

The contracting problem under limited contractibility Given that π is contractible, they can decide the allocation as if they observe π. Therefore, the problem is equivalent to choosing the desired allocation conditional on π:

$$\max_{W(\pi):\mathbb{R}\to\mathbb{R}_+} \mathbb{E}^z \left[u\left(We^{-\pi}\right) + \lambda \left[z - We^{-\pi} - \phi \max\left\{ We^{-\pi} - z, 0 \right\} \right] \middle| \pi \right]$$

Denote $W^{\mathcal{C}}(\pi)$ as the nominal wage from the optimal inflation-contingent contract, and $w^{\mathcal{C}}(\pi)$ the real wage implied by it. The next proposition characterizes the optimal inflation-contingent contract.

Proposition 7. *$w^{\mathcal{C}}(\pi)$ is uniquely defined by*

$$u'(w^{\mathcal{C}}(\pi)) - \lambda - \lambda\phi F^z(w^{\mathcal{C}}(\pi)|\pi) = 0$$

and the corresponding nominal wage is $W^{\mathcal{C}}(\pi) = w^{\mathcal{C}}(\pi)e^{\pi}$. Moreover, $w^{\mathcal{C}}(\pi)$ is decreasing (increasing) in π iff $F^z_\pi(w^{\mathcal{C}}(\pi)|\pi) > 0$ (< 0). If z and π are independent, i.e., $F^z_\pi = 0$, the contract provides a constant real wage.

Proof. Let $w = We^{-\pi}$. We first integrate the ex post value over z conditional on π, and the resulting conditional expectation is denoted by $v^{int}(w, \pi)$.

$$\begin{aligned}
v^{int}(w,\pi) &= \mathbb{E}^z \left[u(w) + \lambda \left[z - w - \phi \max\{w - z, 0\} \right] \middle| \pi \right] \\
&= u(w) + \lambda \left(\mathbb{E}\left[z|\pi \right] - w \right) - \lambda\phi \mathbb{E}\left[\max\{w - z, 0\} \,|\pi \right] \\
&= u(w) + \lambda \left(\mathbb{E}\left[z|\pi \right] - w \right) - \lambda\phi \int_0^w (w - z) dF^z(z|\pi) \\
&= u(w) + \lambda \left(\mathbb{E}\left[z|\pi \right] - w \right) - \lambda\phi \int_0^w F^z(z|\pi) dz
\end{aligned}$$

where the second-to-last line uses integration by parts. Notice

$$\frac{\partial^2 v^{int}(w,\pi)}{\partial w^2} = u''(w) - \lambda\phi f^z(w|\pi) < 0$$

Hence, $v^{int}(w, \pi)$ is a strictly concave function of w. This means the FOC defines the unique maximizer $w^{\mathcal{C}}(\pi)$ which reads

$$u'(w^{\mathcal{C}}(\pi)) - \lambda - \lambda\phi F^z(w^{\mathcal{C}}(\pi)|\pi) = 0$$

[42] The firm always wants to report a low z to reduce wage payment. In contrast, inflation is published by the government and is relatively reliable.

Differentiating with respect to π, we have

$$\frac{dw^{\mathcal{C}}(\pi)}{d\pi} = \frac{\lambda\phi F^z_\pi(w^{\mathcal{C}}(\pi)|\pi)}{u''(w^{\mathcal{C}}(\pi)) - \lambda\phi f^z(w^{\mathcal{C}}(\pi)|\pi)}$$

Since the denominator is negative, the sign of $\frac{dw^{\mathcal{C}}(\pi)}{d\pi}$ must be the opposite of the sign of $F^z_\pi(w^{\mathcal{C}}(\pi)|\pi)$. □

Therefore, in the optimal inflation-contingent contract, the nominal wage is indexed to inflation, as π is a coarser instrument and can act as a signal of z. If π and z are independent, the firm absorbs all the risks and offers a constant real wage to the worker. If higher inflation signals lower firm revenues ($\partial F^z(z|\pi)/\partial\pi > 0$), the worker will absorb some risk, and the real wage will be decreasing with inflation. As a result, the pass-through from price to nominal wage is optimally smaller than one.

C.4 Extending the Contracting Model to Two Periods

We next consider a two-period contracting model. Assume they discount future values by the same factor β. The shock in period t is $(z_t, \pi_t), t = 1, 2$, with an arbitrary correlation structure. We maintain the limited contractibility assumption, namely, they cannot contract on z but can contract on the history of π up to t, denoted as π^t. Therefore, the problem can be written as

$$\max_{W_t(\pi^t)} \mathbb{E}_0 \left[\sum_{t=1}^{2} \beta^{t-1} \left(u\left(W_t e^{-\pi_t}\right) + \lambda \left[z_t - W_t e^{-\pi_t} - \phi \max\left\{ W_t e^{-\pi_t} - z_t, 0 \right\} \right] \right) \right]$$

Intuitively, because the wage function chosen in period 1 does not affect any states in period 2, there are no intertemporal considerations. A deeper reason for this result to hold is the two-sided commitment assumption. Because of that, there is no need for the firm to use wages strategically to retain the worker. Outcomes in period 1 affect choice in period 2 only up to the point that π_1 may contain information about z_2. To be more precise, we can split the objective function by period:

$$\begin{aligned}
&\mathbb{E}_0 \left[\sum_{t=1}^{2} \beta^{t-1} \left(u\left(W_t e^{-\sum_{s=1}^{t}\pi_s}\right) + \lambda \left[z_t - W_t e^{-\sum_{s=1}^{t}\pi_s} - \phi \max\left\{ W_t e^{-\sum_{s=1}^{t}\pi_s} - z_t, 0 \right\} \right] \right) \right] \\
=&\mathbb{E}_0 \left[u\left(W_1 e^{-\pi_1}\right) + \lambda \left[z_1 - W_1 e^{-\pi_1} - \phi \max\left\{ W_1 e^{-\pi_1} - z_1, 0 \right\} \right] \right] \\
&+ \beta \mathbb{E}_0 \left[u\left(W_2 e^{-\pi_2-\pi_1}\right) + \lambda \left[z_2 - W_2 e^{-\pi_2-\pi_1} - \phi \max\left\{ W_2 e^{-\pi_2-\pi_1} - z_2, 0 \right\} \right] \right] \\
=&\mathbb{E} \left[\mathbb{E} \left[u\left(W_1 e^{-\pi_1}\right) + \lambda \left[z_1 - W_1 e^{-\pi_1} - \phi \max\left\{ W_1 e^{-\pi_1} - z_1, 0 \right\} \right] \middle| \pi_1 \right] \right] \\
&+ \beta \mathbb{E} \left[\mathbb{E} \left[u\left(W_2 e^{-\pi_2-\pi_1}\right) + \lambda \left[z_2 - W_2 e^{-\pi_2-\pi_1} - \phi \max\left\{ W_2 e^{-\pi_2-\pi_1} - z_2, 0 \right\} \right] \middle| \pi_1, \pi_2 \right] \right]
\end{aligned}$$

Therefore, the second-best contract allows them to choose $W_1(\pi_1)$ that maximizes the interim value in period 1 conditional on π_1, and $W_2(\pi_1, \pi_2)$ that maximizes the interim value in period 2 conditional on (π_1, π_2). Denote the corresponding second-best real wage by $w_t^{\mathcal{C}}(\pi^t)$, which is characterized by the next proposition.

Proposition 8. *(Second-best contract) $w_t^{\mathcal{C}}(\pi^t)$ is uniquely defined by*

$$u'(w_t^{\mathcal{C}}(\pi^t)) - \lambda - \lambda\phi F^z(w_t^{\mathcal{C}}(\pi^t)|\pi^t) = 0$$

and the corresponding nominal wages are given by

$$W_t^{\mathcal{C}}(\pi^t) = w_t^{\mathcal{C}}(\pi^t)e^{\sum_{s=1}^t \pi_s}$$

This means given π_1, the second-best contract is simply a repetition of two static contracts. It does not matter if we extend the duration to some $T \geq 3$.

C.5 Second-Order Approximation

To make more analytical progress, we expand the problem to second order. We assume firm revenue contains an idiosyncratic and an aggregate part, and the idiosyncratic component is i.i.d. and is independent of everything else.[43] In particular, $z_{it} = \eta_{it}\zeta_t$. Denote by $Q(\zeta, w)$ the expected value of the firm if the aggregate component of its revenue is ζ_t, and the real wage is w_t, i.e.,

$$Q(\zeta_t, w_t) := \mathbb{E}^{\eta}\left[\eta_{it}\zeta_t - w_t - \phi(w_t - \eta_{it}\zeta_t)_+\right]$$

The function Q can be regarded as the effective "utility function" of the firm, which affects risk-sharing with the worker in the contracting problem. It can be easily shown that Q decreases with wage payment w_t, but at any given w_t, the cost of marginally more wage is lower if ζ_t is higher, meaning the firm is less effectively risk averse. However, since ζ_t is not observed by the worker, it is not contractible. Given the observed path of inflation realization and the knowledge of the shock process, they can infer ζ_t from π_t. The next lemma tells us how the predicted revenue risk is shared.

Lemma 4. *Up to 2nd order around the steady state, the optimal COLA real wage is proportional to the forecast of ζ_t, namely*

$$\hat{w}_t^{\mathcal{C}} = \mu\mathbb{E}\left[\hat{\zeta}_t|\mathcal{F}_{t-1}, \pi_t\right]$$

where throughout the paper, $\hat{x} := \log x - \log \bar{x}$ denotes the log deviation of variable x from its steady-state level.

[43] This is more of a technical assumption. Since nothing can predict the idiosyncratic shock, we can first integrate it out. This makes the objective function that matters for wage determination smooth, and therefore, second-order expansion is feasible.

Proof. Denote by $V(\zeta_t, w_t)$ the joint value conditional on ζ_t and w_t. We expand it in log terms around the no-shock steady state up to second order:

$$V(\zeta_t, w_t) \approx \bar{V} + \bar{V}_w \bar{w} \hat{w}_t + \frac{1}{2}\left(\bar{V}_{ww}\bar{w}^2 + \bar{V}_w \bar{w}\right)\hat{w}_t^2 + \bar{V}_{w\zeta}\bar{w}\bar{\zeta}\hat{w}_t\hat{\zeta}_t + t.i.w.$$

where $t.i.w.$ means terms independent of $\hat{w}_t$, $\bar{w}$ and $\bar{\zeta}$ are the steady-state wage and revenue, and $\bar{V}, \bar{V}_w, \bar{V}_{ww}, \bar{V}_{w\zeta}$ are functions/derivatives evaluated at the steady-state level. FOC for $\hat{w}_t$ implies that

$$\bar{V}_w = 0$$

Therefore, the optimal inflation-contingent real wage conditional on inflation realization π_t is

$$\hat{w}_t^{\mathcal{C}} = \frac{\bar{V}_{\zeta w}\bar{\zeta}}{-\bar{V}_{ww}\bar{w}}\mathbb{E}\left[\hat{\zeta}_t | \mathcal{F}_{t-1}, \pi_t\right]$$

Define

$$\mu = \frac{\bar{V}_{\zeta w}\bar{\zeta}}{-\bar{V}_{ww}\bar{w}} = \frac{\bar{Q}_{\zeta w}\bar{\zeta}}{-\left(u''(\bar{w}) + \bar{Q}_{ww}\right)\bar{w}} > 0$$

as the worker's risk exposure. □

In fact, if ζ_t is contractible (but not η_{it}), the optimal real wage should be

$$\hat{w}_t^* = \mu\hat{\zeta}_t$$

That is, the optimal real wage is proportional to the real revenue, and μ measures how much of a unit revenue risk is passed to the worker. When there are no financial frictions so that $\phi = 0$, the firm is truly risk neutral. In that case, we have $\mu = 0$, and the firm would absorb all the risk, while the worker is fully insured. When ζ_t is not contractible, inflation is a coarser instrument when they use inflation-contingent contracts. A fraction μ of the revenue risk predicted by the inflation realization is passed to the worker.

C.6 The Non-contingent contract and the gains from COLA

We now study the optimal non-contingent contracts in a one-period contracting problem. We keep the assumption on the shock process in the last subsection.

Recall $Q(\zeta_t, w_t)$ denotes the firm's expected real cash flow over its idiosyncratic shock. For a one-period non-contingent contract, instead of choosing a nominal wage as a function of ζ and/or π, the contract specifies a single number W for all possible states of the world based on available information at contracting. In particular, the problem can be written as

$$\max_{W\in\mathbb{R}_+} \mathbb{E}\left[u\left(\frac{W}{P_t}\right) + \lambda Q\left(\zeta_t, \frac{W}{P_t}\right)\middle| \mathcal{F}_{t-1}\right]$$

It turns out the optimal non-contingent wage is tightly linked to the optimal COLA wage, which we show next.

Proposition 9. *(1) In nominal terms, the optimal non-contingent wage targets the optimal COLA nominal wage with available information, namely,*

$$\log W_t^{\mathcal{N}} = \mathbb{E}\left[\log W_t^{\mathcal{C}} \middle| \mathcal{F}_{t-1}\right]$$

and in real terms, the optimal non-contingent wage targets the optimal COLA real wage, adjusting for inflation forecast error

$$\hat{w}_t^{\mathcal{N}} = \mathbb{E}\left[\hat{w}_t^{\mathcal{C}} \middle| \mathcal{F}_{t-1}\right] + \mathbb{E}\left[\pi_t \middle| \mathcal{F}_{t-1}\right] - \pi_t$$

(2) The gains from COLA relative to the non-contingent wage are

$$\mathcal{V}_{t-1} = \bar{C}\mathbb{E}_{t-1}\left[\left(\hat{w}_t^{\mathcal{C}} - \mathbb{E}_{t-1}\left[\hat{w}_t^{\mathcal{C}}\right] + \pi_t - \mathbb{E}_{t-1}\left[\pi_t\right]\right)^2 \middle| \pi_t\right]$$

where $\bar{C}$ is a constant that only depends on steady-state values.

Proof. For any realized real wage $\hat{w}_t$ that depends only on π_t given information set $\mathcal{F}_{t-1}$, the expected loss compared to COLA is

$$\begin{aligned}
\mathcal{L}_{t-1}^w =& \mathbb{E}_{t-1}\left[\left(\frac{1}{2}\bar{V}_{ww}\bar{w}^2\left(\hat{w}_t^{\mathcal{C}}\right)^2 + \bar{V}_{w\zeta}\bar{w}\bar{\zeta}\hat{w}_t^{\mathcal{C}}\hat{\zeta}_t\right) - \left(\frac{1}{2}\bar{V}_{ww}\bar{w}^2\hat{w}_t^2 + \bar{V}_{w\zeta}\bar{w}\bar{\zeta}\hat{w}_t\hat{\zeta}_t\right)\right] \\
=& \mathbb{E}_{t-1}\left[\frac{1}{2}\bar{V}_{ww}\bar{w}^2\left(\hat{w}_t^{\mathcal{C}} + \hat{w}_t\right)\left(\hat{w}_t^{\mathcal{C}} - \hat{w}_t\right) + \bar{V}_{w\zeta}\bar{w}\bar{\zeta}\hat{\zeta}_t\left(\hat{w}_t^{\mathcal{C}} - \hat{w}_t\right)\right] \\
=& \mathbb{E}_{t-1}\left[\left[\frac{1}{2}\bar{V}_{ww}\bar{w}^2\left(\hat{w}_t^{\mathcal{C}} + \hat{w}_t\right) + \bar{V}_{w\zeta}\bar{w}\bar{\zeta}\hat{\zeta}_t\right]\left(\hat{w}_t^{\mathcal{C}} - \hat{w}_t\right)\right] \\
=& \mathbb{E}_{t-1}\left[\mathbb{E}_{t-1}\left[\left[\frac{1}{2}\bar{V}_{ww}\bar{w}^2\left(\hat{w}_t^{\mathcal{C}} + \hat{w}_t\right) + \bar{V}_{w\zeta}\bar{w}\bar{\zeta}\hat{\zeta}_t\right]\left(\hat{w}_t^{\mathcal{C}} - \hat{w}_t\right)\middle|\pi_t\right]\right] \\
=& \mathbb{E}_{t-1}\left[\mathbb{E}_{t-1}\left[\left[\frac{1}{2}\bar{V}_{ww}\bar{w}^2\left(\hat{w}_t^{\mathcal{C}} + \hat{w}_t\right) + \bar{V}_{w\zeta}\bar{w}\bar{\zeta}\mathbb{E}_{t-1}\left[\hat{\zeta}_t\middle|\pi_t\right]\right]\left(\hat{w}_t^{\mathcal{C}} - \hat{w}_t\right)\middle|\pi_t\right]\right] \\
=& \mathbb{E}_{t-1}\left[\mathbb{E}_{t-1}\left[\left[\frac{1}{2}\bar{V}_{ww}\bar{w}^2\left(\hat{w}_t^{\mathcal{C}} + \hat{w}_t\right) - \bar{V}_{ww}\bar{w}^2\hat{w}_t^{\mathcal{C}}\right]\left(\hat{w}_t^{\mathcal{C}} - \hat{w}_t\right)\middle|\pi_t\right]\right] \\
=& -\mathbb{E}_{t-1}\left[\mathbb{E}_{t-1}\left[\frac{1}{2}\bar{V}_{ww}\bar{w}^2\left(\hat{w}_t^{\mathcal{C}} - \hat{w}_t\right)^2\middle|\pi_t\right]\right] \\
=& -\mathbb{E}_{t-1}\left[\frac{1}{2}\bar{V}_{ww}\bar{w}^2\left(\hat{w}_t^{\mathcal{C}} - \hat{w}_t\right)^2\right]
\end{aligned}$$

where the third to the fourth line uses the law of iterated expectations, the fourth to the fifth line uses the assumption that $\hat{w}_t$ depends only on π_t given information set $\mathcal{F}_{t-1}$, and the fifth to the sixth line uses the optimality condition of $\hat{w}_t^{\mathcal{C}}$. Since for any non-contingent wage W_t, the

corresponding realized real wage depends only on π_t conditional on $\mathcal{F}_{t-1}$, the expected loss from W_t is

$$\begin{aligned}\mathcal{L}_{t-1}(W_t) &= -\mathbb{E}_{t-1}\left[\frac{1}{2}\bar{V}_{ww}\bar{w}^2\left(\hat{w}_t^{\mathcal{C}} - \hat{w}_t\right)^2\right] \\ &= -\mathbb{E}_{t-1}\left[\frac{1}{2}\bar{V}_{ww}\bar{w}^2\left(\log W_t^{\mathcal{C}} - \log W_t\right)^2\right] \\ \implies \log W_t^{\mathcal{N}} &= \mathbb{E}_{t-1}\left[\log W_t^{\mathcal{C}}\right]\end{aligned}$$

In real terms, we have

$$\begin{aligned}\hat{w}_t^{\mathcal{N}} &= \log W_t^{\mathcal{N}} - \log P_{t-1} - \pi_t \\ &= \mathbb{E}_{t-1}\left[\log W_t^{\mathcal{C}}\right] - \log P_{t-1} - \pi_t \\ &= \mathbb{E}_{t-1}\left[\hat{w}_t^{\mathcal{C}} + \pi_t\right] - \pi_t \\ &= \mathbb{E}_{t-1}\left[\hat{w}_t^{\mathcal{C}}\right] + \mathbb{E}_{t-1}\left[\pi_t\right] - \pi_t\end{aligned}$$

The minimal expected loss from non-contingent wages compared to COLA is therefore

$$\begin{aligned}\mathcal{V}_{t-1} &= -\frac{1}{2}\bar{V}_{ww}\bar{w}^2\mathbb{E}_{t-1}\left[\left(\hat{w}_t^{\mathcal{C}} - \mathbb{E}_{t-1}\left[\hat{w}_t^{\mathcal{C}}\right] + \pi_t - \mathbb{E}_{t-1}\left[\pi_t\right]\right)^2\right] \\ &= -\frac{1}{2}\bar{V}_{ww}\bar{w}^2\left(\mathrm{Var}_{t-1}\left[\hat{w}_t^{\mathcal{C}}\right] + \mathrm{Var}_{t-1}\left[\pi_t\right] + \mathrm{Cov}_{t-1}\left[\hat{w}_t^{\mathcal{C}}, \pi_t\right]\right)\end{aligned}$$

□

Hence, non-contingent wages are even coarser instruments compared to COLA: they try to target COLA wages in nominal terms in expectation conditional on the information at contracting. In real terms, non-contingent wages also try to target COLA, but will also adjust for inflation forecast error. The second part provides a characterization of the loss from non-contingent wages (or the gain from COLA). We cannot say more without specifying the underlying information structure. However, if $\hat{w}_t^{\mathcal{C}}$ is proportional to π_t, the loss function can be greatly simplified. This amounts to saying that the forecast of $\hat{\zeta}_t$ conditional on π_t is linear in π_t. We will derive this proportionality endogenously in the general equilibrium setting.

D General-Equilibrium Environment and Derivations

This section embeds the contracting model into a New Keynesian environment. On the one hand, this endogenizes the inflation and revenue processes, which are key inputs to the contracting problem. On the other hand, it demonstrates how individual contract choices affect aggregate dynamics. The main text lays out the environment. In what follows, I first characterize the no-shock steady state. Then, I expand the economy around the steady state, with the exception that

I leave the contract composition nonlinear. I then identify the novel aggregate state variables introduced by the contracting problem. I will analyze the model in the next two sections.

D.1 The No-Shock Steady State

We first characterize the no-shock steady state. The no-shock steady state is the economy in which

$$\sigma_{ut} = \sigma_{\rho t} \equiv 0$$

In other words, the markup shock and the demand shock are muted, but each individual firm still faces an idiosyncratic shock. Throughout, $\bar{x}$ means the steady-state value of the variable x. I assume there is no trend inflation, i.e., $\bar{\pi} = 0$.

D.1.1 The Contracting Block

We first look at the contracting problem. We analyze the problem backwards. Because firms can reset employment each period, and we assume two-sided commitment to the contract, the employment choice problem can be solved period by period. In each period, given the real wage w and the revenue factor ζ, a firm chooses employment to maximize the expected value before its idiosyncratic shock realizes:

$$\begin{aligned} &\max_{L} \mathbb{E}^{\eta} \left[q - \phi_0 q^2 \mathbb{I}(q < 0) \right] \\ &s.t.\ q = \eta \zeta L^{1-\frac{1}{\epsilon_x}} - wL \end{aligned} \tag{13}$$

Denote by $Q(\zeta, w)$ the indirect value function, and $L^*(\zeta, w)$ the optimal labor choice.[44] Both the policy function and the value function hold in any equilibrium, and in particular, in the steady state, employment is

$$\bar{L} = L^*(\bar{\zeta}, \bar{w}) \tag{14}$$

In the wage-setting stage at each contract renewal, they choose the optimal wage rule conditional on the contract type. Because of (1) the flexible employment decision, (2) the two-sided commitment, and (3) the assumption that they can have a different wage rule each period, we can split the wage-setting problem period by period. For a time-t renewing pair that considers wage rules in period $t+s$, the COLA wage solves

$$W^{\mathcal{C}}_{t+s,s}\left(\pi^{t+s}\right) \in \arg \max_{W_{t+s}(\pi^{t+s})} \mathbb{E}_t \Bigg[L^* \left(\zeta_{t+s}, \frac{W_{t+s}\left(\pi^{t+s}\right)}{P_{t+s}} \right) \left[u\left(\frac{W_{t+s}\left(\pi^{t+s}\right)}{P_{t+s}} \right) - u(b) \right] \tag{15}$$

$$+ \lambda Q \left(\zeta_{t+s}, \frac{W_{t+s}\left(\pi^{t+s}\right)}{P_{t+s}} \right) \Bigg] \tag{16}$$

[44] The service firm's pricing decision has already been incorporated into the employment choice problem, since given the demand curve, choosing employment is equivalent to choosing a service price.

while the non-contingent wage solves

$$W^{\mathcal{N}}_{t+s,s} \in \arg\max_{W_{t+s}} \mathbb{E}_t \left[L^* \left(\zeta_{t+s}, \frac{W_{t+s}}{P_{t+s}} \right) \left[u \left(\frac{W_{t+s}}{P_{t+s}} \right) - u(b) \right] + \lambda Q \left(\zeta_{t+s}, \frac{W_{t+s}}{P_{t+s}} \right) \right]$$

where the subscript t, s means they are chosen at the end of period t, and they are the wage rules in period $t + s$.

We analyze their properties in detail after we expand the problem to second order. For the no-shock steady state, because there is no uncertainty, the two wage rules must be identical.[45] Additionally, the steady-state revenue factor is revealed and publicly observable. Therefore,

$$\bar{W}^{\mathcal{C}} = \bar{W}^{\mathcal{N}} = W^*(\bar{\zeta})$$

Finally, in the contract choice stage, COLA is chosen if its gain exceeds the adoption cost. In the steady state, however, since both contracts prescribe identical wages, they also yield the same joint values. Because COLA requires a positive cost, firm-union pairs will strictly prefer non-contingent contracts in the steady state. In addition, since there is no trend inflation, nominal wages and real wages are always proportional to each other. I will work with the steady-state real wage $\bar{w}$ next, where

$$\bar{w} = w^*(\bar{\zeta}) := \frac{W^*(\bar{\zeta})}{\bar{P}} \tag{17}$$

D.1.2 Other Blocks

I analyze other blocks that pin down the general equilibrium together with the contracting block. Here, I only characterize them in steady state. I leave the characterization of the other blocks out of the steady state to the next subsection.

The revenue factor $\bar{\zeta}$ summarizes a service firm's aggregate demand conditions and is given by

$$\bar{\zeta} = \frac{\bar{P}^X}{\bar{P}} \bar{X}^{\frac{1}{\varepsilon_X}} \tag{18}$$

where $\bar{P}^X$ is the price index of the service bundle, $\bar{P}$ is the price index of the final good, $\bar{X}$ is the downstream demand for the service bundle, all of which are in the steady state. Since all service firms are ex ante identical in the steady state, and their idiosyncratic shocks are i.i.d. over time, their labor inputs are the same. Therefore, the aggregate service bundle is

$$\bar{X} = \bar{L} \tag{19}$$

We next look at those downstream firms. In the no-shock steady state, the Calvo pricing-resetting friction is irrelevant. Their pricing decision is simply a markup over the marginal cost.

[45] This will be true even if we have trend inflation, as the trend is anticipated and will be compensated by non-contingent wages as well.

Given the linear production function, the marginal cost is exactly the price of the service bundle, i.e.,

$$\bar{P} = \frac{\bar{\varepsilon}_y}{\bar{\varepsilon}_y - 1} \bar{P}^X \tag{20}$$

where $\bar{\varepsilon}_y$ is the downstream elasticity of substitution in the steady state.

Equations (14)-(20) already allow us to solve for the wage, employment, and the revenue factor in the steady state. In particular, $\bar{L}$ and $\bar{\zeta}$ jointly solve

$$\bar{L} = \left(\frac{\bar{\varepsilon}_y - 1}{\bar{\varepsilon}_y} \bar{\zeta} \right)^{\varepsilon_x}$$
$$\bar{L} = L^* \left(\bar{\zeta}, w^*(\bar{\zeta}) \right)$$

We also derive some other steady-state quantities for later use. Similar to the upstream firms, downstream firms also set the same price and use the same amount of service bundle in the steady state. Therefore, total output is simply

$$\bar{Y} = \bar{X}$$

Part of the output is used to cover the external financing costs by upstream firms who get low idiosyncratic shocks. Denote by $\bar{\Phi}$ the total resources directed to external financing. This is assumed to be a pure waste and is given by

$$\bar{\Phi} = \phi_0 \int_0^{\frac{\bar{w}\bar{L}^{1/\varepsilon_x}}{\bar{\zeta}}} \left(\bar{w}\bar{L} - \eta \bar{\zeta} \bar{L}^{1-1/\varepsilon_x} \right)^2 dF_\eta(\eta)$$

where F_η is the CDF of the idiosyncratic shock.

Finally, from the aggregate resource constraint, the aggregate consumption of entrepreneurs in the steady state is

$$\bar{C}^{\text{ent}} = \bar{Y} - \bar{w}\bar{L} - \bar{\Phi}$$

In what follows, I assume structural parameters are such that a unique interior steady state exists. I also assume $\varepsilon_x > 1$ and $\bar{\varepsilon}_y > 1$, so both differentiated services and differentiated varieties are substitutes within their groups.

D.2 Expanding the Model

I next expand the model around the no-shock steady state derived in the last subsection. I will keep the contract distribution nonlinear. In the data, we do observe large movements of the COLA adoption rate, making the validity of a full expansion questionable. As one will see shortly, the interaction of the contract distribution and other aggregate variables is important for the transmission of shocks. I label the model as a "semi-linear" model.

D.2.1 Upstream Firms

We first derive the policy functions of an individual firm-union pair after we expand the firm value and the joint value to second order. Based on those policy functions, I characterize the contract type choice, which is intrinsically nonlinear.

Contract choice. We again analyze the problem backwards. Given revenue factor ζ_t and real wage w_t, at the employment choice stage, expanding the objective function (13) to second order, one can get the (log) linear policy function. Alternatively, we can simply expand the nonlinear policy function (14) to first order. Either way, the linear labor input policy is

$$\hat{l}_{it} = \iota_1 \hat{\zeta}_t - \iota_2 \hat{w}_{it} \tag{21}$$

where

$$\iota_1 = \bar{L}^*_\zeta \bar{\zeta} / \bar{L} > 0$$
$$\iota_2 = -\bar{L}^*_w \bar{w} / \bar{L} > 0$$

In the wage-setting stage, we first study COLA. Since COLA allows the nominal wage to be a function of the concurrent inflation history, it is equivalent to choosing a real wage under each inflation realization. In other words, we can recast the objective (15) in real terms:

$$w^{\mathcal{C}}_{t+s,s} \in \arg\max_{w_{t+s}} \mathbb{E}_t \left[\underbrace{L^* \left(\zeta_{t+s}, w_{t+s}\right) \left[u\left(w_{t+s}\right) - u(b)\right] + \lambda Q\left(\zeta_{t+s}, w_{t+s}\right)}_{:= V(\zeta_{t+s}, w_{t+s})} \middle| \pi^{t+s} \right]$$

Following the same procedure as in Section C.5, we have

$$\hat{w}^{\mathcal{C}}_{t+s,s} = \mu \mathbb{E}_t \left[\hat{\zeta}_{t+s} \middle| \pi^{t+s} \right] \tag{22}$$

where

$$\mu = \frac{\bar{V}_{\zeta w} \bar{\zeta}}{-\bar{V}_{ww} \bar{w}}$$

The sign of μ deserves some discussion here. I show in Section C.5 that under fixed employment, the μ following the same definition is positive. This means that a higher revenue factor always raises the wage payment in that case. With variable labor demand, however, there is another force that pushes μ towards the other direction. Specifically, a higher revenue factor induces a higher optimal labor demand. Holding the wage fixed, more employment may push the firm closer to incurring costly external financing even if the revenue factor rises. That may then create an incentive to reduce the wage payment. The union is happy with that because the total wage bill may still rise.

The second force can be weak if, for example, workers have relatively strong outside options (their home production b). In that case, workers are more indifferent between employment and unemployment, so they may not be willing to cut wages in exchange for more employment when the revenue factor rises. Alternatively, if the firm's weight λ in the wage bargaining problem is relatively low, the firm's increased financial stress may be less of a concern in wage bargaining. That can also weaken the second force. To simplify the subsequent analysis, I will impose the condition that $\mu > 0$ throughout the theoretical part. This is the empirically relevant range based on the calibration in the quantitative part.

We then turn to the non-contingent wages. Similar to the derivation in Section C.6, the optimal non-contingent wage tries to target COLA in nominal terms with the available information. Namely,

$$\log W^{\mathcal{N}}_{t+s,s} = \mathbb{E}_t \left[\log W^{\mathcal{C}}_{t+s,s} \right]$$

Or in real terms, the optimal non-contingent wage still tries to target COLA, but adjusts for cumulative inflation forecast errors:

$$\begin{aligned}
\hat{w}^{\mathcal{N}}_{t+s,s} &= \log W^{\mathcal{N}}_{t+s,s} - \log P_t - \sum_{r=1}^{s} \pi_{t+r} \\
&= \mathbb{E}_t \left[\log W^{\mathcal{C}}_{t+s,s} \right] - \log P_t - \sum_{r=1}^{s} \pi_{t+r} \\
&= \mathbb{E}_t \left[\hat{w}^{\mathcal{C}}_{t+s,s} + \sum_{r=1}^{s} \pi_{t+r} \right] - \sum_{r=1}^{s} \pi_{t+r} \\
&= \mathbb{E}_t \left[\hat{w}^{\mathcal{C}}_{t+s,s} \right] + \sum_{r=1}^{s} \left(\mathbb{E}_t \left[\pi_{t+r} \right] - \pi_{t+r} \right)
\end{aligned} \tag{23}$$

Next, at the contract type choice, the loss structure derived in Section C.6 applies to the long-term contract here period by period. In particular, when a firm-union pair renews their contract at the end of period t, the expected loss from non-contingent wages is

$$\mathcal{V}_t = -\frac{1}{2} \bar{V}_{ww} \bar{w}^2 \sum_{s=1}^{N} \beta^s \mathbb{E}_t \left[\left(\hat{w}^{\mathcal{C}}_{t+s,s} - \mathbb{E}_t \left[\hat{w}^{\mathcal{C}}_{t+s,s} \right] + \sum_{r=1}^{s} \left(\pi_{t+r} - \mathbb{E}_t \left[\pi_{t+r} \right] \right) \right)^2 \right] \tag{24}$$

Therefore, the pair from industry i prefers COLA if

$$\mathcal{V}_t > \chi_i$$

Aggregate service firm production decisions. Upstream firms take the revenue factor ζ_t as given, which, in the log deviation term, is

$$\hat{\zeta}_t = \hat{p}^X_t + \frac{1}{\varepsilon_x} \hat{x}_t \tag{25}$$

where $p_t^X := P_t^X / P_t$ is the real price of the service bundle.

Upstream firms are ex ante heterogeneous in the sense that they may choose different contracts, and their wage rules are conditional on different information sets, depending on the time of contracting. As a result, the realized real wages they face, and therefore, the levels of employment, can be different under the same aggregate shock realizations. Denote by $\mathcal{M}_{t,k}$ the mass of firms that adopt COLA in period $t-k$. Then, mass $\frac{1}{N} - \mathcal{M}_{t,k}$ from the $t-k$ renewing cohort adopts non-contingent contracts, following the Taylor staggered contract structure. The vector $\{\mathcal{M}_{t,k}\}_{k=1}^N$, labeled as the contract distribution (or composition), is therefore an aggregate state that one needs to keep track of in order to aggregate service firm production decisions in period t. I take the vector as given in this subsection. I will discuss and highlight the dynamics (law of motion) of the contract distribution in Section D.2.3.

Given the contract distribution $\{\mathcal{M}_{t,k}\}_{k=1}^N$, up to first order, the aggregate production function of the service bundle is

$$\hat{x}_t = \sum_{k=1}^{N} \left[\mathcal{M}_{t,k} \hat{l}_{t,k}^{\mathcal{C}} + \left(\frac{1}{N} - \mathcal{M}_{t,k} \right) \hat{l}_{t,k}^{\mathcal{N}} \right]$$

where $\hat{l}_{t,k}^{\mathcal{C}}$ is the labor input of firms who adopted COLA in period $t-k$, and $\hat{l}_{t,k}^{\mathcal{N}}$ is the same object but for firms who adopted non-contingent wages in period $t-k$. In words, the total production of the service bundle is a weighted average of the labor inputs of different firms weighted by their shares. Combining with equation (21), we can further write it as

$$\begin{aligned}
\hat{x}_t &= \sum_{k=1}^{N} \left[\mathcal{M}_{t,k} \left(\iota_1 \hat{\zeta}_t - \iota_2 \hat{w}_{t,k}^{\mathcal{C}} \right) + \left(\frac{1}{N} - \mathcal{M}_{t,k} \right) \left(\iota_1 \hat{\zeta}_t - \iota_2 \hat{w}_{t,k}^{\mathcal{N}} \right) \right] \\
&= \iota_1 \hat{\zeta}_t - \iota_2 \sum_{k=1}^{N} \left[\mathcal{M}_{t,k} \hat{w}_{t,k}^{\mathcal{C}} + \left(\frac{1}{N} - \mathcal{M}_{t,k} \right) \hat{w}_{t,k}^{\mathcal{N}} \right] \\
&= \iota_1 \hat{\zeta}_t - \iota_2 \hat{\omega}_t
\end{aligned} \tag{26}$$

where I define the aggregate real wage as the weighted average of the realized real wages faced by different firms, i.e.,

$$\hat{\omega}_t := \sum_{k=1}^{N} \left[\mathcal{M}_{t,k} \hat{w}_{t,k}^{\mathcal{C}} + \left(\frac{1}{N} - \mathcal{M}_{t,k} \right) \hat{w}_{t,k}^{\mathcal{N}} \right] \tag{27}$$

Hence, the total available service bundle increases with the revenue factor but decreases with the aggregate real wage.

D.2.2 Downstream Firms, Entrepreneurs, Aggregate Resources, and Monetary Policy

I describe those more standard blocks in this subsection. The Calvo price-resetting friction and the pricing decision of those downstream firms lead to the standard New Keynesian Phillips curve (NKPC) in marginal cost version to first-order:

$$\pi_t = \beta \mathbb{E}_t[\pi_{t+1}] + \kappa \hat{mc}_t + u_t \tag{28}$$

where u_t is the markup shock coming from the changing elasticity of substitution in the downstream. Given the linear production function, the real marginal cost is exactly the real price of the service bundle, i.e.,

$$\hat{mc}_t = \hat{p}_t^X \tag{29}$$

Up to first order, the aggregate production function is linear, i.e.,

$$\hat{y}_t = \hat{x}_t \tag{30}$$

Entrepreneurs' consumption-saving decision leads to the standard Euler equation, which is the dynamic IS curve to first order:

$$\hat{c}_t^{\text{ent}} = \mathbb{E}_t\left[\hat{c}_{t+1}^{\text{ent}}\right] - \gamma(i_t - \mathbb{E}_t[\pi_{t+1}] - \rho_t) \tag{31}$$

where ρ_t is a demand shock coming from the time variation of the entrepreneurs' discount factor. The nominal rate i_t follows the standard Taylor rule

$$i_t = \bar{i} + \phi_\pi \pi_t + \phi_y \hat{y}_t \tag{32}$$

Aggregate resources are divided into three parts: entrepreneur consumption, worker consumption, and payments to the external financing costs. Notice that I microfound the COLA cost as a calculation and planning cost for workers, which is in utility terms. One can think of the COLA cost as directly reducing a worker's value. So it does not enter the aggregate physical resource constraint. The complete resource constraint is

$$Y_t = C_t^{\text{ent}} + \sum_{k=1}^{N}\left[\mathcal{M}_{t,k} w_{t,k}^{\mathcal{C}} L_{t,k}^{\mathcal{C}} + \left(\frac{1}{N} - \mathcal{M}_{t,k}\right) w_{t,k}^{\mathcal{N}} L_{t,k}^{\mathcal{N}}\right] + \sum_{k=1}^{N}\left[\mathcal{M}_{t,k}\Phi_{t,k}^{\mathcal{C}} + \left(\frac{1}{N} - \mathcal{M}_{t,k}\right)\Phi_{t,k}^{\mathcal{N}}\right]$$

where

$$\Phi_{t,k}^{\mathcal{C}} := \phi_0 \int_0^{\frac{w_{t,k}^{\mathcal{C}}\left(L_{t,k}^{\mathcal{C}}\right)^{1/\varepsilon_x}}{\zeta_t}} \left(w_{t,k}^{\mathcal{C}} L_{t,k}^{\mathcal{C}} - \eta_t \zeta_t \left(L_{t,k}^{\mathcal{C}}\right)^{1-1/\varepsilon_x}\right)^2 dF_\eta(\eta_t)$$

and $\Phi_{t,k}^{\mathcal{N}}$ is defined similarly. We now log-linearize this equation. Y_t and C_t^{ent} are easy:

$$Y_t \approx \bar{Y}(1 + \hat{y}_t), \quad C_t^{\text{ent}} \approx \bar{C}^{\text{ent}}(1 + \hat{c}_t^{\text{ent}})$$

The second term on the right-hand-side (RHS) is the total wage payment, which, denoted by C_t^w and up to first order, is

$$\begin{aligned}
C_t^w \approx& \bar{C}^w + \sum_{k=1}^{N} \mathcal{M}_{t,k} \bar{w} \bar{L} \left(\hat{w}_{t,k}^{\mathcal{C}} + \hat{l}_{t,k}^{\mathcal{C}} \right) + \left(\frac{1}{N} - \mathcal{M}_{t,k} \right) \bar{w} \bar{L} \left(\hat{w}_{t,k}^{\mathcal{N}} + \hat{l}_{t,k}^{\mathcal{N}} \right) \\
=& \bar{C}^w + \bar{w} \bar{L} \hat{\omega}_t + \bar{w} \bar{L} \sum_{k=1}^{N} \left[\mathcal{M}_{t,k} \left(\iota_1 \hat{\zeta}_t - \iota_2 \hat{w}_{t,k}^{\mathcal{C}} \right) + \left(\frac{1}{N} - \mathcal{M}_{t,k} \right) \left(\iota_1 \hat{\zeta}_t - \iota_2 \hat{w}_{t,k}^{\mathcal{N}} \right) \right] \\
=& \bar{C}^w + \bar{w} \bar{L} \left(\iota_1 \hat{\zeta}_t + (1 - \iota_2) \hat{\omega}_t \right)
\end{aligned}$$

where the first to the second line uses the labor input policy (21), and the second to the last line uses the definition of the aggregate real wage index (27). The last term on the RHS is the total external financing cost, which, denoted by Φ_t and up to first order, is

$$\begin{aligned}
\Phi_t \approx& \bar{\Phi} + 2\phi_0 \sum_{k=1}^{N} \mathcal{M}_{t,k} \int_0^{\bar{\tau}} (\bar{L}\bar{w} - \eta_t \bar{R}) \left[\bar{L}\bar{w} \hat{l}_{t,k}^{\mathcal{C}} + \bar{L}\bar{w} \hat{w}_{t,k}^{\mathcal{C}} - \eta_t \left(\bar{R} \hat{\zeta}_t + \frac{\varepsilon_x - 1}{\varepsilon_x} \bar{R} \hat{l}_{t,k}^{\mathcal{C}} \right) \right] dF(\eta_t) \\
&+ 2\phi_0 \sum_{k=1}^{N} \left(\frac{1}{N} - \mathcal{M}_{t,k} \right) \int_0^{\bar{\tau}} (\bar{L}\bar{w} - \eta_t \bar{R}) \left[\bar{L}\bar{w} \hat{l}_{t,k}^{\mathcal{N}} + \bar{L}\bar{w} \hat{w}_{t,k}^{\mathcal{N}} - \eta_t \left(\bar{R} \hat{\zeta}_t + \frac{\varepsilon_x - 1}{\varepsilon_x} \bar{R} \hat{l}_{t,k}^{\mathcal{N}} \right) \right] dF(\eta_t) \\
=& \bar{\Phi} + 2\phi_0 \left[I_1 \bar{L}\bar{w} \left(\iota_1 \hat{\zeta}_t - \iota_2 \hat{\omega}_t + \hat{\omega}_t \right) - \bar{R} I_2 \left(\hat{\zeta}_t + \frac{\varepsilon_x - 1}{\varepsilon_x} (\iota_1 \hat{\zeta}_t - \iota_2 \hat{\omega}_t) \right) \right]
\end{aligned}$$

where

$$\begin{aligned}
\bar{\tau} :=& \frac{\bar{w} \bar{L}^{1/\varepsilon_x}}{\bar{\zeta}}, \quad \bar{R} := \bar{\zeta} \bar{L}^{1 - 1/\varepsilon_x} \\
I_1 :=& \int_0^{\bar{\tau}} (\bar{w}\bar{L} - \eta_t \bar{R}) dF(\eta_t) \\
I_2 :=& \int_0^{\bar{\tau}} (\bar{w}\bar{L} - \eta_t \bar{R}) \eta_t dF(\eta_t)
\end{aligned}$$

and we use the labor input policy and the definition of the aggregate real wage similarly. Combining the three parts together, and eliminating the steady state parts, the aggregate resource constraint can be written as

$$\bar{Y} \hat{y}_t = \bar{C}^{\text{ent}} \hat{c}_t^{\text{ent}} + \bar{w} \bar{L} \left(\iota_1 \hat{\zeta}_t + (1 - \iota_2) \hat{\omega}_t \right) + 2\phi_0 \left[I_1 \bar{L}\bar{w} \left(\iota_1 \hat{\zeta}_t - \iota_2 \hat{\omega}_t + \hat{\omega}_t \right) - \bar{R} I_2 \left(\hat{\zeta}_t + \frac{\varepsilon_x - 1}{\varepsilon_x} (\iota_1 \hat{\zeta}_t - \iota_2 \hat{\omega}_t) \right) \right] \tag{33}$$

D.2.3 Contract Distribution as an Aggregate State

In this section, I characterize the dynamics (or the law of motion) of the contract distribution, which is obtained by aggregating the individual firm-union contract choice decisions.

Let $\{\mathcal{M}_{t,k}\}_{k=1}^N$ be the contract distribution in period t. We want to derive the contract distribution in period $t+1$, $\{\mathcal{M}_{t+1,k}\}_{k=1}^N$. At the end of period t, the $t-N$ cohort renews their contracts,

and they will be the $(t+1)-1$ cohort from time $t+1$ perspective. Therefore, given the cost distribution F_χ, the mass of new COLA adoption from the time-t renewing cohort is

$$\begin{aligned}\mathcal{M}_{t+1,1} &= \frac{1}{N}\mathbb{P}\left[\mathcal{V}_t > \chi_i\right] \\ &= \frac{1}{N}F_\chi(\mathcal{V}_t)\end{aligned} \tag{34}$$

where the gain from COLA $\mathcal{V}_t$ is determined by equation (24). All other cohorts only shift their index forward without actually changing their contracts, i.e.,

$$\mathcal{M}_{t+1,k} = \mathcal{M}_{t,k-1}, \quad k \in \{2, ..., N\} \tag{35}$$

The contract distribution is a novel endogenous aggregate state this paper identifies. It affects the aggregate wage and employment composition, as shown in the last subsection, and is affected by the aggregate dynamics through the gain from COLA. The endogenous contract choice makes the model intrinsically nonlinear. It also makes the model complicated: agents have to form their expectations about the future path of the contract distribution, and their expectations should be rational (either coincide with the actual path if it is nonstochastic, or have zero expected forecast error).

D.2.4 The Forecast Problems

Finally, one may have already noticed that many equations from the last few subsections contain forecasts. In particular, for those who adopt COLA, they need to forecast the revenue factor with the information set at contracting and the concurrent inflation history, i.e.,

$$\mathbb{E}\left[\hat{\zeta}_{t+k}\middle|\mathcal{F}_t, \pi^{t+k}\right]$$

For those who choose non-contingent wages, they need to forecast both revenue and inflation using the information at contracting:

$$\mathbb{E}\left[\hat{\zeta}_{t+k}\middle|\mathcal{F}_t\right], \quad \mathbb{E}\left[\pi_{t+k}\middle|\mathcal{F}_t\right]$$

Finally, at the contract choice stage, to calculate the gain from COLA, they need to know the entire future variance-covariance structure:

$$\mathrm{Cov}_t\left[\hat{w}^{\mathcal{C}}_{t+k}, \pi_{t+s}\right], \quad k, s \in \{1, .., N\}$$

All of them pose great challenges to the problem: $\hat{\zeta}_{t+k}, \pi_{t+k}, \hat{w}^{\mathcal{C}}_{t+k}$ are all endogenous stochastic processes that depend on the forecasts and all other equations in the previous subsections, but one needs to know the stochastic processes in order to solve for the forecast problems (e.g., applying

Kalman filters if these processes are linear, or applying nonlinear filters if they are not). In other words, the problem involves a fixed point between forecasts and endogenous variables.

The forecast problems, especially those from firms using non-contingent contracts, also introduce a backward-looking feature to the model, as their wages and labor inputs depend on past information sets. Therefore, strictly speaking, the aggregate states should also contain past information sets $\{\mathcal{F}_{t-k}\}_{k\geq 0}$.

D.3 Simplifying the Semi-Linear System

The main text defines the conditional general equilibrium and the contracting general equilibrium. I summarize here the minimal set of equations with clear economic interpretations that can characterize them.

Lemma 5. *Given the two AR(1) shock processes of u_t and ρ_t, and some initial contract distribution $\{\mathcal{M}_{0,k}\}_{k=1}^N$, the contracting general equilibrium is characterized by*

$$\pi_t = \beta \mathbb{E}_t[\pi_{t+1}] + \kappa \left[\left(1 - \frac{\iota_1}{\varepsilon_x}\right) \hat{\zeta}_t + \iota_2 \frac{1}{\varepsilon_x} \hat{\omega}_t \right] + u_t \tag{36}$$

$$\hat{\omega}_t = \sum_{k=1}^{N} \mathcal{M}_{t,k} \mu \mathbb{E}_{t-k}[\hat{\zeta}_t | \pi^t] - \sum_{k=1}^{N} \left(\frac{1}{N} - \mathcal{M}_{t,k}\right) \left[\mu \mathbb{E}_{t-k}[\hat{\zeta}_t] + \sum_{\tau=0}^{k-1} \left(\pi_{t-\tau} - \mathbb{E}_{t-k}[\pi_{t-\tau}]\right) \right] \tag{37}$$

$$\hat{c}_t^{\text{ent}} = \mathbb{E}_t[\hat{c}_{t+1}^{\text{ent}}] - \gamma(i_t - \mathbb{E}_t[\pi_{t+1}] - \rho_t) \tag{38}$$

$$\hat{c}_t^{\text{ent}} = \frac{\bar{Y}}{\bar{C}^{\text{ent}}} \left[(\iota_1 - \varphi_\zeta)\hat{\zeta}_t - (\iota_2 + \varphi_\omega)\hat{\omega}_t \right] \tag{39}$$

$$i_t = \phi_\pi \pi_t + \phi_y \left[(\iota_1 \hat{\zeta}_t - \iota_2 \hat{\omega}_t) \right] \tag{40}$$

$$\mathcal{V}_t = -\frac{1}{2} \bar{V}_{ww} \bar{w}^2 \sum_{s=1}^{N} \beta^s \mathbb{E}_t \left[\left(\hat{w}_{t+s}^{\mathcal{C}} - \mathbb{E}_t \left[\hat{w}_{t+s}^{\mathcal{C}} \right] + \sum_{r=1}^{s} \left(\pi_{t+r} - \mathbb{E}_t \left[\pi_{t+r} \right] \right) \right)^2 \right]$$

$$\mathcal{M}_{t+1,1} = \frac{1}{N} F_\chi(\mathcal{V}_t)$$

$$\mathcal{M}_{t+1,k} = \mathcal{M}_{t,k-1}, \quad k \in \{2, ..., N\}$$

where

$$\varphi_\zeta = \frac{\bar{w}\bar{L}\iota_1 + 2\phi_0 \left(I_1 \bar{w}\bar{L}\iota_1 - I_2 \bar{R} \left(1 + \iota_1 \frac{\varepsilon_x - 1}{\varepsilon_x} \right) \right)}{\bar{Y}}$$

$$\varphi_\omega = \frac{\bar{w}\bar{L}(1 - \iota_2) + 2\phi_0 \left(I_1 \bar{w}\bar{L}(1 - \iota_2) + I_2 \bar{R} \iota_2 \frac{\varepsilon_x - 1}{\varepsilon_x} \right)}{\bar{Y}}$$

Proof. Combining equations (25), (26), and (29), the real price of the service bundle, which is also the real marginal cost to the downstream, is

$$\hat{p}_t^X = \hat{mc}_t = \left(1 - \frac{\iota_1}{\varepsilon_x}\right) \hat{\zeta}_t + \frac{\iota_2}{\varepsilon_x} \hat{\omega}_t$$

Plugging in equation (28), we have

$$\pi_t = \beta \mathbb{E}_t[\pi_{t+1}] + \kappa \left[\left(1 - \frac{\iota_1}{\varepsilon_x} \right) \hat{\zeta}_t + \iota_2 \frac{1}{\varepsilon_x} \hat{\omega}_t \right] + u_t$$

Equations (26) and (30) lead to a simple aggregate production function

$$\hat{y}_t = \iota_1 \hat{\zeta}_t - \iota_2 \hat{\omega}_t$$

Rearranging the resource constraint (33), we can express $\hat{y}_t$ in terms of entrepreneur consumption, the revenue factor, and the aggregate real wage:

$$\hat{y}_t = \frac{\bar{C}^{\text{ent}}}{\bar{Y}} \hat{c}_t^{\text{ent}} + \underbrace{\frac{\bar{w}\bar{L}\iota_1 + 2\phi_0 \left(I_1 \bar{w}\bar{L}\iota_1 - I_2 \bar{R} \left(1 + \iota_1 \frac{\varepsilon_x - 1}{\varepsilon_x} \right) \right)}{\bar{Y}}}_{:=\varphi_\zeta} \hat{\zeta}_t + \underbrace{\frac{\bar{w}\bar{L}(1-\iota_2) + 2\phi_0 \left(I_1 \bar{w}\bar{L}(1-\iota_2) + I_2 \bar{R} \iota_2 \frac{\varepsilon_x - 1}{\varepsilon_x} \right)}{\bar{Y}}}_{:=\varphi_\omega} \hat{\omega}_t$$

Therefore, the φ_ζ and φ_ω measure the elasticity of output with respect to the revenue factor and the aggregate real wage, respectively, holding the entrepreneur consumption fixed. Eliminating $\hat{y}_t$ using the last two equations, we get

$$\hat{c}_t^{\text{ent}} = \frac{\bar{Y}}{\bar{C}^{\text{ent}}} \left[(\iota_1 - \varphi_\zeta) \hat{\zeta}_t - (\iota_2 + \varphi_\omega) \hat{\omega}_t \right]$$

Replacing $\hat{y}_t$ in the Taylor rule using the aggregate production function, we obtain

$$i_t = \phi_\pi \pi_t + \phi_y \left[(\iota_1 \hat{\zeta}_t - \iota_2 \hat{\omega}_t) \right]$$

The rest of the equations are directly from the last few subsections. □

Notice the first five equations characterize the conditional general equilibrium for any given time path of the contract distribution $\{\mathcal{M}_{t,k}\}_{t \geq 0, k \in \{1,\ldots,N\}}$. The last three equations describe the law of motion of the contract distribution over time. Together, the eight equations characterize the contracting general equilibrium. This system is different from a standard log-linearized DSGE system in three places. First, the contract distribution is kept nonlinear and interacts with other macro variables through the aggregate real wage index. Second, the system involves many forecasts conditional on different information sets, including the forecast of revenue conditional on concurrent inflation history, both of which are endogenous variables. Third, the dynamics of the contract distribution require second-order moments. All features make the model intrinsically nonlinear.

Before I move on, the next lemma gives the relationships between some elasticities from the steady state that I defined earlier. They will be useful to sign terms in the next section.

Lemma 6. *At a positive interior steady state with $\gamma > 0$, $\kappa > 0$, and $\phi_y \geq 0$,*

$$0 < \iota_1 < \varepsilon_x, \qquad \iota_2 = 1 + \frac{\varepsilon_x - 1}{\varepsilon_x}\iota_1 > 1$$

$$\varphi_\omega = -\frac{\varepsilon_x - 1}{\varepsilon_x}\varphi_\zeta, \qquad \iota_1 - \varphi_\zeta > 0$$

Proof. Following the notations in Section D.2, differentiating the firm's objective function for labor demand and evaluating its first-order condition at the steady state gives

$$0 = \mathbb{E}^\eta \left[(1 - 2\phi_0(\eta\bar{R} - \bar{w}\bar{L})\mathbb{I}(\eta < \bar{\tau})) \left(\frac{\varepsilon_x - 1}{\varepsilon_x}\frac{\eta\bar{R}}{\bar{L}} - \bar{w} \right) \right]$$

Using $\mathbb{E}[\eta] = 1$ and multiplying by $\bar{L}$ yields

$$\bar{w}\bar{L}(1 + 2\phi_0 I_1) = \frac{\varepsilon_x - 1}{\varepsilon_x}\bar{R}(1 + 2\phi_0 I_2) \tag{41}$$

So,

$$\frac{\varepsilon_x - 1}{\varepsilon_x} - \bar{\tau} = 2\phi_0\bar{R}\int_0^{\bar{\tau}} (\bar{\tau} - \eta)\left(\bar{\tau} - \frac{\varepsilon_x - 1}{\varepsilon_x}\eta \right) dF_\eta(\eta) > 0 \tag{42}$$

In particular, $0 < \bar{\tau} < (\varepsilon_x - 1)/\varepsilon_x$.

To calculate the derivatives at this steady state, take the first-order expansion of the labor demand condition. For log deviations from $(\bar{L}, \bar{\zeta}, \bar{w})$, the definitions of revenue and the financing threshold imply

$$\widehat{R} = \hat{\zeta} + \frac{\varepsilon_x - 1}{\varepsilon_x}\hat{l}, \qquad \widehat{\tau} = \hat{w} + \frac{\hat{l}}{\varepsilon_x} - \hat{\zeta}$$

The first-order expansion of (42) is

$$\hat{\zeta} + \frac{\varepsilon_x - 1}{\varepsilon_x}\hat{l} = -\left[\frac{\bar{\tau}}{(\varepsilon_x - 1)/\varepsilon_x - \bar{\tau}} + \frac{\bar{\tau}\int_0^{\bar{\tau}} \left[2\bar{\tau} - \left(1 + \frac{\varepsilon_x - 1}{\varepsilon_x} \right)\eta \right] dF_\eta(\eta)}{\int_0^{\bar{\tau}} (\bar{\tau} - \eta)\left(\bar{\tau} - \frac{\varepsilon_x - 1}{\varepsilon_x}\eta \right) dF_\eta(\eta)} \right] \left(\hat{w} + \frac{\hat{l}}{\varepsilon_x} - \hat{\zeta} \right) \tag{43}$$

The coefficient inside the square brackets exceeds one. Its first term is positive, and the numerator of its second term exceeds its denominator because

$$\bar{\tau}\int_0^{\bar{\tau}} \left[2\bar{\tau} - \left(1 + \frac{\varepsilon_x - 1}{\varepsilon_x} \right)\eta \right] dF_\eta(\eta) - \int_0^{\bar{\tau}} (\bar{\tau} - \eta)\left(\bar{\tau} - \frac{\varepsilon_x - 1}{\varepsilon_x}\eta \right) dF_\eta(\eta) = \int_0^{\bar{\tau}} \left(\bar{\tau}^2 - \frac{\varepsilon_x - 1}{\varepsilon_x}\eta^2 \right) dF_\eta(\eta) > 0$$

Notice equation (43) holds for any given $(\hat{\zeta}, \hat{w})$ to first order. Set $\hat{w} = 0$ in (43) and substitute $\hat{l} = \iota_1\hat{\zeta}$. The coefficient on $\hat{\zeta}$ on the left is $1 + (\varepsilon_x - 1)\iota_1/\varepsilon_x$, while the right-hand side is the positive bracket times $1 - \iota_1/\varepsilon_x$. If $\iota_1 \leq 0$, the left coefficient is at most one and the right coefficient exceeds one, a contradiction. If $\iota_1 \geq \varepsilon_x$, the left coefficient is positive and the right coefficient is nonpositive, also a contradiction. Thus $0 < \iota_1 < \varepsilon_x$.

To obtain the relation between the elasticities, set $\hat{\zeta} = (\varepsilon_x - 1)/\varepsilon_x$ and $\hat{w} = 1$ in the linearized condition. The value $\hat{l} = -1$ makes both $\widehat{R}$ and $\widehat{\tau}$ zero and hence solves (43). This solution is unique because, after collecting terms, the coefficient on $\hat{l}$ is $(\varepsilon_x - 1)/\varepsilon_x$ plus the strictly positive bracket divided by ε_x. Substitution into $\hat{l} = \iota_1\hat{\zeta} - \iota_2\hat{w}$ gives, after rearranging,

$$\iota_2 = 1 + \frac{\varepsilon_x - 1}{\varepsilon_x}\iota_1 > 1$$

Finally, substitution of the steady-state labor demand condition (41) into the definition of φ_ζ gives

$$\varphi_\zeta = \frac{\bar{R}}{\bar{Y}}\left(\frac{\varepsilon_x - 1}{\varepsilon_x}\iota_1 - 2\phi_0 I_2\right)$$

Plugging in the definition of φ_ω and using $\iota_2 = 1 + (\varepsilon_x - 1)\iota_1/\varepsilon_x$ that I derived earlier lead to

$$\varphi_\omega = -\frac{\varepsilon_x - 1}{\varepsilon_x}\varphi_\zeta$$

Since $\bar{R}/\bar{Y} = (\bar{\varepsilon}_y - 1)/\bar{\varepsilon}_y < 1$, we have

$$\iota_1 - \varphi_\zeta = \left(1 - \frac{\bar{R}}{\bar{Y}}\frac{\varepsilon_x - 1}{\varepsilon_x}\right)\iota_1 + 2\phi_0\frac{\bar{R}}{\bar{Y}}I_2 > 0$$

□

E An Analytical Benchmark of the Model

This section provides a complete analytical characterization of a simple case of the model. In particular, I assume i.i.d. aggregate shocks and one-period contracts, which break down the dynamic linkages in the model. These assumptions isolate the three stark differences of the current model compared to a standard linear DSGE model, namely, the nonlinearity of the contract distribution, the forecast problem, and the nonlinear evolution of the contract distribution. Standard forces, e.g., forward-looking behavior through the NKPC and the dynamic IS curve, are shut down. Specifically,

1. $\varrho_u = \varrho_\rho = 0$, and $\sigma_{ut} = \sigma_u, \sigma_{\rho t} = \sigma_\rho$ for all $t \in \mathbb{Z}$ where σ_u and σ_ρ are known scalars; contract length $N = 1$.

One immediate implication is that the contract distribution can be summarized by a single COLA share $\mathcal{M}_{t,1}$. To ease notation, I use m_t to replace it, and it is predetermined in $t - 1$. In what follows, I first study the contracting problem with explicit derivations of the forecasts. Then, I solve the conditional general equilibrium given the current COLA share m_t. Finally, I characterize how m_t is determined. The order of the results will be the same as those in the main text.

E.1 The Contracting Problem

We solve the problem backwards. Firm employment choice follows directly from previous results and is independent of the simplifying assumption.

Lemma 7. *(Firm labor input choice) Up to second order, given the realized real wage and the real revenue factor,*

$$\hat{l}_{it} = \iota_1 \hat{\zeta}_t - \iota_2 \hat{w}_{it}$$

where $\iota_1, \iota_2 > 0$ are steady-state elasticities, and for any variable x, $\hat{x}$ denotes its log deviation from steady state.

Proof. See Section D.2.1. □

From Section D.2.1, we know that under the simplifying assumption,

$$\hat{w}_t^{\mathcal{C}} = \mu \mathbb{E}\left[\hat{\zeta}_t \middle| \mathcal{F}_{t-1}, \pi_t\right]$$
$$\hat{w}_t^{\mathcal{N}} = \mu \mathbb{E}\left[\hat{\zeta}_t \middle| \mathcal{F}_{t-1}\right] + \mathbb{E}\left[\pi_t \middle| \mathcal{F}_{t-1}\right] - \pi_t$$

In order to get an explicit expression of these forecasts, we need to know the stochastic processes of π_t and $\hat{\zeta}_t$. But they are endogenous to the DSGE system, to which those wage formulas are inputs to solve the system. This creates a fixed-point problem. I use the guess-and-verify approach to solve the system. With an exogenously fixed contract distribution, this system reduces to a standard linear DSGE model, and we know that the rational expectations equilibrium (REE) solutions must look like

$$\hat{\zeta}_t^{exo} = \zeta_\rho \rho_t + \zeta_u u_t$$
$$\pi_t^{exo} = \pi_\rho \rho_t + \pi_u u_t$$

for some constants $\zeta_\rho, \zeta_u, \pi_\rho, \pi_u$. With the nonlinear contract choice block, we expect that the contract distribution will interact with structural aggregate shocks, and we guess the REE solution

$$\hat{\zeta}_t = \zeta_\rho(m_t)\rho_t + \zeta_u(m_t)u_t$$
$$\pi_t = \pi_\rho(m_t)\rho_t + \pi_u(m_t)u_t \tag{44}$$

In particular, we assume $\zeta_\rho(m_t), \pi_\rho(m_t), \pi_u(m_t) > 0$ and $\zeta_u(m_t) < 0$. Both the functional forms and the signs will be verified in the general equilibrium. Now, a firm-union pair takes them as given, which can already give us explicit wage rules.

Lemma 8. *(Realized real wage) Given the above shock process and an inflation realization π_t, the realized real wages implied by the two contracts are*

$$\hat{w}_t^{\mathcal{C}} = \delta_t \pi_t$$
$$\hat{w}_t^{\mathcal{N}} = -\pi_t$$

where under some regularity conditions,

$$\delta_t = \underbrace{\mu}_{\text{worker's risk exposure (from s.s.)}} \underbrace{\frac{\text{Cov}\left(\hat{\zeta}_t, \pi_t\right)}{\text{Var}(\pi_t)}}_{\text{size of revenue risk}} \in (-1, 0)$$

Proof. The key is to derive those forecasts. Given the guess (44) and i.i.d. Gaussian aggregate shocks, this is a standard linear-Gaussian framework with zero prior means. Therefore,

$$\begin{aligned}
\mathbb{E}\left[\hat{\zeta}_t \middle| \mathcal{F}_{t-1}, \pi_t\right] &= \mathbb{E}\left[\hat{\zeta}_t \middle| \pi_t\right] \\
&= \frac{\text{Cov}\left(\hat{\zeta}_t, \pi_t\right)}{\text{Var}(\pi_t)} \pi_t \\
&= \frac{\zeta_\rho(m_t)\pi_\rho(m_t)\sigma_\rho^2 + \zeta_u(m_t)\pi_u(m_t)\sigma_u^2}{\pi_\rho(m_t)^2\sigma_\rho^2 + \pi_u(m_t)^2\sigma_u^2} \pi_t
\end{aligned}$$

which is the linear projection of $\hat{\zeta}_t$ on π_t, and

$$\mathbb{E}\left[\hat{\zeta}_t \middle| \mathcal{F}_{t-1}\right] = \mathbb{E}\left[\pi_t \middle| \mathcal{F}_{t-1}\right] = 0$$

They imply

$$\hat{w}_t^{\mathcal{C}} = \mu \mathbb{E}\left[\hat{\zeta}_t \middle| \mathcal{F}_{t-1}, \pi_t\right] = \underbrace{\mu \frac{\zeta_\rho(m_t)\pi_\rho(m_t)\sigma_\rho^2 + \zeta_u(m_t)\pi_u(m_t)\sigma_u^2}{\pi_\rho(m_t)^2\sigma_\rho^2 + \pi_u(m_t)^2\sigma_u^2}}_{:=\delta_t} \pi_t$$

$$\hat{w}_t^{\mathcal{N}} = -\pi_t$$

The coefficient δ_t measures price-to-real wage pass-through under COLA contracts. Accordingly, $1 + \delta_t$ measures the price-to-nominal-wage pass-through. $\delta_t \in (-1, 0)$ simply says the nominal pass-through is between zero and one. The regularity conditions based on structural parameters and steady-state quantities will be stated in the next subsection. In partial equilibrium, this will be

the case if inflation is primarily driven by markup shocks, and the magnitude of markup shocks is not extremely large relative to demand shocks. To see this, notice given $\mu > 0$,

$$\begin{aligned}\delta_t < 0 &\iff \zeta_\rho(m_t)\pi_\rho(m_t)\sigma_\rho^2 + \zeta_u(m_t)\pi_u(m_t)\sigma_u^2 < 0 \\ &\iff \frac{\sigma_u^2}{\sigma_\rho^2}\frac{|\zeta_u(m_t)|/\zeta_\rho(m_t)}{\pi_\rho(m_t)/\pi_u(m_t)} > 1\end{aligned}$$

Assume the four coefficients are fixed in PE. Then this is more likely if σ_u is much larger than σ_ρ, i.e., supply shocks are more volatile than demand shocks. On the other hand, notice that if $|\zeta_u(m_t)| > \pi_u(m_t)$, and $\sigma_u^2 \to \infty$, δ_t will be smaller than -1. Hence, as long as the markup shock is not too extreme, COLA would prescribe a higher pass-through than non-contingent wages, which is negative one. □

Given the explicit expressions of the wage rules, we can now calculate the gain from COLA explicitly.

Lemma 9. *(Contract type choice) Given Assumption 1 and the guess $\zeta_\rho, \zeta_u, \pi_\rho, \pi_u$, the gain from COLA compared to non-contingent wages at the end of $t-1$ is*

$$\mathcal{V}_{t-1}(m_t^e) = C\left[1 + \delta_t\left(m_t^e\right)\right]^2 \mathrm{Var}(\pi_t; m_t^e)$$

where $m_t^e = \mathbb{E}_{t-1}[m_t]$ is an individual firm-worker pair's belief about the COLA share in period t. Therefore, the probability that the firm adopts COLA in $t-1$ is

$$\mathbb{P}[\mathbb{I}_{i,t-1}^{\mathcal{C}} = 1] = F_\chi\left(\mathcal{V}_{t-1}(m_t^e)\right)$$

Proof. Using the result of equation (24) from Section D.2.1, the gain from COLA under the simplifying assumption is

$$\mathcal{V}_{t-1}(m_t^e) = -\frac{1}{2}\bar{V}_{ww}\bar{w}^2\beta\mathbb{E}_{t-1}\left[\left(\hat{w}_t^{\mathcal{C}} - \mathbb{E}_{t-1}\left[\hat{w}_t^{\mathcal{C}}\right] + \left(\pi_t - \mathbb{E}_{t-1}\left[\pi_t\right]\right)\right)^2\right]$$

where there are slight changes of the time subscript because we now look at contract type choice at the end of period $t-1$. Using the wage rules derived above, we have

$$\begin{aligned}\mathcal{V}_{t-1}(m_t^e) &= -\frac{1}{2}\bar{V}_{ww}\bar{w}^2\beta\mathbb{E}_{t-1}\left[\left(\delta_t\left(m_t^e\right)\pi_t - \mathbb{E}_{t-1}\left[\delta_t\left(m_t^e\right)\pi_t\right] + \left(\pi_t - \mathbb{E}_{t-1}\left[\pi_t\right]\right)\right)^2\right] \\ &= -\frac{1}{2}\bar{V}_{ww}\bar{w}^2\beta\mathbb{E}_{t-1}\left[\left(\left[1 + \delta_t\left(m_t^e\right)\right]\left(\pi_t - \mathbb{E}_{t-1}\left[\pi_t\right]\right)\right)^2\right] \\ &= \underbrace{-\frac{1}{2}\bar{V}_{ww}\bar{w}^2\beta}_{:=C}\left[1 + \delta_t\left(m_t^e\right)\right]^2 \mathrm{Var}_{t-1}(\pi_t; m_t^e)\end{aligned}$$

where from the second to the last line, I use the fact that δ_t is predetermined at the end of period $t-1$, and it does not comove with variables in period t. Notice when the structural shocks are

i.i.d., past shocks contain no information about today, so we can drop the time subscript of the conditional expectation/variance. Finally, the choice rule given the gain and the COLA cost is straightforward. □

E.2 The Conditional General Equilibrium

I next solve the conditional general equilibrium for any given COLA share and study its properties.

E.2.1 Solving the Conditional GE by the Method of Undetermined Coefficients

Given the guess of the REE solutions (44), one can plug them into the first five equations of the system summarized in Section D.3 and solve for the coefficient functions by the method of undetermined coefficients.

Recall that the aggregate real-wage index is

$$\hat{\omega}_t = m_t \hat{w}_t^{\mathcal{C}}(\pi_t) + (1 - m_t)\hat{w}_t^{\mathcal{N}}(\pi_t) = [m_t \delta_t - (1 - m_t)]\, \pi_t := K(m_t, \delta_t)\pi_t$$

In equilibrium, the optimal COLA wage rule δ_t is itself a function of the aggregate COLA share. Therefore define the equilibrium inflation elasticity of the aggregate real wage by

$$\mathcal{K}(m_t) := K(m_t, \delta(m_t)) = \left(1 + \mu \frac{\zeta_\rho(m_t)\pi_\rho(m_t)\sigma_\rho^2 + \zeta_u(m_t)\pi_u(m_t)\sigma_u^2}{\pi_\rho(m_t)^2\sigma_\rho^2 + \pi_u(m_t)^2\sigma_u^2}\right) m_t - 1 \tag{45}$$

Using $\mathcal{K}$, one can reduce the system that characterizes the conditional GE, i.e., (47)-(51) in Section D.3, to two equations that contain only inflation π and the real revenue factor $\hat{\zeta}$:

$$\left(1 - \underbrace{\kappa \frac{\iota_2}{\varepsilon_x}}_{:=p_\omega} \mathcal{K}(m_t)\right) \pi - \underbrace{\kappa\left(1 - \frac{\iota_1}{\varepsilon_x}\right)}_{:=p_\zeta} \hat{\zeta} = u \tag{46}$$

$$-\left[\underbrace{\left((1 + \gamma_c \phi_y)\iota_2 + \varphi_\omega\right)}_{:=y_\omega} \mathcal{K}(m_t) - \underbrace{\frac{\bar{C}^{\text{ent}}}{\bar{Y}\gamma}}_{:=\gamma_c} \phi_\pi\right] \pi + \left[\underbrace{(1 + \gamma_c \phi_y)\iota_1 - \varphi_\zeta}_{:=y_\zeta}\right] \hat{\zeta} = \gamma_c \rho \tag{47}$$

where, to make the notation simpler and the algebra more transparent, I define a few auxiliary objects, including the aggregate-demand sensitivity γ_c, the two price-setting loadings p_ζ, p_ω from the price-setting block, and the two demand loadings y_ζ, y_ω from the demand block. I suppress the time subscript when no confusion arises and write m for a given current COLA share.

Equations (45)-(47) characterize the conditional GE for any given m. Also notice that the COLA share m affects π and $\hat{\zeta}$ only through $\mathcal{K}(m)$. In what follows, I will first solve for π and $\hat{\zeta}$ as

functions of shocks u and ρ for any given $\mathcal{K}$ from (46) and (47). Then, I plug the derived coefficient functions, as functions of $\mathcal{K}$, into (45), and solve the fixed-point function of $\mathcal{K}$.

For any given $\mathcal{K}$, (46) and (47) are linear, and their determinant is

$$\Delta(\mathcal{K}) := y_\zeta(1 - p_\omega \mathcal{K}) - p_\zeta(y_\omega \mathcal{K} - \gamma_c \phi_\pi)$$

Therefore, conditional on $\mathcal{K}$, when $\Delta(\mathcal{K}) \neq 0$, we have

$$\pi = \frac{y_\zeta}{\Delta(\mathcal{K})} u + \frac{\gamma_c p_\zeta}{\Delta(\mathcal{K})} \rho \tag{48}$$

$$\hat{\zeta} = \frac{y_\omega \mathcal{K} - \gamma_c \phi_\pi}{\Delta(\mathcal{K})} u + \frac{\gamma_c (1 - p_\omega \mathcal{K})}{\Delta(\mathcal{K})} \rho \tag{49}$$

Equations (48) and (49) allow me to calculate the variance of π and the covariance between π and $\hat{\zeta}$ conditional on $\mathcal{K}$, which are needed for the optimal COLA wage rule:

$$\mathrm{Var}(\pi \mid \mathcal{K}) = \frac{(\gamma_c p_\zeta \sigma_\rho)^2 + (y_\zeta \sigma_u)^2}{\Delta(\mathcal{K})^2} := \frac{\mathcal{S}_\pi}{\Delta(\mathcal{K})^2}$$

$$\begin{aligned}
\mathrm{Cov}(\hat{\zeta}, \pi \mid \mathcal{K}) &= \frac{\gamma_c^2 p_\zeta (1 - p_\omega \mathcal{K}) \sigma_\rho^2 + y_\zeta (y_\omega \mathcal{K} - \gamma_c \phi_\pi) \sigma_u^2}{\Delta(\mathcal{K})^2} \\
&= \frac{\left(\gamma_c^2 p_\zeta \sigma_\rho^2 - \gamma_c \phi_\pi y_\zeta \sigma_u^2\right) + \left(-\gamma_c^2 p_\zeta p_\omega \sigma_\rho^2 + y_\zeta y_\omega \sigma_u^2\right) \mathcal{K}}{\Delta(\mathcal{K})^2} \\
&:= \frac{\mathcal{C}_0 + \mathcal{C}_1 \mathcal{K}}{\Delta(\mathcal{K})^2}
\end{aligned}$$

where I define coefficients $\mathcal{S}_\pi, \mathcal{C}_0, \mathcal{C}_1$ to make the structure clearer. One can think of $\mathcal{S}_\pi$ as relating to the variance of π, and $\mathcal{C}_0, \mathcal{C}_1$ as relating to the covariance of π and $\hat{\zeta}$. Plugging this back into equation (45), one can solve for $\mathcal{K}$ in the conditional general equilibrium conditional on m:

$$\begin{aligned}
\mathcal{K} &= \left(1 + \mu \frac{\mathcal{C}_0 + \mathcal{C}_1 \mathcal{K}}{\mathcal{S}_\pi}\right) m - 1 \\
\implies \mathcal{K}(m) &= \frac{\mu m \mathcal{C}_0 - \mathcal{S}_\pi (1 - m)}{\mathcal{S}_\pi - \mu m \mathcal{C}_1}
\end{aligned} \tag{50}$$

given $\mathcal{S}_\pi - \mu m \mathcal{C}_1 \neq 0$ for all $m \in [0, 1]$.

Remarkably, although equation (45) appears to be highly nonlinear because of the inference problem, the fixed-point equation of $\mathcal{K}$, (50), reduces to a linear equation because the determinant terms in the numerator and the denominator cancel out. This property greatly simplifies the analysis below.

In the conditional GE, the determinant, as a function of m, is

$$\Delta(\mathcal{K}(m)) = \frac{\mathcal{S}_\pi}{\mathcal{S}_\pi - \mu m \mathcal{C}_1}\left[y_\zeta + \gamma_c \phi_\pi p_\zeta + (1-m)(y_\zeta p_\omega + p_\zeta y_\omega) - \mu m(y_\omega - p_\omega \gamma_c \phi_\pi)\right] \tag{51}$$

$$:= \frac{\mathcal{S}_\pi}{\mathcal{S}_\pi - \mu m \mathcal{C}_1}\mathcal{D}(m) \tag{52}$$

where I define $\mathcal{D}(m)$, and I sometimes refer to it as the scaled (equilibrium) determinant for convenience.

Before summarizing the above results, I impose the following regularity condition, which ensures the existence of a well-defined conditional GE for any $m \in [0,1]$.

Assumption 3. *(Conditional GE regularity)* $\mathcal{S}_\pi - \mu\mathcal{C}_1 > 0$ *and* $\mathcal{D}(1) > 0$.

The next lemma summarizes the solution of the conditional GE.

Lemma 10. *(Closed-form conditional GE) Under Assumption 3, for every predetermined* $m \in [0,1]$ *there is a unique conditional general equilibrium. It is given by*

$$\pi = \pi_u(\mathcal{K}(m))u + \pi_\rho(\mathcal{K}(m))\rho$$
$$\hat{\zeta} = \zeta_u(\mathcal{K}(m))u + \zeta_\rho(\mathcal{K}(m))\rho$$

where $\mathcal{K}(m)$ *is given in equation (50), and the coefficient functions are*

$$\pi_u(\mathcal{K}) = \frac{y_\zeta}{\Delta(\mathcal{K})}, \qquad \zeta_u(\mathcal{K}) = \frac{y_\omega \mathcal{K} - \gamma_c \phi_\pi}{\Delta(\mathcal{K})}$$
$$\pi_\rho(\mathcal{K}) = \frac{\gamma_c p_\zeta}{\Delta(\mathcal{K})}, \qquad \zeta_\rho(\mathcal{K}) = \frac{\gamma_c(1 - p_\omega \mathcal{K})}{\Delta(\mathcal{K})}$$

Moreover, for all $m \in [0,1]$

$$\pi_u(m) > 0, \qquad \pi_\rho(m) > 0$$

If one additionally assumes $\sigma_u^2 > \frac{\gamma_c p_\zeta \sigma_\rho^2}{y_\zeta \phi_\pi}$*, then, for all* $m \in [0,1]$

$$\zeta_u(m) < 0, \qquad \zeta_\rho(m) > 0$$

Proof. Equations (46)-(47) give the coefficient functions in the conditional inflation and revenue processes given $\mathcal{K}$. Substituting those coefficient functions into the optimal COLA wage rule leads to the linear fixed-point equation (45). To make sure they are well-defined, notice $\mathcal{S}_\pi - \mu m \mathcal{C}_1$ is affine in m. Its endpoint values are $\mathcal{S}_\pi > 0$ and $\mathcal{S}_\pi - \mu\mathcal{C}_1 > 0$ from Assumption 3, so it is positive for all $m \in [0,1]$. Similarly, $\mathcal{D}$ is affine, and

$$\mathcal{D}(0) = y_\zeta + \gamma_c \phi_\pi p_\zeta + y_\zeta p_\omega + p_\zeta y_\omega > 0$$

Together with $\mathcal{D}(1) > 0$ from Assumption 3, $\mathcal{D}(m) > 0$ for all $m \in [0,1]$. Thus, the determinant $\Delta(\mathcal{K}(m)) > 0$, so the two-equation system is nonsingular, and the solution is unique.

To show that the two coefficient functions in the conditional inflation process are positive, it remains to show $y_\zeta > 0$ and $\gamma_c p_\zeta > 0$. Notice $\gamma_c > 0$ for any interior steady state. Since

$$y_\zeta = (\iota_1 - \varphi_\zeta) + \gamma_c \phi_y \iota_1$$
$$y_\omega = 1 + \gamma_c \phi_y + \frac{\varepsilon_x - 1}{\varepsilon_x} y_\zeta$$

Lemma 6 shows that $\iota_1 > \varphi_\zeta$. This implies $y_\zeta > 0$ and $y_\omega > 0$. Combining with $\Delta(\mathcal{K}(m)) > 0$ that I derived earlier, we have $\pi_u(m) > 0$ and $\pi_\rho(m) > 0$ for all $m \in [0,1]$.

Finally, the condition $\sigma_u^2 > \frac{\gamma_c p_\zeta \sigma_\rho^2}{y_\zeta \phi_\pi}$ implies $\mathcal{C}_0 = \gamma_c^2 p_\zeta \sigma_\rho^2 - \gamma_c \phi_\pi y_\zeta \sigma_u^2 < 0$. Then

$$\mathcal{K}(1) = \frac{\mu \mathcal{C}_0}{\mathcal{S}_\pi - \mu \mathcal{C}_1} < 0$$

Combining it with $\mathcal{K}(0) = -1 < 0$ and that $\mathcal{K}(m)$ is affine in m implies $\mathcal{K}(m) < 0$ for $m \in [0,1]$. Using Lemma 6 again, we have

$$p_\zeta = \kappa \left(1 - \frac{\iota_1}{\varepsilon_x}\right) > 0$$
$$p_\omega = \kappa \frac{\iota_2}{\varepsilon_x} > 0$$

Thus, $y_\omega \mathcal{K}(m) - \gamma_c \phi_\pi < 0$, and $\gamma_c (1 - p_\omega \mathcal{K}(m)) > 0$. They together imply that $\zeta_u(m) < 0$ and $\zeta_\rho(m) > 0$ for all $m \in [0,1]$. □

Therefore, the same regularity conditions that guarantee a well-defined conditional GE also imply that both a positive markup shock and a positive demand shock raise inflation. I will maintain Assumption 3 throughout this section without restating it for brevity. However, it does not automatically guarantee $\zeta_u(m) < 0$ and $\zeta_\rho(m) > 0$. The additional assumption basically requires that for any given monetary policy response to inflation ϕ_π, markup shocks should be sufficiently volatile to drive inflation dynamics. In that case, given an inflation spike, policy would be strong enough to suppress aggregate demand and therefore hurt firm revenue. This additional assumption is not necessary for the subsequent results. But to make those analyses more economically interpretable, I also maintain it throughout this section without restating it.

For later use, I define the nominal pass-through of a COLA contract as

$$\Pi(m) := 1 + \delta(m)$$

Recall the definition of $\mathcal{K}(m)$ is

$$\mathcal{K}(m) = m\delta(m) - (1 - m)$$

Combining the two leads to the equilibrium COLA nominal pass-through conditional on m

$$\Pi(m) = \frac{1 + \mathcal{K}(m)}{m} = \frac{\mathcal{S}_\pi + \mu(\mathcal{C}_0 - \mathcal{C}_1)}{\mathcal{S}_\pi - \mu m \mathcal{C}_1} \tag{53}$$

E.2.2 Properties of the Conditional GE

Given the existence of a well-behaved conditional GE for any $m \in [0, 1]$ that I established in the last subsection, I now study other properties of the conditional GE. In the main text, I state all the aggregate implications based on the following two conditions:

1. (Positive nominal COLA pass-through) For every $m \in [0, 1]$, $1 + \delta(m) > 0$.
2. (Positive slope of COLA pass-through) For every $m \in [0, 1]$, $\delta_m(m) \geq 0$.

I show that these two conditions are equivalent to the following assumptions on structural parameters and parameters from the steady state.

Assumption 4. *(Pass-through restriction) Assume (i) $\mathcal{S}_\pi + \mu(\mathcal{C}_0 - \mathcal{C}_1) > 0$ and (ii) $\mathcal{C}_1 \geq 0$.*

Lemma 11. *(Equivalence) Part (i) of Assumption 4 is equivalent to Condition 1. Given part (i), part (ii) is equivalent to Condition 2.*

Proof. The denominator in (53) is strictly positive, so $\Pi(m) > 0$ for every m if and only if $\mathcal{S}_\pi + \mu(\mathcal{C}_0 - \mathcal{C}_1) > 0$. Differentiating (53) gives

$$\Pi_m(m) = \delta_m(m) = \frac{\mu \mathcal{C}_1 \left[\mathcal{S}_\pi + \mu(\mathcal{C}_0 - \mathcal{C}_1)\right]}{(\mathcal{S}_\pi - \mu m \mathcal{C}_1)^2}$$

Given $\mu > 0$ and $\mathcal{S}_\pi + \mu(\mathcal{C}_0 - \mathcal{C}_1) > 0$, this derivative is nonnegative if and only if $\mathcal{C}_1 \geq 0$. □

To understand the condition for positive pass-through, notice that

$$\mathcal{S}_\pi + \mu(\mathcal{C}_0 - \mathcal{C}_1) > 0 \iff \mathcal{S}_\pi + \mu(A_1 \sigma_\rho^2 - A_2 \sigma_u^2) > 0$$

for some $A_1, A_2 > 0$. $\mathcal{S}_\pi > 0$ by definition, so this condition may fail only if μ and σ_u are relatively large. Intuitively, that means (1) workers have large exposure to the revenue risk, which happens if workers' risk aversion is small relative to the firm's financing cost; and (2) markup shocks are very volatile and drive the dynamics of inflation and revenue. In that case, conditional on a COLA, high inflation can lead to a nominal wage cut, meaning workers are actually compensating the firm.

As for the condition that $\delta_m(m) \geq 0$, notice that

$$\mathcal{C}_1 \geq 0 \iff y_\zeta y_\omega \sigma_u^2 - \gamma_c^2 p_\zeta p_\omega \sigma_\rho^2 \geq 0$$

As I will show shortly, a higher COLA share m raises inflation volatility. This dilutes the information about the revenue factor contained in the same amount of inflation. The part of inflation

driven by demand shocks would imply a lower revenue factor, pushing the COLA pass-through downwards, while the part of inflation driven by markup shocks would imply a higher revenue factor, pushing the COLA pass-through upwards. Equilibrium COLA pass-through increases with the COLA share if the second force dominates.

Given Lemma 10, we can derive the following comparative statics in the conditional general equilibrium.

Lemma 12. *(Conditional-GE amplification) Under Condition 1, when the COLA share m rises, the aggregate real wage is less sensitive to inflation surprises, i.e.,*

$$\mathcal{K}_m(m) > 0$$

and inflation in the conditional GE responds more strongly to both structural shocks, i.e.,

$$\pi_{u,m}(m) > 0, \qquad \pi_{\rho,m}(m) > 0$$

Proof. Differentiating $\mathcal{K}(m) = m\Pi(m) - 1$ with respect to m, we have

$$\mathcal{K}_m(m) = \Pi(m) + m\Pi_m(m) \tag{54}$$

Condition 1 directly implies that the first term is positive, i.e., more COLA with positive pass-through makes the aggregate real wage more rigid. To determine the sign of the whole term, using the expression of $\Pi(m)$ in the conditional GE in (53), we have

$$\mathcal{K}_m(m) = \frac{\mathcal{S}_\pi[\mathcal{S}_\pi + \mu(\mathcal{C}_0 - \mathcal{C}_1)]}{(\mathcal{S}_\pi - \mu m\mathcal{C}_1)^2} = \frac{\mathcal{S}_\pi}{\mathcal{S}_\pi - \mu m\mathcal{C}_1}\Pi > 0$$

under Condition 1.

For the two coefficient functions in the conditional inflation process, $\pi_u(\mathcal{K}) = \frac{y_\zeta}{\Delta(\mathcal{K})}$, $\pi_\rho(\mathcal{K}) = \frac{\gamma_c p_\zeta}{\Delta(\mathcal{K})}$,notice

$$\Delta_\mathcal{K}(\mathcal{K}) = -(y_\zeta p_\omega + p_\zeta y_\omega) < 0$$

Since both coefficients have a positive numerator that is independent of $\mathcal{K}$, the signs of their derivatives with respect to m are the same as the sign of

$$-\frac{1}{\Delta(\mathcal{K}(m))^2}\Delta_\mathcal{K}(\mathcal{K}(m))\mathcal{K}_m(m)$$

which is positive. □

The proof shows that when the COLA share m rises, the first term $\Pi(m)$ in equation (54) captures the mechanical increase in the aggregate real rigidity, because more wages are indexed with positive nominal pass-through given Condition 1. However, the optimal COLA pass-through will also change endogenously as the joint distribution of inflation and the revenue factor changes.

That effect is captured by $m\Pi_m(m)$. This lemma implies that Condition 1 is sufficient for the first force to dominate the total response,[46] so more indexation makes the aggregate real wage more rigid. This consequently raises the inflation response to aggregate shocks in the conditional GE.

With the above results, the conditional general equilibrium can be written in the three-equation representation highlighted in the main text.

Proposition 10. *(Three-equation representation) Under Condition 1, for any predetermined $m \in [0,1]$, the conditional general equilibrium has the representation*

$$\begin{aligned} NKPC: \quad & \pi_t = \frac{\kappa \frac{\varepsilon_x - \iota_1}{\iota_1 \varepsilon_x} \hat{y}_t + u_t}{1 - \kappa \frac{\iota_2}{\iota_1} \mathcal{K}(m)} \\ DIS: \quad & \hat{y}_t = \frac{\iota_1}{\iota_1 - \varphi_\zeta} \left[-\frac{\bar{C}^{\text{ent}}}{\bar{Y}} \gamma (i_t - \rho_t) + \left(\varphi_\omega + \frac{\iota_2}{\iota_1} \varphi_\zeta \right) \mathcal{K}(m) \pi_t \right] \\ MP: \quad & i_t = \phi_\pi \pi_t + \phi_y \hat{y}_t \end{aligned}$$

where $\mathcal{K}(m) \in [-1, 0)$ and is increasing in m.

Proof. Equations (47)-(51) in Lemma 5 characterize the conditional GE in the full model. Set $N = 1$ and $\varrho_u = \varrho_\rho = 0$. First use the production function to eliminate $\hat{\zeta}_t$ from the NKPC and the resource constraint. Then substitute $\hat{\omega}_t$ with $\mathcal{K}(m)\pi_t$. After rearranging, one obtains the three equations.

The proof of Lemma 10 already implies $\mathcal{K}(m) < 0$. Since $\mathcal{K}(0) = -1$ and $\mathcal{K}(m)$ is affine in m, it lies between -1 and 0 for $m \in [0,1]$. Lemma 12 shows $\mathcal{K}(m)$ is increasing in m. □

One remark is that $\iota_1 - \varphi_\zeta > 0$ from Lemma 6, so the dynamic IS curve is upward sloped and thus well-behaved.

Corollary 2. *(Inflation volatility) Under Condition 1, inflation volatility increases with the COLA share:*

$$\frac{\partial}{\partial m} \operatorname{Var}(\pi_t \mid m) > 0$$

Proof. The result follows directly from

$$\operatorname{Var}(\pi_t \mid m) = \pi_u(m)^2 \sigma_u^2 + \pi_\rho(m)^2 \sigma_\rho^2$$

and Lemma 12, as both $\pi_u(m)$ and $\pi_\rho(m)$ are positive and increasing in m. □

E.3 Strategic Complementarity and the Contracting General Equilibrium

Recall from Lemma 9 that the gain from COLA is

$$\mathcal{V}(m) = C\Pi(m)^2 \operatorname{Var}(\pi \mid m) \tag{55}$$

[46] Condition 1 does not sign $\Pi_m(m)$. Under Condition 2, it is nonnegative.

where $C = -\bar{V}_{ww}\bar{w}^2\beta/2 > 0$ is a constant from the steady state. It has a particularly convenient expression in the conditional GE.

Lemma 13. *(COLA gain) In the conditional GE,*

$$\mathcal{V}(m) = C\frac{[\mathcal{S}_\pi + \mu(\mathcal{C}_0 - \mathcal{C}_1)]^2}{\mathcal{S}_\pi \mathcal{D}(m)^2} \tag{56}$$

Proof. Notice

$$\text{Var}(\pi \mid m) = \frac{\mathcal{S}_\pi}{\Delta(\mathcal{K}(m))^2} = \frac{(\mathcal{S}_\pi - \mu m \mathcal{C}_1)^2}{\mathcal{S}_\pi \mathcal{D}(m)^2}$$

Multiplying this expression by

$$C\Pi(m)^2 = C\frac{[\mathcal{S}_\pi + \mu(\mathcal{C}_0 - \mathcal{C}_1)]^2}{(\mathcal{S}_\pi - \mu m \mathcal{C}_1)^2}$$

leads to the stated result. □

Proposition 11. *(Strategic complementarity) Under Conditions 1–2, the gain from COLA is increasing in the expected aggregate COLA share:*

$$\mathcal{V}_m(m) > 0$$

Proof. Log differentiating equation (55) with respect to m, we have

$$\frac{\mathcal{V}_m}{\mathcal{V}} = \frac{\partial \log \text{Var}(\pi \mid m)}{\partial m} + 2\frac{\Pi_m}{\Pi}$$

Corollary 2 shows that

$$\frac{\partial \log \text{Var}(\pi \mid m)}{\partial m} > 0$$

while Condition 2 directly implies

$$\Pi_m = \delta_m \geq 0$$

Hence, $\mathcal{V}_m(m) > 0$ as stated. □

The first term captures the effect from higher inflation risk and is positive as long as COLA pass-through is positive: when more wage contracts are indexed, the aggregate real wage is more rigid, so inflation becomes more responsive to both shocks. The second term captures the effect from changes in the optimal COLA pass-through. Condition 2 requires this term to be non-negative. This simplifies the proof, although it is not necessary. Even if it declines, as long as the inflation-risk force dominates, the gain from COLA still rises with the COLA share m.[47]

[47] However, for empirically plausible parameter values, this term is positive. So both the individual pass-through force and the inflation volatility force contribute to a higher gain from COLA.

E.3.1 Contracting Equilibrium

Recall F_χ is the CDF of COLA adoption costs. Unless otherwise stated, I retain the regularity condition in the main text: F_χ is twice continuously differentiable with bounded density f_χ and support $[0, \infty)$. Given an expected COLA share m_t^e, the share of firm-union pairs that choose COLA, namely, the aggregate best response, is

$$\mathrm{BR}(m_t^e) := F_\chi(\mathcal{V}(m_t^e))$$

A contracting equilibrium is a fixed point of this aggregate best response.

Proposition 12. *(Contracting GE) Under Condition 1, a contracting general equilibrium exists. Any equilibrium COLA share solves*

$$m = F_\chi\left(C\frac{[\mathcal{S}_\pi + \mu(\mathcal{C}_0 - \mathcal{C}_1)]^2}{\mathcal{S}_\pi \mathcal{D}(m)^2}\right) \tag{57}$$

If Condition 2 also holds, the aggregate best response is nondecreasing:

$$\mathrm{BR}_m(m) = f_\chi(\mathcal{V}(m))\mathcal{V}_m(m) \geq 0$$

Proof. Continuity of F_χ and $\mathcal{V}$ implies that BR is a continuous map from $[0, 1]$ into itself. Notice

$$\mathrm{BR}(0) - 0 \geq 0, \qquad \mathrm{BR}(1) - 1 \leq 0$$

The intermediate value theorem then implies that there is at least one fixed point. Equation (57) follows from the expression of the COLA gain in the conditional GE in equation (56). Under Condition 2, $\mathcal{V}_m > 0$ from Proposition 11, So the aggregate best-response function is also increasing in m. □

Proposition 12 establishes the existence of at least one contracting general equilibrium. However, it need not be unique, especially in light of the strategic complementarity in contract choice. The next proposition gives a clear example.

Proposition 13. *(Uncertainty and the three regimes with homogeneous COLA costs) Suppose every firm-union pair has the same COLA cost $\chi_0 > 0$, Conditions 1–2 hold, and*

$$\sigma_u = \sigma, \qquad \sigma_\rho = k\sigma$$

for some constant $k > 0$. There exist finite thresholds $0 < \underline{\sigma} < \overline{\sigma} < \infty$ such that the economy has a unique no-COLA equilibrium for $\sigma < \underline{\sigma}$, exactly three equilibrium shares for $\sigma \in (\underline{\sigma}, \overline{\sigma})$, and a unique all-COLA equilibrium for $\sigma > \overline{\sigma}$.

Proof. Substituting $\sigma_u = \sigma$ and $\sigma_\rho = k\sigma$ into the definitions of $\mathcal{S}_\pi$, $\mathcal{C}_0$, and $\mathcal{C}_1$, we have

$$\begin{aligned}\mathcal{S}_\pi &= [(\gamma_c p_\zeta k)^2 + y_\zeta^2]\sigma^2 \\ \mathcal{C}_0 &= (\gamma_c^2 p_\zeta k^2 - \gamma_c \phi_\pi y_\zeta)\sigma^2 \\ \mathcal{C}_1 &= (-\gamma_c^2 p_\zeta p_\omega k^2 + y_\zeta y_\omega)\sigma^2\end{aligned}$$

They are all proportional to σ^2, while $\mathcal{D}(m)$ does not depend on σ. Plug these expressions into the COLA gain in equation (56), so its numerator is proportional to σ^4, and its denominator is proportional to σ^2. To make this dependence explicit, I write the gain as $\mathcal{V}(m;\sigma)$. The analysis implies that $\mathcal{V}(m;\sigma)$ is homothetic to σ^2, i.e.,

$$\mathcal{V}(m;\sigma) = \sigma^2 \mathcal{V}(m;1)$$

Since $\mathcal{V}(m;1) > 0$, for any given m, the gain is strictly increasing in $\sigma > 0$, and

$$\lim_{\sigma\to 0} \mathcal{V}(m;\sigma) = 0, \qquad \lim_{\sigma\to\infty} \mathcal{V}(m;\sigma) = \infty$$

Given Conditions 1–2, Proposition 11 shows that the gain is also strictly increasing in m. In particular, $0 < \mathcal{V}(0;1) < \mathcal{V}(1;1)$. Define

$$\underline{\sigma} = \sqrt{\frac{\chi_0}{\mathcal{V}(1;1)}}, \qquad \overline{\sigma} = \sqrt{\frac{\chi_0}{\mathcal{V}(0;1)}}$$

Then $0 < \underline{\sigma} < \overline{\sigma} < \infty$, and

$$\mathcal{V}(1;\underline{\sigma}) = \mathcal{V}(0;\overline{\sigma}) = \chi_0$$

Given the common cost χ_0, the aggregate best response is

$$\mathrm{BR}(m) = \begin{cases} \{0\}, & \mathcal{V}(m;\sigma) < \chi_0 \\ [0,1], & \mathcal{V}(m;\sigma) = \chi_0 \\ \{1\}, & \mathcal{V}(m;\sigma) > \chi_0 \end{cases}$$

If $\sigma < \underline{\sigma}$, then $\mathcal{V}(m;\sigma) \le \mathcal{V}(1;\sigma) < \chi_0$ for all $m \in [0,1]$. Thus no pair chooses COLA, and $m = 0$ is the unique equilibrium. If $\sigma > \overline{\sigma}$, then $\mathcal{V}(m;\sigma) \ge \mathcal{V}(0;\sigma) > \chi_0$ for all $m \in [0,1]$. Thus all pairs choose COLA, and $m = 1$ is the unique equilibrium.

If $\sigma \in (\underline{\sigma}, \overline{\sigma})$, then

$$\mathcal{V}(0;\sigma) < \chi_0 < \mathcal{V}(1;\sigma)$$

The first inequality implies that choosing no COLA is optimal when $m = 0$, so $m = 0$ is an equilibrium. The second inequality implies that choosing COLA is optimal when $m = 1$, so $m = 1$ is also an equilibrium. Since $\mathcal{V}(m;\sigma)$ is continuous and strictly increasing in m, there is a unique $m^{int} \in (0,1)$ such that $\mathcal{V}(m^{int};\sigma) = \chi_0$. At this share, all pairs are indifferent between the two

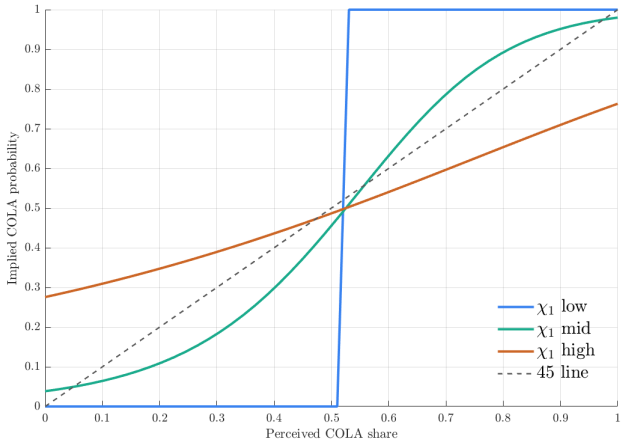


Figure 28: The contracting general equilibrium under different COLA cost structures

contracts. A fraction m^{int} of pairs can choose COLA, so m^{int} is an equilibrium. No other interior share can be an equilibrium. This gives exactly three equilibrium shares. At either threshold, the interior share coincides with one of the two endpoints, and both $m = 0$ and $m = 1$ are equilibria. □

The Role of Random COLA Costs. The homogeneous-cost case makes the source of multiplicity transparent, but its interior equilibrium requires mixing strategies. With heterogeneous costs, the CDF smooths the aggregate best response and interior COLA shares arise in pure equilibria. The strength of the aggregate strategic interaction is

$$\mathrm{BR}_m = f_\chi(\mathcal{V})\mathcal{V}_m$$

Thus, even when the private gain exhibits strong complementarity, a dispersed cost distribution can flatten the aggregate best response because few pairs are close to the adoption margin. Intuitively, if costs are dispersed, private costs become more important in determining individual adoption decisions relative to coordination.

Figure 28 illustrates this point using a logistic distribution for COLA costs (negative costs are bunched to zero). The three cases have the same location parameter but different scale parameters. When costs are concentrated, the best response resembles the correspondence in the homogeneous-cost case and can support low- and high-COLA regimes. As cost dispersion rises, the best response becomes flatter, the extreme equilibria move inward, and multiplicity eventually disappears.

For later use, I define locally stable equilibrium COLA shares.

Definition 2. A fixed point m^* is locally stable if

$$\mathrm{BR}_m(m^*) < 1$$

When the best-response function crosses the 45-degree line from above, its slope is below one and the equilibrium is locally stable. For example, if there are three regular intersections and the best response is increasing, the low and high intersections are stable and the middle intersection is unstable whenever the crossings alternate in the usual way. But the number of fixed points is not necessarily one or three: a sufficiently irregular cost density may generate more than three intersections. Notice when there is a unique equilibrium COLA share m^*, the boundary values derived in Proposition 12 imply that it must be a stable fixed point because it crosses the 45° line from above.

E.4 The Role of Monetary Policy

A more hawkish monetary policy rule affects contract choice in two steps. Holding the contract distribution fixed, it changes inflation risk and the optimal pass-through within COLA contracts. In the contracting general equilibrium, the resulting change in the private COLA gain feeds back into the aggregate contract distribution. I first characterize the conditional effect and then turn to the equilibrium response.

Lemma 14. *(Monetary policy in the conditional GE) Under Condition 1, for every predetermined $m \in [0,1]$ and for local policy changes that remain in the admissible parameter region, a stronger policy response to inflation lowers the aggregate real wage sensitivity to inflation surprises and the optimal COLA pass-through:*

$$\mathcal{K}_{\phi_\pi}(m) \leq 0, \qquad \Pi_{\phi_\pi}(m) < 0$$

where the first inequality is strict if $m > 0$. Moreover, it stabilizes the conditional inflation process and ultimately lowers the gain from COLA:

$$\pi_{u,\phi_\pi}(m) < 0, \qquad \pi_{\rho,\phi_\pi}(m) < 0, \qquad \mathcal{V}_{\phi_\pi}(m) < 0$$

Proof. Log differentiating $\mathcal{V}(m) = C\Pi(m)^2 \operatorname{Var}(\pi \mid m)$ with respect to ϕ_π, we have

$$\frac{\mathcal{V}_{\phi_\pi}}{\mathcal{V}} = \frac{\partial \log \operatorname{Var}(\pi \mid m)}{\partial \phi_\pi} + 2\frac{\Pi_{\phi_\pi}}{\Pi}$$

The first term captures the effect from changes in inflation risk. The second term captures the effect from changes in the optimal COLA pass-through. I show that both terms are negative in steps.

First, notice that $\mathcal{S}_\pi$ and $\mathcal{C}_1$ do not depend on ϕ_π, while

$$\mathcal{C}_{0,\phi_\pi} = -\gamma_c y_\zeta \sigma_u^2 < 0$$

Differentiating the optimal pass-through in equation (53), we have

$$\Pi_{\phi_\pi}(m) = -\frac{\mu\gamma_c y_\zeta \sigma_u^2}{\mathcal{S}_\pi - \mu m \mathcal{C}_1} < 0$$

Since $\mathcal{K}(m) = m\Pi(m) - 1$, we also have

$$\mathcal{K}_{\phi_\pi}(m) = m\Pi_{\phi_\pi}(m) \leq 0$$

with strict inequality when $m > 0$. Thus, holding the COLA share fixed, a more hawkish policy rule makes the aggregate real wage more sensitive to inflation surprises.

Next, recall from the conditional GE that

$$\pi_u(m) = \frac{y_\zeta}{\Delta(\mathcal{K}(m))}, \qquad \pi_\rho(m) = \frac{\gamma_c p_\zeta}{\Delta(\mathcal{K}(m))}$$

Their numerators do not depend on ϕ_π. Differentiating their common denominator, including the change in $\mathcal{K}(m)$, we have

$$\frac{\partial}{\partial \phi_\pi}\left(\frac{1}{\Delta(\mathcal{K}(m))}\right) = -\frac{1}{[\Delta(\mathcal{K}(m))]^2}\left[\gamma_c p_\zeta - (y_\zeta p_\omega + p_\zeta y_\omega) m \Pi_{\phi_\pi}(m)\right] < 0$$

The first term captures the direct stabilizing effect of monetary policy. The second term captures the additional stabilization from the lower COLA pass-through. Hence, both $\pi_{u,\phi_\pi}(m) < 0$ and $\pi_{\rho,\phi_\pi}(m) < 0$.

This directly implies a higher ϕ_π lowers inflation risk because

$$\text{Var}(\pi \mid m) = \pi_u(m)^2\sigma_u^2 + \pi_\rho(m)^2\sigma_\rho^2$$

where both coefficient functions are positive and decreasing with ϕ_π. Hence, a more hawkish policy rule lowers both inflation risk and the optimal nominal pass-through, so it also lowers the gain from COLA, as stated. □

Fixing the aggregate COLA share, the direct effect of a stronger policy response to inflation through the aggregate demand channel is standard. This also triggers a response in the optimal COLA pass-through, which may feed back to aggregate demand because COLA pass-through determines the wage payments those hand-to-mouth workers receive and consume. To understand the effect of policy on the optimal COLA pass-through, notice that for the same amount of inflation driven by demand shocks, a stronger policy response does not change the inferred revenue, while for the same amount of inflation driven by markup shocks, a stronger policy response implies less revenue. This is captured by $\mathcal{C}_{0,\phi_\pi} = -\gamma_c y_\zeta \sigma_u^2 < 0$ in the proof, and thus optimal pass-through

declines with ϕ_π. In the main text, I attribute both effects to the conventional aggregate demand channel, as the COLA share m is fixed.[48]

I next study how monetary policy changes the equilibrium COLA share locally. I focus on locally stable COLA equilibria.

Proposition 14. *(Policy and stable contracting equilibria) Under Conditions 1–2, let m^* be a locally stable interior contracting equilibrium. Then*

$$\frac{dm^*}{d\phi_\pi} \leq 0$$

The inequality is strict if and only if $f_\chi(\mathcal{V}(m^*)) > 0$.

Proof. Differentiating the equilibrium condition $m^* = F_\chi(\mathcal{V}(m^*, \phi_\pi))$ with respect to ϕ_π, we have

$$\frac{dm^*}{d\phi_\pi} = f_\chi(\mathcal{V}(m^*)) \left[\mathcal{V}_{\phi_\pi}(m^*) + \mathcal{V}_m(m^*)\frac{dm^*}{d\phi_\pi}\right]$$

Solving for $dm^*/d\phi_\pi$, we have

$$\frac{dm^*}{d\phi_\pi} = \frac{f_\chi(\mathcal{V}(m^*))\mathcal{V}_{\phi_\pi}(m^*)}{1 - \mathrm{BR}_m(m^*)}$$

Lemma 14 shows that $\mathcal{V}_{\phi_\pi} < 0$, while local stability implies $1 - \mathrm{BR}_m(m^*) > 0$. This leads to the stated inequality. The inequality is strict if and only if $f_\chi(\mathcal{V}(m^*)) > 0$. □

As a remark on the stability requirement, I mentioned earlier that when there is a unique equilibrium COLA share, it must be stable. Therefore, a more hawkish monetary policy rule will reduce the equilibrium COLA share. When there are three equilibria, a downward shift in the best response moves the stable low and high equilibria toward lower COLA shares. The unstable middle equilibrium moves in the opposite direction and approaches the high equilibrium.

The previous results characterize the local response of regular equilibrium COLA shares. The next result gives a global sufficient condition under which a sufficiently hawkish monetary policy response to inflation weakens the strategic complementarity and ensures uniqueness.

Proposition 15. ***(Policy and global uniqueness)*** *Under Conditions 1–2, a more hawkish Taylor rule weakens the strategic complementarity for any expected COLA share, i.e., for all $m \in [0, 1]$,*

$$\frac{\partial^2 \mathcal{V}}{\partial m \partial \phi_\pi} < 0$$

Moreover, starting from any policy coefficient satisfying these restrictions, there is a finite ϕ_π^U below the boundary at which $[\mathcal{S}_\pi + \mu(\mathcal{C}_0 - \mathcal{C}_1)] = 0$ such that every policy coefficient stronger than ϕ_π^U but below that boundary ensures a unique contracting general equilibrium.

[48] Quantitatively, the effect through the pass-through margin is relatively small.

Proof. I first show that a higher ϕ_π weakens strategic complementarity. Log differentiating $\mathcal{V}(m) = C\Pi(m)^2 \operatorname{Var}(\pi \mid m)$ with respect to m again gives

$$\mathcal{V}_m = \mathcal{V}\left[\frac{\partial \log \operatorname{Var}(\pi \mid m)}{\partial m} + 2\frac{\Pi_m}{\Pi}\right] \tag{58}$$

The first term in brackets captures the increase in inflation risk when the COLA share rises. The second term captures the change in the optimal pass-through. Under Conditions 1–2, the first term is positive and the second is nonnegative. Lemma 14 shows that $\mathcal{V}$ falls with ϕ_π. I next analyze how ϕ_π affects the two terms in brackets.

Notice that

$$\frac{\partial \log \operatorname{Var}(\pi \mid m)}{\partial m} = \frac{2(y_\zeta p_\omega + p_\zeta y_\omega)\mathcal{K}_m(m)}{\Delta(\mathcal{K}(m))}$$

I showed in Lemma 12 that

$$\mathcal{K}_m = \frac{\mathcal{S}_\pi}{\mathcal{S}_\pi - \mu m \mathcal{C}_1}\Pi$$

Thus,

$$\mathcal{K}_{m\phi_\pi} = \frac{\mathcal{S}_\pi}{\mathcal{S}_\pi - \mu m \mathcal{C}_1}\Pi_{\phi_\pi} < 0$$

Lemma 14 already shows that $1/\Delta(\mathcal{K}(m))$ decreases with ϕ_π. They together imply that

$$\frac{\partial}{\partial \phi_\pi}\left[\frac{\partial \log \operatorname{Var}(\pi \mid m)}{\partial m}\right] < 0$$

For the other term, differentiating the optimal pass-through in equation (53) leads to

$$\frac{\Pi_m}{\Pi} = \frac{\mu \mathcal{C}_1}{\mathcal{S}_\pi - \mu m \mathcal{C}_1} \geq 0$$

This ratio does not depend on ϕ_π. Therefore

$$\frac{\partial}{\partial \phi_\pi}\left(\frac{\Pi_m}{\Pi}\right) = 0$$

Thus, a more hawkish policy rule lowers (1) the gain from COLA $\mathcal{V}$, (2) the sensitivity of inflation risk in the conditional GE to the COLA share, and (3) the (normalized) sensitivity of the optimal COLA pass-through to the COLA share. In addition, all three terms in equation (58) are positive. They together imply a higher ϕ_π weakens the strategic complementarity, i.e., $\mathcal{V}_{m\phi_\pi} < 0$.

To establish global uniqueness, first notice that positive pass-through reaches zero at the finite policy coefficient

$$\bar{\phi}_\pi = \frac{\mathcal{S}_\pi + \mu(\gamma_c^2 p_\zeta \sigma_\rho^2 - \mathcal{C}_1)}{\mu \gamma_c y_\zeta \sigma_u^2}$$

At this value, $\mathcal{S}_\pi + \mu(\mathcal{C}_0 - \mathcal{C}_1) = 0$. Starting from any admissible policy coefficient, this expression remains positive until $\bar{\phi}_\pi$, while $\mathcal{C}_1$ is unchanged. So, Conditions 1–2 continue to hold until $\bar{\phi}_\pi$. Also, $\mathcal{S}_\pi - \mu m \mathcal{C}_1$ is unchanged and bounded away from zero for $m \in [0,1]$, while

$$\mathcal{D}_{\phi_\pi}(m) = \gamma_c(p_\zeta + \mu m p_\omega) > 0$$

Hence, a well-defined conditional GE for all $m \in [0,1]$ continues to hold. Thus all required restrictions remain valid on the open policy interval leading to $\bar{\phi}_\pi$.

I next show that $\mathcal{V}_m(m)$ approaches zero as ϕ_π rises to $\bar{\phi}_\pi$ by analyzing terms in equation (58) again.

As ϕ_π approaches $\bar{\phi}_\pi$ from below,

$$\Pi(m) = \frac{\mathcal{S}_\pi + \mu(\mathcal{C}_0 - \mathcal{C}_1)}{\mathcal{S}_\pi - \mu m \mathcal{C}_1} \longrightarrow 0$$

uniformly over $m \in [0,1]$, as the numerator converges to zero, while the denominator has a strictly positive lower bound independent of policy. Inflation variance $\mathrm{Var}\,(\pi|m)$ decreases with policy and is uniformly bounded by its maximum over $m \in [0,1]$ at the initial policy coefficient. Hence

$$\lim_{\phi_\pi \uparrow \bar{\phi}_\pi} \sup_{m \in [0,1]} \mathcal{V}(m) \propto \lim_{\phi_\pi \uparrow \bar{\phi}_\pi} \sup_{m \in [0,1]} \Pi(m)^2 \,\mathrm{Var}\,(\pi|m) = 0$$

On the other hand, the two positive terms in brackets in equation (58) both decrease with policy. So, they are also uniformly bounded by their respective maxima over $m \in [0,1]$ at the initial policy coefficient. Combining these bounds, equation (58) implies

$$\lim_{\phi_\pi \uparrow \bar{\phi}_\pi} \sup_{m \in [0,1]} \mathcal{V}_m(m) = 0$$

Since f_χ is bounded, and

$$\mathrm{BR}_m(m) = f_\chi(\mathcal{V}(m))\mathcal{V}_m(m) \geq 0$$

we have

$$\lim_{\phi_\pi \uparrow \bar{\phi}_\pi} \sup_{m \in [0,1]} \mathrm{BR}_m(m) = 0$$

from above. Therefore, by continuity, there is a finite $\phi_\pi^U < \bar{\phi}_\pi$ such that, for every $\phi_\pi \in (\phi_\pi^U, \bar{\phi}_\pi)$,

$$\sup_{m \in [0,1]} \mathrm{BR}_m(m) < 1$$

This means the aggregate best response is a contraction of $[0,1]$ into itself and thus has a unique fixed point. □

E.5 Uncertainty and the Anchoring Threshold

I finally study how uncertainty, measured by the volatility of primitive shocks, affects the private gain from COLA, the aggregate best-response function, and the anchoring role of monetary policy. As in Proposition 13, I let

$$\sigma_u = \sigma, \qquad \sigma_\rho = k\sigma$$

for some constant $k > 0$, and vary $\sigma > 0$. In other words, both shocks become more volatile, while their relative importance remains unchanged. As I will show immediately, this assumption keeps the optimal COLA pass-through unchanged, making the analysis simpler. Conceptually, it isolates the effect of the total amount of aggregate uncertainty from the composition of uncertainty.

Proposition 16. *(Uncertainty and the COLA gain) Under Condition 1, for every predetermined $m \in [0,1]$, higher uncertainty leaves the sensitivity of the aggregate real wage to inflation surprises and the optimal COLA pass-through unchanged:*

$$\mathcal{K}_\sigma(m) = 0, \qquad \Pi_\sigma(m) = 0$$

On the other hand, it raises inflation risk in the conditional GE and the gain from COLA:

$$\frac{\partial \operatorname{Var}(\pi \mid m)}{\partial \sigma} = \frac{2 \operatorname{Var}(\pi \mid m)}{\sigma} > 0, \qquad \mathcal{V}_\sigma(m) = \frac{2\mathcal{V}(m)}{\sigma} > 0$$

Proof. Log differentiating $\mathcal{V}(m) = C\Pi(m)^2 \operatorname{Var}(\pi \mid m)$ with respect to σ, we have

$$\frac{\mathcal{V}_\sigma}{\mathcal{V}} = \frac{\partial \log \operatorname{Var}(\pi \mid m)}{\partial \sigma} + 2\frac{\Pi_\sigma}{\Pi}$$

I first show that the optimal pass-through does not change. Substituting $\sigma_u = \sigma$ and $\sigma_\rho = k\sigma$ into the definitions of $\mathcal{S}_\pi$, $\mathcal{C}_0$, and $\mathcal{C}_1$, we have

$$\begin{aligned}
\mathcal{S}_\pi &= [(\gamma_c p_\zeta k)^2 + y_\zeta^2]\sigma^2 \\
\mathcal{C}_0 &= (\gamma_c^2 p_\zeta k^2 - \gamma_c \phi_\pi y_\zeta)\sigma^2 \\
\mathcal{C}_1 &= (-\gamma_c^2 p_\zeta p_\omega k^2 + y_\zeta y_\omega)\sigma^2
\end{aligned}$$

Hence, they are all proportional to σ^2, implying both the numerator and the denominator in

$$\Pi(m) = \frac{\mathcal{S}_\pi + \mu(\mathcal{C}_0 - \mathcal{C}_1)}{\mathcal{S}_\pi - \mu m \mathcal{C}_1}$$

are proportional to σ^2. So $\Pi_\sigma(m) = 0$. Since $\mathcal{K}(m) = m\Pi(m) - 1$, we also have $\mathcal{K}_\sigma(m) = 0$.

Next, recall from the conditional GE that

$$\pi_u(m) = \frac{y_\zeta}{\Delta(\mathcal{K}(m))}, \qquad \pi_\rho(m) = \frac{\gamma_c p_\zeta}{\Delta(\mathcal{K}(m))}$$

Both coefficients are unchanged when σ rises because $\mathcal{K}(m)$ is unchanged. Therefore

$$\mathrm{Var}(\pi \mid m) = \pi_u(m)^2\sigma^2 + \pi_\rho(m)^2 k^2 \sigma^2$$

implies

$$\frac{\partial \log \mathrm{Var}(\pi \mid m)}{\partial \sigma} = \frac{2}{\sigma} > 0$$

Together with $\Pi_\sigma = 0$, this gives $\mathcal{V}_\sigma = 2\mathcal{V}/\sigma$, leading to the stated inequality.

Finally, to make sure that one always stays in the admissible parameter range as σ varies, notice that $\mathcal{D}(m)$ does not depend on σ. Since $\mathcal{S}_\pi$, $\mathcal{C}_0$, and $\mathcal{C}_1$ all scale with σ^2, the signs of $\mathcal{S}_\pi - \mu\mathcal{C}_1$, $\mathcal{S}_\pi + \mu(\mathcal{C}_0 - \mathcal{C}_1)$, $\mathcal{C}_0$, and $\mathcal{C}_1$ do not change for $\sigma > 0$. Thus, a conditional GE for any $m \in [0,1]$ is well-defined, and Conditions 1-2 remain valid at every $\sigma > 0$. □

Fixing the aggregate COLA share, higher uncertainty raises inflation risk without changing the optimal COLA pass-through. To understand the intuition, notice that while inflation becomes more volatile, its information about the revenue factor does not change because the two shocks are scaled by the same proportion. The optimal wage compensation for a given amount of inflation is thus unchanged.

Corollary 3. *(Uncertainty and stable contracting equilibria) Under Conditions 1–2, let m^* be a regular locally stable interior contracting equilibrium. Then*

$$\frac{dm^*}{d\sigma} \geq 0$$

The inequality is strict if and only if $f_\chi(\mathcal{V}(m^*)) > 0$.

Proof. Local stability implies $1 - \mathrm{BR}_m(m^*) > 0$, so the implicit function theorem gives a locally differentiable equilibrium branch. Differentiating the equilibrium condition $m^* = F_\chi(\mathcal{V}(m^*, \sigma))$ with respect to σ, we have

$$\frac{dm^*}{d\sigma} = f_\chi(\mathcal{V}(m^*)) \left[\mathcal{V}_\sigma(m^*) + \mathcal{V}_m(m^*)\frac{dm^*}{d\sigma}\right]$$

Solving for $dm^*/d\sigma$, we have

$$\frac{dm^*}{d\sigma} = \frac{f_\chi(\mathcal{V}(m^*))\mathcal{V}_\sigma(m^*)}{1 - \mathrm{BR}_m(m^*)} = \frac{2f_\chi(\mathcal{V}(m^*))\mathcal{V}(m^*)}{\sigma[1 - \mathrm{BR}_m(m^*)]}$$

Proposition 16 shows that $\mathcal{V}_\sigma > 0$, while the denominator is positive by local stability. This leads to the stated inequality, which is strict if and only if $f_\chi(\mathcal{V}(m^*)) > 0$. □

Higher uncertainty raises the gain from COLA and shifts the aggregate best response upwards. Along a stable equilibrium branch, more pairs adopt COLA, which further raises inflation risk and the gain from COLA through the contract accelerator.

I finally study how uncertainty affects the anchoring role of monetary policy. When there are three regular equilibrium COLA shares, a more hawkish policy rule shifts the stable low and high fixed points weakly downwards and the unstable middle fixed point upwards. If the middle and high fixed points meet at a nondegenerate fold, the aggregate best response is tangent to the 45° line there. The next proposition shows that higher uncertainty raises the policy coefficient required to reach this tangency.

Proposition 17. *(Policy anchoring threshold) Under Conditions 1–2, consider an interior fold at which an unstable middle equilibrium and a stable high-COLA equilibrium meet. Let the fold $(\tilde{m}, \tilde{\phi}_\pi)$ be characterized by*

$$G(\tilde{m}, \tilde{\phi}_\pi, \sigma) = 0$$
$$G_m(\tilde{m}, \tilde{\phi}_\pi, \sigma) = 0$$

where

$$G(m, \phi_\pi, \sigma) := \mathrm{BR}(m, \phi_\pi, \sigma) - m$$

If $G_{mm} \neq 0$, then the local policy threshold increases with uncertainty

$$\frac{d\tilde{\phi}_\pi}{d\sigma} > 0$$

Proof. Differentiating the first fold equation with respect to σ, we have

$$0 = G_\sigma + G_{\phi_\pi}\frac{d\tilde{\phi}_\pi}{d\sigma} + G_m\frac{d\tilde{m}}{d\sigma}$$

At the fold, $G_m = 0$. Solving for $d\tilde{\phi}_\pi/d\sigma$ and cancelling the positive cost density, we have

$$\frac{d\tilde{\phi}_\pi}{d\sigma} = -\frac{G_\sigma}{G_{\phi_\pi}} = -\frac{\mathcal{V}_\sigma}{\mathcal{V}_{\phi_\pi}} = -\frac{2\mathcal{V}}{\sigma\mathcal{V}_{\phi_\pi}} > 0$$

where all terms are evaluated at the fold. The last inequality follows from $\mathcal{V} > 0$ and $\mathcal{V}_{\phi_\pi} < 0$. □

The result describes an upward-sloping local anchoring threshold in the policy-uncertainty plane. Higher uncertainty raises the value of COLA insurance. A more hawkish policy response is then needed to lower inflation risk (and the optimal pass-through) enough to reach the tangency point at which the high-COLA equilibrium disappears.

As a remark, this result does not say anything about global uniqueness. It only describes properties at a fold. Under some typical cost functions F_χ, there is a unique equilibrium if the aggregate best response moves beyond that fold (like the regular "S"-shaped best response). In contrast, Proposition 15 gives a sufficient condition for global uniqueness without imposing a shape restriction on F_χ, except for those regularity conditions. So, it is often sufficient but not

necessary. The equilibrium can already be unique just beyond a fold even if the slope of the best-response function is still larger than one at some COLA shares.

E.6 AS-AD Diagram

In this section, I show how the conditional GE can be visualized by AS-AD diagrams, and how the equilibrium COLA share affects those curves.

For any given COLA share m, the AS curve is the Phillips curve, and the AD curve combines the DIS curve and the monetary policy equation in the three-equation representation in Proposition 10. Namely, the AS-AD representation can be written as

$$\begin{aligned} AS: &\qquad \pi = \frac{a}{1-b\mathcal{K}(m)}\hat{y} + \frac{1}{1-b\mathcal{K}(m)}u \\ AD: &\qquad \hat{y} = -d_\pi \pi + d_\rho \rho \end{aligned}$$

where a, b, d_π, d_ρ are positive constants, and $\mathcal{K}(m) < 0$ and is increasing with m, as defined in Proposition 10. Strictly speaking, d_ρ should also be a function of $\mathcal{K}(m)$. However, we can treat it as a constant when we analyze the response of inflation to aggregate shocks, as this will not affect the qualitative results. It affects the output response to shocks, which is not my main focus here.

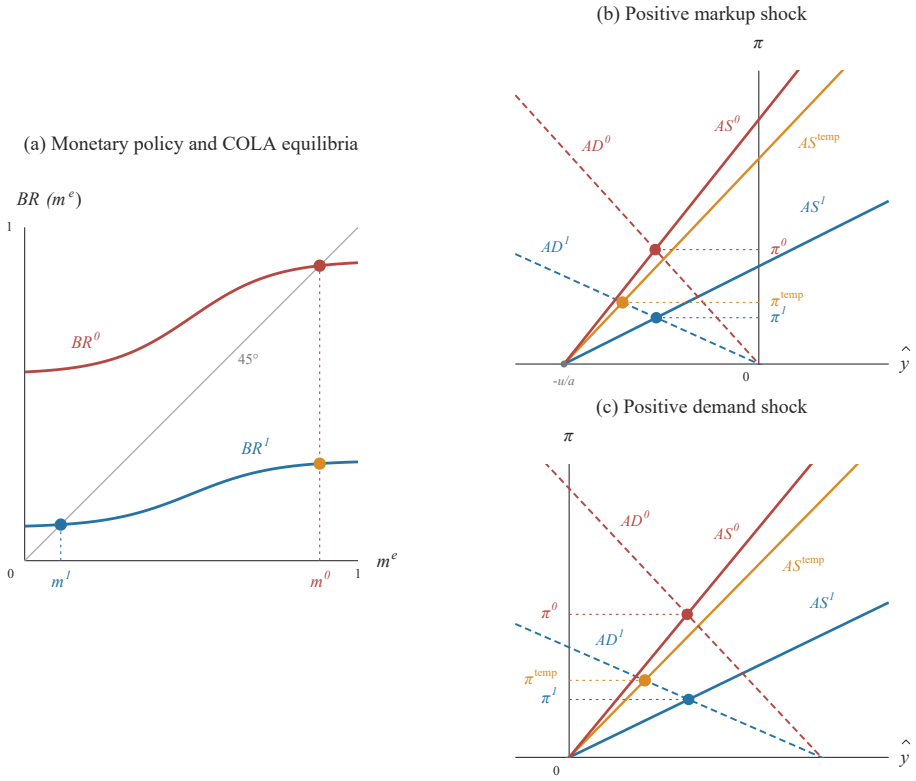
Given the AS-AD system above, the COLA share affects the economy only through the AS curve. For a given markup shock u, we can write the AS curve as

$$\pi = \frac{a}{1-b\mathcal{K}(m)}\left(\hat{y} - \frac{u}{a}\right)$$

Hence, an increase in m raises $\mathcal{K}(m)$, leading to a counterclockwise rotation of the AS curve around $(\frac{u}{a}, 0)$ on the $\hat{y}$-π plane. For a given demand shock, in the simplified case, a higher COLA share m rotates the AS curve counterclockwise around the origin. In neither case does m affect the AD curve directly. I show two types of comparative statics: policy rule hawkishness and primitive shock uncertainty.

Policy Hawkishness. Figure 29 illustrates how a more hawkish monetary policy rule affects the economy through the two spirals. In the main text, I use the inflation response to a given markup shock as an example to explain how policy stabilizes inflation twice. The case for a given demand shock is similar. I stack the AS-AD diagrams in response to a markup shock and a demand shock in panels (b) and (c).

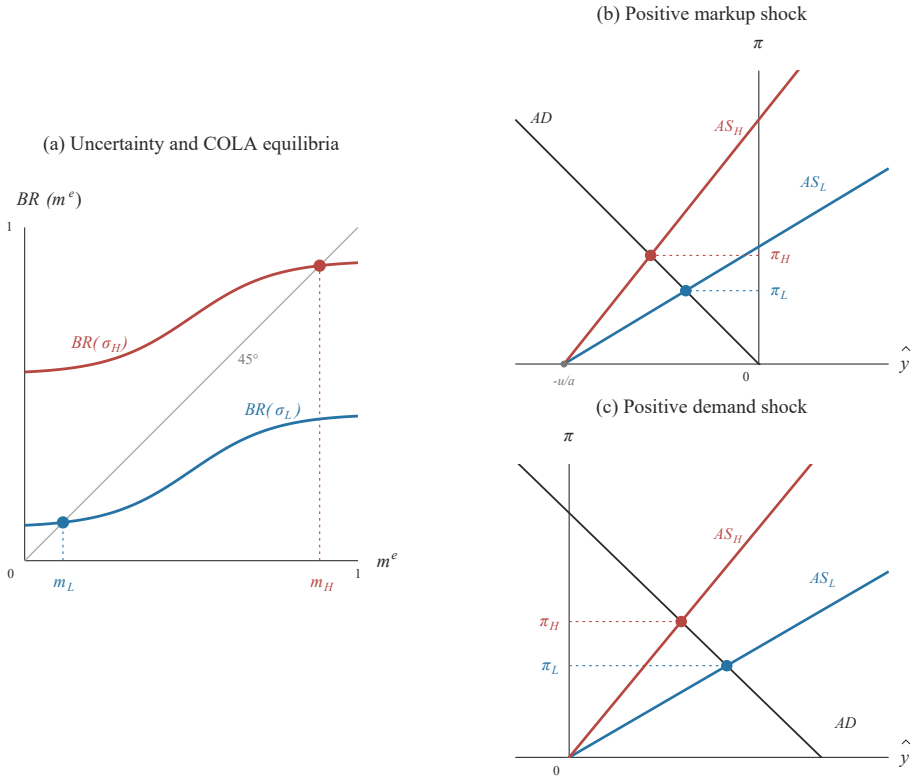
Figure 29: Illustration of a more hawkish Taylor rule



Notes: The left panel illustrates the aggregate best response, where BR^1 corresponds to a higher ϕ_π than BR^0. Colors link each equilibrium to its corresponding AS curve, AD curve, and inflation response. The AS curves are the COLA-regime-specific Phillips curves, and the AD curves combine the DIS and MP equations in Proposition 10 and summarize the demand block. The case for a demand shock is similar and is reported in the appendix. For illustration, I omit the change in the AD curve when m_t changes, which does not affect the main conclusion.

Primitive-Shock Volatility. Figure 30 illustrates how uncertainty, measured by primitive-shock volatility, affects the economy. Without endogenous contract choice, higher uncertainty would not affect the economy at all, which is a standard result in a log-linearized system. However, when private agents choose wage contracts to allocate risks, uncertainty bites. When either markup shocks or demand shocks become more volatile, COLA has a higher insurance value, which shifts the best-response curve up. This induces more firm-union pairs to adopt COLA, which steepens the AS curve (Panels b and c). Thus, a given shock will generate a larger inflation response, which further amplifies inflation volatility and COLA adoption. The economy settles down in a new equilibrium with a higher COLA share and higher inflation responses to shocks.

Figure 30: Illustration of higher volatility of primitive shocks



Notes: The left panel illustrates the aggregate best response, where BR^1 corresponds to higher volatility of primitive shocks than BR^0. Colors link each equilibrium to its corresponding AS curve, AD curve, and inflation response. The AS curves are the COLA-regime-specific Phillips curves, and the AD curves combine the DIS and MP equations in Proposition 10 and summarize the demand block. The case for a demand shock is similar and is reported in the appendix. For illustration, I omit the change in the AD curve when m_t changes, which does not affect the main conclusion.

E.6.1 The Effect of Price Rigidity

This paper endogenizes wage rigidity through contract choice but fixes price rigidity. Prices can also become more flexible in times of high inflation, which is documented by a large body of empirical literature, and may result from menu costs, rational inattention, or other forces. I discuss briefly here how a changing price rigidity can affect COLA adoption.

In my model, price rigidity is captured by the parameter κ. A higher κ means more firms reset their price each period and thus, price is more flexible. If we write κ explicitly in the AS-AD framework, we have

$$AS: \qquad \pi = \frac{a\kappa}{1 - b\kappa\mathcal{K}(m)}\hat{y} + \frac{1}{1 - b\kappa\mathcal{K}(m)}u$$
$$AD: \qquad \hat{y} = -d_\pi \pi + d_\rho \rho$$

We can solve for π for any given m and κ:

$$\pi(m,\kappa) = \frac{u + \kappa a d_\rho \rho}{1 + \kappa\left(b|\mathcal{K}(m)| + a d_\pi\right)}$$

Thus, for its response to a given markup shock,

$$\frac{\partial}{\partial \kappa}\left(\frac{\partial \pi}{\partial u}\right) = -\frac{b|\mathcal{K}(m)| + a d_\pi}{\left[1 + \kappa\left(b|\mathcal{K}(m)| + a d_\pi\right)\right]^2} < 0$$

while for a given demand shock,

$$\frac{\partial}{\partial \kappa}\left(\frac{\partial \pi}{\partial \rho}\right) = \frac{a d_\rho \rho}{\kappa^2\left[1/\kappa + \left(b|\mathcal{K}(m)| + a d_\pi\right)\right]^2} > 0$$

That is to say, when price becomes more flexible, for any given COLA share m, inflation will become less responsive to a given markup shock, but more responsive to a given demand shock. Intuitively, following a markup shock, the initial cost-push shock induces a contraction in demand, which lowers firms' marginal costs and creates an incentive to reduce prices (or raise prices by less). When prices are more flexible, firms can reduce their prices (or raise their prices by less) more quickly, resulting in a smaller overall inflation response and a smaller output loss relative to a more price-rigid economy. By contrast, following a demand shock, greater price flexibility allows firms to translate stronger demand and higher marginal costs into prices more quickly, leading to a larger inflation response but a smaller output response, as higher prices partly offset demand.

How does price rigidity affect COLA adoption? For simplicity, we fix COLA pass-through, whose variation is small in the data. Then, the gain from COLA is proportional to inflation volatility. Using the results above, it is easy to show that

$$\frac{\partial \operatorname{Var}(\pi|m)}{\partial \kappa} = 2\frac{-\left(b|\mathcal{K}(m)| + a d_\pi\right)\sigma_u^2 + \kappa a^2 d_\rho^2 \sigma_\rho^2}{\left[1 + \kappa\left(b|\mathcal{K}(m)| + a d_\pi\right)\right]^3}$$

The sign is ambiguous and depends on the relative primitive-shock volatility. If inflation is primarily driven by markup shocks, higher price flexibility tends to reduce inflation volatility and discourage COLA adoption. Conversely, if inflation is primarily driven by demand shocks, more flexible prices tend to raise inflation volatility and encourage COLA adoption.

Therefore, the level of price rigidity can affect the level of wage rigidity through contract choice, and the direction depends on the nature of the structural shocks. Of course, one should interpret this analysis as in partial equilibrium in the sense that I do not model state-dependent pricing explicitly. I only vary κ explicitly to see its effect on inflation and on contract choice, and the direction is only from price rigidity to wage rigidity. A full account of both endogenous price rigidity and wage rigidity may generate non-trivial interactions of the two types of nominal rigidity. This is an interesting extension and is left for future research.

Finally, it is also worth mentioning that even if one assumes fully flexible prices, endogenous contract choice still matters. This is because Taylor staggered nominal wage contracts still retain some nominal rigidity. To be more precise, letting $\kappa \to \infty$, for a given m, the AS curve becomes

$$\pi = \frac{a}{b|\mathcal{K}(m)|}\hat{y}$$

In this case, a cost-push shock will not have any effects: an initial inflationary pressure reduces demand, and firms reduce prices immediately. As a result, neither price nor output responds in equilibrium. Nevertheless, the AS curve is still positively sloped rather than vertical, and demand shocks will generate inflation and output responses.

F The Full Model and the Solution Algorithm

The analytical convenience in the last section comes at the cost of richer dynamic linkages by shutting down the persistence of exogenous shocks and restricting the contract length to one. Once we relax these assumptions, the model will have much more complicated forward- and backward-looking features. In this section, I look at a particular class of contracting equilibrium: I assume the time path of the primitive-shock volatilities is deterministic and look for equilibria with a deterministic time path of the contract distribution. I then characterize that class of contracting equilibria. Finally, I show some numerical results. I will relax those restrictions in model extensions. There, the time path of the primitive-shock volatilities can be stochastic. Alternatively, I can also generalize the contractual environment so that the endogenous time path of the contract distribution can become stochastic even if primitive-shock volatilities stay constant. Those relaxations are less tractable analytically and require some more advanced computational tools.

F.1 Deterministic Time Path of Primitive-Shock Volatilities

Assumption 5. *The time path of the primitive-shock volatilities* $\left\{\Sigma_t : \Sigma_t = diag(\sigma_{ut}^2, \sigma_{\rho t}^2)\right\}_{t\in\mathbb{Z}}$ *is deterministic.*

When we analyze the stripped-down version of the model under Assumption 1 in the last section, we implicitly make this assumption, because we assume σ_u and σ_ρ are constant over time, and all agents know them. This assumption retains the linearity of the forecasting problem if the revenue and inflation processes are linear in the first place, a standard result in the linear-Gaussian environment. With that observation, we can show that for any given time path of the COLA distribution $\{\mathcal{M}_t\}_t$, where $\mathcal{M}_t = \{\mathcal{M}_{t,k}\}_{k=1}^N$, the REE solution of the conditional general equilibrium has a linear MA representation in terms of structural shocks. In particular, given the assumed information structure that a contracting cohort observes all past shocks at each contract renewal, the conditional equilibrium is actually closed with the class of finite MA representations:

Lemma 15. *Under Assumption 5, if a contracting cohort observes all past shocks at each contract renewal, and $v_t = (u_t, \rho_t)^T$ is AR(1), then given any time path of contract distribution $\{\mathcal{M}_t\}_{t\in\mathbb{Z}}$, the conditional general equilibrium is closed within the class of MA(N+1) representations in terms of v_t, where N is the contract duration, i.e.,*

$$\pi_t = \sum_{s=0}^{N} q_{t,s}^{\pi} v_{t-s} \tag{59}$$

$$\hat{\zeta}_t = \sum_{s=0}^{N} q_{t,s}^{\zeta} v_{t-s} \tag{60}$$

Proof. We analyze the first five equations (36)-(40) in Lemma 5 which characterize the conditional general equilibrium. Given (59)-(60), both the LHS and the RHS of equations (36) and (38)-(40) will only contain $v_{t-N}, ..., v_t$. Therefore, it suffices to show that the RHS of equation (37), the aggregate wage index, depends only on shocks $v_{t-N}, ..., v_t$.

Fix a cohort $k \in \{1, \ldots, N\}$. I first consider the COLA part of the wage index. The contract was chosen at date $t-k$, and by assumption the renewing cohort observes all shocks dated weakly before $t-k$. Because v_t is Markov and the equilibrium under consideration is linear Gaussian, the prior forecast of any date-t variable based on the renewal information is a linear function of v_{t-k}. The additional information contained in the inflation history observed during the life of the COLA contract,

$$\pi^t \setminus \mathcal{F}_{t-k}$$

is generated by the innovations realized between $t-k+1$ and t. Under the linear conjecture, the Kalman projection $\mathbb{E}_{t-k}[\hat{\zeta}_t|\pi^t]$ is therefore a linear function of $(v_{t-k}, v_{t-k+1}, \ldots, v_t)$. Since $k \leq N$, this vector is contained in $(v_{t-N}, \ldots, v_t)$. Hence each term $\mu\mathbb{E}_{t-k}[\hat{\zeta}_t|\pi^t]$ in the COLA component of $\hat{\omega}_t$ depends only on shocks dated from $t-N$ to t.

Now consider the non-contingent wage component for the same cohort. Its relevant term is

$$\mu\mathbb{E}_{t-k}[\hat{\zeta}_t] + \sum_{\tau=0}^{k-1} (\pi_{t-\tau} - \mathbb{E}_{t-k}[\pi_{t-\tau}])$$

The first term is a linear function of v_{t-k} by the same Markov argument. For each $\tau \in \{0, \ldots, k-1\}$, the forecast error $\pi_{t-\tau} - \mathbb{E}_{t-k}[\pi_{t-\tau}]$ removes all components that are measurable with respect to $\mathcal{F}_{t-k}$. The remaining part is driven only by innovations realized after $t-k$ and no later than $t-\tau$. These dates are again contained in $\{t-N, \ldots, t\}$. Thus every non-contingent wage term in equation (37) is a linear function of $(v_{t-N}, \ldots, v_t)$.

The contract distribution $\{\mathcal{M}_{t,k}\}_{k=1}^{N}$ is taken as given and deterministic in the conditional general equilibrium considered here. Therefore the weighted sum in the aggregate wage index is also a linear function of $(v_{t-N}, \ldots, v_t)$. This proves that equation (37) maps the conjectured finite-dimensional representation into itself. The remaining four equilibrium equations are linear and, as

argued above, involve only the same vector of shocks. Hence the conditional general equilibrium can be solved within the class

$$\pi_t = \sum_{s=0}^{N} q_{t,s}^{\pi} v_{t-s}, \qquad \hat{\zeta}_t = \sum_{s=0}^{N} q_{t,s}^{\zeta} v_{t-s}$$

This is the desired $MA(N+1)$ representation. □

In general, if one assumes that v_t is AR(p), the conditional equilibrium will be closed in the class of MA$(N+p)$ representations; if one assumes that past shocks are never revealed, the conditional equilibrium will only be closed in the class of MA(∞) representations. Notice the lemma does not say anything about existence and uniqueness. The resulting equation system of the coefficients q's will be nonlinear. Numerically, however, the conditional equilibrium indeed exists and is almost always unique under standard parameter values. In addition, the difference of the solution under different information structures regarding past shocks is limited.

A useful by-product from this procedure is the inference of $\hat{\zeta}_t$ conditional on the observed inflation history. In particular, for any given $\{\mathcal{M}_t\}_{t\in\mathbb{Z}}$, the forecast of $\hat{\zeta}_t$ can be written as a function of π:

$$\mathbb{E}\left[\hat{\zeta}_t \middle| \mathcal{F}_{t-k}, \pi_{t-k+1}, ..., \pi_t\right] = \sum_{s=0}^{k-1} l_{t,s}\left(\{\mathcal{M}_t\}_{t\in\mathbb{Z}}\right)\pi_{t-s} + f_{t,k}\left(\mathcal{F}_{t-k}\right)$$

where $l_{t,s}$ are terms from the projection coefficient vector $\text{Cov}\left[\hat{\zeta}_t, (\pi_t, ..., \pi_{t-k+1})\right] \text{Var}\left(\pi_t, ..., \pi_{t-k+1}\right)^{-1}$, and $f_{t,k}\left(\mathcal{F}_{t-k}\right)$ is a linear function of past shocks, which is not important for our purpose here. We highlight the dependence of the coefficients $l_{t,s}$ on the path of the contract distribution $\{\mathcal{M}_t\}_{t\in\mathbb{Z}}$.

Next, we can go back to the contract choice problem to find out the path of $\{\mathcal{M}_t\}_{t\in\mathbb{Z}}$ that is consistent with the service firms' endogenous contract choice, which defines the contracting general equilibrium. The key object that matters for an individual service firm's contract choice is its perceived gains from COLA. Similar to the analytical model where the gains from COLA depend on COLA pass-through and inflation volatility from any perceived COLA share, the gains in the general model with long-term contracts hinge on the average COLA-pass-through and the variance-covariance structure of future inflation from any perceived path of the contract distribution. In a rational equilibrium, the perceived path and the actual path must coincide. The next proposition summarizes the above discussion.

Proposition 18. *Under Assumption 5,*

1) for an individual firm-union pair that renews its contract at the end of $t-1$, denote by $\left\{\mathcal{M}_{t+\tau}^{e}\right\}_{\tau\geq 0}$ the firm's perceived path of contract distribution from period t onwards, and $\{\mathcal{M}_{t+\tau}\}_{\tau<0}$ the actual contract distribution up to $t-1$. Then, the perceived gains from COLA for this firm are

$$\mathcal{V}_{t-1}\left(\{\mathcal{M}_{t+\tau}\}_{\tau<0}, \{\mathcal{M}_{t+\tau}^{e}\}_{\tau\geq 0}\right) \propto \sum_{n=1}^{N} \beta^n \mathbb{E}_{t-1}\left[\sum_{k=0}^{n-1}\left(\mu l_{t-1+n,k}+1\right)\left(\pi_{t-1+n-k} - \mathbb{E}_{t-1}[\pi_{t-1+n-k}]\right)\right]^2$$

2) the contracting general equilibrium in the general model can be found by solving the following equation system for the path of the contract distribution $\{\mathcal{M}_t\}_{t\in\mathbb{Z}}$*:*

$$\mathcal{M}_{t,1} = \frac{1}{N} F_\chi \left(\mathcal{V}_{t-1} \left(\{\mathcal{M}_{t+\tau}\}_{\tau<0}, \{\mathcal{M}_{t+\tau}\}_{\tau\geq 0} \right) \right)$$
$$\mathcal{M}_{t,k} = \mathcal{M}_{t-1,k-1}, \quad k \in \{2, ..., N\}$$

for all $t \in \mathbb{Z}$*.*

Proof. Follows immediately from Lemma 5 and the projection discussed in the text above. □

In words, for any perceived path of contract distribution, service firms' choice leads to the actual COLA inflow in each period. The actual path of the contract distribution implied by the path of COLA inflow must be consistent with the perceived path, which is a big fixed point in terms of the entire path of the contract distribution $\{\mathcal{M}_t\}_{t\in\mathbb{Z}}$.

F.2 Solution Algorithm

Given the two layers of fixed points with nonlinearity and the forward- and backward-looking behavior in the contracting general equilibrium, it seems impossible to solve the general model analytically. However, our previous discussion indeed provides us with guidance on how to solve the model numerically. The algorithm contains two parts. First, we discuss how to solve for a stationary contracting equilibrium in the sense that $\Sigma_t = \Sigma_0$ for all t. Then, we study, starting from a stationary contracting equilibrium, given a deterministic path of primitive shock volatilities, what the contracting equilibrium will be. In other words, $\Sigma_t = \Sigma_0$ for $t < 0$, and at the end of period zero, agents know the path of primitive shock volatilities becomes $\{\Sigma_t\}_{t\geq 0}$.[49]

F.2.1 Stationary Contracting Equilibrium

When $\Sigma_t = \Sigma_0$, we guess that the contract distribution is time-invariant, i.e., $\mathcal{M}_t = \mathcal{M}$. In addition, the REE solution of all other variables given $\mathcal{M}_t$ is also stationary, or in other words, independent of the calendar time. This leads to the following algorithm.

1. Guess a stationary contract distribution $\mathcal{M}_t = \mathcal{M}^0$.

2. Guess a stationary REE solution for inflation and real revenue factor of this form

$$\pi_t^0 = \sum_{s=0}^{N} q_s^{\pi 0} v_{t-s}$$
$$\hat{\zeta}_t^0 = \sum_{s=0}^{N} q_s^{\zeta 0} v_{t-s}$$

[49] These two exercises seem analogous to the stationary distribution and the transitional dynamics of an Aiyagari model. But again, at every point in time, the economy is on its stochastic equilibrium path.

where $v_t = (u_t, \rho_t)^T$ contains exogenous shocks. The guess contains $2(N+1)$ vectors $\left\{q_s^{\pi 0}, q_s^{\zeta 0}\right\}_{s=0}^{N}$, each with size 1×2 .

3. Given the guess of π_t^0 and $\hat{\zeta}_t^0$, set up the Kalman filter and obtain the forecast of real revenue factor as a linear function of inflation history.

4. With the forecast from step 3, solve for the real wage index $\hat{\omega}_t^0$ by averaging different wages, solve for entrepreneur consumption $\hat{c}_t^{e0}$ from the aggregate resource constraint, and solve for i_t^0 from the dynamic IS curve. All of them are linear functions of current and past shocks.

5. Update $\hat{\zeta}_t^1$ from the Taylor rule, and update π_t^1 from the Phillips curve, using $\hat{\omega}_t^0, \hat{c}_t^{e0}, i_t^0$ solved in step 4, and the initial guess π_t^0 and $\hat{\zeta}_t^0$.

6. Replace the initial guess of π_t^0 and $\hat{\zeta}_t^0$ with π_t^1 and $\hat{\zeta}_t^1$. Repeat steps 3-6 until convergence.

7. Given $\pi_t^1(\mathcal{M}^0)$ and $\hat{\zeta}_t^1(\mathcal{M}^0)$, compute the gains from COLA according to the first part of Proposition 18. Update the contract distribution according to the second part of Proposition 18 to get $\mathcal{M}^1$.

8. Replace the initial guess of $\mathcal{M}^0$ with $\mathcal{M}^1$. Repeat steps 2-8 until convergence.

F.2.2 Contracting Equilibrium with Time-Varying Contract Distribution

Suppose $\Sigma_t = \Sigma_{old}$ for $t < 0$, and at the end of period zero, agents know that the path of primitive shock volatility becomes $\{\Sigma_t\}_{t \geq 0}$. For simplicity, we assume Σ_t converges to a constant again in the end. That is $\Sigma_t = \Sigma_{new}$ for $t \geq T$ where T is a big number. To compute the dynamics, we first compute the initial and ending stationary equilibria following the last subsection. Denote by $\mathcal{M}_{old/new}, \pi_{t,old/new}, \hat{\zeta}_{t,old/new}$ the solution to them, respectively. Now, we can spell out the algorithm.

1. Guess a path of contract distribution $\left\{\mathcal{M}_t^0\right\}_{t=1}^{T}$, and $\mathcal{M}_t = \mathcal{M}_{old}$ for $t \leq 0$, $\mathcal{M}_t = \mathcal{M}_{new}$ for $t > T$. Because of the staggered contract structure, we only need to guess the path of COLA inflow $\left\{\mathcal{M}_{t,1}^0\right\}_{t=1}^{T}$, and recover $\left\{\mathcal{M}_t^0\right\}_{t=1}^{T}$ from the law of motion. Therefore, we need to guess T numbers.

2. Guess the REE solution for inflation and real revenue factor for each $t \in \{1, ..., T\}$

$$\pi_t^0 = \sum_{s=0}^{N} q_{t,s}^{\pi 0} v_{t-s}$$
$$\hat{\zeta}_t^0 = \sum_{s=0}^{N} q_{t,s}^{\zeta 0} v_{t-s}$$

and for $t \leq 0$ or $t > T$, we use the corresponding stationary REE solution. The guess contains $T \times 2(N+1)$ vectors, each with size 1×2.

3. Given the guess of $\{\pi_t^0, \hat{\zeta}_t^0\}_{t=1}^T$, we do the same thing as in the stationary case to solve for the linear system for the given path of contract distribution, which includes applying the Kalman filter, solving for $\{\hat{\omega}_t^0, \hat{c}_t^{e0}, i_t^0\}_{t=1}^T$, and updating to get $\{\pi_t^1, \hat{\zeta}_t^1\}_{t=1}^T$.

4. Replace the initial guess $\{\pi_t^0, \hat{\zeta}_t^0\}_{t=1}^T$ with $\{\pi_t^1, \hat{\zeta}_t^1\}_{t=1}^T$. Repeat steps 3-4 until convergence.

5. Given $\left\{\pi_t^1\left(\left\{\mathcal{M}_{t,1}^0\right\}_{t=1}^T\right), \hat{\zeta}_t^1\left(\left\{\mathcal{M}_{t,1}^0\right\}_{t=1}^T\right)\right\}_{t=1}^T$, compute the gains from COLA according to the first part of Proposition 18. Update the COLA inflow according to the second part of Proposition 18 to update the path of new COLA inflow $\left\{\mathcal{M}_{t,1}^1\right\}_{t=1}^T$.

6. Replace the initial guess $\left\{\mathcal{M}_{t,1}^0\right\}_{t=1}^T$ with $\left\{\mathcal{M}_{t,1}^1\right\}_{t=1}^T$. Repeat steps 2-6 until convergence. Then, recover $\{\mathcal{M}_t\}_{t=1}^T$ from the converged COLA inflow path.

G Calibration and Quantitative Details

This section provides further explanations and robustness checks to the calibration and the counterfactual exercises. I first discuss the labor-adjustment cost used to discipline firm-side risk exposure. I then describe the construction of two model moments: the value of internal funds and COLA pass-through. The section then explains how the utility cost of COLA is converted into a money-metric number, reports robustness to the extrapolation from union contracts to the broader economy, and documents the volatility decomposition for the Great Moderation.

G.1 Adding the Labor Adjustment Cost

The quantitative model introduces a static labor-adjustment cost on the service firm side. The purpose is narrow. In the data, COLA pass-through is substantial but below one. In the model, pass-through is lower when the firm side is sufficiently exposed to revenue risk: if inflation predicts weak revenue conditions, full indexation would transfer too much risk to the firm. The labor-adjustment cost is a parsimonious way to increase this firm-side exposure while leaving the structure of the contract-choice problem intact.

The service firm's cash flow before external financing is

$$q_{it} = \zeta_t \eta_{it} L_{it}^c - w_{it} L_{it} - \nu (L_{it} - \bar{L})^2$$

where $c = 1 - 1/\varepsilon_x$. The cost is quadratic in the deviation of employment from its steady-state level. This formulation differs from a conventional adjustment cost relative to last period's employment. A lagged-employment cost would make each firm's previous employment a state variable. Because previous employment depends on the age and type of the outstanding contract, the contracting problem would have to track firm-level histories in addition to the aggregate contract

distribution. That would destroy the current tractability of the model, whose inner problem remains a repeated per-period employment choice conditional on the wage rule and the aggregate state.

The static adjustment cost avoids this complication. Since the marginal cost term is zero at the steady state, it does not change steady-state levels. It changes local slopes: employment responds less elastically to the wage and revenue factor, and the worker's risk exposure coefficient μ changes accordingly. Once μ is updated, the rest of the model is unchanged. The same three-equation logic, the Kalman-filter inference problem, and the contract-choice fixed point continue to apply.

G.2 Model Moments Construction

The internally calibrated parameters are chosen to match a set of steady-state and dynamic moments. Two steady-state moments summarize the firm-union side: the value of internal funds and the labor share. Dynamic moments include the average COLA pass-through, inflation and output volatilities in the two regimes, the average autocorrelations of inflation and output. This subsection clarifies two moments that are less directly observable in the model: the value of internal funds and COLA pass-through.

Value of internal funds. Following Gilchrist et al. (2017), I discipline the external-financing friction using the average value of internal funds. In the model, one additional dollar of internal cash is worth one dollar when the firm has enough revenue to pay the wage bill. When the firm would otherwise need external funds, that dollar also saves the marginal external-financing cost. Evaluated ex ante in the steady state, the value of internal funds is therefore the probability-weighted marginal value of cash across idiosyncratic revenue states:

$$1 + \phi_0 \int_0^{\bar{w}\bar{L}/(\bar{\zeta}\bar{L}^c)} \left(\bar{w}\bar{L} - \bar{\zeta}\eta\bar{L}^c\right) dF_\eta(\eta)$$

This is the model counterpart of the empirical target. It is informative about how costly low-revenue states are for the firm, which is precisely the force that prevents optimal COLA pass-through from mechanically equaling one.

COLA pass-through. The model does not contain a primitive parameter called COLA pass-through. Instead, pass-through is an equilibrium object implied by the optimal COLA wage rule and the inference problem. Under COLA, nominal wages respond to current inflation both mechanically and because current inflation changes the inferred revenue factor. In the linear solution, the relevant wage response can be written as

$$\frac{\partial \Delta W_t^{\mathcal{C}}}{\partial \pi_t} = 1 + \mu b_t$$

where b_t is the projection coefficient from inflation to the forecast revision of the real revenue factor, and μ is the worker risk-exposure coefficient from the contracting block. In the code, the projection coefficient is computed from the Kalman-filter solution and averaged over contract horizons. The reported model moment is

$$PT^{model} = \frac{1}{2}\left(1 + \mu \bar{b}_{GI} + 1 + \mu \bar{b}_{GM}\right)$$

which corresponds to the average pass-through targeted in the data. This construction also implies that the pass-through moment is informative about both the firm-side risk exposure and the aggregate shock process.

G.3 Understanding the COLA Costs

The quantitative model assumes heterogeneous COLA costs to generate interior equilibrium COLA shares. Let F_χ denote the distribution of the cost of adopting a state-contingent wage rule. In the calibration, F_χ is logistic with location parameter χ_0 and scale parameter χ_1, with negative draws bunching to zero. For any perceived private gain from COLA $\mathcal{V}$, the mass of renewing contracts that adopts COLA is $F_\chi(\mathcal{V})$. The two cost parameters are chosen so that the observed average COLA shares in the Great Inflation and Great Moderation are fixed points of the stationary contracting equilibrium.

The gain $\mathcal{V}$ is measured in the utility units generated by the second-order approximation to the contracting problem. One first computes the loss constant from the local curvature of the worker-firm objective,

$$C = -\frac{1}{2}\bar{V}_{ww}\bar{w}^2$$

and then expresses the estimated utility cost as $C \times 0.0058$. The number 0.0058 is a cost in the normalized units of the second-order contracting loss. To convert it into consumption units, I divide by the worker's steady-state marginal utility and union employment:

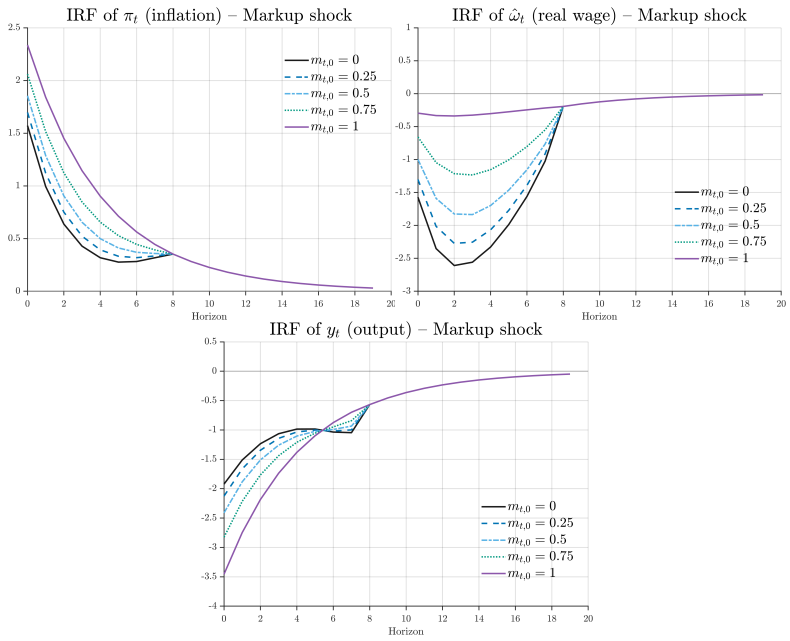
$$cost^{money} = \frac{C \times 0.0058}{u'(\bar{w})\bar{L}}$$

Expressing this number relative to steady-state labor income gives approximately two days of wages over a two-year contract, as stated in the main text. The small money-metric value is useful for interpretation: non-contingent contracts remain common not because COLA is prohibitively costly, but because the private insurance gain from COLA is small in low-inflation-risk environments.

G.4 Impulse Response Functions

Figure 31 plots the impulse response functions of inflation, the aggregate real wage index, and output to a unit markup shock in different conditional general equilibria with different levels of COLA shares.

Figure 31: IRFs of inflation, real wage, and output to a markup shock



*Notes:*The figure plots the impulse response functions (IRFs) of inflation, the aggregate real wage index, and output with respect to a unit markup shock in the stationary conditional general equilibrium with different levels of COLA shares.

G.5 Robustness for Great Inflation Counterfactuals

G.5.1 Robustness for Union COLA Share Extrapolation

The baseline counterfactual in the main text for the Great Inflation extrapolates the inflation-dependent wage setting measured in the union sector by explicit COLA clauses to the broader economy. This amounts to assuming that wage setting in the non-union sector has the same response as in the union sector, through either explicit or implicit indexation that we cannot easily observe. I provide robustness checks using alternative assumptions for extrapolations in this section.

I consider two alternative assumptions. In both cases, I fix the union share at 26 percent (Hirsch, Macpherson, and Even, 2026), and union COLA coverage rises from 15 percent to 50 percent during the Great Inflation. The most conservative assumption is that the non-union sector only uses non-contingent wages, so inflation-contingent wage setting appears only in the union sector. The other assumption uses evidence from the recent post-pandemic inflation episode to infer the response in the non-union sector. I show in Section A.6 that during the recent inflation episode, COLA coverage rose from 10 percent to 17 percent, while the share of workers receiving

off-cycle wage raises rose from 17 percent to 27 percent (Hurst et al., 2026). I assume the ratio of these changes is stable over time. Using a back-of-the-envelope calculation, I translate the 35 percentage point increase in union COLA coverage during the Great Inflation into a 50 percentage point increase in the share of non-union workers receiving off-cycle wage raises. I show in Section A.9 that most COLA in the union sector adjusts wages each quarter, while I do not observe the actual frequency of those off-cycle wage raises. To be more conservative, I assume there is only one off-cycle wage raise each year, indicating that wages in the non-union sector adjust twice a year. Thus, I halve this 50 percentage point increase in the non-union sector. This gives a 25 percentage point increase in the effective rate of indexation in the non-union sector. I refer to this extrapolation as the ADP-informed response.

Table 11 reports the resulting aggregate COLA shares in the respective stationary stochastic equilibria before and after the Great-Inflation volatility shock. In the summation, the union and the non-union sectors are weighted by their respective employment shares.

Table 11: Actual and counterfactual COLA share in different scenarios

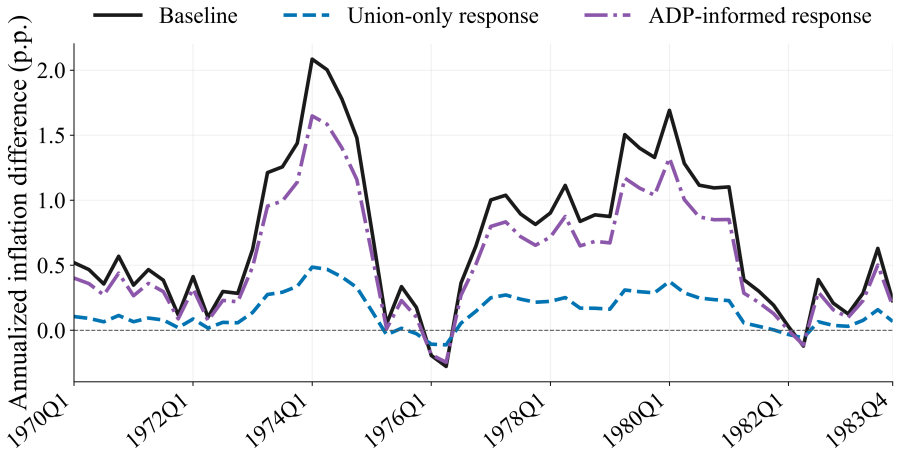
Scenario	Actual GI	Low-inflation counterfactual
	Stationary aggregate COLA share	
Baseline extrapolation	0.5000	0.1500
Union-only response	0.1300	0.0390
ADP-informed response	0.4408	0.1648

Holding the calibrated structural parameters fixed, I repeat the exercise in the baseline counterfactual in the main text under each extrapolation. Notice that for each extrapolation, I back out a different sequence of structural shocks, and then compare how the actual path of inflation and the counterfactual path differ. In other words, each extrapolation per se is not a counterfactual exercise. Instead, each extrapolation amounts to taking a particular stand on wage setting in the non-union sector, given the lack of direct measures. Different assumptions for extrapolation then yield different conclusions about inflation amplification.

Figure 32 plots the different levels of inflation amplification under different extrapolations. The ADP-informed response gives a very similar amplification path to the baseline extrapolation. This provides some reassurance for the quantitative results in the main text. Even under the most conservative assumption that the wage setting in the non-union sector never responds, the effect is still nontrivial: the union-only response raises annualized inflation by about 0.5 percentage points at those peaks.

Table 12 reports the average amplification during the 1973–1975 and 1979–1981 inflation spikes. Across the two episodes, the ADP-informed response raises average annualized inflation by 0.83 percentage points, compared with 1.06 percentage points in the baseline. Even the union-only response raises it by 0.22 percentage points.

Figure 32: The path of inflation amplification under alternative assumptions for extrapolation



G.5.2 Robustness for Firm Weight λ in the Contracting Problem

In my contracting problem, the firm's weight in the total value, λ, has a similar, but not identical, interpretation to the bargaining power in a Nash bargaining in the canonical search-and-match model. There, the bargaining power is notoriously difficult to identify because we often do not have a direct measure of it, and it confounds with other parameters like worker home production during unemployment. The same issue appears in my internal estimation if I include λ in the list of parameters to be calibrated. Therefore, I follow the standard practice in the literature and set λ to one in the baseline calibration. In other words, I assume unions and firms are equally weighted in the total value. I then check the robustness of the quantitative results under different values of λ.

In particular, I vary λ from 0.2 to 3 while keeping all other parameters at the baseline calibration. I do not recalibrate the model under different values of λ. Instead, for each value of λ, I recalculate some steady-state moments, including the ones that I target in the baseline calibration, as well as some important steady-state elasticities. I also redo the GI counterfactual, i.e., the inflation amplification effect of endogenous contract choice for each λ. One can think of the exercise here as a sensitivity check of model quantities with respect to λ. Holding all other parameters fixed makes the result comparable.

Table 13 reports some steady-state moments (Panel A) and some steady-state elasticities and dynamic moments (Panel B). Although λ changes by a factor of 15 from 0.2 to 3, the variation of these moments and elasticities is relatively small. The two elasticities ι_1 and ι_2, representing labor demand elasticities with respect to the revenue factor ζ and real wage w, have the largest variation: by a factor of three. Dynamic moments like inflation standard deviation and COLA share are relatively stable.

Table 12: Inflation amplification during the two spikes

Scenario	1973–1975	1979–1981	Both spikes
Baseline extrapolation	1.10	1.02	1.06
Union-only response	0.24	0.21	0.22
ADP-informed response	0.86	0.79	0.83

How do the changes of these moments and elasticities alter the quantitative results? Figure 33 plots the main quantitative results, namely the inflation amplification effect of endogenous contract choice during GI. The figure shows that changing λ from 0.2 to 3 almost has no impact on the amplification effect. This supports the interpretation that the main counterfactual is driven by the interaction between inflation risk, contract adoption, and inflation propagation, instead of a knife-edge choice of the firm weight.

Figure 33: Inflation amplification during GI under different values of λ



G.6 Great Moderation Counterfactuals

The "policy vs. luck" debate on the Great Moderation period concerns the relative contribution of the more responsive monetary policy and the exogenous decline in shock sizes to the more stabilized inflation and output. However, when contract choice is endogenous, to the extent that a more hawkish monetary policy also discourages COLA adoption and therefore shock propagation, the contribution of monetary policy may be underestimated. I study this channel from different angles.

G.6.1 Decomposing the Decline in On-Impact Inflation Response to Shocks

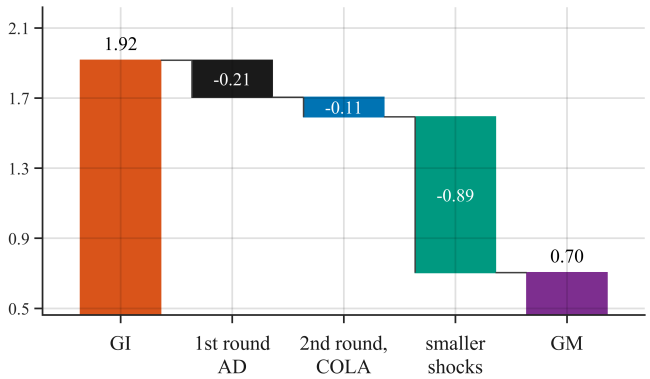
The main text reports the decomposition of the decline in inflation response to typical markup shocks from the Great Inflation to Great Moderation. I report other robustness exercises here.

Table 13: Sensitivity check for λ

	Panel A. Steady-state moments				
λ	Cash value	Labor share	Financing prob.	Financing cost	UI replacement
0.2	1.030	0.634	0.209	0.0019	0.804
0.5	1.037	0.633	0.206	0.0023	0.806
1.0	1.052	0.630	0.200	0.0032	0.810
2.0	1.096	0.622	0.183	0.0058	0.822
3.0	1.182	0.608	0.156	0.0104	0.845
	Panel B. Elasticities and dynamic diagnostics				
λ	μ	ι_1	ι_2	Std. π	Implied COLA
0.2	0.164	3.415	-3.469	0.979	0.513
0.5	0.164	3.172	-3.221	0.972	0.507
1.0	0.155	2.766	-2.813	0.963	0.502
2.0	0.136	1.946	-1.974	0.945	0.498
3.0	0.192	1.132	-1.183	0.924	0.481

One caveat is that there could be multiple COLA equilibria when we turn on COLA response to the policy change but fix shock volatility at the Great Inflation level. The main text reports the decomposition results based on the equilibrium in which the initial conjecture is that no one uses COLA after the policy change. As a robustness check, Figure 34 reports the decomposition based on the equilibrium in which the initial conjecture of the COLA path is fixed at the Great Inflation level. In that equilibrium, COLA share drops by less, and the contribution of monetary policy to inflation stabilization through the contract accelerator is also smaller. However, even in that equilibrium, the contribution of monetary policy increases by one half.

Figure 34: Decomposing the decline in on-impact inflation IRF to a one-std markup shock

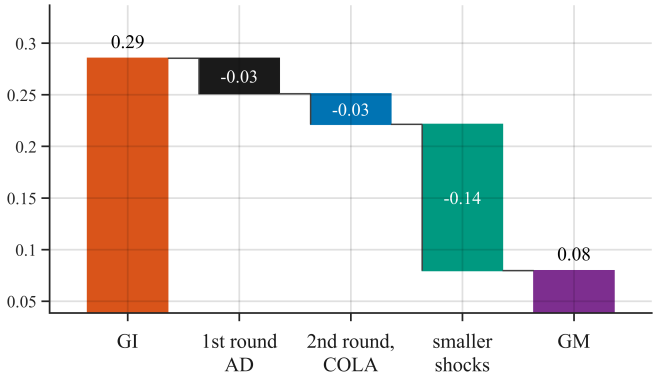


Notes: This figure shows the decomposition results for typical markup shocks based on the equilibrium in which the initial conjecture is that the COLA adoption rate is the same as the Great Inflation after the policy change.

Figures 35-37 show the decomposition results for typical demand shocks, again based on the two respective equilibria. Demand shocks are overall less inflationary than markup shocks. But

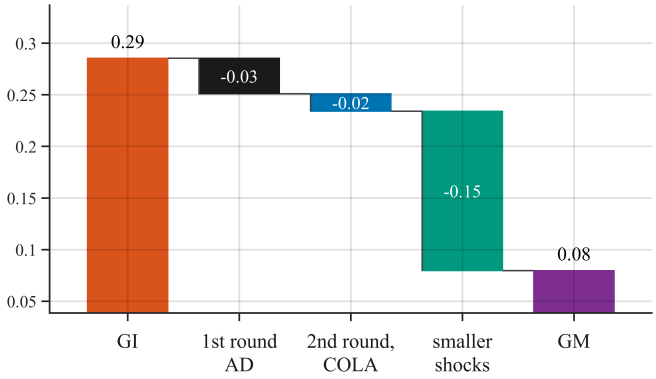
the relative magnitudes of the contribution to inflation stabilization from different channels are roughly comparable to those of markup shocks.

Figure 35: Decomposing the decline in on-impact inflation IRF to a one-std demand shock



Notes: This figure shows the decomposition results for typical demand shocks based on the equilibrium in which the initial conjecture is that no one uses COLA after the policy change.

Figure 36: Decomposing the decline in on-impact inflation IRF to a one-std demand shock



Notes: This figure shows the decomposition results for typical demand shocks based on the equilibrium in which the initial conjecture is that the COLA adoption rate is the same as the Great Inflation after the policy change.

G.6.2 Decomposing Inflation Volatility

The main text emphasizes the decomposition of inflation responses to shocks. Table 14 reports the analogous decomposition for unconditional volatilities. The columns compare the Great Inflation economy, the Great Moderation economy, a counterfactual with the more hawkish Great Moderation monetary-policy rule but Great Inflation shock volatilities and fixed contract composition, and two endogenous-contract counterfactuals in which monetary policy also changes the equilibrium COLA share.

The model-implied annualized inflation volatility falls from 0.92 in the Great Inflation to 0.38 in the Great Moderation. Moving only to the more hawkish monetary-policy rule lowers inflation volatility from 0.92 to 0.84, so the conventional aggregate-demand channel accounts for about 15 percent of the total decline. Allowing contract choice to respond lowers inflation volatility

further, to 0.80 or 0.76 depending on the belief/conjecture used to initialize the contracting fixed point. Thus the policy contribution rises to roughly 25 percent once the contract-choice channel is active. The remaining decline is mainly due to the lower primitive-shock volatilities in the Great Moderation.

This result is consistent with the stronger role of policy in the impulse-response decomposition. The response exercise asks how much a given shock is propagated. The volatility exercise also depends on the size of shocks. Even if a hawkish rule and the induced decline in COLA substantially weaken propagation, unconditional inflation volatility remains high when the economy is still fed Great Inflation shock variances. This is why policy explains a large share of the decline in inflation responses, while lower primitive-shock volatility still explains most of the decline in unconditional inflation volatility.

The other rows show that this decomposition has real-side implications. Output and employment volatility fall from 1.96 to 1.02 across regimes, but they would rise under the more hawkish rule if shock volatilities and the high-COLA contract distribution stayed at Great Inflation levels. During the Great Inflation, inflation is heavily driven by markup shocks, and a high COLA share makes the output-inflation tradeoff sharper. A more hawkish rule therefore stabilizes inflation partly by moving output more. The real-wage row moves in the opposite direction once contract choice is active: lower COLA adoption weakens aggregate inflation propagation, but it also leaves more workers exposed to real-wage risk.

Table 14: Decomposition of unconditional volatility

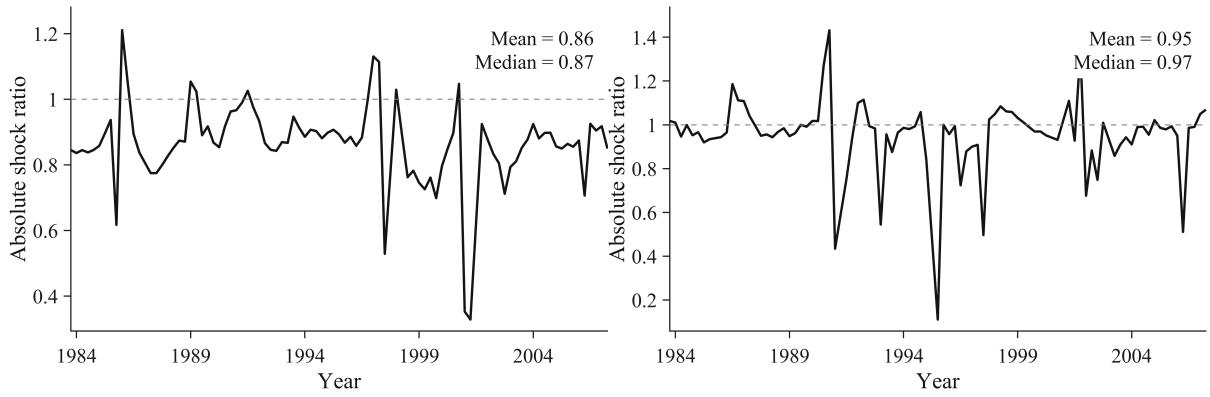
	GI	GM	AD	COLA min	COLA max
π_t	0.92	0.38	0.84	0.80	0.76
y_t, l_t	1.96	1.02	2.42	2.39	2.36
ω_t	1.13	0.67	1.06	1.18	1.26

G.6.3 Recovering First-Order Shocks with and without COLA Response

From the observed time path of inflation and output, the model can uniquely recover the underlying realized aggregate shock sequence, namely the sequence of markup and demand shocks. Compared to the economy in which COLA adoption endogenously shifts downwards in response to the more hawkish monetary policy rule, the counterfactual economy in which COLA does not respond requires smaller shocks to account for the observed path of inflation and output. In other words, the counterfactual economy needs more luck to explain the more tranquil macroeconomic dynamics. Figure 37 illustrates the gap for the required shocks in the two economies. In particular, it plots the ratio of the markup shock and the demand shock when COLA does not respond to those when COLA responds. Quantitatively, without endogenous COLA response, the economy

requires the realized markup shocks to be 14 percent smaller on average, and the realized demand shocks to be 5 percent smaller on average, in order to explain the more stabilized inflation and output.

Figure 37: The ratio of the required markup (left) and demand (right) shock without and with COLA response



H Welfare Analysis and Policy Experiments

COLA improves private risk sharing but can raise aggregate inflation volatility by strengthening the wage-price feedback loop. The welfare effect is therefore ambiguous and depends on the relative importance of the two forces. I discuss this tension and use the Great-Inflation calibration as a laboratory for policy analysis.

H.1 Deriving the Social Welfare Function

This subsection derives the welfare object used in the quantitative exercises. The planner is utilitarian over workers and assigns Pareto weight Λ_e to the representative entrepreneur. COLA adoption costs are utility costs paid by workers, while external financing costs and price dispersion are resource costs.

For a firm-union pair with real revenue factor ζ_t and real wage w_{it}, let $L^*(\zeta_t, w_{it})$ denote the firm's optimal employment choice. Ignoring the additive constant from non-employed workers, the worker-side payoff from this pair is

$$U(\zeta_t, w_{it}) = L^*(\zeta_t, w_{it})[u(w_{it}) - u(b)]$$

The aggregate worker payoff is the contract-share-weighted sum of this object across contract cohorts.

The period social payoff is therefore

$$\mathcal{W}_t = \sum_{k=1}^{N} \left[\mathcal{M}_{t,k} U(\zeta_t, w_{t,k}^{\mathcal{C}}) + \left(\frac{1}{N} - \mathcal{M}_{t,k} \right) U(\zeta_t, w_{t,k}^{\mathcal{N}}) \right] + \Lambda_e u(C_t^{\text{ent}}) - \Xi_t$$

Here Ξ_t is the aggregate utility cost of COLA adoption among newly renewing contracts. If a renewing firm-union pair adopts COLA whenever its cost χ is below the private gain $\mathcal{V}_t$, then

$$\Xi_t = \frac{1}{N} \int_0^{\mathcal{V}_t} \chi dF_\chi(\chi)$$

This term is absent from the resource constraint because it captures cognitive calculation costs in utility units.

I next expand the worker payoff around the nonstochastic steady state. Let hats denote log deviations, define $\bar{U} = u(\bar{w}) - u(b)$, and write derivatives of $L^*(\zeta, w)$ at $(\bar{\zeta}, \bar{w})$ as L_ζ, L_w, $L_{\zeta\zeta}$, $L_{\zeta w}$, and L_{ww}. Since the expansion is in log deviations, $w_{it} = \bar{w} \exp(\hat{w}_{it})$ and $\zeta_t = \bar{\zeta} \exp(\hat{\zeta}_t)$. Keeping terms up to second order gives

$$E[U(\zeta_t, w_{it}) - U(\bar{\zeta}, \bar{w})] = A_3 \operatorname{Var}(\hat{\zeta}_t) + A_4 \operatorname{Cov}(\hat{\zeta}_t, \hat{w}_{it}) + A_5 \operatorname{Var}(\hat{w}_{it})$$

where

$$A_3 = \frac{1}{2} \left(\bar{\zeta}^2 \bar{U} L_{\zeta\zeta} + \bar{\zeta} \bar{U} L_\zeta \right)$$

$$A_4 = \left(\bar{U} L_{\zeta w} + L_\zeta u'(\bar{w}) \right) \bar{\zeta} \bar{w}$$

and

$$A_5 = \frac{1}{2} \left[\left(\bar{U} L_{ww} + 2 L_w u'(\bar{w}) + \bar{L} u''(\bar{w}) \right) \bar{w}^2 + \left(\bar{U} L_w + \bar{L} u'(\bar{w}) \right) \bar{w} \right]$$

The terms involving first derivatives arise from the curvature of the exponential mapping from log deviations to levels.

Aggregate worker welfare can then be written in terms of two sufficient second moments for a given contract distribution. The first is the covariance between revenue and the aggregate real wage index, where

$$\hat{\omega}_t = \sum_{k=1}^{N} \left[\mathcal{M}_{t,k} \hat{w}_{t,k}^{\mathcal{C}} + \left(\frac{1}{N} - \mathcal{M}_{t,k} \right) \hat{w}_{t,k}^{\mathcal{N}} \right]$$

The second is the contract-share-weighted wage variance,

$$\mathcal{S}_t^w = \sum_{k=1}^{N} \left[\mathcal{M}_{t,k} \operatorname{Var}(\hat{w}_{t,k}^{\mathcal{C}}) + \left(\frac{1}{N} - \mathcal{M}_{t,k} \right) \operatorname{Var}(\hat{w}_{t,k}^{\mathcal{N}}) \right]$$

Thus, after omitting the steady-state constant,

$$\Delta W_t^W = A_3 \operatorname{Var}(\hat{\zeta}_t) + A_4 \operatorname{Cov}(\hat{\zeta}_t, \hat{\omega}_t) + A_5 \mathcal{S}_t^w$$

This expression captures the worker-side insurance trade-off. COLA changes individual wage risk through $\mathcal{S}_t^w$, but widespread COLA also changes the general-equilibrium process for inflation faced by those who do not adopt COLA.

Entrepreneur welfare has two components. First, entrepreneur consumption fluctuates around steady state. With $C_t^{\text{ent}} = \bar{C}^{\text{ent}} \exp(\hat{c}_t^{\text{ent}})$, its second-order contribution is

$$\Delta W_t^{E,disp} = \frac{1}{2}\left[u''(\bar{C}^{\text{ent}})(\bar{C}^{\text{ent}})^2 + u'(\bar{C}^{\text{ent}})\bar{C}^{\text{ent}}\right] \text{Var}(\hat{c}_t^{\text{ent}})$$

Second, wage dispersion and price dispersion use up resources due to resource misallocation.[50] To see the wage-dispersion term, notice that cost minimization implies that producing a given amount of the aggregate service input requires an input bundle multiplied by a dispersion index. Around the symmetric steady state, if $\hat{q}_{it}^x$ denotes the log relative service price, this index satisfies

$$\Delta_t^x - 1 = \frac{\varepsilon_x}{2} \text{Var}_i(\hat{q}_{it}^x) + o(\|\hat{q}_t^x\|^2)$$

In the model, service prices move one-for-one with the relevant real wage to first order. Thus the resource loss from wage dispersion is governed by the cross-sectional variance of relative real wages:

$$D_t^w = \frac{\varepsilon_x}{2} \text{Var}_i(\hat{w}_{it} - \hat{\omega}_t)$$

The object $\mathcal{S}_t^w$ before captures wage risk relative to the nonstochastic steady state faced by workers, while the resource wedge here depends on relative wages across firms at a point in time because relative input prices distort the allocation of labor services across producers.

The price-dispersion term follows the same logic from the final-goods CES aggregator under Calvo pricing. Let Δ_t^p be the standard price-dispersion index multiplying the amount of intermediate goods needed to produce final output. Log-linearizing the Calvo law of motion around zero steady-state inflation and retaining second-order terms gives

$$E(\Delta_t^p - 1) = \Psi_p \,\text{Var}(\pi_t), \qquad \Psi_p \equiv \frac{\varepsilon_y \theta}{2} \frac{\varepsilon_y - 1 + 2\theta}{(1-\theta)^2}$$

Therefore the total output-equivalent resource loss, as a fraction of steady-state final output, is

$$D_t^Y = \frac{\varepsilon_x}{2} \text{Var}_i(\hat{w}_{it} - \hat{\omega}_t) + \Psi_p \,\text{Var}(\pi_t)$$

If this residual resource loss is borne by entrepreneurs, the utility loss is

$$\Delta W_t^{E,mis} = -u'(\bar{C}^{\text{ent}})\bar{Y} D_t^Y$$

Equivalently, if the same resource loss is expressed as a fraction of entrepreneur consumption, $D_t^C = (\bar{Y}/\bar{C}^{\text{ent}}) D_t^Y$, then the loss is $-u'(\bar{C}^{\text{ent}})\bar{C}^{\text{ent}} D_t^C$.

[50] Because firms have linear technology, resource misallocation comes from the CES aggregator.

Combining the worker, entrepreneur, and COLA-cost terms gives the second-order social welfare criterion

$$E[\mathcal{W}_t - \bar{\mathcal{W}}] = \Delta W_t^W + \Lambda_e \left(\Delta W_t^{E,disp} + \Delta W_t^{E,mis} \right) - \Xi_t$$

A rule-based planner or central bank choosing a policy vector Θ solves

$$\max_{\Theta} E \sum_{t=0}^{\infty} \beta^t \left[\Delta W_t^W(\Theta) + \Lambda_e \left(\Delta W_t^{E,disp}(\Theta) + \Delta W_t^{E,mis}(\Theta) \right) - \Xi_t(\Theta) \right]$$

subject to private contract choice, wage rules conditional on contract type, and the aggregate equilibrium conditions. The policy vector may parameterize a cap on COLA pass-through, a simple Taylor rule, or a Taylor rule augmented with the contract distribution.

Different from a standard New Keynesian model, where inflation volatility only generates losses but does not change any transmission mechanism, inflation volatility affects the demand for COLA and alters inflation transmission in equilibrium, which feeds back into social losses.

H.1.1 The One-Period i.i.d. Case

The trade-off between micro insurance and macro volatility becomes more transparent in the analytical case with i.i.d. shocks and one-period contracts. In that case there is only one contracting cohort, so denote the COLA share by m. The two real wage rules are

$$\hat{w}_t^{\mathcal{C}} = \delta \pi_t, \qquad \hat{w}_t^{\mathcal{N}} = -\pi_t$$

where $1 + \delta$ is the price-to-nominal-wage pass-through under COLA. In equilibrium, δ is itself determined by the risk-sharing condition and by the aggregate stochastic environment. Hence the aggregate real wage index is

$$\hat{\omega}_t = \mathcal{K}(m)\pi_t, \qquad \mathcal{K}(m) \equiv m\delta(m) - (1-m) = (1 + \delta(m))m - 1$$

This is exactly the statistic that enters marginal cost and the state-dependent New Keynesian Phillips curve.

For any given m, substituting these wage rules into worker welfare gives

$$\Delta W^W(m) = A_3 \operatorname{Var}(\hat{\zeta}_t; m) + A_4 \mathcal{K}(m) \operatorname{Cov}(\hat{\zeta}_t, \pi_t; m) + A_5 \left[m\delta(m)^2 + (1-m) \right] \operatorname{Var}(\pi_t; m)$$

Equivalently, since the optimal COLA risk-sharing rule satisfies $\delta(m) = \mu \operatorname{Cov}(\hat{\zeta}_t, \pi_t; m) / \operatorname{Var}(\pi_t; m)$, this can be written as

$$\Delta W^W(m) = A_3 \operatorname{Var}(\hat{\zeta}_t; m) + \left[\frac{A_4}{\mu} \mathcal{K}(m)\delta(m) + A_5 \left[m\delta(m) + (1-m) \right] \right] \operatorname{Var}(\pi_t; m)$$

The last term displays the private insurance force. Holding the inflation process fixed, replacing non-contingent wages with COLA changes the worker wage-risk term from $\operatorname{Var}(\pi_t)$ to $\delta^2 \operatorname{Var}(\pi_t)$

for the marginal adopting worker. When $|\delta| < 1$, COLA lowers individual real-wage exposure to inflation risk. This is the micro-insurance benefit.

The same substitution gives a different object for the resource wedge. Cross-sectional wage dispersion is

$$\mathrm{Var}_i(\hat{w}_{it} - \hat{\omega}_t; m) = m(1-m)(1+\delta(m))^2 \,\mathrm{Var}(\pi_t; m)$$

To derive this, note that a COLA wage deviates from the aggregate wage by $[\delta - K(m,\delta)]\pi_t = (1-m)(1+\delta)\pi_t$, while a non-contingent wage deviates by $[-1 - K(m,\delta)]\pi_t = -m(1+\delta)\pi_t$. Weighting these squared deviations by the two contract shares yields the expression above. Therefore, in the one-period case,

$$D^Y(m) = \left[\frac{\varepsilon_x}{2} m(1-m)(1+\delta(m))^2 + \Psi_p\right] \mathrm{Var}(\pi_t; m)$$

Putting these pieces together, the one-period social welfare criterion can be summarized as

$$\begin{aligned}\mathbb{E}\left[\mathcal{W}_t(m) - \bar{\mathcal{W}}\right] = &\underbrace{A_3 \,\mathrm{Var}(\hat{\zeta}_t; m) + B_W(m, \delta(m)) \,\mathrm{Var}(\pi_t; m) + \Lambda_e B_e \,\mathrm{Var}(\hat{c}_t^{\text{ent}}; m)}_{\text{loss from consumption volatility}} \\ &\underbrace{-\Lambda_e \Psi_{wp}(m, \delta(m)) \,\mathrm{Var}(\pi_t; m)}_{\text{loss from misallocation due to wage and price dispersion}} \\ &\underbrace{-\Xi_t}_{\text{direct COLA cost}}\end{aligned}$$

where

$$\begin{aligned} B_W(m,\delta) &\equiv \frac{A_4}{\mu} K(m,\delta)\delta + A_5\left[m\delta^2 + (1-m)\right] \\ B_e &\equiv \frac{1}{2}\left[u''(\bar{C}^{\text{ent}})(\bar{C}^{\text{ent}})^2 + u'(\bar{C}^{\text{ent}})\bar{C}^{\text{ent}}\right] \\ \Psi_{wp}(m,\delta) &\equiv u'(\bar{C}^{\text{ent}})\bar{Y}\left[\frac{\varepsilon_x}{2} m(1-m)(1+\delta)^2 + \Psi_p\right] \end{aligned}$$

This is the closest analogue to the standard New Keynesian welfare expression. The micro vs. macro tradeoff of COLA adoption can be seen by differentiating the social welfare function with respect to the COLA share m:

$$\frac{d\mathbb{E}\left[\mathcal{W}_t(m)-\bar{\mathcal{W}}\right]}{dm} = \left[\underbrace{|A_5|(1-\delta(m))+\frac{A_4}{\mu}(1+\delta(m))\delta(m)}_{\text{micro insurance gain from more COLA}} - \Lambda_e \frac{d\Psi_{wp}(m,\delta(m))}{dm}\right] \text{Var}(\pi_t;m)$$
$$\underbrace{-\left|A_3+B_W(m,\delta(m))-\Lambda_e\Psi_{wp}(m,\delta(m))\right|\frac{d\,\text{Var}(\pi_t;m)}{dm}}_{\text{macro inflation volatility externality from more COLA}} + \Lambda_e B_e \frac{d\,\text{Var}(\hat{c}_t^{\text{ent}};m)}{dm}$$
$$+\frac{\partial\mathbb{E}[\mathcal{W}_t]}{\partial\delta}\frac{d\delta}{dm}-\frac{d\Xi_t}{dm}$$

The first term in the bracket on the first line of the RHS captures the micro-level direct mutual insurance gain from COLA when the contract of a marginal worker changes from non-contingent wages to COLA. Under plausible parameter values (including the calibrated parameter values), this term dominates, and the first line is positive.[51] The second line on the RHS captures the macro inflation volatility externality generated by a higher COLA share. I have shown that

$$\frac{d\,\text{Var}(\pi_t;m)}{dm} > 0$$

Under the calibrated parameter values, the second line is negative, meaning that a higher COLA share raises the equilibrium variance of inflation, which worsens the price-dispersion loss, the wage-dispersion loss, and the welfare of workers under non-contingent contracts whose real wages move one-for-one against inflation.

The first term on the third line captures the effect from the intensive margin adjustment. A change in m will also induce a change in the pass-through parameter δ. Although δ is optimally chosen privately, it will affect the aggregate pass-through and change inflation transmission. This macro effect is not internalized by private agents. Therefore, a standard envelope-type argument does not apply. In practice, however, the variation in the COLA pass-through is small, and the size of this term is limited.

H.2 COLA and Social Welfare During the Great Inflation

We now calculate the welfare levels for workers and the entrepreneur numerically. For all the subsequent analyses and policy experiments, I use the calibration from the Great Inflation period during which COLA is especially salient. To make the welfare numbers more interpretable, I

[51] The second term measures the changes in the revenue-inflation adjustment to the extent that inflation signals revenue. This term is small under the calibrated parameter values compared to the direct insurance gain. The third term captures the changes in wage dispersion as m changes. $m = \frac{1}{2}$ generates the largest wage dispersion. This term is also relatively small.

convert the changes in utility relative to the steady state to consumption-equivalent values (CEV). In particular, I first divide the change in utils ΔV by the corresponding steady-state marginal utility $u'(\bar{C})$ of the two groups to get the change in consumption in physical units; I then divide it by the corresponding consumption level $\bar{C}$ in the steady state and multiply it by 100 to further express it in percentage terms. In addition, since the transitional dynamics of the COLA response are relatively short, as shown in the GI/GM exercise, I only compare different stationary stochastic equilibria in what follows.

Table 15 reports five different cases: pre-GI shocks with no COLA, pre-GI shocks with equilibrium COLA, GI shocks with no COLA, GI shocks with pre-GI COLA, and GI shocks with equilibrium COLA. It not only isolates the general-equilibrium effect of higher COLA adoption conditional on Great Inflation shocks, but also shows whether COLA is desirable relative to an economy with no COLA, and separates the direct welfare cost of larger shocks from the amplification caused by the endogenous contract distribution.

The comparison shows that no COLA delivers the highest welfare in both shock environments. Under pre-GI shocks, moving from no COLA to equilibrium COLA (15 percent) lowers average union welfare from -0.032 to -0.035 and entrepreneur welfare from -0.212 to -0.227. Under GI shocks, moving from no COLA to the pre-GI COLA share and then to the GI equilibrium COLA share (50 percent) further lowers average union welfare from -0.081 to -0.090 and -0.099, and entrepreneur welfare from -0.513 to -0.554 and -0.670. Thus, the direct increase in primitive shock volatility explains a large part of the welfare deterioration, but COLA adoption adds a separate amplification margin. COLA privately insures covered workers, yet its aggregate effect raises inflation and real-wage risk for workers under non-contingent contracts and increases entrepreneurial-consumption risk and misallocation losses; in the calibrated economy, this externality is large enough that zero COLA is welfare-superior.

Table 15: Welfare effects of endogenous contract choice during the Great Inflation

	pre-GI shock		GI shock		
	no COLA	eqm COLA	no COLA	pre-GI COLA	eqm COLA
Individual union if COLA	0.011	0.010	0.029	0.025	0.018
Individual union if scheduled	-0.031	-0.045	-0.079	-0.114	-0.211
Average union	-0.032	-0.035	-0.081	-0.090	-0.099
Entrepreneur	-0.212	-0.227	-0.513	-0.554	-0.670

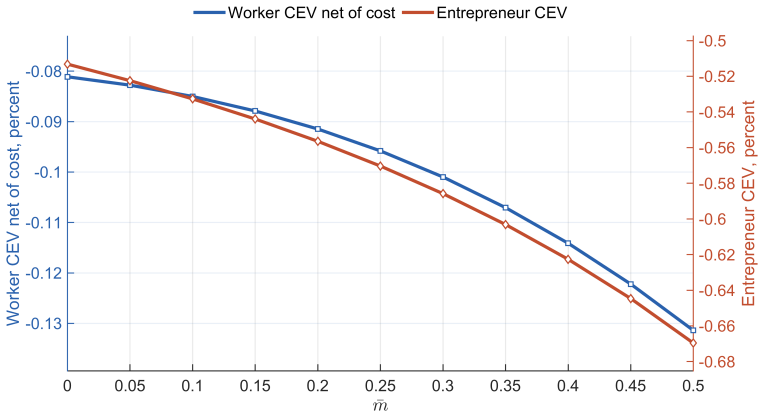
Notes: Entries are consumption-equivalent values in percent of steady-state consumption. The table excludes direct COLA adoption costs. "eqm COLA" in the "pre-GI shock" period is set to 15 percent, and that in the "GI shock" period is set to 50 percent.

H.3 COLA Rationing

The most straightforward policy experiment in response to the welfare calculation in the last subsection is COLA rationing. The government can set a cap $\bar{m}$ on COLA. In every period, firm-union

pairs can apply to include COLA clauses in their contracts. The government grants COLA randomly to those applicants if the cap has not been met. Figure 38 plots the consumption-equivalent value for workers and the entrepreneur under different COLA caps $\bar{m}$, where the worker value is the net value after COLA costs. For all levels of $\bar{m}$ from 0 to 0.5, the cap is binding in the stationary stochastic equilibrium. A striking feature of this figure is that the values of both groups drop when $\bar{m}$ increases. This implies that no matter what the Pareto weights are, the best policy is always to set the COLA cap $\bar{m}$ to zero, or banning COLA.

Figure 38: CEV effects under different levels of COLA rationing



Notes: placeholder

H.4 Compulsory COLA Adoption

The above exercise rations COLA, and the optimal policy turns out to be a COLA ban. One may wonder what if the policy goes the other direction. Namely, is it welfare-improving if COLA becomes compulsory? The rationale is that if everyone is covered by COLA, the macro externality does not matter anymore, at least for workers. However, those pairs who do not adopt COLA in the decentralized equilibrium must face an adoption cost that is higher than the gain from COLA. If the government forces them to adopt, it is not privately optimal. Moreover, compulsory adoption raises inflation volatility, which makes the representative entrepreneur worse off. Both margins imply that compulsory adoption cannot be welfare improving. To validate this argument, the next figure plots the expected value for an average worker under different COLA coverage, before and after the adoption cost is considered. Instead of random assignment, this figure assumes that the government can observe the adoption cost, so those with lower adoption costs adopt first.[52]

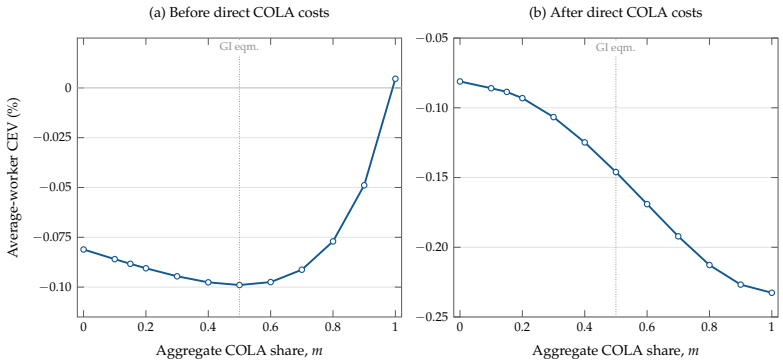
The figure is consistent with the reasoning. Without the adoption cost, the average worker starts to become better off as the COLA coverage increases beyond the equilibrium point. When all workers are covered by COLA, inflation risk becomes the least relevant. However, the picture is very different after the adoption cost is taken into account. Those who do not adopt in the decentralized equilibrium are exactly those who face high adoption costs. After paying the costs,

[52]This assumption is not essential.

the welfare of an average worker is monotonically decreasing with the COLA share, let alone the representative entrepreneur. Therefore, the optimal COLA share is indeed zero COLA for the Great-Inflation calibration.

This exercise also implies the nature and the heterogeneity of the adoption cost are important. In the main text, I interpret it as a cognitive calculation cost and calibrate it from the observed COLA share under some parametric assumptions. It may also come from administrative costs, bargaining costs, or even represent some deeper contractual frictions. A more direct measurement will be important for policymaking, which is left for future research.

Figure 39: CEV effects under different levels of compulsory COLA adoption



Notes: placeholder

H.5 Taylor Rule

The last few subsections discuss direct policy restrictions on COLA itself. This section discusses the optimal monetary policy. This discussion can be useful if direct restrictions are not feasible due, for example, to political reasons or implementability.[53] In particular, I discuss the optimal hawkishness of the Taylor rule. As I have shown in the theory and quantitative part, a more hawkish policy rule can reduce COLA adoption, which further stabilizes inflation. But they do not tell us how hawkish the Taylor rule should be. I answer this question here.

Compared to direct COLA restriction, the transmission of monetary policy is more complicated. I first clarify them briefly here, which also helps to understand the quantitative results. Let ϕ_π denote the responsiveness to inflation of the Taylor-type policy rule. Because adoption is strategic, policy generally induces an equilibrium correspondence rather than a single function:

$$\mathcal{E}(\phi_\pi, \sigma) = \{m : m = F_\chi(\mathcal{V}(m; \phi_\pi, \sigma))\},$$

[53] This is especially true for implicit inflation-contingency, which is not easy to track and regulate.

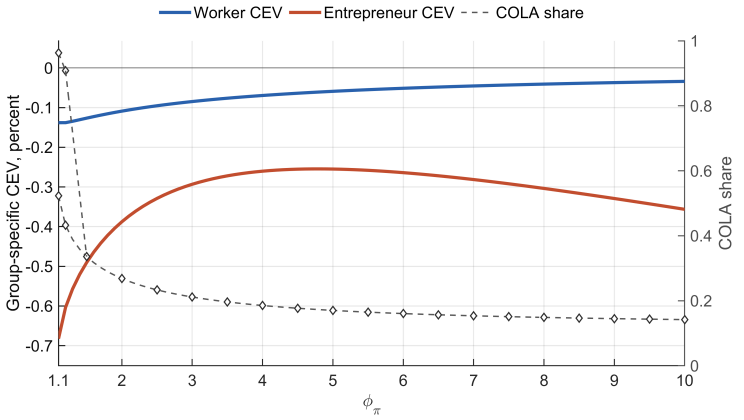
where σ denotes the primitive shock volatility. Monetary policy therefore has two distinct roles. First, it changes the equilibrium set itself and can act as an equilibrium-selection device by eliminating or selecting high-COLA fixed points. Second, conditional on a locally differentiable selected branch $m_j(\phi_\pi) \in \mathcal{E}(\phi_\pi, \sigma)$, it changes welfare along that branch according to

$$\frac{d\mathcal{W}_j}{d\phi_\pi} = \frac{\partial \mathcal{W}_j}{\partial \phi_\pi} + \frac{\partial \mathcal{W}_j}{\partial m_j} \frac{dm_j}{d\phi_\pi}$$

The first term is the direct stabilization effect familiar from New Keynesian policy analysis. For small values of ϕ_π, this term is typically positive, as inflation is stabilized, and consumption dispersion and output loss are both smaller. When ϕ_π is large, this term can be negative, especially when inflation is driven by cost-push shocks. In that case, an even higher ϕ_π shifts inflation loss to output loss. The second term is the contract-composition channel: monetary policy changes the private value of COLA and hence the equilibrium allocation of wage risk. The tradeoff between micro insurance and macro externality is embodied in the second term when a more hawkish rule reduces COLA share.

Figure 40 plots the CEV for workers (net of COLA costs) and the entrepreneur under different values of ϕ_π. It also plots the equilibrium COLA share. For small values of ϕ_π, multiple equilibria appear, and the reported CEV comes from the low-COLA equilibrium. Workers always benefit from a more hawkish Taylor rule. This is consistent with our COLA rationing experiment: an average worker (or all workers as a whole) wants zero COLA, and a more hawkish rule reduces COLA share. The CEV for the entrepreneur is not monotonic. It increases at first, as inflation is stabilized. However, it starts to drop when ϕ_π keeps increasing. This is exactly the inflation-output tradeoff we discussed before: when ϕ_π is large, the marginal gain from inflation reduction is small, but the marginal loss of output is large, and the entrepreneur is the residual claimer after paying relatively stable wages to workers.

Figure 40: CEV effects under different levels of Taylor rule hawkishness

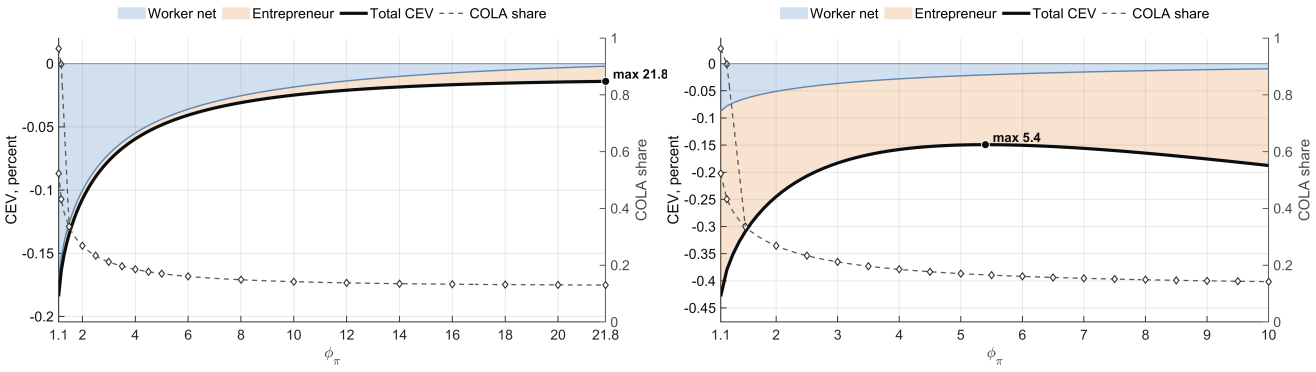


Notes: placeholder

Although the group-specific CEV is clear, the optimal policy based on them depends heavily on the Pareto weights in the social welfare function. If one assumes equal Pareto weights and transforms the SWF from the utility space to the CEV space, the weight on the entrepreneur will be extremely small. This is because in the calibrated steady state, the entrepreneur is much richer than an individual worker.[54] So, the marginal value of the entrepreneur is much lower. Therefore, if different individuals are equally weighted, the change in social welfare function will be driven by agents with high marginal utilities, which are workers in this case. Alternatively, one can assume the planner values a 1% increase in consumption the same for workers and the entrepreneur. In other words, the planner is attaching equal weights to the two groups in the CEV space. This will make the contribution of workers and the entrepreneur to the SWF closer.

Figure 41 plots the respective contribution of workers and the entrepreneur to the SWF for different ϕ_π when they have equal Pareto weights (left) and when they are equally weighted in the CEV space (right). The shaded area is their original CEV times their weights in the CEV space. For the equal-Pareto-weight case, the CEV weight for the entrepreneur is only 0.02. Therefore, workers' CEV drives the results, and they want the central bank to stabilize inflation as much as possible. For the equal-CEV-weight case, the entrepreneur matters more, and the optimal hawkishness equalizes the marginal gain of workers from less inflation risk and the marginal loss of the entrepreneur from more output risk.

Figure 41: Optimal hawkishness under different social welfare weights



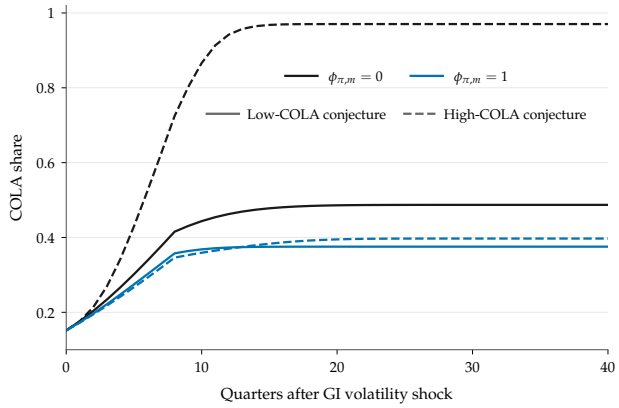
Notes: placeholder

H.6 COLA-Augmented Taylor Rule

The theory suggests that a more hawkish Taylor rule is needed to anchor the wage-setting institution when uncertainty rises and COLA becomes more attractive. While uncertainty is not easy to measure, the contract distribution, or the even coarser aggregate COLA share, is more tangible. This motivates a COLA-augmented Taylor rule: the policy rule automatically becomes

[54] Although the labor share is about 60%, the average wage is low relative to profit.

Figure 42: Anchoring the contract regime with a COLA-augmented Taylor rule



Notes: The figure plots equilibrium COLA paths following the Great Inflation volatility shock from low- and high-COLA initial conjectures. Black lines correspond to the standard Taylor rule ($\phi_{\pi 0} = 1.1, \phi_{\pi m} = 0$). Blue lines correspond to the COLA-augmented rule with$\phi_{\pi 0} = 1.1$ and$\phi_{\pi m} = 1$.

more hawkish if COLA share in this economy goes up. This class of policy rule does not become unconditionally hawkish but is more targeted to the wage-setting regime. In particular, I consider

$$i_t = (\phi_{\pi 0} + \phi_{\pi m} m_t)\, \pi_t + \phi_y \hat{y}_t$$

I use the calibration from the Great Inflation period and set $\phi_{\pi 0} = \phi_{\pi,GI} = 1.1$. Then, I vary the slope $\phi_{\pi m}$. $\phi_{\pi m} = 0$ reduces to the baseline model.

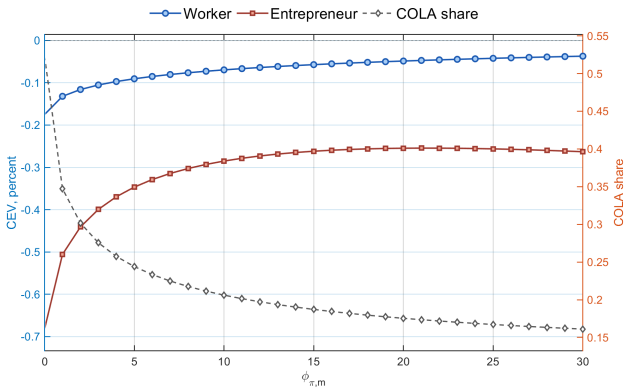
Before discussing its welfare implications, Figure 42 illustrates the effect of this rule on the equilibrium COLA path after the volatility shock at the onset of the Great Inflation. The two paths where $\phi_{\pi m} = 0$ replicate the actual path and the high-COLA path in Figure 9, which are solved from two different initial conjectures about future COLA adoption rates. They lead to two sharply different COLA shares: 50 percent and 96 percent, respectively. When $\phi_{\pi m} = 1$, the paths generated by both initial conjectures converge to the same COLA share of about 36 percent, up to small numerical error. Hence, the augmented rule eliminates the multiple dynamic paths and anchors the economy in a lower-COLA regime.

I next compare the welfare implications for the COLA-augmented Taylor rule with those of the standard Taylor rule in stationary stochastic equilibrium.

H.6.1 Welfare Comparison

Figure 43 plots the group-specific CEV for workers and the entrepreneur under different values of $\phi_{\pi m}$. The general pattern is similar to the standard Taylor rule: as the central bank becomes less tolerant to COLA share, welfare for both groups increases. However, when $\phi_{\pi m}$ becomes sufficiently large, the entrepreneur starts to get hurt, while an average worker still benefits from less inflation risk.

Figure 43: Group-specific CEV for different values of $\phi_{\pi m}$



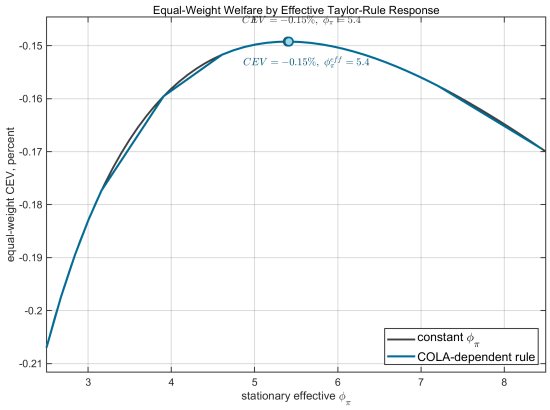
Notes: placeholder

For any given welfare weights, we can also search for the optimal $\phi_{\pi m}$ that maximizes social welfare. Denote by ϕ_π^{eff} the effective policy rule hawkishness in equilibrium, namely

$$\phi_\pi^{eff} = \phi_{\pi 0} + \phi_{\pi m} m^*$$

where m^* is the equilibrium COLA share. It turns out that the effective hawkishness induced by the optimal $\phi_{\pi m}$ is the same as the optimal hawkishness from the standard Taylor rule. Figure 44 illustrates the comparison under equal-CEV weights.

Figure 44: Welfare from the standard and COLA-augmented Taylor rule



Notes: placeholder

I Model Extensions

This final section collects material that is conceptually useful but secondary to the main argument. I discuss implicit contingency more formally, showing that to some extent, it is isomorphic to explicit COLA formulas. I also study the extension of endogenous stochastic inflation volatility.

I.1 Implicit Inflation-Contingent Contracts

The benchmark model focuses on explicit COLA clauses in union contracts. This is the cleanest object to measure and the most direct source of inflation-contingent wage adjustment. However, explicit COLA clauses are not the only way in which wages can become contingent on inflation. In high-inflation environments, firms and workers can reopen their contracts, start renegotiations, and have more frequent wage reviews. Although the micro-level implementations differ, they all make wages effectively more responsive to inflation surprises. This subsection formalizes this idea in a reduced-form extension. The goal is not to build a second model of state-dependent wage setting, but to show how an implicit contingency margin can be embedded in the wage-index block of the benchmark economy.

I split the economy into two wage-setting sectors. The first sector is the union sector studied in the benchmark model. Its labor share is s_U, and its aggregate real wage index is $\hat{\omega}_t^U$. The second sector is an implicit-contingency sector with labor share $1 - s_U$. Contracts in this sector are not explicit COLA contracts: they start as non-contingent contracts. However, during the life of a contract, the wage schedule can be reopened with a state-dependent probability. Denote the aggregate real wage index in this sector by $\hat{\omega}_t^I$. The aggregate real wage index is therefore

$$\hat{\omega}_t = s_U \hat{\omega}_t^U + (1 - s_U)\hat{\omega}_t^I$$

Consider an implicit-sector contract that was signed at the end of period v and is in its k-th period in period $t = v + k$. If the contract is not reopened, its realized real wage follows the non-contingent wage rule chosen at v. I denote this stale non-contingent wage by $\hat{w}_{t,k}^{stale}$. If the contract is reopened, the wage is reset using the information available at t. I denote the reset wage by $\hat{w}_{t,k}^{reset}$. A convenient reduced-form representation of the realized wage is the mixture

$$\hat{w}_{t,k}^I = (1 - h_{t,k})\hat{w}_{t,k}^{stale} + h_{t,k}\hat{w}_{t,k}^{reset} \tag{61}$$

where $h_{t,k} \in [0, 1]$ is the reopening probability for contracts of age k. This equation is meant to capture reopening, catch-up adjustment, and informal cost-of-living adjustment in a parsimonious way.

The stale wage is the same object as a non-contingent wage in the benchmark model.[55] It is fixed in nominal terms at the time of contracting and therefore its realized real value falls with unexpected inflation. The reset wage uses the current information set. A simple way to discipline it is to set it equal to the wage a newly negotiated non-contingent contract would choose at t for the remaining life of the relationship. Hence, with one-period contracts, a reopened contract is equivalent to COLA. With multi-period contracts, they are the same in the period when the contract is reopened.

[55] One can relax this assumption and impose different contracting environments in the implicit-sector. The two wage accounting equations above will not change.

The reopening probability is reduced-form.[56] Let ψ_i be a contract-specific reopening cost, drawn from a distribution with CDF H. The contract is reopened when the expected benefit from reopening exceeds the cost:

$$h_{t,k} = H(B_{t,k}) \tag{62}$$

The object $B_{t,k}$ summarizes the value of escaping the stale nominal schedule. For example, one can write

$$B_{t,k} = \Gamma_1 \operatorname{Var}_t \left(\sum_{r=t+1}^{v+N} \pi_r \right) + \Gamma_2 \left| \sum_{r=v+1}^{t} (\pi_r - \mathbb{E}_v \pi_r) \right| \tag{63}$$

where $\Gamma_1, \Gamma_2 \geq 0$. The first term captures the option value of reopening when future inflation risk is high. The second term captures accumulated real-wage misalignment under the original nominal schedule. This specification is deliberately reduced-form: it is a sufficient-statistic representation of the forces that would appear in a full state-dependent wage renegotiation model.

Suppose cohorts are also evenly distributed in the implicit sector. The implicit-sector real wage index is

$$\hat{\omega}_t^I = \sum_{k=1}^{N} \left[\frac{h_{t,k}}{N} \hat{w}_{t,k}^{reset} + (\frac{1}{N} - \frac{h_{t,k}}{N}) \hat{w}_{t,k}^{stale} \right] \tag{64}$$

The structure resembles explicit COLA where

$$\hat{\omega}_t^U = \sum_{k=1}^{N} \left[\mathcal{M}_{t,k} \hat{w}_{t,k}^{\mathcal{C}} + (\frac{1}{N} - \mathcal{M}_{t,k}) \hat{w}_{t,k}^{\mathcal{N}} \right]$$

In other words, the reset probability plays a similar role to the COLA distribution.

Equations (61)-(64) show that implicit contingency enters the aggregate economy through the same object as explicit COLA: the aggregate real wage index. A higher reopening probability shifts the implicit-sector wage away from a stale non-contingent wage and toward a wage based on more recent inflation information.

Proposition 19. *Suppose the reset wage is weakly more inflation-contingent than the stale non-contingent wage, i.e., $\partial_\pi \hat{w}_{t,k}^{reset} \geq \partial_\pi \hat{w}_{t,k}^{stale}$. Then, holding the conditional general equilibrium fixed, an increase in the reopening probability $h_{t,k}$ weakly raises the inflation sensitivity of the implicit-sector real wage index. If $B_{t,k}$ is increasing in inflation risk or accumulated inflation surprises, then any force that raises inflation risk or inflation surprises also raises effective inflation contingency outside explicit COLA contracts.*

Proof. Differentiate (61) with respect to inflation:

$$\frac{\partial \hat{w}_{t,k}^I}{\partial \pi_t} = (1 - h_{t,k}) \frac{\partial \hat{w}_{t,k}^{stale}}{\partial \pi_t} + h_{t,k} \frac{\partial \hat{w}_{t,k}^{reset}}{\partial \pi_t} + \frac{\partial h_{t,k}}{\partial \pi_t} \left(\hat{w}_{t,k}^{reset} - \hat{w}_{t,k}^{stale} \right)$$

[56] One can derive it from the micro-level decision problem. It resembles the optimal stopping problem in a menu cost model. Here, I only discuss its qualitative feature. Formal derivation is left for future research.

For a fixed reopening probability, the derivative with respect to $h_{t,k}$ is the difference between the reset-wage and stale-wage inflation sensitivities, which is weakly positive by assumption. Aggregating across cohorts using (64) preserves the sign. Finally, since $h_{t,k} = H(B_{t,k})$ and H is weakly increasing, any increase in $B_{t,k}$ raises the reopening probability. □

This proposition clarifies the connection with the benchmark mechanism. In the benchmark model, a higher measured COLA share makes inflation more responsive to aggregate shocks and raises inflation volatility. In the extension above, this inflation-risk channel also raises the benefit from reopening stale nominal schedules in the implicit sector. To some extent, the contract distribution in the explicit union sector is isomorphic to the distribution of the reset probability in the implicit nonunion sector.

Moreover, explicit and implicit indexation interact through the aggregate inflation risk. A higher explicit COLA can make implicit adjustment more frequent, while more implicit adjustment feeds back into inflation volatility and can in turn change the value of explicit COLA. Still, one caveat is that the discussion here retains the benchmark assumption that workers cannot move across unions or firms. In other words, there is no direct labor market interaction. Workers in the labor market only (strategically) interact indirectly through the wage setting and the general equilibrium transmission.

I.2 Endogenous Stochastic Inflation Volatility

This section describes an extension in which inflation volatility is endogenous even though the primitive shock variances are constant. The motivation is the following. In the benchmark model, the gain from COLA depends on conditional second moments. This delivers a useful separation: first-order shocks move inflation and real activity, while the contract-choice margin is disciplined by inflation risk. As a result, for a given deterministic path of primitive shock volatilities, the equilibrium paths of COLA adoption and inflation volatility are both deterministic (but endogenous). This is analytically convenient, but it also removes a potentially important mechanism. In the micro data, COLA adoption rises with both inflation volatility and the level of inflation.

The exercise here asks what happens if inflation levels enter the gain from COLA. In that case, a first-order inflationary shock raises the value of COLA and thus COLA adoption. Since contracts are long-term, this changes the volatility of future inflation. In other words, the model is able to create stochastic inflation volatility from the extensive margin of contract choice, without imposing stochastic volatility directly on the primitive shocks. The mechanism differs from a reduced-form GARCH assumption, because the state-dependent volatility is structural: it comes from the endogenous contract distribution rather than from an assumed law of motion for shock variances.

Conceptually, this may help explain why inflation and inflation volatility often comove in the data: high inflation triggers more indexation, which amplifies inflation volatility. This may also

provide a transparent way to explain inflation-regime inertia: once the economy enters a high-COLA state, the same primitive shock is transmitted more strongly, making the high-COLA state more persistent and making exit from the high-volatility regime harder.

The main caveat is that I do not micro-found the inflation-level term in the COLA gain here. One possible foundation is that non-contingent contracts have less frequent wage reviews than COLA contracts.[57] Another is behavioral: when inflation is salient, workers and firms may perceive nominal wage schedules as more costly even holding conditional variances fixed. In the computation below, I take a reduced-form route and add inflation-level terms to the gain from COLA. I also shut down the intensive inference margin to simplify the computation. Conditional on a contract type, wage rules use the relevant rational forecasts, while COLA provides full pass-through to realized inflation.

I.2.1 Two-period contracts

I use a two-period-contract version of the model, which is suitable for demonstrating the numerical complexity of the problem. Each period a mass of one half of firms renews its contract. Let $\mathcal{M}_{t,1} \in [0, 1/2]$ denote the mass of newly active COLA contracts signed at the end of $t-1$, and let $\mathcal{M}_{t,2} = \mathcal{M}_{t-1,1}$ denote the older COLA cohort. The primitive shocks are a cost-push shock and a demand shock,

$$u_t = \varrho_u u_{t-1} + \sigma_u \varepsilon_{ut}, \qquad \rho_t = \varrho_\rho \rho_{t-1} + \sigma_\rho \varepsilon_{\rho t}$$

where $(\varepsilon_{ut}, \varepsilon_{\rho t})$ is i.i.d. standard normal and σ_u, σ_ρ are constant.

The aggregate real wage index is

$$\hat{\omega}_t = \frac{\mu}{2}\left(\mathbb{E}_{t-1}\hat{\zeta}_t + \mathbb{E}_{t-2}\hat{\zeta}_t\right) - \left(\frac{1}{2} - \mathcal{M}_{t,1}\right)(\pi_t - \mathbb{E}_{t-1}\pi_t) - \left(\frac{1}{2} - \mathcal{M}_{t,2}\right)(\pi_t - \mathbb{E}_{t-2}\pi_t + \pi_{t-1} - \mathbb{E}_{t-2}\pi_{t-1}) \tag{65}$$

The first term is the real target embedded in both types of contracts. The second and third terms are forecast errors borne by non-contingent contracts. COLA contracts are protected from these inflation forecast errors, so their mass subtracts from the non-contingent wage exposure.

The remaining aggregate block is the same as in the main text, where I redefine some notation for brevity:

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \kappa \hat{mc}_t + u_t \tag{66}$$

$$\hat{mc}_t = \alpha_1 \hat{y}_t + \alpha_2 \hat{\omega}_t \tag{67}$$

$$\hat{y}_t = \mathbb{E}_t \hat{y}_{t+1} - \frac{1}{\gamma}(i_t - \mathbb{E}_t \pi_{t+1} - \rho_t) \tag{68}$$

$$\hat{y}_t = \iota_1 \hat{\zeta}_t - \iota_2 \hat{\omega}_t \tag{69}$$

$$i_t = \phi_\pi \pi_t + \phi_y \hat{y}_t \tag{70}$$

[57] In practice, non-contingent contracts often specify annual wage adjustments, while most COLA contracts review each quarter or half a year.

Given the current state and the conditional expectation rules, equations (65)-(70) are linear in $(\pi_t, \hat{\zeta}_t)$ and can be solved period by period.

The gain from COLA for a newly signed contract is specified as

$$\mathcal{V}_t = C\left[\text{Var}_t\left(\pi_{t+1}\right) + \beta\left\{\text{Var}_t\left(\pi_{t+1}\right) + \text{Var}_t\left(\pi_{t+2}\right) + 2\,\text{Cov}_t\left(\pi_{t+1}, \pi_{t+2}\right)\right\}\right] + \theta_1\pi_t + \theta_2\pi_t^2. \quad (71)$$

The first bracket is the fundamental loss from non-contingent contracts' exposure to inflation forecast errors. The reduced-form terms $\theta_1\pi_t + \theta_2\pi_t^2$ allow inflation levels to enter contract choice. The new COLA inflow is

$$\mathcal{M}_{t+1,1} = \frac{1}{2}\Lambda\left(\frac{\mathcal{V}_t - \chi_0}{\chi_1}\right), \qquad \Lambda(z) = \frac{1}{1+\exp(-z)} \quad (72)$$

where χ_0 is the location parameter and χ_1 controls cost dispersion.

In the baseline model in the main text, we can solve for those expectations, variances, and covariances explicitly given the path of the contract distribution, past shocks, and a conjecture of the endogenous data generating process. Here, however, we do not have analytical formulas for them. This is because a conjecture that posits endogenous variables as linear functions of exogenous shocks cannot be verified in general. That means we have to keep track of them and add them to the state space. Therefore, the state vector is

$$x_t = \left(u_t, \rho_t, \pi_{t-1}, \mathbb{E}_{t-1}\left[\pi_t\right], \mathbb{E}_{t-1}\left[\hat{\zeta}_t\right], \mathbb{E}_{t-1}\left[\pi_{t+1}\right], \mathbb{E}_{t-1}\left[\hat{\zeta}_{t+1}\right], \mathbb{E}_{t-2}\left[\pi_{t-1}\right], \mathbb{E}_{t-2}\left[\pi_t\right], \mathbb{E}_{t-2}\left[\hat{\zeta}_t\right], \mathcal{M}_{t,1}, \mathcal{M}_{t,2}\right) \quad (73)$$

The endogenous expectation objects are

$$G(x_t) = \left(\mathbb{E}_t\left[\pi_{t+1}\right], \mathbb{E}_t\left[\hat{\zeta}_{t+1}\right], \mathbb{E}_t\left[\pi_{t+2}\right], \mathbb{E}_t\left[\hat{\zeta}_{t+2}\right], \text{Var}_t\left[\pi_{t+1}\right], \text{Var}_t\left[\pi_{t+2}\right], \text{Cov}_t(\pi_{t+1}, \pi_{t+2})\right)$$

Once G is known, current inflation, revenue, the wage index, the loss from non-contingent contracts, and next period's state are all known.

I.2.2 Solution method

The solution uses parameterized expectations. I approximate the elements of $G(x)$ as functions of the state vector. The approximation is structured to match the analytical form of the benchmark model: non-COLA states enter linearly, while their loadings are flexible Chebyshev polynomials in the two COLA masses $(\mathcal{M}_{t,1}, \mathcal{M}_{t,2})$. Variances are approximated in logs to guarantee positivity. Given a guess for G, I solve the current equilibrium, integrate over next period's Gaussian shocks using Gauss-Hermite quadrature, compute updated conditional means, variances, and covariances, and project these updated objects back onto the basis. Iterating this mapping gives a parameterized-expectations fixed point. I then refine the fixed point by minimizing the projected residual $T(G) - G$, where T is the expectation-updating operator.

In particular, there are 10 non-COLA states plus a constant term:

$$s_t = \left\{1, u_t, \rho_t, \pi_{t-1}, \mathbb{E}_{t-1}\left[\pi_t\right], \mathbb{E}_{t-1}\left[\hat{\zeta}_t\right], \mathbb{E}_{t-1}\left[\pi_{t+1}\right], \mathbb{E}_{t-1}\left[\hat{\zeta}_{t+1}\right], \mathbb{E}_{t-2}\left[\pi_{t-1}\right], \mathbb{E}_{t-2}\left[\pi_t\right], \mathbb{E}_{t-2}\left[\hat{\zeta}_t\right]\right\}$$

Each of the two COLA masses corresponds to degree-2 Chebyshev:

$$T_0(\mathcal{M}_{t,1}), T_1(\mathcal{M}_{t,1}), T_2(\mathcal{M}_{t,1}), T_0(\mathcal{M}_{t,2}), T_1(\mathcal{M}_{t,2}), T_2(\mathcal{M}_{t,2})$$

For each of the seven endogenous expectation objects, for instance, $\mathbb{E}_t\left[\pi_{t+1}\right]$, I guess

$$\mathbb{E}_t\left[\pi_{t+1}\right] = \sum_{j=0}^{10}\sum_{a=0}^{2}\sum_{b=0}^{2} c_{jab} T_a(\mathcal{M}_{t,1}) T_b(\mathcal{M}_{t,2}) s_{jt}$$

The iteration is over the coefficients $\left\{c_{jab}\right\}$. For the two-period contract case, $\mathbb{E}_t\left[\pi_{t+1}\right]$ alone involves $11 \times 3 \times 3 = 99$ coefficients, and all of the seven moments in $G(x_t)$ require solving for $7 \times 99 = 693$ coefficients as the fixed point.

For the benchmark simulation below, the COLA inflow has mean 0.355, standard deviation 0.184, and 10th, 50th, and 90th percentiles $(0.017, 0.462, 0.500)$. The 90/10 ratio of the one-step conditional inflation variance is 1.39, corresponding to a 90/10 ratio of conditional standard deviations of 1.18. The simulated correlation between inflation and the COLA inflow is 0.76, and the simulated correlation between inflation and conditional inflation variance is 0.89.

Figure 45: Conditional inflation variance by inflation in the endogenous-volatility benchmark

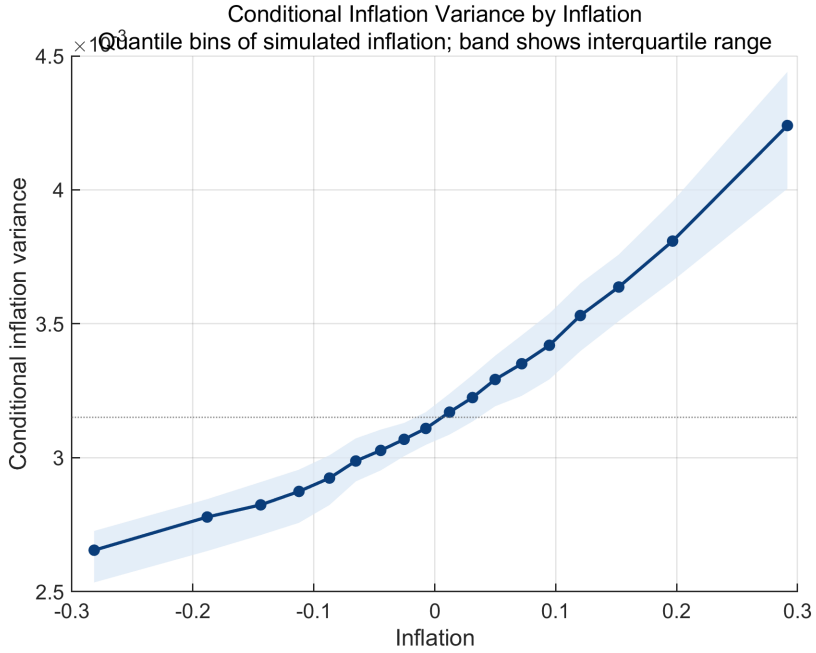


Figure 45 plots the average one-step conditional inflation variance across quantile bins of simulated inflation. The relationship is increasing: high-inflation states are also high conditional-volatility states. Importantly, the primitive shock variances are constant in the simulation. The positive relation is generated by contract choice: high inflation raises the gain from COLA, which raises COLA adoption, and a higher COLA mass amplifies inflation volatility.

Figure 46: A simulated high-COLA episode

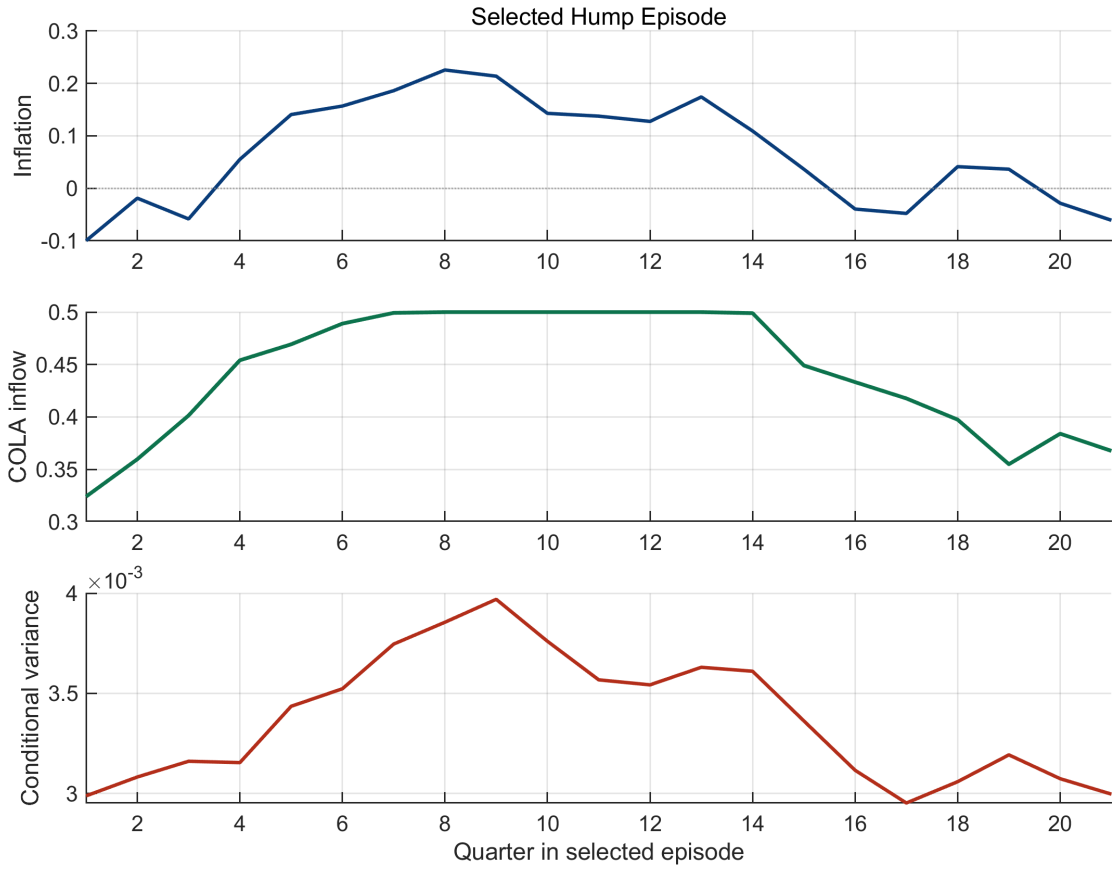


Figure 46 shows a short segment of the simulated path. The middle panel shows a hump-shaped increase and subsequent decline in new COLA inflows. The bottom panel shows that conditional inflation variance is elevated during the high-COLA part of the episode.

I.2.3 Longer contracts

Extending the exercise to longer contracts is conceptually straightforward but computationally difficult. With N-period contracts, the state must track N COLA cohorts, the relevant forecast histories for non-contingent wages, and cumulative forecast errors for each active vintage. The expectation vector also expands because the loss function for a newly signed contract depends on conditional variances and covariances over the entire contract horizon. Therefore the dimensionality of both the state space and the expectation objects grows quickly. The two-period problem above has 12 state variables. With the structured basis used here, the fixed point is represented

by 693 coefficients. A comparable four-period full-state problem has 30 state variables: two exogenous shocks, four active COLA cohorts, four cumulative forecast-error states, and 20 forecast-history states. If I used the same full tensor Chebyshev structure in the four COLA masses, the basis would have 2,187 terms. Since the four-period loss requires three expectation objects for each horizon--conditional means of inflation and revenue and a conditional risk object--there are 12 expectation objects, or 26,244 coefficients in total. This is before adding any extra basis terms for interactions among non-COLA states. Thus, a full tensor basis similar to the two-period approximation becomes impractical for four-period or eight-period contracts. Compressed-state or shock-history approximations are possible, but they tend to attenuate the amplification mechanism because they restrict how much independent contract-history information can enter the state. For this reason, I interpret the two-period exercise as an illustrative mechanism and leave a full quantitative assessment for future research.