Sticky Discount Rates*

Masao Fukui Niels Joachim Gormsen Kilian Huber

February 2025

Abstract

We show that firms' nominal required returns to capital (i.e., their discount rates) are sticky with respect to expected inflation. Such nominally sticky discount rates imply that increases in expected inflation directly lower firms' real discount rates and thereby raise real investment. We analyze the macroeconomic implications of sticky discount rates using a New Keynesian model. The model naturally generates investment-consumption comovement in response to household demand shocks and higher investment in response to government spending. Sticky discount rates imply that inflation has real effects, even absent other nominal rigidities, making them a distinct source of monetary non-neutrality. At the same time, sticky discount rates make the short-term interest rate less effective at stimulating investment. Optimal monetary policy in the model attempts to steer long-run inflation expectations and permanently, even when shocks are temporary.

Keywords: nominal rigidity, monetary non-neutrality, New Keynesian, discount rates

*We thank Adrien Auclert, Gabriel Chodorow-Reich, Mikhail Golosov, Veronica Guerrieri, Greg Kaplan, Rohan Kekre, Chen Lian, Emi Nakamura, Lasse Pedersen, Valerie Ramey, David Romer, Jón Steinsson, Ludwig Straub, Stephen Terry, Thomas Winberry, and Christian Wolf for comments. We are grateful to Thomas Bourany, Sonali Mishra, and Esfandiar Rouhani for excellent research support. This research was supported by the Becker Friedman Institute, the Fama-Miller Center, the Fama Faculty Fellowship, and the Lee Economics Program at the University of Chicago. Fukui is at Boston University, mfukui@bu.edu. Gormsen and Huber are at the University of Chicago, niels.gormsen@chicagobooth.edu and kilianhuber@uchicago.edu.

In modern models of business cycles, the real effects of monetary and fiscal shocks are driven by the assumption that nominal prices or wages are sticky. In this paper, we argue that there exists another form of nominal rigidity relevant to macroeconomic outcomes: firms' nominal required returns to capital, also known as firms' discount rates, do not respond to expected inflation. This stickiness influences firms' investment decisions and thereby aggregate dynamics. In a New Keynesian model augmented with sticky discount rates, aggregate consumption shocks generate investment-consumption comovement and government spending shocks "crowd in" investment. Sticky discount rates also constitute a new source of monetary non-neutrality: with sticky discount rates, shocks to expected inflation influence real investment, consumption, and output, even when wages and prices are fully flexible.

Standard macroeconomic models assume that discount rates are flexible and that firms move their nominal discount rates one-to-one with a relevant nominal interest rate. These assumptions ensure that firms' real discount rates, which are the difference between nominal discount rates and expected inflation, only depend on the real interest rate and do not move only because inflation has changed. As a result, the marginal return to capital, which is pinned down by real discount rates, also depends only on real interest rates and not directly on inflation. More generally, real investment in standard models does not change unless there are changes in real interest rates or investment opportunities.

If nominal discount rates are sticky, however, this logic reverses. Inflation now directly influences real discount rates. If expected inflation, for instance, increases and nominal discount rates are unchanged, real discount rates mechanically decrease. The decrease in real discount rates then implies that firms reduce their required marginal returns to capital and invest more until their marginal returns reach the lower discount rates. The positive shock to expected inflation thereby increases real investment, even if the economy is otherwise unchanged in real terms.

We present evidence that, in contrast to typical models, nominal discount rates are indeed sticky with respect to expected inflation. The evidence is based on a panel dataset of firms' discount rates for the years 2002 to 2023. The data are collected from corporate conference calls where firms themselves discuss their discount rates with investors and financial analysts. Firms communicate their discount rates as their minimum required returns from new investment projects, in nominal

terms. The reported discount rates are verified by the fact that firm-level changes in discount rates strongly predict future firm-level investment and realized returns, consistent with theory.

The data reveal that breakeven inflation in asset markets, a measure of expected inflation, is not associated with changes in firms' nominal discount rates. This is most easily seen in the recent inflationary period. Between 2020 and 2023, breakeven inflation increased by 1.5 percentage points. Over the same period, firms kept their nominal discount rates almost constant. As a result, firms' real discount rates decreased by around 1.5 percentage points, even though real interest rates did not decrease. The impact of inflation on real discount rates generalizes throughout our sample. Between 2002 and 2023, firms did not, on average, move their nominal discount rates with expected inflation, so real discount rates moved inversely to expected inflation. The negative comovement between real discount rates and expected inflation is not driven by fluctuations in real interest rates and risk premia. Rather, the comovement reflects that the real discount rates used by firms deviate from what standard models assume.[1]

We further study the nature of this stickiness in nominal discount rates. A large share of firms do not adjust their discount rates over time at all, with around 80 percent reporting unchanged discount rates within two years. However, for the few firms that do adjust their discount rates, there exists a large mass of small as well as some very large changes and, on average, the changes are close to contemporaneous changes in the cost of capital (the appropriate measure of firm funding cost, analogous to the interest rate in simple models). Taken together, these facts suggest that a friction in the spirit of Calvo (1983) may be a reasonable starting point to modeling sticky discount rates. Our initial aim is to explore whether sticky discount rates can have macroeconomic implications, so we do not attempt to micro-found the frictions underlying sticky discount rates. However, we will discuss several theories of organizational frictions that could explain why discount rates are sticky.

Since expected inflation lowers real discount rates, it should also directly raise real investment by lowering the real required return to capital. We indeed find evidence consistent with this mechanism. For firms that have not changed their

[1]Previous work by Gormsen and Huber (2023a) already shows that the average relation between discount rates and the cost of capital is statistically different from both zero and one. However, this work does not analyze the effect of inflation and only analyzes the 2002-2021 period with relatively stable inflation.

discount rate over a given period, we find that their investment rates increase with breakeven inflation, consistent with the real discount rates of these firms falling with expected inflation. In contrast, there is no such relation for firms that changed their discount rates over the same period.

We study how sticky discount rates influence macroeconomic dynamics using a textbook New Keynesian model augmented with sticky discount rates. In the textbook, firms simply set their nominal discount rate equal to the long-run nominal interest rate. In contrast, in our model, firms adjust their discount rates over time subject to a Calvo friction: only a random share of firms are allowed to update their discount rate every period. Firms that can update set their nominal discount rate to maximize the financial market value of the firm. Firms that cannot update simply use the previous period's nominal discount rate. Once firms have their nominal discount rate, they optimally choose investment to maximize firm value, subject to that given discount rate. In line with the empirical evidence, this nominal discount rate process implies that the average real discount rate and investment are directly affected by shocks to expected inflation. If we instead allow all firms to choose discount rates flexibly, the model generates the textbook assumption that discount rates reflect the long-run interest rate. We calibrate the stickiness of discount rates using the fact that roughly five percent of firms change their discount rate every quarter.

The remainder of the model is deliberately similar to a New Keynesian textbook model with hand-to-mouth households. The central bank follows a nominal interest rate rule and the government raises taxes and issues discretionary spending. To be consistent with the empirical literature on marginal propensities to consume, we model two types of households: permanent-income households with access to financial markets and hand-to-mouth households consuming only labor income. In extensions, we find that sticky discount rates generate similar mechanisms in a model with richer household heterogeneity (Kaplan et al. 2018, Auclert et al. 2023) and in a real business cycle model with only permanent-income households.

The presence of sticky discount rates has first-order implications for how investment reacts to aggregate demand shocks. We first show that a decrease in household patience (i.e., a negative shock to the household discount factor) raises both consumption and investment if discount rates are sticky. In contrast, the textbook model with flexible discount rates generates increased consumption but decreased

investment. Sticky discount rates generate the opposite investment pattern through two mechanisms. The first is that the patience shock increases consumer demand and expected inflation. In turn, inflation lowers the real discount rates of firms, thereby increasing their investment demand. This direct link between expected inflation and investment is the main mechanism highlighted throughout the paper. The second mechanism relates to the fact that central banks respond to the consumer demand shock by increasing interest rates, which lowers investment in the textbook model. But this increase in nominal interest rates has weaker effects when discount rates are sticky, in line with recent empirical evidence (e.g., see Koby and Wolf 2020). The two mechanisms together generate procyclical comovement between consumption and investment. Such comovement has been a long-standing challenge in the literature since Barro and King (1984) and the model with sticky discount rates generates it naturally.

Sticky discount rates also influence the effect of government spending on investment. With flexible discount rates, government spending raises the short-term interest rate and "crowds out" investment. In contrast, government spending "crowds in" investment when nominal discount rates are sticky. Government spending is therefore more powerful in stimulating the economy when discount rates are sticky. Again, the direct link between real discount rates and expected inflation as well as the lower interest rate sensitivity of investment jointly account for the different finding in the model with sticky discount rates. More generally, we find that aggregate demand shocks (e.g., to patience or government spending) generate comovement patterns that are qualitatively consistent with US business cycle data when discount rates are sticky. This result hints at the possibility that aggregate demand shocks can account for US business cycles more consistently in models with sticky discount rates than in textbook models.

We next focus on the effects of monetary policy on investment. We first highlight a distinct form of monetary non-neutrality due to sticky discount rates by examining changes in the central bank's long-run inflation target. A permanent increase in the inflation target immediately raises inflation expectations. With sticky nominal discount rates, this lowers real discount rates and raises investment in the subsequent few quarters. These effects are consistent with time series evidence on the real effects of the inflation target (e.g., Mumtaz and Theodoridis 2017). We also find that the inflationary response is much larger than the initial shock to the inflation

target, suggesting that inflation can endogenously amplify through sticky discount rates, consistent with historical episodes of runaway inflation. With flexible discount rates, the effects of the inflation target on inflation and investment are much smaller (and exactly zero when allowing firms to index prices to inflation).

We also examine monetary policy shocks to the short-term interest rate when discount rates are sticky. The investment response is qualitatively similar to the textbook model, but muted by a factor of ten because of the lower interest rate sensitivity of investment. This finding in a general equilibrium model is consistent with the reduced-form analysis in Gormsen and Huber (2023a).

The monetary policy results imply that changes to the central bank's long-run inflation target have stronger effects under sticky discount rates than in the textbook model, whereas the short-term interest rate has weaker effects. These results raise the possibility that monetary policy should rely more on inflation expectations when discount rates are sticky. We explore this idea by studying the optimal conduct of monetary policy. We find that a central bank that can credibly commit to future policies permanently changes its long-run inflation target in response to transitory shocks. By changing the inflation target, the central bank can directly affect inflation expectations, reduce the extent of misallocation caused by sticky discount rates, and prevent sticky discount rates from causing large welfare losses. In contrast, if the central bank cannot commit, it can only rely on the short-term interest rate as a policy tool and the welfare losses caused by sticky discount rates are large. Taken together, through the lens of the model with sticky discount rates, long-run inflation expectations become an important policy tool. A natural next step would be to explore this conclusion in a model where sticky discount rates are micro-founded.

The mechanisms generated by sticky discount rates operate independently from other nominal frictions. We show that they are present in models with and without nominally rigid wages and prices. Sticky discount rates also have different implications from the "money illusion" described by Modigliani and Cohn (1979), which predicts that the real cost of capital rises with expected inflation. Our data and model instead suggest that the real discount rate falls with expected inflation.

While firms' nominal discount rates do not move with expected inflation, the nominal cost of additional funding that firms calculate themselves and report on conference calls (i.e., the perceived cost of capital) is associated with expected infla-

tion. Similarly, firm surveys suggest that cash flow forecasts and prices move with expected inflation (e.g., Bunn et al. 2022). These patterns suggest that a model of general inattention to financial prices or inflation would not produce the empirical findings.

Apart from discount rates, cash flow forecasts are the other key input into firm investment decisions. Models with irrational expectations can explain why firms make systematic errors when forecasting their cash flows (e.g., Angeletos et al. 2021, Bordalo et al. 2022), yet do not directly explain why firms' discount rates should be sticky and not their perceived cost of capital. The macroeconomic predictions due to sticky discount rates are distinct and complementary to the existing literatures on expectation errors and inattention.[2]

1 Conceptual Overview

Both in theory and practice, firms make investment decisions based on discount rates. A firm's discount rate is the minimum return that a firm is willing to accept on new investment projects. In the context of macroeconomic models, it is also known as a firm's required return to capital and equals the expected marginal revenue product of capital. Almost all large firms report using methods based on a discount rate in their investment decisions (Trahan and Gitman 1995, Graham 2022). Firms either use the discount rate in a net present value (NPV) calculation or employ it as a "hurdle" rate, a threshold for the minimum rate of return that a project must meet. The NPV and threshold methods lead to equivalent investment decisions as long as the NPV of the firm's investment projects declines smoothly in the discount rate, which is the case in standard macroeconomic models.[3] Firms can choose their discount rate relatively freely, in particular if they have some degree of market power. The discount rate affects a firm's total investment because a lower

[2]For instance, a distinct firm-level prediction due to sticky discount rates is that expected inflation is directly linked to the real discount rates and investment rates of firms with unchanged nominal discount rates, but not of other firms. A distinct macroeconomic prediction is that monetary policy is not neutral even when wages and prices are flexible.

[3]See Brealey et al. (2011), pages 109–113 for details. This condition can be violated if projects involve large lending transactions early in the lifetime of the project or have multiple internal rates of return. Investment problems based on the stochastic discount factor can also be represented using a discount rate. We explain the relation between discount rate methods (NPV and hurdle), the stochastic discount factor, and the cost of capital in the Supplemental Material.

discount rate implies that the NPV of a standard project is higher and therefore the project is more likely to be accepted by the firm. We denote a firm's nominal discount rate by δ.[4]

Standard models assume that firms' discount rates are equal to their cost of capital. The cost of capital is the return required by financial investors in exchange for providing capital to the firm. In models without risk, the cost of capital is simply the risk-free interest rate. In more realistic settings with risk and different types of liabilities (i.e., debt and equity), the appropriate cost of capital incorporates risk premia and is the weighted average cost of debt and equity (known as WACC, Modigliani and Miller 1958). This weighted average cost is not directly observed because it depends on the unobserved risk perceptions of financial investors and because returns to debt and equity are usually not paid out by the firm to investors but earned through changes in financial prices (Fama and French 1997). Firms themselves cannot directly observe their cost of capital either and instead estimate a "perceived cost of capital," typically using financial market data. We denote a firm's nominal cost of capital by i.

Models typically assume that firms use their cost of capital as discount rate because this assumption implies that firms optimize their financial market value in standard models. For instance, in a simple model without risk, the firm maximizes its value by equating the marginal revenue product of capital with the interest rate. In more complex models with risk and multiple liabilities, this decision is analogous to equating the discount rate with the cost of capital.[5] The expected paths of future inflation and the cost of capital only matter for the firm's discount rate decision insofar as they influence the current cost of capital: in standard models, firms maximize the current value in financial markets by using the current cost of capital as the discount rate, even if one expects the cost of capital to change in the future.

[4]Textbooks recommend that firms should use multiple discount rates that vary with the risk of the projects they are considering. In practice, the vast majority of firms in the conference call data and in previous surveys report using just one discount rate that is based on a firm's typical project (Graham 2022). The empirical or theoretical analysis of this paper does not depend on the number of discount rates used by firms, since we focus on changes in the firm's representative discount rate over time.

[5]Using the cost of capital as discount rate leads to the same investment decision as a complex decision rule based on the stochastic discount factor, as long as the projects under consideration have the same risk as the firm's existing investments and the model is otherwise standard, as we explain in the Supplemental Material. This holds, for instance, in macroeconomic models with one type of capital.

The stylized assumption has implications for how expected inflation affects firm discount rates and investment. The change in the nominal cost of capital is approximately $\Delta i = \Delta r + \Delta\pi$, the sum of changes in the real cost of capital r and the expected inflation rate π. Empirically, expected inflation tends to affect the nominal cost of capital because it does not perfectly negatively comove with the real cost of capital. According to the stylized assumption, firms should incorporate expected inflation into their nominal discount rate exactly to the same extent as it affects their nominal cost of capital, so that $\Delta\delta = \Delta r + \Delta\pi$. Under this stylized benchmark, the change in the real discount rate (which is given by $\Delta\delta^{\text{real}}$) only depends on the real cost of capital and not directly on expected inflation (i.e., $\Delta\delta^{\text{real}} \equiv \Delta\delta - \Delta\pi = \Delta r$). As a result, this benchmark implies no direct effect of expected inflation on real investment through the real discount rate channel.

However, in practice, firms may not set their discount rates in line with the stylized assumption. In the extreme, imagine that a firm chooses a "sticky" nominal discount rate, in the sense that it does not incorporate changes in expected inflation at all. This would imply that the change in the nominal discount rate does not depend on $\Delta\pi$ and is simply $\Delta\delta = \Delta r$. In turn, the real discount rate now depends directly on expected inflation: $\Delta\delta^{\text{real}} \equiv \Delta\delta - \Delta\pi = \Delta r - \Delta\pi$. With sticky discount rates, there is thus a direct link between expected inflation and real investment: increases in expected inflation lead to lower real discount rates and therefore greater real investment.

2 Evidence on Sticky Discount Rates and Investment

We present evidence that discount rates are sticky with respect to expected inflation and that the investment rates of firms with sticky discount rates move with expected inflation.

2.1 Data

Typical datasets do not contain firms' discount rates and perceived cost of capital. We therefore use the deanonymized panel data assembled by Gormsen and Huber (2023a), which is based on corporate conference calls. Relative to that paper, we extend the sample period to additionally include the high-inflation years 2021 to 2023.

	(1) N	(2) mean	(3) p5	(4) p95
Discount rate	3,286	15.4	8.00	25
Perceived cost of capital	3,139	8.68	5.30	13
Breakeven inflation	2,100	2.15	1.19	3.06

Table 1: Summary Statistics

Details on the data collection and evidence that discount rates capture required returns are in that paper. We provide a brief summary here.

Listed firms typically organize quarterly conference calls where they inform analysts and investors about their investment strategy. On these calls, firms occasionally report their discount rate and perceived cost of capital. Discount rates are reported as minimum required internal rates of return on new investments, whereas the perceived cost of capital is the firm's estimate of its weighted average cost of capital. Most firms report only one discount rate for the whole firm. In case there are multiple reported discount rates, the data contain the discount rate that is most representative for the firm's projects.

The data are based on all call transcripts on the databases Refinitiv and FactSet for the period 2002q1 to 2023q1. A team of research assistants manually read through a total of 110,000 paragraphs from the calls (each of which contained at least one relevant keyword) to identify relevant numbers. The final data are based only on non-hypothetical statements by firm managers (i.e., excluding statements such as "imagine that the cost of capital is x percent"). The perceived cost of capital includes only statements referring to the cost of capital for the firm's total debt and equity (i.e., excluding statement such as "the yield for this bond issuance was x percent"). Firms report nominal numbers, except for a handful of cases where regulated energy companies report the government-approved real cost of capital. The average discount rate is 15.4 percent and the average perceived cost of capital is 8.9 percent, as shown in Table 1.[6]

Several pieces of evidence suggest that the discount rates and perceived cost of capital reported on the calls capture firms' investment behavior. For instance,

[6]Some firms report discount rates that are adjusted upward to compensate for the fact that certain overhead costs, such as the costs to the headquarters of administering new projects, are omitted from the cash flow analyses. For a detailed analysis of the level of discount rates, see Gormsen and Huber (2023a).

within-firm changes in reported discount rates predict changes in future investment and within-firm changes in the perceived cost of capital broadly reflect time variation in expected returns on debt and equity from financial models (see Gormsen and Huber 2023a and Table 4). Statements from calls often appear as evidence in securities lawsuits (Rogers et al. 2011) and analysts and investors ask detailed questions about previously reported rates, incentivizing managers to report these ex-post verifiable numbers accurately. We merge firm-level investment rates from Compustat (CAPEX_{t+1} / PPEN_t) by manually matching firm names.

We also merge breakeven inflation rates in the country of the firm's headquarter (annualized over a ten-year horizon). Breakeven inflation is the difference in yields between nominal and real risk-free bonds, thereby measuring the financial market's expected inflation rate. The ten-year horizon is appropriate because firms typically discuss long-run investments and use ten-year risk-free yields as basis when estimating their perceived cost of capital.

We study breakeven inflation since a consistent, quarterly measure exists for several countries (Australia, Germany, France, the UK, Japan, and the US), which we access on Bloomberg. In the recent inflationary period, breakeven inflation moved closely with firm survey expectations. For instance, US breakeven inflation rose by 1 percentage point between 2020 and 2021 and US firms' long-run expected inflation in the Coibion-Gorodnichenko survey rose by 0.9 percentage points.[7]

To understand whether firms in the sample have similar characteristics to other listed firms, we calculate the average percentile rank of firms in the sample for different firm characteristics, relative to the population of firms in Compustat in the same year and country. The results are in Table 2. The main difference is that firms in the sample are relatively large: the average market value rank of firms in the discount rate sample is 87. However, the sample is representative in terms of other characteristics, as the average rank is close to 50 for the investment rate, book-to-market ratio, profits/assets, leverage, and Z-score (bankruptcy risk).

[7]The survey contains annual data on long-run firm expectations from 2018 to 2022. In this period, the correlation between changes in US breakeven inflation and changes in the survey is 0.8. The level of firm expectations is higher than breakeven inflation, but changes, which are the focus of our paper, are highly correlated. Similarly, firms' perceived inflation tends to be higher than true inflation, but their changes are highly correlated (Savignac et al. 2022).

	Firms with observed disc. rate			Firms with observed CoC		
	mean	min	max	mean	min	max
Market value	87.34	26.79	100.00	84.83	5.18	100.00
Investment rate	50.48	3.11	99.78	51.08	0.60	99.88
Book-to-market ratio	46.52	0.46	94.15	47.73	0.12	99.90
Profits / assets	51.91	5.87	99.46	49.82	0.18	99.83
Leverage	59.91	4.92	99.47	61.22	3.03	100.00
Z-score (bankruptcy risk)	50.25	7.54	93.33	47.69	1.13	100.00

Table 2: Representativeness of Firms in the Sample

The table reports characteristics of firms for two samples: firms for which we observe at least one discount rate and firms for which we observe at least one perceived cost of capital. Characteristics are measured in percentile ranks relative to the universe of firms in Compustat in the same year and same country of listing. A mean value close to 50 indicates that the average rank of firms in our dataset is close to the average rank of firms in the Compustat year-country population.

2.2 Evidence on Discount Rates and Expected Inflation

We now present evidence that discount rates hardly respond to expected inflation, even though the cost of capital is sensitive to inflation. As motivating evidence, we focus on the "soaring 20s," the inflationary period after 2020. This period featured a relatively sharp shock to breakeven inflation, a measure of inflation expectations in financial markets. As shown in Figure 1, breakeven inflation was declining between 2018 and 2020 and then started rising sharply in early 2020. Breakeven inflation measures long-run expected inflation, which is the sufficient statistic for how inflation should affect the cost of capital and discount rates.

Alongside breakeven inflation, we plot time series of the average discount rate and cost of capital, measured using the conference calls. We isolate within-firm variation over time, to ensure that differences in sample composition in different quarters do not affect the results. Specifically, we regress firm discount rates on quarter and firm fixed effects and then measure the average discount rate in every quarter by adding the estimated quarter fixed effect to the unconditional mean discount rate. Similarly, we construct a within-firm time series of the perceived cost of capital by estimating quarter fixed effects. We plot the two within-firm time series in Figure 1 as three-quarter moving averages.

The nominal perceived cost of capital was falling until mid-2020 and started rising sharply in 2021. This suggests that firms know that breakeven inflation is associated with financial returns and adjust their nominal cost of capital accordingly,

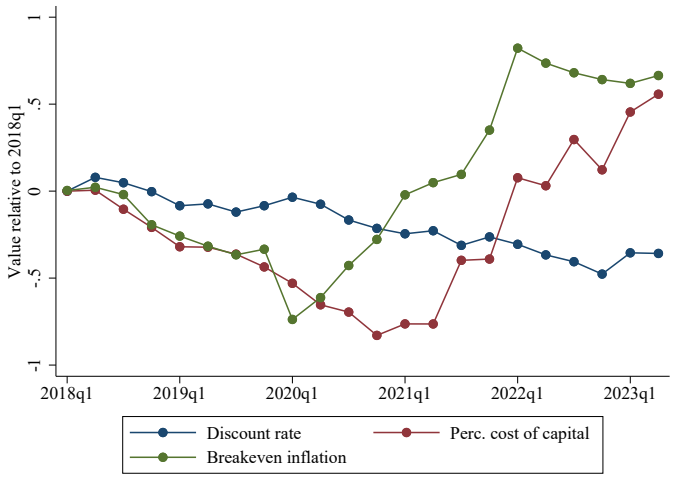


Figure 1: Discount Rates, Cost of Capital, and Inflation During the Soaring 20s

The figure plots the average quarterly perceived cost of capital, discount rate, and breakeven inflation rate between 2018q1 and 2023q1. We substract the 2018q1 value so that each series starts at 0. We isolate variation over time in the average perceived cost of capital and discount rate, which are reported on firms' conference calls, by controlling for time-invariant differences across firms. Specifically, we regress firm discount rates on quarter and firm fixed effects. We then measure the average "within-firm" discount rate in every quarter by adding the estimated quarter fixed effect to the unconditional mean and plotting a three-quarter moving average. We follow an analogous procedure to create the average "within-firm" perceived cost of capital in every year. Breakeven inflation is the annualized country-level breakeven rate (ten-year horizon), averaged across firms in our sample.

although with a delay of roughly 3 quarters relative to breakeven inflation.[8]

In contrast, firms' nominal discount rates continued to trend downward until mid-2022 and marginally rose thereafter. The combination of relatively stable nominal discount rates and rising inflation expectations implies that the real discount rates of firms decreased, even though the real perceived cost of capital stayed relatively constant. This decrease in real discount rates raised firm investment demand through the soaring 20s, relative to what it would have been if firms had moved nominal discount rates with the nominal cost of capital.

This pattern generalizes throughout the sample, as shown in Table 3. We include firm fixed effects in all regressions, which implies that we rely only on changes within firms to identify the coefficients. In column 1, we find that, on average, a 1 percentage point increase in breakeven inflation is associated with a statistically significant 0.2 percentage point increase in the perceived cost of capital. The average

[8] Gormsen and Huber (2023b) show that firms adjust their perceived cost of capital in line with traditional asset pricing factors, but do not analyze inflation.

VARIABLES	(1) CoC	(2) DiRa	(3) Wedge	(4) Wedge
Breakeven inflation	0.20**	-0.040	-0.28***	-0.23**
	(0.093)	(0.086)	(0.099)	(0.11)
Observations	1,547	1,546	1,546	1,546
FE	Firm	Firm	Firm	Firm/year
Within R^2	0.0054	0.00020	0.0087	0.0033

Table 3: Discount Rates and Expected Inflation

The table reports the relation between country-level expected inflation and firm-level rates reported on conference calls. Breakeven inflation is the annualized country-level breakeven rate (ten-year horizon). CoC is the firm's perceived weighted average cost of capital. DiRa is the firm's discount rate. The wedge is the difference between between the firm's DiRa and a predicted cost of capital. The predicted cost of capital is a generated measure of the firm's perceived cost of capital, predicted using a Lasso estimation procedure and using financial prices as inputs. The predicted cost of capital is closely associated with the perceived cost of capital. When using the predicted cost of capital as outcome, the coefficient on breakeven inflation is also 0.2***. The dataset is at the firm-quarter level and runs from 2002 to 2023. Standard errors (in parentheses) are clustered by firm. Statistical significance is denoted by *** p<0.01, ** p<0.05, * p<0.1.

relation is weaker than during the soaring 20s due to offsetting movements in real rates and risk premia during other periods of the sample. During the 2008-09 crisis, for instance, inflation decreased while real rates and risk premia increased. The advantage of our approach is that we account for such offsetting movements by first establishing the effect of breakeven inflation on the cost of capital and subsequently the effect of breakeven inflation on discount rates.

We find no evidence that discount rates change with breakeven inflation. The coefficient in column 2 is small, negative, and insignificant. Because breakeven inflation raises the cost of capital but not the discount rate, breakeven inflation affects the wedge between the two, as shown in column 3. A 1 percentage point increase in breakeven inflation is associated with a 0.3 percentage point decrease in the wedge. The coefficient is significant at the 1 percent level. Since firms often report only one of the discount rate or the perceived cost of capital in a given quarter, we calculate the wedge as the difference between the discount rate and a firm-level predicted value for the perceived cost of capital, following Gormsen and Huber (2023a). The coefficient is robust to including year fixed effects in column 4. We find similar results when analyzing only the pre-2020 period in Table A2.

Existing evidence in Gormsen and Huber (2023a) already documents time-varying wedges between discount rates and the perceived cost of capital for the years 2002

to 2021, thereby rejecting the stylized assumption that discount rates and the cost of capital always move one-to-one. However, this existing work does not directly examine how nominal discount rates respond to expected inflation and does not study the high-inflation period after 2020. It is plausible that firms incorporate changes in expected inflation into their discount rates to a smaller or larger extent than other changes in the cost of capital, especially when there are strong and salient changes in inflation. It is therefore not clear from previous work to what extent nominal discount rates are "sticky" with respect to expected inflation.

2.3 Additional Facts on Sticky Discount Rates

We now examine additional facts suggesting that a Calvo (1983)-style friction may be a reasonable starting point to modeling sticky discount rates. First, many firms maintain unchanged discount rates for several years. We plot in Figure 2 the share of firms with an unchanged discount rate relative to the first discount rate observation for the same firm. Around 80 percent of discount rates reported within two years remain unchanged. Half of firms report an unchanged discount rate after five years.

The second fact is that, once firms do change their discount rates, the changes are relatively dispersed, with a large mass of small as well as some very large changes, as shown in Figure 3. And third, the average change in the discount rate is close to the contemporaneous change in the perceived cost of capital.[9] These facts suggest that firms face constraints in adjusting their discount rates at all, but are able to make relatively large changes once they do adjust. The frictions underlying sticky discount rates may therefore be modeled using an environment featuring infrequent but dispersed changes in discount rates, analogously to the Calvo (1983) model of sticky prices (see also Alvarez et al. 2016).

A detailed micro-foundation for the frictions generating sticky discount rates is beyond the scope of this paper. Our focus will instead be on the implications of sticky discount rates for investment in a New Keynesian model. Our approach is therefore analogous to the large macroeconomics literature that assumes wage or

[9]To be precise, we find that, for firms that changed their discount rate, a 1 percentage point increase in the predicted perceived cost of capital is associated with a 0.71 percentage point (standard error: 0.38) increase in the discount rate, conditional on firm fixed effects. See Gormsen and Huber 2023a, Section 4.2 for related evidence.

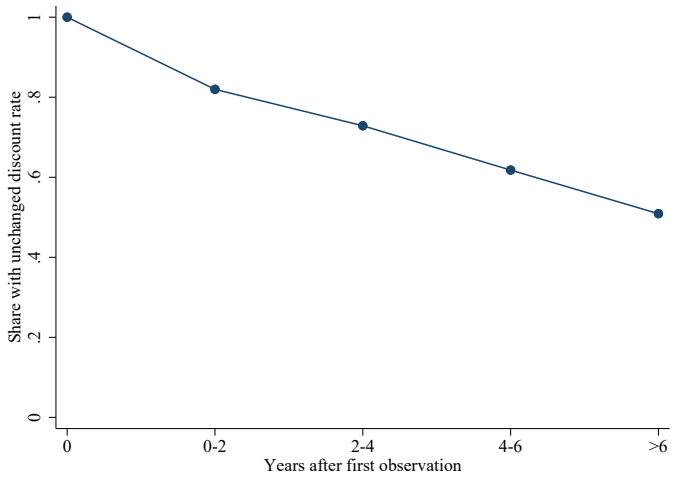


Figure 2: The Share of Firms With Unchanged Discount Rates

The figure plots the share of firms that have not changed their discount rate, relative to each firm's first observed discount rate in the dataset, by the number of years that have passed since the first observation. The starting point 0 on the horizontal axis indicates the quarter of the first observation, where by definition all observations are unchanged relative to the first observation. The second point on the horizontal axis (labeled 0-2) indicates observations from the two years following the first observation. The third point on the horizontal axis (labeled 2-4) indicates observations from the subsequent two years, and so on.

price stickiness without micro-founding such stickiness. Nonetheless, we point out organizational frictions discussed in the literature that could explain sticky discount rates.

A first potential explanation is that financial investors perceive stable discount rates as valuable signals about the quality of firms. Investors often worry that managers invest irresponsibly, for example, to eclectically build "empires" (Jensen 1986). Stable discount rates could serve as commitment devices and signal that firms deploy capital conscientiously. This view is supported by firm statements on the conference calls: 59 percent of firms with a positive discount rate wedge argue that stable discount rates benefit investors and firm value. Under this view, firms would ex-ante commit to only changing their discount rates rarely.[10]

A second potential explanation is that firms use sticky discount rates as organizational disciplining devices. If the firm's discount rate is "sacred" (Graham 2022)

[10] Alternatively, firms could commit to indexing their discount rate to the cost of capital. However, the cost of capital cannot be unambiguously observed and its estimation involves subjective choices and statistical uncertainty (Fama and French 1997), so indexing may be difficult to contract in advance.

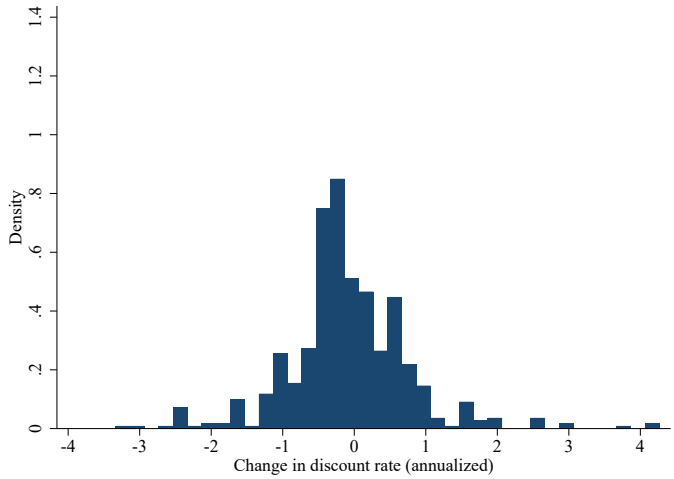


Figure 3: Non-Zero Changes in Discount Rates

The figure plots a histogram of the difference between a firm's discount rate in a given quarter and the firm's first observed discount rate. The plotted difference is in percentage points and annualized (i.e., normalized by the years between the quarter of observation and the quarter of the first observation). The sample includes only observations with non-zero changes (i.e., observations where the firm's discount rate in the given quarter differs from the first observed discount rate). The sample runs from 2002 to 2023.

and cannot be changed, different divisions within a firm may be less tempted to engage in rent-seeking and power struggles to tilt internal capital allocation in their favor (as in, e.g., Scharfstein and Stein 2000, Rajan et al. 2000). Instead, divisions will have to focus on developing high-return projects that meet the firm-wide discount rate in order to receive internal capital. In that sense, inflexible discount rates can serve as incentivizing devices set by top management. Under this view, many firms would maintain unchanged discount rates in order to avoid opening the door to divisional rent-seeking and power struggles.

Third, firms may believe that changes in the cost of capital do not accurately reflect changes in the true returns required by financial investors and firms may not want to maximize the current market value of their firm (but a future value). If both these assumptions hold, firms may not want to adjust discount rates with the cost of capital (Stein 1996). Instead, firms may only move discount rates infrequently, once they are convinced that the cost of capital accurately reflects required returns. Finally, sticky discount rates could reflect inattention or optimization failures by firms (Reis 2006, Maćkowiak et al. 2023). However, the type of inattention would need to be very specific and apply only to how firms set their nominal discount rate

and not to their perception of the nominal cost of capital, since the perceived cost of capital moves with expected inflation in the data.

2.4 Evidence on Investment and Inflation

If firms have sticky nominal discount rates, breakeven inflation affects their real discount rates. And since real discount rates influence real investment, breakeven inflation may in turn also affect the investment rates of firms with sticky discount rates. We examine this hypothesis by testing whether the investment rates of firms with unchanged discount rates behave differently when breakeven inflation increases.

We construct a firm-level dataset measuring the change in a firm's discount rate relative to the last observed discount rate in an earlier year. We construct indicators for whether the firm has the same discount rate in the two years or whether it has a different discount rate. Over the same period, we measure the change in the firm's annual investment rate and the change in breakeven inflation in the firm's country.

We regress the change in the investment rate on the change in breakeven inflation, interacted with the two indicators (for unchanged and changed discount rate). A 1 percentage point increase in breakeven inflation is associated with a 2.4 percentage point increase in the investment rate for firms with unchanged discount rate, as shown in column 1. The coefficient is statistically significant at the 5 percent level. In column 2, we additionally control for the change in the firm's nominal discount rate. We find a negative and statistically significant coefficient on the change in the discount rate. This confirms that discount rate changes are negatively associated with subsequent investment changes, consistent with the view that the discount rates reported on calls reflect firm investment behavior.

In column 3, we also control for the change in the return on equity (to proxy for idiosyncratic cash flow shocks). The coefficient on breakeven inflation for firms with unchanged discount rates remains robustly positive and significant. In contrast, the coefficient on breakeven inflation for firms with changed discount rate is negative and statistically indistinguishable from zero. The coefficient for firms with changed discount rate is also statistically different at the 10 percent level from the coefficient for firms with unchanged discount rate. This result suggests that changes in breakeven inflation are not generally positively associated with investment, but only for firms with unchanged discount rates.

VARIABLES	(1)	(2)	(3)
Change in breakeven inflation	2.36**	2.37**	2.21*
* 𝟙[unchanged discount rate]	(1.09)	(1.09)	(1.21)
Change in breakeven inflation	-0.74	-0.40	-1.81
* 𝟙[changed discount rate]	(1.56)	(1.48)	(1.99)
Change in discount rate		-0.63**	-0.72*
		(0.30)	(0.43)
Change in ROE			0.14***
			(0.050)
Observations	545	545	426
FE	Country	Country	Country
Within R^2	0.0092	0.017	0.034

Table 4: Investment of Firms With Sticky Discount Rates

The table shows that firms with stable discount rates increase their investment rate when expected inflation is high, but firms with an adjusted discount rates do not. We calculate whether a firm's discount rate, observed in year t, has changed relative to the last observed discount rate for the same firm, observed in a previous year t-j. The main regressors measure the change in the financial breakeven inflation rate in the firm's country between years t-j and t interacted with indicators for whether the firm has the same discount rate in t-j and t or whether it has a different discount rate. The outcome variable is the change in the firm's investment rate between years t-j and t, where we define the investment rate for year t as the firm's capital expenditure in the year after the discount rate is set ($CAPEX_{t+1}$) over the firm's capital stock ($PPEN_t$). We measure the investment rate using Compustat and trim at the 2.5^{th} and 97.5^{th} percentiles. The dataset is at the firm-year level and runs from 2002 to 2023. This dataset contains fewer observations because we use a first-differenced (not a panel) specification and annual (not quarterly) data to match the investment data. Standard errors (in parentheses) are clustered by firm. Statistical significance is denoted by *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

The results reveal an important difference between firms with and without unchanged discount rates in how breakeven inflation is associated with investment. Of course, nominal discount rates are not randomly assigned across firms, so the evidence does not offer conclusive causal proof. Nonetheless, the results are consistent with the view that sticky discount rates affect investment through a real discount rate channel.

While firms' discount rates are sticky with respect to inflation, their cash flows projections do not appear to be sticky (see Koller et al. 2015, chapter 22). For instance, firms report in surveys that their future prices and revenue depend on expected inflation (Meyer et al. 2021, Bunn et al. 2022). Moreover, Coibion et al. (2020) and Andrade et al. (2022) find that firms raise prices when expected inflation is

higher. Similarly, Coibion et al. (2018) find lower expected prices when expected inflation is lower (although the effect is imprecisely estimated).

Related to our findings, Coibion et al. (2018) implement an information treatment that lowered the inflation expectations of New Zealand firms but did not not shift their macroeconomic growth expectations. Treated firms subsequently reduced their investment by more. This finding is consistent with the view that lower expected inflation raises real discount rates.[11]

3 New Keynesian Model With Sticky Discount Rates

Inspired by the evidence, we augment a textbook New Keynesian model with sticky discount rates. In this section, we first describe a firm's investment problem with sticky discount rates. We illustrate two key mechanisms generated by sticky discount rates: a direct link from expected inflation to investment and a lower interest rate sensitivity of investment. We then present the full general equilibrium model and show that a calibrated model matches the firm-level evidence.

Our baseline model contains both permanent-income and hand-to-mouth households in order to capture realistic marginal propensities to consume. We show in the Supplemental Material that the main mechanisms also hold in a representative-agent real business cycle model and in a richer heterogeneous-agent model.

3.1 Firm Problem With Sticky Discount Rates

The basic setup of the firm problem is taken from the textbook. Time is discrete and the time horizon is infinite. There is a continuum of firms that own capital and hire labor to produce intermediate goods using the Cobb-Douglas production function

$$y_t = F_t(k_t, l_t) = A_t(k_t)^{\alpha}(l_t)^{1-\alpha}, \tag{1}$$

[11] In comparison, the information treatment in Coibion et al. (2020) raised Italian firms' inflation expectations but also worsened their macroeconomic growth expectations and uncertainty, which ultimately lowered their investment.

where k_t denotes the capital stock, l_t denotes labor, A_t is Hicks-neutral technology, and y_t is output at time t. The capital stock evolves according to the law of motion

$$k_{t+1} = (1 - \xi)k_t + I_t, \tag{2}$$

where $\xi \in [0, 1]$ is the depreciation rate and I_t is investment.

Firms hire labor services at nominal wage W_t and sell their output at price P_t in a competitive goods market. A firm's static profit is given by

$$\Omega_t(k_t) = \max_{l_t} P_t F_t(k_t, l_t) - W_t l_t. \tag{3}$$

Firms invest in capital subject to adjustment costs $\Phi(I_t, K_t) = \varphi(I_t/K_t)K_t$ and the elasticity in steady state is $\phi \equiv \frac{\varphi''(I/K)(I/K)}{\varphi'(I/K)}$. Investment and adjustment costs are both paid in the final good and P_t denotes the final good price.

Firms maximize the value of discounted future profits,

$$V_t^I(k, d_t) = \max_{k', I} \Omega_t(k) - P_t(I + \Phi(I, k)) + \frac{1}{1 + d_t}\mathbb{E}_t V_{t+1}^I(k', d_{t+1}) \tag{4}$$
$$\text{s.t.} \quad k' = (1 - \xi)k + I,$$

where d_t is the rate that a firm uses to discount its next-period value. The textbook model without risk assumes that the sequence of discount rates $\{d_s\}_{s=t}^{\infty}$ equals the sequence of expected short-term interest rates $\{i_s\}_{s=t}^{\infty}$. This assumption implies that firms maximize their value in financial markets.

The investment problem (4) is often written using just one discount rate δ_t that applies to all future periods, rather than a sequence $\{d_s\}_{s=t}^{\infty}$. This formulation is useful because the discount rate δ_t can now be interpreted as the firm's long-run required return, the key metric that firms use in their internal investment decisions and report on conference calls. The formulation using just one δ_t is:

$$V_t^I(k, \delta_t) = \max_{k', I} \Omega_t(k) - P_t(I + \Phi(I, k)) + \frac{1}{1 + \delta_t}\mathbb{E}_t V_{t+1}^I(k', \delta_t) \tag{5}$$
$$\text{s.t.} \quad k' = (1 - \xi)k + I.$$

This problem leads to the identical investment decision as (4) in log-linearized mod-

els, as long as the discount rate δ_t is appropriately chosen.[12] The log-linearized textbook model without risk assumes that the long-run discount rate δ_t equals the cost of capital, which is a weighted average of expected short-term interest rates. This assumption implies that the firm maximizes its value in financial markets in that model.

We deviate from the textbook by assuming that firms' long-run discount rates are sticky. The textbook assumes that firms can change their discount rates every period. We instead assume that only a randomly drawn fraction of firms $1-\theta$ can adjust their discount rate δ_t in a given period, while the remaining firms have to maintain unchanged discount rates relative to the previous period. Since firms cannot change their discount rates every period, they no longer maximize their value by simply using the current cost of capital as their discount rate. Instead, a firm that gets to change its discount rate has to solve for the optimal discount rate, given it has to use this discount rate until it gets to adjust it again by random chance. The optimal discount rate maximizes the value of the firm in financial markets, taking as given the firm's investment policy $I_t(k, \delta_t)$, the optimal solution to the firm problem (5). A firm that gets to change its discount rate thus solves:

$$\begin{aligned} V_t^a(k) = \max_{\delta_t} \Omega_t(k) - P_t(I + \Phi(I,k)) + \frac{1}{1+i_t}\mathbb{E}_t\left[\theta V_{t+1}^n(k', \delta_t) + (1-\theta)V_{t+1}^a(k')\right] \\ \text{s.t.} \quad k' = (1-\xi)k + I, \\ I = I_t(k, \delta_t), \end{aligned} \tag{6}$$

where $V_t^n(k, \delta_t)$ denotes the financial market value of a firm that cannot adjust in t and is given by

$$\begin{aligned} V_t^n(k, \delta_t) = \Omega_t(k) - P_t(I + \Phi(I,k)) + \frac{1}{1+i_t}\mathbb{E}_t\left[\theta V_{t+1}^n(k', \delta_t) + (1-\theta)V_{t+1}^a(k')\right] \\ \text{s.t.} \quad k' = (1-\xi) + I, \\ I = I_t(k, \delta_t). \end{aligned} \tag{7}$$

In problem (6), the firm maximizes its financial market value. This value is calcu-

[12] In Appendix A, we make a more general argument for the claim that the solutions to (4) and (5) are identical. Another way to understand this equivalence is that problems (4) and (5) yield the same log-linearized solution, where optimal investment solely depends on "marginal Q," as we show in Section 3.2.

lated using the sequence of short-term interest rates $\{i_s\}_{s=t}^{\infty}$ that investors in financial markets use to discount. The firm thus incorporates the full path of expected short-term interest rates when choosing its optimal discount rate. The expected next-period value of the firm in (6) is the weighted average of firm value if discount rates can be adjusted (V_{t+1}^a) and if discount rates cannot be adjusted (V_{t+1}^n).

Our formulation of sticky discount rates is analogous to Calvo's (1983) price stickiness and motivated by the evidence presented above. The nominal discount rates (long-run required returns) reported by firms often remain unchanged for several years and diverge from the nominal cost of capital. However, once the reported nominal discount rates do adjust, the observed changes can be relatively large and quantitatively close to changes in the cost of capital. The Calvo formulation similarly generates infrequent but dispersed changes in discount rates.

3.2 Log-Linearized Solution to the Firm Problem

Log-linearizing the firm problem around a deterministic steady state gives a tractable characterization.[13] Let $1 + r_t = \frac{1+i_t}{1+\pi_{t+1}}$ denote the gross short-term real interest rate, where π_{t+1} is inflation. Variables without time subscripts denote steady-state values. We define $\hat{i}_t = \log\left(\frac{1+i_t}{1+i}\right)$, $\hat{\delta}_t = \log\left(\frac{1+\delta_t}{1+\delta}\right)$, and $\hat{r}_t = \log\left(\frac{1+r_t}{1+r}\right)$. All other variables with hats represent log-deviations from steady-state values, for example, $\hat{I}_t = \log\left(\frac{I_t}{I}\right)$.

Proposition 1 (Optimal Discount Rates). *To a first-order approximation around the steady state, the optimal nominal discount rate for firms that can adjust their discount rate is given by*

$$\hat{\delta}_t^* = \frac{1 + r - \theta}{1 + r}\widehat{coc}_t + \frac{\theta}{1 + r}\hat{\delta}_{t+1}^*, \tag{8}$$

where

$$\widehat{coc}_t = \frac{r}{1 + r}\hat{i}_t + \frac{1}{1 + r}\mathbb{E}_t\widehat{coc}_{t+1} \tag{9}$$

defines the nominal cost of capital coc_t (alternatively known as long-run interest rate) and $1 + r$ is the steady-state real interest rate. The average nominal discount rate δ_t evolves

[13] A non-linear characterization of the firm problem is in the Supplemental Material.

according to

$$\hat{\delta}_t = \theta \hat{\delta}_{t-1} + (1-\theta)\hat{\delta}_t^*. \tag{10}$$

The proof is in Appendix B.

The results provide an intuitive characterization of discount rate dynamics. Let us first focus on the textbook case where discount rates are fully flexible, $\theta = 0$. In this case, (8) and (10) imply that

$$\hat{\delta}_t = \hat{\delta}_t^* = \widehat{coc}_t,$$

which says that nominal discount rates track the current nominal cost of capital. This condition implies that firms discount future values at the same rate as financial markets.

With sticky discount rates, $\theta > 0$, the discount rates chosen by firms reflect not only the current cost of capital but also future ones. Intuitively, firms internalize that they may not be able to adjust their discount rates in future, which leads them to introduce a wedge between the discount rate and the current cost of capital. In this case, firms therefore discount future values differently than financial markets.

Given that firms set discount rates as in Proposition 1, the aggregate investment rate is

$$\hat{I}_t - \hat{K}_t = \frac{1}{\xi\phi}\hat{q}_t^{adj}, \tag{11}$$

where q_t^{adj} is an "adjusted Q," the marginal value of capital in the eyes of the firm (marginal Q), and given by

$$\hat{q}_t^{adj} = \hat{q}_t - \frac{1+r}{r} \underbrace{(\hat{\delta}_t - \widehat{coc}_t)}_{\text{discount rate wedge}}. \tag{12}$$

Adjusted Q depends on two terms: the textbook Tobin's Q, given by q_t, and the "discount rate wedge," given by the difference between the average discount rate and the cost of capital. Tobin's Q is the marginal value of capital in financial markets:

$$\hat{q}_t = \mathbb{E}_t\left[\hat{\pi}_{t+1} - i_t + \frac{1}{1+r}\left((r+\xi)\hat{\omega}_{t+1} + \hat{q}_{t+1}\right)\right], \tag{13}$$

where

$$\hat{\omega}_t = \frac{1}{\alpha}\hat{A}_t + \frac{1-\alpha}{\alpha}(\hat{P}_t - \hat{W}_t) \tag{14}$$

are real profits per unit of capital. A higher Tobin's Q is associated with greater returns to investment and therefore raises firms' investment demand, as in the textbook. In contrast, a greater discount rate wedge implies that firms value future returns less than financial markets (i.e., firms' discount rates exceed the cost of capital), which in turn implies lower investment.

With flexible discount rates, $\theta = 0$, the discount rate wedge is always zero because firms set the discount rate equal to the cost of capital. Aggregate investment therefore only depends on Tobin's Q, as in the textbook. In contrast, when discount rates are sticky, $\theta > 0$, the average discount rate wedge set by firms is non-zero, adjusted Q does not equal Tobin's Q, and investment differs from the textbook model.

To understand whether investment in the model with sticky discount rates differs from the textbook, we need to understand the discount rate wedge. For this purpose, it is useful to define an average real discount rate δ_t^{real} that evolves as

$$\hat{\delta}_t^{\text{real}} = \hat{\delta}_t - \hat{\pi}_t^l, \tag{15}$$

where $\hat{\pi}_t^l$ is the long-run inflation rate, recursively given by $\hat{\pi}_t^l = \frac{r}{1+r}\hat{\pi}_t + \frac{1}{1+r}\mathbb{E}_t\hat{\pi}_{t+1}^l$. Similarly, we can define the real cost of capital $\widehat{coc}_t^{\text{real}}$ that evolves as

$$\widehat{coc}_t^{\text{real}} = \widehat{coc}_t - \hat{\pi}_t^l. \tag{16}$$

The discount rate wedge can therefore be written as a function of only real objects: $\hat{\delta}_t - \widehat{coc}_t = \hat{\delta}_t^{\text{real}} - \widehat{coc}_t^{\text{real}}$. Simply put, whenever the real average discount rate does not move one-to-one with the real cost of capital, the discount rate wedge changes and the model with sticky discount rates generates a different investment response than the textbook model.

3.3 Key Mechanisms with Sticky Discount Rates

In this section, we highlight two theoretical mechanisms that generate different investment responses in the model with sticky discount rates compared to the textbook model with flexible discount rates. The first mechanism is that expected infla-

tion directly raises investment when nominal discount rates are sticky, but not with flexible nominal discount rates. The second mechanism is that a real interest rate shock has weaker effects on investment when nominal discount rates are sticky.

Mechanism 1: Expected inflation and real discount rates We illustrate the first mechanism by simulating an exogenous 1 percentage point shock to short-run expected inflation π_{t+1}. We do not shock any other aggregates and conduct a partial equilibrium analysis. We therefore also assume that the short-run real interest rate r_t and the real cost of capital coc_t^{real} remain unchanged. Our simple aim here is to compare the effect of a pure inflationary shock in the model with sticky discount rates and the textbook. The exact model calibration is described in Section 3.5, but is not relevant for the qualitative mechanisms we illustrate here.

The left-hand panel of Figure 4a shows that the inflationary shock raises the short-term nominal interest rate i_t one-to-one, since $i_t = r_t + \pi_{t+1}$ and r_t is assumed constant. The middle panel shows that the average real discount rate δ_t^{real} remains unchanged in the textbook model because firms fully incorporate changes in nominal interest rates into their nominal discount rates. However, with sticky nominal discount rates, firms adjust their discount rates by less, as we derived in Proposition 1. As a result, the average real discount rate falls when discount rates are sticky.[14]

The discount rate wedge, given by the difference between the average discount rate and the cost of capital, rises when discount rates are sticky whereas it stays constant with flexible discount rates. The investment response to inflation therefore differs across the two models: investment remains unchanged in the textbook model but rises with sticky discount rates. The model with sticky discount rates thus features a direct link between expected inflation and real investment that does not exist with flexible discount rates.

Mechanism 2: Limited interest rate transmission We illustrate the second mechanism by simulating an exogenous 1 percentage point shock to the short-term real and nominal interest rates, r_t and i_t. We keep expected inflation π_{t+1} and other aggregates constant, as shown in the left-hand panel of Figure 4b.

[14]The decrease in the average real discount rate is smaller than the rise in expected inflation because the real discount rate captures the long-run required return and thus depends on current and future inflation.

The shock raises the real and nominal cost of capital, but by less than i_t, since the cost of capital depends on future interest rates as well. In the textbook model, firms always equate the real discount rate and the current real cost of capital, so the real discount rate rises in response to the shock. In the model with sticky discount rates, the real discount rate depends on both the current and future real cost of capital, so firms raise the discount rate by less than in the textbook model. As a result, the discount rate wedge becomes smaller and the investment response in the model with sticky discount rates is weaker. In general, sticky discount rates dampen the investment sensitivity to the interest rate.

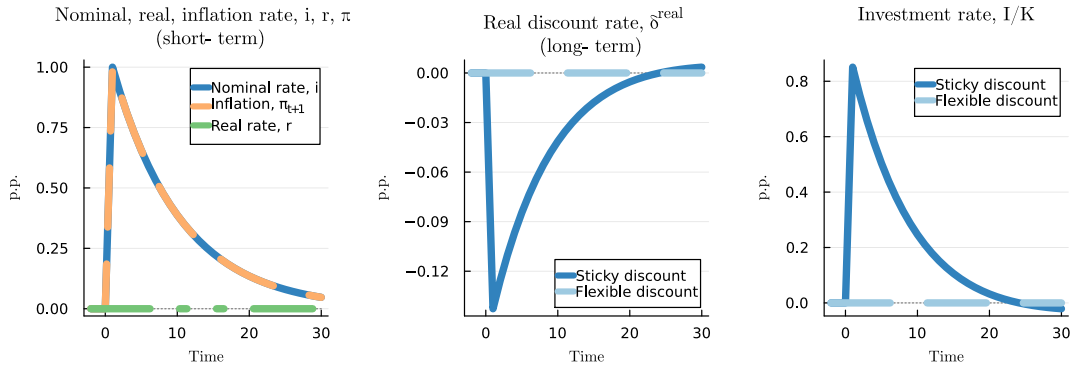
An influential literature has shown that the investment sensitivity to the short-term rate plays an important role in macroeconomic models (e.g., House and Shapiro 2008, Khan and Thomas 2008, Ottonello and Winberry 2020, Winberry 2021). Models with flexible discount rates often generate an investment sensitivity that is higher than the sensitivity implied by empirical estimates (e.g., see the discussion in Koby and Wolf 2020). The existing literature typically relies on large adjustment costs to match the data. Sticky discount rates are an empirically supported alternative that generates a lower short-run sensitivity without relying on large adjustment costs.

3.4 Rest of the New Keynesian Model

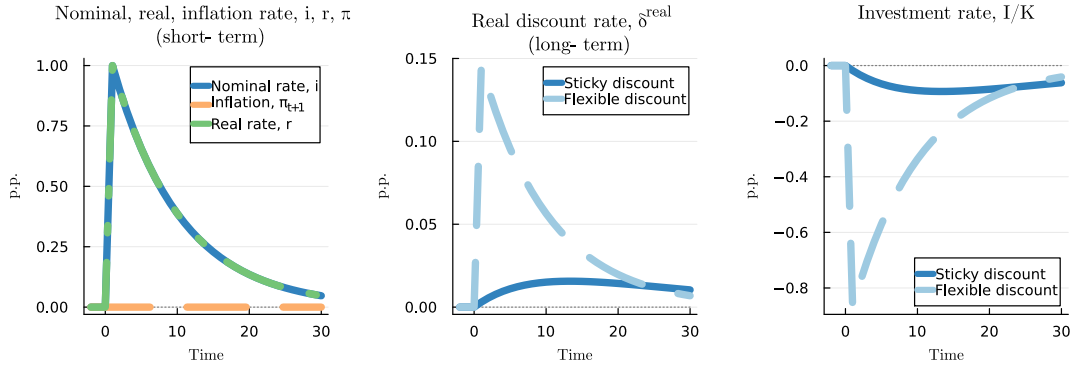
The rest of the model is deliberately close to the textbook New Keynesian model. Since the features are standard, we relegate details and robustness exercises (e.g., adding sticky wages and inflation indexation) to Appendix C and Appendix D.

In the baseline model, there are two types of households: a share $\mu^p \equiv 1 - \mu \in (0, 1]$ are permanent-income households and a share $\mu^h \equiv \mu \in [0, 1)$ are hand-to-mouth households. Permanent-income households have access to financial markets, whereas hand-to-mouth households do not. These features generate the large marginal propensity to consume out of transitory income shocks documented in the empirical literature. Superscript $i = p$ denotes permanent-income households and superscript $i = h$ denotes hand-to-mouth households. Both have the same preferences over consumption C_t^i and labor l_t^i, given by

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \prod_{s=0}^{t} \beta_t [u(C_t^i) - v(l_t^i)],$$



(a) Response to an expected inflation shock (keeping real interest rate constant)



(b) Response to a real interest rate shock (keeping expected inflation constant)

Figure 4: Partial Equilibrium Responses to Expected Inflation and the Interest Rate

Panel a plots the partial equilibrium responses to a shock to the expected inflation rate π_{t+1}, keeping the real interest rate unchanged. Panel b plots the partial equilibrium response to a shock to the nominal and real interest rate, keeping expected inflation unchanged. We use the calibration described in Table 5: $\theta = 0.95, r = 0.017, \phi_I = 10$, although the calibration does not affect the qualitative insights.

where $u(C) \equiv \frac{(C)^{1-\sigma}}{1-\sigma}$ and $v(l) = \frac{l^{1+\nu}}{1+\nu}$. Households face a stochastic shock to their patience parameter β that follows an AR(1) process:

$$\ln \beta_{t+1} = \ln \beta + \varrho_\beta \ln(\beta_t/\beta) + \varepsilon_t^\beta, \tag{17}$$

where $\varrho_\beta \in [0,1)$ governs the persistence of the shock and $\beta \in [0,1)$ is the steady-state household discount factor.

The budget constraint of permanent-income households is

$$C_t^p + b_{t+1} = (1 + r_{t-1}^p)b_t + (1 - \tau^l)W_t L_t/P_t - T_t, \tag{18}$$

where T_t is a lump-sum tax, τ^l is a labor income tax rate, and b_{t+1} are holdings of firms' debt and equity. The rate of return on these assets is $1 + r_t^p$, which in equilibrium gives the same expected rate of return as the real interest rate, $\mathbb{E}_t[1 + r_t^p] = 1 + r_t$. We introduce labor income taxation to offset the steady-state distortion from the market power of labor unions. Hand-to-mouth households earn labor income only.[15] The budget constraint of hand-to-mouth households is

$$C_t^h = (1 - \tau^l) W_t L_t / P_t - T_t. \tag{19}$$

Households supply labor through a continuum of labor unions, which convert labor into differentiated labor services. Aggregate labor services are given by a CES basket with elasticity of substitution of $\varepsilon_w > 1$. Labor unions set wages for differentiated labor services to maximize the weighted average utility of the two types of households. Wages are fully flexible.

Retailers purchase intermediate goods from firms at price p_t, convert them into differentiated products with an elasticity of substitution $\varepsilon > 1$, and set prices. A price can be adjusted with probability $1 - \gamma_p$. The government subsidizes the production of final goods at a rate $1 + \tau^p = \frac{\varepsilon}{\varepsilon - 1}$. After combining optimality conditions for unions and retailers and taking the first-order approximation around the zero-inflation steady state, we obtain the following New Keynesian Phillips curve:

$$\hat{\pi}_t = \psi_p \left[\sigma \left(\sum_i \vartheta^i \frac{1}{\chi^i} d\chi_t^i \right) + \sigma \hat{C}_t + \nu \hat{L}_t - \hat{A}_t - \alpha (\hat{K}_t - \hat{L}_t) \right] + \beta \mathbb{E}_t \hat{\pi}_{t+1}, \tag{20}$$

where $\psi_p \equiv (1 - \gamma_p)(1 - \beta\gamma_p)/\gamma_p$. Aggregate productivity follows an AR(1) process

$$\hat{A}_t = \varrho_A \hat{A}_{t-1} + \varepsilon_t^A, \tag{21}$$

where $\varrho_A \in [0, 1)$. We normalize steady-state productivity to 1.

The central bank sets the nominal interest rate according to the interest rate rule:

$$\hat{i}_t - \bar{\pi}_t = \ln R + \phi_\pi (\hat{\pi}_t - \bar{\pi}_t) + \varepsilon_t^m, \tag{22}$$

where $\bar{\pi}_t$ is the inflation target of the central bank, $R \equiv 1/\beta$ is the steady-state

[15] The results are similar if hand-to-mouth households also earn capital income (see Appendix C).

gross real interest rate, ϕ_π governs the strength of the monetary policy response to deviations from the inflation target, and ε_t^m is the monetary policy shock.

The government sets spending on the final good, G_t, collects labor income and lump-sum taxes, and issues government debt. The government budget constraint in real terms is

$$G_t + b_t^g = (1 + r_{t-1}) b_{t-1}^g + \tau^l W_t L_t / P_t + T_t, \tag{23}$$

where b_t^g are government bond holdings and government spending evolves as

$$G_t = \varrho_G G_{t-1} + \varepsilon_t^G. \tag{24}$$

Following Auclert et al. (2023), government bonds evolve as

$$b_t^g = \varrho_B (b_{t-1}^g + G_t). \tag{25}$$

Government spending is thus financed by an initial deficit and gradually raising taxes. The parameter $\varrho_B \in [0, 1)$ captures the degree of deficit financing.

Throughout, we work with a log-linear approximation around the steady state where the government runs a balanced budget, $b_t^g = 0$, and government spending is zero, $G_t = 0$, in steady state. The goods market clearing condition is, to a first-order approximation,

$$\frac{C}{Y}\hat{C}_t + \frac{I}{Y}\hat{I}_t + G_t = \hat{A}_t + \alpha \hat{K}_t + (1 - \alpha)\hat{L}_t. \tag{26}$$

Combining these assumptions, we can derive the first-order aggregate consumption Euler equation, as in Dupraz (2023). Given the paths of $\{\hat{I}_t, G_t, T_t, \hat{\beta}_{t+1}, \hat{r}_{t+1}\}$, the path of aggregate consumption, $\{\hat{C}_t\}$, satisfies

$$\hat{C}_t = \mathbb{E}_t \left[\hat{C}_{t+1} - \frac{\zeta_r}{\sigma} \left[\hat{\beta}_{t+1} + \hat{r}_{t+1} \right] + \zeta_I \left[\hat{I}_t - \hat{I}_{t+1} \right] + \zeta_G^0 G_t - \zeta_G^1 G_{t+1} - \zeta_T^0 T_t + \zeta_T^1 T_{t+1} \right], \tag{27}$$

where $\zeta_r = 1 - \frac{(1-\alpha)\mu\frac{I}{C}}{1-\mu+\mu\alpha}$, $\zeta_I = \frac{\mu(1-\alpha)\frac{I}{C}}{1-\mu+\mu\alpha}$, $\zeta_T^1 = \frac{1}{1-\mu+\mu\alpha}\frac{\mu}{C}$, $\zeta_G^1 = \frac{1}{1-\mu+\mu\alpha}[1-\alpha]\,\mu\frac{1}{C}$,

$$\zeta_G^0 = \frac{1}{1-\mu+\mu\alpha}\frac{1}{\beta}\frac{1}{C}\left(1-\beta\left[1-\mu+\mu\alpha\right]\right) - \zeta_r\frac{1-\beta}{\beta}\frac{1}{C},$$
$$\zeta_T^0 = \frac{1}{1-\mu+\mu\alpha}\frac{1}{\beta}\frac{1}{C}\left(1-\beta(1-\mu)\right) - \zeta_r\frac{1-\beta}{\beta}\frac{1}{C}.$$

Without hand-to-mouth households (i.e., $\mu = 0$), the above expression collapses to a standard consumption Euler equation: $\hat{C}_t = \mathbb{E}_t[\hat{C}_{t+1} - (1/\sigma)(\hat{\beta}_{t+1} + \hat{r}_{t+1})]$. When $\mu > 0$, we see that $\zeta_I > 0$. In this case, a temporary increase in investment will stimulate consumption, as emphasized by Auclert et al. (2020).

The relative consumption of hand-to-mouth households is, to a first order,

$$\chi_t^h - \chi^h = (1-\alpha-\chi^h)\hat{C}_t + (1-\alpha)\frac{1}{C}G_t - \frac{1}{C}T_t + (1-\alpha)\frac{I}{C}\hat{I}_t, \tag{28}$$

and the consumption share of permanent-income households satisfies the adding-up constraint:

$$0 = \mu(\chi_t^h - \chi^h) + (1-\mu)(\chi_t^p - \chi^p). \tag{29}$$

Given a sequence of shocks, $\{\varepsilon_t^A, \varepsilon_t^G, \varepsilon_t^m, \varepsilon_t^\beta\}$, the initial capital stock K_{-1}, and initial price level P_{-1}, the equilibrium of this economy consists of quantities $\{\hat{C}_t, \hat{I}_t, \hat{\imath}_t, \hat{L}_t, \hat{K}_t, \chi_t^h, \chi_t^p, T_t, b_t^g\}$, prices $\{\hat{\delta}_t^*, \hat{\delta}_t, \widehat{cocc}_t, i_t, \hat{W}_t, \hat{P}_t, \hat{\pi}_t\}$, and shock processes $\{\beta_t, G_t, A_t\}$ such that (8)-(14), (17), and (20)-(29) hold.

3.5 Calibration

We will use a calibrated version of the model to study how an economy with sticky discount rates responds to shocks. Table 5 summarizes the calibration, which is at the quarterly frequency. We set the discount rate stickiness parameter θ to 0.95 because the data in Figure 2 suggest that around 80 percent of nominal discount rates remain unchanged within two years, implying a quarterly value of roughly 5 percent. As benchmark, we also present results for an alternative textbook calibration with flexible discount rates where $\theta = 0$.

The remaining calibration choices are mostly standard in the business cycle literature. We set the household's patience parameter to match a steady-state annual

Parameter	Description	Value
θ	Stickiness in discount rates	0.95
β	Household patience parameter	0.98
ξ	Depreciation rate	0.025
α	Capital share	0.3
$1/\sigma$	Elasticity of intertemporal substitution	0.2
$1/\nu$	Frisch elasticity of labor supply	0.5
ϕ_π	Taylor rule inflation coefficient	1.5
γ_p	Stickiness in prices	0.75
ϕ	Capital adjustment cost	10
μ	Hand-to-mouth share	0.3
ϱ	Shock persistence	0.9

Table 5: Parameter Values in the Calibration

real cost of capital of 7%, the average real perceived cost of capital in our dataset. This gives $\beta = 0.98$. The capital share parameter is $\alpha = 0.3$. We set the inverse elasticity of intertemporal substitution to $\sigma = 5$ and the inverse Frisch elasticity of labor supply to $\nu = 2$. The coefficient on inflation in the interest rule is $\phi_\pi = 1.5$. The elasticity of capital adjustment costs is $\phi = 10$, following Philippon (2009).

The price stickiness parameter is $\gamma_p = 0.75$, which implies a quarterly 25% adjustment probability, in line with Nakamura and Steinsson (2008). The share of hand-to-mouth households is $\mu = 0.3$, corresponding to the fraction of liquidity-constrained households in Kaplan et al. (2014). Shocks to technology, government spending, and patience have the same persistence of 0.9, $\varrho_A = \varrho_G = \varrho_\beta = 0.9$. The speed at which the government adjusts its deficit is $\varrho_B = 0.982$ (0.93 at the annual frequency), following Auclert et al. (2023). We set $\vartheta^i = \mu^i$ so that the consumption distribution does not affect labor supply.

3.6 Expected Inflation and Investment in Data and Model

In the data, the investment rates of firms that have maintained unchanged discount rates increase by more in response to expected inflation than those of firms that have adjusted their discount rates. We now show that the model generates a similar pattern.

In both data and model, we compare firms with and without unchanged discount rates. This cross-sectional comparison holds constant all general equilibrium

effects that affect both firm types evenly. The advantage of this approach is that a shock generating breakeven inflation may affect all firms and that therefore it is not clear to what extent sticky discount rates drive the investment response.[16] In contrast, the cross-sectional comparison isolates the impact of sticky discount rates on firm investment when breakeven inflation changes.

In column 3 of Table 4, we showed the empirical estimate: when breakeven inflation increases by 1 percentage point, firms with an unchanged discount rate increase their investment rate by 4 percentage points more than firms that have adjusted their discount rate (statistically significant at 10 percent). To generate an analogous coefficient in the model, we simulate a shock to breakeven inflation. We assume that the shock affects the inflation rate expected in all future periods evenly, in line with the high persistence implied by the term structure of breakeven inflation, and that long-run nominal interest rates change twice as much as breakeven inflation, in line with the estimated relationship between expected inflation in the Survey of Professional Forecasters and long-run rates.

Using data on investment generated by the model, we then regress the annual change in the firm-level investment rate on the annual change in breakeven inflation, interacted with an indicator for whether the firm changed its discount rate in the previous year. We find that the coefficient in the model is 3.1. Since the coefficient depends only on the direct comparison of firms with and without unchanged discount rates, it does not vary with general equilibrium forces and is similar in models with flexible prices or richer household heterogeneity (see the Supplemental Material). In sum, the key mechanism relating expected inflation to investment thus operates similarly in both data and model.

4 Macroeconomic Dynamics With Sticky Discount Rates

We use the model to study the macroeconomic implications of sticky discount rates. We show that sticky discount rates alter the macroeconomic effects of aggregate demand shocks, including shocks to household patience, government spending, and monetary policy. We also consider implications for optimal monetary policy.

[16] For example, breakeven inflation may be the result of an optimal monetary policy response to a household patience shock. The path of investment would conflate the impacts of the patience shock, breakeven inflation, and the monetary policy response.

The main insight of this section is that the two mechanisms highlighted in Section 3.3 can be influential in general equilibrium.

4.1 Household Patience and Investment

A challenge faced by business cycle models is to generate the empirically observed positive comovement between consumption and investment in response to aggregate demand shocks (Barro and King 1984). In standard models, positive demand shocks, such as decreases in household patience, raise consumption and the real interest rate, thereby reducing investment.

We argue that sticky discount rates can naturally generate comovement in response to aggregate demand shocks. We consider a positive shock to consumption demand (a negative shock to households' patience parameter β_t). Figure 5 displays the impulse responses to a unit increase in ε_t^β at $t = 0$. In a benchmark model with fully flexible discount rates, consumption rises and investment falls. In contrast, the model with sticky discount rates generates increases on impact in both consumption and investment. Two mechanisms generate this stark difference. First, the patience shock raises consumption and thus expected inflation. In turn, inflation lowers real discount rates and raises investment when discount rates are sticky. Second, sticky discount rates adjust more slowly to higher nominal interest rates, so higher interest rates reduce investment by less. After around five quarters, firms reduce investment because inflation has fallen and firms have invested more than optimal in the long run.

We explore additional comovement patterns in Appendix E. We generally find that the comovement patterns generated by the model with sticky discount rates are qualitatively consistent with the US business cycle data when we simulate aggregate demand shocks, including household patience, government spending, and inflation target shocks. For these shocks, the model with sticky discount rates generates the correct sign for key correlations, such as investment-consumption and output-consumption. In contrast, the model with flexible discount rates generates the opposite sign in these cases. The results hint that models with sticky discount rates could imply that aggregate demand shocks have played a more important role in US business cycles than suggested by previous research (e.g., Smets and Wouters 2007). Estimating a full quantitative model with richer shocks is beyond the scope of this paper, but may be a natural next step for future research.

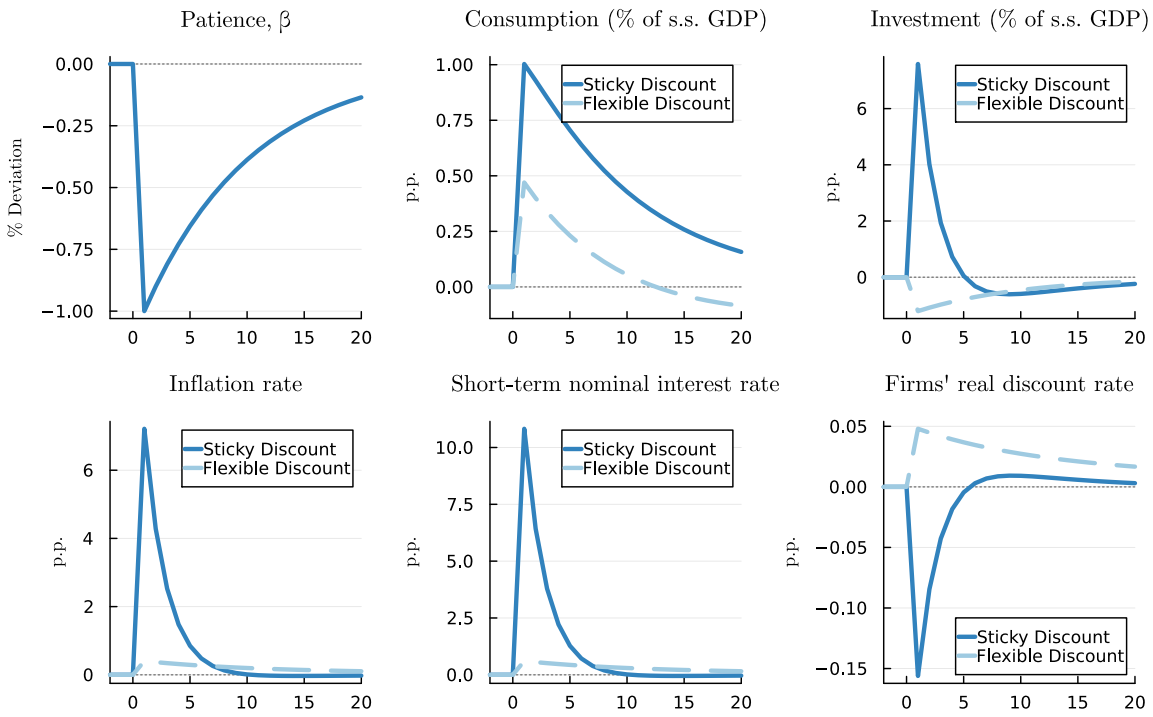


Figure 5: Impulse Responses to a Household Patience Shock

The figure plots impulse responses to a household patience shock for two different values of discount rate stickiness, $\theta \in \{0, 0.95\}$.

4.2 Government Spending and Investment

In Figure 6, we consider another demand shock, namely, government spending financed by an initial deficit and gradually increasing taxes. In the textbook, government spending raises the short-term rate and thereby "crowds out" investment. However, in the model with sticky discount rates, investment rises in response to a government spending shock. The two driving mechanisms are similar to above. Government spending raises expected inflation, thus lowering real discount rates and increasing investment. In addition, the interest rate sensitivity of investment is muted when discount rates are sticky.

The response of consumption depends on the calibration of the hand-to-mouth households. In our calibration, consumption increases on impact because greater investment raises labor income and thus the consumption of hand-to-mouth households. After about ten quarters, the tax increases reduce consumption. With fewer hand-to-mouth households, consumption could also fall sooner (due to the rise in the short-term interest rate). We choose a relatively standard calibration and refer to

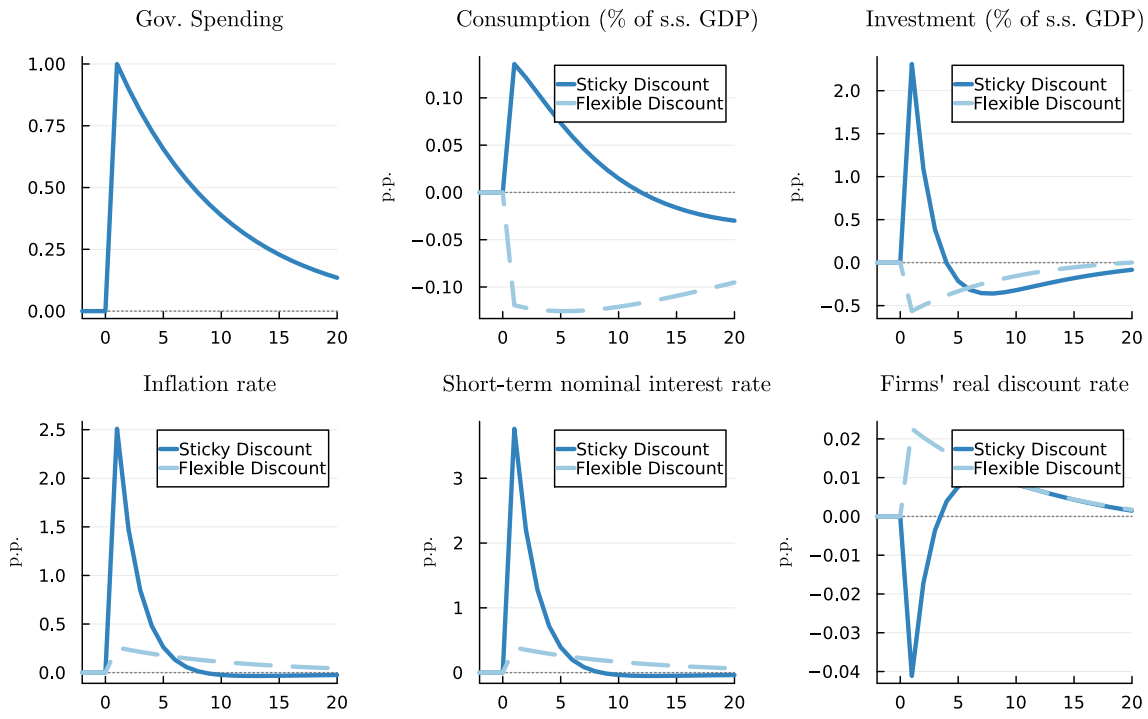


Figure 6: Impulse Responses to a Government Spending Shock

The figure plots impulse responses to a government spending shock for two different values of discount rate stickiness, $\theta \in \{0, 0.95\}$.

the heterogeneous-agent literature for detailed analyses (e.g., see the Supplemental Material and Kaplan et al. 2018, Auclert et al. 2023, Wolf 2023).

The related time series evidence on how government spending affects investment and consumption is mixed. For instance, Ramey (2016) reports that investment rises when identifying spending shocks using the Ben Zeev–Pappa series but falls using the Blanchard–Perotti and Ramey news series.

4.3 Monetary Policy and Investment

We show that sticky discount rates are an independent source of monetary nonneutrality because expected inflation directly affects real investment. However, shocks to the short-term interest rate are less effective than in the textbook model.

We illustrate the effects of expected inflation by studying unanticipated changes in the long-run inflation target of the central bank. In Figure 7, we consider a permanent increase in the inflation target of 0.1 percentage points. With flexible discount rates, the inflation target shock has only small effects on investment and these ef-

fects are exactly zero when prices can be indexed to inflation, as shown in Appendix D. In contrast, the increased inflation target has strong effects on investment in the model with sticky discount rates, because of the direct link from expected inflation to real discount rates. The resulting increase in labor income boosts the consumption of hand-to-mouth households.

The increases in investment and consumption in turn feed back to inflation. In Figure 7, the initial inflation response is 5 percentage points, even though the initial shock to the inflation target was only 0.1 percentage points. Therefore, sticky discount rates can endogenously amplify even relatively small inflationary shocks.

The real effects of changes in the inflation target are not driven by other nominal rigidities, as the inflation target affects investment similarly in a model with flexible prices (see the Supplemental Material). This finding underscores that sticky discount rates constitute a distinct source of monetary non-neutrality, even in otherwise frictionless models.

Our model leaves out other potential effects of high inflation, such as increases in uncertainty, that may ultimately harm investment. The main aim of this subsection is to illustrate that the mechanism linking expected inflation to investment through sticky discount rates can matter in general equilibrium, but we do not argue that inflation is always beneficial to investment. Nonetheless, the effects of an inflation target shock in the model are consistent with time series evidence showing that modest shocks to the inflation target can generate large increases in investment (Mumtaz and Theodoridis 2017) and output (De Michelis and Iacoviello 2016, Uribe 2022, Lukmanova and Rabitsch 2023). Moreover, Ireland (2007) argues that the US Federal Reserve has adjusted its inflation target repeatedly since the 1950s, suggesting that the inflation target can be a policy tool.

In Figure 8, we analyze a temporary shock to the short-term interest rate. We display impulse responses to a unit increase in the Taylor rule innovation, ε_t^m, at $t = 0$ that decays with an autocorrelation of 0.45. Sticky discount rates dampen the short-run effect on investment by about two-third. This result follows from the fact that discount rates respond substantially more sluggishly to the short-term rate.

4.4 Implications for Optimal Monetary Policy

The results above show that changes to the central bank's long-run inflation target are more powerful under sticky discount rates than in the textbook model, whereas

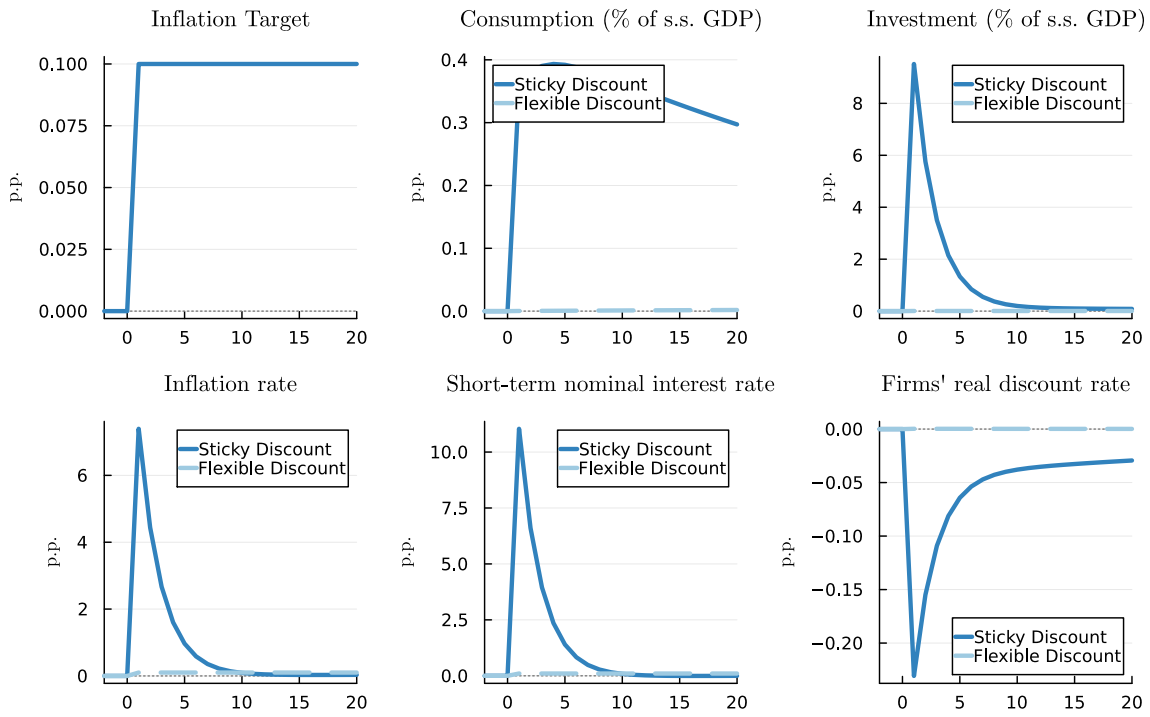


Figure 7: Impulse Responses to an Inflation Target Shock

The figure plots impulse responses to an increase in the long-run inflation target of 0.1 percentage points for two different values of discount rate stickiness, $\theta \in \{0, 0.95\}$.

the short-term interest has weaker effects. We now explore how these insights could affect the optimal conduct of monetary policy. Through the lens of our model, we find that the central bank's inflation target should be actively employed as a policy tool.

The central bank's objective is to maximize the welfare of households with Pareto weight Γ^i attached to each household type $i = p, h$. We follow McKay and Wolf (2022) and choose Pareto weights $\{\Gamma^i\}$ as well as tax rates $\{\tau^l, \tau^p\}$ so that the steady state is efficient. As is common in the literature, we study the linear-quadratic approximation of the policy problem around the steady state with zero inflation. We present the main findings here and additional details in Appendix F.

The linear-quadratic approximation of the optimal policy problem is to minimize the loss function

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \frac{1}{2} \mathbb{L}_t, \tag{30}$$

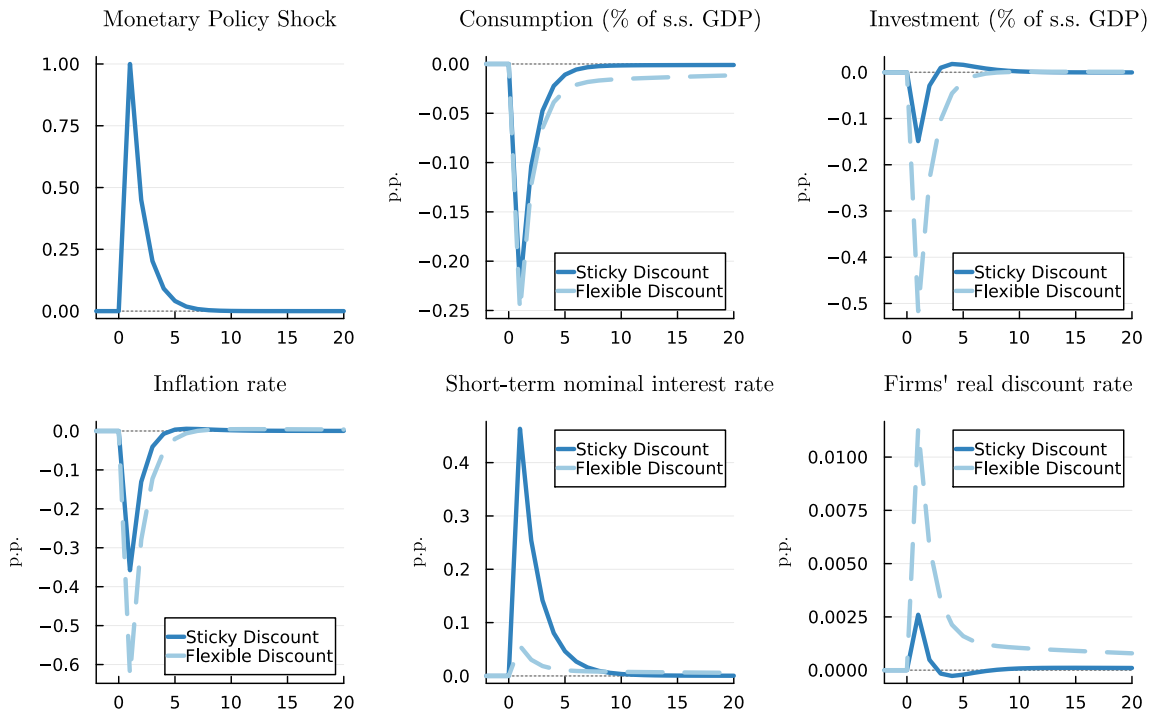


Figure 8: Impulse Responses to a Monetary Policy Shock

The figure plots impulse responses to a monetary policy shock for two different values of discount rate stickiness, $\theta \in \{0, 0.95\}$.

where

$$\mathbb{L}_t = \omega_{KL}\left[(\hat{L}_t - \hat{K}_t) - (\hat{L}_t^n - \hat{K}_t^n)\right]^2 + \omega_{IK}\left[(\hat{I}_t - \hat{K}_t) - (\hat{I}_t^n - \hat{K}_t^n)\right]^2 + \omega_C(\hat{C}_t - \hat{C}_t^n)^2$$
$$+\omega_L(\hat{L}_t - \hat{L}_t^n)^2 + \omega_\pi(\hat{\pi}_t)^2 + \omega_\chi \sum_{i\in\{p,h\}} \mu^i \frac{1}{\chi^i}(\chi_t^i - \chi^i)^2 + \omega_\delta(\hat{\delta}_t - \hat{\delta}_{t-1})^2$$

and the coefficients on the objective functions are given by

$$\omega_{KL} = \alpha(1-\alpha)Y, \quad \omega_{IK} = \phi\xi^2 K, \quad \omega_C = \sigma C, \quad \omega_L = \nu(1-\alpha)Y, \tag{31}$$

$$\omega_\pi = \varepsilon(1-\alpha)Y\frac{1}{\psi_p}, \quad \omega_\chi = \omega_C, \quad \omega_\delta = \frac{1}{\phi}K\left(\frac{1+r}{r}\right)^2 \frac{\theta}{(1-\theta)(1-\beta\theta)}, \tag{32}$$

and $\{C_t^n, L_t^n, K_{t+1}^n, I_t^n\}$ denotes the first-best allocation of consumption, labor, capital, and investment, respectively. The first-best allocation maximizes welfare subject to the resource constraint, as defined in the Supplemental Material.

The policy problem is to choose $\{\hat{\imath}_t, \widehat{coc}_t, \hat{C}_t, \hat{L}_t, \hat{K}_{t+1}, \hat{q}_t, \hat{\omega}_t, \hat{I}_t, \hat{\pi}_t, \hat{\delta}_t, \chi_t^p, \chi_t^h\}_{t=0}^{\infty}$ to minimize (30) subject to (8)-(14), (17), (20)-(21), and (23)-(27).

Sticky discount rates introduce two new sources of inefficiency. First, there is misallocation of capital across firms because two firms with the same marginal rate of transformation between consumption goods and investment end up with different investment rates. This is captured by the final term in the objective function. Second, there is misallocation between investment and consumption because the aggregate marginal rate of transformation between investment and consumption deviates from the aggregate marginal rate of substitution. This is captured by the discount rate wedges in the equilibrium condition (12).

The optimal long-run inflation target under commitment is generally non-zero whenever discount rates are sticky, as shown in Appendix F.3. In particular, the central bank's optimal response to temporary shocks is to permanently change its inflation target. The non-zero inflation target leads to long-run misallocation due to the presence of sticky prices and therefore comes at a welfare cost. Hence, with flexible discount rates, there is no reason to deviate from the zero inflation target. However, with sticky discount rates, committing to a non-zero inflation target in response to a temporary shock brings welfare benefits. Specifically, sticky discount rates introduce wedges between the marginal rates of substitution and transformation between consumption and investment goods. Changes in the long-run inflation target are highly effective at closing the wedges because long-run inflation directly affects firms' real discount rates. In contrast, moving the short-term interest rate temporarily in response to a shock is not effective at closing the wedges. Therefore, the central bank optimally reacts to temporary shocks by committing to a different inflation target when discount rates are sticky.

If the central bank cannot credibly commit, it has no incentive to deviate from the zero inflation target after the initial shocks have died out. As a result, the central bank without commitment does not permanently change its inflation target and instead uses the short-run interest rate as its policy tool.

We illustrate the findings using quantitative simulations in Table 6. The first row of Panel A shows that, under sticky discount rates with commitment, the central bank lowers the inflation target in response to an expansionary household patience shock. The lower inflation target raises real discount rates and thereby effectively closes discount rate wedges. As a consequence, the welfare loss caused by sticky discount rates (relative to flexible discount rates) is minimal in the commitment solution, as shown in the first column of Panel B of Table 6. In contrast, without

	Sticky Discount		Flexible Discount	
	Commitment	Discretion	Commitment	Discretion
A. Changes in Inflation Target				
β shock	-0.080 pp	0.000 pp	0.000 pp	0.000 pp
G shock	-0.036 pp	0.000 pp	0.000 pp	0.000 pp
A shock	0.024 pp	0.000 pp	0.000 pp	0.000 pp
B. Welfare Loss				
β shock	0.014%	0.044%	0.014%	0.014%
G shock	0.007%	0.021%	0.007%	0.007%
A shock	0.004%	0.015%	0.004%	0.004%

Table 6: Inflation Target and Welfare Loss under Optimal Monetary Policy

Panel A describes the optimal changes in the long-run inflation target in response to a one standard deviation expansionary shock to the given parameter. Panel B describes the welfare loss relative to the first best that results from a one standard deviation expansionary shock to the given parameter. The standard deviation of the shock is chosen to match GDP volatility in the data. The first best is a model with complete markets and flexible prices and wages. The welfare loss is expressed in consumption equivalent units.

commitment, the central bank can only rely on the short-term interest rate and is therefore less effective at closing discount rate wedges. As a result, the welfare loss caused by sticky discount rates is large without commitment, as shown in the second column of Table 6. Our analyses of government spending and TFP shocks lead to similar conclusions (see the impulse responses under optimal monetary policy in Figures A2, A4, and A6).

In sum, long-run inflation expectations, which in the model can be steered through the central bank's inflation target, may become part of the monetary policy toolbox when discount rates are sticky. In the model, a central bank with the power to commit can use the inflation target to largely undo the distortions caused by sticky discount rates, but a central bank without commitment fails to do so.

5 Conclusion

We present evidence that firms' discount rates (i.e., their required returns to capital) are nominally sticky, so that they do not change with expected inflation. Changes in expected inflation therefore directly impact firms' real discount rates and investment. Consistent with this mechanism, we find that expected inflation is associated with higher investment rates for firms that maintained unchanged discount rates.

Hence, sticky discount rates imply that expected inflation directly increases investment by lowering real discount rates.

Using a New Keynesian model with sticky discount rates, we study the macroeconomic implications of sticky discount rates. First, household patience shocks lead to procyclical investment-consumption comovement, in contrast to textbook models. Second, government spending does not crowd out investment but raises it. Third, changes in the central bank's long-run inflation target raise investment, but changes in the short-term interest rate are less effective than in the textbook.

We also explore the implications for optimal monetary policy. If the central bank can commit to future policy, permanent changes in the long-run inflation target are the optimal responses to temporary demand shocks, whereas the short-term interest rate plays a smaller role. This optimal policy analysis is a first step toward understanding how trade-offs change when nominal discount rates are sticky. A natural next step would be to explore how many of the conclusions carry over to a model that micro-founds the frictions generating sticky discount rates.

References

ALVAREZ, F., H. LE BIHAN, AND F. LIPPI (2016): "The Real Effects of Monetary Shocks in Sticky Price Models: A Sufficient Statistic Approach," *American Economic Review*, 106, 2817–2851.

ANDRADE, P., O. COIBION, E. GAUTIER, AND Y. GORODNICHENKO (2022): "No Firm is an Island? How Industry Conditions Shape Firms' Expectations," *Journal of Monetary Economics*, 125, 40–56.

ANGELETOS, G.-M., Z. HUO, AND K. A. SASTRY (2021): "Imperfect Macroeconomic Expectations: Evidence and Theory," *NBER Macroeconomics Annual*, 35, 1–86.

AUCLERT, A., B. BARDÓCZY, M. ROGNLIE, AND L. STRAUB (2021): "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models," *Econometrica*, 89, 2375–2408.

AUCLERT, A., M. ROGNLIE, AND L. STRAUB (2020): "Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model," NBER Working Paper 26647.

——— (2023): "The Intertemporal Keynesian Cross," *Journal of Political Economy, forthcoming*.

BARRO, R. J. AND R. G. KING (1984): "Time-Separable Preferences and Intertemporal-Substitution Models of Business Cycles," *Quarterly Journal of Economics*, 99, 817–839.

BORDALO, P., N. GENNAIOLI, AND A. SHLEIFER (2022): "Overreaction and Diagnostic Expectations in Macroeconomics," *Journal of Economic Perspectives*, 36, 223–244.

BREALEY, R. A., S. C. MYERS, AND F. ALLEN (2011): "Principles of Corporate Finance, 10th Edition," .

BUNN, P., L. S. ANAYI, N. BLOOM, P. MIZEN, G. THWAITES, AND I. YOTZOV (2022): "Firming up Price Inflation," .

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

COIBION, O., Y. GORODNICHENKO, AND S. KUMAR (2018): "How Do Firms Form Their Expectations? New Survey Evidence," *American Economic Review*, 108, 2671–2713.

COIBION, O., Y. GORODNICHENKO, AND T. ROPELE (2020): "Inflation Expectations and Firm Decisions: New Causal Evidence," *Quarterly Journal of Economics*, 135, 165–219.

DE MICHELIS, A. AND M. IACOVIELLO (2016): "Raising an Inflation Target: The Japanese Experience With Abenomics," *European Economic Review*, 88, 67–87.

DUPRAZ, S. (2023): "The Dynamic IS Curve When There Is Both Investment and Savings," Banque de France Working Paper.

FAMA, E. F. AND K. R. FRENCH (1997): "Industry Costs of Equity," *Journal of Financial Economics*, 43, 153–193.

GORMSEN, N. J. AND K. HUBER (2023a): "Corporate Discount Rates," .

——— (2023b): "Firms' Perceived Cost of Capital," .

GRAHAM, J. R. (2022): "Presidential Address: Corporate Finance and Reality," *Journal of Finance*, 77, 1975–2049.

HOUSE, C. L. AND M. D. SHAPIRO (2008): "Temporary Investment Tax Incentives: Theory with Evidence from Bonus Depreciation," *American Economic Review*, 98, 737–768.

IRELAND, P. N. (2007): "Changes in the Federal Reserve's Inflation Target: Causes

and Consequences," *Journal of Money, Credit and Banking*, 39, 1851–1882.

JENSEN, M. C. (1986): "Agency Costs of Free Cash Flow, Corporate Finance, and Takeovers," *American Economic Review*, 76, 323–329.

KAPLAN, G., B. MOLL, AND G. L. VIOLANTE (2018): "Monetary Policy According to HANK," *American Economic Review*, 108, 697–743.

KAPLAN, G., G. L. VIOLANTE, AND J. WEIDNER (2014): "The Wealthy Hand-to-Mouth," *Brookings Papers on Economic Activity*, 77–138.

KHAN, A. AND J. K. THOMAS (2008): "Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics," *Econometrica*, 76, 395–436.

KOBY, Y. AND C. WOLF (2020): "Aggregation in Heterogeneous-Firm Models: Theory and Measurement," .

KOLLER, T., M. GOEDHART, D. WESSELS, ET AL. (2015): *Valuation: Measuring and Managing the Value of Companies*, John Wiley and Sons.

LUKMANOVA, E. AND K. RABITSCH (2023): "Evidence on Monetary Transmission and the Role of Imperfect Information: Interest Rate Versus Inflation Target Shocks," *European Economic Review*, 158, 104557.

MAĆKOWIAK, B., F. MATĚJKA, AND M. WIEDERHOLT (2023): "Rational Inattention: A Review," *Journal of Economic Literature*, 61, 226–273.

MCKAY, A. AND C. WOLF (2022): "Optimal Policy Rules in HANK," .

MEYER, B., N. PARKER, AND X. SHENG (2021): "Unit Cost Expectations and Uncertainty: Firms' Perspectives on Inflation," FRB Atlanta Paper 2021-12a.

MODIGLIANI, F. AND R. A. COHN (1979): "Inflation, Rational Valuation and the Market," *Financial Analysts Journal*, 24–44.

MODIGLIANI, F. AND M. H. MILLER (1958): "The Cost of Capital, Corporation Finance and the Theory of Investment," *American Economic Review*, 48, 261–297.

MUMTAZ, H. AND K. THEODORIDIS (2017): "The Federal Reserve's Implicit Inflation Target and Macroeconomic Dynamics: A SVAR analysis," *International Economic Review*, 64, 1749–1775.

NAKAMURA, E. AND J. STEINSSON (2008): "Five Facts about Prices: A Reevaluation of Menu Cost Models," *Quarterly Journal of Economics*, 123, 1415–1464.

OTTONELLO, P. AND T. WINBERRY (2020): "Financial Heterogeneity and the Investment Channel of Monetary Policy," *Econometrica*, 88, 2473–2502.

PHILIPPON, T. (2009): "The Bond Market's Q," *Quarterly Journal of Economics*, 124,

1011–1056.

RAJAN, R., H. SERVAES, AND L. ZINGALES (2000): "The Cost of Diversity: Diversification Discount and Inefficient Investment," *Journal of Finance*, 55, 35–80.

RAMEY, V. A. (2016): "Macroeconomic Shocks and Their Propagation," *Handbook of Macroeconomics*, 2, 71–162.

REIS, R. (2006): "Inattentive Producers," *Review of Economic Studies*, 73, 793–821.

ROGERS, J. L., A. VAN BUSKIRK, AND S. L. ZECHMAN (2011): "Disclosure Tone and Shareholder Litigation," *Accounting Review*, 86, 2155–2183.

SAVIGNAC, F., E. GAUTIER, Y. GORODNICHENKO, AND O. COIBION (2022): "Firms' Inflation Expectations: New Evidence from France," *Journal of European Economic Association, forthcoming*.

SCHARFSTEIN, D. S. AND J. C. STEIN (2000): "The Dark Side of Internal Capital Markets: Divisional Rent-Seeking and Inefficient Investment," *Journal of Finance*, 55, 2537–2564.

SMETS, F. AND R. WOUTERS (2007): "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 97, 586–606.

STEIN, J. C. (1996): "Rational Capital Budgeting in an Irrational World," *Journal of Business*, 69, 429–455.

TRAHAN, E. A. AND L. J. GITMAN (1995): "Bridging the Theory-Practice Gap in Corporate Finance: A Survey of Chief Financial Officers," *Quarterly Review of Economics and Finance*, 35, 73–87.

URIBE, M. (2022): "The Neo-Fisher Effect: Econometric Evidence from Empirical and Optimizing Models," *AEJ: Macroeconomics*, 14, 133–162.

WINBERRY, T. (2021): "Lumpy Investment, Business Cycles, and Stimulus Policy," *American Economic Review*, 111, 364–396.

WOLF, C. K. (2023): "The Missing Intercept: A Demand Equivalence Approach," *American Economic Review*, 113, 2232–2269.

Online Appendix to "Sticky Discount Rates"

Appendix A Details on Discount Rate Formulations

The investment problem of a firm with discount rate δ_t can equivalently be represented in sequential form, where the firm solves:

$$V_t^I(k_t, \delta_t) = \max_{\{k_{s+1}, I_s\}_{s=0}^{\infty}} \sum_{s=t}^{\infty} \frac{1}{(1+\delta_t)^s} \left[\Omega_{t+s}(k_{t+s}) - P_{t+s}(I_{t+s} + \Phi(I_{t+s}, k_{t+s}))\right] \quad \text{(A1)}$$

$$\text{s.t.} \quad k_{t+s+1} = (1-\xi)k_{t+s} + I_{t+s}.$$

The optimality condition for investment in period t is:

$$\frac{1}{1+\delta_t} \sum_{s=1}^{\infty} \left(\frac{1-\xi}{1+\delta_t}\right)^{s-1} \left\{\Omega'_{t+s}(k_{t+s}) - P_{t+s}\Phi_k(I_{t+s}, k_{t+s})\right\} = P_t(1 + \Phi_I(I_t, k_t)). \quad \text{(A2)}$$

We now formally explain why using one discount rate δ_t that applies to all future periods allows the firm to generate the same investment decision as using a sequence of discount rates $\{d_s\}_{s=t}^{\infty}$, conditional on investment being positive, $\iota_s \equiv I_s/K_s \geq 0$ for all $s \geq t$. With constant returns to scale in production and in adjustment costs, $\Omega_t(k) = \omega_t k$ and $\Phi(I, k) = \varphi(I/k)k$, one can rewrite the optimality condition (A3) as

$$\frac{1}{1+\delta_t} \sum_{s=1}^{\infty} \left(\frac{1-\xi}{1+\delta_t}\right)^{s-1} \left\{\omega_{t+s} - P_{t+s}\left(-\varphi'(\iota_{t+s})\iota_{t+s} + \varphi(\iota_{t+s})\right)\right\} = P_t(1 + \varphi'(\iota_t)). \quad \text{(A3)}$$

Note that the left-hand side is strictly decreasing in δ_t as long as the left-hand side is in the relevant range from 0 to infinity. Assume that $\varphi(\iota)$ satisfies $\varphi'(0) = 0$ and $\lim_{\iota \to \infty} \varphi'(\iota) = \infty$. Then, for any given $\{\iota_{t+s}\}_{s=1}^{\infty}$, firms can achieve any desired investment rate $\iota_t \geq 0$ by choosing an appropriate δ_t. Firms can therefore also achieve the investment rate that would be chosen using a sequence of discount rates $\{d_s\}_{s=t}^{\infty}$.

The intuitive interpretation of a discount rate is that $1 + \delta_t$ is the marginal return to capital in period t. We now present the optimality condition in steady

state, which intuitively illustrates that the steady-state discount rate δ defines the marginal return. After imposing steady state, (A3) simplifies to:

$$1 + \delta = 1 - \xi + \omega, \tag{A4}$$

where the right-hand side is the marginal return to capital that consists of capital in the next period net of depreciation $(1 - \xi)$ plus the per-period real net profits from a unit of capital $\omega \equiv \frac{1}{P_t k} \Omega(k)$.

Appendix B Proof of Proposition 1

Due to the constant returns to scale assumptions, each individual firm's problem is independent of its size, k_t. Let $\iota_t = I_t / k_t$ be the investment rate of the firm. It is easy to verify that the firm's value functions is linear in the capital stock:

$$V_t^I(k_t, \delta) = v_t^I(\delta) P_t k_t, \tag{A5}$$

where $v_t^d(\delta)$ denotes the real marginal value of unit capital. It solves the following Bellman equation:

$$v_t^I(\delta) = \max_{\iota_t} \omega_t - \iota_t - \varphi(\iota_t) + \mathbb{E}_t \frac{1 + \pi_{t+1}}{1 + \delta} \left\{ (1 - \xi) + \iota_t \right\} v_{t+1}^I(\delta), \tag{A6}$$

where $1 + \pi_{t+1} \equiv P_{t+1} / P_t$ is the gross inflation rate, ω_t are real profits from unit capital,

$$\omega_t \equiv \max_l \frac{1}{P_t} (p_t F_t(1, l) - W_t l), \tag{A7}$$

and $\varphi(\iota) \equiv \Phi(\iota, 1)$. The first-order optimality condition for investment is

$$1 + \varphi'(\iota_t) = \mathbb{E}_t \frac{1 + \pi_{t+1}}{1 + \delta} v_{t+1}^I(\delta) \tag{A8}$$

Likewise, the firm's financial market value is also linear in capital:

$$V_t^a(k) = v_t^a P_t k, \quad V_t^n(k, \delta) = v_t^n(\delta) P_t k, \tag{A9}$$

where v_t^a and v_t^n solve the following recursion

$$v_t^a = \max_{\delta^*} \omega_t - \iota_t - \varphi(\iota_t) + \mathbb{E}_t \frac{1+\pi_{t+1}}{1+i_t} \{(1-\xi)+\iota_t\} [\theta v_{t+1}^n(\delta^*) + (1-\theta) v_{t+1}^a] \tag{A10}$$

$$\text{s.t.} \quad \iota_t = \bar{\iota}_t(\delta^*) \tag{A11}$$

and

$$v_t^n(\delta^*) = \omega_t - \iota_t - \varphi(\iota_t) + \mathbb{E}_t \frac{1+\pi_{t+1}}{1+i_t} \{(1-\xi)+\iota_t\} [\theta v_{t+1}^n(\delta^*) + (1-\theta) v_{t+1}^a] \tag{A12}$$

$$\text{s.t.} \quad \iota_t = \bar{\iota}_t(\delta^*) \tag{A13}$$

The first-order optimality condition for the choice of the discount rate is

$$\begin{aligned} \mathbb{E}_t \left[\frac{1+\pi_{t+1}}{1+i_t} [\theta v_{t+1}^n(\delta^*) + (1-\theta) v_{t+1}^a] - (1+\varphi'(\iota_t)) \right] \frac{d\bar{\iota}_t}{d\delta_t^*} \\ + \mathbb{E}_t \frac{1+\pi_{t+1}}{1+i_t} \{(1-\xi)+\iota_t\} \theta \frac{d v_{t+1}^n(\delta^*)}{d\delta^*} = 0. \end{aligned} \tag{A14}$$

A first-order approximation of (A14) around the deterministic steady state gives

$$\begin{aligned} \mathbb{E}_t \Bigg[\frac{d\bar{\iota}}{d\delta^*} \Bigg(v \frac{1}{1+r} \Bigg[d\pi_{t+1} - d\ln(1+i_t) + d\ln v_{t+1}^a + \theta \underbrace{v^{n\prime}(\delta^*)}_{=0} \frac{1}{1+r} d\ln(1+\delta_t^*) \Bigg] - \varphi''(\iota) d\iota_t \Bigg) \\ + \underbrace{\left[\frac{1}{1+r} v - (1+\varphi'(\iota)) \right]}_{=0} d\left(\frac{d\bar{\iota}_t}{d\delta_t^*} \right) \\ + \underbrace{v^{n\prime}(\delta^*)}_{=0} d\left(\frac{1+\pi_{t+1}}{1+i_t} \{(1-\delta^*)+\iota_t\} \theta \right) \\ + \frac{1}{1+r} \theta d\left(\frac{d v_{t+1}^n(\delta^*)}{d\delta^*} \right) \Bigg] = 0, \end{aligned} \tag{A15}$$

which, in turn, simplifies as follows, since many terms disappear owing to envelope

conditions:

$$\mathbb{E}_t\left[\frac{d\bar{\iota}}{d\delta^*}\left(v\frac{1}{1+r}\left[d\pi_{t+1}-d\ln(1+i_t)+d\ln v^a_{t+1}\right]-\varphi''(\iota)d\iota_t\right)\right.$$
$$\left.+\frac{1}{1+r}\theta d\left(\frac{dv^n_{t+1}(\delta^*)}{d\delta^*}\right)\right]=0. \tag{A16}$$

Linearizing (A8) gives

$$\varphi''(\iota)d\iota=\frac{1}{1+r}v\left(d\pi_{t+1}-d\ln(1+\delta^*_t)+d\ln v^d_{t+1}(\delta^*)\right). \tag{A17}$$

Combining (A16) and (A17),

$$\mathbb{E}_t\left[\frac{d\bar{\iota}}{d\delta^*}v\frac{1}{r}\left[d\ln(1+\delta^*_t)-d\ln(1+coc_t)\right]+\frac{1}{1+r}\theta d\left(\frac{dv^n_{t+1}(\delta^*)}{d\delta^*}\right)\right]=0, \tag{A18}$$

where

$$d\ln(1+coc_t)=\frac{r}{1+r}\mathbb{E}_t\sum_{s=t}^{\infty}\frac{1}{(1+r)^{s-t}}d\ln(1+i_s) \tag{A19}$$

denote the long-run cost of capital, which is a weighted average of future short-term interest rates with weights that depend on the steady-state interest rate.

Following the same steps as above, we can show that

$$d\left(\frac{dv^n_t(\delta^*_t)}{d\delta^*}\right)=\mathbb{E}_t\left[\frac{d\bar{\iota}}{d\delta^*}v\frac{1}{r}\left[d\ln(1+\delta^*_t)-d\ln(1+coc_t)\right]+\frac{1}{1+r}\theta d\left(\frac{dv^n_{t+1}(\delta^*_t)}{d\delta^*}\right)\right]. \tag{A20}$$

Using (A20) and hat notation, we can iterate (A18) forward to obtain the expression for the optimal choice of discount rate:

$$\hat{\delta}^*_t=\frac{1+r-\theta}{1+r}\sum_{s=t}^{\infty}\left(\frac{\theta}{1+r}\right)^{s-t}\widehat{coc}_s, \tag{A21}$$

which we can write recursively as

$$\hat{\delta}_t^* = \frac{1+r-\theta}{1+r}\widehat{coc}_t + \frac{\theta}{1+r}\hat{\delta}_{t+1}^*. \tag{A22}$$

The average discount rate in the economy, which we denote as δ_t, evolves according to

$$\hat{\delta}_t = \theta\hat{\delta}_{t-1} + (1-\theta)\hat{\delta}_t^*. \tag{A23}$$

Appendix C Details on the New Keynesian Model

We describe the details of the model here (apart from the firm problem). We present a more general model by allowing both prices and wages to be sticky and letting hand-to-mouth households earn non-labor income. Throughout, we log-linearize around the steady state of the economy, where the gross interest rate is $1/\beta$ and household bond holdings are zero.

Let ω^p and ω^h denote real per-capita labor and dividend income for permanent-income households and hand-to-mouth households, respectively. The budget constraint of the permanent-income household in log-linearized form is

$$\sum_{s=0}^{\infty}\beta^s\hat{C}_{t+s}^p = \frac{1}{\beta C}b_{t-1} + \sum_{s=0}^{\infty}\beta^s\hat{\omega}_{t+s}^p, \tag{A24}$$

where b_{t-1} denotes the bond-holdings. The Euler equation is

$$\hat{C}_t^p = -(1/\sigma)\hat{\beta}_{t+1} - (1/\sigma)\hat{r}_{t+1} + \hat{C}_{t+1}^p. \tag{A25}$$

Combining (A24) and (A25),

$$\hat{C}_t^p = -(1/\sigma)\beta\sum_{s=0}^{\infty}\beta^s\left[\hat{\beta}_{t+1+s} + \hat{r}_{t+1+s}\right] + (1-\beta)\frac{1}{\beta C}b_{t-1} + (1-\beta)\sum_{s=0}^{\infty}\beta^s\hat{\omega}_{t+s}^p. \tag{A26}$$

The consumption of a hand-to-mouth household is

$$\hat{C}_t^h = \hat{\omega}_t^h. \tag{A27}$$

Adding up (A26) and (A27), aggregate consumption can be expressed as

$$\hat{C}_t = \frac{\mu\omega^h}{C}\hat{C}_t^h + \frac{(1-\mu)\omega^p}{C}\hat{C}_t^p \tag{A28}$$

$$= \frac{(1-\mu)\omega^p}{C}(1-\beta)\sum_{s=0}^{\infty}\beta^s\hat{\omega}_{t+s}^p + \frac{\mu\omega^h}{C}\hat{\omega}_t^h + \frac{(1-\mu)\omega^p}{C}\left(-(1/\sigma)\sum_{s=0}^{\infty}\beta^s\left[\hat{\beta}_{t+1+s} + \hat{r}_{t+1+s}\right] + (1-\beta)\frac{1}{\beta C}b_{t-1}\right). \tag{A29}$$

Since profits are zero and the production function is Cobb-Douglas in labor and capital, labor income for each household in each period is $(1-\alpha)(C_t + I_t + G_t)$, and aggregate dividend income in each period is $\alpha(Y_t + I_t + G_t) - I_t$. Therefore,

$$\omega^p = (1-\alpha)\left(C_t + G_t + I_t\right) + \frac{1}{1-\mu}(1-\mu s^h)(\alpha\left(C_t + G_t + I_t\right) - I) - T_t \tag{A30}$$

$$\omega^h = (1-\alpha)\left(C_t + G_t + I_t\right) + s^h(\alpha\left(C_t + G_t + I_t\right) - I_t) - T_t, \tag{A31}$$

which comes from the fact that the hand-to-mouth households with measure μ own s^h of firm shares and the remaining dividend income accrues to permanent income households. Log-linearizing,

$$\omega^p\hat{\omega}_t^p = \left[(1-\alpha) + \frac{1}{1-\mu}(1-\mu s^h)\alpha\right](C\hat{C}_t + G_t) + (1-\alpha)\left[1 - \frac{1}{1-\mu}(1-\mu s^h)\right]I\hat{I}_t - T_t \tag{A32}$$

$$\omega^h\hat{\omega}_t^h = \left[(1-\alpha) + \alpha s^h\right](C\hat{C}_t + G_t) + (1-\alpha)(1-s^h)I\hat{I}_t - T_t. \tag{A33}$$

Substituting (A32) and (A33) into (A29),

$$\hat{C}_t = (1-\beta)\sum_{s=0}^{\infty}\beta^s\left[(1-\alpha)+\frac{1}{1-\mu}(1-\mu s^h)\alpha\right](1-\mu)\left(\hat{C}_{t+s}+\frac{1}{C}G_{t+s}\right) \tag{A34}$$

$$+\left[(1-\alpha)+\alpha s^h\right]\mu\left(\hat{C}_t+\frac{1}{C}G_t\right) \tag{A35}$$

$$+\frac{1-\mu}{\omega}(1-\beta)\sum_{s=0}^{\infty}\beta^s\left((1-\alpha)\left[1-\frac{1}{1-\mu}(1-\mu s^h)\right]I\hat{I}_{t+s}-T_{t+s}\right) \tag{A36}$$

$$+\frac{\mu}{C}\left((1-\alpha)(1-s^h)I\hat{I}_t-T_t\right) \tag{A37}$$

$$+\frac{(1-\mu)\omega^p}{C}\left(-(1/\sigma)\beta\sum_{s=0}^{\infty}\beta^s\left[\hat{\beta}_{t+1+s}+\hat{r}_{t+1+s}\right]+(1-\beta)\frac{1}{\beta C}b_{t-1}\right). \tag{A38}$$

The government budget constraint and the market clearing constraint for bonds imply that $b_{t-1}-\beta b_t = T_t - G_t$. Computing $\hat{C}_t - \beta\hat{C}_{t+1}$ and rearranging, we obtain

$$\begin{aligned}\hat{C}_t = \mathbb{E}_t\Big[&\hat{C}_{t+1}-\zeta_r(1/\sigma)\left[\hat{\beta}_{t+1}+\hat{r}_{t+1}\right]+\zeta_I\left[\hat{I}_t-\hat{I}_{t+1}\right]\\&+\zeta_G^0 G_t-\zeta_G^1 G_{t+1}-\zeta_T^0 T_t+\zeta_T^1 T_{t+1}\Big],\end{aligned} \tag{A39}$$

where

$$\zeta_r = 1-\frac{(1-\alpha)\mu(1-s^h)\frac{I}{C}}{1-\mu+\mu\alpha(1-s^h)} \tag{A40}$$

$$\zeta_I = \frac{\mu(1-\alpha)(1-s^h)\frac{I}{C}}{1-\mu+\mu\alpha(1-s^h)} \tag{A41}$$

and

$$\zeta_G^0 = \frac{1}{1-\mu+\mu\alpha(1-s^h)}\frac{1}{\beta}\frac{1}{C}\left(1-\beta\left[1-\mu+\mu\alpha(1-s^h)\right]\right)-\zeta_r\frac{1-\beta}{\beta}\frac{1}{C} \tag{A42}$$

$$\zeta_G^1 = \frac{1}{1-\mu+\mu\alpha(1-s^h)}\left[1-\alpha(1-s^h)\right]\mu\frac{1}{C} \tag{A43}$$

$$\zeta_T^0 = \frac{1}{1-\mu+\mu\alpha(1-s^h)}\frac{1}{\beta}\frac{1}{C}\left(1-\beta(1-\mu)\right)-\zeta_r\frac{1-\beta}{\beta}\frac{1}{C} \tag{A44}$$

$$\zeta_T^1 = \frac{1}{1-\mu+\mu\alpha(1-s^h)}\frac{\mu}{C}. \tag{A45}$$

Setting $s_h = 0$ yields the expressions in the main text.

The retailer r purchases intermediate goods from firms at price p_t, converts them into differentiated products, and sells them at price $P_t(r)$. The final good is a CES basket of differentiated products

$$Y_t = \int_0^1 \left(y_t(r)\right)^{\frac{\varepsilon-1}{\varepsilon}}\right)^{\frac{\varepsilon}{\varepsilon-1}} dr, \tag{A46}$$

where $\varepsilon > 1$ is the elasticity of substitution across products. Demand for each variety is

$$y_t(r) = \left(\frac{P_t(r)}{P_t}\right)^{-\varepsilon} Y_t, \tag{A47}$$

where $P_t(r) = \left[\int_0^1 P_t(r)^{1-\varepsilon} dr\right]^{\frac{1}{1-\varepsilon}}$ is the price index. The price can be adjusted with probability $1-\gamma_p$, as in Calvo (1983). The government subsidizes the production of the final goods producer at a rate $1+\tau^p = \frac{\varepsilon}{\varepsilon-1}$. The retailer with an opportunity to adjust its price faces the following price-setting problem:

$$\max_{P_t(r)} \mathbb{E}_t \sum_{s=0}^{\infty} M_{t,s} (\gamma_p)^s \left[P_t(r) - p_s\right] y_s(r) \tag{A48}$$

where $M_{t,s} = \prod_{s=0}^{\infty} \frac{1}{1+r_{t+s}}$ subject to (A47). Linearzing the optimality conditions and aggregating across retailers result in the standard New Keynesian Phillips curve:

$$\hat{\pi}_t = \frac{(1-\gamma_p)(1-\beta\gamma_p)}{\gamma_p}(\hat{p}_t - \hat{P}_t) + \beta \mathbb{E}_t \hat{\pi}_{t+1}, \tag{A49}$$

where $\hat{\pi}_t = \hat{P}_t - \hat{P}_{t-1}$. Firm labor demand is

$$\hat{A}_t + \alpha(\hat{K}_t - \hat{L}_t) = \hat{W} - \hat{P}_t. \tag{A50}$$

Using (A50), we can rewrite (A49) as

$$\hat{\pi}_t = \frac{(1-\gamma_p)(1-\beta\gamma_p)}{\gamma_p}(\hat{W}_t - \hat{P}_t - \alpha(\hat{K}_t - \hat{L}_t)) + \beta \mathbb{E}_t \hat{\pi}_{t+1} \tag{A51}$$

Households supply labor through a continuum of labor unions, $\ell \in [0,1]$. Each

labor union converts labor into differentiated labor services, $L_t(\ell)$. Aggregate labor services are given by a CES basket with an elasticity of substitution $\varepsilon_w > 1$:

$$L_t = \left[\int_0^1 L_t(\ell)^{\frac{\varepsilon_w - 1}{\varepsilon_w}} d\ell \right]^{\frac{\varepsilon_w}{\varepsilon_w - 1}}. \tag{A52}$$

The demand for each labor variety is

$$L_t(\ell) = (W_t(\ell)/W_t)^{-\varepsilon_w} L_t. \tag{A53}$$

Labor unions set wages of differentiated labor services but can only adjust with probability $1 - \gamma_w \in [0, 1]$. They maximize the weighted average of the utility of permanent income households and hand-to-mouth households. A labor union chooses wage $W_t(\omega)$ to maximize the weighted average:

$$\mathbb{E}_t \sum_{s=0}^{\infty} (\beta \gamma_w)^{s-t} \left[\tilde{\vartheta}^h u(C_s^h) + \tilde{\vartheta}^p u(C_s^p) - v(l_s) \right],$$

with $\sum_{i \in \{p,h\}} \vartheta^i = 1$ subject to $l_s = \int_0^1 l_s(\ell) d\ell$, $W_t L_t = \int_0^1 W_t(\ell) L_t(\ell) d\ell$, and (18)-(A53).

Solving the above and taking a first-order approximation around the steady state yields the following New Keynesian wage Phillips Curve:

$$\hat{\pi}_t^w = \psi_w \left[\sigma \hat{C}_t + \sigma \sum_i \vartheta^i (\chi_t^i - \chi^i)/\chi^i + \nu \hat{L}_t - \hat{W}_t + \hat{P}_t \right] + \frac{1}{1+r} \mathbb{E}_t \hat{\pi}_{t+1}^w, \tag{A54}$$

where $\vartheta^i = \frac{\tilde{\vartheta}^i u'(C^i)}{\tilde{\vartheta}^h u'(C^h) + \tilde{\vartheta}^p u'(C^p)}$, $\chi_t^i \equiv C_t^i / C_t$ is the consumption of household type $i \in \{p, h\}$, relative to aggregate consumption, and $\psi_w \equiv (1 - \gamma_w)(1 - \beta \gamma_w)/\gamma_w$.

To a first-order approximation, the goods market clearing condition can be expressed as

$$\frac{C}{Y} \hat{C}_t + \frac{I}{Y} \hat{I}_t + \frac{1}{Y} \hat{G}_t = \hat{A}_t + \alpha \hat{K}_t + (1 - \alpha) \hat{L}_t. \tag{A55}$$

Appendix D Model with Inflation Indexation

We modify the baseline model by allowing retailers to index prices to the central bank's inflation target. We assume that retailers adjust prices one-to-one with the inflation target, $\bar{\pi}_t$, irrespective of whether the Calvo friction gives them an opportunity to adjust (which happens with probability $1-\gamma_p$). The New Keynesian Phillips curve becomes

$$\hat{\pi}_t - \bar{\pi}_t = \psi_p \left[\sigma(\sum_i \vartheta^i \frac{1}{\chi^i} d\chi_t^i) + \sigma \hat{C}_t + \nu \hat{L}_t - \hat{A}_t - \alpha(\hat{K}_t - \hat{L}_t) \right] + \frac{1}{1+r} \mathbb{E}_t[\hat{\pi}_{t+1} - \bar{\pi}_{t+1}]. \tag{A56}$$

The equilibrium definition of this economy replaces (20) with (A56).

Starting from a steady state with a zero inflation target, we consider a shock to the inflation target at $t = 0$. Specifically, $\bar{\pi}_t = d\bar{\pi}$ at $t \geq 0$ and $\bar{\pi}_t = 0$ at $t < 0$. With flexible discount rates, $\theta = 0$, it is straightforward to verify that a permanent increase in the inflation target at $t = 0$ does not affect the allocation, while the nominal interest rate and firms' discount rates adjust one-to-one with the inflation target, $i_t = \delta_t = d\bar{\pi}$ for $t \geq 0$. With sticky discount rates, $\theta > 0$, the shock still has substantial real consequences, as shown in Figure A1.

Appendix E Business Cycle Statistics

Table A1 compares business cycle statistics generated by the model to post-1959 US data. In the model, we allow only one type of shock to generate the moments in each column. We alternate between a shock to household patience (β), government spending (G), the inflation target ($\bar{\pi}$), the short-term nominal interest rate (i), and TFP (A).

In the model with sticky discount rates, patience, government spending, and inflation target shocks all generate volatility and correlation statistics that are of the same sign as the US data. In contrast, the model with flexible discount rates produces the opposite sign for some key correlations, including the output-consumption and output-inflation correlations for patience shocks and the output-consumption and output-investment correlation for government spending shocks. With sticky discount rates, short-term interest rate shocks have more muted effects on the real

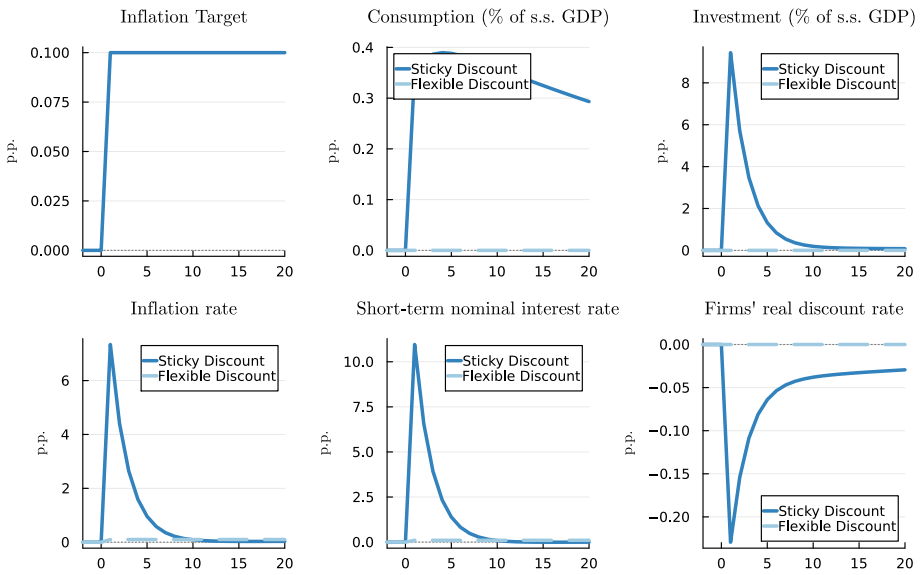


Figure A1: Inflation Indexation: Impulse Responses to an Inflation Target Shock

The figure plots the impulse response to an inflation target shock for two different values of discount rate stickiness, $\theta \in \{0, 0.95\}$.

economy and TFP shocks generate a negative correlation between output and consumption.

	Data	Sticky Discount Rates					Flexible Discount Rates				
		β	G	$\bar{\pi}$	i	A	β	G	$\bar{\pi}$	i	A
A. Volatility											
std(ΔY) $\times 100$	0.81	0.81	0.81	0.81	0.81	0.81	0.06	0.07	0.00	1.59	0.23
std(ΔC)/std(ΔY)	0.80	0.13	0.05	0.04	0.71	0.07	0.80	0.45	0.11	0.39	0.50
std(ΔI)/std(ΔY)	4.88	5.06	4.02	5.45	2.38	5.90	9.29	10.06	5.38	3.80	3.32
std(ΔL)/std(ΔY)	0.98	1.45	1.44	1.46	1.44	1.99	1.47	1.40	1.43	1.45	0.55
std($\Delta \pi$)/std(ΔY)	0.64	0.83	0.74	0.76	0.95	0.98	0.55	0.81	17.95	0.83	0.15
std(Δi)/std(ΔY)	1.08	1.25	1.10	1.13	1.14	1.47	0.83	1.22	18.07	0.07	0.23
B. Correlation											
corr(ΔY, ΔC)	0.65	0.93	0.95	0.87	1.00	-0.94	-0.99	-0.97	0.49	1.00	1.00
corr(ΔY, ΔI)	0.81	1.00	0.99	1.00	0.99	1.00	1.00	-1.00	1.00	1.00	1.00
corr(ΔY, ΔL)	0.66	1.00	1.00	1.00	1.00	1.00	0.99	0.99	0.98	1.00	-1.00
corr(ΔY, $\Delta \pi$)	0.15	1.00	1.00	1.00	1.00	1.00	-1.00	1.00	1.00	1.00	-1.00
corr(ΔY, Δi)	0.31	1.00	1.00	1.00	-0.99	1.00	-1.00	1.00	1.00	-1.00	-1.00

Table A1: Business Cycle Moments

Panel A describes the standard deviation of each variable in the data, in the model with sticky discount rates, and in the model with flexible discount rates. Panel B describes correlations between the variables. $\Delta Y, \Delta C, \Delta I, \Delta L$ denote log changes in real GDP, real consumption, real investment, and total hours worked. $\Delta\pi$ and Δi denote changes in the inflation rate and the nominal interest rate. The columns β, G, $\bar{\pi}$, i, and A correspond to a model with only a patience shock, a government spending shock, an inflation target shock, a monetary policy shock, and a TFP shock, respectively. We choose the volatility of each shock to match the volatility of GDP in the model with sticky discount rates. The reported model moments are analytical second moments based on the first-order impulse response function for each shock, as in Auclert et al. (2021). We use quarterly US data from 1959 to 2019. Real GDP is BEA code A191RX, real consumption is BEA code DPCERX, real investment is BEA code A006RX, total hours worked are from the BLS, inflation is the CPI provided by the BLS, and the nominal interest rate is the federal funds rate.

Appendix F Optimal Monetary Policy

Appendix F.1 Interpretation of the Optimal Policy Problem

The objective function (30) in the main text consists of several standard terms and a new term $\omega_\delta(\hat{\delta}_t - \hat{\delta}_{t-1})^2$. We formally derive the loss function in the Supplemental Material and explain the terms here.

The first five terms are standard. They imply that the central bank would like to close the gap between the equilibrium and the first best frictionless economy in terms of the aggregate capital-labor ratio (the first term), investment rate (the second term), consumption (the third term), labor (the forth term), and wage inflation (the fifth term). The sixth term is new and comes from the misallocation in investment across firms due to sticky discount rates. When discount rates are sticky, a change in firms' average discount rate introduces dispersion in discount rates and thereby investment rates. Since all firms have the same technology, the dispersion in investment rates implies misallocation and a welfare loss. The central bank therefore has an incentive to smooth discount rates. The seventh term $\omega_C \sum_i \mu^i \frac{1}{\chi^i}(\chi_t^i - \chi^i)^2$ is standard in the literature on optimal monetary policy with heterogeneous households and captures the central bank's incentive to smooth relative consumption across the two types of households.

What are the sources of inefficiency that the central bank aims to overcome in this economy? As in the textbook New Keynesian model, wage stickiness implies two sources of inefficiency. First, the dispersion in wages creates misallocation in the labor services provided by different households for a given total labor supply, which is captured by the fifth term in the objective function (30). Second, wage stickiness implies that the marginal rate of transformation (MRT) between labor and consumption goods and the marginal rate of substitution (MRS) between the two are not equalized. This leads to an inefficient level of aggregate labor supply. With only wage stickiness and a representative agent, monetary policy can entirely undo these inefficiencies at the same time, which is known as the "divine coincidence."[A1] While the textbook New Keynesian model often abstracts from investment, the divine coincidence remains true in textbook environments with investment. In addition, the model contains an incomplete financial market, which is another source of

[A1] In our environment with investment and sticky wages, divine coincidence means that wage inflation is zero, $\pi_t^w = 0$, but that price inflation may not be zero.

inefficiency and which leads to heterogeneous consumption movements across the two different household types. This inefficiency is reflected in the seventh term in the objective function mentioned above.

Sticky discount rates introduce two new sources of inefficiency: first, misallocation of investment across firms, and second, misallocation between investment and consumption. Misallocation across firms comes from the fact that two firms with the same marginal rate of transformation between consumption goods and investment goods end up with different investment rates when discount rates are sticky. Misallocation between investment and consumption comes from the fact that the aggregate marginal rate of transformation between investment and consumption deviates from the aggregate marginal rate of substitution between consumption and investment when discount rates are sticky. This is reflected in the movements of discount rate wedges.

Appendix F.2 Optimal Monetary Policy Solution

The central bank conducts monetary policy trying to address all of the above inefficiencies at the same time. In the commitment solution, the planner takes into account the effect of its future actions on future private sector expectations. In the discretion solution, the central bank takes private sector expectations as given.

We show in Appendix F.3 that inflation satisfies

$$\omega_\pi \hat{\pi}_t + \lambda_{p,t} - \lambda_{p,t-1} - r\lambda_{l,t} = 0, \tag{A57}$$

where $\lambda_{p,t}$ is the Lagrangian multiplier on the New Keynesian Phillips curve, (A54), and $\lambda_{l,t}$ is the Lagrangian multiplier on the frictionless nominal discount rate, (9), which in turn follows

$$\lambda_{l,t} = \lambda_{l,t-1} + \tilde{\theta}\varsigma_t, \tag{A58}$$

where ς_t is a function of a sequence of exogenous shocks. We show that $\tilde{\theta} \neq 0$ if and only if $\theta \neq 0$. This implies that the Lagrangian multiplier, $\lambda_{l,t}$ is a random walk process whenever discount rates are sticky, $\theta \neq 0$. In turn, (A57) implies that inflation does not converge back to zero as $t \to \infty$, even if the shocks $\boldsymbol{\varepsilon}_t \equiv (\varepsilon_t^\beta, \varepsilon_t^A, \varepsilon_t^G) \to \mathbf{0}$. Therefore, optimal monetary policy involves permanent changes

in the inflation target even when shocks are temporary.

In contrast, with discretion, we show in Appendix F.3 that the long-run inflation target is always zero regardless of the discount rate stickiness parameter, θ.

What is the economic intuition behind this result? The higher inflation target comes at a welfare cost because permanent inflation leads to more misallocation across households and more inefficient aggregate labor supply in the long run. This welfare cost implies that the inflation target reverts to zero in the textbook New Keynesian model. However, in the model with sticky discount rates, this cost is undone by the benefits of a higher inflation target. Specifically, sticky discount rates introduce wedges between the MRS and MRT between consumption and investment goods. Moving short-term interest rates temporarily in response to a shock is not effective at closing the wedges. In contrast, changes in the long-run inflation target are highly effective at closing the wedges because long-run inflation directly affects firms' real discount rates. Therefore, the central bank optimally reacts to a shock by committing to a different inflation target. This argument crucially relies on the central bank's ability to commit. Without commitment, after shocks have died out, there is no reason to deviate from the zero long-run inflation target.

The result that changing the permanent inflation target is optimal is not driven by any modifications to the central bank's objective function. Rather, it comes from the presence of discount rate wedges that enter the equilibrium conditions. We illustrate this point by showing that even a dual-mandate policymaker, who minimizes a loss function that only depends on aggregate quantities and inflation, reaches qualitatively the same optimal policy solution. The dual-mandate policymaker minimizes (30) where $\mathbb{L}_t$ is replaced with

$$\begin{aligned}\mathbb{L}_t^d = \omega_{KL}\left[(\hat{L}_t - \hat{K}_t) - (\hat{L}_t^n - \hat{K}_t^n)\right]^2 + \omega_{IK}\left[\hat{I}_t - \hat{K}_t - (\hat{I}_t^n - \hat{K}_t^n)\right]^2 \\ + \omega_C(\hat{C}_t - \hat{C}_t^n)^2 + \omega_L(\hat{L}_t - \hat{L}_t^n)^2 + \omega_\pi(\hat{\pi}_t)^2.\end{aligned} \tag{A59}$$

With the dual-mandate objective, the central bank no longer directly cares about discount rates or inequality. We show below in Appendix F.3 that the optimal long-run inflation rate remains non-zero under commitment.

Panel A of Table 6 reports welfare losses for each shock, relative to the first-best outcome. The welfare loss from sticky discount rates is minimal as long as the central bank can commit. In contrast, without commitment, sticky discount rates induce substantial welfare losses, particularly for a patience shock or a government

spending shock. Panel B explains why. After a patience or government spending shock, the central bank would like to commit to changing its inflation target by a sizable amount, while the desired change is small after a TFP shock. But since the central bank cannot commit to a changed inflation target int the discretion solution, the welfare loss is substantial.

Appendix F.3 Formal Characterization of Optimal Monetary Policy Solution

We first describe the case with commitment. The optimal monetary policy problem is

$$\min_{\{\hat{i}_t,\hat{C}_t,\hat{L}_t,\hat{K}_{t+1},\hat{I}_t,\hat{\pi}_t,\hat{\delta}_t,\chi_t^p,\chi_t^h,\widehat{coc}_t\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \frac{1}{2} \mathbb{L}_t, \tag{A60}$$

where $\mathbb{L}_t$ is defined in (31) subject to (with the relevant Lagrangian multiplier given on the left-hand side)

$$\lambda_{rc}: \quad 0 = C\hat{C}_t + I\hat{I}_t + Y\hat{G}_t - Y\hat{A}_t - \alpha Y\hat{K}_t - (1-\alpha)Y\hat{L}_t \tag{A61}$$

$$\lambda_K: \quad \hat{K}_{t+1} = (1-\xi)\hat{K}_t + \xi\hat{I}_t \tag{A62}$$

$$\lambda_I: \quad \hat{I}_t - \hat{K}_t = \frac{1}{\xi\phi}\left(\hat{q}_t - \frac{1+r}{r}(\hat{\delta}_t - \widehat{coc}_t)\right) \tag{A63}$$

$$\lambda_\delta: \quad \hat{\delta}_t = \theta\hat{\delta}_{t-1} + (1-\theta)\hat{\delta}^*_t \tag{A64}$$

$$\lambda_*: \quad \hat{\delta}^*_t = \frac{1+r-\theta}{1+r}\widehat{coc}_t + \frac{\theta}{1+r}\hat{\delta}^*_{t+1} \tag{A65}$$

$$\lambda_q: \quad \hat{q}_t = -i_t + \hat{\pi}_{t+1} + \frac{r+\xi}{1+r}\left[\hat{A}_{t+1} - (1-\alpha)(\hat{K}_{t+1} - \hat{L}_{t+1})\right] + \frac{1}{1+r}\hat{q}_{t+1} \tag{A66}$$

$$\lambda_l: \quad \widehat{coc}_t = \frac{r}{1+r}\hat{i}_t + \frac{1}{1+r}\widehat{coc}_{t+1} \tag{A67}$$

$$\lambda_C: \quad \hat{C}_t = \mathbb{E}_t\left[\hat{C}_{t+1} - \xi_r(1/\sigma)\left[\hat{\beta}_{t+1} + \hat{r}_{t+1}\right] + \xi_I\left[\hat{I}_t - \hat{I}_{t+1}\right]\right] \tag{A68}$$

$$\lambda_\pi: \quad \hat{\pi}_t = \psi_p\left[\sigma(\sum_i \vartheta^i \frac{1}{\chi^i} d\chi^i_t) + \sigma\hat{C}_t + \nu\hat{L}_t - \alpha(\hat{K}_t - \hat{L}_t)\right] + \beta\hat{\pi}_{t+1} \tag{A69}$$

$$\lambda_{\chi^h}: \quad \chi^h_t - \chi^h = (1-\alpha-\chi^h)\hat{C}_t + (1-\alpha)\frac{1}{C}G_t - \frac{1}{C}T_t + (1-\alpha)\frac{I}{C}\hat{I}_t \tag{A70}$$

$$\lambda_{\chi^p}: \quad 0 = \mu(\chi^h_t - \chi^h) + (1-\mu)(\chi^p_t - \chi^p). \tag{A71}$$

The first-order condition with respect to $\hat{\delta}_t$ is

$$\omega_\delta(\hat{\delta}_t - \hat{\delta}_{t-1}) - \beta\omega_\delta(\hat{\delta}_{t+1} - \hat{\delta}_t) + \lambda_{I,t}\frac{1+r}{r}\frac{1}{\xi\phi} + \lambda_{\delta,t} - \beta\theta\lambda_{\delta,t+1} = 0. \tag{A72}$$

Solving for $\lambda_{\delta,t}$ by iterating forward,

$$\lambda_{\delta,t} = \sum_{s=t}^{\infty}(\beta\theta)^{s-t}\left\{\omega_\delta(\hat{\delta}_s - \hat{\delta}_{s-1}) - \beta\omega_\delta(\hat{\delta}_{s+1} - \hat{\delta}_s) + \lambda_{I,s}\frac{1+r}{r}\frac{1}{\xi\phi}\right\}. \tag{A73}$$

The first-order condition for $\hat{\delta}^*_t$ is

$$-(1-\theta)\lambda_{\delta,t} + \lambda_{*,t} - \theta\lambda_{*,t-1} = 0. \tag{A74}$$

Iterating backwards,

$$\lambda_{*,t} = \sum_{s=0}^{t} \theta^{t-s}(1-\theta)\lambda_{\delta,s}. \tag{A75}$$

The first-order condition for $\widehat{coc}_t$ is

$$\lambda_{l,t} - \lambda_{l,t-1} - \frac{1+r-\theta}{1+r}\lambda_{*,t} - \frac{1}{\xi\phi}\frac{1+r}{r}\lambda_{I,t} = 0. \tag{A76}$$

Combining (A73), (A75), and (A76),

$$\begin{aligned}\lambda_{l,t} = \lambda_{l,t-1} &- \frac{1}{\xi\phi}\frac{1+r}{r}\lambda_{I,t} \\ &+ \sum_{s=0}^{t}\theta^{t-s}(1-\theta)\sum_{u=s}^{\infty}(\beta\theta)^{u-s}\left\{\omega_\delta(\hat{\delta}_u - \hat{\delta}_{u-1}) - \beta\omega_\delta(\hat{\delta}_{u+1} - \hat{\delta}_u) + \lambda_{I,u}\frac{1+r}{r}\frac{1}{\xi\phi}\right\}.\end{aligned} \tag{A77}$$

Note that when $\theta = 0$ (flexible discount rates), the above expression collapses to

$$\lambda_{l,t} = \lambda_{l,t-1}, \tag{A78}$$

which would imply $\lambda_{l,t} = 0$ for all t given $\lambda_{l,-1} = 0$. Away from such a case, (A77) shows that the process for $\lambda_{l,t}$ is nonstationary. This means that, in general,

$$\lambda_{l,\infty} \neq 0. \tag{A79}$$

The first-order conditions with respect to π_t and i_t are

$$\omega_\pi\hat{\pi}_t - \beta^{-1}\lambda_{q,t-1} - \beta^{-1}\lambda_{C,t-1}(1/\sigma)\zeta_r + \lambda_{p,t} - \lambda_{p,t-1} = 0 \tag{A80}$$

$$\lambda_{q,t} - \lambda_{l,t}\frac{r}{1+r} + \lambda_{C,t}(1/\sigma)\zeta_r = 0. \tag{A81}$$

Combining the two,

$$\omega_\pi\hat{\pi}_t - r\lambda_{l,t-1} + \lambda_{p,t} - \lambda_{p,t-1} = 0. \tag{A82}$$

In the limit as $t \to \infty$,

$$\lim_{t\to\infty} \pi_t = \frac{r}{\omega_\pi} \lambda_{l,\infty} \tag{A83}$$

Equations (A83) and (A79) show $\lim_{t\to\infty} \pi_t \neq 0$ in general under commitment.

We now turn to the case with discretion. Consider a central bank optimizing at $t \geq T$, where all the shocks have died out: $(\varepsilon_t^\beta, \varepsilon_t^G, \varepsilon_t^A) = 0$ for all $t \geq T$. In this case, starting at such time t, the first-best allocation is a feasible solution. This implies $\pi_t = 0$ for all $t \geq T$.

All the above arguments hold after setting $\omega_\delta = \omega_\chi = 0$. Therefore, the same set of results apply to the optimal monetary policy problem with a dual-mandate objective function.

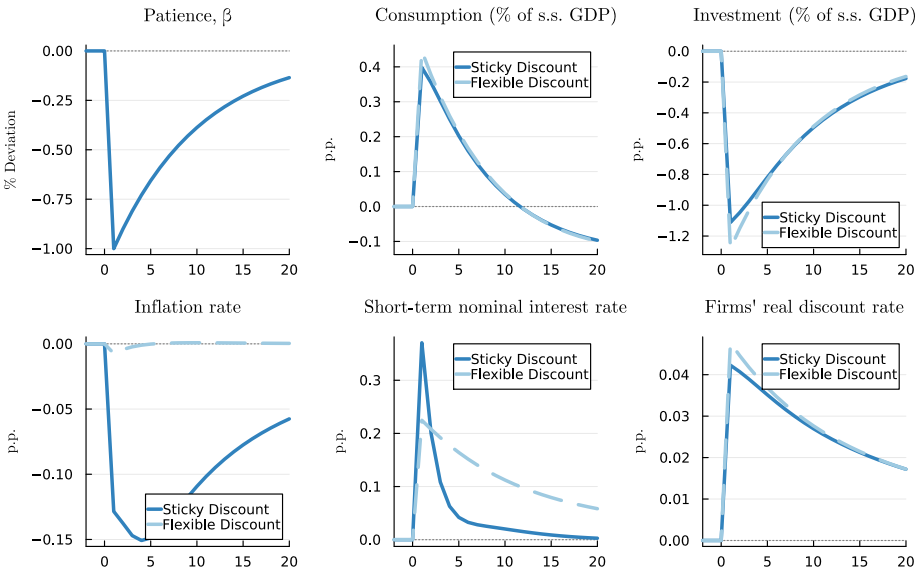


Figure A2: Optimal Response to a Patience Shock under Commitment

The figure plots the optimal monetary policy response to a discount factor shock under commitment for two different values of discount rate stickiness, $\theta \in \{0, 0.95\}$.

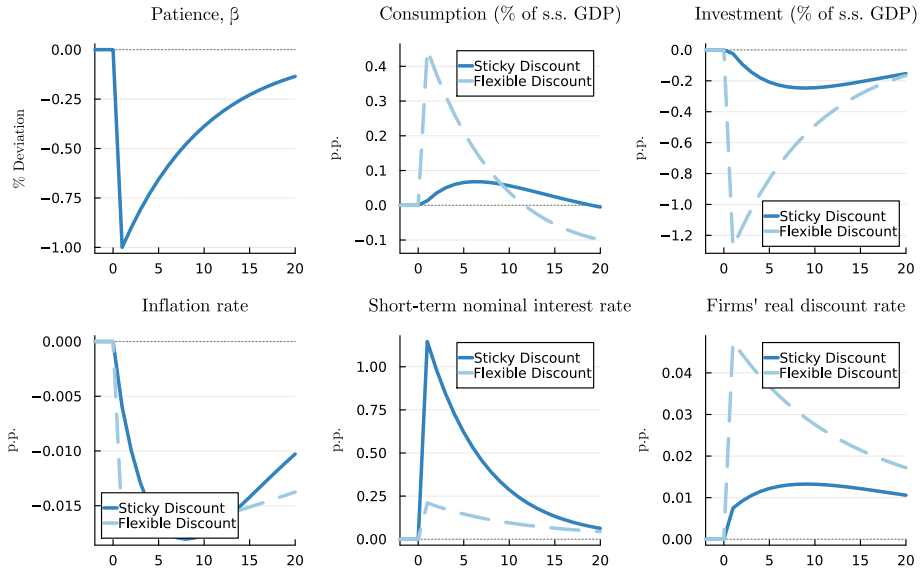


Figure A3: Optimal Response to a Patience Shock under Discretion

The figure plots the optimal monetary policy response to a discount factor shock under discretion for two different values of discount rate stickiness, $\theta \in \{0, 0.95\}$.

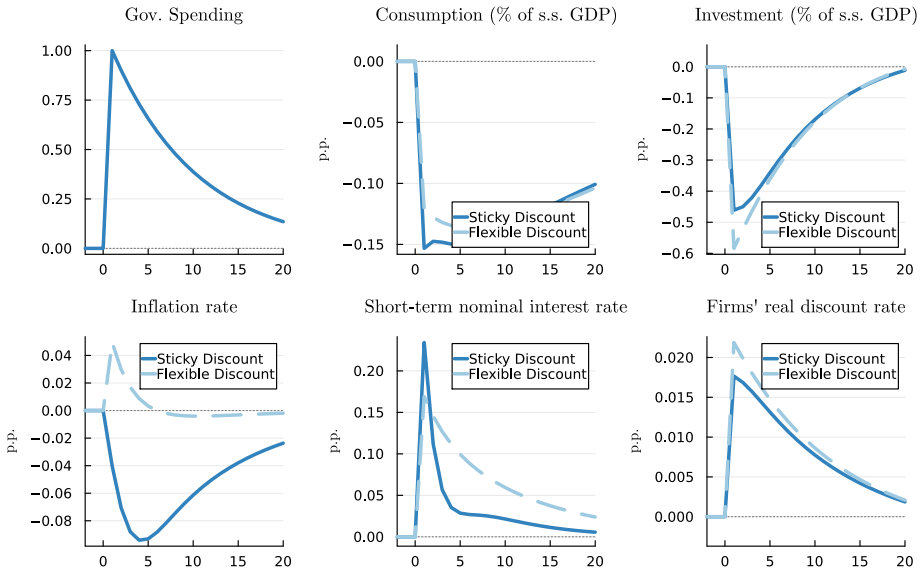


Figure A4: Optimal Response to a Government Spending Shock under Commitment

The figure plots the optimal monetary policy response to a discount factor shock under commitment for two different values of discount rate stickiness, $\theta \in \{0, 0.95\}$.

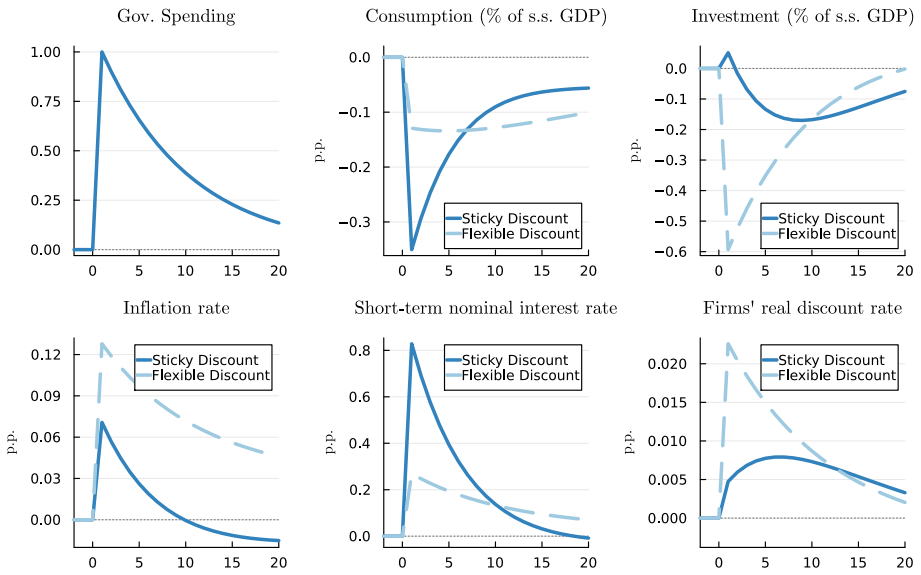


Figure A5: Optimal Response to a Government Spending Shock under Discretion

The figure plots the optimal monetary policy response to a discount factor shock under discretion for two different values of discount rate stickiness, $\theta \in \{0, 0.95\}$.

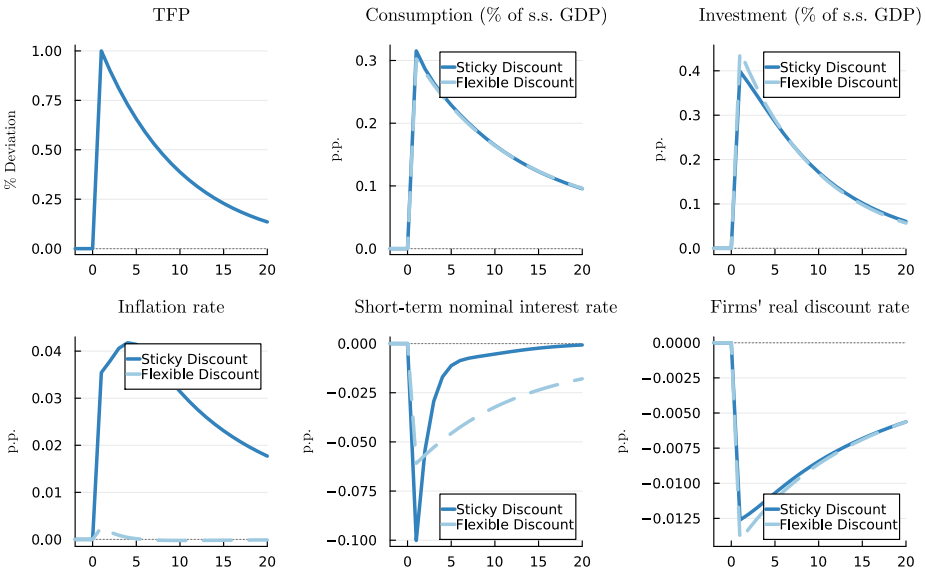


Figure A6: Optimal Response to a TFP Shock under Commitment

The figure plots the optimal monetary policy response to a discount factor shock under commitment for two different values of discount rate stickiness, $\theta \in \{0, 0.95\}$.

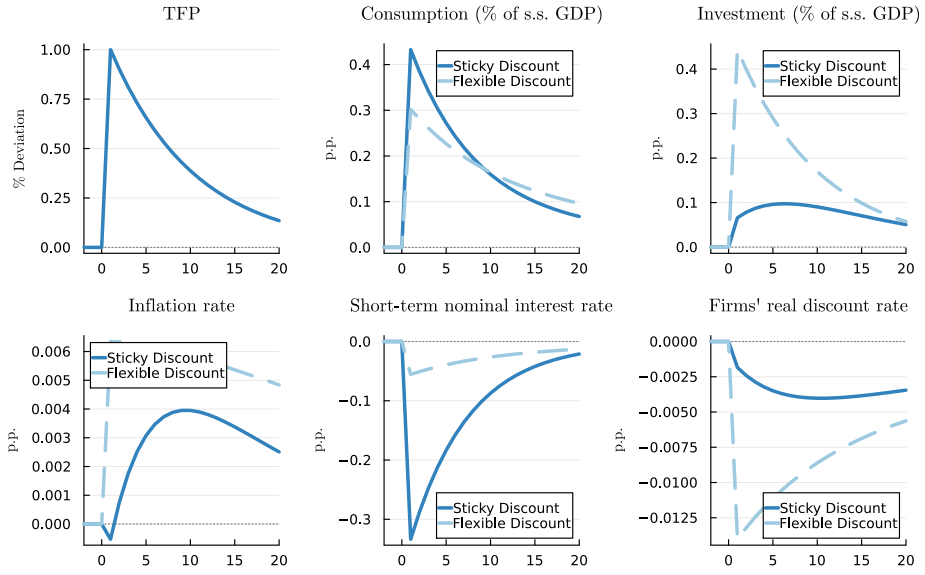


Figure A7: Optimal Response to a TFP Shock under Discretion

The figure plots the optimal monetary policy response to a discount factor shock under discretion for two different values of discount rate stickiness, $\theta \in \{0, 0.95\}$.

Appendix G Additional Exhibits

VARIABLES	(1) CoC	(2) Pred. CoC	(3) DiRa	(4) Wedge	(5) Wedge
Breakeven inflation	0.30**	0.30***	0.054	-0.25**	-0.29**
	(0.13)	(0.056)	(0.11)	(0.11)	(0.12)
Observations	1,414	1,405	1,405	1,405	1,405
FE	Firm	Firm	Firm	Firm	Firm/year
Within R^2	0.011	0.047	0.00032	0.0061	0.0058

Table A2: Discount Rates and Expected Inflation 2002-2020

The table reports the relation between country-level expected inflation and firm-level rates reported on conference calls for the years 2002 to 2020. Breakeven inflation is the annualized country-level breakeven rate (ten-year horizon). CoC is the firm's perceived weighted average cost of capital. Pred. CoC is a generated measure of the firm's perceived cost of capital, predicted using a Lasso estimation procedure and using financial prices as inputs. DiRa is the firm's discount rate. The wedge is the difference between between the firm's DiRa and pred. CoC. The dataset is at the firm-quarter level and runs from 2002 to 2020. Standard errors (in parentheses) are clustered by firm. Statistical significance is denoted by *** $p<0.01$, ** $p<0.05$, * $p<0.1$.