Essays in Monetary Policy and Growth

by

Basil Halperin

B.S., B.A. University of Chicago (2015)

Submitted to the Department of Economics
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2024

Authored by: Basil Halperin
Department of Economics
May 15, 2024

Certified by: George-Marios Angeletos
Professor of Economics, Thesis Supervisor

Certified by: Iván Werning
Professor of Economics, Thesis Supervisor

Accepted by: Isaiah Andrews
Professor of Economics
Chairman, Departmentmental Committee on Graduate Studies

Essays in Monetary Policy and Growth

by

Basil Halperin

Submitted to the Department of Economics
on May 15, 2024 in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

ABSTRACT

This thesis is composed of three essays studying both monetary policy as well as economic growth. The first two chapters study optimal monetary (and fiscal) policy. The third chapter studies the relationship between transformative artificial intelligence, economic growth, and asset pricing.

The first chapter (joint with Daniele Caratelli) studies optimal monetary policy in a world with menu costs. We analytically characterize optimal monetary policy in a multisector economy with menu costs and show that inflation and output should move inversely following sectoral shocks. That is, after negative productivity shocks, inflation should be allowed to rise, and vice versa. In a baseline parameterization, optimal policy stabilizes nominal wages. This *nominal wage targeting* contrasts with inflation targeting, the optimal policy prescribed by the textbook New Keynesian model in which firms are permitted to adjust their prices only randomly and exogenously. The key intuition is that stabilizing inflation causes shocks to spill over across sectors, needlessly increasing the number of firms that must pay the fixed menu cost of price adjustment compared to optimal policy. Finally, we show in a quantitative model that, following a sectoral shock, nominal wage targeting reduces the welfare loss arising from menu costs by 81% compared to inflation targeting.

The second chapter offers a reexamination of optimal monetary and fiscal policy at the zero lower bound. I make five conceptual points about optimal monetary and fiscal policy at the zero lower bound (ZLB) in representative agent New Keynesian models, using the simplest possible version of such a model. (1) Monetary policy is typically described as facing a time consistency problem at the zero lower bound; but if ZLB episodes are a *repeated* game rather than a one-shot game – as is empirically realistic – then the time consistency problem can be easily overcome by reputational effects. (2) The ZLB is not special, in terms of the constraint it creates for monetary policy: an *intra*temporal rigidity, such as the minimum wage or rent control, creates exactly the same kind of constraint on monetary policy as the *inter*temporal rigidity of the ZLB. (3) Austerity is stimulus: in the representative agent New Keynesian model, fiscal stimulus works through the *change* in government spending. Promising to cut future spending – committing to austerity – has precisely the same effect on inflation and the output gap as a decision to raise spending today. (4) Fiscal stimulus can be contractionary, when targeted heterogeneously: if fiscal

spending is targeted at certain sectors, this can in fact lower inflation and deepen the output gap. (5) Fiscal policy faces a time consistency problem at the ZLB, just as monetary policy does. Overall, I suggest that – in this class of models – the power of monetary policy at the ZLB has been underrated, and the power of fiscal policy has been overrated.

The third chapter (joint with Trevor Chow and J. Zachary Mazlish) studies how asset prices can be used to forecast the pace of development of AI technology. We study the implications of transformative artificial intelligence for asset prices, and in particular, how financial market prices can be used to forecast the arrival of such technology. We take into account the double-edged nature of transformative AI: while advanced AI could lead to a rapid acceleration in economic growth, some researchers are concerned that building a superintelligence misaligned with human values could create an existential risk for humanity. We show that under standard asset pricing theory, either possibility – aligned AI accelerating growth or unaligned AI risking extinction – would predict a large increase in *real interest rates*, due to consumption smoothing. The simple logic is that, under expectations of either rapid future growth or future extinction, agents will save less, increasing real interest rates. We contribute a variety of new empirical evidence confirming that, contrary to some recent work, higher growth expectations cause higher long-term real interest rates, as measured using inflation-linked bonds and rich cross-country survey data on inflation expectations. We conclude that monitoring real interest rates is a promising framework for forecasting AI timelines.

Thesis supervisor: George-Marios Angeletos
Title: Professor of Economics

Thesis supervisor: Iván Werning
Title: Professor of Economics

Acknowledgments

I am eternally thankful for and grateful to my family: Erik, Whitney, my grandparents, my extended family, and above all my parents. Thank you for the encouragement to think for myself; and also for the discipline to allow myself do so. Thank you for the unlimited supply of books and the "why can't you just do it all"s. I am grateful for you every day.

I am indebted to my advisors: Marios Angeletos, Iván Werning, and Christian Wolf. The culture and much of the method of academia is passed down by oral tradition, and I am grateful for all of the wisdom passed down to me. I would also like to thank the other faculty, the organizations, and the MIT staff who have supported my research.

A special thank you to my academic coauthors: Daniele Caratelli, Zach Mazlish, Adam Baybutt, and Trevor Chow. I hope that our ideas live on forever and that our friendships last even longer.

Beyond my direct collaborators in research, I owe much to the many other fantastic economists and people – the friends – with whom I've interacted during grad school. Words do not suffice for expressing my gratitude.

To Jett Pettus, who somehow manages to be the same exact person as me but also totally different: thank you for making me a better person. I expect to both admire and learn from your talents, your self-discipline, and your ice cream consumption until the stars burn out – and then in the darkness. I looked for perfection and found something better.

Finally, I am especially grateful for the internet; for all those willing to share their ideas in public for me to learn from; for the "free entry, competitive, methodologically eclectic, wisdom of crowds-based global hive mind". I am lucky to have been born when I was.

Contents

Chapter 1

Optimal monetary policy under menu costs

This chapter is jointly authored with Daniele Caratelli.

1 Introduction

Many central banks around the world have adopted some form of inflation targeting over the past three decades. The textbook formulation of the New Keynesian model provides theoretical grounding for such policies: in the Calvo formulation of the New Keynesian model, where firms are only randomly given the opportunity to change prices, optimal policy in response to efficient shocks is strict inflation targeting. This is true in the textbook one-sector New Keynesian model (Woodford 2003) as well as in heterogeneous multisector versions of the model, for an appropriately-defined price index (Rubbo 2023). The Calvo assumption of random price changes upon which these models are built is mathematically convenient, but arguably comes at the cost of realism. A natural but notoriously less tractable alternative is the "menu cost" model in which firms can *choose* to change their prices at any time, but must pay a fixed menu cost to do so.

We analytically and without linearization characterize optimal monetary policy in a multisector economy with menu costs and show that optimal policy ensures that inflation and output move inversely after sectoral productivity shocks. That is, following negative productivity shocks, inflation should be allowed to rise, and vice versa. We show this by developing a model in which the economy is made up of sectors, where firms are subject to sector-specific productivity shocks and can change their price at any point by paying a menu cost. In baseline parameterizations, optimal policy in response to such shocks is precisely nominal wage targeting: nominal wages should be stabilized, but inflation should not be. This is despite wages themselves being completely flexible. More generally, the optimal policy response to sectoral productivity shocks ensures that the nominal marginal costs of unshocked firms are not affected by shocks.

Intuition. The intuition for this result is that stabilizing inflation causes shocks to spill over across sectors and therefore leads the economy to incur unnecessary menu costs. Consider, for example, a positive productivity shock affecting only firms in sector 1. If the shock is sufficiently large, then it is efficient and desirable for firms in sector 1 to cut their *relative* prices, compared to firms in other sectors of the economy. Under a policy of inflation targeting, the overall price level must be unchanged. To simultaneously have the relative price fall *and* the price level be stable requires not only that sector-1 firms cut their nominal prices, but also that firms in all other sectors raise their nominal prices. As a result, *all* sectors are forced to adjust their prices and pay a menu cost.

A natural alternative – which we show to be optimal – is instead to simply allow firms in sector 1 to cut their nominal prices and to ensure that firms in other sectors do not want to adjust their prices. As a result, relative prices are correct, and *only* sector-1 firms must pay a menu cost. Thus optimal policy economizes on wasteful menu costs compared to inflation targeting, while still achieving the efficient allocation. Optimal policy "looks through" the shock in the sense that *aggregate* inflation is allowed to adjust in response to the sectoral shock, instead of the central bank acting to ensure aggregate inflation is unaffected.

Firms only want to adjust their prices if their *nominal marginal costs* change, and so optimal policy seeks to ensure that nominal marginal costs do not change in unshocked sectors. In a baseline model where wages and productivity are the only factors affecting marginal cost, optimal policy stabilizes nominal wages, since this ensures marginal costs are unchanged for those firms whose productivity does not change. More generally, optimal policy causes inflation and output to move inversely: the positive productivity shock causes output to rise, and the price decrease in sector 1 causes aggregate inflation to fall. In contrast, inflation targeting would be optimal in the Calvo version of this model, as noted above.[1]

Analytical model. We begin in sections 2 and 3 with an off-the-shelf multisector menu cost model and analyze a one-sector shock, as described above, which is the minimum necessary machinery to highlight the core economic logic. The framework is general and can allow "menu costs" to capture a broad conception of any *fixed* costs of price adjustment, whether they be physical, informational, or behavioral "menu" costs.

We go on to show in sections 4 and 5 that the logic generalizes to several extensions. We characterize optimal policy when shocks affect multiple sectors simultaneously and show that for a broad class of such shocks stabilizing the nominal marginal costs of unshocked firms continues to be optimal policy. Next, we extend the model to allow for production networks, and we show that the same characterization of optimal policy holds. In the roundabout economy of Basu (1995), there is a particularly simple characterization

[1] Rubbo (2023) shows this in a multisector model with a general input-output structure and the textbook Calvo friction. Woodford (2003), Aoki (2001), and Benigno (2004) show the same in models without the general network structure, as in the environment presented here, again under the Calvo friction.

of optimal policy: stabilizing nominal marginal costs of unshocked firms means stabilizing a weighted average of nominal wages and prices. The weight on prices corresponds to the production share of intermediate inputs. Indeed, if intermediate inputs are more important than labor in production, then optimal policy attaches more weight to inflation stabilization than to nominal wage stabilization.

We also consider a model variant where *wages* are sticky due to some fixed 'menu' cost (e.g. a fixed cost of contract renegotiation), while prices are flexible. Surprisingly, in response to the same sectoral shocks, optimal policy continues to be stabilize nominal marginal costs of unshocked sectors. Such a policy minimizes expenditure on wasteful menu costs.

Additionally, we use the menu cost model to shed new light on the standard Calvo result: we show that inflation targeting is optimal in the multisector Calvo model only for rather subtle reasons that have not been fully understood. Consider within the Calvo model the previously-described shock raising productivity only in sector 1. Firms in sector 1 want to cut their price as a result. Under Calvo, it is optimal for the central bank to induce every other sector to raise their prices slightly, so that aggregate inflation is zero.[2] Because of the Calvo friction, some firms within each sector are exogenously prevented from adjusting their price; this causes within-sector price dispersion. If instead of all sectors adjusting, *only* sector-1 firms adjusted, then the within-sector price dispersion would be *only* in sector 1. Optimal policy instead forces *all* sectors to adjust to the shock only affecting sector 1, causing within-sector price dispersion in *all* sectors, but with the benefit of reducing the severity of price dispersion within sector 1.

In short, the *convexity of the welfare costs of relative price dispersion* under Calvo makes it optimal to smooth dispersion across sectors instead of concentrating it in one sector. In contrast, menu costs are *nonconvex*, implying that it is optimal for only the shocked firms to adjust, instead of smoothing adjustments across all firms in the economy. This is because, under menu costs, shocked firms can choose whether or not to adjust prices, preventing arbitrarily-severe within-sector price dispersion – unlike under Calvo.

We interpret these results as support for the idea that central banks should "look through" sectoral shocks and allow them to affect aggregate inflation (Powell 2023; Brainard 2022; Schnabel 2022). For example, instead of tightening monetary policy in response to a negative oil supply shock in order to hold down inflation – e.g. as implemented by the European Central Bank explicitly in 2011 or the Federal Reserve implicitly in 2008 – this menu cost logic would suggest that it is efficient to allow energy prices to adjust without

[2] In a model with symmetry across sectors, "inflation" here refers to the standard price index. In a model with additional heterogeneity, this refers to inflation in the "divine coincidence index" of Rubbo (2023).

forcing other prices to decrease in order to compensate.

Quantification. We quantify the welfare loss of inflation targeting under menu costs in a dynamic version of the model. This model is calibrated to US data and includes idiosyncratic, firm-level shocks, a second major source of price changes, on top of sectoral shocks.

We evaluate the performance of inflation and nominal wage targeting by comparing the welfare loss relative to a corresponding flexible-price economy, measured in units of consumption. Following a sectoral shock, the welfare loss due to nominal rigidities is 80.6% smaller under nominal wage targeting than under inflation targeting. We also find that nominal wage targeting is the optimal rule in the class of monetary policy rules that stabilize a weighted average of wages and prices.

We further decompose the welfare loss caused by menu costs into "direct costs", from the labor required for price adjustment, and "efficiency losses", from incorrect relative prices. We show that the welfare gains from nominal wage targeting reflect not only a substantial reduction in direct costs but also a reduction in efficiency losses.

These results follow in part from the fact that the empirical literature estimates menu costs to be fairly large. Based on directly-measured costs alone, the existing literature finds that between 0.6% and 1.2% of firm revenue is spent per year on costs related to price adjustment (Levy et al. 1997; Dutta et al. 1999; Zbaracki et al. 2004).[3] We calibrate the quantitative model to the findings of this literature.

Position in literature. To our knowledge, we are the first to fully characterize optimal monetary policy in the face of fixed menu costs when firms have a motive to adjust relative prices. On the one hand, without changes in productivity between firms, there is no motive for relative-price changes and so optimal policy under menu costs is trivially zero inflation: prices never need to move and price stickiness is irrelevant (see e.g. Nakov and Thomas 2014).[4] On the other hand, several papers allow for relative-price movements under menu costs but take as given that the central bank targets inflation, and simulate numerically how the presence of menu costs affects the optimal level of inflation (Blanco 2021, Nakov and Thomas 2014, Wolman 2011). Adam and Weber (2023) study optimal

[3]The important measurement work of Levy et al. (1997), Dutta et al. (1999), and Zbaracki et al. (2004) is extensively cited in the menu cost literature. These measurements are, as we emphasize, endogenous to the monetary policy regime.

[4]Other papers that analyze optimal monetary policy in a sectoral setting, i.e. a setting with relative price movements, besides those already cited include the vertical chain model of Huang and Liu (2005), again under the Calvo friction; Kreamer (2022) as well as Erceg and Levin (2006), who study optimal monetary policy in sectoral models with fixed prices and durable goods; and Guerrieri et al. (2021), who study optimal monetary policy in a sectoral model with downward nominal wage rigidity.

monetary policy at steady state under menu cost frictions with deterministic productivity trends, to a first order approximation, which is complementary to our study of optimal policy in response to stochastic productivity shocks. Adam and Weber (2023) explicitly turn off consideration of minimizing resource costs (their Assumption 1), which we highlight as an important factor in optimal policy. A larger literature makes assumptions on monetary policy – i.e. does not analyze optimal policy – and asks how the presence of menu costs affects macroeconomic dynamics (among others, Caplin and Spulber 1987; Golosov and Lucas 2007; Gertler and Leahy 2008; Nakamura and Steinsson 2010; Midrigan 2011; Alvarez, Lippi and Paciello 2011; Auclert et al. 2023). That is, these papers conduct a positive analysis, while we conduct a normative analysis. There is also a large empirical literature on menu costs, finding that menu cost models fit the micro data very well.[5]

Our model formalizes and extends the insightful, literary argument made by Selgin (1997) (chapter 2, section 3) that nominal income targeting, or something like it, is optimal in a world with menu costs. Relative to Selgin's elegant informal discussion, we are able to introduce the role of state dependence, which is natural in the context of menu costs and does affect optimal policy. Additionally, we are able to formalize and be precise about the argument in the context of a standard macro model. This formalization allows us to connect our results to prior modeling work, to characterize precisely the nature of optimal policy, and to take the model to the data to quantify the welfare costs of inflation targeting.

The bigger picture. We see our paper as helping unify the literature on optimal monetary policy. In the last decade a number of papers across a variety of classes of models have found that optimal policy should cause inflation to be countercyclical, not constant: the price level P should move inversely with real output Y. However, sticky price models – the workhorse model of modern macro – had conspicuously held out for the optimality of inflation targeting.

First, Koenig (2013) and Sheedy (2014) show in heterogeneous agent models that when financial markets are incomplete and debt is written in nominal, non-state contingent terms, then nominal income targeting is optimal and inflation targeting is suboptimal. That is, $P \times Y$ should be stabilized, and P should move in response to shocks. Werning (2014) notes that if additional heterogeneity is added to the model, then P and Y should move inversely but not one-for-one. This echoes our results.[6] Second, Angeletos and La'O (2020) show that in a world where agents have incomplete information about the

[5] Among many others: Alvarez et al. (2019); Nakamura et al. (2018); Cavallo (2018); Cavallo and Rigobon (2016); Klenow and Kryvtsov (2008); Gautier and Le Bihan (2022).

[6] These ideas have been developed further in Bullard and DiCecio (2019) and Bullard et al. (2023).

economy, optimal policy should again ensure the price level P and real output Y move inversely, in order to minimize monetary misperceptions.[7] Third, when wages are sticky due to a Calvo-type friction, optimal monetary policy is to stabilize nominal wages (Erceg, Henderson and Levin 2000), a policy which also results in countercyclical price inflation.

Despite these results in three highly important classes of models – incomplete markets, information frictions, and sticky wages – it may have been easy to set them aside and nonetheless consider inflation targeting as the proper baseline for optimal monetary policy due to its optimality in the workhorse sticky price model (e.g. Woodford 2003).[8] We hope our paper helps to conceptually integrate these results from across the incomplete market, information friction, sticky wage, and sticky price models. Our results suggest that *countercyclical inflation*, not stable inflation, is a robustly-optimal policy prescription.

Outline. We first illustrate the optimal policy result in sections 2 and 3 in a baseline setting as described above: an off-the-shelf sectoral model, augmented with menu costs, hit by an unanticipated sectoral productivity shock. In section 4, we use our setup to shed new light on the conventional New Keynesian model. In section 5, we show that the optimality of nominal wage targeting continues to hold under a number of generalizations. In section 6, we generalize further by building a quantitative model in order to incorporate dynamics and calculate the welfare gains of adopting nominal wage targeting. Section 7 concludes with a discussion of practical implementation.

2 Baseline model

Our baseline framework is a two-period model starting at steady state. There are S sectors, each consisting of a continuum of monopolistically competitive intermediate firms which are aggregated into a sectoral good by a competitive sectoral packager. A competitive final goods producer combines the output of each of the S sectors into a final good, sold to the household. The model and the functional forms we use are the same as Golosov and Lucas (2007), except that productivity shocks are sectoral rather than firm-specific and we analyze optimal monetary policy instead of exogenous monetary shocks. In section 5, we generalize the functional forms.

[7] The nominal contracts and incomplete information literatures also were preceded and discussed very clearly by Selgin (1997).

[8] In the textbook sticky price model, countercyclical inflation can be optimal if there is a binding zero lower bound constraint on the nominal interest rate (Eggertsson and Woodford 2003; Werning 2011; Woodford 2012).

2.1 Household

The representative household's utility function is given by

$$\mathbb{W} = \ln C - N + \ln\left(\frac{M}{P}\right) \tag{1}$$

Utility is a function of consumption C, labor N, and real money holdings $\frac{M}{P}$.[9] The household chooses C, N and M to maximize its utility, subject to its budget constraint:

$$PC + M = WN + D + M_{-1} - T$$

To fund expenditures, the household uses labor income from wages W, firm dividends D, and previous period money balances M_{-1}, less taxes T. The first order conditions imply:

$$PC = M \tag{2}$$

$$W = M \tag{3}$$

Our particularly simple assumptions on preferences – again matching those of Golosov and Lucas (2007) – result in two simple optimality conditions: an equation of exchange (2) and an equation (3) stating that in equilibrium the nominal wage W is directly determined by the money supply M.

2.2 Final good producer

The representative final good producer aggregates sectoral goods y_i of price p_i across S sectors, using Cobb-Douglas technology, into the final good Y consumed by the household. Operating under perfect competition, its problem is:

$$\max_{\{y_i\}_{i=1}^S} PY - \sum_{i=1}^{S} p_i y_i$$

$$\text{s.t. } Y = \prod_{i=1}^{S} y_i^{1/S} \tag{4}$$

[9]We follow Woodford (1998) in ignoring the welfare effects of real balances when analyzing optimal monetary policy. Additionally, we need not use a money-in-utility framework – the results generalize to any framework where the central bank controls some nominal variable – but it allows us to depart minimally from existing literature.

The resulting demand for sectoral goods is:

$$y_i = \frac{1}{S} \frac{PY}{p_i} \tag{5}$$

The zero profit condition gives the price P for the final good:

$$P = S \prod_{i=1}^{S} p_i^{1/S} \tag{6}$$

In section 5.1 we discuss how generalizing the Cobb-Douglas functional form used here has no impact on the optimal policy result.

2.3 Sectoral goods producers

In sector i, a representative sectoral goods producer packages the continuum of intermediate goods, $y_i(j)$, produced within the sector using CES technology. Note that for notational clarity, we will consistently use j to identify an intermediate firm and i to identify a sector. The problem of the sectoral packager for sector i is:

$$\begin{aligned} \max_{[y_i(j)]_{j=0}^{1}} \; & p_i y_i - \int_0^1 p_i(j) y_i(j) dj \\ \text{s.t. } y_i &= \left[\int_0^1 y_i(j)^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}} \end{aligned} \tag{7}$$

This results in a demand function $y_i(j)$ and a sectoral price index p_i:

$$y_i(j) = y_i \left(\frac{p_i(j)}{p_i} \right)^{-\eta} \tag{8}$$

$$p_i = \left[\int_0^1 p_i(j)^{1-\eta} dj \right]^{\frac{1}{1-\eta}} \tag{9}$$

2.4 Intermediate goods producers

In each sector there is a unit mass of monopolistically competitive firms, each producing a different variety of the sectoral good. Their technology is linear, and all firms within a sector i share a common productivity level A_i.[10] The linearity of technology signifi-

[10] Note that it is standard in the optimal policy literature on sectors and networks to only consider the optimal policy response to *sector*-level productivity shocks: see for example Rubbo (2023), Woodford (2003), Aoki (2001), or Benigno (2004). In particular, these papers do not consider idiosyncratic, firm-level pro-

cantly simplifies the exposition and is important for generating the optimality of nominal wage targeting; we generalize this in section 5.

Intermediate firms are subject to menu costs: if they choose to adjust their price, they must hire an extra ψ units of labor at the wage rate W. This fixed cost of price adjustment, $W\psi$, is what we refer to as a "menu cost". The menu cost itself is simply a transfer from firm profits to household labor income; the *welfare cost* of menu costs comes from the fact that households must supply extra labor in order for prices to be adjusted, and there is a disutility cost associated to this additional labor. This is motivated by the idea of firms needing to employ workers for extra hours to physically walk around and update price stickers in a store, but more generally can be thought of as a modeling device to stand in for *any* fixed costs of price adjustment. For example, if menu costs are information costs, $W\psi$ represents the opportunity cost of the labor time spent thinking about what the optimal price adjustment should be.. Modeling menu costs in other ways does not affect the optimal policy conclusions.[11]

Firm j in sector i thus maximizes profits, including the menu cost if choosing to adjust its price, subject to its demand curve and its production technology, taking as given the inherited price from the previous period $p_i^{\text{old}}(j)$:

$$
\begin{aligned}
&\max_{p_i(j)} p_i(j)y_i(j) - Wn_i(j)(1-\tau) - W\psi\chi_i(j) \\
&\text{s.t. } \chi_i(j) = \begin{Bmatrix} 1 & \text{if } p_i(j) \neq p_i^{\text{old}}(j) \\ 0 & \text{else} \end{Bmatrix} \\
&\quad y_i(j) = y_i\left(\frac{p_i(j)}{p_i}\right)^{-\eta} \\
&\quad y_i(j) = A_i n_i(j)
\end{aligned}
\tag{10}
$$

The objective function defines firm profits, $D_i(j)$. The variable $\chi_i(j) \in \{0,1\}$ is a dummy indicating whether or not the firm chooses to adjust its price, $p_i(j)$. If it does, it incurs

ductivity differences. In contrast, in the separate literature on menu costs, it is common to consider such idiosyncratic shocks (Golosov and Lucas 2007). We analyze the case of both sectoral and idiosyncratic shocks using the quantitative model in section 6.

[11] One example of an alternate modeling method is when menu costs burn real resources (as in Alvarez, Lippi and Paciello 2011), and therefore lower the level of profits transferred to households. Another, more behavioral, modeling method would be to model menu costs as directly inflicting a utility penalty on households, as in Auclert, Rognlie and Straub (2018) and as could be motivated by the literature on fairness in pricing (e.g. Eyster, Madarász and Michaillat 2021). It is straightforward to show that optimal policy is the same if menu costs are modeled in either of these ways. This is because the core intuition remains unchanged: optimal policy still seeks to stabilize the desired price of unshocked firms in order to minimize menu costs.

the menu cost $W\psi$. Otherwise, the price remains at the level inherited from the previous period, denoted $p_i^{\text{old}}(j)$. The term τ in the firm's problem is the standard labor subsidy provided by the fiscal authority to undo the markup distortion from monopolistic competition, $\tau = \frac{1}{\eta}$, for each unit of labor used in production, $n_i(j)$.

If the firm chooses to pay the menu cost and adjust its price, then – from the firm's first order condition – the optimal reset price equals the nominal marginal cost:

$$p_i(j) = \frac{W}{A_i} \tag{11}$$

Notice that, because productivity is sector-specific, all firms j within a sector i face the same decision problem, and thus all make the same decision on whether to adjust and choose the same reset price. Because of this equivalence, we will often refer interchangeably to firm-specific versus sector-specific prices and quantities, e.g. $p_i(j)$ versus p_i and $n_i(j)$ versus n_i.[12]

2.5 The intermediate firm's adjustment decision

We now turn to the question of whether a given intermediate firm will pay the menu cost to adjust its price. The firm makes it decision to adjust by comparing profits under the new optimal price $\frac{W}{A_i}$ net of the menu cost $W\psi$, versus profits under the inherited price p_i^{old} without the loss from menu costs. Plugging in the respective prices as well as constraints into the profit function, we arrive at the price-adjustment condition: firm j in sector i will adjust if and only if

$$\left(\frac{W}{A_i}\right)^{1-\eta} p_i^{\eta} y_i \left[\frac{1}{\eta}\right] - W\psi > \left(p_i^{\text{old}}(j)\right)^{1-\eta} p_i^{\eta} y_i \left[1 - \frac{W/A_i}{p_i^{\text{old}}(j)} \cdot \frac{\eta - 1}{\eta}\right] \tag{12}$$

This nonlinear adjustment condition implies an inaction region Λ, a standard result in menu cost models. The following lemma describes the inaction region.

Lemma 1 (Inaction region)**.** There exists an inaction region Λ in (W, A_i) space such that a firm in sector i will not adjust its price if and only if the value of (W, A_i) remains within this inaction region:

$$(W, A_i) \in \Lambda \tag{13}$$

The larger the menu cost ψ, the larger is this inaction region. The locus of points that

[12]Where sectoral labor is defined naturally as $n_i \equiv \int_0^1 n_i(j)dj$.

result in the new optimal price equaling the inherited price, $\{(W, A_i) | \frac{W}{A_i} = p_i^{\text{old}}\}$, always lies within the inaction region. The inaction region is a connected set.

Proof: See Appendix A.1. □

To interpret this, note that the desired reset price W/A_i depends on two factors:

1. The sectoral productivity A_i, which is exogenous.
2. The level of nominal wages W, which we saw from (3) is completely determined by the central bank, $W = M$, in equilibrium.

Thus, firms are more likely to adjust after either a large productivity shock or a large monetary action, all else equal.

2.6 Market clearing

Labor market clearing implies that total labor supplied by the household, N, equals labor demanded in production, $\sum_i n_i$, plus the amount of labor required to adjust prices, which is $\psi \sum_i \chi_i$:

$$N = \sum_{i=1}^{S} n_i + \psi \sum_{i=1}^{S} \chi_i \tag{14}$$

This market clearing condition is key to the welfare costs of menu costs. Since labor supply N enters the household utility function negatively, larger menu costs ψ requiring the household to work more to adjust prices will lower household welfare.

The remaining equilibrium conditions are standard. The government budget constraint is: $T + (M - M_{-1}) = \tau W \sum_{i=1}^{S} n_i$. Finally, the aggregate resource constraint implies that consumption equals aggregate output:

$$C = Y \tag{15}$$

2.7 Steady state

The economy begins in a symmetric, flexible-price steady state (steady state variables are denoted with a superscript *ss*) in which sectoral productivities A_i^{ss} for $i \in \{1, \ldots, S\}$ are taken as given and nominal wages are normalized to $W^{ss} = 1$. Without loss of generality, we can set $A_i^{ss} = 1$ for all i.

The money supply from (3) is then $M^{ss} = 1$. Firms set prices at their flexible levels (11), $p_i^{ss} = 1$. The aggregate price level (6) is $P^{ss} = S$. From money demand (2), consumption and therefore output are equal to aggregate productivity, $C^{ss} = Y^{ss} = M^{ss}/P^{ss} = 1/S$. From demand equations (8) and (5), sectoral output is $y_i^{ss} = \frac{1}{S}$. From intermediate production technology (10) we recover labor in sector i as $n_i^{ss} = \frac{1}{S}$ and aggregate labor from market clearing (14) as $N^{ss} = 1$.

3 Optimal policy after a productivity shock

As our baseline exercise, we consider the optimal response to an unexpected shock to sector 1 alone. For concreteness, consider a positive productivity shock, which we denote as $\gamma > A_1^{ss} = 1$. How should monetary policy optimally set the money supply M?

Because in the initial steady state all sectors have the same productivity normalized to one, firms in all unshocked sectors $i > 1$ face precisely the same problem after the shock to sector 1 and make the same decision on whether and how to adjust. As a result, for our purposes in this section there are effectively two sectors of different sizes, sector 1 (with productivity $A_1 = \gamma > 1$ and size 1) and sectors k (with productivity $A_k = 1$ and size $S - 1$). Section 5.3 discusses how this generalizes to shocking multiple sectors. We will consistently identify variables for these unshocked sectors with a k. The relative price between the shocked and unshocked sectors, p_1/p_k, will be a key object of analysis.

Proposition 1 characterizes optimal monetary policy in response to this shock.

Proposition 1 (Optimal monetary policy)**.** For a fixed level of menu costs ψ, there exists a threshold level of productivity $\overline{\gamma} > 1$, such that:

1. If the productivity shock to sector 1 is above the threshold, $\gamma \geq \overline{\gamma}$, then optimal policy is exactly nominal wage targeting: monetary policy should ensure $W = W^{ss}$. This results in firms in sector 1 adjusting their prices, while firms in other sectors k leave prices unchanged. This is implemented by leaving the money supply unchanged, $M = M^{ss}$.

2. If the shock is below the threshold, $\gamma \in [1, \overline{\gamma})$, then optimal policy is to ensure that prices remain unchanged and no firm in any sector adjusts.

Additionally, the productivity threshold $\overline{\gamma}$ is increasing in the size of menu costs ψ.

Proof: Lemma 2 and lemma 3 below directly imply the proposition. □

First, we review the economic intuition, which was previewed in the introduction, before proving the proposition. For a sufficiently large productivity shock $\gamma \geq \overline{\gamma}$, it is efficient for the relative price of sector 1, p_1/p_k, to update. To achieve this while simultaneously minimizing the number of sectors which must incur a wasteful menu cost, it is only necessary that firms in sector 1 update their price p_1 – firms in other sectors do not need to update p_k. To ensure that firms in other sectors have no desire to update, the central bank wants to stabilize the level of nominal wages, W, so that the nominal marginal cost of firms in these other sectors is unchanged and these firms have no motive to adjust their prices. On the other hand, for a small productivity shock $\gamma \in [1, \overline{\gamma})$, the benefit of updating the relative price p_1/p_k does not outweigh the welfare loss from the menu cost necessary to do so. It is therefore optimal to ensure that prices remain unchanged across all sectors.

We next step through the math behind this intuition in more detail. We build up to lemma 2 and lemma 3, which together prove proposition 1.

3.1 Allocations in four possible regimes

In this subsection, we characterize the four possibilities for equilibrium that monetary policy can implement. In the next subsection, we will compare welfare across them.

Because there are two types of firms (those in sector 1 hit with productivity shock γ, and those in other sectors k with unchanged productivity) and each type has a binary choice (adjusting or not adjusting its price), there are 2×2 possibilities for what may occur in equilibrium:

1. Both sector 1 and sectors k adjust prices
2. Only sector 1 adjusts its price; sectors k do not adjust
3. Only sectors k adjusts their prices; sector 1 does not adjust
4. Neither sector 1 nor sectors k adjusts price

Furthermore, the central bank determines which of these regimes occurs in equilibrium by manipulating the money supply, M. Whether a firm in some sector i decides to adjust its price depends solely on whether its target price, $\frac{W}{A_i}$, is outside its inaction region (13). Because the central bank can move nominal wages W by its choice of money supply M, it controls which equilibrium is implemented. (Note that there is always a unique equilibrium for a given choice of M – this is not a choice of *equilibrium* selection, but a choice by monetary policy of how much to increase aggregate demand.)

The optimal policy problem thus consists of:

1. Considering each of these regimes individually, and choosing M to maximize welfare *conditional* on the given regime;
2. Then, choosing the regime among the four which has the highest welfare, and implementing the associated optimal M.

This optimal policy problem, formalized in (60) in appendix A.2, is necessarily piecewise due to the sharp discontinuities created by the discontinuous pricing rules of (13), themselves the result of the fixed menu costs.

We now consider each of these four regimes individually, after discussing a benchmark of flexible prices. The ensuing subsection compares across the four.

Flexible price benchmark. As a benchmark, first consider the flexible price allocation, where the menu cost $\psi = 0$. Nominal wages are determined by the money supply, $W = M$, so that from (11) the flexibly-adjusted prices are $p_1 = \frac{M}{\gamma}$ and $p_k = M$. Observe that under flexibility, the relative price across types $\frac{p_1}{p_k}$ is:

$$\left(\frac{p_1}{p_k}\right)_{\text{flex}} = \frac{1}{\gamma}$$

This is an important object. This flexible relative price results in aggregate output and consumption equal to $Y = C = \frac{\gamma^{1/S}}{S}$. Total labor, as in steady state, is $N = 1$. Plugging these quantities into the household utility function (1), we have a flexible-price benchmark for welfare of:

$$\mathbb{W}_{\text{flex}} = \ln\left(\frac{\gamma^{1/S}}{S}\right) - 1 \tag{16}$$

The flexible-price level of welfare is the first-best, efficient benchmark to which policy should be compared.

All sectors adjust. Next, return to the world where there are nonzero menu costs, and consider the case where all sectors pay the menu cost to adjust. Because all firms adjust to the flexible levels of $p_1 = \frac{M}{\gamma}$ and $p_k = M$, the relative price achieves the flexible price level:

$$\left(\frac{p_1}{p_k}\right)_{\text{all adjust}} = \frac{1}{\gamma} = \left(\frac{p_1}{p_k}\right)_{\text{flex}}$$

However, despite prices adjusting, the equilibrium differs from the flexible-price equilibrium because additional labor is required to pay the menu costs of price adjustment. This is where the assumptions on preferences plays a useful simplifying role: the fact that the Golosov-Lucas preferences are quasilinear in labor ensure that all income effects affect labor supply. As a result, the additional labor required for menu costs has no effect on equilibrium *except* to increase the amount of labor used.[13] That is, prices and quantities are the same as the flexible price equilibrium, *except* for the additional labor which must be hired to pay for the menu costs: $N = 1 + S\psi$, where $S\psi$ reflects that there are S sectors which must hire ψ units of labor each to adjust prices.

All together, this means that – conditional on all sectors adjusting – welfare is independent of monetary policy and it is equal to the flexible-price level minus the S sectors' worth of menu costs:

$$\begin{aligned} \mathbb{W}_{\text{all adjust}} &= \ln\left(\frac{\gamma^{1/S}}{S}\right) - [1 + S\psi] \\ &= \mathbb{W}_{\text{flex}} - S\psi \end{aligned} \tag{17}$$

Only sector 1 adjusts. Next consider if only sector 1 updates to $p_1 = \frac{M}{\gamma}$ and sectors k leave their prices unchanged at the steady state level of $p_k = 1$. This results in aggregate output of $Y = \frac{\gamma^{1/S}}{S} M^{\frac{S-1}{S}}$. The total level of labor is $N = \left[\frac{1}{S} + (S-1)\frac{M}{S}\right] + \psi$, reflecting one sector's worth of menu costs ψ, since only sector 1 is adjusting. Thus, household welfare is a function of the money supply decision:

$$\mathbb{W}_{\text{only 1 adjusts}}(M) = \ln\left(\frac{\gamma^{\frac{1}{S}}}{S} M^{\frac{S-1}{S}}\right) - \left[\frac{1}{S} + (S-1)\frac{M}{S} + \psi\right]$$

Conditional on being in this regime, optimal monetary policy chooses M to maximize this expression, which can be found from the first order condition to be:

$$M^*_{\text{only 1 adjusts}} = 1$$

where asterisks denote objects under optimal policy.[14] The optimal money supply in this case is left unchanged at the steady state level, $M^{ss} = 1$. Importantly, this ensures that the

[13] Without preferences ensuring no income effects on consumption, optimal policy would need to account for the fact that the labor required for menu costs affects the marginal rate of substitution between consumption and leisure. Under optimal policy, production would therefore be slightly tilted away from the flexible-price level. An alternative approach would be to model menu costs as a utility penalty affecting the household directly (c.f. Auclert, Rognlie and Straub 2018 among others), in which case the flexible-price allocation is replicated exactly; we model this in appendix B. In general, as long as the income effects of menu costs are quantitatively small, then preferences which are quasilinear in labor are a good benchmark.

[14] In subsection 3.3, we discuss the incentive compatibility of this choice of money supply: i.e. if this choice of M ensures only sector 1 wants to adjust.

relative price across sectors, $\frac{p_1}{p_k} = \frac{M}{\gamma}$, equals the flex-price level:

$$\left(\frac{p_1}{p_k}\right)^*_{\text{only 1 adjusts}} = \frac{1}{\gamma} = \left(\frac{p_1}{p_k}\right)_{\text{flex}}$$

Why does this policy result in the efficient relative price? Setting $M = 1$ ensures nominal wages are $W = 1$, since $M = W$ from (3), which means that nominal wages are unchanged from steady state $W^{ss} = 1$. As a result, the optimal reset price $\frac{W}{A_k} = W$ *coincides with* the inherited price, $p_k^{ss} = 1$, and the optimal pricing is achieved without a need to adjust.

Thus, optimal monetary policy is able to replicate the flexible-price allocation by ensuring that all prices are at the correct level despite sectors k not adjusting, aside from the extra labor required for menu costs. As a result, welfare under optimal policy is equal to the flexible-price level, minus one sector's worth of menu costs from sector 1 adjusting:

$$\begin{aligned} \mathbb{W}^*_{\text{only 1 adjusts}} &= \ln\left(\frac{\gamma^{1/S}}{S}\right) - [1+\psi] \\ &= \mathbb{W}_{\text{flex}} - \psi \end{aligned} \tag{18}$$

Only sectors k adjust. If only sectors $k = 2, ..., S$ adjust, the logic is similar to the prior case, except $S-1$ sectors adjust, instead of only one sector adjusting. The flexible-price allocation is again achievable aside from the extra labor required to pay for menu costs, this time by ensuring that the desired price in sector 1 equals the inherited price. This is implemented by the central bank increasing the money supply, inflating nominal wages to the point where firms in sector 1 have no desire to adjust, $W = \gamma$, and causing firms in other sectors to have a motive to adjust. The optimized level of welfare is thus the flexible-price level minus $S-1$ sectors' worth of menu costs:

$$\mathbb{W}^*_{\text{only } k \text{ adjust}} = \mathbb{W}_{\text{flex}} - (S-1)\psi \tag{19}$$

No sector adjusts. Finally consider the possibility that no firm in any sector adjusts. Sectoral prices are thus unchanged from steady state, $p_i = p_i^{ss} = 1 \;\; \forall i$, and consequently so is the aggregate price level, $P = P^{ss} = S$. Within this regime, this is as if all prices were fully rigid: aggregate output is determined by monetary policy, $Y = C = \frac{M}{S}$. Total labor is $N = \frac{1}{\gamma}\frac{M}{S} + (S-1)\frac{M}{S}$, noting no labor is required for menu costs because no prices are adjusted. Household welfare as a function of the chosen level of the money supply M is:

$$\mathbb{W}_{\text{none adjust}}(M) = \ln\left(\frac{M}{S}\right) - \left[\frac{1}{\gamma}\frac{M}{S} + (S-1)\frac{M}{S}\right]$$

Conditional on being in this regime, optimal monetary policy chooses M to maximize this expression, which can be found from the first order condition to be $M^*_{\text{none adjust}} = \left[\frac{1}{\gamma}\frac{1}{S} + \frac{S-1}{S}\right]^{-1}$. Under this, the optimized level of welfare is:

$$\mathbb{W}^*_{\text{none adjust}} = -\ln\left(S - 1 + \frac{1}{\gamma}\right) - 1 \tag{20}$$

To understand this, note that the relative price $\frac{p_1}{p_k}$ is stuck at the steady state level of 1 instead of being updated to the flexible price level of $\frac{1}{\gamma}$:

$$\left(\frac{p_1}{p_k}\right)_{\text{none adjust}} = 1 \neq \left(\frac{p_1}{p_k}\right)_{\text{flex}}$$

It is because this relative price is stuck at a distorted level that monetary policy is unable to achieve the flexible-price allocation.

3.2 Comparing across regimes

In this subsection, we compare welfare across the four possible regimes just derived. We can immediately observe that only two of the four are worth considering for optimal policy.

Lemma 2 (If adjusting, only the shocked sector should adjust)**.** Welfare when only sector 1 adjusts, $\mathbb{W}^*_{\text{only 1 adjusts}}$, is strictly higher than welfare when all sectors adjust, $\mathbb{W}^*_{\text{all adjust}}$, and welfare when only sectors k adjust, $\mathbb{W}^*_{\text{only } k \text{ adjust}}$.

Proof: This follows immediately from comparing (18) with (17) and (19). □

Lemma 2 follows from the idea that it is better to have fewer firms incur menu costs, together with the fact that optimal policy can implement the efficient relative price $\left(\frac{p_1}{p_k}\right)_{\text{flex}}$ by having *either* sector 1 only adjust, *or* sectors k only adjust, *or* all sectors adjust. Thus, if any firms at all are going to adjust, it is best to have sector-1 firms only adjust.

What remains is to compare welfare if "only sector 1 adjusts" versus if "none adjust". The next lemma compares these two.

Lemma 3 (Only adjust prices beyond a threshold)**.** There is a threshold $\overline{\gamma}$ such that $\mathbb{W}^*_{\text{only 1 adjusts}}$ dominates $\mathbb{W}^*_{\text{none adjust}}$ if and only if the productivity shock exceeds the threshold, $\gamma \geq \overline{\gamma}$. Furthermore, the threshold $\overline{\gamma}$ is increasing in the menu cost ψ.

Proof: Define $f(\gamma) \equiv \mathbb{W}^*_{\text{none adjust}} - \mathbb{W}^*_{\text{only 1 adjusts}}$. Observe that if $\gamma = 1$, then $f(\gamma) =$

$\psi > 0$. Additionally, as $\gamma \to \infty$, then $f(\gamma) \to -\infty$. Finally, f is strictly monotonically decreasing in γ, with $f'(\gamma) = \frac{1}{\gamma}\left[\frac{1}{\gamma(S-1)+1} - \frac{1}{S}\right] < 0$. Since f is continuous in γ, by the intermediate value theorem there exists a $\overline{\gamma} > 1$ such that $f(\overline{\gamma}) = 0$. To see that $\overline{\gamma}$ is increasing in ψ, observe that increasing ψ shifts the entire $f(\gamma)$ curve up, i.e. $\frac{\partial f}{\partial \psi} > 0$. $\square$

Lemma 3 says that there is an important threshold level $\overline{\gamma}$ for the productivity shock. Below this threshold, household welfare is maximized by ensuring that no firm in any sector adjusts; above this threshold, it is maximized by ensuring that sector-1 firms adjust. The intuition for this, as emphasized, is that the welfare loss from menu costs is fixed in size. For a sufficiently small improvement in productivity, the benefit to adjusting prices does not outweigh the fixed welfare loss from the menu cost that is required to adjust. It is only worthwhile to pay this fixed cost above the threshold. The proof follows this same logic.

Additionally, the productivity threshold $\overline{\gamma}$ is increasing in the size of the menu cost ψ. The intuition for this is that for a larger menu cost, the productivity shock must be bigger for it to be worthwhile to adjust.

In the case where none adjust, the level of nominal wages is $W_{\text{none adjust}} = M^*_{\text{none adjust}} = \left[\frac{1}{\gamma}\frac{1}{S} + \frac{S-1}{S}\right]^{-1}$. Observe that for $\gamma = 1$, then nominal wages are exactly unchanged from the steady state level of $W^{ss} = 1$. For small shocks, $1 < \gamma < \overline{\gamma}$, nominal wages are also approximately unchanged. In the quantitative model of section 6, we discuss how close this approximation is.

Denote welfare under optimal policy as $\mathbb{W}^*$, where lemma 2 and lemma 3 together imply $\mathbb{W}^* = \max\left\{\mathbb{W}^*_{\text{only 1 adjusts}}, \mathbb{W}^*_{\text{neither adjust}}\right\}$. Lemma 2 and lemma 3 together also prove proposition 1.

3.3 Adjustment externalities

When discussing the regime where only sector 1 adjusts, we derived equilibrium household welfare as a function of the money supply choice, $\mathbb{W}_{\text{only 1 adjusts}}(M)$, by *assuming* that only firms in sector 1 adjusted prices. We then found the optimal $M^*_{\text{only 1 adjusts}}$ by simply taking the first order condition of this function.

More precisely, however, a central bank would choose the money supply M that maximizes welfare $\mathbb{W}_{\text{only 1 adjusts}}(M)$ *subject to the implementability constraint* that such a choice of M *induces* sector 1 to adjust and other sectors k not to adjust. We term this as "constrained" optimal. The choice of M would need to be incentive-compatible with the assumption on who is adjusting price. The same is true for the case where none adjust: the choice of optimal M must not push any firm outside its inaction region. (The same is true of the case where only sectors k adjust, though this is less important because of lemma 2.) These constraints can be written formally as in equations (53)-(55) in appendix A.2.

In proposition 1 and throughout the body of this paper, we have endowed the social planner with the power to force firms to adjust prices – or equivalently, to subsidize price adjustment – so that these implementability constraints are always nonbinding. This is written out explicitly in (57)-(59) in appendix A.2. In appendix B, however, we show that if the planner does not have this instrument, then it is possible for these implementability constraints to bind.

We term the case where the unconstrained-optimal choice of M is not feasible as "adjustment externalities", and discuss these in detail in appendix B. It may be the case that it is socially optimal for firms in sector 1 to adjust their prices, but it may not be privately optimal to adjust: prices are "too sticky", and there is a positive externality to price adjustment. It is also possible, however, that it is socially optimal for firms in either sector 1 or in sectors k to leave their price unchanged, but it is privately optimal to adjust: prices are "too flexible", and there is a negative externality to price adjustment. The nature of the externality depends on the size of the shock and the size of menu costs, as detailed in appendix B. For the purposes of proposition 1 and for our other analytical results, we have endowed the planner with the choice to overcome adjustment externalities by selectively subsidizing price adjustment.

This issue does not arise in the Calvo literature since there is no *choice* of whether or not to adjust – if given the chance to freely adjust under Calvo, a firm will always do so – and therefore these adjustment externalities do not arise. As a result, there is limited precedent in the literature, with a handful of important exceptions. Ball and Romer (1989*a*) find that menu costs create negative externalities after a *monetary policy shock*. Our setting instead studies whether *efficient* (productivity) shocks create externalities, when monetary policy is set optimally, and finds the possibility of not just negative externalities but also the possibility of positive adjustment externalities. Other related studies include Ball (1987) on negative externalities in the length of labor contracts; Ball and Romer (1989*b*) on externalities in the timing of staggered price setting; and Ball and Romer (1991) on the possibility of menu cost-induced multiple equilibria. All of these papers study economies where monetary policy is not set optimally; our results show that, *even* when monetary policy is set optimally, adjustment externalities may arise.

Finally, Angeletos and La'O (2020) study optimal monetary policy under information frictions, and in the case of endogenous information acquisition studied in their online appendix A, they find that there are no externalities to information acquisition in price-setting if technology is specified as Dixit-Stiglitz as long as monetary policy is set optimally.[15] In our setting with optimal policy under menu costs, rather than information frictions, we show the possibility of externalities even under the Dixit-Stiglitz specification.

3.4 Discussion: "Menu costs" in the model can be interpreted broadly

The term "menu costs" originates with the *physical* resource costs of updating posted prices: restaurants needing to print new menus, or retailers needing to pay workers to replace price stickers on their shelves. These physical resource costs are sizeable and underappreciated; we review the literature in section 4.2, where direct measurement shows these to be between 0.6% and 1.2% of firm revenues in key industries of the economy.

However, menu costs can be conceptualized more broadly than simply physical resource costs, both in reality and through the lens of our model. Consider four possible sources of menu costs:

1. **Physical adjustment costs.** This is the baseline interpretation of our model and proposition 1. For example, retail firms needing to pay to print new price stickers and employ workers in updating these stickers on store shelves.

2. **Information costs.** "Menu costs" may represent fixed costs of *information acquisition* or *information processing*. Suppose firm managers must invest time and attention to the state of the economy before updating prices. These costs operate through the same mechanism as in the model above: these costs require more labor. Optimal policy in this environment is unchanged: monetary policy should minimize unnecessary price adjustments, to reduce resources expended on information acquisition.

3. **Behavioral costs.** "Menu costs" could, alternatively, be interpreted more behaviorally: perhaps consumers have an intrinsic *preference* for stable prices, and changing prices has a direct psychological cost on consumers. In appendix B, we model menu costs as directly impinging on household welfare. Optimal policy is unchanged, since the core intuition of proposition 1 carries through: monetary policy should minimize unnecessary price adjustments, to reduce psychological costs caused by price adjustment.[16]

[15] Gorodnichenko (2008) studies a model with both menu costs and information frictions and numerically studies an information externality that results from their interaction. Caplin and Leahy (2010) also speculate informally about such a phenomenon; see also Caplin and Leahy (1994).

[16] If these psychological costs are asymmetric and consumers only react negatively to price increases, then optimal policy would be altered. In response to a positive productivity shock, nominal wage targeting would remain optimal. In response to a negative productivity shock, optimal policy would stabilize the nominal marginal costs of the adversely affected sectors.

4. **Zero-sum backlash.** Optimal policy would be altered if "menu costs" are zero sum. Suppose "menu costs" in the firm's problem simply represent a behavioral *backlash* by consumers to changes in prices: in response to a change in prices, consumers shift their purchases from one firm to another. In this case, there need not be *aggregate* consequences of "menu costs", since one firm's loss may be another's gain. Note that taking this seriously as the *only* reason for menu costs would entail strong and unlikely conclusions: optimal policy would be to hyperinflate every period, to ensure that firms are induced to adjust, so that relative prices are always correct – without any cost. However, if menu costs are at all non-zero, then the optimal policy results here go through.

The general principle is that any fixed cost of price adjustment which is not zero sum operates through the same logic as proposition 1. In reality, "menu costs" likely are a combination of both zero-sum and non-zero-sum costs. To the extent that menu costs at all have a non-zero-sum component, the logic of proposition 1 goes through.

4 The welfare loss of inflation targeting & comparison with the NK model

In this section, we contrast our optimal policy results under menu costs to those of the canonical sticky price New Keynesian model, where prices are sticky exogenously due to the Calvo assumption.[17]

1. We first discuss how stabilizing inflation under menu costs generates a welfare loss, in contrast to the standard Calvo model where stabilizing inflation is optimal. This is also directly policy-relevant, because leading central banks today describe their policy goals in terms of inflation targeting. This welfare loss is determined by the size of menu costs.

2. Second, we review the direct estimates available in the empirical literature on the size of *physical* menu costs. They are large: at least 0.5% of firm revenue and plausibly much more.

It is not clear whether consumers have an asymmetric distaste for price increases. Anderson and Simester (2010) conduct a field experiment with a publisher and find that customers are significantly antagonized by price *decreases*.

[17]Throughout the paper, we refer to "Calvo" sticky pricing, for the sake of space and following the vast majority of the literature. A fuller accounting would refer to the "Calvo-Yun assumption", in reference to the important work of Yun (1996).

3. Third, we contrast optimal policy in our setting of *nonconvex* menu costs with optimal policy in a model of quadratic *convex* menu costs (Rotemberg 1982). In the convex setting, inflation targeting is optimal, like the Calvo model.

4. Finally, we use the preceding discussion to explain why the standard Calvo model prescribes inflation targeting and not nominal wage targeting, in contrast to the model we present here, shedding new light on the canonical model.

4.1 The welfare loss of inflation targeting under menu costs

The standard New Keynesian model of sticky prices, built on the Calvo exogenous sticky pricing framework, implies that a policy of zero inflation is optimal policy; but in our menu cost setting, such inflation targeting would be strictly suboptimal. This subsection shows a simple result characterizing, quantitatively, how suboptimal inflation targeting is.

In order to implement inflation targeting – i.e. in order to ensure that the price level is unchanged with $P = P^{ss} = 1$ – the central bank has two possibilities:

1. "All adjust": It may force all firms to adjust, and set M to ensure that the increase in price in sectors k to $p_k = M$ exactly offsets the fall in price in sector 1 to $p_1 = M/\gamma$.

2. "None adjust": It may ensure that no firm in any sector adjusts.

Although per proposition 1 it is optimal to ensure no sector adjusts in the case of small productivity shocks, $\gamma < \overline{\gamma}$, it would be unusual to conceptualize a policy of inflation targeting as aiming for a world in which relative prices *never* change. If maintained indefinitely, such a policy of aiming to prevent all relative price changes would seem to be self-evidently unreasonable, since it would shut down the price system. Thus it is more natural to characterize "inflation targeting" in this context as referring to the policy that would ensure all firms adjust.

Thus, inflation targeting requires all firms to adjust – but we saw above that forcing all sectors to adjust prices results in unnecessary menu costs. Thus, the welfare loss of inflation targeting relative to optimal policy is directly captured by the welfare loss caused by the unnecessary menu costs paid by the $S - 1$ unshocked sectors. The empirical size of menu costs ψ together with the number of unshocked sectors $S - 1$ are sufficient statistics for the welfare gains that would come from moving from inflation targeting to nominal wage targeting.

The next proposition summarizes this discussion.

Proposition 2 (The welfare loss of inflation targeting)**.** Denote a policy of "inflation targeting" as a rule for monetary policy ensuring that $P = P^{ss}$ while having correct relative prices, and suppose $\gamma \geq \overline{\gamma}$. Then:

1. Inflation targeting requires that all sectors adjust their prices. It is implemented by increasing the money supply to $M = \gamma^{1/S} > M^{ss}$.

2. Welfare under inflation targeting, denoted $\mathbb{W}^{\mathrm{IT}}$, is strictly less than welfare under the optimal policy described in proposition 1, $\mathbb{W}^*$. The welfare loss is determined by size of menu costs ψ and the number of sectors unaffected by the shock, $S-1$:

$$\mathbb{W}^{\mathrm{IT}} - \mathbb{W}^* = (S-1)\psi$$

Proof: The second claim comes from formulas (17) and (18). For the first claim, suppose the central bank tried to both achieve correct relative prices by only having sector 1 adjust – in which case, $P = SM^{1/S}\gamma^{-1/S}$ – and simultaneously setting M such that the price level was unchanged, $P = P^{ss} = S$. This would require $M = \gamma$. However, if $M = \gamma$ then the optimal price for firms in sector 1 is $p_1 = W/\gamma = 1$, which would mean that firms in sector 1 leave prices unchanged, a contradiction. Similarly, if the central bank tried to achieve correct relative prices while only having sectors k adjust, in which case $P = SM^{\frac{S-1}{S}}$, then this would require $M = 1$, which would cause firms in sectors k to not adjust, again a contradiction. Finally, if no sector adjusts, then it is impossible to achieve correct relative prices, since $p_1/p_k = 1$. It is only by having firms in all sectors adjust, in which case $P = SM\gamma^{-1/S}$, that the central bank can achieve both correct relative prices and ensure that $P = P^{ss}$, by setting $M = \gamma^{1/S}$. □

4.2 Empirical estimates of the size of menu costs are sizeable

In this subsection, we review the literature measuring menu costs.

A reaction to proposition 2 may be the idea that the welfare costs imposed directly from updating prices could be relatively small. The literature on menu costs often builds on the idea that 'second-order menu costs can result in first-order output fluctuations' (Mankiw 1985), in which case the welfare loss of inflation targeting compared to optimal policy would be second-order.

However, it is important to note that in the textbook New Keynesian model with the exogenous Calvo friction, the welfare loss of price stickiness is *also* only second-order (see e.g. Gali 2008).

Additionally, estimates of the real resource cost of menu costs from the empirical literature in fact are quite sizeable and arguably underappreciated: at least 0.5% of total firm revenues annually. These estimates come from two sources: calibrated models and direct measurement.

Calibrated models. One method for estimating the size of menu costs is to measure the frequency of price adjustment, build a model of price adjustment with menu costs, and calibrate the magnitude of menu costs to fit the microdata on the frequency of price adjustment. A number of papers perform this exercise, such as Nakamura and Steinsson (2010), who estimate the size of menu costs to be around 0.5% of revenue per year (their Table II). The more recent work of Blanco et al. (2022) finds that fitting the data would require 2.4% of firm revenue to be paid as menu costs. These estimates thus should be interpreted to capture total "menu costs", very broadly construed, as described in section 3.4.

Direct measurement. A more direct and model-free approach to measuring menu costs is to simply measure them directly. However, because measuring all forms of menu costs – physical adjustment costs, information costs, psychological costs – is difficult, the extant measurement literature focuses on physical adjustment costs alone. These numbers should therefore be interpreted as a lower bound on the total size of "menu costs" for the relevant firms.

To our knowledge, only three papers directly measure menu costs. Levy et al. (1997) directly measures the physical costs of price adjustment for five large grocery store chains across the US. They directly measure the time spent by workers manually changing price stickers on grocery store shelves, using a stopwatch. Such time maps directly the menu cost parameter ψ in our model. They find such menu costs to be 0.7% of firm revenue on average. Dutta et al. (1999) use a similar approach to examine a large drugstore chain, with a narrower conception of menu costs, and find menu costs to be 0.6% of firm revenue. Finally, Zbaracki et al. (2004) examine an industrial manufacturer and, using a broader conception of menu costs, find such costs to be 1.2% of firm revenue.[18]

In short, the literature has found that menu costs – even when only examining the most measurable such costs – are no small matter.

[18] Of course, any measurement of the level of expenditure on menu costs is endogenous to the existing monetary policy regime, as proposition 1 emphasizes.

4.3 Nonconvex menu costs vs. convex (Rotemberg) menu costs

In this subsection, we continue the comparison of optimal policy under menu costs to optimal policy in standard models. This will not only clarify our results but also will shed new light on existing literature.

Throughout this paper, we have used a model of "nonconvex menu costs": the cost of price adjustment is *fixed* and does not scale with the size of a price change (Barro 1972; Sheshinki and Weiss 1977). That is, the menu cost facing a firm is a *nonconvex* function of the size of its price change. Contrast this model with models of *convex* menu costs, where the cost of price adjustment depends on the magnitude of the desired price change, and this cost grows at an increasing rate.

Consider the canonical model of convex menu costs, the Rotemberg (1982) model of quadratic menu costs, where the menu cost scales with the square of the size of the price change:

$$\psi \cdot (p_i - p_i^{ss})^2$$

In contrast, in the model we present above, the menu cost is constant as a function of the size of the price change: where $\mathbb{I}$ represents the indicator function,

$$\psi \cdot \mathbb{I}\{p_i \neq p_i^{ss}\}$$

It is well known that the single-sector Rotemberg convex menu cost model is isomorphic in its structural equations, to a first-order approximation, to the textbook New Keynesian model built on the Calvo time-dependent friction.[19] Furthermore, the model is isomorphic to a second-order approximation in its optimal policy implications to the Calvo model (Nisticò 2007), i.e. inflation targeting not wage targeting is optimal.

Why do Rotemberg *convex* menu costs imply inflation targeting is optimal, while our *nonconvex* menu costs imply nominal wage targeting is optimal?

The difference comes directly from the convex nature of the Rotemberg menu costs: due precisely to the convexity, it is better to have all sectors adjust prices a little than to have one sector do all of the adjustment. With the nonconvex menu costs of our model, it is instead optimal to minimize the number of sectors which choose to adjust at all.

[19] Do note however that the *mechanism* of the Rotemberg model is very different from the Calvo model. Under Calvo, the welfare loss of monetary instability is the resulting relative price dispersion: total factor productivity is effectively lower. Under Rotemberg – and in our model – instead, the loss comes from the real resource cost of menu costs. If the quadratic menu cost requires extra labor, this comes from the additional labor required to adjust prices. If the quadratic menu cost is a real resouce cost, then this is a wedge between consumption and output (see also footnote 11).

The key intuition is in the labor market clearing condition. The labor market clearing condition in the multisector Rotemberg model is:

$$N = \sum_{i=1}^{S} n_i + \psi \sum_{i=1}^{S} (p_i - p_i^{ss})^2 \tag{21}$$

This contrasts with the labor market clearing condition under nonconvex menu costs, our equation (14):

$$N = \sum_{i=1}^{S} n_i + \psi \sum_{i=1}^{S} \mathbb{I}\{p_i \neq p_i^{ss}\}$$

Under both Rotemberg and nonconvex menu costs, it is desirable to minimize the amount of menu costs because of the disutility of labor they create. Due to the *convex* nature of the Rotemberg menu costs in (21), it is better to *smooth* the price changes over all sectors: it is better to have a small price change in every sector, rather than a large price change in one sector. Under nonconvex menu costs, it is instead better to minimize the *number* of sectors which experience any price change. It this difference – convex versus nonconvex costs of adjustment – which explains the differing optimal policy prescriptions.

4.4 Reexamining optimal policy in the Calvo model

Finally, we come to why optimal policy in the menu cost setting differs from optimal policy in the textbook Calvo model (Woodford 2003; Rubbo 2023).

Consider the analogy to the Rotemberg model. As discussed in the prior subsection, it is known that optimal policy under the Rotemberg and Calvo frictions are the same: inflation targeting. We also explained that the Rotemberg model and the fixed menu cost model have differing optimal policy implications because Rotemberg assumes a *convex* menu cost, implying that price changes should be spread over many sectors rather than concentrated in one sector.

Similarly, in the Calvo model, *the welfare cost of price dispersion is convex*. While a perfectly clean comparison with the menu cost model cannot be made due to the differing nature of the models – unlike the comparison of (21) versus (14) – the intuition can still be seen in the welfare-theoretic notion of price dispersion from the Calvo model. As standardly defined, price dispersion under Calvo is defined as:

$$\Delta \equiv \sum_{i=1}^{S} \int_0^1 \left[\frac{p_i(j)}{p_i} \right]^{-\eta} dj \tag{22}$$

Here, $\eta > 1$ continues to be the within-sector elasticity of substitution; and recall $p_i(j)$ is the nominal price of firm j in sector i, which now is heterogeneous within a sector thanks to the Calvo friction.

In the Calvo price dispersion formula, the convexity can be seen mechanically from the fact that the within-sector elasticity of substitution is positive, $\eta > 0$. The intuition follows directly. Due to the Dixit-Stiglitz assumption of *complementarity across goods*, i.e. that $\eta > 0$, it is better to have many goods with slightly distorted prices, rather than to have few goods with highly distorted prices.

Illustration. This discussion is illustrated in figure 1. The figure depicts the level of sectoral prices p_i in a three-sector Calvo model where production technology is constant returns to scale and sectors have symmetric parameters. Marginal cost is, as in (11), W/A_i.

In subfigure 1a, the economy is at steady state, where by assumption all sectoral productivities and prices are equal. That is, if we normalize $W^{ss} = 1$ and $A_i^{ss} = 1$ as in section 2.7, then $p_i^{ss} = 1$.

We then consider an increase in productivity in sector 1 to $A_1 = \gamma > 1$. Subfigure 1b shows what would happen under flexible prices: firms in sector 1 would cut prices, while firms in other sectors remain unchanged. That is, $p_1 = 1/\gamma < 1$ and $p_k = 1$.[20]

Subfigure 1c shows what would happen to sectoral prices in the Calvo sticky-price world under nominal wage targeting. Firms in sector 1 want to cut their prices to the flexible-price level. However, only some fraction of firms in that sector may do so, thanks to the Calvo friction. Other firms remain stuck at the steady state price, and have the wrong price. This creates within-sector price dispersion – the blue area in sector 1 is not uniformly the same height – as well as incorrect relative prices between the unchanged sector-1 firms and firms in other sectors. However, there is no within-sector price dispersion in other sectors: the green area in sector 2 has a uniform height, as does the purple area in sector 3. Firms in unshocked sectors are thus not affected by the shock.

Subfigure 1d shows what happens to sectoral prices in the Calvo world under inflation targeting, which is optimal policy in this setting. Now, monetary policy seeks to ensure that on average prices are unchanged. This requires the central bank to induce an increase

[20]Of course, with flexible prices, the price level is indeterminate. The equilibrium described here can be characterized as an equilibrium refinement where the Calvo parameter is taken to zero.

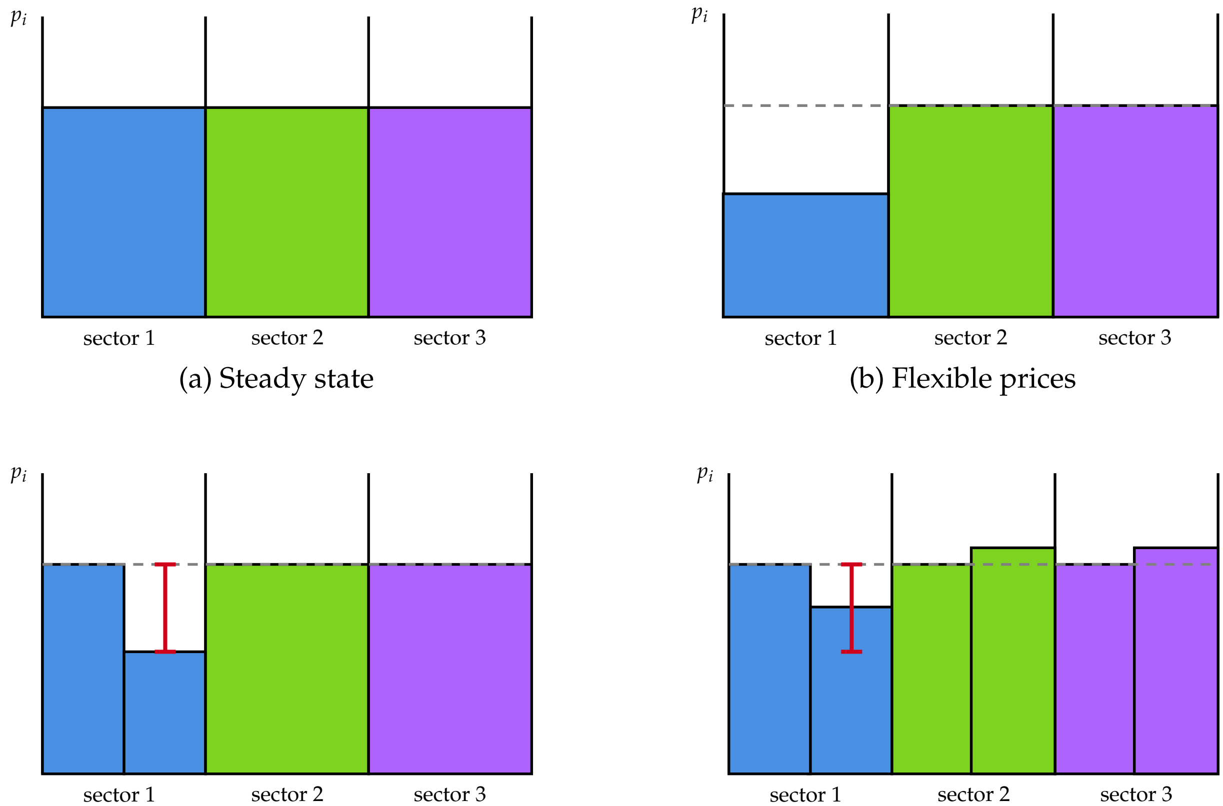

(a) Steady state

(b) Flexible prices

(c) Nominal wage targeting under Calvo

(d) Inflation targeting under Calvo

Figure 1: Sectoral prices under Calvo

in nominal wages. The result is that firms in sector 1 want to cut their prices, but not as much as under stable nominal wages; and firms in other sectors want to increase their prices. Because of the Calvo friction, only a fraction of firms in each sector is able to set the optimal price. As a result, there is within-sector dispersion in *all* sectors: every sector contains firms with differing prices.

As described above, the benefit of disturbing sectors 2 and 3 under inflation targeting is that *the sector-1 price dispersion is lessened*. That is, the gap in heights between the two blue bars is lessened compared to nominal wage targeting, as indicated by the red line. Because of the convexity of the welfare cost of this gap, the welfare benefit from this decrease outweighs the incorrect prices induced in other sectors.

Future work. This suggests an important target for empirical work: *how convex are the costs of price changes* as a function of the size of the change? To our knowledge, this question has received little or no quantitative attention in the empirical literature cited in section 4.2. Since the convexity of price adjustment costs has direct implications for the optimal target for monetary policy, this seems like an important gap to fill.

Additionally, future work could consider coordination frictions that dampen the ability of monetary policy to smooth shocks across sectors. The exogenous nature of the Calvo friction means that firms in unshocked sectors respond symmetrically to a movement in nominal wages induced by monetary policy as do firms in the shocked sector in response to the shock. More realistically, firms may more easily adjust prices following a productivity shock to their own sector, since they are more likely to be aware of it. This would be natural if productivity changes are in fact *endogenous*, rather than an exogenous shock; or in models of rational inattention where sectoral shocks typically receive more attention, endogenously, than aggregate shocks (Maćkowiak, Matějka and Wiederholt 2023).

5 Extensions to the benchmark model

We now consider several natural extensions to the model, which we continue to solve analytically. These showcase the robustness of our results, and are useful for reinforcing the intuition built above, regarding the mechanism of the results.

The core intuition argued above is that optimal monetary policy seeks to ensure that nominal marginal costs are unchanged for firms who do not receive a productivity shock, so that they have no desire to adjust their prices. This intuition is preserved even as the model is extended in various directions. What does change, however, is the formula for the nominal marginal cost for an unshocked firm.

5.1 Functional form generalizations

Consider the baseline model of section 2, but allow for the following generalizations:

1. Any constant returns to scale production technology for final goods, with (4) becoming $Y = F(y_1, ..., y_S)$ with F homogenous of degree 1
2. Potentially decreasing returns to scale in intermediate goods production technology, with (10) becoming $y_i(j) = A_i n_i(j)^{1/\alpha}$ with $1/\alpha \in (0, 1]$
3. Any household preferences quasilinear in labor, with (1) becoming $\mathbb{W} = U\left(C, \frac{M}{P}\right) - N$

Denote this as the generalized model.[21]

[21] In the generalized model, the functional form for utility remains quasilinear in labor, as with Golosov-Lucas preferences. As discussed in section 3.1 and footnote 13, this is necessary to ensure that the income effects induced by menu costs do not distort the consumption-leisure margin. In the model of appendix B, where menu costs are modeled as a utility penalty and therefore do not reduce household income, pref-

As in the baseline model, nominal marginal costs are an important object. Nominal marginal costs can be derived as:

$$MC_i(j) = \left[\alpha \frac{W}{A_i^\alpha} \left(y_i p_i^\eta\right)^{\alpha-1}\right]^\theta \tag{23}$$

$$\theta \equiv \left[1 - \eta(1-\alpha)\right]^{-1} \tag{24}$$

Observe that if α were equal to one, then $\theta = 1$ and $MC_i(j) = W/A_i$, as in equation (11) of the baseline model. In this more general model, marginal cost depends on not just wages and productivity, but also on demand. This is due to the decreasing returns to scale of the production technology.

As a result, optimal policy in this general model does not exactly stabilize nominal wages, but instead stabilizes nominal marginal costs (23) for unshocked sectors, as the next proposition describes. The intuition remains the same: ensuring that only the shocked sector adjusts and others do not minimizes the menu costs incurred by the economy.

Proposition 3 (Functional form generalizations)**.** Consider again the positive productivity shock γ affecting sector 1, in the generalized model. For a fixed level of menu costs ψ, there exists a threshold level of productivity $\overline{\gamma} > 1$, such that:

1. If the productivity shock to sector 1 is above the threshold, $\gamma \geq \overline{\gamma}$, then optimal policy is to ensure the nominal marginal costs of firms outside sector 1 are unchanged. This results in firms in sector 1 adjusting their price, while firms in other sectors k leave prices unchanged.

2. If the shock is below the threshold, $\gamma \in [1, \overline{\gamma})$, then optimal policy is to ensure that prices remain unchanged and no firm in any sector adjusts.

Additionally, the productivity threshold $\overline{\gamma}$ is increasing in the size of menu costs ψ.

Proof: The proof follows exactly the same steps as in the proof of proposition 1. □

The statement of proposition 3 is precisely the same as the statement of proposition 1, except that instead of stabilizing *nominal wages* to ensure unshocked firms do not adjust, the central bank stabilizes the *nominal marginal costs* of unshocked firms. These marginal costs depend on demand, which is affected by the shock to sector 1.

erences need not be quasilinear in labor and can be fully general. Alternatively, consider if menu costs are modeled as a loss of profits without transfer to the household (as in e.g. Alvarez et al. 2019). Then, we would want to ensure there is no income effect on *labor supply* rather than no income effect on consumption. In other words, preferences would have to be of the Greenwood, Hercowitz and Huffman (1988) form.

An example. In the most abstract case, optimal policy cannot be characterized more sharply than proposition 3. However, it is instructive to consider a special case that is common in much of the optimal policy literature.

Suppose preferences over consumption and real balances are isoelastic; aggregation technology is Cobb-Douglas; and continue to allow for decreasing returns to scale technology:

$$\begin{aligned}
\mathbb{W} &= \frac{1}{1-\gamma}C^{1-\gamma} + \frac{1}{1-\gamma}\left(\frac{M}{P}\right)^{1-\gamma} - N \\
Y &= \prod_{i=1}^{S} y_i^{1/S} \\
y_i(j) &= A_i n_i(j)^{1/\alpha}
\end{aligned}$$

In particular, this differs from the baseline model by allowing for isoelastic preferences and also decreasing returns to scale.

In this example, it can be shown that marginal cost is a weighted average of nominal wages and the aggregate price level:

$$MC_i(j) = k\frac{W^{\xi}P^{1-\xi}}{A_i} \tag{25}$$

$$\xi \equiv \frac{\gamma+\alpha-1}{\gamma\alpha} \tag{26}$$

where k is an unimportant constant.[22]

Thus, in this specification, following a shock to sector-1 productivity alone, the central bank should stabilize $W^{\xi}P^{1-\xi}$. This means manipulating nominal wages W to offset any change in the aggregate price level P caused by the change in prices in sector 1. By holding steady $W^{\xi}P^{1-\xi}$, the central bank ensures nominal marginal costs are unchanged for firms outside sector 1, causing them to leave prices unchanged.

How much weight should the central bank give to stabilizing wages versus prices? We offer a more complete assessment in the quantitative model of section 6, but we may here consider a back of the envelope calculation. A typical calibration for γ, the inverse of the elasticity of intertemporal substitution, is $\gamma = 2$; a typical calibration for α might be $1/\alpha = 0.6$ to match the labor share. Plugging these values into (26) would result in the central bank putting a weight of $\xi = 0.8$ on stabilizing nominal wages, and a weight of $1 - \xi = 0.2$ on stabilizing inflation.

[22] $k = \frac{1}{S}(\alpha S)^{1/\alpha}$.

5.2 Sectoral heterogeneity and a monetary "least-cost avoider" principle

We now return to the baseline setting and consider two kinds of heterogeneity, in sector size and menu cost magnitude, that lead us to a "least-cost avoider" interpretation of optimal monetary policy. We offer economic intuition on each in turn, and then state the result formally.

5.2.1 Heterogeneity in sector size

Suppose that sectors are of different sizes. Instead of each sector containing a continuum of firms on $[0, 1]$, allow sector i to range over $[0, S_i]$ for some finite S_i. We then denote $S \equiv \sum_i S_i$. Let everything else in the model remain as in the baseline model.

Heterogeneity in sector size affects the labor market clearing condition:

$$N = \sum_i n_i + \psi \sum_i S_i \chi_i \tag{27}$$

Heterogeneity in sector size interacts with menu costs, since larger sectors require hiring more labor to adjust prices.

With this change, optimal policy is *nearly* the exact same as characterized in proposition 1. What differs is only in the extreme case when sector 1 is larger than all other sectors put together, $S_1 > \sum_{k>1} S_k$. Then if relative prices are to adjust it is actually optimal to have firms *outside* sector 1 adjust their price in response to a shock affecting sector 1 itself. That is because although it is only sector 1 which is affected by the shock, the combined mass of firms outside sector is smaller than sector 1 itself. Therefore the menu costs burned by having all other firms adjust price is less than the menu costs burned by having "just" sector 1 adjust.

Thus under the (extreme) assumption that sector 1 is larger than the combined mass of all other sectors, implementing the regime where 'only sectors k adjust' is preferable to implementing the regime where 'only 1 adjusts'. This has the same intuition that it achieves the correct relative prices while economizing on menu costs – where, here, economizing on menu costs means having sector 1 not adjust. This policy would not stabilize nominal wages (or inflation).

We interpret this case as an illustration of the logic of our results, rather than an empirically-relevant case in general. Only in the case where *more than half* of the economy is *homogeneously* affected by the same shock does this result carry through. Otherwise, the optimal policy prescriptions of proposition 1 carry through exactly.

5.2.2 Heterogeneity in menu cost size

Introducing heterogeneity in menu cost size by sector is mostly similar. If the menu cost of sector i is ψ_i, the labor market clearing (14) becomes:

$$N = \sum_i n_i + \sum_i S_i \psi_i \chi_i$$

Observe that the direct effect of heterogeneity in menu costs (ψ_i) on welfare is isomorphic to that of heterogeneity in sector size (S_i). But heterogeneity in menu cost size, unlike that in sector size, *also* affects the size of inaction regions given in (13). However, this additional complication has somewhat limited impact.

First – analogous to the possibility just discussed that sector 1 is very large in size – if weighted the menu costs of sector 1 are extremely large relative to those of other sectors, $S_1\psi_1 > \sum_{k>1} S_k\psi_k$, then it again is optimal to have all firms outside sector 1 adjust rather than those in sector 1 if relative prices are to change. Second, variation in ψ_1 does affect when it is optimal to allow prices to go unchanged, i.e. affects the value of the threshold $\overline{\gamma}$.

5.2.3 Interpretation: a monetary "least-cost avoider" principle

We summarize both the above results in the following proposition.

Proposition 4 (Sectoral heterogeneity)**.** Suppose sector i is of size S_i and has menu cost ψ_i. Suppose further that the size-weighted menu cost of sector 1 is smaller than the combined weighted sum of menu costs for other sectors, $S_1\psi_1 < \sum_{k>1} S_k\psi_k$. Then optimal monetary is exactly the same as characterized in proposition 1 modulo changes in the constant $\overline{\gamma}$.

Proof: Under the assumption about the magnitude of weighted menu costs, the proof follows exactly as in the proof of proposition 1. □

These two results on sectoral size and menu cost heterogeneity, summarized in proposition 4, can be interpreted as a "least-cost avoider" theory of optimal monetary policy. In the economic analysis of law, the least-cost avoider principal states that when considering assignment of liability between parties, it is efficient to assign liability to the party who has the lowest cost of avoiding harm (Calabresi 1970). Similarly, the generalized principle of optimal monetary policy under menu costs is: *the agents for whom it is least costly to adjust their price are the agents who should do so.*

More closely to the monetary economics literature, this is also very related to the idea that 'monetary policy should target the stickiest price' (e.g. Mankiw and Reis 2003 and

Aoki 2001). Under menu costs, the central bank should minimize adjustment by the firms with the most expensive menu costs – i.e. it should stabilize the stickiest prices.

5.3 Multiple shocks

In the baseline exercise analyzed in proposition 1, we consider a shock to the productivity of sector 1 alone. One motivation for this is the idea that in reality productivity shocks arrive as Poisson shocks, separated by spans of time, with no two sectors ever being shocked at precisely the same time. Such a motivation introduces an intertemporal aspect which we do not study in the analytical model, however – though we will consider the role of dynamics in the quantitative results of section 6.

In this section, we consider the case where an arbitrary set of sectors is shocked, including possibly every sector. Start again at steady state, where every sector has productivity of $A_i^{ss} = 1$. We consider the exercise of shocking every sector to productivity A_i, where A_i could be potentially above, below, or equal to 1: it may be a positive shock, it may be a negative shock, or the sector may be unshocked.

Equilibrium. It is illustrative to consider a generic equilibrium when some fixed subset of sectors $\Omega \subseteq \{1, ..., S\}$ adjusts, while the remaining sectors do not adjust. Denote the cardinality of Ω as $\omega \equiv |\Omega|$. Sectors which adjust update their price to $p_i = M/A_i$, whereas others remain at the steady state value of $p_i^{ss} = 1$. The aggregate price level thus aggregates from (6) to:

$$P = \frac{SM^{\omega/S}}{\prod_{i\in\Omega} A_i^{1/S}}$$

From the quantity equation (2), this gives consumption and output as:

$$C = Y = \frac{1}{S}\left[\prod_{i\in\Omega} A_i^{1/S}\right] M^{\frac{S-\omega}{S}}$$

For comparison, the flexible-price level of output is $Y_{\text{flex}} = \frac{1}{S}\prod_{i=1}^{S} A_i^{1/S}$. Using sectoral demand, production technology, and labor market aggregation, the amount of aggregate labor is:

$$N = \frac{\omega}{S} + \frac{M}{S}\sum_{i\notin\Omega}\frac{1}{A_i} + \psi\omega$$

Thus, welfare – conditional on the set Ω of sectors adjusting – as a function of the choice of money supply, is:

$$\mathbb{W}_{\Omega}(M) = \ln\left[\frac{1}{S}\left[\prod_{i\in\Omega} A_i^{1/S}\right] M^{\frac{S-\omega}{S}}\right] - \left[\frac{\omega}{S} + \frac{M}{S}\sum_{i\notin\Omega}\frac{1}{A_i} + \psi\omega\right] \tag{28}$$

As before allowing the social planner to overcome adjustment externalities, the optimal choice of money supply (conditional on Ω sectors adjusting) is found directly from the first order condition of (28).[23] The first order condition gives:

$$M^*_{\Omega} = \frac{S-\omega}{\sum_{i\notin\Omega}\frac{1}{A_i}} \tag{29}$$

Welfare under optimal policy, conditional on the set Ω of sectors adjusting, is:

$$\mathbb{W}^*_{\Omega} = \ln\left[\frac{1}{S}\left[\prod_{i\in\Omega} A_i^{1/S}\right]\left(\frac{S-\omega}{\sum_{i\notin\Omega}\frac{1}{A_i}}\right)^{\frac{S-\omega}{S}}\right] - [1+\psi\omega] \tag{30}$$

Replicating the flexible-price allocation. Now, it is only possible to replicate the flexible-price allocation – aside from the extra labor required for menu costs – in two cases. First, as before, if all sectors adjust, i.e. $\Omega = \{1, ..., S\}$ and $\omega = S$, then naturally this ensures the flexible-price allocation. This comes at the cost of S sectors' worth of menu costs.

The flexible-price allocation can, however, be achieved also if $\omega = S-1$, so all but one sector adjust. The one non-adjusting sector simply may be any arbitrary sector r. In this case, the central bank would set the money supply at $M = \frac{1}{A_r}$. This ensures that the desired price of sector r, that is $\frac{M}{A_r}$, equals the steady-state level of $p_r^{ss} = 1$, so that despite not changing its price, sector r has the price that it would choose if it were to adjust. This comes at the cost of $S-1$ sectors' worth of menu costs.

An immediate implication is that it is never optimal to have all S sectors to adjust, since the same allocation can be achieved if $S-1$ sectors adjust, but with less labor required for menu costs. In other words, it is always best to peg at least one sector and ensure that at least that one sector need not adjust its price. This sector may be arbitrarily chosen in the baseline model. If sectors were of heterogeneous size or had heterogeneous menu cost sizes, as in proposition 4, then it would be optimal to choose the sector with the largest size-weighted menu costs. This reinforces the "least-cost avoider principle"

[23] As long as the number of sectors adjusting is not all S sectors, i.e. $\omega < S$. If all sectors adjust, i.e. if $\omega = S$ and $\Omega = \emptyset$, then welfare is independent of the choice of the money supply M.

interpretation, or the "stabilize the stickiest price" interpretation described in that proposition.

If more than one sector leaves their price unchanged, i.e. $\omega < S - 1$, then it is not possible to achieve the flexible-price allocation. This is the standard result that when relative prices change, if there is sufficient nominal rigidity, the flexible-price allocation cannot be achieved: there is more than one target (the many relative prices), but only one instrument, M (Poole 1970).

Interpreting conditionally-optimal policy. Since it is still the case that nominal wages are determined by monetary policy, $W = M$, it follows from (29) that nominal wages *conditional on sectors* Ω *adjusting* are:[24]

$$W^*_\Omega = \frac{S - \omega}{\sum_{i \notin \Omega} \frac{1}{A_i}} \tag{31}$$

To emphasize, the equilibrium nominal wage under optimal policy depends on the central bank's choice of the set of adjusting firms Ω. However, for any fixed choice of Ω, this is the equilibrium nominal wage.

From (31), we can see that optimal policy will stabilize nominal wages, $W^*_\Omega = W^{ss}$, if all firms who do not adjust are unshocked (since $A_i = 1$ for unshocked sectors, and the cardinality of Ω is $S - \omega$).

Thus we can summarize optimal policy under multiple shocks as follows.

Proposition 5 (Optimal policy with multiple shocks)**.** Consider an arbitrary set of productivity shocks to the baseline model, $\{A_1, ..., A_S\}$.

1. Conditional on sectors $\Omega \subseteq \{1, ..., S\}$ adjusting, optimal policy is given by setting $M = M^*_\Omega$ defined in (29).

2. The optimal set of sectors that should adjust, Ω^*, is given by comparing welfare under the various possibilities for Ω, using $\mathbb{W}^*_\Omega$ defined in (30).

3. Nominal wage targeting is exactly optimal if the set of sectors which should not adjust, are unshocked: $A_i = 1 \;\; \forall i \notin \Omega^*$.

Even when optimal policy does not exactly stabilize nominal wages, it may nonetheless be considered to approximately do so. For $A_i \approx 1$, it is the case that $\frac{1}{A_i} \approx 1$; this implies that $\sum_{i \notin \Omega} \frac{1}{A_i} \approx \sum_{i \notin \Omega} 1 = S - \omega$, and so by (31) nominal wages are approximately

[24] For any Ω with $\omega < S$. With $\omega = S$, nominal wages are indeterminate and may be anything. We may then for simplicity choose the level given in (31).

unchanged, $W^*_\Omega \approx \frac{S-\omega}{S-\omega} = 1$. As in proposition 1, it will only be optimal to adjust for sectors which experience larger shocks. As a result, for unadjusting sectors, $i \notin \Omega$, it is particularly true that $A_i \approx 1$.

Ultimately, the performance of exact nominal wage targeting in the face of multiple shocks depends on an empirical question – the distribution of shocks, and how tightly centered around 1 they are – and a quantitative question – how well exact nominal wage targeting performs compared to the analytically optimal policy. We answer this question in the quantitative model of section 6.

5.4 Production networks

In this subsection, we consider a variant of the baseline model where intermediate firms use not just labor as a factor of production, but also use other goods as an input. We consider the symmetric roundabout economy of Basu (1995).

We need only alter the intermediate firm's problem. Suppose the production function of firm j in sector i is now:

$$y_i(j) = A_i n_i(j)^\beta I_i(j)^{1-\beta} \tag{32}$$
$$I_i(j) = \prod_{k=1}^{S} I_{ki}(j)^{1/S}$$

Here, $I_i(j)$ is the composite intermediate good used by firm j in sector i to produce output. It has weight $1-\beta$ in production compared to β for labor; setting $\beta = 1$ returns us to the baseline model. The composite intermediate good is a bundle of outputs from every other sector: $I_{ki}(j)$ is output *from* sector k purchased by the firm j in sector i. Output from every other sector is bundled into the symmetric Cobb-Douglas composite, which is used in production.

With input-output linkages, the marginal cost of a given firm now depends on not just wages but also on the price of every other good in the economy. The firm's problem now consists not only of choosing the price $p_i(j)$, but also choosing the input demand $n_i(j)$ and $I_i(j)$. These demand functions can be found from cost minimization.

It can then be shown that nominal marginal cost is a weighted average of nominal wages and the aggregate price index:

$$MC_i(j) = k\frac{W^\beta P^{1-\beta}}{A_i} \tag{33}$$

where k is an unimportant constant.[25] The marginal cost has a weight β on wages, reflecting labor's share of β in the production function (32). The price index has a weight $1-\beta$, reflecting that intermediates have a share of $1-\beta$ in the production function, and the average price of intermediates is precisely the average price of goods in the economy: i.e., the price index.

In response to the baseline shock affecting only sector 1, optimal policy thus is to stabilize this weighted average of wages and prices, with the weights determined by β. This parallels the example in section 5.1, where optimal policy also stabilized a weighted average of wages and prices. To get a sense of the implications for policy, we can consider off-the-shelf calibrations for β: Basu (1995) suggests β is at most 50%; Nakamura and Steinsson (2010) estimate that β is roughly 30%.

Proposition 6 (Roundabout economy)**.** Consider again the positive productivity shock γ affecting sector 1, in the baseline model augmented with the roundabout production technology of (32). For a fixed level of menu costs ψ, there exists a threshold level of productivity $\overline{\gamma} > 1$, such that:

1. If the productivity shock to sector 1 is above the threshold, $\gamma \geq \overline{\gamma}$, then optimal policy is to ensure the nominal marginal costs of firms outside sector 1 are unchanged. This is implemented by stabilizing $W^{\beta}P^{1-\beta}$. This results in firms in sector 1 adjusting their price, while firms in other sectors k leave prices unchanged.

2. If the shock is below the threshold, $\gamma \in [1, \overline{\gamma})$, then optimal policy is to ensure that prices remain unchanged and no firm in any sector adjusts.

Proof: The proof follows exactly the same steps as in the proof of proposition 1. □

5.5 Sticky wages

Throughout the paper thus far, we have considered a model with sticky prices and flexible wages; we now consider the case of flexible prices and sticky wages. Optimal policy in response to the same sectoral productivity shock continues to stabilize nominal wages, again in order stabilize the nominal marginal cost of unshocked sectors and thus minimize menu cost expenditure.

Setup. Consider the baseline model of section 2, but now allow for heterogenous types of labor organized into a union with wage-setting power. Suppose there are S labor sectors, and each goods-producing sector $i = 1, ..., S$ hires labor exclusively from the corre-

[25] $k = \frac{1}{(1-\alpha)^{1-\alpha}\alpha^{\alpha}}$.

sponding labor sector $i = 1, ..., S$. Within each labor sector, there is a continuum of differentiated worker types, indexed on $[0, 1]$. Continue to denote the total amount of labor used by goods-producing firm j in sector i as $n_i(j)$, and continue to endow intermediate firms with technology $y_i(j) = A_i n_i(j)$.

Firm-level labor input $n_i(j)$ is now, unlike in the baseline model, composed of a CES bundle of workers from labor sector i:

$$n_i(j) = \left[\int_0^1 \left(n_i^k(j) \right)^{\frac{\varepsilon-1}{\varepsilon}} dk \right]^{\frac{\varepsilon}{\varepsilon-1}}$$

where $n_i^k(j)$ denotes the quantity of labor of type k in labor sector i hired by firm j in goods-producing sector i; and ε is the elasticity of substitution across labor types. Cost minimization produces the standard demand curve, $n_i^k(j) = \left(\frac{W_i(k)}{W_i} \right)^{-\varepsilon} n_i(j)$, where $W_i(k)$ is the nominal wage of labor type k in sector i and W_i is the wage index for sector i labor. This results in a profit function of $D_i(j) = p_i(j) y_i(j) - W_i n_i(j)$.

For firms in sector i, the optimal reset price (i.e. the nominal marginal cost) now depends on the *sector-specific* wage and sectoral productivity:

$$p_i^{\text{flex}}(j) = MC_i(j) = \frac{W_i}{A_i} \tag{34}$$

Because prices are now flexible – that is, firms face $\psi = 0$ – firms always set price equal to this optimal reset price.

The nominal wage for worker type j in sector i, $W_i(j)$, is set by a union which must pay a fixed cost for nominal wage changes: the "menu cost". This can be motivated by a fixed cost of contract renegotiation. The union's problem can be considered as part of the representative household's problem, which is written as:

$$\max_{C, M, \{W_i(j)\}_{i,j}} \ln C + \ln \left(\frac{M}{P} \right) - \sum_{i=1}^{S} \int_0^1 N_i(j) dj$$

$$\text{s.t. } PC + M = \sum_{i=1}^{S} \int_0^1 W_i(j) N_i(j) (1 + \tau^W) dj + M_{-1} + D - T - \psi^W \chi^W$$

$$N_i(j) = \left(\frac{W_i(j)}{W_i} \right)^{-\varepsilon} N_i$$

This utility maximization problem differs from that in section 2.1 in a few ways. The household sets nominal wages, given the demand curves for labor types, analogous to the price-setting problem of intermediate firms. The household receives a labor subsidy

τ^W to offset the monopoly distortion from its wage-setting power, analogous to the labor subsidy that firms are given to offset the monopoly distortion from firm price-setting power. Finally, the household must pay a menu cost ψ^W if it wishes to change any wage, analogous to the menu cost facing firms. The variable χ^W measures the mass of wages which are changed:

$$\chi^W \equiv \sum_i \int_j \chi_i^W(j) dj$$
$$\chi_i^W(j) \equiv \mathbb{I}\{W_i(j) \neq W_i^{\text{old}}(j)\}$$

where $W_i^{\text{old}}(j)$ is the inherited nominal wage for type j in sector i, analogous to $p_i^{\text{old}}(j)$ in the firm's problem.

The optimality conditions of this optimization problem include the equation of exchange (2) but also a new condition for optimal wage-setting. Under the optimal labor subsidy of $\tau^W = \frac{1}{\varepsilon - 1}$, the optimal wage-setting condition *conditional on adjusting* is determined by the marginal rate of substitution between consumption and leisure:

$$W_i^{\text{flex}}(j) = PC = M \quad \forall i, j \tag{35}$$

Note that under flexible wages, wages across types and sectors are all equalized.

In equilibrium, the wage menu cost paid by the household creates a wedge between consumption and output:

$$C = Y - \psi^W \chi^W \tag{36}$$

This aggregate resource constraint is derived from the household budget constraint and market clearing conditions.

We summarize the important differences with the baseline model. First, the nominal marginal costs of the intermediate goods producers (34) now depend on a sector-specific wage, $MC_i(j) = W_i / A_i$. Firms always set price equal to this level because prices are flexible. Second, in the efficient flexible-wage equilibrium, all nominal wages are equalized from (35). Third, wage menu costs result in a wedge between consumption and output, from (36), and lower welfare.

Optimal policy after a sectoral shock. Consider again the same exercise: starting from a steady state with $A_i^{ss} = 1$ and $M = 1$, shock productivity of sector-1 goods producers, $A_1 = \gamma > 1$. For clarity, consider the case where γ is sufficiently large, so that we do not need to discuss an analog to the $\overline{\gamma}$ of proposition 1.

The optimal reset prices and wages – i.e., those that would prevail in the frictionless equilibrium – are:

$$\begin{aligned} p_1^{\text{flex}}(j) &= \frac{W_1}{\gamma} \\ p_k^{\text{flex}}(j) &= W_k \quad \forall k > 1 \\ W_i^{\text{flex}}(j) &= M \quad\ \ \forall i, j \end{aligned}$$

To minimize the amount of menu costs and simultaneously achieve correct relative prices, it is again desirable to leave M unchanged, thereby stabilizing nominal wages. This ensures that no wages need to be adjusted and no wage menu costs need to be paid: that is, $W_i = W_i^{ss}$ for all i. Meanwhile, given the shock was assumed to be sufficiently large, goods-producing firms in sector 1 can update their prices, thus ensuring all relative prices are correct.

In short, optimal policy stabilizes all nominal wages, which ensures correct relative prices and causes only sector-1 firms to adjust.

5.6 Optimal policy without selection effects

The existence (or not) of selection effects in menu cost models is an important question in the literature, due to the argument that selection effects reduce monetary non-neutrality relative to models with time-dependent pricing like the Calvo model (Golosov and Lucas 2007; Caballero and Engel 2007; Carvalho and Kryvtsov 2021; Karadi, Schoenle and Wursten 2022; Gautier et al. 2022). The question this literature generally considers is: in response to a *monetary policy shock*, how much is real output affected? On the other hand, under optimal monetary policy naturally there are no monetary shocks.

In this subsection, we show that the existence or not of selection effects plays little role. We demonstrate this by briefly characterizing a "CalvoPlus" variant of the baseline model (Nakamura and Steinsson 2010). In the CalvoPlus variant, the setup is precisely as in section 2, except that menu costs are now idiosyncratic at the firm level, $\psi_i(j)$. In particular, a random fraction $\zeta \in (0, 1)$ of firms in each has an opportunity to adjust price for free, $\psi_i(j) = 0$; other firms face the nonzero menu cost, $\psi_i(j) = \psi$.

As the next proposition describes, optimal policy is in essence the same as in the baseline result of proposition 1, though the threshold $\overline{\gamma}$ is altered.

Proposition 7 (Optimal policy under CalvoPlus)**.** Consider the baseline model, modified so that a random fraction $\zeta \in (0, 1)$ of firms in each sector face no menu cost, $\psi_i(j) = 0$, and remaining firms face $\psi_i(j) = 1$. For a fixed level of menu costs ψ, there exists a threshold level of productivity $\overline{\gamma}_{\text{CalvoPlus}} > 1$, such that:

1. If the productivity shock to sector 1 is above the threshold, $\gamma \geq \overline{\gamma}_{\text{CalvoPlus}}$, then optimal policy is exactly nominal wage targeting: monetary policy should ensure $W = W^{ss}$. This results in all firms in sector 1 adjusting their price, while all firms in other sectors k leave prices unchanged regardless of whether they have a free adjustment.

2. If the shock is below the threshold, $\gamma \in [1, \overline{\gamma}_{\text{CalvoPlus}})$, then optimal policy is to ensure that prices remain unchanged for firms with a nonzero menu cost, $\psi_i(j) > 0$.

3. The threshold $\overline{\gamma}_{\text{CalvoPlus}}$ is smaller than in the baseline model:

$$\overline{\gamma}_{\text{CalvoPlus}} < \overline{\gamma} \tag{37}$$

We sketch the proof to provide economic intuition. The logic follows very closely to the proof of proposition 1.

By stabilizing W, the central bank ensures that the nominal marginal cost of unshocked firms remains unchanged. Unshocked firms thus have no desire to change price, even if they have a free opportunity to do so. In this case, correct relative prices are achieved, at a welfare cost of $\zeta\psi$ quantity of menu costs. For large enough shocks, this is optimal, because the welfare gains from correct relative prices outweighs the loss from paying the $\zeta\psi$ menu cost.

For small shocks, however, menu costs are too costly to be worthwhile. In this case, firms with a free adjustment will adjust – creating within-sector price dispersion – but it is optimal to ensure that firms with a nonzero menu cost do not adjust. In fact, *within* this case, the equilibrium is as-if Calvo: an exogenous fraction of firms in each sector is allowed to update prices for free, and the remaining firms do not adjust – here, endogenously, unlike Calvo. Optimal policy thus precisely replicates the Calvo optimum, *if* $\gamma < \overline{\gamma}_{\text{CalvoPlus}}$.

Observe that in the case where firms in sector 1 adjust, the welfare loss from menu costs is $\zeta\psi$, which is smaller than the level of ψ in the non-CalvoPlus economy. Thus, the fixed welfare loss from menu costs in this case is smaller. As a result, the increase in productivity needed to make price adjustment overcome this welfare loss is correspondingly smaller. This explains the third part of the proposition.

6 Quantitative model

In this section, we develop a dynamic version of our baseline model, augmented with idiosyncratic shocks as well as more general functional forms, calibrated to the US economy.

6.1 Model description and solution method

Our dynamic multisector model of menu costs is similar to that in Nakamura and Steinsson (2010). The main difference is that we include sector-specific productivity shocks on top of idiosyncratic, firm-level shocks. To solve for the aggregate dynamics of the economy, we use the sequence-space Jacobian approach of Auclert et al. (2021).

Household. The household chooses paths for consumption, C_t, labor N_t, money balances, M_t, and bonds, B_t to maximize the present discounted value of utility. The problem faced by the household is:

$$\begin{aligned} &\max_{\{C_t,N_t,B_t,M_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \left[\frac{C_t^{1-\gamma}}{1-\gamma} - \omega \frac{N_t^{1+\varphi}}{1+\varphi} + \ln\left(\frac{M_t}{P_t}\right)\right] \qquad (38) \\ \text{s.t.} \quad & P_t C_t + B_t + M_t \leq R_t B_{t-1} + W_t N_t + M_{t-1} + D_t - T_t \end{aligned}$$

This problem represents a dynamic version of the static household problem presented in section 2.1, with more general preferences. To consume and save (the left hand side of the budget constraint) the household uses gross wealth, money holdings, labor earnings, and firm dividends net of the lump sum tax imposed by the government (the right hand side of the budget constraint). R_t is the nominal interest rate on bonds.

Firms. The final good producer and sectoral good producers behave the same as in the baseline analytical model of section 2. These firms operate in competitive environments and aggregate the goods produced by the corresponding lower-tier firms according to the technologies:

$$Y_t = \prod_{i=1}^{S} y_{it}^{1/S} \qquad (39)$$

$$y_{it} = \left[\int_0^1 y_{it}(j)^{\frac{\eta-1}{\eta}}\, dj\right]^{\frac{\eta}{\eta-1}} \qquad (40)$$

Intermediate firms. A first key difference in the quantitative model compared to the baseline model of section 2 is that intermediate firms are now subject not only to sector-level productivity shocks, A_{it}, but also to idiosyncratic, firm-level shocks, $a_{it}(j)$. A second difference is that firms' production technology displays potentially decreasing returns to labor with parameter $\alpha \leq 1$:

$$y_{it}(j) = a_{it}(j) A_{it} n_{it}(j)^{\alpha} \tag{41}$$

Firm-level idiosyncratic shocks follow an AR(1) process with persistence ρ_{idio}. The innovations in these processes are Gaussian, $\varepsilon_{it}^{\text{idio}}(j) \sim \mathcal{N}\left(0, \sigma_{\text{idio}}^2\right)$.

$$\log\left(a_{it}(j)\right) = \rho_{\text{idio}} \log\left(a_{it-1}(j)\right) + \varepsilon_{it}^{\text{idio}}(j) \tag{42}$$

The firm maximizes the present discounted value of real profits. In any given period the firm chooses whether to update its price, $\chi_{it}(j) = 1$, or to keep its price unchanged, $\chi_{it}(j) = 0$. If the firm decides to change its price $p_{it}(j)$, it must pay a menu cost worth ψ hours of labor. The problem intermediate firms face is:

$$\begin{aligned}
\max_{p_{it}(j)} \quad & \sum_{t=0}^{\infty} \mathbb{E}\left[\frac{1}{R^t P_t}\left\{p_{it}(j)\, y_{it}(j) - W_t n_{it}(j)\,(1-\tau) - \chi_{it}(j)\psi W_t\right\}\right] \qquad (43)\\
\text{s.t.} \quad & \chi_{it}(j) = \begin{cases} 1 & \text{if } p_{it}(j) \neq p_{it-1}(j) \\ 0 & \text{otherwise} \end{cases} \\
& y_{it}(j) = y_{it}\left(\frac{p_{it}(j)}{p_{it}}\right)^{-\eta} \\
& y_{it}(j) = A_{it} a_{it}(j)\, n_{it}(j)^{\alpha} \\
& R^t = \prod_{\tau=0}^{t} R_{\tau}
\end{aligned}$$

where τ is again the labor subsidy.

Market clearing. As in the baseline model, labor market clearing requires that labor supplied by the household equals labor used in production plus labor used in menu costs

$$N_t = \sum_i \int_0^1 n_{it}(j) dj + \psi \chi_t \tag{44}$$

where χ_t is the mass of firms that adjust prices in period t:

$$\chi_t \equiv \sum_i \int_0^1 \chi_{it}(j)dj \tag{45}$$

Solution method. To solve the model numerically, we use the sequence-space Jacobian method developed in Auclert et al. (2021). This method provides linearized general equilibrium responses with respect to perfect-foresight shocks to aggregate variables, while allowing agents' decision rules to be nonlinear with respect to idiosyncratic variables. Note that, by certainty equivalence, the linearized perfect-foresight transition paths that we show here are equal to the first-order perturbation solution of the model with aggregate risk.

6.2 Calibration

The model is calibrated to match salient micro and macro moments of the US economy at the quarterly frequency. There are two sets of parameters. The first set of parameters is standard and taken from the macroeconomics literature. The second is calibrated to match price-adjustment behavior by US firms. The model parameters are listed in table 1.

The preference parameters are set to standard values. The discount factor $\beta = 0.99$ is chosen to match a 4% annualized interest rate. The disutility of labor is $\omega = 1$, the inverse Frisch elasticity is $\varphi = 0$ as in Golosov and Lucas (2007), and the inverse elasticity of intertemporal substitution is $\gamma = 2$.

Following Nakamura and Steinsson (2010), we choose $S = 6$ sectors. We assume the sectors are identical in their structural parameters: firms in different sectors are subject to the same idiosyncratic productivity processes and face the same menu cost. However, sectors are subject to heterogeneous sectoral shocks. The baseline elasticity of substitution across goods within sector is $\eta = 5$ and the decreasing return to scale parameter is $\alpha = 0.6$. The labor subsidy $\tau = \frac{1}{\eta} = 0.2$ is set to offset the markup distortion.

We select three parameters to match the price-changing behavior of US firms. These parameters are the standard deviation and persistence of idiosyncratic productivity shocks, σ_{idio} and ρ_{idio}, as well as the size of the menu cost ψ, that is the hours of labor required to change prices. These parameters are set to match two targets. First, per the literature cited in section 4.2, we target a menu cost expenditure as a share of firm revenue of 1%. Second, following Nakamura and Steinsson (2010), we target a median quarterly frequency of price change of 26.1%.[26] The resulting estimated parameters are $\sigma_{\text{idio}} = 0.13$,

[26] They find a monthly median of 8.7%.

$\rho_{\text{idio}} = 0.86$, and $\psi = 0.016$ leading to a menu cost to revenue share of 1.06% and a share of firms changing price in each quarter of 25.7%.

Table 1: Model parameters and baseline calibration

	Parameter (quarterly frequency)	Value	Target
β	Discount factor	0.99	standard
ω	Disutility of labor	1	standard
φ	Inverse Frisch elasticity	0	Golosov and Lucas (2007)
γ	Inverse EIS	2	standard
S	Number of sectors	6	Nakamura and Steinsson (2010)
η	Elasticity of subst. between sectors	5	standard value
α	Returns to scale	0.6	standard value
τ	Labor subsidy	0.2	$1/\eta$
σ_{idio}	Standard deviation of idio. shocks	0.13	menu cost expenditure / revenue $\sim$ 1%
ρ_{idio}	Persistence of idio. shocks	0.86	and
ψ	Menu cost	0.016	share of price changers $\sim$ 26.1%

6.3 Comparing inflation targeting and nominal wage targeting

We begin by comparing the performance of nominal wage targeting versus inflation targeting after a temporary but persistent shock to sector-1 productivity, i.e. the same sort of shock analyzed in section 3. We study the two policies as simple monetary policy rules, motivated by the analytical model, before turning to more general policies in later subsections.

Prices and quantities. Consider a shock to sector-1 productivity A_1 of 5%, which decays exponentially, while sectoral productivity is unshocked for all other sectors. Figure 2 displays the time paths of A_1, consumption, labor, and prices in percentage point deviations from steady state. Three cases are depicted: nominal wage targeting, inflation targeting, and flexible prices.

Observe in the first column that under nominal wage targeting, besides nominal wages W remaining stable, the aggregate price level P falls in response to the positive productivity shock, driven almost entirely by the fall in sector-1 price p_1. Other sectoral prices move, because of decreasing returns to scale, but minimally so. Meanwhile, under inflation targeting in the second column, P is constant and W rises, driven by a smaller fall in p_1 and an increase in all other sectoral prices p_k.

As a benchmark for comparison, the third column shows the impulse responses under flexible prices, in which $\psi = 0$. In this economy, we also renormalize steady state sectoral

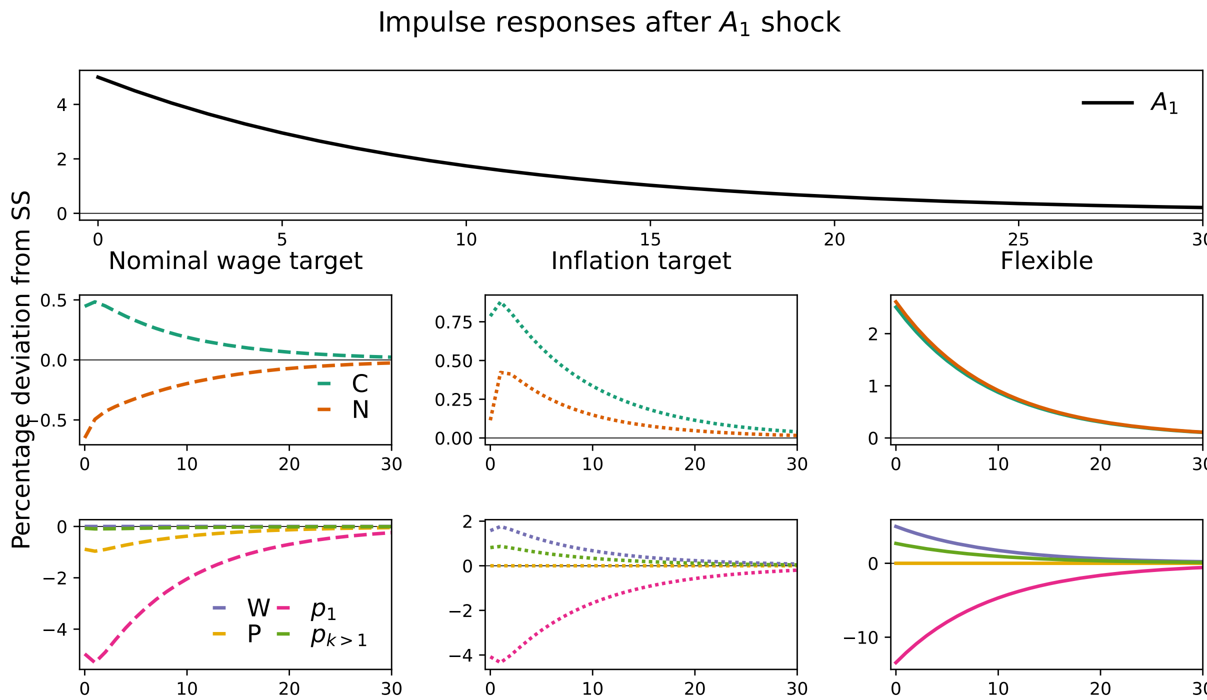

Figure 2: Sector-1 sectoral productivity is increased by a 5% shock which decays at rate ρ_{sect}. The first column shows outcomes under nominal wage targeting; the second under inflation targeting; and the third column under the flexible-price benchmark.

productivity, A_i^{ss} so that steady state welfare in this flex-price economy exactly matches steady state welfare in the menu cost model.[27]

Observe that output follows a hump-shaped response under menu costs. Under nominal wage targeting, aggregate inflation also follows a hump shape. Under either monetary policy rule, sector-level prices are hump-shaped.

Menu cost expenditure. The total quantity of labor N depicted above can be decomposed into labor used in production and menu cost expenditure. We denote labor used in production as $N_y \equiv \sum_i \int_j n_i(j)dj$, while labor used in menu costs is $N - N_y = \psi\chi$.

The left panel of figure 3 shows this decomposition of labor for both inflation targeting (dotted) and nominal wage targeting (dashed). The gap between N and N_y represents the labor used in menu costs. This gap is notably larger and more persistent 'under inflation targeting than under nominal wage targeting, indicating that, following the productivity shock, the menu costs expended under inflation targeting are considrably larger than those expended under nominal wage targeting. This is depicted in the right panel of the

[27] $A_i^{ss} = 1$ in the menu cost economy and $A_i^{ss} = 0.98125$ in the flex-price economy.

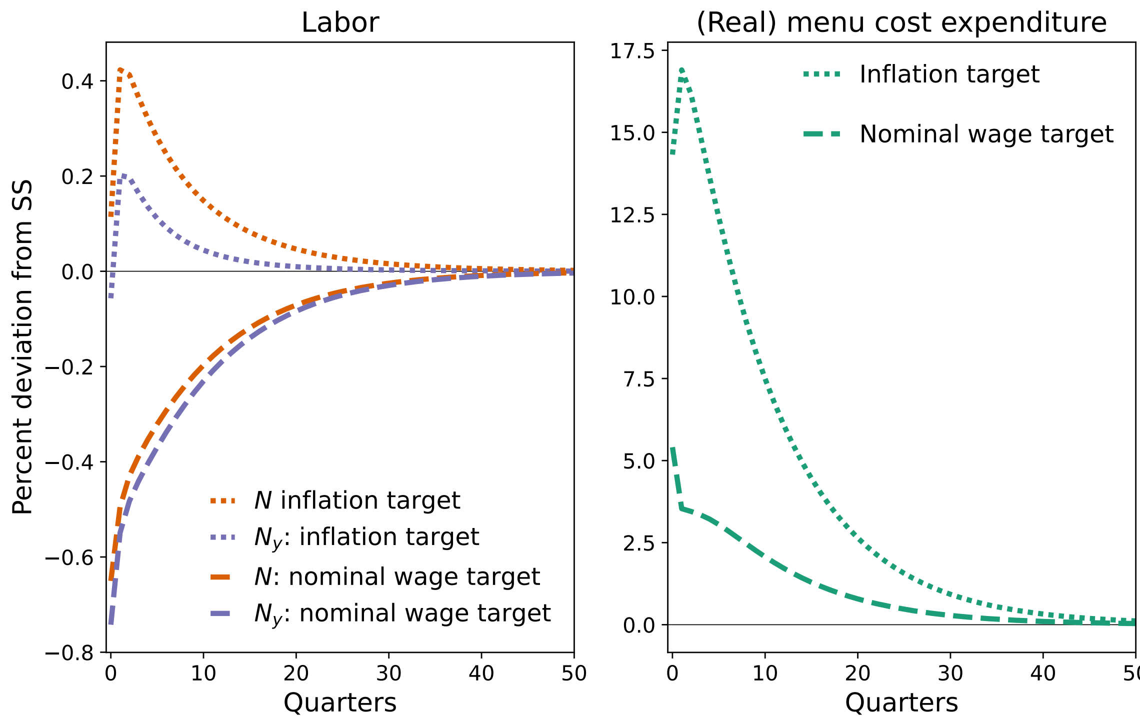

Figure 3: Percent deviation from steady state in labor used on menu costs after the persistent 5% increase in A_1, under nominal wage targeting versus inflation targeting.

figure which displays *real* menu cost expenditures in both regimes. This result agrees with the takeaway from the static model that wage targeting economizes on menu cost expenditure.

Welfare. The menu cost expenditure described above, along with consumption and labor used in production, determine the welfare response to the productivity shock, which is shown in figure 4.

To make interpretable the gain in welfare from moving from an inflation targeting regime to a nominal wage targeting regime, we convert the welfare differences to consumption units. Denote the paths for consumption and labor as $\{C_t^x\}$ and $\{N_t^x\}$, respectively. x refers to one of the three possible economies: the nominal wage targeting, inflation targeting, or flexible economies, with $x \in \{W, P, \text{flex}\}$. Denote the household period utility function from (38) as $U\left(C_t^x, N_t^x\right)$.

We ask: with menu costs, how much higher would the path for consumption need to be, such that lifetime household welfare equals that in the flex-price world. That is, to a

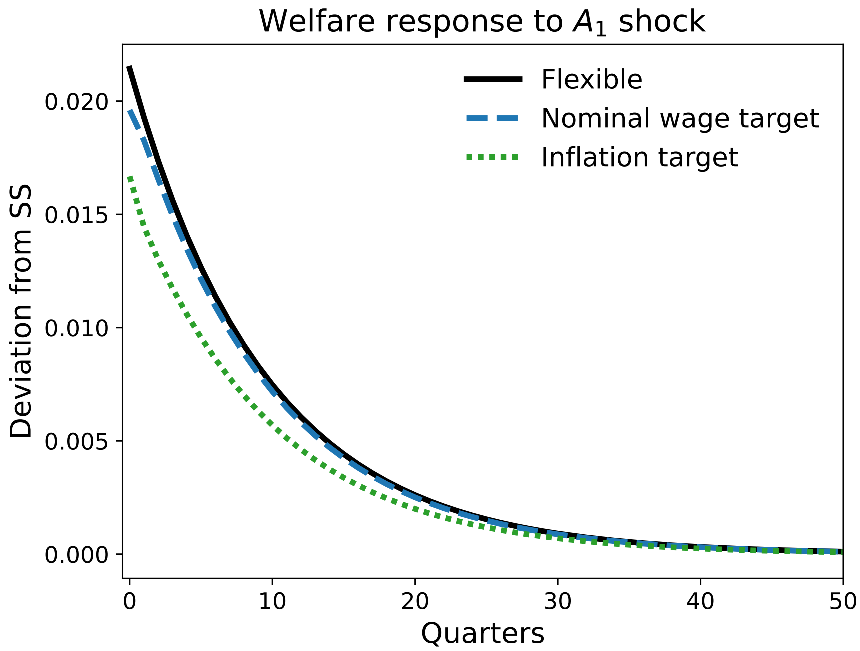

Figure 4: Deviation from steady state in welfare after the persistent 5% increase in A_1 under nominal wage targeting, under inflation targeting targeting, and in the flex-price benchmark.

first order approximation, what is the λ^x which solves:

$$\sum_{t=0}^{\infty} \beta^t U\left(\left(1+\lambda^x\right) C_t^x, N_t^x\right) = \sum_{t=0}^{\infty} \beta^t U\left(C_t^{\text{flex}}, N_t^{\text{flex}}\right) \tag{46}$$

We solve for the λ^W measuring by how much consumption needs to be scaled up under nominal wage targeting to ensure welfare matches flex-price welfare; λ^P is defined analogously for inflation targeting.[28]

We find that $\lambda^W = 0.004\%$ and $\lambda^P = 0.02\%$. This implies that moving from inflation to nominal wage targeting reduces 80.6% of the welfare loss caused by sticky prices after

[28] Since we are working with first-order approximations, λ has an analytical formula. First, define Δ as the first order approximation to the difference between welfare under sticky prices and flexible prices, i.e. the first order approximation to $\sum_{t=0}^{T} \beta^t \left[U\left(C_t^x, N_t^x\right) - U\left(C_t^{\text{flex}}, N_t^{\text{flex}}\right)\right]$. Then:

$$\begin{aligned} \lambda &= \frac{1}{\sum_{t=0}^{T} \beta^t} \frac{1}{C^{ss} U_C\left(C^{ss}, N^{ss}\right)} \Delta \\ &\to (1-\beta) \frac{1}{C^{ss} U_C\left(C^{ss}, N^{ss}\right)} \Delta \end{aligned}$$

where C^{ss}, N^{ss} refer to the steady state variables in the sticky price economy. The interpretation of λ can be seen here: it takes the utility difference Δ, converts it to consumption units by dividing by steady state marginal utility U_C, and takes the ratio with steady state consumption C^{ss}.

the sectoral shock.

This number should be interpreted carefully. As described above, we renormalized steady state sectoral productivity in the flex-price world such that *steady state* welfare is the same under flexible and sticky prices. (Note that steady state welfare is not affected by the choice of monetary policy regime.) Without this adjustment to flex-price productivity, the welfare *responses* cannot be easily compared across the two cases, since the responses would be on top of different baselines. Much of the welfare loss due to menu costs arises from this gap in *steady state welfare*. The λ results describe the reduction in the welfare loss moving from inflation to nominal wage targeting that is *purely* due to the shock.

6.4 Decomposing welfare: direct effects vs. efficiency effects

To understand the welfare effects of different monetary regimes we decompose welfare losses into *direct* losses arising from menu costs and *efficiency* losses arising from incorrect relative prices. We start by explaining the modified model that will produce this decomposition and then we define each of these terms.

"Modified model" with no direct costs. We build an alternative version of the model where menu costs have no *direct* effect on welfare, in the following sense.

Firms solve for their policy functions as if menu costs were nonzero, but aggregate labor demand, N, and firm profits, $D_i(j)$, are computed setting $\psi = 0$. In other words, menu costs are "rebated" to the household in the form of lower labor and consequently lower disutility. Much of the non-normative literature rebates menu costs to households through lump-sum transfers (e.g. Auclert, Rognlie and Straub (2018) or Guerrieri et al. (2021), in settings with convex menu costs), while here the rebate is in units of labor.[29] Note that because the wage payment for menu costs is simply a transfer from firm profits to household labor income and in equilibrium households receive firm profits, the net effect of this modification is *only* to reduce labor demand.

To summarize, this modified model is the same as the benchmark model in section 6.1, except that the labor market clearing condition becomes simply

$$\tilde{N}_t = \sum_i \int \tilde{n}_{it} j dj \tag{47}$$

[29] We also normalize steady state sectoral productivities in this model with rebates so that steady state welfare matches that of the standard model, just as we did for the flex-price model, as described in the previous section. In particular, denoting with a tilde variables in the model with rebates, $\tilde{A}_i^{ss} = 0.988$.

whereas before (44) was

$$N_t = \sum_i \int n_{it}(j)dj + \psi\chi_t$$

where variables with tildes are those in the modified model.

Direct costs vs. efficiency costs. Furthermore, the dynamics of this modified model are *precisely* the same as those of the benchmark model except that welfare is higher by the amount of menu cost labor which would otherwise be required, $\psi\chi$. This follows from $\varphi = 0$, à la Golosov and Lucas (2007): all income effects on the household accrue to labor, so the income effects of reduced labor demand in this model do not affect household consumption.

As a result, we can define "direct loss" from menu costs in period t as precisely $\psi\chi_t$. This is the reduction in welfare, in the benchmark model, that comes *directly* from higher labor demand due to menu costs.

In turn, we can define the "efficiency costs" as the gap between the flexible-price response and the modified model without direct costs. This gap reflects only incorrect relative prices, which we term efficiency costs.

To make these terms precise, define

$$\text{direct losses}_t = \psi\chi_t \tag{48}$$

$$\text{efficiency losses}_t = U\left(C_t^{\text{flex}}, N_t^{\text{flex}}\right) - U\left(\tilde{C}_t, \tilde{N}_t\right) \tag{49}$$

where variables with tildes are those in the modified model.

Decomposition. Figure 5 decomposes the welfare gap into the "efficiency costs" and the "direct costs". These gaps are shown for both nominal wage targeting (left panel) and inflation targeting (right panel). The dark shaded areas reflect the efficiency costs: under either monetary regime, how much do incorrect relative prices impact welfare. The translucent shaded areas reflect the direct costs of menu costs: how much lower is welfare due to households mechanically needing to supply more labor for price adjustment.

To interpret these findings quantitatively, we convert to consumption units. We ask, analogously to before: how much higher would the path for consumption need to be such that lifetime household welfare *without the direct effects of menu costs* equals welfare in the

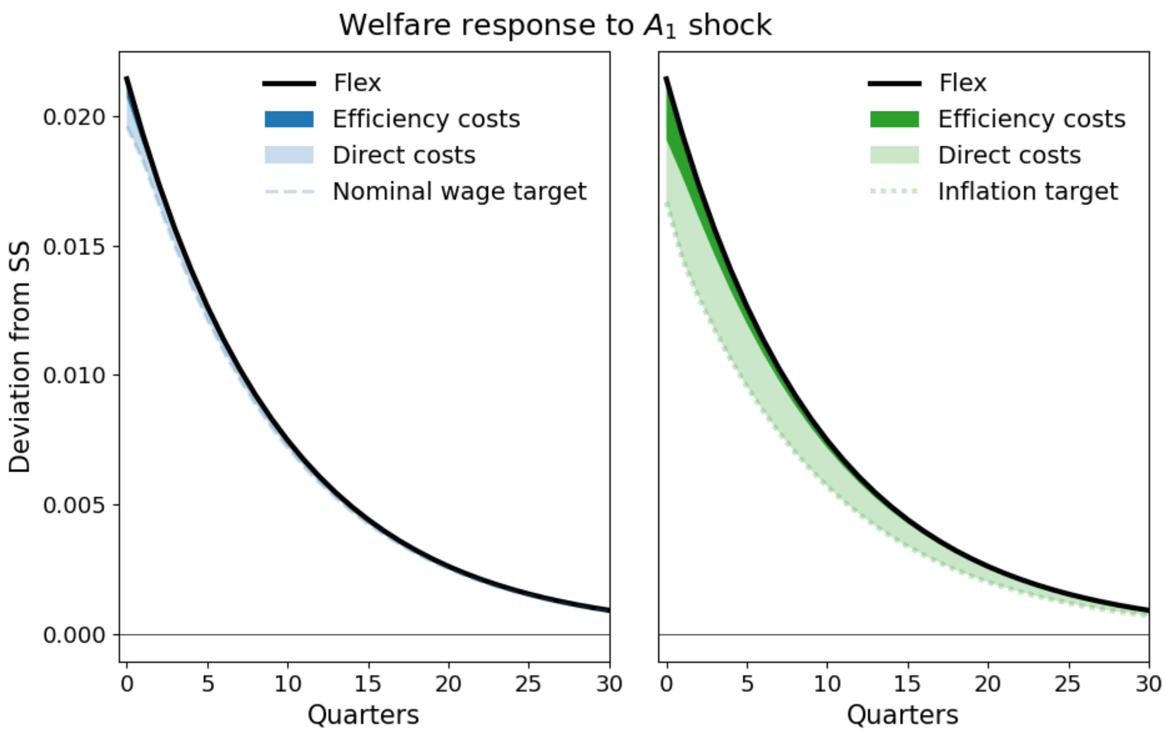

Figure 5: Welfare response after a persistent 5% increase in sector-1 productivity, A_1. The black line shows the response under flexible prices, which is the first best. The colored lines show the welfare response in the presence of menu costs, under a nominal wage target (left) and an inflation target (right). The dark shaded area represents "efficiency costs", that is, the welfare loss from incorrect relative prices only. The translucent shaded area represents the "direct costs" of menu costs, that is, the mechanical welfare loss from additional labor required to adjust prices.

flex-price world? This quantity is given by $\tilde{\lambda}^x$ in the following equation:

$$\sum_{t=0}^{\infty} \beta^t U\left(\left(1+\tilde{\lambda}^x\right) \tilde{C}_t^x, \tilde{N}_t^x\right) = \sum_{t=0}^{\infty} \beta^t U\left(C_t^{\text{flex}}, N_t^{\text{flex}}\right) \tag{50}$$

We solve to find $\tilde{\lambda}^W = 0.0007\%$ and $\tilde{\lambda}^P = 0.0060\%$, reflecting the welfare losses without direct costs. Recall that, for comparison, $\lambda^W = 0.0040\%$ and $\lambda^P = 0.0200\%$, reflecting total welfare losses.

These results make it clear that the welfare losses from direct menu costs are substantially larger under inflation targeting than under nominal wage targeting. Furthermore, while efficiency costs are a relatively small contributor to the welfare loss under both regimes, they are also considerably larger under inflation targeting.

6.5 Optimal policy within a class of simple monetary policy rules

In this section, we solve numerically for optimal policy within the class of rules targeting a weighted average of nominal wages and prices. More specifically, we consider monetary policy rules of the form

$$W_t^{\xi} P_t^{1-\xi} = (W^{ss})^{\xi} (P^{ss})^{1-\xi} \tag{51}$$

and solve over a grid of $\xi \in [0, 1]$ for the value that maximizes welfare following the same shock to A_1 described in the previous subsections.

For each ξ, we solve for λ of equation (46) that expresses the welfare loss from the flex-price economy in units of consumption. We also use equation (50) to decompose this loss into direct costs and efficiency costs.

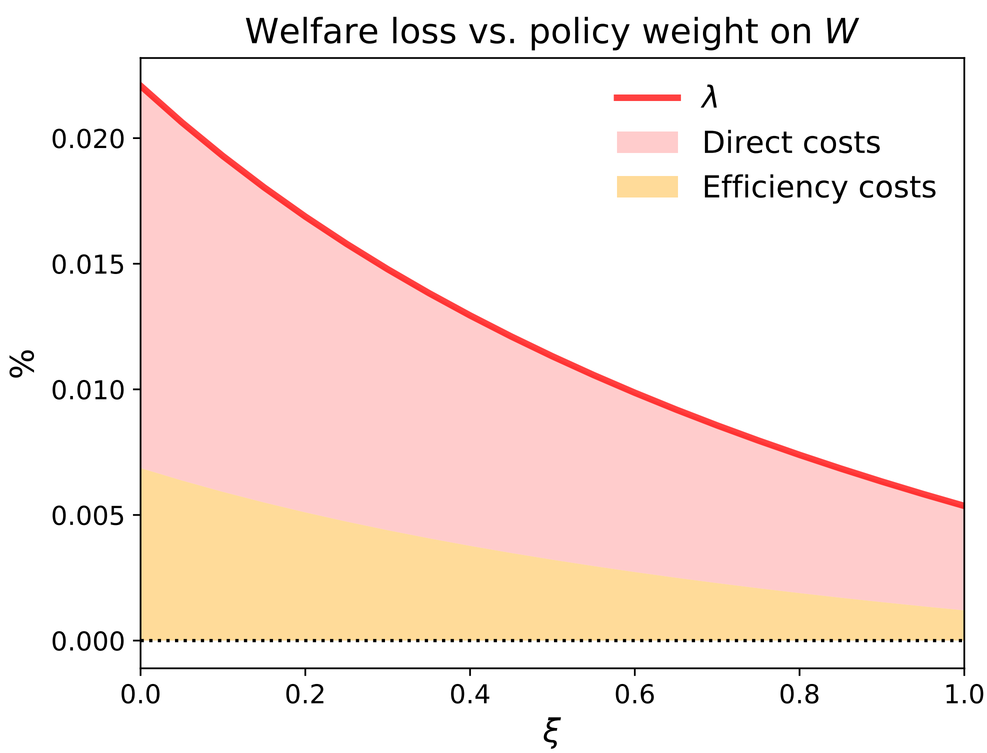

Figure 6: The welfare loss after the A_1 shock, in units of consumption, in an economy under a monetary policy rule stabilizing a weighted average of nominal wages and prices, with a weight of ξ on nominal wages.

Figure 6 shows that, in this class of rules, nominal wage targeting ($\xi = 1$, the rightmost point) minimizes the welfare loss after the A_1 shock and inflation targeting ($\xi = 0$, the leftmost point) maximizes the loss. Furthermore, the welfare loss strictly decreases with the weight on nominal wages; and this is true for both the direct and efficiency components of the welfare loss.

7 Conclusion

Summarizing. Consider an economy with S sectors, where firms within each sector are subject to sector-specific productivity shocks. As an example, suppose firms in sector 1 are hit by a positive productivity shock. If the shock is sufficiently large, then it is efficient and desirable for firms in this sector to cut their *relative* prices, compared to other firms in other sectors of the economy. Under a monetary policy that stabilizes the nominal marginal costs of firms outside of sector 1, such firms have no desire to adjust their prices. Meanwhile, firms in sector 1 choose to adjust their nominal prices because of the productivity shock. As a result, relative prices between sector-1 firms and other firms are undistorted and *only* the one shocked sector has incurred wasteful menu costs. This has a natural "least-cost avoider" interpretation, reflecting the desire to stabilize the stickiest type of price – where, thanks to menu costs, the stickiest type is endogenous to the shock.

This logic is formalized and generalized in our analytical model and explored quantitatively in our computational results. We revisit the question of optimal monetary under sticky prices using a more realistic microfoundation for sticky prices – menu costs – than the benchmark New Keynesian model. Our analytical approach shows, without linearization, that the textbook prescription for inflation stabilization is not optimal under the more realistic foundation of menu costs, in response to sectoral shocks. Instead, the central bank should stabilize the nominal marginal cost of unshocked sectors. In our quantitative model, the welfare loss from implementing inflation targeting rather than nominal wage targeting is large, due to the sizable estimates in the literature for the empirical magnitude of menu costs.

Practical implementation. A full analysis of practical considerations for implementing optimal policy is beyond the scope of this paper and is a topic ripe for future research. One important question for policymakers may be the choice of which empirical measure of nominal wages to track in order to best track nominal marginal costs. A similar issue arises for inflation targeting, where policymakers and analysts frequently debate the relevance of "headline" versus "core" inflation; or debate the use of a consumer price index versus personal consumption expenditures; or consider "trimming" components of any of these price indices. A second question may be the role of revisions in nominal wage statistics, although once again there is a similar issue for tracking inflation since inflation revisions are often quite sizeable (Audoly et al. 2023). Third and more generally, optimal policy under menu costs may be affected by information constraints facing the central bank. Indeed, this possibility helps motivate our quantitative analysis of simple monetary policy rules in section 6.5. This is plausibly quite important: for example, a lack of knowledge of the true level of the output gap and the natural rate of interest have been an important challenge for monetary policy historically (Orphanides 2003; Gorodnichenko and Shapiro 2007), as emphasized conceptually by Friedman (1968).

Future work. The conclusion that optimal monetary policy in response to sectoral shocks should result in countercylical inflation resonates with the results of other studies in the broader optimal monetary policy literature away from Calvo sticky prices, as discussed in the introduction (Sheedy 2014; Angeletos and La'O 2020; Selgin 1997).[30] Integrating these varied approaches into a unified theory of optimal monetary policy is an open question for future research.

[30] More generally, the optimality of countercyclical inflation is also discussed in the literature on nominal income targeting. For informal such discussion beyond the works already cited, see Sumner (2012), Beckworth (2019), Binder (2020), and Hall and Mankiw (1994).

References

Adam, Klaus, and Henning Weber. 2023. "Estimating the optimal inflation target from trends in relative prices." *American Economic Journal: Macroeconomics*.

Alvarez, Fernando, Francesco Lippi, and Luigi Paciello. 2011. "Optimal price setting with observation and menu costs." *The Quarterly Journal of Economics*, 126(4): 1909–1960.

Alvarez, Fernando, Martin Beraja, Martin Gonzalez-Rozada, and Pablo Andrés Neumeyer. 2019. "From hyperinflation to stable prices: Argentina's evidence on menu cost models." *The Quarterly Journal of Economics*, 134(1): 451–505.

Anderson, Eric T, and Duncan I Simester. 2010. "Price stickiness and customer antagonism." *The Quarterly Journal of Economics*, 125(2): 729–765.

Angeletos, George-Marios, and Jennifer La'O. 2020. "Optimal monetary policy with informational frictions." *Journal of Political Economy*, 128(3): 1027–1064.

Aoki, Kosuke. 2001. "Optimal monetary policy responses to relative-price changes." *Journal of Monetary Economics*, 48(1): 55–80.

Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub. 2021. "Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models." *Econometrica*, 89(5): 2375–2408.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub. 2018. "The Intertemporal Keynesian Cross." *National Bureau of Economic Research*, , (August): 1–58.

Auclert, Adrien, Rodolfo Rigato, Ludwig Straub, and Matthew Rognlie. 2023. "New Pricing Models, Same Old Phillips Curve?" *The Quarterly Journal of Economics*, 138(4): qjad041.

Audoly, Richard, Martin Almuzara, Richard K Crump, Davide Melcangi, and Roshie Xing. 2023. "How Large Are Inflation Revisions? The Difficulty of Monitoring Prices in Real Time." *The Difficulty of Monitoring Prices in Real Time (September 7, 2023). Liberty Street Economics*.

Ball, Laurence. 1987. "Externalities from Contract Length." *The American Economic Review*, 77(4): 615–629.

Ball, Laurence, and David Romer. 1989*a*. "Are Prices too Sticky?" *Quarterly Journal of Economics*, 104(3): 507–524.

Ball, Laurence, and David Romer. 1989*b*. "The equilibrium and optimal timing of price changes." *Review of Economic Studies*, 56(2): 179–198.

Ball, Laurence, and David Romer. 1991. "Sticky Prices as Coordination Failure." *The American Economic Review*, 81(3).

Barro, Robert J. 1972. "A theory of monopolistic price adjustment." *Review of Economic Studies*, 39(1): 17–26.

Basu, Susanto. 1995. "Intermediate Goods and Business Cycles: Implications for Productivity and Welfare." *American Economic Review*, 85: 512.

Beckworth, David. 2019. "Facts, Fears, and Functionality of NGDP Level Targeting: A Guide to a Popular Framework for Monetary Policy."

Benigno, Pierpaolo. 2004. "Optimal monetary policy in a currency area." *Journal of International Economics*, 63(2): 293–320.

Binder, Carola. 2020. "NGDP targeting and the public." *Cato Journal*, 40(2): 321–342.

Blanco, Andres. 2021. "Optimal inflation target in an economy with menu costs and a zero lower bound." *American Economic Journal: Macroeconomics*, 13(3): 108–41.

Blanco, Andres, Corina Boar, Callum Jones, and Virgiliu Midrigan. 2022. "Nonlinear Inflation Dynamics in Menu Cost Economies." *Manuscript, December*.

Brainard, Lael. 2022. "What Can We Learn from the Pandemic and the War about Supply Shocks, Inflation, and Monetary Policy?: At the 21st BIS Annual Conference Central Banking after the Pandemic: Challenges Ahead, Bank for International Settlements, Basel, Switzerland November 28th 2022."

Bullard, James, and Riccardo DiCecio. 2019. "Optimal monetary policy for the masses." *FRB St. Louis Working Paper*, , (2019-9).

Bullard, James, Riccardo DiCecio, Aarti Singh, and Jacek Suda. 2023. "Optimal Macroeconomic Policies in a Heterogeneous World."

Caballero, Ricardo J, and Eduardo MRA Engel. 2007. "Price stickiness in Ss models: New interpretations of old results." *Journal of Monetary Economics*, 54: 100–121.

Calabresi, Guido. 1970. *The Cost of Accidents.*

Caplin, Andrew, and John Leahy. 1994. "Business as usual, market crashes, and wisdom after the fact." *The American economic review*, 548–565.

Caplin, Andrew, and John Leahy. 2010. "Economic Theory and the World of Practice: A Celebration of the (S, s) Model." *Journal of Economic Perspectives*, 24(1): 183–202.

Caplin, Andrew S, and Daniel F Spulber. 1987. "Menu costs and the neutrality of money." *The Quarterly Journal of Economics*, 102(4): 703–725.

Carvalho, Carlos, and Oleksiy Kryvtsov. 2021. "Price selection." *Journal of Monetary Economics*, 122: 56–75.

Cavallo, Alberto. 2018. "Scraped data and sticky prices." *Review of Economics and Statistics*, 100(1): 105–119.

Cavallo, Alberto, and Roberto Rigobon. 2016. "The billion prices project: Using online prices for measurement and research." *Journal of Economic Perspectives*, 30(2): 151–178.

Dutta, Shantanu, Mark Bergen, Daniel Levy, and Robert Venable. 1999. "Menu costs, posted prices, and multiproduct retailers." *Journal of Money, Credit, and Banking*, 683–703.

Eggertsson, Gauti, and Michael Woodford. 2003. "The Zero Bound On Interest Rates and Optimal Monetary Policy." *Brookings Papers on Economic Activity*, 34.

Erceg, Christopher J., and Andrew T. Levin. 2006. "Optimal Monetary Policy with Durable Consumption Goods." *Journal of Monetary Economics*, 53(3): 631–653.

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

Eyster, Erik, Kristóf Madarász, and Pascal Michaillat. 2021. "Pricing under Fairness Concerns." *Journal of the European Economic Association*, 19(3): 1853–1898.

Friedman, Milton. 1968. "The Role of Monetary Policy." *The American Economic Review*, 58(1): 1–17.

Gali, Jordi. 2008. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications.*

Gautier, Erwan, and Hervé Le Bihan. 2022. "Shocks versus Menu Costs: Patterns of Price Rigidity in an Estimated Multisector Menu-Cost Model." *Review of Economics and Statistics*, 104(4): 668–685.

Gautier, Erwan, Cristina Conflitti, Riemer P Faber, Brian Fabo, Ludmila Fadejeva, Valentin Jouvanceau, Jan-Oliver Menz, Teresa Messner, Pavlos Petroulas, Pau Roldan-Blanco, et al. 2022. "New facts on consumer price rigidity in the euro area."

Gertler, Mark, and John Leahy. 2008. "A Phillips curve with an Ss foundation." *Journal of Political Economy*, 116(3): 533–572.

Golosov, Mikhail, and Robert E Lucas. 2007. "Menu costs and Phillips curves." *Journal of Political Economy*, 115(2): 171–199.

Gorodnichenko, Yuriy. 2008. "Endogenous information, menu costs and inflation persistence." National Bureau of Economic Research.

Gorodnichenko, Yuriy, and Matthew D Shapiro. 2007. "Monetary policy when potential output is uncertain: Understanding the growth gamble of the 1990s." *Journal of Monetary Economics*, 54(4): 1132–1162.

Greenwood, Jeremy, Zvi Hercowitz, and Gregory W Huffman. 1988. "Investment, Capacity Utilization, and the Real Business Cycle." *The American Economic Review*, 78(3): 402–417.

Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, and Ivan Werning. 2021. "Monetary Policy in Times of Structural Reallocation." *Proceedings of the 2021 Jackson Hole Symposium*.

Hall, Robert E, and N Gregory Mankiw. 1994. "Nominal income targeting." In *Monetary policy*. 71–94. The University of Chicago Press.

Huang, Kevin XD, and Zheng Liu. 2005. "Inflation targeting: What inflation rate to target?" *Journal of Monetary Economics*, 52(3).

Karadi, Peter, Raphael Schoenle, and Jesse Wursten. 2022. "Price selection in the microdata." Tech. rep., Mimeo. Oct. 10.

Klenow, Peter J, and Oleksiy Kryvtsov. 2008. "State-dependent or time-dependent pricing: Does it matter for recent US inflation?" *The Quarterly Journal of Economics*, 123(3): 863–904.

Koenig, Evan F. 2013. "Like a good neighbor: Monetary policy, financial stability, and the distribution of risk." *International Journal of Central Banking*.

Kreamer, Jonathan. 2022. "Sectoral Heterogeneity and Monetary Policy." *American Economic Journal: Macroeconomics*, 14(2): 123–159.

Levy, Daniel, Mark Bergen, Shantanu Dutta, and Robert Venable. 1997. "The magnitude of menu costs: direct evidence from large US supermarket chains." *The Quarterly Journal of Economics*, 112(3): 791–824.

Maćkowiak, Bartosz, Filip Matějka, and Mirko Wiederholt. 2023. "Rational inattention: A review." *Journal of Economic Literature*, 61(1): 226–273.

Mankiw, N. Gregory. 1985. "Small Menu Costs and Large Business Cycles: A Macroeconomic Model of Monopoly." *Quarterly Journal of Economics*, 100(2): 529–537.

Mankiw, N. Gregory, and Ricardo Reis. 2003. "What measure of inflation should a central bank target?" *Journal of Monetary Economics*.

Midrigan, Virgiliu. 2011. "Menu costs, multiproduct firms, and aggregate fluctuations." *Econometrica*, 79(4): 1139–1180.

Nakamura, Emi, and Jon Steinsson. 2010. "Monetary non-neutrality in a multisector menu cost model." *The Quarterly Journal of Economics*, 125(3): 961–1013.

Nakamura, Emi, Jón Steinsson, Patrick Sun, and Daniel Villar. 2018. "The elusive costs of inflation: Price dispersion during the US great inflation." *The Quarterly Journal of Economics*, 133(4): 1933–1980.

Nakov, Anton, and Carlos Thomas. 2014. "Optimal monetary policy with state-dependent pricing." *International Journal of Central Banking*, 10(3): 49–94.

Nisticò, Salvatore. 2007. "The welfare loss from unstable inflation." *Economics Letters*, 96(1): 51–57.

Orphanides, Athanasios. 2003. "Historical monetary policy analysis and the Taylor rule." *Journal of Monetary Economics*, 50(5): 983–1022.

Poole, William. 1970. "Optimal choice of monetary policy instruments in a simple stochastic macro model." *Quarterly Journal of Economics*.

Powell, Jerome. 2023. "Opening Remarks, at "Monetary Policy Challenges in a Global Economy"."

Rotemberg, Julio J. 1982. "Sticky prices in the United States." *Journal of Political Economy*, 90(6): 1187–1211.

Rubbo, Elisa. 2023. "Networks, Phillips curves, and monetary policy." *Econometrica*, 91(4): 1417–1455.

Schnabel, Isabel. 2022. "Looking through Higher Energy Prices? Monetary Policy and the Green Transition."

Selgin, George. 1997. *Less Than Zero: The Case for a Falling Price Level in a Growing Economy.* Institute for Economic Affairs.

Sheedy, Kevin D. 2014. "Debt and incomplete financial markets: A case for nominal GDP targeting." *Brookings Papers on Economic Activity*, 2014(1): 301–373.

Sheshinki, Eytan, and Yoram Weiss. 1977. "Inflation and Costs of Price Adjustment." *The Review of Economic Studies*, 44(2): 287–303.

Sumner, Scott. 2012. "The Case for Nominal GDP Targeting." *Mercatus Center working paper*.

Werning, Ivan. 2011. "Managing a liquidity trap: Monetary and fiscal policy." National Bureau of Economic Research.

Werning, Ivan. 2014. "Comment: Debt and Incomplete Financial Markets." *Brookings Papers on Economic Activity*, 368–372.

Wolman, Alexander L. 2011. "The optimal rate of inflation with trending relative prices." *Journal of Money, Credit and Banking*, 43(2-3): 355–384.

Woodford, Michael. 1998. "Doing without Money: Controlling Inflation in a Post-Monetary World." *Review of Economic Dynamics*, 1(1): 173–219.

Woodford, Michael. 2003. "Interest and prices: Foundations of a theory of monetary policy."

Woodford, Michael. 2012. "Methods of policy accommodation at the interest-rate lower bound."

Yun, Tack. 1996. "Nominal price rigidity, money supply endogeneity, and business cycles." *Journal of Monetary Economics*, 37(2): 345–370.

Zbaracki, Mark J, Mark Ritson, Daniel Levy, Shantanu Dutta, and Mark Bergen. 2004. "Managerial and customer costs of price adjustment: direct evidence from industrial markets." *Review of Economics and Statistics*, 86(2): 514–533.

A Additional proofs

A.1 Proof of lemma 1

Proof. Equation (12) showed that a firm j in sector i with inherited price p_i^{old} adjusts if and only if:

$$\left(\frac{W}{A_i}\right)^{1-\eta} p_i^\eta y_i \left[\frac{1}{\eta}\right] - W\psi > \left(p_i^{\text{old}}\right)^{1-\eta} p_i^\eta y_i \left[1 - \frac{W/A_i}{p_i^{\text{old}}} \cdot \frac{\eta-1}{\eta}\right]$$

Define:

$$f(W, A_i) \equiv \left(\frac{W}{A_i}\right)^{1-\eta} p_i^\eta y_i \left[\frac{1}{\eta}\right] - W\psi - \left(p_i^{\text{old}}\right)^{1-\eta} p_i^\eta y_i \left[1 - \frac{W/A_i}{p_i^{\text{old}}} \cdot \frac{\eta-1}{\eta}\right] \tag{52}$$

The firm will adjust iff $f(W, A_i) \geq 0$.

Observe first that for the locus of (W, A_i) such that $W/A_i = p_i^{\text{old}}$, it is the case that $f(W, A_i) = -W\psi < 0$ and the firm will not adjust. That is, this locus is a subset of the inaction region $\Lambda \equiv \{(W, A_i) | f(W, A_i) < 0\}$. Thus Λ is nonempty.

In A_i space. Observe that

$$\frac{\partial f}{\partial A_i} = p_i^\eta y_i W(A_i^{-2}) \left(\frac{\eta-1}{\eta}\right) \left[\left(\frac{W}{A_i}\right)^{-\eta} - (p_i^{\text{old}})^{-\eta}\right]$$

This is positive iff $W/A_i < p_i^{\text{old}}$. Additionally, $\lim_{A_i \to 0} f(\cdot, A_i) = \lim_{A_i \to \infty} f(\cdot, A_i) = \infty$ and $f(\cdot, A_i)$ is continuous in A_i on $(0, \infty)$.

Now consider any fixed W^0 such that there exists some A_i^0 with $f(W^0, A_i^0) < 0$. Then by the intermediate value theorem there exists an inaction interval $(\underline{\lambda}, \overline{\lambda})$ around W^0 such that $f(W^0, \underline{\lambda}) = f(W^0, \overline{\lambda}) = 0$, and $f(W^0, A_i) < 0$ iff $A_i \in (\underline{\lambda}, \overline{\lambda})$. To see that $\overline{\lambda}$ is increasing in ψ and $\underline{\lambda}$ is decreasing in ψ, observe that increasing ψ shifts the entire $f(x)$ curve down, i.e. $\frac{\partial f}{\partial \psi} < 0$.

If for a fixed W^0 there is no A_i with $f(W, A_i) < 0$, then by construction there is no point in Λ along W^0.

In W space. Similarly, observe that

$$\frac{\partial f}{\partial W} = p_i^\eta y_i A_i^{-1} \left(\frac{\eta-1}{\eta}\right) \left[(p_i^{\text{old}})^{-\eta} - \left(\frac{W}{A_i}\right)^{-\eta}\right] - \psi$$

This is zeroed for the locus of (W, A_i) such that

$$\left(\frac{W}{A_i}\right)^{-\eta} = (p_i^{\text{old}})^{-\eta} - \psi p_i^{-\eta} y_i^{-1} A_i \left(\frac{\eta - 1}{\eta}\right)^{-1} \equiv \zeta^{-\eta}$$

Observe that $f_1 < 0$ iff $W/A_i < \zeta$. Additionally, $\lim_{W \to 0} f(W, \cdot) = \infty$.[31] Additionally, $f(W, \cdot)$ is continuous in W on $(0, \infty)$. Thus, as above, consider any fixed A_i^0 such that there exists some W^0 with $f(W^0, A_i^0) < 0$. By the intermediate value theorem there exists (abusing notation) an inaction interval $(\underline{\lambda}, \overline{\lambda})$ around A_i^0 such that $f(\underline{\lambda}, A_i^0) = f(\overline{\lambda}, A_i^0) = 0$, and $f(W, A_i^0) < 0$ iff $W \in (\underline{\lambda}, \overline{\lambda})$, where $\overline{\lambda}$ is potentially infinite. To see that $\overline{\lambda}$ is increasing in ψ and $\underline{\lambda}$ is decreasing in ψ, observe that increasing ψ shifts the entire $f(x)$ curve down, i.e. $\frac{\partial f}{\partial \psi} < 0$. $\square$

[31] The second limit comes from using L'Hopital's rule, together with the natural parameter restriction that $\psi < \frac{1}{S\eta}$. Without this maximum bound on ψ, firms would *always* earn negative profits after adjusting – i.e., firms would *never* adjust.

A.2 Formal statement of planner's problem

Recall we derived that equilibrium welfare in each of the four regimes as a (potentially constant) function of the social planner's choice of the money supply:

$$\mathbb{W}_{\text{all adjust}} = \ln\left(\frac{\gamma^{1/S}}{S}\right) - [1 + S\psi]$$
$$\mathbb{W}_{\text{only 1 adjusts}}(M) = \ln\left(\frac{\gamma^{\frac{1}{S}}}{S} M^{\frac{S-1}{S}}\right) - \left[\frac{1}{S} + (S-1)\frac{M}{S} + \psi\right]$$
$$\mathbb{W}_{\text{only } k \text{ adjust}}(M) = \ln\left(\frac{1}{S} M^{1/S}\right) - \left[\frac{S-1}{S} + \frac{1}{S}\frac{M}{S} + \frac{S-1}{S}\psi\right]$$
$$\mathbb{W}_{\text{none adjust}}(M) = \ln\left(\frac{M}{S}\right) - \left[\frac{1}{\gamma}\frac{M}{S} + (S-1)\frac{M}{S}\right]$$

Constrained planner's problem. We define the constrained planner as the planner who chooses M in each regime to maximize welfare, *constrained* by the fact that the choice of M must be incentive compatible with whether various sectors actually adjust or not:

$$M^{*,\text{ constrained}}_{\text{only 1 adjusts}} \equiv \arg\max_M \mathbb{W}_{\text{only 1 adjusts}}(M) \quad \text{s.t. } f(M,\gamma) \geq 0 \text{ and } f(M,1) \leq 0 \tag{53}$$

$$M^{*,\text{ constrained}}_{\text{only } k \text{ adjust}} \equiv \arg\max_M \mathbb{W}_{\text{only } k \text{ adjust}}(M) \quad \text{s.t. } f(M,\gamma) \leq 0 \text{ and } f(M,1) \geq 0 \tag{54}$$

$$M^{*,\text{ constrained}}_{\text{none adjust}} \equiv \arg\max_M \mathbb{W}_{\text{none adjust}}(M) \quad \text{s.t. } f(M,\gamma) \leq 0 \text{ and } f(M,1) \leq 0 \tag{55}$$

where $f(M, A_i)$ refers to the function defined in (29) which is positive if and only if firms in sector i want to adjust. This defines, for example, $M^{*,\text{ constrained}}_{\text{only 1 adjusts}}$ as the level of money supply which maximizes welfare in equilibrium when only sector 1 adjusts $\mathbb{W}_{\text{only 1 adjusts}}(M)$, *subject to the constraint that* it is indeed incentive-compatible for sector-1 firms to adjust, $f(M,\gamma) \geq 0$, and incentive-compatible for firms in sectors k to not adjust, $f(M,1) \leq 0$. Denote the associated constrained-optimal levels of welfare in each regime as:

$$\mathbb{W}^{*,\text{ constrained}}_{\text{only 1 adjusts}} = \mathbb{W}_{\text{only 1 adjusts}}\left(M^{*,\text{ constrained}}_{\text{only 1 adjusts}}\right)$$

$$\mathbb{W}^{*,\text{ constrained}}_{\text{only } k \text{ adjust}} = \mathbb{W}_{\text{only } k \text{ adjust}}\left(M^{*,\text{ constrained}}_{\text{only } k \text{ adjust}}\right)$$
$$\mathbb{W}^{*,\text{ constrained}}_{\text{none adjust}} = \mathbb{W}_{\text{none adjust}}\left(M^{*,\text{ constrained}}_{\text{none adjust}}\right)$$

The constrained social planner's problem is then to select among these, or to implement the regime where all adjust (in which case the choice of M is irrelevant, as long as it is incentive-compatible):

$$\max\left\{\mathbb{W}^{*,\text{ constrained}}_{\text{only 1 adjusts}}, \mathbb{W}^{*,\text{ constrained}}_{\text{only } k \text{ adjust}}, \mathbb{W}^{*,\text{ constrained}}_{\text{none adjust}}, \mathbb{W}_{\text{all adjust}}\right\} \tag{56}$$

which she implements with the associated incentive-compatible choice for the money supply.

Unconstrained social planner's problem. In the body of the paper and in this subsection, we endow the planner with the instrument of subsidizing menu costs, so that the constraints on (53), (54), and (55) never bind. Because taxation is lump sum and wholly non-distortionary, subsidies to offset the menu cost are equivalent to endowing the planner with the power to change prices directly (but, if doing so, still incurring a menu cost for affected firms). The unconstrained social planner's problem is thus the same as the constrained planner's problem, but without any of the implementability constraints.

$$M^{*}_{\text{only 1 adjusts}} \equiv \arg\max_{M} \mathbb{W}_{\text{only 1 adjusts}}(M) \tag{57}$$

$$M^{*}_{\text{only } k \text{ adjust}} \equiv \arg\max_{M} \mathbb{W}_{\text{only } k \text{ adjust}}(M) \tag{58}$$

$$M^{*}_{\text{none adjust}} \equiv \arg\max_{M} \mathbb{W}_{\text{none adjust}}(M) \tag{59}$$

Since the objective functions in all of these arg maxes are strictly concave, the solution is found from the first order condition, as presented in the text. We denoted the associated unconstrained-optimal levels of welfare in each regime in equations (18), (19), (20) as:

$$\mathbb{W}^{*}_{\text{only 1 adjusts}} = \mathbb{W}_{\text{only 1 adjusts}}\left(M^{*}_{\text{only 1 adjusts}}\right)$$
$$\mathbb{W}^{*}_{\text{only } k \text{ adjust}} = \mathbb{W}_{\text{only } k \text{ adjust}}\left(M^{*}_{\text{only } k \text{ adjust}}\right)$$
$$\mathbb{W}^{*}_{\text{none adjust}} = \mathbb{W}_{\text{none adjust}}\left(M^{*}_{\text{none adjust}}\right)$$

The constrained social planner's problem is then to select among these, or to implement the regime where all adjust (in which case the choice of M is irrelevant):

$$\max\left\{\mathbb{W}^*_{\text{only 1 adjusts}}, \mathbb{W}^*_{\text{only } k \text{ adjust}}, \mathbb{W}^*_{\text{none adjust}}, \mathbb{W}_{\text{all adjust}}\right\} \tag{60}$$

It is this maximization problem that produces lemma **??** and lemma 3, which in turn produce proposition 1. □

B Adjustment externalities

In this section, we work with a slightly modified version of the baseline model, where menu costs are modeled as a utility penalty rather than as a labor cost. This highlights the way in which the results do not depend on how menu costs are modeled, and also facilitates an analysis of the role of adjustment externalities.

B.1 Model setup

The final goods producer and sectoral goods producer are exactly the same as the baseline model.

Intermediate goods producers. The intermediate goods producers again are a unit mass of monopolistically competitive firms in each sector with linear technology and productivity that is common to the sector. They again face a menu cost if adjusting prices. Here, unlike the baseline model, the adjustment cost is not ψ units of extra labor, but a penalty $(1-\psi)$ that scales down the firm manager's objective function (but not profits). Firm i in sector j faces the following maximization program:

$$\begin{aligned}
&\max_{p_i(j)} D_i(j)\left(1-\psi\chi_i(j)\right)\\
&\text{s.t. } D_i(j) = p_i(j)y_i(j) - Wn_i(j)(1-\tau)\\
&\quad \chi_i(j) = \begin{Bmatrix} 1 & \text{if } p_i(j) \neq p_i^{\text{old}} \\ 0 & \text{else} \end{Bmatrix}\\
&\quad y_i(j) = y_i\left(\frac{p_i(j)}{p_i}\right)^{-\eta}\\
&\quad y_i(j) = A_i n_i(j)
\end{aligned}$$

This menu cost is a utility penalty that is passed on to households, but does not affect physical profits. As before, all firms within a sector face the same problem, and we drop the (j) notation when the context is clear.

Households. The representative household is precisely as in the baseline model, except that the utility function (1) is modified to be:

$$\mathbb{W} = \ln C - N + \ln\left(\frac{M}{P}\right) - \psi\sum_i \chi_i$$

The household has the same preferences over consumption, labor, and real balances; but now is directly penalized in terms of welfare when firms adjust prices. The benefit of this modeling technique is that it turns off the income effects caused by menu costs, as discussed in section 5.1, and is a technique that has been used by e.g. Auclert, Rognlie and Straub (2018) or Guerrieri et al. (2021).

B.2 Shock and equilibrium

We run the same exercise, shocking the productivity of sector 1 from $A_1 = 1$ to $A_1 = \gamma > 1$. The equilibrium allocations in the four regimes 3.1 is *exactly* the same as in the body of the paper, except that the level of aggregate labor in each of the four regimes is no longer affected by menu costs. The equilibrium level of welfare in each of the four regimes as a function of the choice of money supply is:

$$\begin{aligned}
\mathbb{W}_{\text{flex}} &= \ln\left(\frac{\gamma^{1/S}}{S}\right) - 1 \\
\mathbb{W}_{\text{all adjust}} &= \ln\left(\frac{\gamma^{1/S}}{S}\right) - 1 \\
\mathbb{W}_{\text{only 1 adjusts}}(M) &= \ln\left(\frac{\gamma^{1/S}}{S} M^{\frac{S-1}{S}}\right) - \frac{1}{S}\left[1 + M(S-1)\right] - \psi \\
\mathbb{W}_{\text{only } k \text{ adjust}}(M) &= \ln\left(\frac{1}{S} M^{\frac{1}{S}}\right) - \frac{1}{S}\left[S - 1 + \frac{M}{\gamma}\right] - (S-1)\psi \\
\mathbb{W}_{\text{none adjust}}(M) &= \ln\left(\frac{M}{S}\right) - \frac{M}{S}\left[S - 1 + \frac{1}{\gamma}\right]
\end{aligned}$$

B.3 Adjustment decision

The firm compares its objective function under price adjustment versus under the inherited price. The adjustment condition can be simplified to be written as: adjust if and only if

$$\frac{1}{\eta}(1-\psi) > \left[\frac{W/A_i}{p_i^{\text{old}}}\right]^{\eta} \left(\left[\frac{W/A_i}{p_i^{\text{old}}}\right]^{-1} - \frac{\eta-1}{\eta}\right)$$

For additional analytical tractability, we make the following assumption in this section:

Assumption 1. The elasticity of substitution is $\eta = 2$.

This assumption allows for a closed form solution to the inaction region, using the quadratic formula: do not adjust if and only if

$$\frac{W}{A_i} \in \left(p_i^{\text{old}}(1-\sqrt{\psi}), p_i^{\text{old}}(1+\sqrt{\psi})\right) \tag{61}$$

Clearly this has the same properties as the Λ inaction region described in lemma 1.

When starting from the steady state where p_i^{old} for all sectors i, and using the equilibrium $W = M$ condition, then we have the following. The inaction region for sector 1 is

$$M \in (\gamma(1-\sqrt{\psi}), \gamma(1+\sqrt{\psi}))$$

The inaction region for sectors k is

$$M \in (1-\sqrt{\psi}, 1+\sqrt{\psi})$$

B.4 The planner's problem

The planner's problem – importantly, without the ability to subsidize menu costs and so denoted "constrained" – written in full is:

$$\max_{\substack{\text{all adjust, only 1 adjusts}\\ \text{only } k \text{ adjust, none adjust}}} \left\{\mathbb{W}_{\text{all adjust}}, \mathbb{W}^{*,\text{constrained}}_{\text{only 1 adjusts}}, \mathbb{W}^{*,\text{constrained}}_{\text{only } k \text{ adjust}}, \mathbb{W}^{*,\text{constrained}}_{\text{none adjust}}\right\} \tag{62}$$

s.t.

$$\mathbb{W}_{\text{all adjust}} = \left\{\ln\left(\frac{\gamma^{1/S}}{S}\right) - 1\right\}$$

$$\mathbb{W}^{*,\text{constrained}}_{\text{only 1 adjusts}} = \left\{\begin{matrix} \max\limits_{M} \ \ln\left(\frac{\gamma^{1/S}}{S} M^{\frac{S-1}{S}}\right) - \frac{1}{S}[1+M(S-1)] - \psi \\ \text{s.t. } M \in (1-\sqrt{\psi}, 1+\sqrt{\psi}) \\ M \notin (\gamma(1-\sqrt{\psi}), \gamma(1+\sqrt{\psi})) \end{matrix}\right\}$$

$$\mathbb{W}^{*,\text{constrained}}_{\text{only } k \text{ adjust}} = \left\{\begin{matrix} \max\limits_{M} \ \ln\left(\frac{1}{S} M^{\frac{1}{S}}\right) - \frac{1}{S}\left[S-1+\frac{M}{\gamma}\right] - (S-1)\psi \\ \text{s.t. } M \notin (1-\sqrt{\psi}, 1+\sqrt{\psi}) \\ M \in (\gamma(1-\sqrt{\psi}), \gamma(1+\sqrt{\psi})) \end{matrix}\right\}$$

$$\mathbb{W}^{*,\text{constrained}}_{\text{none adjust}} = \left\{\begin{matrix} \max\limits_{M} \ \ln\left(\frac{M}{S}\right) - \frac{M}{S}\left[S-1+\frac{1}{\gamma}\right] \\ \text{s.t. } M \in (1-\sqrt{\psi}, 1+\sqrt{\psi}) \\ M \in (\gamma(1-\sqrt{\psi}), \gamma(1+\sqrt{\psi})) \end{matrix}\right\}$$

B.5 Interior optima

The interior optima for each regime, found from the first order conditions, are the same as the baseline model:

$$
\begin{aligned}
M^{\text{interior}}_{\text{only 1 adjusts}} &= 1 \\
M^{\text{interior}}_{\text{only } k \text{ adjust}} &= \gamma \\
M^{\text{interior}}_{\text{none adjust}} &= \frac{S}{S - 1 + 1/\gamma}
\end{aligned}
$$

B.6 Only sector 1 adjusts: The possibility of *positive* adjustment externalities

Suppose the unconstrained social planner – i.e., one who could subsidize menu costs and ignore the implementability constraints – would want to implement the regime where only sector 1 adjusts, and she therefore wants to set $M = 1$. We now examine whether this is incentive compatible: does it result in sector-k firms being within their inaction region and sector-1 firms being outside it?

First observe that $M = 1$ indeed ensures that sector-k firms are in their inaction region, since $M = 1 \in (1 - \sqrt{\psi}, 1 + \sqrt{\psi})$ always.

However, it is possible that $M = 1$ could leave sector-1 firms inside their inaction region, if the following condition holds:

$$
\gamma < \frac{1}{1 - \sqrt{\psi}} \equiv \gamma_1 \tag{63}
$$

As an existence proof, it is possible to come up with numerical examples for parameters satisfying the above where it would be, in fact, socially optimal to implement this regime if there were no implementability constraints. When this is the case, then the best the central bank can do within this regime is to set $M = \gamma(1 - \sqrt{\psi})$. This is a case of *positive* adjustment externalities: the social planner would prefer that sector 1 adjusts its prices, even though it is individually rational to not do so.

B.7 No sectors adjust: The possibility of *negative* adjustment externalities

Now suppose the unconstrained social planner would prefer that no sector adjusts (i.e. $\gamma < \overline{\gamma}$). The interior optimum level of the money supply, as previously noted, would be $M^{\text{interior}}_{\text{none adjust}} = \frac{S}{S-1+1/\gamma}$. Is this incentive-compatible?

To be incentive-compatible requires that both $\frac{S}{S-1+1/\gamma} > \gamma(1-\sqrt{\psi})$ and $\frac{S}{S-1+1/\gamma} < 1+\sqrt{\psi}$. Thus, there is a *negative* adjustment externality – where its privately optimal for firms in a sector to adjust even when its not socially optimal to do so – if either:

$$\gamma < \frac{1+\frac{1}{S-1}\sqrt{\psi}}{1-\sqrt{\psi}} \equiv \gamma_2 \tag{64}$$

or

$$\gamma > \frac{1+\sqrt{\psi}}{S-(S-1)(1+\sqrt{\psi})} \equiv \gamma_3 \tag{65}$$

As an existence proof, it is possible to come up with numerical examples for parameters satisfying either of the above where it would be, in fact, socially optimal to implement this regime if there were no implementability constraints. When this is the case, then the best the central bank can do within this regime is to set M at the respective boundary.

B.8 Summarizing the possibilities for welfare

A similar analysis the above can be done for the case when only sectors k adjust, where a constraint will bind if $\gamma > \gamma_4 \equiv 1+\sqrt{\psi}$. We summarize the results from above and this additional case in the following:

$$\mathbb{W}_{\text{all adjust}} = \mathbb{W}_{\text{flex}} - S\psi$$

$$\mathbb{W}^{*,\text{constrained}}_{\text{only 1 adjusts}} = \begin{Bmatrix} \mathbb{W}_{\text{flex}} - \psi & \text{if } \gamma \geq \gamma_1 \\ \ln\left(\frac{\gamma^{1/S}}{S}\gamma_1^{\frac{S-1}{S}}\right) - \frac{1}{S}\left[1+\gamma_1(S-1)\right] - \psi & \text{else} \end{Bmatrix}$$

$$\mathbb{W}^{*,\text{constrained}}_{\text{only } k \text{ adjust}} = \begin{Bmatrix} \mathbb{W}_{\text{flex}} - (S-1)\psi & \text{if } \gamma \leq \gamma_4 \\ \ln\left(\frac{1}{S}\gamma_4^{1/S}\right) - \frac{1}{S}\left[S-1+\frac{\gamma_4}{\gamma}\right] - (S-1)\psi & \text{else} \end{Bmatrix}$$

$$\mathbb{W}^{*,\text{constrained}}_{\text{none adjust}} = \begin{Bmatrix} -\log\left[S-1+1/\gamma\right]-1 & \text{if } \gamma \in [\gamma_3, \gamma_2] \\ \ln\left(\frac{\gamma_2}{S}\right) - \frac{\gamma_2}{S}\left[S-1+\frac{1}{\gamma}\right] & \text{if } \gamma > \gamma_2 \\ \ln\left(\frac{\gamma_3}{S}\right) - \frac{\gamma_3}{S}\left[S-1+\frac{1}{\gamma}\right] & \text{if } \gamma < \gamma_3 \end{Bmatrix}$$

Optimal monetary policy considers which of these achieves the highest welfare, and sets the money supply M to implement.

Chapter 2

Reexamining optimal policy in the New Keynesian "liquidity trap"

1 Introduction

This paper reevaluates optimal monetary and fiscal policy at the zero lower bound (ZLB) in the workhorse representative agent New Keynesian (RANK) model. I write down a simple and transparent two-period version of the RANK model as a baseline, which I extend minimally as the paper develops, to highlight a set of five underappreciated or novel implications of the model for optimal monetary and fiscal policy. The unifying theme is that monetary policy is less constrained at the ZLB than is commonly thought, whereas fiscal policy in turn is more constrained than typically thought.

I begin with the canonical infinite-horizon New Keynesian model, as described in the Woodford (2003) or Gali (2008) textbooks. I transform the model into a simple two-period model by assuming that after the first two periods, prices are completely flexible. This technique, as in Krugman (1998), creates a "simplest possible version of the model" that is usefully transparent for presenting the conceptual points in the rest of the paper.

This setup is useful for highlighting a baseline result – well-known in the literature since the seminal work of Krugman (1998), if often misunderstood in popular discourse – that at the ZLB, monetary policy is not "pushing on a string". Indeed, as long as the ZLB is only temporarily binding, a credible central bank can achieve any desired level of inflation under standard assumptions. The typical concern, however, is that doing so requires the central bank to commit to expansionary monetary policy *after* the ZLB has ceased to bind, and that this commitment to future expansionary policy is not a time-consistent promise (Eggertsson and Woodford 2003). In the memorable language of Krugman (1998), it is difficult for central banks to "credibly promise to be irresponsible." In this view, the ZLB is less of a "liquidity trap" than it is an "*expectations* trap".

I refer to the model just described as the "baseline model" for the rest of the paper, and I offer a sequence of extensions to it to illustrate my five main points.

Repeated ZLB episodes. The first proposition of the paper, also shown in the important work of Nakata (2018), is that if ZLB episodes are a repeated game – rather than a one-time event, as is typically modeled – then the promises of a ZLB-constrained central bank can easily be credible. During any given single ZLB episode, the central bank recognizes that if it deviates from its promised commitment, this will harm its reputation and credibility for *future* promises during *future* ZLB episodes. If ZLB episodes are sufficiently frequent, then the reputational cost will outweigh the gains of reneging on the promise, and the central bank will choose to honor its promise.

The intuition for this result comes directly from the result in basic game theory that actions which are not credible in one-shot games may easily be credible in repeated games. I write down a simple modification of the baseline model to illustrate and make precise this claim, and I show how it is merely the deflationary mirror image of the Barro and Gordon (1982) critique that reputation can help central banks overcome the *inflationary* time consistency problem of Kydland and Prescott (1976). To paraphrase Krugman (1998): it is easy to credibly promise to be 'irresponsible' *if* your reputation is on the line.

Intertemporal vs. intratemporal distortions. Even if central banks can overcome the time consistency problem at the ZLB, they still cannot achieve the first best in the New Keynesian model – which absent the ZLB they could do – because the ZLB constraint means that there are fewer instruments than targets. The second proposition of the paper is to show and make precise the idea that this limitation is not special: the ZLB is qualitatively no different from any other exogenous nominal rigidity.

The real interest rate is the relative price of consumption today versus consumption tomorrow, and the zero lower bound is merely a nominal price floor on that relative price; a nominal price floor (or any nominal rigidity) on *any* relative price between *any* two goods would constrain monetary policy in the same way as the ZLB.

This point is typically obscured by the fact that the baseline RANK model works with a single, representative consumption good. I modify the baseline model to include two consumption goods at any given point in time, and I show that the *inter*temporal nominal rigidity of the ZLB constrains monetary policy in a way that is isomorphic to an *intra*temporal nominal rigidity on the relative price of the two consumption goods. There are countless examples of such heterogeneous intratemporal rigidities in reality. The restrictions imposed by minimum wage laws are one of many examples (Minton and Wheaton 2021).

Austerity is stimulus. Turning to fiscal policy, the third result of the paper is to show that the welfare-relevant effects of fiscal stimulus in the baseline RANK model depend exclusively on the *change* in government spending, rather than the level. That is, government spending in RANK only affects inflation and the output gap through *the expected growth in government spending from today to tomorrow*. As a result, austerity is stimulus: the promise of spending cuts tomorrow is isomorphic to temporarily raising spending today, in terms of the effect on inflation and the output gap. I show this by extending the baseline model to include government purchases of and production of public goods (as in Werning 2012, Woodford 2011, and Eggertsson 2001).

The intuition for this result comes from consumption smoothing: if the government is going to buy fewer goods tomorrow than today, then tomorrow there will be more goods available for private consumption than today. Under flexible prices, households would therefore want to smooth their consumption over time and borrow to consume more today, in order to match their higher expected consumption tomorrow more closely. This demand for borrowing causes the real rate to rise in flexible price models – or in sticky price models, causes the *natural* rate of interest to rise. Given a nominal interest rate fixed at the zero lower bound and an otherwise negative natural rate, the rise in the natural real rate is stimulative, as it reduces the gap between the policy rate and the natural rate.

A second piece of intuition for this result is that the mechanism through which fiscal stimulus operates is perfectly isomorphic (in terms of the effect on inflation and the output gap) to the stimulative effect of a temporary negative productivity shock at the ZLB. It is well-known that temporary negative productivity shocks are expansionary at the ZLB.[1] Both negative productivity shocks and positive fiscal stimulus lower current consumption relative to future consumption, pushing up the natural rate, which is stimulative at the ZLB. In short, government spending in RANK is – in terms of its effects on the output gap and inflation – a negative productivity shock to the economy.

Contractionary stimulus. This analogy to negative productivity shocks points to the fourth result: if positive fiscal stimulus is targeted at certain specific sectors rather than spread evenly throughout the economy, then it can in fact be *contractionary*. For example, if fiscal stimulus

[1] The benefits of negative productivity shocks at the ZLB are discussed, among many other papers, in Eggertsson (2011, 2012). Wieland (2019) as well as Garín, Lester, and Sims (2019) empirically evaluate this prediction unfavorably. Kiley (2016) and Eggertsson and Garga (2019) discuss this issue under the sticky information friction (Mankiw and Reis 2002) versus the Calvo-Yun sticky price friction.

consists solely of increased purchases in the goods sector – rather than being spread evenly across every sector, from goods to agriculture to services and beyond – then this may cause deflation and create an output gap.

The intuition for this comes from the work of Guerrieri, Lorenzoni, Straub, and Werning (2022), who show that while a negative productivity shock in a *one-sector* model always raises the natural rate, it may lower the natural rate in a *multisector* model when the shock doesn't affect all sectors equally. As noted above, positive fiscal stimulus in RANK is isomorphic to a negative productivity shock. Thus, analogously to the result of Guerrieri-Lorenzi-Straub-Werning, while increased government purchases of a single representative good are always expansionary at the ZLB, government purchases which are heterogeneous across multiple sectors may be contractionary by lowering the natural rate. I demonstrate this in a multisector extension to the baseline model.[2]

Time inconsistency of fiscal policy. The emphasis above on fiscal stimulus as inherently *intertemporal* brings us to the fifth and final result of the paper: optimal fiscal policy faces a time consistency problem at the zero lower bound. To my knowledge, this is the first paper to analytically characterize the time consistency of optimal fiscal policy at the ZLB.

The intuition for this time inconsistency is that optimal fiscal policy would like to spread changes in government consumption over time, for reasons of optimal provision of public goods. As a consequence, optimal fiscal stimulus – which, recall, is driven by the *change* in spending – involves *both* an increase in fiscal spending during the ZLB *and* a cut in spending after the ZLB, rather than simply a large increase during the ZLB. Even a fully beneficent fiscal authority has an incentive to deviate from this policy after the ZLB and to not follow through on the promised austerity, because the austerity is distortionary for public goods provision. Of course, fiscal policy may overcome its time consistency problem in the same way as monetary policy – through a reputational mechanism – though plausibly this is more challenging for fiscal authorities who operate under relatively more severe political economy constraints.

Layout. The literature on optimal monetary and fiscal policy in the New Keynesian liquidity trap is vast, and rather than give a literature review I cite papers throughout where they are relevant. The rest of the paper proceeds directly as it was summarized above. In section 2, I introduce the baseline model and review the idea that, in modern models, monetary policy is able to affect inflation even at the ZLB. Sections 3 and 4 advance the two theses on optimal monetary

[2] Aoki (2001), Benigno (2004), Woodford (2003, ch. 6 sec. 4.3), and Rubbo (2023) study optimal monetary policy in multisector New Keynesian models away from the ZLB. My analysis on this issue differs by studying optimal fiscal policy, and (crucially) by allowing for intra-sector and inter-sector elasticities of substitution to differ, as in Guerrieri, Lorenzoni, Straub, and Werning (2022).

policy; sections 5 through 7 advance the three theses on optimal fiscal policy. I conclude with some remarks on the limitations of the representative agent New Keynesian model.

2 Baseline model

In this section, I set up a minimalistic New Keynesian model of the ZLB. It will serve as the baseline model for the rest of the paper, as we extend it in various directions and use it to conduct policy experiments. Readers familiar with Eggertsson and Woodford (2003) or Werning (2012) may wish to look at equations (5)-(12), the discussion in section 2.3, and to skip to section 3.

2.1 Setup

For the baseline model, I start directly from the canonical log-linearized New Keynesian model that results from the basic setup with Calvo pricing and no capital (c.f. Gali 2008 or Bergholt 2012, for example).[3]

The core equations are the New Keynesian Phillips Curve (NKPC),

$$\pi_t = kx_t + \beta E_t \pi_{t+1} \tag{1}$$

and the Euler equation (EE) written in terms of the output gap,

$$x_t = E_t x_{t+1} - \sigma[i_t - E_t \pi_{t+1} - r_t^n] \tag{2}$$

The notation is standard: π_t is inflation; x_t is the output gap; k is a constant that depends on preference and technological parameters; β is the steady state rate of time preference; σ is the inverse of the elasticity of intertemporal substitution; i_t is the gross nominal interest rate.

r_t^n is the gross natural interest rate, which in the baseline framework only depends on the exogenous rate of time preference ρ_t:[4]

$$r_t^n = \rho_t \tag{3}$$

The central bank's per-period welfare loss function is, to the usual second-order approximation:

$$\mathbb{W}_t = \pi_t^2 + \lambda x_t^2 \tag{4}$$

Here, λ is a constant. This is the welfare-theoretically correct loss function for the central bank to minimize in order to maximize the household's utility.

Once we make an assumption about how the central bank uses the objective function (4) to set the path of interest rates $\{i_t\}$, then equations (1) through (4) fully specify the model, in the absence of the zero lower bound.

[3] I assume throughout that the fiscal authority implements the usual production subsidy to offset the distortion from market power, and that it offsets time variation in the wedge between flexible-price output and first-best output so that there are no shocks to the NKPC (i.e. so-called "cost push" shocks).

[4] In particular, this assumes an absence of productivity shocks. Section 5 incorporates productivity shocks.

Note that I choose to write the entire model in terms of the output gap x_t, rather than in terms of the actual level of output itself. While, of course, writing the model the two ways is completely equivalent, it will turn out that writing the model in terms of the output gap is quite valuable for generating sharp intuition for the results. This is because the output gap is directly the *welfare-relevant* object from the loss function – unlike output or consumption – and therefore makes discussion of *optimal*, welfare-maximizing policy especially clear.

I transform the infinite-horizon model to a simplified two-period model by assuming that: from period 2 onwards, all prices are completely flexible – so the output gap is forever zero – and assume that the central bank then chooses to set inflation to zero.[5] As a result, the model collapses to a small system of equations:

1. The NKPC in period 0 and in period 1:

$$\pi_0 = kx_0 + \beta E_0 \pi_1 \tag{5}$$

$$\begin{aligned} \pi_1 &= kx_1 + \beta E_1 \pi_2 \\ &= kx_1 \end{aligned} \tag{6}$$

2. The EE in period 0 and in period 1:

$$x_0 = E_0 x_1 - \sigma[i_0 - E_0 \pi_1 - r_0^n] \tag{7}$$

$$\begin{aligned} x_1 &= E_1 x_2 - \sigma[i_1 - E_1 \pi_2 - r_1^n] \\ &= -\sigma[i_1 - r_1^n] \end{aligned} \tag{8}$$

3. The welfare loss function of the central bank:

$$\mathbb{W} = \pi_0^2 + \lambda x_0^2 + \beta[\pi_1^2 + \lambda x_1^2] \tag{9}$$

We assume that, due to the existence of non-interest bearing cash as an outside option, there is a zero lower bound on nominal interest rates.[6]

$$i_0, i_1 \geq 1 \tag{10}$$

This could easily be generalized to an "effective lower bound" at some exogenous level below zero.

Finally, since we want the ZLB to bite in period 0 but not in period 1, we assume that the natural rate equals its steady state value of $1/\beta$ in period 1; but that in period 0 it is equal to some level $\phi < 1$ so that the ZLB constraint is relevant:

$$r_0^n = \phi \tag{11}$$

$$r_1^n = 1/\beta \tag{12}$$

2.2 Equilibrium

[5] The choice of zero inflation for $t \geq 2$ is merely for convenience.

[6] As Koning (2013) highlights, this ZLB constraint can be viewed as an implication of Gresham's Law.

Substituting in the natural rate terms and working backwards from period 1 to period 0, we can express all four endogenous objects – the output gap and inflation in both periods, $\{x_1, \pi_1, x_0, \pi_0\}$ – in terms of the two exogenous interest rates set by the central bank, $\{i_1, i_0\}$.

I highlight the terms of particular interest. At period 1:

$$x_1 = -\sigma\left[i_1 - \tfrac{1}{\beta}\right] \tag{13}$$

$$\pi_1 = -\sigma k\left[i_1 - \tfrac{1}{\beta}\right] \tag{14}$$

At period 0:

$$x_0 = -\sigma(1 + k\sigma)\left[E_0 i_1 - \tfrac{1}{\beta}\right] - \sigma[i_0 - \phi] \tag{15}$$

$$\pi_0 = -\sigma k[1 + \sigma k + \beta]\left[E_0 i_1 - \tfrac{1}{\beta}\right] - \sigma k[i_0 - \phi] \tag{16}$$

In period 1, if the policy rate i_1 is kept below the natural rate $1/\beta$, then there is inflation ($\pi_1 > 0$) and an economic boom ($x_1 > 0$). In period 0, whether inflation and the output gap are positive or negative depends on the entire path of policy rates $\{i_0, i_1\}$ compared to each period's natural nominal rates $\{\phi, 1/\beta\}$.

Under the welfare loss function (9), the central bank would like to ensure that inflation and output gap are zero in both periods, $\pi_t = x_t = 0$. In the absence of the ZLB, this could easily be done by setting the policy rate in each period equal to the natural rate, $i_0 = \phi$ and $i_1 = 1/\beta$. When $\phi < 1$ so that the ZLB binds in period 0, however, this is not feasible; the policy rate is stuck at $i_0 = 1$.

2.3 Central banks are not "pushing on a string" at the ZLB

In popular discourse, it is common to hear claims of total monetary policy impotency under a binding ZLB constraint: "central banks are pushing on a string" in terms of ability to raise inflation, and similar cliched analogies.

It is worth emphasizing that this is simply not the case in our standard textbook models, New Keynesian or otherwise, if the economy is expected to ever have any probability of leaving the ZLB.[7] This point, first made formally in the seminal Krugman (1998) paper, can be seen through the lens of the model here in equations (15) and (16). Supposing the nominal rate in period 0 is constrained by the ZLB so that $i_0 = 1$, then inflation and the output gap during this "liquidity trap" are:

$$x_0 = -\sigma(1 + k\sigma)\left[E_0 i_1 - \tfrac{1}{\beta}\right] - \sigma[1 - \phi] \tag{17}$$

$$\pi_0 = -\sigma k[1 + \sigma k + \beta]\left[E_0 i_1 - \tfrac{1}{\beta}\right] - \sigma k[1 - \phi] \tag{18}$$

[7] For the case when the economy is never expected to have even epsilon probability of leaving the ZLB, see the literature on secular stagnation (e.g. Eggertsson, Summers, and Mehrotra 2016).

Clearly, the level of inflation during the "liquidity trap", π_0, can be raised by the central bank's choice of *future* policy i_1. Thus, the notion of "forward guidance": if the central bank at time 0 can make credible promises about future policy, i_1, they can manipulate inflation during the liquidity trap, π_0.

Is there any sense, then, in which the ZLB a "problem" for central banks? There are two:

(1) Promises for future policy action may not be believed: forward guidance may be time inconsistent.

(2) The "fewer-instruments-than-targets" problem: although the central bank can affect inflation during the liquidity trap π_0, the ZLB constraint means that it cannot achieve the first best and set inflation and the output gap to zero in both periods.

I now lay out these two issues, before discussing in sections 3 and 4 the results which show caveats to the importance of each.

2.4 The "expectations trap": the time consistency of monetary policy at the ZLB

Consider first a central bank which is unable to credibly commit to its promises and operates under discretion. Such a central bank minimizes the welfare loss function period-by-period. In period 1 after the ZLB episode has ended, the central bank seeks to set the interest rate i_1 to minimize the welfare loss, subject to the equilibrium conditions (7) and (8):[8]

$$\min_{i_1} \pi_1^2 + \lambda x_1^2 \tag{19}$$

$$\text{s.t.} \quad \pi_1 = -\sigma k\left[i_1 - \tfrac{1}{\beta}\right]$$

$$x_1 = -\sigma\left[i_1 - \tfrac{1}{\beta}\right]$$

Observe that the central bank can perfectly minimize its loss function in period 1 by setting the nominal rate equal to the natural rate $r_1^n = 1/\beta$:

$$i_1 = \tfrac{1}{\beta}$$

By setting the nominal rate equal to the natural rate, this ensures zero inflation, $\pi_1 = 0$, which ensures that the underlying Calvo price rigidity has no consequence for real output or the output gap, $x_1 = 0$.

Note that this optimal policy problem and its implication – "set the nominal rate equal to the natural rate" – is entirely independent of the ZLB constraint. That is, whether or not the ZLB were to bind at time 0, this would be optimal policy at time 1 for a discretionary central bank.

Now compare this to optimal policy as set by a central bank which is able to credibly commit to its promises. Such a central bank maximizes lifetime welfare (9), by choosing the interest rate *path* $\{i_0, i_1\}$, subject to the ZLB constraint (10) and the equilibrium conditions (13)-(16):

[8] The ZLB does not bind at time 1, so for clarity I ignore it here.

$$\max_{i_0,i_1} \pi_0^2 + \lambda x_0^2 + \beta[\pi_1^2 + \lambda x_1^2] \tag{20}$$
$$\text{s.t.} \quad i_0 \geq 1$$
$$\text{and (13)-(16)}$$

One can then substitute out the endogenous objects $\{\pi_t, x_t, r_t^n\}$ using the equations (13)-(16) derived above and take the first order conditions on the central bank's problem under commitment. After tedious algebra, and again under assumption (7) that the ZLB binds at $t = 0$, optimal commitment policy is derived as:

$$i_0 = 1$$
$$i_1^* = \frac{1}{\beta} - \alpha(1-\phi) \tag{21}$$

Here, α is an unimportant positive constant.[9]

What is the interpretation? In period zero, the central bank sets the nominal rate as low as possible, at the zero lower bound. In period one, the central bank commits itself to follow through on *expansionary forward guidance*: the optimal interest rate after the ZLB under commitment, i_1^*, is *below* the natural rate $r_1^n = 1/\beta$ since $1 - \phi > 0$ by assumption and $\alpha > 0$. The more binding is the ZLB during the liquidity trap, then the stronger is the forward guidance, i.e. the smaller is ϕ the smaller is i_1^*. Observe from equations (13) and (14) that this policy of setting $i_1 < r_1^n = 1/\beta$ results in positive inflation and a boom in the output gap during period 1.

Thus, under optimal commitment policy, the central bank promises to have a lower interest rate in the future than would otherwise be optimal without the ZLB, which creates a boom in period 1 after the ZLB episode. Expectations during the liquidity trap at $t = 0$ for that future boom at $t = 1$ help to ameliorate the deflationary recession during the liquidity trap.

Finally, we can observe the time consistency problem of the ZLB: optimal policy under commitment differs from optimal policy under discretion. The discretionary central bank, unable to keep its promises, would like to promise to keep the nominal rate low after the ZLB episode ends; but it is unable to commit to follow through on its forward guidance, and it reverts to setting $i_1 = 1/\beta$ and does not create the post-ZLB boom. By definition of the max operator, such a discretionary policy achieves lower lifetime utility: the world would be better off if the central bank could commit. Hence, the idea that there is an "expectations trap" at the ZLB due to the time consistency problem, rather than a "liquidity trap" where the central bank is mechanically unable to affect inflation (Krugman 1998; Eggertsson and Woodford 2003).

---

[9] $\alpha = \frac{B}{A} > 0$, where $B \equiv k^2(1 + \sigma k + \beta) + \lambda(1 + k\sigma) > 0$ and $A \equiv k^2(1 + \sigma k + \beta)^2 + \lambda(1 + k\sigma)^2 + \beta k^2 + \lambda\beta > 0$. We also clearly need to assume a configuration of parameters such that the ZLB constraint does not bind at $t = 1$, i.e. $i_1^* \geq 1$.

2.5 The fewer-instruments-than-targets problem

The expectations trap is the first challenge created by the ZLB, but there is another: even if the central bank can commit, it still cannot achieve the first best. In the first-best world, the central bank would be unconstrained by the ZLB, and would be able to set the policy interest rate equal to the natural rate in both periods – $i_0 = \phi$ and $i_1 = 1/\beta$ – which would ensure both zero inflation and zero output gap in both periods, $\pi_t = x_t = 0$, which would be optimal given the objective function (9).

The reason the central bank cannot achieve the first best at the ZLB, even under commitment, is that the ZLB constraint means that there are in effect fewer instruments than targets, a la Poole (1970). In effect, there are two targets – zero inflation in each period, $\pi_0 = 0$ and $\pi_1 = 0$, since zero inflation would also ensure zero output gap – but only one free instrument i_1, because the instrument at time 0 is locked by the ZLB, $i_0 = 1$. As a result, there is the optimal policy rate $i_1^* < r_1^n$ derived above, which trades off: the benefit of bringing inflation during the ZLB, π_0, closer to the desired level of zero *versus* the cost of raising inflation afterwards, π_1, above zero.

As an aside to the main framework used in this paper, I note that while this second "fewer-instruments-than-targets" problem arises in the mainline New Keynesian model, it does not appear in all microfounded business cycle models. Under the staggered pricing of the Calvo (1983)-Yun (1996) friction used in the New Keynesian framework, any amount of inflation creates inefficient price dispersion, because otherwise-identical firms are forced to make different pricing decisions by the exogenous Calvo fairy. This means that a promise by the central bank during a liquidity trap for inflation after the episode ends necessarily creates distortions. These distortions lower welfare and prevent the central bank from achieving the first best even when it can commit.

Under other, non-Calvo nominal rigidities where pricing decisions are more synchronized, forward guidance is not distortionary, and thus credible monetary policy can achieve the first best.[10] Mankiw and Weinzierl (2011) study optimal monetary and fiscal policy at the ZLB in such an environment, where all prices are completely rigid in period 0 and completely flexible in period 1; see also Auerbach and Obstfeld (2005).[11] One-period information frictions have the same property, such as in the model of Lucas (1972) or in simple versions of the "sticky information" model of Mankiw and Reis (2002). The necessary condition for achieving efficiency in these and other

[10] As just one example, in the menu cost world of Caratelli and Halperin (2024), sufficiently strong forward guidance would cause all firms to adjust, and thus would not be distortionary except for the direct welfare loss from menu costs themselves.

[11] The justification for the staggering assumption of the Calvo-Yun formulation and other time-dependent pricing rules is the claim that price dispersion is present in the data; see Taylor (2016, sec. 3) for a review.

models is that *pre-announced* inflation is not distortionary, so forward guidance is not distortionary.

Having set out the baseline ZLB model, I now turn to my five theses, beginning with two on optimal monetary policy relating to the *two* challenges created by the ZLB just described.

3 Thesis 1: the time consistency problem for monetary policy at the ZLB can be easily overcome by reputational effects

As is well-known in game theory, repeated games are fundamentally different from one-shot games. For instance, although the prisoner's dilemma is truly a dilemma if played once, in a repeated game, cooperation has the possibility of being sustained.

In this section, I extend the baseline ZLB model developed above from a one-shot game where the ZLB binds only once, to a repeated game where there is always a chance that the ZLB may bind again in the future. I show that if the frequency of future ZLB episodes is high enough, then the desire by a discretionary central bank to renege on its forward guidance is outweighed by its desire to maintain its reputation for future episodes, and there is no time consistency problem.

Nakata (2018) has previously highlighted this point explicitly in a larger-scale model, with calibration to data. I also discuss how this result is analogous to the way that Barro and Gordon (1983) showed that desire for an inflation-fighting reputation can overcome the inflationary time consistency problem that central banks face in the model of Kydland and Prescott (1976).[12]

3.1 A model of ZLB cycles

Consider an infinite-horizon model, where each period consists of two sub-periods. Those two sub-periods map to periods $t = 0$ and $t = 1$ of the baseline model; there is no connection between sub-periods of one period with sub-periods of a separate period. Thus, the model consists of playing the baseline model, repeatedly. We can conceive of a single period as a "business cycle", where subperiod 0 is the potential ZLB event and subperiod 1 is the successive recovery until the next business cycle, with connections between business cycles severed for the sake of clarity.

For a given variable X, denote $X_{t(0)}$ the value of the variable at time t in subperiod 0, and $X_{t(1)}$ likewise for subperiod 1.

[12] Stokey (1989, 1991) generalizes this point to other seemingly time inconsistent government policies.

It is now interesting to allow the natural rate during the ZLB to be stochastic. Suppose with probability p the ZLB binds in subperiod 0; otherwise the natural rate is at its steady-state level and the ZLB does not bind.

$$r^n_{t(0)} = \begin{Bmatrix} \phi & \Pr = p \\ 1/\beta & \Pr = 1-p \end{Bmatrix} \tag{22}$$

In subperiod 1, the natural rate is always equal to the steady-state level, $r^n_{t(1)} = 1/\beta$, as in the baseline model.

The central bank's lifetime loss function is:

$$\mathbb{W} = \sum_{t=0}^{\infty} \delta^t \left[\pi^2_{t(0)} + \lambda x^2_{t(0)} + \beta\left[\pi^2_{t(1)} + \lambda x^2_{t(1)}\right]\right] \tag{23}$$

β is now the discount rate between *sub*periods, whereas δ is the (potentially different) discount rate between periods. By differentially setting these discount factors, we effectively allow for the two subperiods to be of different lengths: recessions and recoveries may be of different duration.

Exactly as in the baseline model of section 2, we can solve for the four equilibrium objects of each period $\{\pi_{t(0)}, x_{t(0)}, \pi_{t(1)}, x_{t(1)}\}$. These equations have exactly the same form as equations (13)-(16), except merely for the notational difference to account for the existence of subperiods.

$$x_{t(1)} = -\sigma\left[i_{t(1)} - \tfrac{1}{\beta}\right] \tag{24}$$

$$\pi_{t(1)} = -\sigma k\left[i_{t(1)} - \tfrac{1}{\beta}\right] \tag{25}$$

$$x_{t(0)} = -\sigma(1+k\sigma)\left[E_{t(0)} i_{t(1)} - \tfrac{1}{\beta}\right] - \sigma\left[i_{t(0)} - \phi\right] \tag{26}$$

$$\pi_{t(0)} = -\sigma k[1+\sigma k+\beta]\left[E_{t(0)} i_{t(1)} - \tfrac{1}{\beta}\right] - \sigma k\left[i_{t(0)} - \phi\right] \tag{27}$$

where $E_{t(0)}$ denotes expectations made at time t in subperiod 0.

3.2 Optimal policy under discretion, without reputation

As a baseline, consider the case where the central bank is myopically operating under discretion, without considerations for reputation. Optimal policy under discretion then is exactly the same as in the baseline model with discretion. In subperiod 0 of every period, the central bank sets the nominal rate equal to the natural rate if possible, otherwise to zero; and in subperiod 1 sets it equal to the natural rate.

$$i_{t(0)} = \begin{Bmatrix} 1 & \text{if } r^n_{t(0)} < 1 \\ r^n_{t(0)} & \text{else} \end{Bmatrix}$$

$$i_{t(1)} = 1/\beta$$

If the ZLB binds in subperiod 0, there is a recession and deflation.

$$x_{t(0)} = \begin{Bmatrix} -\sigma(1-\phi) < 0 & \text{if } r^n_{t(0)} < 1 \\ 0 & \text{else} \end{Bmatrix}$$

$$\pi_{t(0)} = \begin{Bmatrix} -\sigma k(1-\phi) < 0 & \text{if } r^n_{t(0)} < 1 \\ 0 & \text{else} \end{Bmatrix}$$

In subperiod 1, optimality is achieved, with the output and inflation at zero.

$$x_{t(1)} = 0$$
$$\pi_{t(1)} = 0$$

As discussed in section 2, the central bank could achieve higher welfare in periods where the ZLB binds if it were able to commit to a lower future interest rate, $i_{t(1)} < 1/\beta$, in order to ameliorate the recession and deflation in subperiod 0.

3.3 Optimal policy under commitment

Now consider the case where the central bank can operate under commitment.

The optimal commitment problem involves choosing a state-contingent plan for interest rates, depending on if the natural rate in subperiod 0 realizes as $r^n_{t(0)} = \phi$ and the ZLB binds, or if not and $r^n_{t(0)} = 1/\beta$. Denote these as states $s_{t(0)} \in \{L, H\}$ where the low state $s_{t(0)} = L$ is the case where the ZLB binds and $r^n_{t(0)} = \phi < 1$; and the high state $s_{t(0)} = H$ otherwise. Denote the interest rate in subperiod 0 of time t in each state s as $i_{t(0)}(s)$ and in subperiod 1 as $i_{t(1)}(s)$.

The optimal commitment problem then is to choose a state-contingent sequence for the interest rate to minimize the expected lifetime loss function (14):

$$\min_{\{i_{t(0)}(s),\ i_{t(1)}(s)\}_{s\in\{L,H\}}} \mathbb{E}_0 \sum_{t=0}^{\infty} \delta^t \left[\pi^2_{t(0)} + \lambda x^2_{t(0)} + \beta\left[\pi^2_{t(1)} + \lambda x^2_{t(1)}\right]\right]$$
$$\text{s.t. } i_{t(0)}(s), i_{t(1)}(s) \geq 1 \qquad \forall s, t$$
$$\text{and (24)-(27)} \qquad \forall s, t$$

It is easy to see that optimal policy under commitment policy is, in each cycle, the same as the baseline model, since there is no connection between policy in one cycle and outcomes in another. That is:

1. If the ZLB is binding in subperiod 0 since $r^n_{t(0)} = \phi$, then it is optimal to set the nominal rate in subperiod 0 at the ZLB, $i_{t(0)} = 1$, and to set the nominal rate in subperiod 1 equal to the optimal level from commitment policy in the baseline model, $i_{t(1)} = i^*_1$.
2. If the ZLB is not binding in subperiod 0 since $r^n_{t(0)} = 1/\beta$, then it is optimal to set the nominal rate in both subperiods equal to the natural rate, $i_{t(0)} = i_{t(1)} = 1/\beta$, so that there is neither inflation nor any output gap.

Written formally, optimal policy under commitment is:

$$\begin{aligned} &i_{t(0)}(L) = 1, && i_{t(1)}(L) = i^*_1 \\ &i_{t(0)}(H) = 1/\beta, && i_{t(1)}(H) = 1/\beta \end{aligned} \tag{28}$$

where i^*_1 was defined in equation (21).

3.4 Optimal policy under discretion, with reputation

So far we have considered: optimal policy under discretion without reputation; and optimal policy under commitment.

Now consider the case where the central bank operates under discretion – but suppose that the public plays a one-period punishment when setting their expectations, so that there is a reputational mechanism, which I will show to be rational. In particular, suppose that the public expects during a ZLB episode today that:

1. If during the last business cycle the central bank did *not* follow through on its forward guidance, then it will again not do so during this business cycle.
2. If, on the other hand, during the last business cycle the central bank did not deviate from any promises, then it will keep any promise made during this business cycle.

In math, this expectations rule is written as:

$$\mathbb{E}_{t(0)}\left[i_{t(1)}(L)\right] = \begin{Bmatrix} i_t^* & \text{if} \quad i_{t-1,(1)} = E_{t-1,(0)}\left[i_{t-1,(1)}\right] \\ \frac{1}{\beta} & \text{else} \end{Bmatrix} \tag{29}$$

This brings us to proposition 1, the main result of this section.

Proposition 1: If the probability of a ZLB episode p is above a threshold $\overline{p}$, then a discretionary central bank playing the optimal commitment policy (28) together with the public playing one-period punishment (29) is an equilibrium. That threshold is:

$$\overline{p} \equiv \frac{\beta}{\delta}\left(1 + \sigma k + \beta \frac{[k^2[1+\sigma k+\beta]+\lambda+k^2]}{[k^2[1+\sigma k+\beta]+\lambda+k\sigma\lambda]}\right)^{-2} \tag{30}$$

Restated, theorem 1 says that if ZLB episodes are sufficiently frequent, then even a central bank which does not have commitment power can nonetheless successfully implement the optimal commitment policy, in the face of the one-period punishment strategy by the public.[13] The intuition here is, as emphasized above, that the central bank trades off the gain from reneging on its promise and ensuring zero inflation, versus the cost of losing its credibility if it faces another ZLB episode in the next period.

The proof reflects that intuition, as the proof is directly from the single-period deviation criterion. Per that criterion, the conjectured equilibrium is indeed an equilibrium if and only if

[13] An alternative, equivalent framing would be: if the central bank is sufficiently patient, i.e. β is sufficiently high, then optimal commitment policy can be sustained.

$$\underbrace{\beta\left(\left[\pi^{\text{commit}}_{t(1)}\right]^2+\lambda\left[x^{\text{commit}}_{t(1)}\right]^2\right)}_{\text{gain from cheating today}} \tag{31a}$$

$$<$$

$$\underbrace{\delta p\left(\left[\pi^{\text{disc}}_{t+1,(0)}\right]^2+\lambda\left[x^{\text{disc}}_{t+1,(0)}\right]^2\right)-\delta p\left(\left[\pi^{\text{commit}}_{t(0)}\right]^2+\beta\lambda\left[x^{\text{commit}}_{t(0)}\right]^2+\beta\left[\pi^{\text{commit}}_{t(1)}\right]^2+\beta\lambda\left[x^{\text{commit}}_{t(1)}\right]^2\right)}_{\text{discounted expected loss from punishment tomorrow}} \tag{31b}$$

$\pi^{\text{commit}}, x^{\text{commit}}$ denote outcomes under commitment play in the low state; $\pi^{\text{disc}}, x^{\text{disc}}$ denote outcomes under discretionary play without reputation in the low state, as this is what occurs under punishment.[14] The term in (31a) is the welfare loss in period t, in subperiod 1, which is avoided by cheating and playing $i_{t(1)} = 1/\beta$. The first term in (31b) is the expected discounted welfare in $t+1$ under punishment when the central bank is forced to act without forward guidance. The second term in (31b) is the expected discounted value of what welfare would have been in $t+1$ if the central bank did not deviate. The chance that the ZLB is binding in $t+1$, captured by the probability parameter p, is what creates the chance for punishment. Simplifying the algebra gives the condition (30) of proposition 1.

The result in proposition 1 can be strengthened further if the public plays a multi-period – or permanent grim-trigger – punishment strategy. That is, under these stronger punishments, the threshold probability for ZLB episodes necessary to sustain forward guidance is even lower than that given in theorem 1.

3.5 Relation to Barro and Gordon (1983) and discussion of post-Great Recession policy

As noted in the introduction to this section, it is a basic implication of microeconomic game theory that commitments which are noncredible in one-shot games can be credible in repeated games. This is simply an implication of the folk theorem. The math above merely formalizes this in our setting.

The logic of theorem 1 – that reputation can overcome the *deflationary* time consistency problem of the ZLB – is also directly analogous to the work of Barro and Gordon (1983) on how reputation can overcome an *inflationary* time consistency problem facing central banks. Kydland and Prescott (1976) had set up an environment where the central bank is continually tempted to implement

[14] That is:
$\pi^{\text{commit}}_{t(1)} = \sigma k\alpha(1-\phi)$
$x^{\text{commit}}_{t(1)} = \sigma\alpha(1-\phi)$
$\pi^{\text{commit}}_{t(0)} = \sigma k[1+\sigma k+\beta]\alpha(1-\phi)-\sigma k[1-\phi]$
$x^{\text{commit}}_{t(0)} = \sigma(1+k\sigma)\alpha(1-\phi)-\sigma[1-\phi]$
$\pi^{\text{disc}}_{t+1,(0)} = -\sigma k[1-\phi]$
$x^{\text{disc}}_{t+1,(0)} = -\sigma[1-\phi]$

inflationary policy to goose the economy; however, the public recognizes this temptation, and merely raises their inflation expectations to account for this. The result is higher inflation, without any of the benefits of stimulative monetary policy. Central banks in this setup thus have an inflationary bias caused by a time consistency problem, where they would like to commit to low inflation, but then would be tempted to deviate to goose the economy.

Barro and Gordon (1983) point out in reply that central banks are playing a repeated game with the public, and concern about reputation can overcome the Kydland-Prescott time consistency problem. Although the central bank may be continually tempted to goose the economy for the sake of monetary stimulus, it also cares that the public trust it in the future not to inflate. That trust is necessary to keep the economy steady in the future. It thus has to trade off the small temporary gain from goosing the economy today with the potentially permanent loss to its reputation, which could make it worse off for the entire future. For small enough temporary gains or large enough punishments, together with a sufficiently low discount rate, low inflation can be sustained as an equilibrium and the inflationary bias overcome.

This time consistency of policy seems to be very much the relevant case empirically. Since the 1980s, developed-world central banks have achieved low inflation. Central banks like the Federal Reserve or European Central Bank spend a substantial portion of their energy on attempting to communicate clearly with the public and continuously stress the importance of maintaining their inflation-fighting reputations.

I do not here attempt to do a full quantitative analysis of whether the Federal Reserve or European Central Bank have implemented the commitment-optimal policy during the multiple ZLB episodes of the last 15 years. Speaking qualitatively, in the recovery from the Great Recession, there seems not to have been any period of above-target inflation or economic boom, which is what such policy would prescribe.

The seeming failure to implement optimal commitment policy was despite an overwhelming body of research during the period arguing that the ZLB would be highly likely to bind in years to come (e.g. Kiley and Roberts 2017) as well as explicit acknowledgment of such probabilities by monetary policymakers, such as Jerome Powell.[15] This is suggestive of an important policy failure during the aftermath of the Great Recession. Reiterating my paraphrase of the Krugman (1998) slogan from the introduction: it should be easy – for a central bank which properly has concern

[15] As reported in WSJ (2019), "'The next time policy rates hit the [lower bound] — and there will be a next time — it will not be a surprise,' Mr. Powell said". Nakata and Sunakawa (2018) appendix H also collects quotes from an international set of monetary policymakers expressing concern about their time consistency.

about its reputation – to credibly promise to be 'irresponsible' at the ZLB, since its reputation is on the line.

4 Thesis 2: the ZLB is not special

As highlighted in section 2, the central bank faces *two* problems at the zero lower bound in the canonical New Keynesian model. The first problem is the time consistency facing central banks trying to implement forward guidance, addressed in section 3. The second problem is that the central bank cannot implement the first best at the ZLB, in the New Keynesian model. That is, the central bank cannot implement the flexible price equilibrium, which is the efficient equilibrium.

The purpose of this section is to make the precise the notion that, although indeed the ZLB prevents the central from achieving the first best, it constrains the central bank no differently than does a nominal rigidity on any other relative price.

In the two-period ZLB model developed above, there are two types of goods: C_0 and C_1. There are in effect three nominal prices – P_0, P_1, and the nominal interest rate $1 + i_0$ – which together jointly determine two relative prices, P_0/P_1 and $1 + r_0 = (1 + i_0)\frac{P_0}{P_1}$, which both must be correct in order to implement the efficient equilibrium.

In a world without the ZLB, even if P_0 is completely rigid and there is a shock, the efficient equilibrium can be achieved by moving a combination of P_1 and i_0. On the other hand, if *both* P_0 is rigid *and* i_0 is constrained, e.g. by a zero lower bound, then after a shock the two correct relative prices cannot be achieved.

This inefficiency due to a second rigid *inter*temporal relative price, i_0, is completely analogous to the case of a second rigid *intra*temporal price. For example, imagine an economy of three goods: C_0, C_a, C_b. At time 0, there is only one good C_0 available; at time 1, there are two types of goods C_a, C_b with associated nominal prices P_a, P_b. Suppose again P_0 is rigid but suppose there are no constraints on the nominal interest rates. If there is a lower bound instead on the relative price between type a and type b goods, then shocks shifting the efficient relative price P_a/P_b may result in a "liquidity trap" in precisely the same sense as the ZLB. Whereas in the canonical New Keynesian model, the central bank has two relative prices to get right and can only move one of them, in this setup the central bank has three relative prices to get right and can only move two of them.

The ZLB is a price floor on one particular relative nominal price, and the inefficiency due to the ZLB is conceptually the same as the inefficiency resulting from any price floor on any relative nominal price, including intratemporal relative prices. Indeed, the same inefficiency arises from symmetric intratemporal relative price rigidities. This includes different frequencies of price adjustment across sectors, which have been well-documented empirically (Rubbo 2023).

5 Thesis 3: austerity is stimulus

5.1 Setup

I now turn to three theses on fiscal policy in the representative agent New Keynesian model (RANK). I emphasize that these conclusions are derived in the representative agent version of the New Keynesian framework in order to highlight that Ricardian equivalence holds, so that "fiscal stimulus" is modeled as government consumption. The recent heterogeneous agent New Keynesian (HANK) literature often analyzes "stimulus" as *transfers* (e.g. Wolf 2021), which in the Ricardian setting of the RANK model analyzed here have no effect.[16] That said, the insights for fiscal purchases analyzed here are still relevant to understanding the effect of fiscal purchases in HANK models.

To analyze optimal fiscal stimulus, I extend the baseline model of section 2 to include government spending on public goods, g_t, exactly as in Werning (2012), Woodford (2011), and Eggertsson (2001). I highlight terms related to this government spending in red.

Since I have written the model in terms of the output gap rather than in terms of output itself, neither the NKPC equation (1) nor the EE equation (2) of the baseline model need change their structure.[17,18] The introduction of government spending only changes the path for the natural rate (3) and the welfare loss function (4).

With government spending, the per-period welfare loss function is:[19]

$$\mathbb{W}_t = \pi_t^2 + \lambda x_t^2 + \lambda_g g_t^2 \tag{32}$$

The log-deviation of government spending, g_t, enters the welfare loss function quadratically: *either* a decrease *or* an increase in government spending g_t directly reduces welfare. This is an

[16] The conclusion section expands on these points and offers further relevant citations.

[17] The output gap is consistently defined as the difference between log-deviations in output y_t and log-deviations in the flex-price level of output, y_t^f: that is, $x_t \equiv y_t - y_t^f$. The introduction of government spending does affect y_t^f, the underlying flexible-price level of output. As the appendix shows, $y_t^f = \Gamma g_t$ where $\Gamma > 0$ is a constant discussed in footnote 20 and the appendix. However, the dynamics of the output gap and inflation (equations 1 and 2) are not affected by this change.

[18] I abuse notation, however, since the parameter σ now needs to be normalized. Whereas without government spending σ was the inverse of the intertemporal elasticity of substitution, now it is this inverse elasticity divided by the steady state ratio of consumption to output. See the appendix for further detail.

[19] This, as before in equation (3), is a second-order approximation around the efficient steady state. In particular, this is taken around the steady state where the fiscal authority has implemented the Samuelson rule for the optimal quantity of public goods. See the appendix for further detail.

implication of microeconomic efficiency concerns. There is some level of government production of public goods which is optimal, as determined by the canonical Samuelson rule for public goods provision. If the production of public goods is below this level, then public goods are underprovided. But also, symmetrically, if the production of public goods is above this level, then households need to overwork in order to produce these public goods. This overproduction is costly due to the disutility of the required excess labor. Hence, the quadratic loss term, weighted by the parameter $\lambda_g > 0$ which reflects the relative importance of this policy goal.

5.2 The natural rate under fiscal stimulus

Critically for the next two sections, the natural rate process is also affected by government purchases. In the baseline model equation (3), the natural rate merely reflected time preference: $r_t^n = \rho_t$ where ρ_t would either be the shock of $\phi < 1$ or the steady state level of $1/\beta$.

With government spending, the natural rate becomes:

$$r_t^n = \rho_t - \gamma E_t \Delta g_{t+1} \quad (33)$$

Here, $\gamma \geq 0$ is a parameter detailed in a footnote[20], which is strictly nonzero unless *both* of the following hold:

1. The household has GHH preferences, so that there are no income effects on labor supply; *and*
2. Production technology is constant returns to scale, rather than decreasing returns to scale.

If both of these conditions hold, then fiscal stimulus has *no* effect on the natural rate, inflation, or the output gap in RANK. I will assume this parameter is nonzero, $\gamma > 0$, to ensure that discussing fiscal policy is even interesting in the first place:

Assumption 1: Either the household does not have GHH preferences, or production technology is not constant returns to scale, or both.

5.3 The baseline NK model with fiscal stimulus

Collecting equations (1), (2), (32), and (33), the baseline New Keynesian model with fiscal stimulus is:

$$\pi_t = kx_t + \beta E_t \pi_{t+1}$$
$$x_t = E_t x_{t+1} - \sigma[i_t - E_t \pi_{t+1} - r_t^n]$$
$$r_t^n = \rho_t - \gamma E_t \Delta g_{t+1}$$

[20] The parameter γ is defined in the appendix: to summarize, $\gamma = \sigma(1 - \Gamma)$. To define Γ in turn, suppose that household preferences over consumption C labor N and public goods G are $u(C) - v(N) + H(G)$ and that production technology of any firm i is $Y_i = f(N_i)$. Additionally, define $\tilde{v}(Y) \equiv v\big(f^{-1}(Y)\big)$. Additionally, let C and Y be the steady state levels of consumption and output. Now finally define $\Gamma \equiv \frac{\eta_u}{\eta_u - \eta_{\tilde{v}}}$, and $\eta_u \equiv -\frac{u''}{u'}C$ and $\eta_{\tilde{v}} \equiv \frac{\tilde{v}''}{\tilde{v}'}Y$. See Woodford (2011) for additional discussion.

$$\mathbb{W}_t = \pi_t^2 + \lambda x_t^2 + \lambda_g g_t^2$$

Under assumption 1, fiscal stimulus affects the output gap and inflation, and does so exclusively through its effect on the natural rate (33). The interpretation of equation (33) for the mechanism of effect of fiscal stimulus is directly apparent from the math itself:

Proposition 3: Only the *change* in fiscal spending Δg_{t+1} matters for the output gap x_t and inflation π_t, not the *level* of fiscal spending.

This equation (24) for the natural rate is not intrinsically new, but the statement of proposition 3 – that *only* the change in spending matters from a welfare perspective – is framed especially clearly by formulating the model in terms of the output gap, rather than in terms of consumption or output.

Raising the level of government spending today, $g_t = g > 0$, while leaving government spending unchanged tomorrow, $g_{t+1} = 0$, is indeed stimulatory. In this case the change in government spending is negative, $\Delta g_{t+1} = -g < 0$, so that the natural rate today r_t^n is pushed up by the amount $\gamma \cdot g$, as can be read off of equation (24). If the *nominal* interest rate is fixed by the ZLB, $i_t = 1$, then this raising of the natural rate is expansionary, since the interest rate gap $i_t - r_t^n$ falls.

But note that it would be *just* as stimulatory to promise to cut future spending by an equivalent amount, $g_{t+1} = -g$, and simultaneously not raise spending today, $g_t = 0$. This would have the same effect on the *change* in spending, $\Delta g_{t+1} = -g < 0$, and therefore the same effect on the natural rate.[21]

Thus, the titular claim of this section: austerity is stimulus. Raising spending today, $g_t \uparrow$, is *just* as stimulatory as promising to cut spending tomorrow, $g_{t+1} \downarrow$.

5.4 Intuition: consumption smoothing

The result that fiscal stimulus in RANK only depends on the change in government spending, rather than the level, is counterintuitive if one thinks of the New Keynesian model as formalizing the Old Keynesian logic for fiscal stimulus, where such spending *mechanically* boosts the economy, without any intertemporal mechanism. What is the intuition for this result?

[21] Since fiscal spending must return to steady state, this will require (say) $g_{t+2} = 0$, and therefore r_{t+1}^n would be changed by $-\gamma g$. If, as in the baseline model, the ZLB does not bind at $t + 1$, then this can easily be offset by monetary policy, and there is precisely zero effect on the output gap and inflation at $t + 2$.

The key intuition comes from consumption smoothing (Rowe 2012). Consider the aggregate resource constraint (without log-linearization), $Y = C + G$. If assumption 1 does *not* hold, then it can be shown that an increase in government spending G increases output Y one-for-one, leaving consumption C unaffected as a result.[22] Supposing instead (and realistically) that assumption 1 does hold, then an increase in government spending G increases output Y – but strictly less than one-for-one. As a result, consumption C must fall: government consumption crowds out private consumption.

Now suppose the representative agent knows that government spending tomorrow is going to be lower than it is today, holding all else constant. From the resource constraint (under assumption 1), this implies that *consumption* tomorrow will be *higher* than it is today: more of total output Y is available for private consumption C.

With the knowledge that government spending tomorrow will be lower than today and that therefore *consumption* will be *higher* than today, the representative agent would like to save less today in order to consume more today, to better smooth consumption over time. This reduced desire to save pushes up the underlying natural rate.

5.4 Further intuition: positive fiscal stimulus is isomorphic to a negative productivity shock

A further intuition pump to understand fiscal policy in the New Keynesian model is to compare the effect of government spending to the effect of productivity shocks. So far, we have assumed that productivity is constant. With time-varying productivity a_t, neither the NKPC (1) nor the Euler equation (2) are affected, but the formula for the natural rate becomes:

$$r_t^n = \rho_t - \gamma E_t \Delta g_{t+1} + \psi \Delta a_{t+1} \tag{34}$$

Here, I have highlighted the new terms introduced by productivity shocks in blue; $\psi > 0$ is a constant parameter.[23]

A temporary negative productivity shock means that productivity is lower today than tomorrow, $a_t < a_{t+1}$, and so $\Delta a_{t+1} > 0$. Meanwhile, positive fiscal stimulus means that government spending is higher today than tomorrow, $g_t > g_{t+1}$, and so $\Delta g_{t+1} < 0$. Paying careful attention to the signs of these various inequalities, we can see in equation (34) that a negative productivity shock and a positive fiscal stimulus have the same sign of an effect on the natural rate. Since neither type of

[22] Woodford (2011) has a very clear exposition.

[23] As in footnote 15, the definition of flex-price output y_t^f changes once we incorporate time-varying productivity, which feeds into the output gap $x_t = y_t - y_t^f$ where y_t is (the log deviation of) output. With time-varying productivity, $y_t^f \equiv \Gamma g_t + \psi\sigma^{-1} a_t$, where again ψ is a positive constant defined in the appendix.

shock directly affects inflation or the output gap in *any* other equation of the model, the two kinds of shocks have the exact same effect on the equilibrium output gap and inflation of the model.[24]

Meanwhile, it is well-appreciated that negative productivity shocks at the ZLB are expansionary. Among many other papers, Eggertsson (2012) discusses this mechanism in theory. Wieland (2019) offers a critical empirical evaluation; he looks at the effects of the Great East Japan Earthquake and oil supply shocks and finds them not to be expansionary.[25]

The analogy between fiscal stimulus and negative productivity shocks also provides a less flattering view of what fiscal stimulus is, conceptually, in RANK. Government spending in effect *is* a negative productivity shock here, by making the economy less efficient at transforming labor inputs into household welfare. Both government spending and negative productivity shocks operate by making households worse off today – shifting in the utility possibilities frontier of the shadow flex-price economy – which in turn pushes up the natural rate through consumption smoothing.

6 Thesis 4: heterogeneous fiscal stimulus can be contractionary

In this section, I show that when fiscal stimulus is targeted heterogeneously across sectors, it can *lower* the natural rate, and therefore is contractionary if the nominal interest rate is at the zero lower bound. This is in contrast to the standard one-sector version of the New Keynesian model, where fiscal stimulus always raises the natural rate and therefore is expansionary.

In this section, I briefly step away from the simple log linearized New Keynesian model used in the rest of the paper in order to illustrate the results more clearly. The result can be shown with a two sector model, with the representative household consuming C_A from sector A and C_B from sector B, as in Guerrieri, Lorenzoni, Straub, and Werning (2022).

6.1 Setup and steady state

The representative household has the following period utility:

$$U(C_{At}, C_{Bt}) = \tfrac{\sigma}{\sigma-1}\left[\phi^{\frac{1}{\epsilon}}C_{At}^{\frac{\epsilon-1}{\epsilon}} + (1-\phi)^{\frac{1}{\epsilon}}C_{Bt}^{\frac{\epsilon-1}{\epsilon}}\right]^{\frac{\epsilon}{\epsilon-1}\frac{\sigma-1}{\sigma}}$$

[24] Positive government spending shocks and negative productivity shocks have the same effect on the output gap and inflation; but negative productivity shocks have an additional, direct, policy-independent negative effect on welfare. In other words, negative productivity shocks additionally directly lower welfare, in such a way that neither fiscal nor monetary policy can affect. This is swept into the typical "terms independent of policy (t.i.p.)" term of the second-order welfare approximation.

[25] Wieland (2019) argues that both the earthquake and the identified oil supply shocks were temporary, rather than permanent, shocks, which is necessary for the logic here to apply.

Thus, the household has a constant elasticity of substitution ϵ between goods from the two sectors and a constant elasticity of intertemporal substitution σ.[26] I assume for simplicity that the household inelastically supplies $N_{At} = \phi$ units of labor to sector A and $N_{Bt} = 1 - \phi$ units of labor to sector B, earning a common nominal wage W_t, which is flexible.

Goods production is perfectly competitive, and each sector has a representative firm producing with linear technology: for sector $i \in \{A, B\}$,

$$Y_{it} = N_{it}$$

When prices are flexible, due to perfect competition, prices equal marginal cost, $P_{it} = W_t$.

Define the real interest rate in terms of good B:

$$1 + r_t \equiv (1 + i_t) \frac{P_{Bt}}{P_{Bt+1}}$$

The Euler equation in terms of good B then is:

$$1 + r_t = \frac{1}{\beta} \frac{U_2(C_{At}, C_{Bt})}{U_2(C_{At+1}, C_{Bt+1})}$$

The fiscal authority can purchase goods from either sector, denoted G_i. Government purchases could be included additively in the utility function, but it does not matter for the purposes of the discussion here. The resource constraint for each good is:

$$Y_{it} = C_{it} + G_{it}$$

6.2 The natural rate of interest

In the flexible price economy, under inelastic labor supply, equilibrium is quickly obtained as follows. Prices equal the nominal wage in both sectors, $P_{it} = W_t$. Output is pinned down in each sector by the inelastic labor supply:

$$Y_{At} = \phi$$
$$Y_{Bt} = 1 - \phi$$

Consumption by sector is in turn pinned down by the resource constraint and the exogenous stream of government purchases:

$$C_{At} = \phi - G_{At}$$
$$C_{Bt} = (1 - \phi) - G_{Bt}$$

The real interest rate in this flexible economy is the natural rate of the corresponding sticky-price economy, and it is found from the Euler equation:

$$1 + r_t = \frac{1}{\beta} \frac{U_2(\phi - G_{At}, 1 - \phi - G_{Bt})}{U_2(\phi - G_{At+1}, 1 - \phi - G_{Bt+1})}$$
$$= \frac{1}{\beta} \left[\frac{1 - \phi - G_{Bt}}{1 - \phi - G_{Bt+1}} \right]^{-\frac{1}{\epsilon}} \left\{ \frac{\phi^{\frac{1}{\epsilon}} (\phi - G_{At})^{\frac{\epsilon - 1}{\epsilon}} + (1 - \phi)^{\frac{1}{\epsilon}} (1 - \phi - G_{Bt})^{\frac{\epsilon - 1}{\epsilon}}}{\phi^{\frac{1}{\epsilon}} (\phi - G_{At+1})^{\frac{\epsilon - 1}{\epsilon}} + (1 - \phi)^{\frac{1}{\epsilon}} (1 - \phi - G_{Bt+1})^{\frac{\epsilon - 1}{\epsilon}}} \right\}^{\frac{\epsilon}{\epsilon - 1} \frac{\sigma - 1}{\sigma} - 1}$$

[26] We need to assume $\epsilon > 1$ to handle the case with $C_{it} = 0$; but $\epsilon \leq 1$ works with $C_{it} \to 0$.

For illustration, suppose we start at a steady state where the government purchases zero output, $G_A^{ss} = G_B^{ss} = 0$. At time $t = 0$, there is fiscal stimulus: the government purchases some amount of output. At time $t = 1$, government purchases return to 0.

We are thus interested in the real rate – the natural rate – at time 0, which is given by:

$$1 + r_0 = \frac{1}{\beta}\left[\frac{1-\phi-G_{B0}}{1-\phi}\right]^{-\frac{1}{\epsilon}}\left\{\phi^{\frac{1}{\epsilon}}(\phi - G_{A0})^{\frac{\epsilon-1}{\epsilon}} + (1-\phi)^{\frac{1}{\epsilon}}(1-\phi - G_{B0})^{\frac{\epsilon-1}{\epsilon}}\right\}^{\frac{\epsilon}{\epsilon-1}\frac{\sigma-1}{\sigma}-1}$$

To illustrate the result that positive fiscal stimulus – i.e., $G_{A0}, G_{B0} \geq 0$ – can *lower* the natural rate r_0, it is sufficient to consider an extreme case. Consider the extreme case where the government targets no stimulus in sector B and purchases the entire output of sector A, so that $G_{B0} = 0$ and $G_{A0} = \phi$. The following proposition states the conditions under which this positive fiscal stimulus lowers the natural rate.

Proposition 4: Consider the two-sector model described above, with $G_i^{ss} = 0$ for $i \in \{A, B\}$. Then $r^{ss} = \frac{1}{\beta} - 1$. Suppose at time $t = 0$, there is fiscal stimulus only in sector A: $G_{A0} > 0$ and $G_{B0} = 0$. Then the natural rate r_0 falls if and only if:

$$\sigma > \epsilon$$

Proof: Plugging $G_{B0} = 0$ into the expression for r_0, we have

$$1 + r_0 = \frac{1}{\beta}\{1-\phi\}^{\frac{\epsilon}{\epsilon-1}\frac{\sigma-1}{\sigma}-1}$$

Expanding the exponent, we have $\frac{\epsilon}{\epsilon-1}\frac{\sigma-1}{\sigma} - 1 = \frac{\sigma-\epsilon}{\sigma(\epsilon-1)}$, and the proposition immediately follows. $\square$

This proposition shows that fiscal stimulus targeting one sector but not another can be contractionary if the *inter*temporal elasticity of substitution is larger than *intra*temporal elasticity of substitution. The intuition for this result is the following. When the government purchases the entirety of the output of sector A at time 0, the household is unable to consume any of good A. Because A and B are complementary, this reduces the marginal utility of consuming B at time 0, relative to consuming at time 1: thus, the desire to save increases. The strength of this force is governed by the intertemporal elasticity of substitution, σ. Meanwhile, because of the drop in consumption of A, this increases the marginal utility of consuming at time 0 relative to time 1: thus, the desire to save decreases. The strength of this force is governed by the intratemporal elasticity of substitution, ϵ. When $\sigma > \epsilon$, the former substitution motive dominates, and the real rate falls to clear the market for savings.

The above intuition is precisely analogous to the intuition of Guerrieri, Lorenzoni, Straub, and Werning (2022), who show that a negative productivity shock shutting down one sector can raise the natural rate under the same condition on preferences. This analog is a result of the fact that fiscal stimulus in this model is isomorphic to a negative productivity shock, exactly as discussed in section 5.4 in the context of a single-sector model.

7 Thesis 5: fiscal policy is time inconsistent at the ZLB

While the modern liquidity trap literature began with Krugman's (1998) analysis of the time consistency of monetary policy at the zero lower bound, existing work has not characterized the time consistency of *fiscal policy* at the zero lower bound.[27] Eggertsson and Woodford (2003) analyze optimal monetary policy under discretion versus commitment at the ZLB without fiscal policy; Werning (2012) analyzes optimal fiscal policy under commitment, except in the special case where fiscal policy has no stimulative effect. As far as I am aware, no paper has characterized the time consistency of optimal fiscal policy at the ZLB.

Optimal fiscal policy is time inconsistent whether monetary policy operates under commitment or discretion. Because it is simpler and more intuitive, I work with the case where monetary policy operates under discretion.[28]

We now solve for optimal fiscal policy under commitment, and then do so under discretion. We can then compare the two to demonstrate the time inconsistency.

7.1 Optimal fiscal policy under commitment

We want to first consider the fiscal authority's problem under commitment, in our baseline two-period environment with the addition of fiscal spending. In line with the assumptions of the baseline setting, we assume that after period 1, government spending is at its steady state level, $g_t = 0 \quad \forall t \geq 2$.

As stated above, we will take monetary policy to be operating under discretion. As a result, in period 1 when the ZLB does not bind, the discretionary central bank will always set the nominal interest rate equal to the natural rate; this ensures zero inflation and zero output gap, thereby maximizing its objective function (a). That natural rate is, from equation (33):

$$\begin{aligned} i_1 = r_1^n &= \tfrac{1}{\beta} - \gamma(0 - g_1) \\ &= \tfrac{1}{\beta} + \gamma g_1 \end{aligned}$$

[27] See also Sumner (1993), who emphasized the importance of distinguishing between temporary and permanent changes in the money supply in explaining the lack of inflation during a period of American colonial history.

[28] I also assume that fiscal policy and monetary policy are playing a Nash game (rather than e.g. Stackelberg), as is standard. The time consistency of fiscal policy does not depend on this.

That is, the central bank will always engage in full "monetary offset", offsetting the effect of any fiscal spending g_1 on inflation and the output gap.[29] As we saw in section 2.4, this will ensure that period-1 inflation and the output gap are always zero, $\pi_1 = x_1 = 0$.

The fact that period-1 inflation and the output gap are zero considerably simplifies the fiscal authority's optimization problem, since it removes π_1 and x_1 from consideration. Under commitment, the objective is to minimize

$$\begin{aligned} \mathbb{W} &= [\pi_0^2 + \lambda x_0^2 + \lambda_g g_0^2] + \beta[\pi_1^2 + \lambda x_1^2 + \lambda_g g_1^2] \\ &= [\pi_0^2 + \lambda x_0^2 + \lambda_g g_0^2] + \beta \lambda_g g_1^2 \end{aligned}$$

The full optimization problem under commitment is to maximize this objective, subject to the EE and NKPC of period 0:

$$\max_{g_0, g_1} [\pi_0^2 + \lambda x_0^2 + \lambda_g g_0^2] + \beta \lambda_g g_1^2 \tag{34}$$

$$\text{s.t. } x_0 = -\sigma\big[1 - \phi + \gamma E_0[g_1 - g_0]\big] \tag{35}$$

$$\pi_0 = k x_0 \tag{36}$$

Here, the constraints (35)-(36) are the same as equations (15) and (16) after having substituted in the binding ZLB constraint in period 0, $i_0 = 1$, and including the effect of government spending on the natural rate from equation (33).

Examining (34)-(36), we can see that optimal fiscal policy trades off two goals.

1. From equation (35), fiscal policy would like to raise the period-0 natural rate, $r_0 = \phi - \gamma E_0[g_1 - g_0]$, up to the level of the nominal rate of $i_0 = 1$. This would cause the ZLB to not bind, allowing monetary policy to ensure zero output gap and inflation in period 0, which is desirable.
2. On the other hand, from the objective (34), the fiscal authority would like to keep government spending close to its optimal level, to minimize deviations from the Samuelson rule for public goods provision, $g_0 = g_1 = 0$.

It is not possible to satisfy both of these goals simultaneously, and optimal policy must trade them off. Lemma 1 solves for the optimal path of spending directly from the first order conditions of the linear-quadratic program (34)-(36).

Lemma 1. If fiscal policy can commit and the central bank operates under discretion, then optimal fiscal policy is:

$$g_0^* = \xi[1 - \phi] \tag{37}$$

$$g_1^* = -\tfrac{1}{\beta} \cdot \xi[1 - \phi] \tag{38}$$

Where $\xi > 0$ is a constant defined in the appendix.

[29] As long as the level of government spending is not so negative as to force the central bank back to the ZLB, i.e. $g_1 > -\frac{1}{\gamma}\left(\frac{1}{\beta} - 1\right)$. For more conceptual discussion of the idea of monetary offset, see Sumner (2021).

The interpretation of lemma 1, i.e. optimal fiscal policy under commitment, is that optimally there should be stimulus during the liquidity trap, $g_0^* > 0$, and austerity during the recovery $g_1^* < 0$. This maximizes the impact of fiscal policy on the natural rate in equation (35), while simultaneously smoothing the distortions to the level of government spending in the convex loss function (26). Indeed with no time discounting ($\beta = 1$) then optimally there is perfect smoothing, $g_1^* = -g_0^*$: the level of stimulus during the liquidity trap is the same as the level of austerity during the recovery.

This matches the Old Keynesian logic of "austerity during the boom, stimulus during the bust" but with a very different motive from the traditional logic. The motive here is to smooth distortions to public goods provision – together with a desire to create a negative productivity shock, which is desirable at the ZLB.

7.2 Optimal fiscal policy under discretion

To see the time consistency problem facing fiscal policy, now consider the fiscal authority's problem in period 1 if it were to reoptimize. The period-1 loss function is:

$$\mathbb{W}_1 = \pi_1^2 + \lambda x_1^2 + \lambda_g g_1^2$$

This can again be simplified since the inflation (π_1) and output gap (x_1) terms drop out: because, as above, the discretionary central bank offsets any effect on inflation and the output gap in order to peg them at zero. The period-1 loss function becomes simply:

$$\mathbb{W}_1 = \lambda_g g_1^2 \tag{39}$$

Since objective function (39) is unconstrained, the discretionarily-optimal action at $t = 1$ is simply to set fiscal spending at the steady state level, $g_1 = 0$. Compared to the commitment level of $g_1^* < 0$, the discretionary fiscal authority is spending too high. That is, although the discretionary fiscal authority would ex ante like to commit to austerity in period 1, when that time rolls around, it does not want to engage in the promised austerity.

7.3 The time consistency problem for fiscal policy

Proposition 5 summarizes the above discussion.

Proposition 5. Optimal fiscal policy at the ZLB is not time consistent. Under commitment, the fiscal authority would like to commit to stimulus $g_0 = g_0^* > 0$ during the liquidity trap and austerity $g_1 = g_1^* < 0$ afterwards. However, without commitment power, it will renege on this promise and set $g_1 = 0$.

Even under discretion, the fiscal authority should still engage in stimulus during the liquidity trap; in fact, it is easy to show that under discretion it should engage in a level of stimulus g_0^{disc} even higher than what is optimal under commitment: $g_0^{\text{disc}} > g_0^* > 0$. This is because knowing that it is not able to raise the natural rate $r_0^n = \phi - \gamma E_0[g_1 - g_0]$ by lowering g_1, since it will be unable to follow through on a promise for austerity, it is better off distorting spending in period 0, g_0, even more than under commitment in order to raise r_0^n somewhat more to reduce deflation and the output gap.

We have seen that a fully beneficent, aligned fiscal authority would ideally like to commit to stimulus during a liquidity trap and austerity afterwards, but has a motive to renege on its promised austerity. This time inconsistency problem could be worsened if there are political economy motives for fiscal decisionmakers to avoid austerity. Although I do not model it formally here, if reducing fiscal spending harms legislators' reelection prospects, then this could further exacerbate the time inconsistency problem. Of course, just as reputation in a repeated game can help central banks overcome their apparent time inconsistency problem at the ZLB, so can such a reputation mechanism help fiscal policymakers.[30]

8 Conclusion

I summarize and reiterate the theses of this paper for optimal monetary and fiscal policy at the zero lower bound in the representative agent New Keynesian model.

1. Monetary policy is not pushing on a string at the ZLB because of the power of forward guidance and expectations; and the force of the time consistency problem in blunting this power is significantly lessened by the fact that central banks want to maintain their credibility and reputation.

2. The ZLB is simply a nominal rigidity on one particular relative price. Any nominal rigidity on *any* relative price limits monetary policy in qualitatively the same way: the ZLB is not special.

3. In RANK, fiscal policy affects inflation and the output gap through the *change* in government spending, not the level. This is because positive fiscal stimulus acts as a negative supply shock.

[30] Although, quite plausibly, because central banks exist as an institution which carries a reputation, it may be easier for central banks to overcome the time inconsistency problem compared to the fiscal authority, which typically is thought of as carrying less of an institutional reputation but instead varying with the politicians in office.

4. In models with heterogeneous goods, positive fiscal stimulus targeted at specific sectors can be contractionary.

5. Optimal fiscal policy at the ZLB is time inconsistent, because the fiscal authority will wish to renege on a promise for austerity after the liquidity trap.

I have discussed these points through a series of small models, meant to serve as intuition pumps for the economic logic – rather than as a quantitatively-accurate model that can be taken to the data.

An important and obvious restriction on the generalizability of the results on fiscal policy is that RANK is a Ricardian model, where the power of fiscal policy is fairly limited. For example, stimulus checks – an important component of policy internationally during the COVID-19 recession – have no impact in the RANK framework (see e.g. Wolf 2021 for an analysis). There is also no distinction in the RANK model between deficit-financed spending versus balanced-budget spending (see e.g. Auclert, Rognlie, and Straub 2018 for an analysis). The burgeoning and active literature on optimal policy in heterogeneous agent New Keynesian models is in the process of analyzing these issues (e.g. Dávila and Schaab 2022; McKay and Wolf 2022; Bilbiie 2021; Acharya, Challe, and Dogra 2023; Bhandari, Evans, Golosov, and Sargent 2021; Le Grand, Martin-Baillon, and Ragot 2021).

Additionally, and as noted in section 2.5, the supply side of the RANK model is based on the canonical Calvo-Yun staggered pricing friction ubiquitous in the New Keynesian literature. Alternative forms of nominal rigidities would have different implications for optimal policy. For example, if nominal rigidities are the result of one-period information frictions as in Lucas (1972) – where imperfect pricing is in a sense synchronized rather than staggered – then forward guidance during a liquidity trap is not distortionary after the ZLB ceases to bind, because *pre-announced* monetary policy actions are not distortionary.[31] Under this nominal rigidity, optimal monetary policy both can achieve the first best and is time consistent, and fiscal stimulus is unnecessary. The same would hold if prices of all firms were uniformly set one period in advance without the staggering of the Calvo-Yun friction.[32]

[31] For a more general treatment of optimal policy under information frictions, without the ZLB constraint, see Angeletos and La'O (2020) or Iovino, La'O, and Mascarenhas (2021).
[32] Woodford (2003) refers to the resulting equilibrium of the supply block of such a model as the "New Classical Phillips Curve".

References

Acharya, S., Challe, E., & Dogra, K. (2023). Optimal monetary policy according to HANK. *American Economic Review*, *113*(7), 1741-1782.

Angeletos, G. M., & La'o, J. (2020). Optimal monetary policy with informational frictions. *Journal of Political Economy*, *128*(3), 1027-1064.

Aoki, K. (2001). Optimal Monetary Policy Responses to Relative-Price Changes. *Journal of Monetary Economics*, 48(1), 55-80.

Auclert, A., Rognlie, M., & Straub, L. (2018). The intertemporal Keynesian cross (No. w25020). National Bureau of Economic Research.

Auerbach, A. J., & Obstfeld, M. (2005). The Case for Open-Market Purchases in a Liquidity Trap. *American Economic Review*, 95(1), 110-137.

Barro, Robert, & Gordon, David. (1982). A Positive Theory of Monetary Policy in a Natural Rate Model. *Journal of Political Economy*, 90(4), 589-610.

Barro, Robert, & Gordon, David. (1983). Rules, Discretion and Reputation in a Model of Monetary Policy. *Journal of Monetary Economics*, 12(1), 101-121.

Benigno, P. (2004). Optimal Monetary Policy in a Currency Area. *Journal of International Economics*, 63(2), 293-320.

Bhandari, A., Evans, D., Golosov, M., & Sargent, T. (2021). Efficiency, Insurance, and Redistribution Effects of Government Policies.

Bilbiie, F. O. (2018). Monetary policy and heterogeneity: An analytical framework.

Caratelli, D., & Halperin, B. (2023). Optimal monetary policy under menu costs.

Dávila, E., & Schaab, A. (2023). Optimal Monetary Policy with Heterogeneous Agents: Discretion, Commitment, and Timeless Policy (No. w30961). National Bureau of Economic Research.

Eggertsson, G. B. (2012). Was the New Deal Contractionary? *American Economic Review*, 102(1), 524-555.

Eggertsson, G. B., & Garga, V. (2019). A Model of Secular Stagnation: Theory and Quantitative Evaluation. *American Economic Journal: Macroeconomics*, 11(1), 1-48.

Eggertsson, G. B., Summers, L. H., & Mehrotra, A. (2016). Secular Stagnation in the Open Economy. *American Economic Review*, 106(5), 503-507.

Eggertsson, Gauti, & Woodford, Michael. (2003). The Zero Bound on Interest Rates and Optimal Monetary Policy. *Brookings Papers on Economic Activity*, 2003(1), 139-211.

Garín, J., Lester, B., & Sims, E. R. (2019). Are Supply Shocks Contractionary at the ZLB? Evidence from Utilization-Adjusted TFP Data. *Review of Economics and Statistics*, 101(4), 619-633.

Guerrieri, V., Lorenzoni, G., Straub, L., & Werning, I. (2022). Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages? *American Economic Review*, 112(5), 1437-1471.

Iovino, L., Jennifer, L. O., & Mascarenhas, R. (2022). Optimal monetary policy and disclosure with an informationally-constrained central banker. *Journal of Monetary Economics*, *125*, 151-172.

Kiley, M. T. (2016). Policy Paradoxes in the New Keynesian Model. *Review of Economic Dynamics*, 21, 1-15.

Kiley, Michael, & Roberts, John M. (2017). Monetary Policy in a Low Interest Rate World. *Brookings Papers on Economic Activity*, 2017(1), 317-372.

Koning, J. P. (2016, November 18). A modern example of Gresham's Law. Retrieved from http://jpkoning.blogspot.com/2016/11/a-modern-example-of-greshams-law.html.

Krugman, Paul. (1998). It's baaack: Japan's slump and the return of the liquidity trap. *Brookings Papers on Economic Activity*, 1998(2), 137-205.

Kydland, Finn, & Prescott, Edward (1976). Rules Rather than Discretion: The Inconsistency of Optimal Plans. *Journal of Political Economy*, 85(3), 473-491.

Le Grand, F., Martin-Baillon, A., & Ragot, X. (2021). Should monetary policy care about redistribution? Optimal fiscal and monetary policy with heterogeneous agents. *Science Po Working Paper*.

Lucas, R. E. Jr. (1972). Expectations and the Neutrality of Money. *Journal of Economic Theory*, 4(2), 103-124.

Mankiw, N. G., & Reis, R. (2002). Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve. *Quarterly Journal of Economics*, 117(4), 1295-1328.

Mankiw, N. G., & Weinzierl, M. (2011). An Exploration of Optimal Stabilization Policy. *Brookings Papers on Economic Activity*, 2011(1), 209-272.

McKay, A., & Wolf, C. (2022). Optimal policy rules in HANK. Working Paper, FRB Minneapolis.

Minton, R., Wheaton, B., & Lunch, H. E. P. (2022). Minimum Wages and the Rigid-Wage Channel of Monetary Policy.

Nakata, T. (2018). Reputation and Liquidity Traps. *Journal of Money, Credit and Banking*, 50(2-3), 299-334.

Poole, W. (1970). Optimal choice of monetary policy instruments in a simple stochastic macro model. *The Quarterly Journal of Economics*, *84*(2), 197-216.

Rowe, N. (2012, January 30). Fiscal Policy, and the Macroeconomics of Doing Nothing, Redux. *Worthwhile Canadian Initiative*.

Rubbo, E. (2023). Networks, Phillips curves, and monetary policy. *Econometrica*, *91*(4), 1417-1455.

Stokey, Nancy. (1989). Reputation and Time Consistency. *The American Economic Review*, 79(2), 134-139.

Stokey, Nancy. (1991). Credible Public Policy. *Journal of Economic Dynamics and Control*, 15(4), 627-656.

Sumner, S. (1993). Colonial currency and the quantity theory of money: a critique of Smith's interpretation. *The Journal of Economic History*, *53*(1), 139-145.

Sumner, S. (2021). *The Money Illusion: Market Monetarism, the Great Recession, and the Future of Monetary Policy*. University of Chicago Press.

Taylor, J. B. (2016). Staggered Price and Wage Setting in Macroeconomics. In J. B. Taylor & H. Uhlig (Eds.), *Handbook of Macroeconomics* (Vol. 2, pp. 1009-1050). Elsevier.

Timiraos, N. (2019, June 4). Powell Suggests Fed Could Cut Rates If Trade Spat Escalates. *The Wall Street Journal*.

Werning, I. (2011). *Managing a liquidity trap: Monetary and fiscal policy* (No. w17344). National Bureau of Economic Research.

Wieland, J. F. (2019). Are Negative Supply Shocks Expansionary at the Zero Lower Bound? *Journal of Political Economy*, 127(3), 973-1007.

Wolf, C. K. (2021). Interest rate cuts vs. stimulus payments: An equivalence result (No. w29193). National Bureau of Economic Research.

Woodford, M. (2003). Interest and Prices: Foundations of a Theory of Monetary Policy. Princeton University Press.

Woodford, M. (2011). Simple Analytics of the Government Expenditure Multiplier. *American Economic Journal: Macroeconomics*, 3(1), 1-35.

Chapter 3

Transformative AI, existential risk, and asset pricing

This chapter is jointly authored with Trevor Chow and J. Zachary Mazlish.

1 Introduction

Background. Recent rapid progress in generative artificial intelligence has highlighted the possibility that humanity may soon develop "transformative AI": AI technology that precipitates a transition comparable to the agricultural or industrial revolutions. Leading research labs like OpenAI and Google DeepMind bluntly declare their mission to build "artificial general intelligence" that can perform at or above human level on all tasks (OpenAI 2023; DeepMind 2023). The possibility of relatively short timelines for AGI is taken seriously by leading machine learning researchers, who in a 2023 survey gave a 10% chance that by 2027 AI will outperform humans at all tasks and a median forecast for such capability by 2047 (Grace, Stewart, et al. 2024).

The prospect of such transformative AI is a "double-edged sword", in the language of Jones (2023). On the one hand, continued AI innovations like those which have occurred in protein folding or text generation could accelerate economic growth and improve well-being. In the same way that growth increased by roughly an order of magnitude with the industrial revolution, some have predicted that transformative AI automating all tasks would increase growth by another order of magnitude, with GDP growth rising to 30% or more per year (Davidson 2021). Indeed, standard models of economic growth extended to include human-level AI can predict even economic singularities: infinite output in finite time (Aghion, Jones, and Jones 2018; Trammell and Korinek 2020).

On the other hand, many in the AI research community and in the broader public are concerned that such powerful AI technology could create severe risks, even an "existential risk" for the human species. This concern is driven by the challenge of ensuring that smarter-than-human AI technology pursues goals matching human values, rather than pursuing unintended and undesirable goals: the "AI alignment problem" (Ngo 2022; Yudkowsky 2016). The 2023 survey of machine learning researchers found that – among those who chose to respond – the median believed there to be a 5% chance that human-level AI results in "human extinction or similarly permanent and severe disempowerment of the human species" (Grace, Stewart, et al. 2024). This scenario is referred to as *unaligned* AI, in contrast to the growth-enhancing scenario with *aligned* AI.

Most economists, meanwhile, have been notoriously less likely to agree that transformative AI will be developed soon, less optimistic that aligned AI would radically accelerate economic growth, and less pessimistic that unaligned AI could pose an existential risk to human survival, on average (Korinek et al forthcoming).

This paper. We study the implications of transformative AI for asset prices and show how financial market prices can be used to forecast the arrival of such technology. In particular, we show that the prospect of transformative AI would predict a large increase in *real interest rates*, and would do so under expectations of either aligned or unaligned AI. As a result, to the extent that financial markets are efficient information aggregators, the level of long-term real interest rates can be used to help forecast the development of transformative AI.

This predicted rise in real interest rates is a basic implication of all modern asset pricing models, and is simply an application of the logic of consumption smoothing. Consider the case of *aligned* transformative AI: with the prospect of growth-induced high consumption and low marginal utility in the future, agents today would want to save less or borrow more, pushing up real interest rates at the relevant maturity. Similarly, if the market were forecasting future AI to be *unaligned* and to extinguish humanity, then there would also be no desire to save for the future (due to future extinction), again pushing up real interest rates at the appropriate maturity.

Empirical results on real rates. We offer new empirical evidence confirming that indeed, in the data, higher future growth increases real interest rates.

Measuring real interest rates is challenging. Existing work estimates real interest rates by using the nominal yields on nominal bonds and attempting to construct a measure of expected inflation to subtract from the nominal yields. The estimation of expected inflation needed for this, however, is difficult. We tackle this difficulty in two ways.

First, we use real yields from *inflation-linked bonds*, which provide a cleaner direct measurement of real rates compared to estimation based on nominal yields. To our knowledge, prior literature on the topic has not used real rates from inflation-linked bonds only because these bonds are comparatively new, with 20 or 30 years of data available.

Using these real yields directly from inflation-linked bonds, we show that higher real rates today indeed predict higher future GDP growth. Figure 1 shows the correlations for the US, UK, Australia, and Canada at the 10-, 15-, and 20-year horizons, comparing real interest rates over the relevant horizon with future GDP growth at the same horizon. While this data is merely correlational, and the data points are not independent of each other, it is suggestive evidence that growth and real interest rates are significantly linked.

Second, we use rich survey data on long-term inflation expectations from across 89 countries over the last 30 years to construct real interest rates from nominal bonds. The survey data is a unique dataset of forecasts from professional forecasters collected by Consensus Economics. By using forward-looking forecasts of inflation – rather than backwards-looking statistical measures of expected inflation, as in much of the literature – we are able to construct a large panel of real interest rate data. Our results using

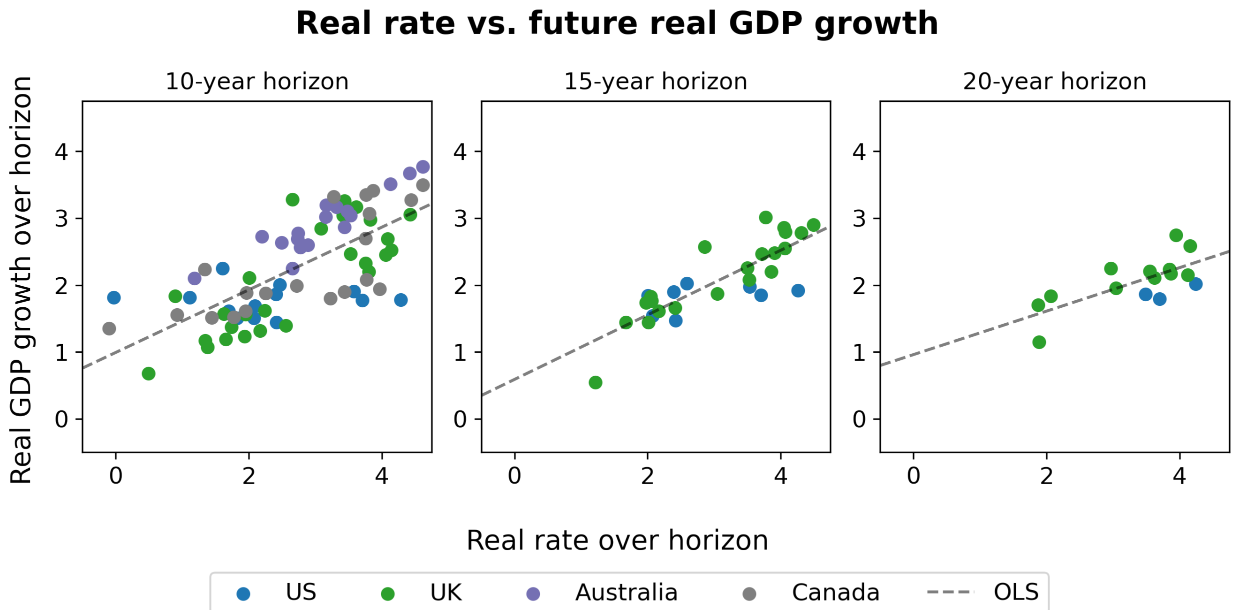

Figure 1: Real interest rates from inflation-linked bonds versus future GDP growth. Each subfigure plots a scatterplot of real interest rates of the titular maturity on the x-axis versus *future* annual GDP growth over the same horizon on the y-axis. Real interest rates are measured using yields on inflation-linked bonds on the last trading day of each year. The scatter plots show all available data up through 2022, for the US (since 1999), the UK (since 1985), Australia (since 1995), and Canada (since 1991), where the end date of the data depends on the time horizon. More details on data sources are given in section 3.

this data are similar to the results using inflation-linked bonds, and combined, provide the best evidence we are aware of about the link between ex-ante real rates and expected growth.

Other asset prices. We also briefly discuss the implications of transformative AI for other asset prices. We highlight that the implications of transformative AI for *equity* prices are much more ambiguous than for real interest rates. Among other issues, while the prospect of *aligned* AI leading to rapid growth may increase equity valuations, expectations of *unaligned* AI on the other hand would lower valuations. The net effect is qualitatively ambiguous, making stocks more difficult to use as a barometer for market expectations for AI timelines without taking a quantitative stand on magnitudes. Moreover, even whether higher expected future growth from aligned AI raises or lowers overall equity valuations is itself unclear, and depends critically on the intertemporal elasticity of substitution: larger future cashflows due to economic acceleration may be more than offset by the higher discount rate previously discussed. Even setting these two issues to the side, additionally it is not obvious that AI companies will capture *profits* from developing advanced AI – which is necessary for the expectation of AI to show up in equity prices – or that any companies which do capture profits are currently publicly traded. Finally, we also touch on the implications of transformative AI for land and commodity prices.

Outline. The structure of this paper is as follows. In section 2, we define the "transformative AI" scenario under consideration, and provide a brief overview of existing related work, which may be less familiar to many readers with a background in economics. We also briefly review the relevant and burgeoning literature on the economics of AI. In section 3, we demonstrate the simple result that growth and death risk raise real interest rates in a very broad set of models. Section 4 presents evidence that higher growth expectations raise real rates today, and offers some commentary on existing analysis of this topic. Section 5 reviews relevant literature finding that mortality and savings behavior is related. Section 6 discusses the implications of transformative AI for equities, land, and commodities. Section 7 concludes.

2 Defining transformative AI and relevant literature

In this section, we define the "transformative AI" scenario under consideration and provide context on existing research on the topic.[1] Much of this work may be unfamiliar to economists; familiar readers may wish to skip to section 3 after reviewing the definition of transformative AI in section 2.1, which is referenced throughout the paper.

2.1 Defining transformative AI

For the purposes of this paper, we consider the prospect of "transformative AI" as defined informally by Karnofsky (2016): artificial intelligence technology that has at least as profound an impact on the human trajectory as did the industrial revolution or agricultural revolution. As Karnofsky (2016) discusses, this term is similar to other concepts such as "artificial general intelligence" and "superintelligence", but is intended to be more inclusive – capturing technology which is transformative, even if such technology is not able to match all human abilities.

We operationalize this definition of transformative AI by dividing two cases.

Definition (Aligned transformative AI)**.** Aligned transformative AI is technology that causes growth in global GDP in excess of 30% per year.

Definition (Unaligned AI)**.** Unaligned AI is technology that causes the extinction of humanity.

Our definition of aligned transformative AI follows Davidson (2021), who defines "explosive growth" as growth in gross world product of at least 30%, i.e. an increase in

[1] Later sections also discuss relevant economics literature in the context of our analysis: section 4 reviews existing work on the relationship between real interest rates and growth; section 5 reviews existing work on the relationship between real interest rates and mortality or catastrophe risk; and section 6 reviews relevant work on other asset prices.

growth rates by an order of magnitude.[2] He discusses the possibility that transformative AI could cause such explosive growth. We take this as our benchmark for the effect of aligned AI, though given the unprecedented magnitude under consideration, these numbers should clearly be taken as rough approximations rather than as precise predictions.[3] Our definition of the unaligned AI scenario follows the literature on the topic, which is summarized in section 2.4.

2.2 Forecasting transformative AI

Analysis of the possibilities for artificial intelligence has a long history. Good (1965) originated the concept of an "intelligence explosion", a hypothesized phenomenon where AI systems gain the ability to improve their own algorithms and architectures, leading to recursive improvement and rapid increases in intelligence and power. Vinge (1993) originated and Kurzweil (2005) popularized the related concept of a "technological singularity", referring to an acceleration in technological progress occurring so quickly that it would be difficult to predict ex ante how the world would look after. While these earlier analyses were mostly speculative, rapid progress in machine learning over the last decade has resulted in analysis more grounded in the reality of modern AI.

Cotra (2020) provides an influential benchmark forecast for the development of transformative AI. Her framework is based on estimating the number of computations the human brain can perform per second. She then forecasts forward trends in the computational power of computers, using long-run trends like Moore's Law. She combines these to estimate the date by which computing power could match that of the human brain. Her analysis generates a distribution of estimates, with Cotra (2020) estimating a median arrival date of 2050 for transformative AI, and the updated analysis in Cotra (2022) forecasting a median of 2040. These estimates, however, are highly uncertain: the analysis of Cotra (2020) showed a 10% probability of transformative AI before 2030 and a 20% probability that transformative AI is not developed until after 2100.

Surveys of machine learning researchers are not too far off from the Cotra (2020) estimates. Grace, Salvatier, et al. (2018) survey 352 AI researchers on "when unaided machines can accomplish every task better and more cheaply than human workers" and find a median of 2061. Stein-Perlman, Weinstein-Raun, and Grace (2022) run an updated version of this survey with 738 respondents, and find a median of 2058 for the same question; Grace, Stewart, et al. (2024), in the latest iteration of the same survey with 2,778 published researchers, find a median of 2047. These results again come with significant dispersion.

Davidson (2023) uses a large-scale semi-endogenous growth model, a la Jones (1995), to forecast timelines for the development of transformative AI, and has a median forecast

[2] See also Hanson (2000).

[3] One alternative would be to use a definition in terms of task models (Zeira 1998). For example, aligned transformative AI could be defined as technology which can perform 100% of "cognitive" tasks that human perform, as in Davidson (2023). Because such technology would plausibly rapidly accelerate GDP growth – see Davidson (2023) for structural estimates – for the purposes of this paper the distinction in definitions is not important.

of 2043 for the development of transformative AI. This approach to forecasting the path of AI is analogous to the "dynamic integrated-climate economy" (DICE) modeling approach used in the climate literature: it is a computational integrated assessment model with an economics foundation.

Economists generally have been more cautious about forecasting the development of transformative AI. Korinek et al (forthcoming) survey economists and AI researchers about the probability of the development of "human-level machine intelligence". The median response of AI researchers in this survey was before 2050; for economists, the median response was after 2070.[4] Another survey by the Centre for Macroeconomics among European economists asked about the implications of progress in AI for global economic growth over the next decade. 64% of respondents answered growth would "increase (to 4-6%)", though in open-ended comments many of these respondents noted that they thought the increase would be smaller; 36% of respondents answered growth would "remain unchanged" (CFM 2023).

2.3 The economics of transformative AI

A small but important economics literature has analyzed the economics of transformative AI. A larger literature studies the economics of prosaic AI more broadly.[5]

The seminal contribution to this literature is Aghion, Jones, and Jones (2018), who consider a range of possible scenarios for the effects of artificial intelligence on economic growth. Of particular interest is their result that if AI automates tasks in the *ideas* production function (rather than the goods production function), then speeding up the rate of automation is equivalent to speeding up the rate of population growth. It is well-known that in semi-endogenous growth models, the rate of growth on the balanced growth path is proportional to population growth. Under some conditions, they show there can be a 'singularity' in the sense of reaching infinite output in finite time. They also highlight the critical role of potential bottleneck tasks in preventing (long-run) growth explosions.[6]

Trammell and Korinek (2020) offer a literature review of how the many permutations of different assumptions on the role of AI in growth models have differential implications for growth and for macroeconomic aggregates like the labor share. Their review highlights a wide range of possibilities, depending on what is assumed about the structure of production.[7]

[4]These results were sensitive to how the question was asked.

[5]See, for example: Agrawal, McHale, and Oettl (2018), Acemoglu and Restrepo (2018), Beraja, Kao, et al. (2023), and Brynjolfsson and McAfee (2014).

[6]Clancy (2022) offers a readable summary. Korinek and Stiglitz (2018), in the same volume as Aghion, Jones, and Jones (2018), analyzes how the development of AI could affect the income distribution (see also Korinek 2019; Korinek and Suh 2024).

[7]Erdil and Besiroglu (2023) offers another review. Nordhaus (2021) builds a model of one specific set of assumptions of these many permutations, and considers whether the empirical predictions of those assumptions are born out in the data, as an attempt to forecast "are we approaching an economic singularity?". Korinek (2019) analyzes the Malthusian implications of AI that can substitute for humans. Besiroglu, Emery-Xu, and Thompson (2022) show that if the capital share in the ideas production function increases,

2.4 The alignment problem and the economics of existential risk

Concern over risks from artificial intelligence technology are widespread not just among the public and in fiction, but also among many scientists across many fields. This has recently been captured by the "Statement on AI Risk" signed by a long list of AI scientists and public figures, stating, "Mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war" (Center for AI Safety 2023).

The basic concern is that it may be technically challenging to successfully program artificial intelligence technology in such a way that it behaves in line with human values. Just as software bugs can have large negative consequences in more mundane computer systems, software bugs in very powerful artificial intelligence systems could have correspondingly impactful negative consequences.

Ngo (2022) provides an overview of this "AI alignment problem" from the perspective of modern deep learning methods. Bostrom (2014) provides a book-length treatment from just before the deep learning revolution. Yudkowsky (2016) provides a conceptual argument for why the task of ensuring agents which are more intelligent than humans will act in line with human values should be perceived as challenging. Karnofsky (2021) offers a comprehensive and updated summary of these arguments. There is limited analysis of the AI alignment problem from an economics perspective. Hadfield-Menell and Hadfield (2019), Gans (2018), and Ely and Szentes (2023) are three exceptions.

In the economics literature, an important set of papers has analyzed how we should think about the tradeoffs between technology which brings positive benefits but creates existential risks. The important work of Aschenbrenner (2020) builds a model of directed technical change, extending the work of Jones (2016), where society can invest either in technology that increases consumption or technology that reduces the risk of death. Aschenbrenner shows that under reasonable parameters, optimally, existential risk follows a Kuznets-style curve: first rising, as society values consumption, and then falling, as the diminishing marginal utility of additional consumption is outweighed by the benefit of lower existential risk. Trammell (2020) shows a similar result in the context of an exogenous growth model. Jones (2023) summarizes these frameworks. Acemoglu and Lensman (2023) study the optimal regulation of technology adoption when that technology poses a (non-existential) disaster risk, motivated by AI technology. They show that if adoption is irreversible, then the path of adoption should be gradual, taking a 'wait-and-see' approach.[8]

then the long-run growth rate also increases. They show that the capital share of deep learning for computer vision is substantially higher than the capital share for prior research technologies, and estimate that if the capital share of the economy-wide ideas production function rose to the level of that for computer vision, then growth would be three to eight times higher.

[8] See also Gans (2024), Beraja and Zorzi (2022), and Lehr and Restrepo (2022)

3 Real interest rates, growth, and mortality in theory

In this section, we demonstrate that real interest rates are connected to both expected economic growth and mortality across a broad range of modeling frameworks. The connection is driven by the same, simple economic logic across all modeling frameworks: higher expected growth and higher mortality risk both reduce the supply of savings, which pushes up real interest rates. We show this logic holds in the three classes of models, covering the modern asset pricing modeling frameworks:

(i) Representative agent models

(ii) Incomplete markets models

(iii) Overlapping generations models

Since these results are known in the literature, we focus on results and intuition, and refer interested readers to relevant papers for full derivations. We also consider extensions with behavioral frictions (rule-of-thumb behavior and household myopia); nonstandard preferences (recursive preferences and habit formation); and models with new goods, which create nonstationary utility functions.

We conclude the section by emphasizing how the relationship between real rates and growth depends critically on the time horizon. In the short run – at business cycle horizons – *real interest rates that are too high cause low growth*, due to nominal rigidities. In the long run, nominal rigidities fade, and *high growth causes high real rates*. Hence, our empirical analyses in sections 4 and 5 focus on long-term real rates.

3.1 Representative agent models

It is well-known that in the canonical infinitely-lived representative agent model that the real interest rate is closely tied to growth and death.

The Ramsey rule. In the deterministic case with time-separable utility over the level of consumption, the real interest rate has a particularly simple expression – the canonical "Ramsey rule":

$$r = \rho + \frac{1}{\sigma} g \tag{1}$$

Here, r is the real interest rate over some time horizon, $\rho > 0$ is the rate of pure time preference, $\sigma > 0$ is the elasticity of intertemporal substitution, and g is the growth rate of consumption. With σ usually calibrated somewhere between 0.2 and 2 – an issue to which we return – we see that higher growth implies a higher real rate. Higher existential risk shows up here as a higher rate of time discounting ρ, thus also implying a higher real rate.

Benchmark calibration under transformative AI. Consider briefly a benchmark calibration at the annual frequency with $\sigma = 1$ (log utility) and $\rho = 0.01$ for a back-of-the-envelope calculation. Then, a growth rate of 1% per capita would imply a real interest rate of 2% under the Ramsey rule – not far off the level seen in advanced economies today. Meanwhile, a transformative AI-induced growth explosion causing the growth rate g to rise to 30% (as defined in section 2.1) would raise real interest rates to 31%. This would be an unprecedentedly high level.

The Euler equation. The Ramsey rule analysis above importantly assumed away uncertainty, among other things. Consider an infinitely-lived household with expected, discounted, time-separable utility over the level of consumption. Denote the period utility function $u(C_t)$ where C_t is consumption and u has diminishing marginal utility, $u' > 0$ and $u'' < 0$. Denote the subjective discount rate as β, and the probability of death in period t as δ. In this representative agent framework, the probability of death δ is equivalent to the probability of extinction.

The resulting intertemporal optimality condition is the well-known Euler equation:

$$1 = \beta\delta\mathbb{E}_t\left[\frac{u'(C_{t+1})}{u'(C_t)}\right](1 + r_t) \tag{2}$$

Suppose the path of consumption does not adjust, as in an endowment economy, for simplicity.

First, observe that higher death risk causes a higher real rate. A higher death risk is a lower probability of surviving until the next period, δ. A lower δ in (2) requires a higher real rate r_t. The intuition is that a higher probability of death shifts in the willingness to supply savings.

Second, observe that higher consumption growth, all else equal, also raises the real rate. Consider a shock that increases next-period consumption in at least one state of the world and shrinks it in none of them. Then, due to diminishing marginal utility, expected marginal utility $\mathbb{E}_t[u'(C_{t+1})]$ decreases, requiring a higher real rate r_t.

A risky shock – one which increases next-period consumption in some states but lowers it in others – does not unambiguously increase the real interest rate. The Euler equation (2) shows that what matters is *expected* growth *in marginal utility* – i.e. growth expectations taken over the risk-neutral measure. Thanks to diminishing marginal utility, this means that low-growth states of the world are weighted more highly.

For example, if consumption growth is lognormally distributed around the mean g and variance Var, then:

$$r_t = \rho + \frac{1}{\sigma}g - \frac{1}{2\sigma^2}\text{Var} \tag{3}$$

Here, $\rho \equiv -\ln(\beta\delta)$ and σ is once again the intertemporal elasticity of substitution.

The fact that the real rate is decreasing in the variance term shows how a shock which increases expected consumption growth could still push down the real rate if the shock also increases the variance of growth sufficiently. This again is due to the fact that what matters is expected growth in marginal utility. Diminishing marginal utility ensures that low-growth states of the world are weighted more highly. We return to this when discussing inequality.

3.2 Incomplete markets models and heterogeneous agents

The analysis above of the representative agent model demonstrates the importance of savings and borrowing decisions for understanding the effect of expected growth and mortality on real interest rates. This suggests it is worth considering how including realistic borrowing frictions affects the analysis.

Werning (2015) provides a benchmark analysis. In a world where idiosyncratic income risk does not covary with aggregate output, assuming isoelastic utility, and taking the "zero-liquidity limit" so that all agents are hand-to-mouth, then the *slope* of the relationship between growth and the real interest rate is the same, but the *level* is lowered. The analog to the Ramsey equation (1) is:

$$r = \rho + \frac{1}{\sigma} g - \gamma_1 \tag{4}$$

All the terms are as before, with the addition of $-\gamma_1$. $\gamma_1 > 0$ reflects the idiosyncratic risk facing the "marginal saver", which is the agent who *most* wants to save. The slope of the relationship between real rates and growth is still governed by the inverse of the intertemporal elasticity of substitution. Thus the real rate still increases with growth g and the existential risk probability embedded in ρ.

Moving away from the Werning (2015) benchmark, if idiosyncratic risk does covary with aggregate output, the relationship is more complicated. Auclert, Rognlie, and Straub (2018) show that an analog of the Ramsey equation can be written, for a particular form of idiosyncratic risk, as:

$$r_t = \rho + \frac{1}{\sigma \gamma_2} \left[C_{t+1} - \gamma_3 C_t \right] - \gamma_1 \tag{5}$$

Recall C_t is aggregate consumption at time t. If $\gamma_2 = \gamma_3 = 1$, then (5) is the same as (4), but they diverge due to cyclical income risk. γ_2 is the cyclicality of income of the marginal saver: the elasticity of individual income to aggregate income. γ_3 is the ratio of average cyclicality of income across all types, relative to the cyclicality of the marginal saver. Holding all else equal, higher C_{t+1} unambiguously increases the real rate.

Summarizing. While the relationship between output or consumption growth and real rates in these models is more complicated, a positive shock to growth still causes higher real rates. An increase in mortality risk has the same effect on the real rate as previously.

3.3 Overlapping generations models

The overlapping generations (OLG) framework is closely related to the incomplete markets framework of the prior section. Consider a simple version of this framework, where each agent lives for two periods and has log utility. There is Cobb-Douglas production technology with capital share α, population growth of n, and exogenous Hicks-neutral productivity growth of g. Then it can be shown that the analog of the Ramsey rule is:

$$r = \rho + g + \gamma_4 \tag{6}$$

Here, the coefficient on growth is 1, since log utility implies that the elasticity of intertemporal substitution is 1. The new term is γ_4, which is a function of the capital share α and population growth n.[9] Once again, the slope of the relationship between the real interest rate and growth is governed by the intertemporal elasticity of substitution; and the relationship between the real rate and mortality risk is direct.[10]

3.4 Recursive preferences and habit formation

Flynn, Schmidt, and Toda (2023) study the relationship between consumption growth and real interest rates under recursive preferences, such as the form studied in Epstein and Zin (1991) and Weil (1989). They show that the relationship is again determined by the elasticity of intertemporal substitution, where this elasticity must be defined appropriately given the recursive nature of preferences. The relationship betwen real rates and existential risk is unaffected by recursive preferences.

Bhamra and Uppal (2014), Hamilton et al. (2016), and Dennis (2009) study the relationship between consumption growth and real rates under habit formation. They show that with internal habits, the real rate is increasing in consumption growth. On the other hand, consider the extreme case with external habit where utility is determined entirely by the difference between individual consumption and average consumption. In such a world, a rapid acceleration in growth that lifts the consumption of all equally would not lower future marginal utility at all, and would not provide any incentive to save less or borrow more today. The real interest rate would be unaffected by the prospect of *aligned* transformative AI under this assumption, though it would still rise under the prospect of misaligned, extinction-causing AI. However, to the extent that preferences are not *purely* based on external habit, then rapid growth caused by transformative AI would still raise the real rate. This discussion emphasizes the importance of whether transformative AI will decrease marginal utility, rather than growth rates per se.

[9] Population growth does not affect the real rate in the canonical representative agent model, unlike the OLG model. Baker, De Long, and Krugman (2005) discuss how under imperfect altruism, population growth increases the real rate even in the representative agent model.

[10] In general without more assumptions than we have made so far, the OLG framework can lead to multiple or degenerate equilibria (see Acemoglu 2009, ch. 9).

3.5 Myopic consumers

If all agents in the economy are fully myopic and do not recognize an impending acceleration in growth or extinction event, then real interest rates are unaffected by such prospects. However, even if *consumers* are fully myopic, as long as *financial markets* can foresee these events, then these prospects will be priced in to real interest rates. Dupraz, Le Bihan, and Matheron (2022) consider a model where consumers are myopic but financial markets are fully forward-looking.

3.6 New goods

Scanlon (2019) and Trammell (2023) both show that the introduction of new goods can keep marginal utility perpetually high, even as consumption grows without bound. In this case, there would not be any incentive to save less or borrow more today in response to higher expected growth. The real interest rate would be unaffected by the prospect of *aligned* transformative AI under this assumption, though it would still rise under the prospect of misaligned, extinction-causing AI.

3.7 Distinguishing the short and long run

The time horizon in question is critical to understanding the relationship between real interest rates and growth. In the short run, nominal rigidities play an important role, while they dissipate at a long enough horizon.

The business cycle literature shows that – under nominal rigidities – when real interest rates are "too" high, short-term growth is lower. Think about central banks raising the nominal policy rate. In a flexible price world, then all nominal prices and wages immediately jump to offset this increase. Inflation rises one-for-one with the nominal rate increase, leaving the real rate unchanged. However, in a world with nominal rigidities, then price or wage inflation cannot immediately jump to fully offset the increase in the nominal rate, so the real interest rate rises. Under standard assumptions about the transversality condition (Cochrane 2017), this jump in the real rate requires the *level* of consumption to jump down immediately, so that the *growth rate* of consumption can be higher, per the Euler equation (2).[11]

To be clear, at business cycle horizons, it is not high real rates *per se* which cause low growth, but real rates which are *higher than they would be* in a world without nominal rigidities. (The real rate in such a flexible world – the "natural rate of interest" – is of course unobservable, however.)

In the long run, on the other hand, it is thought that nominal rigidities relax as prices and wages have time to adjust, and it is higher consumption growth that causes higher real rates. Thus, in the short run, too high real rates cause low growth; in the long run, high growth causes high real rates.

[11] The textbook references on this topic under the canonical New Keynesian model are Galí (2015) and Woodford and Walsh (2005).

Due to this flipping of signs, empirical work must carefully distinguish between short-run analysis and long-run analysis. This is an important issue with existing literature, an issue to which we turn in the next section – on the relationship between real rates and growth in the data.

3.8 Summarizing

Real rates and growth. The common thread across models is: if growth lowers the marginal utility of consumption, then growth increases real interest rates. We showed that this holds broadly across models, and highlighted two ways in which it might not. First, for a shock which increases consumption in some states of the world but lowers it in others, real interest rates could fall depending on how these net out. Second, marginal utility could stay high even with rapid consumption growth if utility is a function of relative consumption (i.e. external habit) or if the introduction of new goods keeps marginal utility high.

Real rates and mortality risk. Across all models, higher expected mortality risk raises real interest rates.

4 Empirical evidence on real rates versus growth: r vs. g

In the last section, we presented theoretical intuition for why higher expected growth would result in higher interest rates: expectations for such high growth would lead people to want to save less and borrow more today. In this section, we provide some simple empirical evidence that the predicted relationship holds in the available data.

First, we use cross-country survey data from Consensus Economics to show that, indeed, when long-term growth expectations are higher, long-term real interest rates are higher. Furthermore, *changes* in growth expectations are positively associated with *changes* in real rates. Second, we show that "market expectations" are rational: when 10-year real rates are higher, subsequent 10-year *realized* growth is also higher. Finally, we present simple evidence that quasi-exogenous shocks resulting in higher growth expectations *causally* increase real rates.

Before presenting our results, we explain how our measures of real rates differ from those use previously used in the literature.

4.1 Measuring real rates

The traditional approach. Most bonds historically have been *nominal*, where the yield is not adjusted for changes in inflation. Therefore, the vast majority of research studying the relationship between real interest rates and growth starts with nominal interest rates and attempts to construct real rates from the nominal rates.

Recalling that real interest rates are nominal interest rates minus expected inflation, such methodology requires estimating inflation expectations to subtract from the nominal rates. However, measures of historical inflation expectations do not exist for many countries or only have short histories – especially for measures of historical *long-term* inflation expectations. Therefore, most papers in this literature have attempted to construct ex-ante inflation expectations using available data, rather than using a direct measure of expectations.

Papers with this approach typically construct inflation expectations using a backwards-looking statistical model – usually simply a rolling AR(1) forecast based on past inflation. This is the approach used in the careful archival work of Schmelzing (2019); and in the analyses of Lunsford and West (2019) or Borio et al. (2022).

However, such an approach is inherently backwards-looking and fails to capture the forward-looking nature of inflation expectations. For example, consider the environment as of September 2023 where inflation was falling rapidly from its highs of the previous year. An AR(1) forecast using US data would assume that the year-over-year 3.7% CPI inflation rate was an appropriate forecast for the year ahead. However, more direct measurements of inflation expectations show substantially lower inflation expectations. For example, the Survey of Professional Forecasters shows a consensus inflation forecast of 2.7%. In short, crude autoregressive statistical models often diverge sharply from more direct measures of inflation expectations, for time periods when such ground truth is available.

Finally, another – sometimes severe – problem with measurement of real rates using historical bonds is credit risk. While modern sovereign bonds from countries like the US are closer to risk-free, this is not the case for all sovereign bonds, and especially was not always so historically. This is relevant, for example, in the long-run historical trends estimated by Schmelzing (2019). He estimates a steady long-run decline in real rates using historical sovereign nominal bonds. Besides also finding an explanation in declining time preference (Clark 2007; Stefanski and Trew 2022), this plausibly reflects a long-run decline in credit risk. For example, the estimates of Schmelzing show a sharp rise in real rates during the Napoleonic wars. It seems natural to suspect this reflects heightened credit risk on sovereign bonds during the conflict, rather than a true increase in risk-free real interest rates.

Our approach. We take a more direct approach.

For our primary analysis, we use two different sources to construct ex-ante real rates. 1) Where available, we use market real interest rates from *inflation-linked bonds*. These are bonds for which the payout is adjusted for realized inflation. As a result, the yields on these bonds directly reflect ex ante *real* interest rates, not nominal interest rates. This allows us to avoid needing to estimate inflation expectations.

2) If TIPS are not available, we use nominal rates less inflation expectations as measured in the Consensus Economics survey. Consensus Economics data covers 89 countries and directly asks professional forecasters who work at banks for their 10-year inflation

forecasts. As such, these survey expectations are a direct measure of the appropriate horizon inflation expectation, and no further assumptions are needed to construct ex-ante real rates.

To our knowledge, prior literature analyzing the determinants of real interest rates has not used data from inflation-linked bonds only because these bonds are comparatively new. In the United States, inflation-linked bonds (known as Treasury Inflation-Protected Securities, or TIPS) only began to be issued in 1997, for example. Many countries do not issue such bonds at all. The Consensus Economics data has been used in other work — for example, Engel and Rogers (2009) — but as far as we know, no other paper using this data has included a similarly large time span and country sample.

The yields on inflation-linked bonds do not perfectly reflect real interest rates because of various wedges, but these wedges are plausibly not too large. First, yields on such bonds also reflect term premia: risk compensation for uncertainty about the future path of real rates (Christensen, Lopez, and Rudebusch 2010). Additionally, such bonds are commonly thought to offer a (negative) convenience yield due to their relative illiquidity; and they have some embedded optionality due to deflation floors, among other complexities related to the bond structuring. See D'Amico, Kim, and Wei (2018), Christensen, Lopez, and Rudebusch (2010), and Fleckenstein, Longstaff, and Lustig (2014) for more. However, the literature finds that such wedges are not too large, and the advantage these bonds offer in terms of being direct measures of ex-ante real rates is substantial.

Inflation-linked bond data. For the US, we use the fitted real yield curve from Gürkaynak, Sack, and Wright (2007), as updated by the Federal Reserve and available online. These fitted rates are available at 10-year, 15-year, and 20-year horizons, since January 1999.[12] For the UK, we use the fitted real yield curve produced by the Bank of England. The 10-year and 15-year horizons that we use are available since January 1985; the 20-year horizon since June 1986.[13] For Australia and Canada, we use fitted 10-year real rates from Augur Labs. These data are available since 1995 and 1991, respectively. We also use Bloomberg's "generic 10-year inflation indexed bond" series for France, Israel, Sweden, Chile, Mexico, South Africa, Brazil, Japan, Germany, and Italy (with varying start dates).

Nominal bond data. When TIPS are unavailable and we subtract Consensus surveyed inflation expectations from nominal bonds, we use two different data sources. Some 10-year nominal rates come from Bloomberg; others come from the OECD's "long-term interest rates" database. Appendix X gives a precise breakdown.

When subtracting inflation expectations from our measures of 10-year nominal rates, the dates of Consensus surveys do not always perfectly align with the dates on which we have 10-year nominal rate data. We always subtract our inflation expectations from the

[12] Indeed, they are available at any yearly horizon up to 20 years.

[13] Since the 25-year horizon is only available since January 1998, there is not yet enough data to include it in our analysis.

closest possible measured rate, and only keep data points where the gap between survey and rate measure is less than one month.

4.2 Real rates and *expected* growth

All theoretical results presented above demonstrate the positive relationship between (risk-free) real interest rates and expected *consumption* growth. In addition to the inflation forecasts already discussed, Consensus Economics also asks for GDP growth and consumption growth expectations. For the results we present in the main body of this text, we use GDP growth expectations, rather than consumption growth. This is because the sample of GDP growth expectations is 34% bigger than the sample of consumption growth expectations and the two measures track each other closely (0.95 correlation). The fact that the two measures track each other so closely is consistent with a failure of international risk-sharing while an own-country aggregate Euler still holds. Appendix A1 shows that all our main results hold when using consumption growth expectations instead.

Consensus surveys of 10-year expectations — for GDP, consumption, and inflation — are conducted twice a year before 2014 and quarterly since then.

Given that our interest rate data does not always align with the exact dates of survey dates and that our default-risk controls are imperfect, our measures of risk-free rates around times of extreme volatility are noisy. To avoid outliers, we remove observations where the ex-ante real rate is > 10 pp or the change in the rate is > 10 pp. Our main results hold regardless, as shown in appendix A2.

Proceeding to the results, we run panel regressions of the following forms:

$$r_{i,t} = \alpha + \beta_1 \mathbb{E}_t(g_{i,t}) + \beta_2 X_{i,t} + \epsilon_{i,t} \tag{7}$$

$$\Delta r_{i,t} = \alpha + \beta_1 \mathbb{E}_t(\Delta g_{i,t}) + \beta_2(\Delta X_{i,t}) + \epsilon_{i,t} \tag{8}$$

The dependent variable is either the level or the change in country i's 10-year real interest rate; the primary independent variable of interest is either the level or the change in average expected GDP growth from five years ahead to ten years ahead. We use Δ to denote the change in a variable's value across Consensus survey dates and present results below using one, three, and five-year changes. Standard errors are Newey-West with appropriate lags to account for overlapping samples.

X is a vector of controls which includes three variables: the standard deviation of the Consensus five-to-ten year ahead growth forecast, Consensus surveyed average expected growth from zero-to-five years, and credit default swap (CDS) rates on the country's ten-year debt. Using CDS rates allows us to control for country default risk, an important issue even for advanced economies like the US, as shown by Chernov, Schmid, and Schneider (2020), which many other papers in the literature have neglected. CDS rates come from either Bloomberg or Longstaff et al. (2011).

The reason we use the "long-term" GDP growth forecast (five-to-ten year average) rather than simply ten-year average expected growth is because short-term expected growth can be confounded by monetary factors. In the short-run, expansionary monetary policy can lower real interest rates while increasing growth expectations. We isolate the relationship between long-term growth and real rates by using the five-to-ten year horizon expected growth. Controlling for zero-to-five year expected growth allows us to get rid of any business cycle dynamics that lead to correlation between short-term monetary driven growth and long-term growth expectations.

Table 1 shows the results from (7), where the regression is in levels. Figure 2 shows a raw scatterplot of the same data, illustrating the relationship absent any controls or fixed effects. Figure 3 aggregates across time to show country averages.[14]

Table 1: Expected growth vs. real rate

	Dependent variable: 10-yr real rate			
	(1)	(2)	(3)	(4)
5-10yr GDP forecast	0.69***	1.21***	1.41***	1.36***
	(0.04)	(0.07)	(0.10)	(0.09)
SD(5-10yr GDP forecast)			-1.18***	-0.54***
			(0.23)	(0.15)
0-5yr GDP forecast			-0.80***	-0.52***
			(0.09)	(0.07)
CDS			0.007***	0.005***
			(0.001)	(0.001)
Observations	2985	2985	2193	2193
Adjusted R^2	0.15	0.53	0.48	0.73
F-stat	248***	63***	185***	924***
Country FE	No	Yes	No	Yes
Note:	$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$			

The first column simply regresses the level of real rates on the long-term growth forecast; the second column adds country fixed-effects; the third column adds controls with no fixed-effects; and the fourth column uses controls and fixed-effects.

The primary result is that the coefficient on long-term growth expectations is uniformly positive and highly significant, with a magnitude greater than one in all specifications with controls or fixed effects. A coefficient of one would imply that when long-term GDP growth is expected to be one percentage point higher, real rates are correspondingly one percentage point higher.

All controls are also highly significant. The coefficient on the standard deviation of the long-term growth forecast matches the model in (3) where higher expected consumption volatility pushes down real rates. The coefficient on 0-5 year growth expectations is

[14]Observations where real rates are greater than 10% are trimmed from figures for visual readability (but included in regressions).

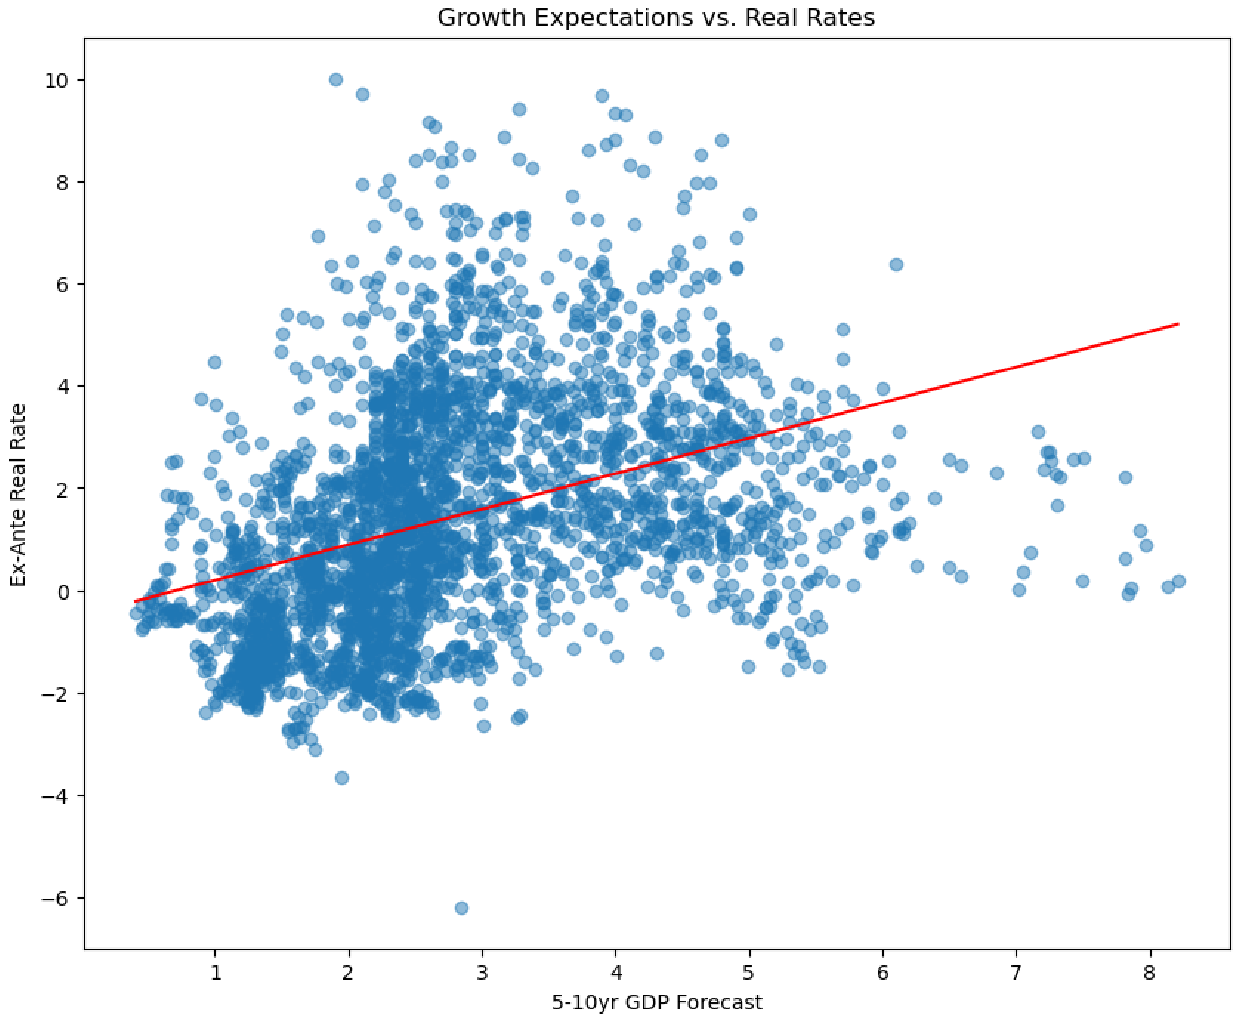

Figure 2: Ex ante real interest rates versus expected GDP growth, both at 10-year horizons. Real interest rates are measured using inflation-linked bonds when available, otherwise using benchmark nominal interest rates minus expected inflation; expected inflation is measured using the consensus of professional forecasters from Consensus Economics. Expected GDP growth is also the consensus of professional forecasters. More details on data construction are given in the text.Appendix figure 5 shows the same, where real interest rates are adjusted for credit risk.

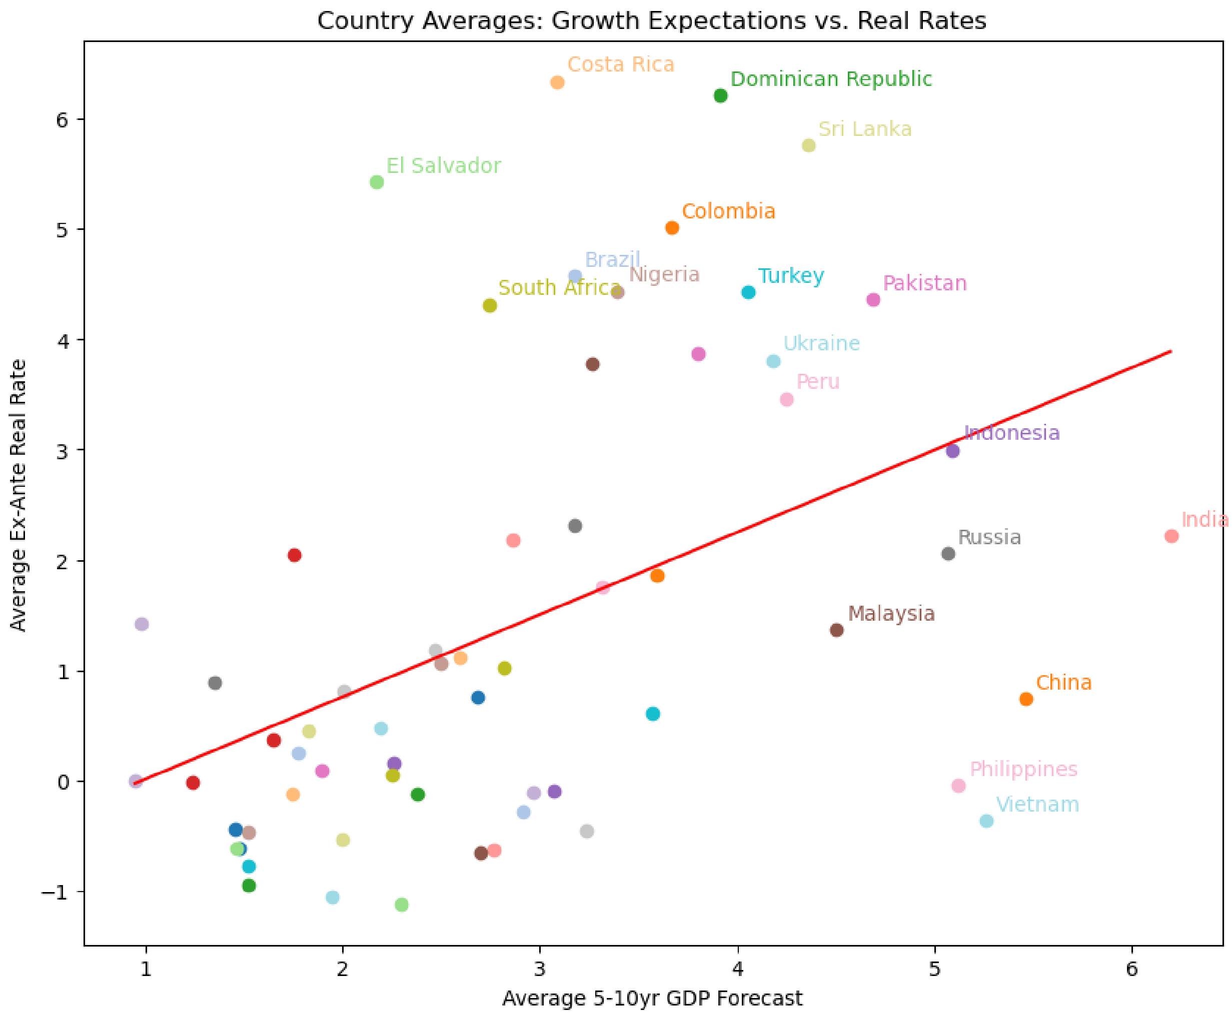

Figure 3: Average by country in sample: ex ante real interest rates versus expected GDP growth, both at 10-year horizons. Real interest rates are measured using inflation-linked bonds when available, otherwise using nominal interest rates minus expected inflation; expected inflation is measured using the consensus of professional forecasters from Consensus Economics. Expected GDP growth is also the consensus of professional forecasters. More details on data construction are given in the text.

negative, likely due to short-run monetary factors, as previously discussed. Finally, the coefficient on CDS represents that an 100 basis point higher CDS rate implies a 50-70 basis point higher ex-ante real rate; an economically large relationship.

The R^2 values are also quite large. Note that in column (3) we do not use any country fixed effects, but still explain almost half (48%) of the variation in ex-ante real rates, across over 60 countries and 2000 observations.

Rather than estimating things as a panel, in appendix A4, we run country-by-country regressions. Using the same controls as above we get that the median coefficient (across 61 countries) on long-term growth is 1.48, with 75% of individual country regression coefficients being positive. Many individual country samples are quite small, so we don't expect perfectly consistent results. Appendix A4 presents more details on these results.

In addition to our regressions using levels, we also estimate regressions in changes. Table 2 below presents the results when the change in dependent and independent variables are either 1, 3, or 5 year changes. Figure 4 plots the raw data, i.e. without controls, in the case of 5-year changes. One potential advantage of using changes is that it avoids stationarity concerns. Another advantage is that it more directly reflects our paper's question: how is a *change* in growth expectations reflected in real rates? The issue with estimating things in changes is that it reduces our sample size and is potentially biased by other sources of noise. For example, short-term liquidity issues in the TIPS market during times of crisis could cause measured real rates to rise while growth expectations are falling.

Similar to the question of changes versus levels, there are pros and cons to focusing on the shorter versus longer horizon changes. The advantage of looking at shorter horizon changes is that it gives us the largest sample. The advantage of looking at longer horizon changes is that they are more likely to purge the short-term noise issues just mentioned.

Table 2: Change in expected growth vs. change in real rates

	Dependent variable: Δ10-yr real rate		
	Δ1yr	Δ3yr	Δ5yr
Δ(5-10yr GDP forecast)	0.40^{**}	0.74^{***}	0.91^{***}
	(0.16)	(0.21)	(0.21)
ΔSD(5-10yr GDP forecast)	0.00	-0.35	-0.36
	(0.16)	(0.22)	(0.24)
Δ(0-5yr GDP forecast)	-0.42^{***}	-0.41^{***}	-0.46^{***}
	(0.08)	(0.12)	(0.13)
Δ(CDS)	0.001^{***}	0.001^{***}	0.004^{***}
	(0.000)	(0.001)	(0.001)
Observations	1911	1507	1157
Adjusted R^2	0.13	0.14	0.27
F Stat	12^{***}	9^{***}	25^{***}
Note:	$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$		

The first column of the above table presents results where independent variables are one-year changes; the second column with three-year changes; and the third column with five-year changes.

Once again, the coefficient on long-term growth expectations is always positive and significant. Its magnitude is noticeably smaller, but increasing in horizon. Both the three and five year change specifications include a point estimate of $\beta_1 = 1$ in their 95% CI. We only present results including controls here, since if the change in CDS is not controlled for, we get that $\beta_1 < 0$.

The coefficient on the change in the standard deviation of long-term growth forecasts is no longer significant, though it is consistently ≤ 0. We should note that our results only shed light on the relationship between *aggregate* risk and real rates. Idiosyncratic risk may have a much stronger negative relationship with real rates, as incomplete markets models would suggest.

Once again, the coefficient on the change in 0-5 growth expectations is negative and significant. We view our results as a potential explanation of the "puzzle" in Duffee (2023) where upward changes in one-year US GDP forecasts (from the Fed's Greenbook) are associated with downward changes in interest rates. Monetary factors account for this short-term inverse relationship, while traditional consumption smoothing logic dominates on longer-horizons. Appendix A3 presents further results in this vein, showing that in both levels and changes real rates are more strongly positively associated with 5-10yr growth expectations than 0-5yr or 0-10yr growth expectations.[15]

The coefficient on CDS remains highly significant, though of smaller magnitude, now implying that an 100 bp change in CDS is associated with a 10-40 bp change in ex-ante 10-year real rates.

Since we are regressing in changes, we do not use country fixed-effects, but still achieve meaningfully large R^2 values. Almost one-third of the variation in five-year changes in real rates is explained by our expectation variables and movements in default risk.

Appendix A5 shows that all results above about the sign and magnitude of the β_1 coefficient are robust to only looking at G7 countries. The coefficient on the regression in changes loses significance, however, which is likely due to the smaller sample.

Altogether, our results show a clear and reliable connection between higher long-term growth expectations and higher long-term real rates. We do not believe such a robust relationship has been shown before, and we see our wide cross-country sample — which uses either market-based inflation expectations or surveyed inflation expectations to construct real rates — as the best available evidence on this foundational macroeconomic relationship.

[15]In fact, in five-year changes, the positive relationship is only observed between 5-10 year growth expectations and real rates; the relationship is significantly negative for the two other horizon growth expectations.

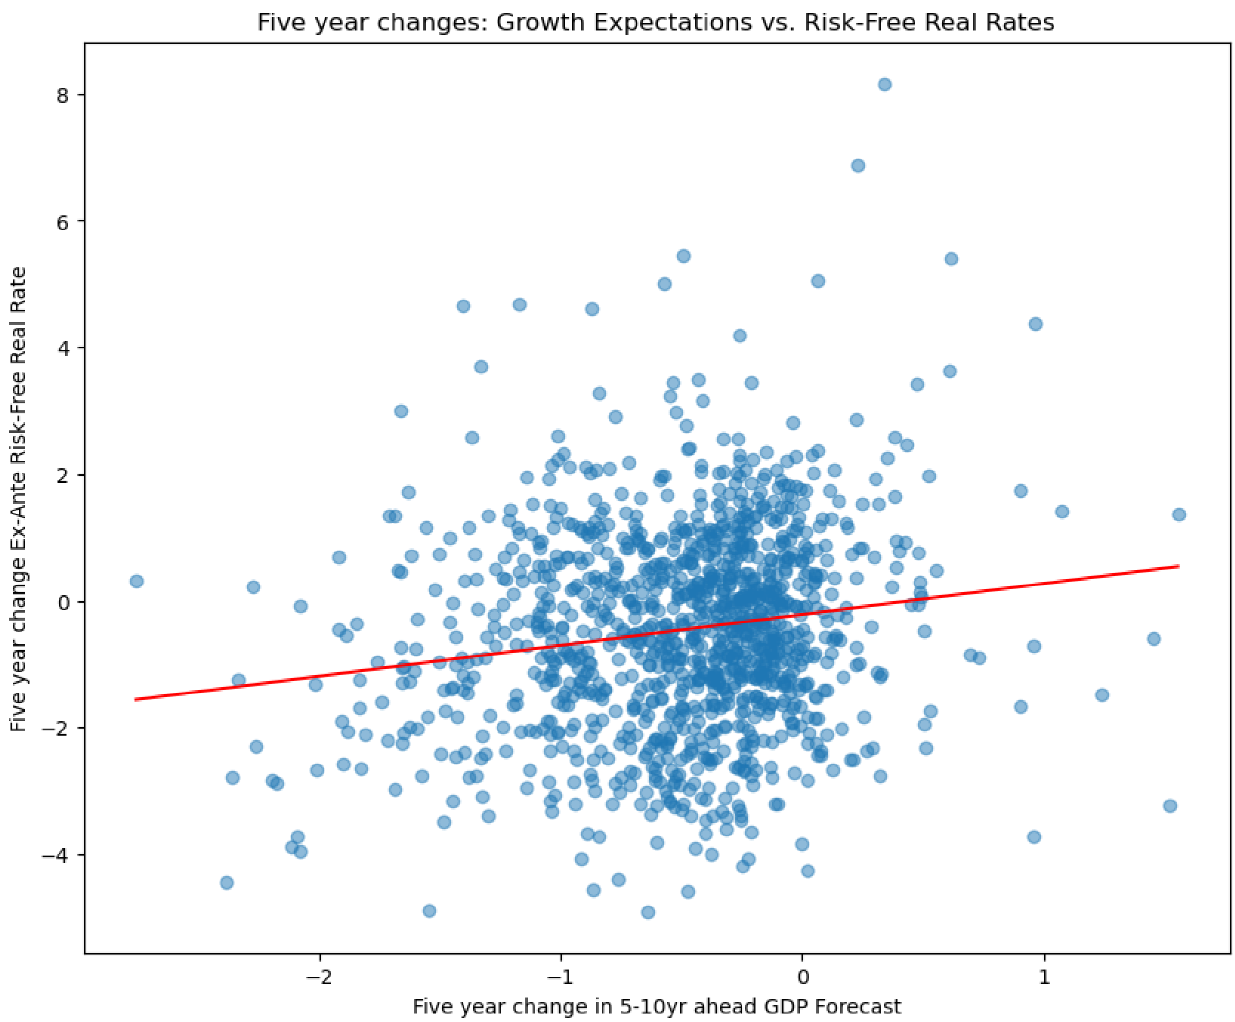

Figure 4: Ex ante real interest rates versus expected GDP growth, both at 10-year horizons. Real interest rates are measured using inflation-linked bonds when available, otherwise using benchmark nominal interest rates minus expected inflation; expected inflation is measured using the consensus of professional forecasters from Consensus Economics. Expected GDP growth is also the consensus of professional forecasters. More details on data construction are given in the text.

4.3 Real rates and *realized* growth

In this subsection, we present some brief correlational evidence showing that real rates and future *realized* growth are also linked in the available data. The link between real rates and *realized* growth relies on growth expectations representing rational forecasts. Therefore, given the above evidence that real rates respond to changes in expected growth, the evidence we now provide is evidence that those growth expectations were indeed rational.

In figure 1, we showed that for those countries with the longest-available real-rate data, there is an evident relationship between the real rate today and future realized GDP growth.

In table 3, to be consistent with the previous section, we run regressions of realized average 5-10 year later GDP growth ($g_{i,t}$) on the 10-year real interest rate today ($r_{i,t}$):

$$g_{i,t} = \alpha + \beta_1 r_{i,t} + \beta_2 X_{i,t} + \epsilon_{i,t} \tag{9}$$

Table 3: Realized growth vs. real rate

	Dependent variable: 5-10yr realized GDP growth	
	(1)	(2)
10yr real rate	0.28***	0.16**
	(0.04)	(0.07)
SD(5-10yr GDP forecast)		0.07
		(0.18)
0-5yr GDP forecast		-0.20***
		(0.05)
CDS		-0.001***
		(0.000)
Observations	1092	478
Adjusted R^2	0.56	0.75
F Stat	681***	981***
Country FE	Yes	Yes
Note:		$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$

Row 1 shows that, indeed, higher real rates today are significantly associated with higher realized long-term (5-10 years ahead) GDP growth, whether or not we include our previous batch of controls. The magnitude of the coefficient is smaller than that between real rates and expected growth. Such a difference is consistent with the fact that in a Euler equation framework, this coefficient is the IES, while the previous coefficient is the CRRA. However, we do not put too much emphasis on this interpretation, as the coefficients are by no means inverses of each other, and the fact that we use 5-10 year ahead growth makes this relationship a bit more complicated to disentangle.

In appendix X we show the robustness of these results to excluding Covid-19 from the sample, restricting to G7 countries, and using 0-10 year realized GDP growth.

5 Empirical evidence on real rates versus mortality risk

In section 3, we presented theoretical intuition for why higher expected mortality or existential risk would result in higher real interest rates: a heightened probability of death tomorrow would lead agents to want to save less and borrow more today. In this section, we present a review of already-existing work from a disparate set of literatures which provide evidence in support of the theory.

As a preliminary comment, we clarify that we study the relationship between real interest rates and the probability of truly *existential* risks – the probability of human extinction. We contrast this with the large literature on "rare disasters", which studies events that have a differential impact on risky assets like equities versus on risk-free bonds. "Disaster risk" thus provides a potential explanation for the equity premium puzzle (Rietz 1988; Barro 2006; Gourio 2008; Gabaix 2012; Pindyck and Wang 2013).[16] While disaster risk is about events that *differentially* affect stocks versus bonds, existential risk is about events that eliminate agents, thus affecting the return on stocks and bonds equally (and therefore cannot contribute to explaining the equity premium puzzle): existential risk sets the return on both to -100%.[17,18]

5.1 Mortality risk and savings behavior

In the theory reviewed in section 3, the mechanism by which higher expected mortality risk increased the real interest rate is by reducing savings. With higher probably of nonexistence in the future, agents have lower incentive to save for the future, and this reduced supply of savings increases the real interest rate.

In this subsection we provide evidence on the mechanism: we review existing work showing that reduced mortality risk causally increases savings (or equivalently, increases investment). While this does not provide direct evidence that extinction risk increases real interest rates, it does provide evidence for the hypothesized *mechanism* through extinction risk would increase interest rates.

[16] For a slightly different but related analysis, see Caldara and Iacoviello (2022) on measuring geopolitical risk.

[17] There is also a literature on the relationship between violent, non-existential conflict and asset prices. Hirshleifer, Mai, and Pukthuanthong (2023) use natural language processing techniques to study war discourse in newspaper articles and the relationship with equity prices. Ferguson (2008) and Bialkowski and Ronn (2017) provide narrative evidence of the effect of the world wars on financial markets. Rexer, Kapstein, Rivera, et al. (2022) study the relationship between violent conflict and sovereign nominal bonds. Leigh, Wolfers, and Zitzewitz (2003) as well as Wolfers and Zitzewitz (2009) study the relationship between the Iraq War and financial markets.

[18] Another literature that studies a non-existential disaster risk and financial markets is the climate literature. Giglio, Kelly, and Stroebel (2021) provide a review.

One example comes from testing for Huntington's disease, a disease which causes a meaningful drop in life expectancy to around 60 years, in Oster, Shoulson, and Dorsey (2013). Using variation in when people are diagnosed with Huntington's, the authors find that those who learn they carry the gene for Huntington's earlier are 30 percentage points less likely to finish college, which is a significant fall in their human capital investment – i.e., savings in the form of human capital investment decrease.

A second example comes from the informational experiment, in Malawi, of Ciancio et al. (2020). The authors provide information to correct pessimistic priors about life expectancy, and find that higher life expectancy directly caused more savings, via investment in agriculture and livestock.

Another set of papers study how the rollout of medical innovations, increasing life expectancy, led to increased savings and investment. Baranov and Kohler (2018) study the provision of a new AIDS therapy (also in Malawi) which caused a 13-year increase in life expectancy. Using spatial and temporal variation in where and when these therapeutics were rolled out, they find that increased life expectancy results in more financial savings and more human capital investment. Jayachandran and Lleras-Muney (2009) study the sudden drop in maternal mortality in Sri Lanka between 1946 to 1953. They find that for every additional year of life expectancy, years of education increase by 0.11 – i.e., savings in the form of human capital investment increased. Hansen (2013) and Hansen and Strulik (2017) argue that difference-in-difference evidence shows that improvements in antibiotics and cardiovascular disease treatment led to increased human capital investment, with a similar elasticity to the other studies.

Finally, there is tentative correlational evidence from surveys during the Cold War that a higher perceived risk of nuclear war was associated with a higher savings rate.[19] Russett and Slemrod (1993) find this in a 1990 survey data based on $n = 431$ American respondents. Slemrod (1982) as well as Russett, Cowden, et al. (1994) look at the timeseries correlation over the course of the Cold War between the U.S. private savings rate and the average of public opinion surveys on nuclear war risk (as well as the correlation with the Bulletin of Atomic Scientists "doomsday clock") and find positive correlations. Finally, Slemrod (1990) finds a suggestive negative correlation between the national savings rate and the survey average of perceived nuclear war risk in a cross-section of 19 OECD countries in the 1980s.[20] Contemporaneously, Heimer, Myrseth, and Schoenle (2019) find,

[19] We note that even a full-blown nuclear war, while the gravest catastrophe in history, need not be a true *existential* risk in the sense of wiping out the entirety of the human population. Besides a two-sided nuclear exchange possibly being limited, it is still a matter of scientific debate just how much damage such a war and the resulting nuclear winter would cause. Reisner et al. (2018) provides a full-scale analysis; Rodriguez (2019) offers an opinionated summary of the literature. As a result, the literature reviewed here on nuclear war risk is not necessarily comparable to the truly *existential* risk postulated by unaligned AI. It *may* be closer in nature to the "rare disasters" literature mentioned above.

[20] There is also work on the relationship between nuclear war risk and *equities*, with particular focus on the Cuban Missile Crisis. Finer (2021) studies the cross-section of US equities during the Cuban Missile Crisis. He compares companies with headquarters that are more or less exposed to Cuban missiles, as assessed by secret (at the time) intelligence assessments. He finds that the more exposed stocks fell by more during the crisis. Burdekin and Siklos (2022) study the Cuban Missile Crisis. They hand collect data on

cross-sectionally in US survey data, that pessimistic survival beliefs are correlated with a lower savings rate. This is true even after controlling for risk preferences, cognitive ability, and socioeconomic factors.

6 Other asset prices

In this section, we consider the possibilities for how transformative AI may affect asset prices other than real interest rates. Our main message is that the sign of the impact on real rates is much clearer the sign of the impact on other asset prices.

6.1 Transformative AI and equity prices

It may be tempting to use for forecasting AI timelines the market capitalization of companies like Alphabet (owner of DeepMind, a leading AI research lab) or that of chipmakers like Nvidia and TSMC. However, extracting AI-related expectations from stock prices is a challenging exercise for four reasons.

Aligned versus unaligned AI. First, and most importantly, AI-related companies will only have the possibility of high profits if transformative AI is aligned. Under *unaligned* AI where humanity is extinguished, the value of stocks along with everything else is converted to zero.

Profiting versus not. Second, it is not obvious that even in the aligned case that these companies will earn high profits. For instance, OpenAI has committed to a capped profit model, and other AI labs may sign on to a similar 'Windfall Clause' promising ex ante to donate profits beyond some threshold (OpenAI 2023; O'Keefe et al. 2020). Beyond corporate altruism, it is plausible that if a private company develops truly transformative AI technology, then the local government may nationalize and expropriate it (or at least attempt to do so) to distribute the benefits more broadly, preventing profits.

Public versus private companies. Third, when considering equity valuations, there is the question of which stock or stocks to consider. Critically, even if one takes a basket of tech companies and averages over them, then this only includes existing public companies. If the market expects transformative AI very soon, but only because it will be developed by a company which is not traded publicly (e.g. leading labs OpenAI or Anthropic) then this will not necessarily show up in any index of publicly-traded equities, depending on the affect of such technology on the distribution of firm profits.[21]

daily equity prices in Canada and Mexico, and together with US data, conclude "markets assigned a very small risk to the crisis leading to the use of nuclear arsenals".

[21] For example, although the development of the automobile transformed the United States, it has been argued that investing in public car companies in 1900 would not have been profitable (Locke 2021).

Higher growth may lower stock prices. Fourth, and quite importantly, it is not obvious whether expectations of transformative AI would raise or lower average equity prices. This is because stock prices reflect the present-discounted value of future profits; and transformative AI may raise those future profits, but – as emphasized throughout this paper – transformative AI would also raise the interest rate used to discount those profits. The net effect on average stock prices is ambiguous, without making more assumptions.

In particular, higher growth causes lower average stock prices if the *intertemporal elasticity of substitution* is greater than one, rather than less than one. This parameter – denoted as σ in section 3 – is subject to significant debate. In particular, while macroeconomics papers often calibrate to $\sigma < 1$, typically asset pricing papers calibrate to $\sigma > 1$. For example, Best et al. (2020) use bunching at mortgage notches to estimate $\sigma = 0.1$, and Crump et al. (2022) use directly-measured subjective expectations data to estimate $\sigma = 0.5$. These estimates would imply that stock prices fall strongly, on average, with news about higher future growth (and that real interest rates are very sensitive to changes in growth expectations).

6.2 The price of land and commodities

To the extent that advanced AI is able to substitute for labor but not for land or commodities in production, then the value of land and commodities could skyrocket in the case of aligned AI. However, this does require the auxiliary assumption about the shape of the production function – regarding the non-substitutability with land or commodities – which was not needed previously, and which is highly uncertain.

Additionally, again the value of land and commodities are (of course) directly sensitive themselves to real interest rates.[22] This complicates interpretation of their valuation for the same reason as stock valuations.

Finally, the value of land and commodities are hurt by the prospect of *unaligned* AI. As with equities, the net effect of higher valuation from the prospect of aligned AI versus lower valuation from the prospect of human extinction makes the prices of these assets difficult to use for forecasting AI timelines.

7 Conclusion

In this paper, we do not use any detailed inside knowledge of artificial intelligence technology to provide a forecast of the likely timeline for the development of transformative AI. That is, we do not present an 'inside view' on AI timelines (Kahneman 2011).

Instead, we argue that market efficiency provides an 'outside view' for forecasting AI timelines. The straightforward economic logic of intertemporal optimization, backed up by simple empirical evidence, shows that the prospect of transformative AI would

[22]Relatedly, Giglio, Maggiori, and Stroebel (2015) estimate 999-year risky, nominal discount rates using features of housing market contracts. See also Andersen (2022), Bracke, Pinchbeck, and Wyatt (2018), Fesselmeyer, Liu, and Salvo (2016), and Giglio, Kelly, and Stroebel (2021).

predict high long-term real interest rates. Such rates can be measured using the yields on long-term inflation linked bonds or by subtracting a measure of expected inflation from nominal bonds, and used to inform forecasts of transformative AI.

References

Acemoglu, Daron (2009). *Introduction to Modern Economic Growth*. Princeton University Press.

Acemoglu, Daron and Todd Lensman (2023). *Regulating Transformative Technologies*. Tech. rep. National Bureau of Economic Research.

Acemoglu, Daron and Pascual Restrepo (2018). "The race between man and machine: Implications of technology for growth, factor shares, and employment". In: *American Economic Review* 108.6, pp. 1488–1542.

Aghion, Philippe, Benjamin F Jones, and Charles I Jones (2018). "Artificial intelligence and economic growth". In: *The economics of artificial intelligence: An agenda*. University of Chicago Press, pp. 237–282.

Agrawal, Ajay, John McHale, and Alexander Oettl (2018). "Finding needles in haystacks: Artificial intelligence and recombinant growth". In: *The economics of artificial intelligence: An agenda*. University of Chicago Press, pp. 149–174.

Andersen, Carsten (2022). "Do distant rent flows matter? Inferring discount rates from leasehold apartments in Denmark". In: *Economics Letters* 219, p. 110798.

Aschenbrenner, Leopold (2020). "Existential risk and growth". In: *Global Priorities Institute. https://globalprioritiesinstitute. org/leopold-aschenbrenner-existential-risk-and-growth*.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub (2018). *The intertemporal keynesian cross*. Tech. rep. National Bureau of Economic Research.

Baker, Dean, J Bradford De Long, and Paul R Krugman (2005). "Asset returns and economic growth". In: *Brookings Papers on Economic Activity* 2005.1, pp. 289–330.

Baranov, Victoria and Hans-Peter Kohler (2018). "The impact of AIDS treatment on savings and human capital investment in Malawi". In: *American Economic Journal: Applied Economics* 10.1, pp. 266–306.

Barro, Robert J (2006). "Rare disasters and asset markets in the twentieth century". In: *The Quarterly Journal of Economics* 121.3, pp. 823–866.

Beraja, Martin, Andrew Kao, et al. (2023). "AI-tocracy". In: *The Quarterly Journal of Economics* 138.3, pp. 1349–1402.

Beraja, Martin and Nathan Zorzi (2022). *Inefficient automation*. Tech. rep. National Bureau of Economic Research.

Besiroglu, Tamay, Nicholas Emery-Xu, and Neil Thompson (2022). "Economic impacts of AI-augmented R&D". In: *arXiv preprint arXiv:2212.08198*.

Best, Michael Carlos et al. (2020). "Estimating the elasticity of intertemporal substitution using mortgage notches". In: *The Review of Economic Studies* 87.2, pp. 656–690.

Bhamra, Harjoat S and Raman Uppal (2014). "Asset prices with heterogeneity in preferences and beliefs". In: *The Review of Financial Studies* 27.2, pp. 519–580.

Bialkowski, Jedrzej and Ehud I Ronn (2017). "Financial markets facing apocalypse in World War II". In: *The Journal of European Economic History* 46.2, pp. 9–49.

Borio, Claudio et al. (2022). "Why so low for so long? A long-term view of real interest rates". In: *International Journal of Central Banking* 18.3, pp. 47–87.

Bostrom, Nick (2014). *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press.

Bracke, Philippe, Edward W Pinchbeck, and James Wyatt (2018). "The time value of housing: Historical evidence on discount rates". In: *The Economic Journal* 128.613, pp. 1820–1843.

Brynjolfsson, Erik and Andrew McAfee (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. WW Norton & Company.

Burdekin, Richard CK and Pierre L Siklos (2022). "Armageddon and the stock market: US, Canadian and Mexican market responses to the 1962 Cuban missile crisis". In: *The Quarterly Review of Economics and Finance* 84, pp. 112–127.

Caldara, Dario and Matteo Iacoviello (2022). "Measuring geopolitical risk". In: *American Economic Review* 112.4, pp. 1194–1225.

Center for AI Safety (2023). *Statement on AI Risk*.

CFM (2023). *Artificial Intelligence and the Economy*. URL: https://cfmsurvey.org/surveys/artificial-intelligence-and-economy.

Chernov, Mikhail, Lukas Schmid, and Andres Schneider (2020). "A macrofinance view of US sovereign CDS premiums". In: *The Journal of Finance* 75.5, pp. 2809–2844.

Christensen, Jens HE, Jose A Lopez, and Glenn D Rudebusch (2010). "Inflation expectations and risk premiums in an arbitrage-free model of nominal and real bond yields". In: *Journal of Money, Credit and Banking* 42, pp. 143–178.

Ciancio, Alberto et al. (2020). "Mortality risk information, survival expectations and sexual behaviors". In.

Clancy, Matt (2022). *What if we could automate invention? Growth theory in the shadow of artificial general intelligence*. URL: https://www.newthingsunderthesun.com/pub/2ek4d4s3/release/3.

Clark, Gregory (2007). *A Farewell to Alms: A brief economic history of the world*. Princeton University Press.

Cochrane, John H (2017). "The new-Keynesian liquidity trap". In: *Journal of Monetary Economics* 92, pp. 47–63.

Cotra, Ajeya (2020). *Forecasting TAI with biological anchors*. https://www.lesswrong.com/posts/KrJfoZzpSDpnrv9va/draft-report-on-ai-timelines.

— (2022). *Two-year update on my personal AI timelines*. https://www.lesswrong.com/posts/AfH2oPHCApdKicM4m/two-year-update-on-my-personal-ai-timelines.

Crump, Richard K et al. (2022). "Subjective intertemporal substitution". In: *Journal of Monetary Economics* 126, pp. 118–133.

D'Amico, Stefania, Don H Kim, and Min Wei (2018). "Tips from TIPS: The Informational Content of Treasury Inflation-Protected Securities". In: *Journal of Financial and Quantitative Analysis* 53.1, pp. 395–436.

Davidson, Tom (2021). "Could advanced AI drive explosive economic growth". In: *Open Philanthropy*.

— (2023). *What a compute-centric framework says about AI takeoff speeds*. https://www.lesswrong.com/posts/Gc9FGtdXhK9sCSEYu/what-a-compute-centric-framework-says-about-ai-takeoff. Accessed on July 7th, 2023.

DeepMind (2023). *About DeepMind*. Accessed: 2023-07-07. URL: https://www.deepmind.com/about.

Dennis, Richard (2009). "Consumption habits in a new Keynesian business cycle model". In: *Journal of Money, Credit and Banking* 41.5, pp. 1015–1030.

Duffee, Gregory R (2023). "Macroeconomic news in asset pricing and reality". In: *The Journal of Finance* 78.3, pp. 1499–1543.

Dupraz, Stéphane, Hervé Le Bihan, and Julien Matheron (2022). "Make-up strategies with finite planning horizons but forward-looking asset prices". In.

Ely, Jeffrey C. and Balazs Szentes (2023). *Natural Selection of Artificial Intelligence*. Working Paper.

Engel, Charles and John H Rogers (2009). "Expected consumption growth from cross-country surveys: implications for assessing international capital markets". In: *IMF Staff Papers* 56.3, pp. 543–573.

Epstein, Larry G and Stanley E Zin (1991). "Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis". In: *Journal of political Economy* 99.2, pp. 263–286.

Erdil, Ege and Tamay Besiroglu (2023). "Explosive growth from AI automation: A review of the arguments". In: *arXiv preprint arXiv:2309.11690*.

Ferguson, Niall (2008). "Earning from history?: Financial markets and the approach of world wars". In: *Brookings Papers on Economic Activity* 2008.1, pp. 431–477.

Fesselmeyer, Eric, Haoming Liu, and Alberto Salvo (2016). "How do households discount over centuries? Evidence from Singapore's private housing market". In.

Finer, David Andrew (2021). *No Shock Waves through Wall Street? Market Responses to the Risk of Nuclear War*. The University of Chicago.

Fleckenstein, Matthias, Francis A Longstaff, and Hanno Lustig (2014). "The TIPS-treasury bond puzzle". In: *the Journal of Finance* 69.5, pp. 2151–2197.

Flynn, Joel P, Lawrence DW Schmidt, and Alexis Akira Toda (2023). "Robust comparative statics for the elasticity of intertemporal substitution". In: *Theoretical Economics* 18.1, pp. 231–265.

Gabaix, Xavier (2012). "Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance". In: *The Quarterly journal of economics* 127.2, pp. 645–700.

Galí, Jordi (2015). *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*. Princeton University Press.

Gans, Joshua S (2018). *Self-regulating artificial general intelligence*. Tech. rep. National Bureau of Economic Research.

— (2024). *How Learning About Harms Impacts the Optimal Rate of Artificial Intelligence Adoption*. Tech. rep. National Bureau of Economic Research.

Giglio, Stefano, Bryan Kelly, and Johannes Stroebel (2021). "Climate finance". In: *Annual Review of Financial Economics* 13, pp. 15–36.

Giglio, Stefano, Matteo Maggiori, and Johannes Stroebel (2015). "Very long-run discount rates". In: *The Quarterly Journal of Economics* 130.1, pp. 1–53.

Good, IJ (1965). “Franz L. Alt and Morris Rubinoff, ed.,” Speculations Concerning the First Ultraintelligent Machine”. In: *Advances in computers*, pp. 31–88.

Gourio, Francois (2008). “Disasters and recoveries”. In: *American Economic Review* 98.2, pp. 68–73.

Grace, Katja, John Salvatier, et al. (2018). “When will AI exceed human performance? Evidence from AI experts”. In: *Journal of Artificial Intelligence Research* 62, pp. 729–754.

Grace, Katja, Harlan Stewart, et al. (2024). “Thousands of AI authors on the future of AI”. In: *arXiv preprint arXiv:2401.02843*.

Gürkaynak, Refet S, Brian Sack, and Jonathan H Wright (2007). “The US Treasury yield curve: 1961 to the present”. In: *Journal of monetary Economics* 54.8, pp. 2291–2304.

Hadfield-Menell, Dylan and Gillian K Hadfield (2019). “Incomplete contracting and AI alignment”. In: *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 417–422.

Hamilton, James D et al. (2016). “The equilibrium real funds rate: Past, present, and future”. In: *IMF Economic Review* 64, pp. 660–707.

Hansen, Casper Worm (2013). “Life expectancy and human capital: Evidence from the international epidemiological transition”. In: *Journal of Health Economics* 32.6, pp. 1142–1152.

Hansen, Casper Worm and Holger Strulik (2017). “Life expectancy and education: evidence from the cardiovascular revolution”. In: *Journal of Economic Growth* 22, pp. 421–450.

Hanson, Robin (2000). “Long-term growth as a sequence of exponential modes”. In: *Working manuscript*.

Heimer, Rawley Z, Kristian Ove R Myrseth, and Raphael S Schoenle (2019). “YOLO: Mortality beliefs and household finance puzzles”. In: *The Journal of Finance* 74.6, pp. 2957–2996.

Hirshleifer, David, Dat Mai, and Kuntara Pukthuanthong (2023). *War Discourse and the Cross Section of Expected Stock Returns*. Tech. rep. National Bureau of Economic Research.

Jayachandran, Seema and Adriana Lleras-Muney (2009). “Life expectancy and human capital investments: Evidence from maternal mortality declines”. In: *The Quarterly Journal of Economics* 124.1, pp. 349–397.

Jones, Chad (2023). “The AI Dilemma: Growth versus Existential Risk”. In.

Jones, Charles I (1995). “R & D-based models of economic growth”. In: *Journal of political Economy* 103.4, pp. 759–784.

— (2016). “Life and growth”. In: *Journal of political Economy* 124.2, pp. 539–578.

Kahneman, Daniel (2011). *Thinking, fast and slow*. macmillan.

Karnofsky, H (2021). “The “most important century” blog post series”. In: *Cold Takes*.

Karnofsky, Holden (2016). *Some Background on Our Views Regarding Advanced Artificial Intelligence*. https://www.openphilanthropy.org/research/some-background-on-our-views-regarding-advanced-artificial-intelligence/.

Korinek, Anton (2019). *The rise of artificially intelligent agents*. Tech. rep. working paper, University of Virginia.

Korinek, Anton and Joseph E Stiglitz (2018). "Artificial intelligence and its implications for income distribution and unemployment". In: *The economics of artificial intelligence: An agenda*. University of Chicago Press, pp. 349–390.

Korinek, Anton and Donghyun Suh (2024). *Scenarios for the Transition to AGI*. Tech. rep. National Bureau of Economic Research.

Kurzweil, Ray (2005). "The singularity is near". In: *Ethics and emerging technologies*. Springer, pp. 393–406.

Lehr, Nils Haakon and Pascual Restrepo (2022). *Optimal Gradualism*. Tech. rep. National Bureau of Economic Research.

Leigh, Andrew, Justin Wolfers, and Eric Zitzewitz (2003). *What do financial markets think of war in Iraq?*

Locke, Taylor (May 2021). "3 investing lessons Warren Buffett shared at the 2021 Berkshire Hathaway meeting". In: URL: https://www.cnbc.com/2021/05/03/investing-lessons-from-warren-buffett-at-berkshire-hathaway-meeting.html.

Longstaff, Francis A et al. (2011). "How sovereign is sovereign credit risk?" In: *American Economic Journal: Macroeconomics* 3.2, pp. 75–103.

Lunsford, Kurt G and Kenneth D West (2019). "Some evidence on secular drivers of US safe real rates". In: *American Economic Journal: Macroeconomics* 11.4, pp. 113–139.

Ngo, Richard (2022). "The alignment problem from a deep learning perspective". In: *arXiv preprint arXiv:2209.00626*.

Nordhaus, William D (2021). "Are we approaching an economic singularity? information technology and the future of economic growth". In: *American Economic Journal: Macroeconomics* 13.1, pp. 299–332.

O'Keefe, Cullen et al. (2020). "The windfall clause: Distributing the benefits of AI for the common good". In: *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, pp. 327–331.

OpenAI (2023). *About OpenAI*. Accessed: 2023-07-07. URL: https://openai.com/about/.

Oster, Emily, Ira Shoulson, and E Ray Dorsey (2013). "Optimal expectations and limited medical testing: Evidence from Huntington disease". In: *American Economic Review* 103.2, pp. 804–830.

Pindyck, Robert S and Neng Wang (2013). "The economic and policy consequences of catastrophes". In: *American Economic Journal: Economic Policy* 5.4, pp. 306–339.

Reisner, Jon et al. (2018). "Climate impact of a regional nuclear weapons exchange: An improved assessment based on detailed source calculations". In: *Journal of Geophysical Research: Atmospheres* 123.5, pp. 2752–2772.

Rexer, Jonah M, Ethan B Kapstein, Andres F Rivera, et al. (2022). *Pricing Conflict Risk: Evidence from Sovereign Bonds*. Tech. rep. Empirical Studies of Conflict Project.

Rietz, Thomas A (1988). "The equity risk premium a solution". In: *Journal of monetary Economics* 22.1, pp. 117–131.

Rodriguez, Luisa (2019). “How many people would be killed as a direct result of a US-Russia nuclear exchange?” In: *Rethink Priorities, June* 29.

Russett, Bruce, Jonathan Cowden, et al. (1994). “Did Americans’ expectations of nuclear war reduce their savings?” In: *International Studies Quarterly* 38.4, pp. 587–603.

Russett, Bruce and Joel Slemrod (1993). “Diminished expectations of nuclear war and increased personal savings: Evidence from individual survey data”. In: *The American Economic Review* 83.4, pp. 1022–1033.

Scanlon, Paul (2019). “New goods and asset prices”. In: *Journal of Financial Economics* 132.3, pp. 140–157.

Schmelzing, Paul (2019). “Eight centuries of global real interest rates, RG, and the ‘suprasecular’decline, 1311–2018”. In: *RG, and the ‘Suprasecular’Decline*, pp. 1311–2018.

Slemrod, Joel (1982). *Post-war capital accumulation and the threat of nuclear war*. Tech. rep. National Bureau of Economic Research.

— (1990). “Fear of nuclear war and intercountry differences in the rate of saving”. In: *Economic inquiry* 28.4, pp. 647–657.

Stefanski, Radoslaw and Alex Trew (2022). *Selection, Patience, and the Interest Rate*. Adam Smith Business School.

Stein-Perlman, Zach, Benjamin Weinstein-Raun, and Katja Grace (2022). “expert survey on progress in AI”. In: *AI Impacts. Available online at: https://aiimpacts. org/2022-expert-survey-on-progress-in-ai (accessed December 7, 2022)*.

Trammell, Philip (2020). “Existential Risk and Exogenous Growth”. In.

— (2023). “New Products and Long-term Welfare”. In.

Trammell, Philip and Anton Korinek (2020). “Economic growth under transformative AI”. In: *Global Priorities Institute. https://globalprioritiesinstitute. org/philip-trammell-and-anton-korinek-economic-growth-under-tra nsformative-ai*.

Vinge, Vernor (1993). “Technological singularity”. In: *VISION-21 Symposium sponsored by NASA Lewis Research Center and the Ohio Aerospace Institute*, pp. 30–31.

Weil, Philippe (1989). “The equity premium puzzle and the risk-free rate puzzle”. In: *Journal of monetary economics* 24.3, pp. 401–421.

Werning, Iván (2015). *Incomplete markets and aggregate demand*. Tech. rep. National Bureau of Economic Research.

Wolfers, Justin and Eric Zitzewitz (2009). “Using markets to inform policy: The case of the Iraq war”. In: *Economica* 76.302, pp. 225–250.

Woodford, Michael and Carl E Walsh (2005). “Interest and prices: Foundations of a theory of monetary policy”. In: *Macroeconomic Dynamics* 9.3, pp. 462–468.

Yudkowsky, Eliezer (2016). “The AI alignment problem: why it is hard, and where to start”. In: *Symbolic Systems Distinguished Speaker* 4.

Zeira, Joseph (1998). “Workers, machines, and economic growth”. In: *The Quarterly Journal of Economics* 113.4, pp. 1091–1117.

A Appendix

A.1 Consumption Growth Expectations

Table 4: Expected consumption growth vs. real rate

	Dependent variable: 10-yr real rate			
	(1)	(2)	(3)	(4)
5-10yr Consumption forecast	0.628***	1.002***	1.017***	1.118***
	(0.045)	(0.081)	(0.097)	(0.098)
SD(5-10yr Consumption forecast)			-0.905***	-0.110
			(0.142)	(0.121)
0-5yr Consumption forecast			-0.464***	-0.361***
			(0.078)	(0.077)
CDS			0.009***	0.007***
			(0.001)	(0.001)
Observations	2513	2513	1818	1818
Adjusted R^2	0.122	0.468	0.437	0.700
F Stat	191***	35***	126***	53***
Country FE	No	Yes	No	Yes
Note:		$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$		

Table 5: Change in expected consumption growth vs. change in real rates

	Dependent variable: Δ10-yr real rate		
	(1)	(2)	(3)
Δ(5-10yr Consumption forecast)	0.268***	0.499***	0.517***
	(0.103)	(0.153)	(0.170)
ΔSD(5-10yr Consumption forecast)	0.026	-0.194	-0.270*
	(0.099)	(0.131)	(0.143)
Δ(0-5yr Consumption forecast)	-0.269***	-0.138	-0.198**
	(0.078)	(0.085)	(0.090)
Δ(CDS)	0.005***	0.008***	0.006***
	(0.001)	(0.001)	(0.001)
Observations	1611	1304	1036
Adjusted R^2	0.167	0.240	0.222
F Stat	36.202***	36.226***	28.072***
Note:		$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$	

A.2 Including Outliers

Table 6: Expected growth vs. real rate

	Dependent variable: 10-yr real rate			
	(1)	(2)	(3)	(4)
5-10yr GDP forecast	0.763***	1.195***	1.796***	1.579***
	(0.054)	(0.103)	(0.113)	(0.116)
SD(5-10yr GDP forecast)			-0.741***	-0.523***
			(0.227)	(0.189)
0-5yr GDP forecast			-1.096***	-0.806***
			(0.095)	(0.108)
CDS			0.002***	0.002***
			(0.001)	(0.001)
Observations	3017	3017	2208	2208
Adjusted R^2	0.081	0.331	0.464	0.723
F Stat	202.816***	55.350***	168.140***	851.699***
Country FE	No	Yes	No	Yes
Note:		$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$		

Table 7: Change in expected growth vs. change in real rates

	Dependent variable: Δ10-yr real rate		
	(1)	(2)	(3)
Δ(5-10yr GDP forecast)	0.471***	0.753***	1.249***
	(0.163)	(0.223)	(0.273)
ΔSD(5-10yr GDP forecast)	-0.214	-0.328	-0.196
	(0.218)	(0.233)	(0.233)
Δ(0-5yr GDP forecast)	-0.414***	-0.526***	-0.857***
	(0.080)	(0.146)	(0.158)
CDS	0.001**	0.002***	0.001***
	(0.000)	(0.000)	(0.000)
Observations	1915	1516	1164
Adjusted R^2	0.162	0.286	0.291
F Stat	10.525***	13.272***	33.077***
Note:	$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$		

A.3 5-10 year growth vs. other horizons

Table 8: Horse race of growth horizons: levels

	Dependent variable: 10-yr real rate					
	(1)	(2)	(3)	(4)	(5)	(6)
5-10yr GDP forecast	0.69***	1.21***				
	(0.04)	(0.07)				
0-10yr GDP forecast			0.488***	0.861***		
			(0.041)	(0.066)		
0-5yr GDP forecast					0.300***	0.442***
					(0.037)	(0.062)
Observations	2985	2985	2985	2985	2994	2994
Adjusted R^2	0.145	0.534	0.076	0.462	0.034	0.413
F Stat	247.750***	63.337***	138.228***	70.910***	66.414***	80.736***
Country FE	No	Yes	No	Yes	No	Yes

Note: $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$

Table 9: Horse race of growth horizons: changes

	Dependent variable: $\Delta_5$10-yr real rate		
	(1)	(2)	(3)
Δ_5(5-10yr GDP forecast)	0.388***		
	(0.139)		
Δ_5(0-10yr GDP forecast)		-0.407***	
		(0.132)	
Δ_5(0-5yr GDP forecast)			-0.393***
			(0.094)
Δ_5SD(5-10yr GDP forecast)	-0.209		
	(0.251)		
Δ_5SD(0-10yr GDP forecast)		1.514***	
		(0.353)	
Δ_5(0-5yr GDP forecast)			1.495***
			(0.288)
Δ_5(CDS)	0.002***	0.001***	0.001***
	(0.001)	(0.001)	(0.000)
Observations	1162	1162	1166
Adjusted R^2	0.124	0.159	0.205
F Stat	4.530***	12.514***	20.393***
Note:	$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$		

A.4 Country-by-country regressions

The first parenthesis reports the percent of coefficients with the "correct" sign; the second reports the percent that are correctly signed and significant.

The "median" columns report median coefficients, observations per country regression, and adjusted R^2 per regression, and similarly the "mean" columns report means. There are 61 countries in the regressions.

Table 10: By country: expected growth vs. real rate

	Dependent variable: 10-yr real rate			
	Median	Mean	Median	Mean
5-10yr GDP forecast	1.40	1.41	1.48	1.92
	(75%) (62%)	(75%) (62%)	(75%) (64%)	(75%) (64%)
SD(5-10yr GDP forecast)			-0.76	-1.59
			(69%) (36%)	(69%) (36%)
0-5yr GDP forecast			-0.15	0.06
			(53%) (18%)	(53%) (18%)
CDS			0.004	0.003
			(66%) (44%)	(66%) (44%)
Observations	51	49	43	36
Adjusted R^2	0.34	0.32	0.56	0.55
Note:				$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$

In the below table, 58 countries appear in the 1-year change regressions, while 46 countries appear in the 5-year change regressions. The first-two columns report medians/means for 1-year changes; the last-two columns for 5-year changes.

Table 11: By country: change in expected growth vs. change in real rate

	Dependent variable: Δ10-yr real rate			
	Median Δ_1	Mean Δ_1	Median Δ_5	Mean Δ_5
Δ(5-10yr GDP forecast)	0.36	0.59	1.23	1.34
	(72%) (43%)	(72%) (43%)	(72%) (52%)	(72%) (52%)
ΔSD(5-10yr GDP forecast)	-0.30	-0.39	-0.32	0.46
	(62%) (26%)	(62%) (29%)	(59%) (30%)	(59%) (30%)
Δ(0-5yr GDP forecast)	-0.19	-0.21	-0.17	-0.26
	(69%) (26%)	(69%) (26%)	(65%) (30%)	(65%) (30%)
Δ(CDS)	0.003	0.005	0.005	0.005
	(79%) (52%)	(79%) (52%)	(70%) (54%)	(70%) (54%)
Observations	42	33	30	25
Adjusted R^2	0.27	0.29	0.41	0.38
Note:		$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$		

A.5 G7 regressions

Table 12: G7: Expected growth vs. real rate

	Dependent variable: 10-yr real rate			
	(1)	(2)	(3)	(4)
5-10yr GDP forecast	1.896***	2.574***	1.098***	2.420***
	(0.129)	(0.154)	(0.170)	(0.224)
SD(5-10yr GDP forecast)			-1.582	-1.171
			(1.052)	(0.761)
0-5yr GDP forecast			-0.673***	-0.666***
			(0.171)	(0.128)
CDS			0.008***	0.002
			(0.001)	(0.001)
Observations	598	598	297	297
Adjusted R^2	0.372	0.596	0.272	0.575
F Stat	215.876***	50.816***	34.311***	37.502***
Country FE	No	Yes	No	Yes
Note:		$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$		

Table 13: G7: Expected change in growth vs. change in real rate

	Dependent variable: Δ10-yr real rate		
	Δ_1	Δ_3	Δ_5
Δ(5-10yr GDP forecast)	0.758	1.023	0.395
	(0.685)	(0.750)	(0.740)
ΔSD(5-10yr GDP forecast)	-0.193	-0.202	-2.141***
	(0.623)	(0.781)	(0.631)
Δ(0-5yr GDP forecast)	-0.378***	-0.492***	-0.388**
	(0.130)	(0.152)	(0.189)
Δ(CDS)	0.006***	0.004*	0.002
	(0.001)	(0.003)	(0.002)
Observations	258	201	151
Adjusted R^2	0.182	0.095	0.100
F Stat	8.880***	6.139***	4.935***
Note:	$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$		

A.6 Additional figures

Figure 5: Ex ante real risk-free rates versus expected GDP growth, both at 10-year horizons. Risk-free rates are calculated using the real interest rates, described in figure 2, adjusted for credit risk as described in section 4. Observations where risk-free real rates are > 10 or < -5 are trimmed for visual readability.