Complexity Aversion: Macroeconomic Consequences*

Xavier Gabaix

March 26, 2026

Abstract

This paper explores the consequences of "complexity aversion" for macroeconomic outcomes, using a very tractable model. (i) As inflation and real-income volatility make planning more complex, complexity aversion generates first-order costs of inflation and the business cycle, which are much larger than the traditional second-order costs. (ii) This allows to calculate an optimal long run level of inflation, taking also into account that workers also dislike nominal wage cuts (in addition to real wage cuts). This yields an optimal inflation rate in the 1.5–4% range. Surprisingly, a zero lower bound does not change this optimal inflation much—it would multiply it by a factor of 1.1. (iii) These considerations in turn alter the optimal monetary policy rules, which should often ensure completely stable inflation, to the exclusion of other goals, except in times of very deep recessions.

*xgabaix@fas.harvard.edu. This is work in progress. Challenges, comments and suggestions solicited; don't hesitate to suggest missing citations, e.g. to your own papers. For helpful comments, I thank Abhijit Banerjee, John Cochrane, Gauti Eggertsson, Ben Enke, Ed Glaeser, Thomas Graeber, Brian Hall, Olivier Hart, Louis Kaplow, David Laibson, Giacomo Lanzani, Shengwu Li, Greg Mankiw, Florencia Marotta-Wurgler, Eric Maskin, Emi Nakamura, Indira Puri, Ricardo Reis, Josh Schwarzstein, Andrei Shleifer, Bennett Smith-Worthington, Kathy Spiers, Jeremy Stein, Jón Steinsson, Dmitry Taubinsky, Silvana Tenreyro, Dhara Yu, Ariel Zetlin-Jones, and participants at various seminars. I gratefully acknowledge financial support from Ferrante Economics Research Fund.

1 Introduction

In this paper, I explore the consequences for dynamic macroeconomics of "complexity aversion", building on the largely static microeconomic theory developed in Gabaix (2025). Inflation, and business cycle volatility, makes decision-making more "complex", some thing that the agents dread. This allows to propose a theory of the complexity cost of inflation and the business cycle. This also affects monetary policy. The main results are as follows.
The cost of income volatility is much higher than in the traditional model – first order rather than second order. Volatile income makes planning more error prone, which creates a first order loss from volatile income. This gives high costs of the business cycle, overcoming a paradox set out by Lucas (1987), that the cost of the business cycle are minuscule in the workhorse neoclassical model.

The cost of inflation is much higher – it is first order. One often mentioned cost is that inflation makes things complicated, as people need to juggle between prices that follow inflation (such as the price of most goods), and other streams that do follow it only infrequently (e.g. the wage, being adjusted to inflation only once a year, or a fixed-interest rate mortgage, that is not indexed). The theory leads to first order losses of inflation $|\pi|$, rather than, as in the traditional model (Woodford (2003)), second order losses π^2. In calibration, this is much bigger, and indeed much more in line with people's intuition.

This changes the optimal level of inflation. One can revisit the important issue of the optimal inflation target: it trades the first order complexity costs of inflation (which militates for an inflation of zero), the first order fact that workers are averse to nominal wages cuts; so that a bit of inflation is useful (and indeed, creates first order gains) in the labor market, and derive an optimal target. A rough calibration yields 1.5-4%. A chief payoff is to have a calibratable framework in which those estimates can be made and further refined.

*The ZLB doesn't change much the optimal level of inflation.*This optimal inflation is first derived in an economy without a zero lower bound. Initially, I thought that the threat of the zero lower bound (ZLB), i.e. the inability of the central bank to set nominal interest below zero, would increase considerably the level of optimal steady state inflation: the intuition was that having higher normal-times inflation hence nominal rates leaves "more room" to cut interest rates to 0. However, a calibration reveals that the ZLB does not change optimal inflation by much at all: it multiplies it by about 1.1. The reason is that the impact on interests and output gap on inflation is not so important quantitatively, and crises (warranting the entry in a very low interest rate zone) are rare enough, that the ZLB is not so important.

This changes monetary policy. We just saw that complexity aversion leads to "first order" rather than "second order" losses from inflation, and, under some conditions, of the output gap. Hence, the central bank's objective function has a loss of inflation proportional

to$|\pi_t|$ rather than π_t^2. This implies that exactly zero inflation is optimal in a large range of circumstances, even when this is at the cost of an output gap (this is in the spirit of the Lasso).

Related literature There is relatively little work on explaining why contracts are simple. One is Tirole (2009), which offers a cognitive interpretation of the incomplete contracts with renegotiation à la Hart (1995), but not why simple, static contracts are "simple". Another literature, on the lack of indexation, emphasizes asymmetric information: the person who proposes to index the contract may have more information than the one would has to accept or reject the indexation (Spier (1992), Hartman-Glaser and Hébert (2020)).[1] The present mechanism is very different, and centered on complexity, rather than asymmetric information: the model presented here predicts that even without asymmetric information, there will still be a lack of indexation. Carroll (2015) proposes conditions under which optimal contracts are linear, using robustness under the worse-case behavior, which is a form of infinite risk aversion. In contrast, this paper uses "smoother" preferences, and emphasizes complexity aversion.

On the theoretical side, this work relates to a number of attacks of complexity. Some are conceptual interesting, but sadly not very tractable, e.g. using discrete automata (Rubinstein, 1998), or sequences of discrete decisions (Gabaix et al., 2006). This paper is closer to the literature on limited attention, and crucially it adds a new, "complexity aversion" cost to the psychological description of the economic agent. I draw on earlier work (Gabaix (2014)), whose emphasis on "sparsity" I share — but the technology in that paper does not generally generate simple contracts. This sort of model is meant to be reasonable compromise between naive bounded rationality and fully Bayesian information processing. Other models of attention or skewed perception of the world include Bordalo et al. (2013), Gagnon-Bartsch et al. (2018), Schwartzstein and Sunderam (2021). At the opposite end, a fully Bayesian processing approach is developed in a series of papers by Sims (2003), Matějka and McKay (2015), Caplin et al. (2020), Woodford (2020), Khaw et al. (2021). The analysis we propose here could be applied to those styles of models. A cumulative body of papers show inattention (in particular in its behavioral flavor) is a good candidate for one unifying force in behavioral economics (see Gabaix (2019) for a survey/manifesto, and Enke et al. (2024) for new systematic evidence).

There is a recent line of work on complexity, which is largely empirical. A growing body of work measures how "intuitively complex" setups lead to more "mistakes", or more pre-

[1] Those asymmetric information theories predict "stickiness of institutions" rather than simplicity per se. For instance, if the social normal was, for some reason, to index contracts by 2/3 inflation, then any deviation from that would be "suspicious": that arrangement would be hard to get away from, and society would be stuck with indexing at 2/3 inflation.

cisely, deviations from a baseline simple economic model (Abeler and Jäger (2015), Martínez-Marquina et al. (2019), Oprea (2020), Dean and Neligh (2023), Enke et al. (2023)). Huck and Weizsäcker (1999), Sonsino et al. (2002), Fudenberg and Puri (2023), Enke and Shubatt (2023), Puri (2025) find lotteries are less tempting when they have more outcomes (so are more "complex"). Oprea (2024b) finds that this behavior may be linked to complexity rather risk per se. Those papers however do not trace this phenomenon to some "mistake aversion". It would be interesting to revisit them with that angle, especially as "complexity aversion" and mistake aversion might be (in part) a way to generate risk aversion.

It is clear intuitively a high complexity is a key determinants in mistakes, and this has been directly or indirectly shown in a number of papers, in all parts of economics. For instance, in macro, losses from lack of consumption planning can be very high (Augenblick et al. (2024)), and where the non-indexation of debt contract amplify financial crises (Gertler and Karadi (2011)).[2] In law and economics, contracts tend to be "simple", standardized and sticky (Hart (1995), Kahan and Klausner (1997), Bar-Gill (2012), Gulati and Scott (2012), Dari-Mattiacci and Marotta-Wurgler (2022)).[3] This is true for the "shape" of the contracts, but also in some important economic features. For instance, many contracts are "rigid in percentage space": Historically, most hedge funds have charged "2 and 20"— 2% of the assets under management, plus 20% of the upside, in a rigid manner. Likewise, sharecropping contracts tend to have a rigid split (often, 50%) between landlord and tenants (Young and Burke (2001)). In finance, firms and analysts often use simple, quite rigid rules of thumb to assess cost of capital, and valuations (Graham (2022); Gormsen and Huber (orth); Ben-David and Chinco (2024)).

This paper could be useful in those parts of economics, as a way model the taste for simplicity, and its equilibrium consequences.

Outline The paper first lays out the basic model, then explores successively how complexity aversion leads to first order costs of income fluctuations, hence of the business cycle, overturning the Lucas (1987) conundrum; the complexity costs of inflation, and optimal monetary policy. Most proofs are in the appendix.

[2] If debt contracts were indexed on aggregate economic activity, typically contracts would automatically channel funds to financial firms in crises, as in contingent convertible bonds—but we scarcely see that in practice, arguably because investors want simple, standard bonds.

[3] This literature emphasizes the "legal certainty" of using a conventional "boilerplate" contract—prior litigation has stabilized its interpretation; and the transaction costs of simply using an existing template.

2 Complexity aversion: overview

I summarize the complexity aversion developed in Gabaix (2025) in a largely static context, because stating it in a more dynamic one. I first review an existing model of thinking cost, then add the key new ingredient: first order complexity aversion.

2.1 Background material: A model of imperfect thinking

I first review model does not feature complexity aversion yet (it is laid out in Gabaix (2014) and Gabaix and Graeber (2024), with an easy review in Gabaix (2019), but the presentation here is self-contained). The task is to maximize over a continuous action a an objective function $u(a, x)$, where u is smooth (three times continuously differentiable) and concave in a. The vector of disturbances x is drawn from a distribution with mean normalized to 0, and its component are perceived by the agent to be uncorrelated. Action a is just a scalar for now, but it is easy to extend to a multidimensional a. I call $a^r(x) = \text{argmax}_a u(a, x)$ the rational action. I call default action $a^d = a^r(0)$ the optimal action "at the default", i.e. when all x are equal to 0, and I normalize it to 0; $a^r_{x_i}$ is the partial derivative at a default. The rational answer is thus (after linearization, so up to second order terms in x_i, assuming that the deviations x_i are "small")

$$a^r = \sum_i a^r_{x_i} x_i \tag{1}$$

The agent's objective is to cognitively construct those $y_i \coloneqq a^r_{x_i} x_i$, which indicates by how much dimension i of the problem should change the rational action. I assume that the prior mean for x_i is 0 (otherwise, the same reasoning applies to $x_i - x_i^d$, where x_i^d is a default value).

People receive noisy signals x_i^s about x_i:[4]

$$x_i^s = m_i x_i + \eta_i \tag{2}$$

where $m_i \in [0, 1]$ is the precision of the signal, and $\sigma^2_{\eta_i} = m_i (1 - m_i) \sigma^2_{x_i}$.[5] If all the shocks are jointly Gaussian, we have $\mathbb{E}[x_i | x_i^s] = x_i^s$. The rational case corresponds to $m_i = 1$, and the extremely inattentive case, to $m_i = 0$.[6] The s signifies "using subjective perceptions",

[4] This involves a choice of the "basis" or "frame" for the x_i's. Fortunately, this is recoverable from data on principle, see Section **??**.

[5] This comes from a more conventional "signal = truth + noise" framework, but here I reason directly with a post-shrinkage signal (so "signal = $m_i \times$truth + noise") which gives the functional form for $\sigma^2_{\eta_i}$. Section **??** gives the derivation.

[6] I do not assume that agents are Bayesian (as traditional information economics) – instead, I simply use that benchmark as an inspiration for the model (as e.g in Gabaix (2014), Woodford (2020)). For instance,

or sometimes "as a sparsity-seeking agent" (who economizes on attention m). Accordingly, I posit that if the decision maker sees x^s, the subjectively perceived value of x, she takes the decision $a^r(x^s)$. Linearizing:

$$a^s = \sum_i a^r_{x_i} x^s_i.$$

The traditional utility losses from an imperfect decision are then:

$$\mathbb{E}\left[u\left(a^r, x\right)\right] - \mathbb{E}\left[u\left(a^s, x\right)\right] = \frac{1}{2}\left|u''\left(a^d\right)\right| \sum_i \left(a^r_{x_i}\sigma_{x_i}\right)^2 (1 - m_i) \tag{3}$$

Cognition cost

Typically, in calculations, we will assume a production function of thought (or precision) as follows: to reach a precision m_i, some effort cost $C_i^{\text{Cognition}}(m_i)$ is needed. The total effort cost is then $C^{\text{Cognition}}(m) = \sum_i C_i^{\text{Cognition}}(m_i)$.[7] I normalize $C^{\text{Cognition}}(0) = 0$.

2.2 Complexity aversion

The consumer's full utility, given attention policy m, is:

$$V(m) = \mathbb{E}\left[u\left(a^s, x\right)\right] - C^{\text{Cognition}}(m) - C^{\text{CA}}(m) \tag{4}$$

where $a^s = a^r\left(x^s(x, m)\right)$ is the action given the perceived x (with $a^r(x)$ the rational action given x) and where $C^{\text{Cognition}}(m)$ is a cognitive cost function, which is the effort needed to achieve precision m.

The key new term in this paper is the "complexity aversion" or "mistake aversion" term $C^{\text{CA}}(m)$. It will is:

$$C^{\text{CA}} = \xi\left|\sigma_a u_{aa}\right| \sum_i \left|a^r_{x_i}\right| \sigma_{x_i}(1 - m_i) \tag{5}$$

I call $\xi \geq 0$ the parameter of complexity aversion; it is unitless, so potentially portable across settings.

Let me motivate it slowly, starting with a simple, intuitive version. Suppose that the agent feels she's making a error of size e (in absolute value: the error could be positive or negative): this is an error in action space, i.e. an sense of the distance between her action

when y_1 and y_2 are correlated, a Bayesian agent would use $\mathbb{E}\left[y_1 | y_1^s, y_2^s\right]$ to infer y_1, but instead we model the agent as being a "limited Bayesian", who simply performs $\mathbb{E}\left[y_1 | y_1^s\right]$.

[7] Obvious variants include functions of that sum.

and the optimal action. Then, the CA is

$$C^{\text{CA},1} = \bar{\xi} e$$

for some parameter $\bar{\xi}$ we'll detail shortly. This is "first order mistake aversion" (hence complexity aversion), because the penalty is proportional to e, rather than e^2.

Let us deepen this thought, first with just one disturbance x. The rational answer is (normalizing the default values to 0, taking Taylor expansion, and using only the leading term) $a^r = a^r_x x$; while the actual action is $a^s = a^r_x x^s$. So the average error is

$$a^\delta \coloneqq \mathbb{E}\left[a^r - a^s | x\right] = a^r \left(1 - m\right)$$

and the typical mistake size is[8]

$$e \coloneqq \sigma_{a^\delta} = \left|a^r_x\right| \sigma_x \left(1 - m\right)$$

Now, we let us turn to the units. In (4), C^{CA} is in utils, $[u]$ (where brackets mean "units of") while the error is in units of action, $[a]$. So, to keep the formulation as invariant as possible,I set $\bar{\xi} = \xi \kappa^u_a$, where ξ is a unitless that can be the same across problems, and κ^u_a fixes the units: it has the units of utils over action. To proceed at a high level of generality, there are not many choices. I propose to take $\kappa^u_a = |\sigma_a u_{aa}|$, where σ_a is the "natural scale" of a. Typically, it will be the typical "magnitude" of the action—e.g. the typical size of consumption of the good. In the analysis below, the specific value of σ_a will play almost no role, except to fix the orders of magnitude.

Hence, we arrive at the specification $C^{\text{CA}} = \bar{\xi} e = \xi \kappa^u_a \sigma_{a^\delta}$, hence:

$$C^{\text{CA}} = \xi \left|\sigma_a u_{aa}\right| \left|a^r_x\right| \sigma_x \left(1 - m\right) \tag{6}$$

The "degree complexity aversion" ξ is unitless. Later I will propose to calibrate it to $\xi \simeq 1/2$. Importantly for portability, C^{CA} is invariant if we change the units of action a, u, or x.[9] In that sense, we achieve a formulation that gives a a reasonable degree of invariance.

To build the key intuition, consider a very boundedly rational agent that has $m = 0$. Then, the cognition cost is 0 ($C^{\text{Cognition}}(m) = 0$ in (4)), but the dread cost C^{CA} is strictly positive (see (6)), as the agent is aware enough that he is making a mistake.[10] Hence, there is an incentive for the planner to create a "simple scheme", that will not create too

[8] One could also consider other measures, e.g. the average distance between the rational and behavioral action. This results in largely similar analytics.

[9] For instance, if we change x units of meters, to units of centimeters, then x is multiplied by 100, but $|a^r_x|$ is divided by 100, so that $|a^r_x| \sigma_x$ is invariant by the change in units of x.

[10] See Enke and Graeber (2023) for evidence that people have a good sense of their cognitive uncertainty.

much "dread", i.e. complexity aversion. Indeed, adding complexity will not guide the very boundedly rational agent towards greater allocative efficiency (as $m = 0$, so she doesn't respond to incentives), but it diminishes her utility, via the dread term. Adding complexity just adds deadweight loss.

On the other hand, when the consumer understands fully the situation (so $m_i = 1$), there is no room mistake, and the dread cost becomes 0 in (5): we recover the rational case, at least for "easy" decisions.

To go from a one-dimensional to a multidimensional x, we take (5): the complexity aversion is simply the sum over each dimension x_i.

The complexity cost (5) is a Lasso-like (Tibshirani (1996)) regularization penalty, that forces simplicity: this is the term $\left|a^r_{x_i}\right|$. It enriches this Lasso core with features that make it appropriate in economic situations. The term $(1 - m_i)$ ensures that complexity aversion is higher for "poorly understood" features—obvious or familiar features do not exhibit a complexity-aversion. The term $|\sigma_a u_{aa}| \sigma_{x_i}$ makes the penalty scaled correctly with incentives of the problem.

2.3 Dynamic and Multidimensional actions

The following is advanced material, and best skipped at a first reading. If $\tilde{a}$ represents the randomness of the action across situations, the complexity aversion is:

$$C^{\text{CA}} = \xi \sum_i \left(\mathbb{E}\left|\tilde{a} u_{aa} a_{x_i} x_i\right|^2\right)^{1/2} (1 - m_i) \tag{7}$$

If there are several dates, and utility is separable, it may be useful to sum separately across periods:

$$C^{\text{CA}} = \xi \sum_{i,t} \left(\mathbb{E}\left[\left|\tilde{a}_t u_{a_t a_t} a_{t,x_i} x_i\right|^2\right]\right)^{1/2} (1 - m_i)\,. \tag{8}$$

This is the formulation that I will use for macroeconomics.

3 The complexity costs of volatile income, and the high costs of the business cycle

Intuitively, business cycle costs, hence volatile incomes, makes things more complicated. This is a straightforward application of the general framework, but it is a good warm-up and it usefully overturns the conclusion of Lucas (1987).

Suppose that we have T periods, in an endowment economy. The endowment is $y_t =$

$\bar{y}(1+\hat{y}_t)$, where $\hat{y}_t$ has mean 0 and variance σ_y^2. The rational utility is:

$$V = \frac{1}{T}\sum_{t=1}^{T} u(y_t)$$

With a traditional utility, this is

$$W = V^{\text{trad}} = \mathbb{E}[V] = u(\bar{y}) - \frac{u'(\bar{y})\,\bar{y}}{2}\gamma\sigma_y^2$$

with $\gamma = -\frac{u''}{\bar{y}u'}$ relative risk aversion. To see the loss, define the proportionate loss as ℓ such that $V((1-\ell)\bar{y}) = V(\bar{y}, \sigma_\varepsilon)$. We verify the conclusion of Lucas (1987) that $\ell^{\text{Lucas}} = \frac{1}{2}\gamma\sigma_y^2$.

However, with complexity aversion, this is

$$W = V^{\text{trad}} - \xi\bar{y}^2\,|u''|\,\sigma_y(1-m_y) \tag{9}$$

where I use $\bar{y}$ for the typical size of consumption. The next proposition summarizes the situation.

Proposition 1. (Losses from the business cycle) *The consumption loss equivalent is:*

$$\ell = \ell^{Lucas} + \ell^{CA}, \qquad \ell^{Lucas} = \frac{1}{2}\gamma\sigma_y^2, \qquad \ell^{CA} = \omega_y\sigma_y, \qquad \omega_y \coloneqq \gamma\xi(1-m) \tag{10}$$

where ℓ^{Lucas} is the classic Lucas term (second order in σ_y), and ℓ^{CA} is the new complexity-aversion term (first order in σ_y).

Let us calibrate the effects. Conventionally, I also take $\sigma_y = 1\%$ and $\gamma = 4$. Lucas found $\ell^{\text{Lucas}} = 0.02\%$, which is minuscule, leading to the notion that business cycles do not matter.

For the complexity perspective, I take $\xi = \frac{1}{2}$, and $m_y = 0.8$ (as income is better perceived than inflation). So we get

$$\omega_y \coloneqq \gamma\xi(1-m) = 0.5$$

so a $\ell^{\text{CA}} = \omega_y\sigma_y = 0.4\%$ loss from the CA term. This is a very reasonable cost for business cycle with typical size 1%. The CA term is 20 times larger than the Lucas term. I conclude that the complexity perspective appears to give a straightforward resolution of the Lucas conundrum.[11]

In the online appendix (Section A.1), I detail some potential variants with idiosyncratic risks. With only large idiosyncratic risks, the loss becomes again quadratic in σ_y. But with

[11] I note that loss aversion would behave quite similarly. Hence, this resolution is not the only one – it's just natural from a complexity perspective.

enough agents with small or 0 idiosyncratic risks, the losses are "large" again, i.e. much greater that σ_y^2 (indeed, they're in σ_y^β for some $\beta \in [1,2)$).

4 The complexity costs of inflation, its implications for monetary policy

4.1 The complexity cost of inflation are first order, so much larger than the usual second order costs

There is a large literature on the welfare cost of inflation, which includes many traditional determinants. But typically, they are very small. For instance, a prominent cost (Woodford (2003)) is the fact that inflation adds dispersion between firm prices: for instance, if Coca-Cola last reset its nominal price τ period before Pepsi, and inflation is π, the relative price between Coca-Cola and Pepsi is off by $\pi\tau$; that distortion yields a welfare loss (Harberger triangle) order $(\pi\tau)^2$. This is typically very small (less than 0.1% per year).

However, one very plausible cost of inflation is that it adds "complexity" for the decision makers (Binetti et al. (2024)). Let us model it precisely.

Time t is continuous. The log price level is $p_t = \ln P_t$. I write $t = n + h$ with n an integer represents is "year" and continuous variable $h \in [0,1)$ is the "calendar date within the year). The log price increase continuously within the year, as $p_{n+h} = p_n + \pi_{n+1}h$, where inflation in year n is π_{n+1}. The agent gets a raise at the beginning of the year. Income is paid continuously as

$$y^{\$}_{n+h} = p_n + \frac{1}{2}\pi_{n+1} \tag{11}$$

as $\frac{1}{2}\pi_{n+1}$ compensates for the lack of raise within the year ($t \in [n, n+1)$). Inflation has expected value $\bar{\pi}$.

For starkness, let us assume that inflation is can known in advance: for instance, it is constant at $\pi_n = \bar{\pi}$. Here, a rational agent would face no cost of inflation.[12]

Figure 1 plots prices (for a constant inflation), and income. The infrequent adjustment of income creates an extra complexity, as income and prices do not move in sync — inflation makes real income less regular. This is also true even if inflation is constant. In that sense, inflation adds extra complexity.

For a rational agent, optimal nominal consumption is (in logs)

$$c^{\$}_{n+h} = y^{\$}_{n+h} - \left(\frac{1}{2} - h\right)\pi_{n+1}$$

[12]It would be easy to make inflation uncertain – then in the income process (11) we'd replace π_n by its expectation.

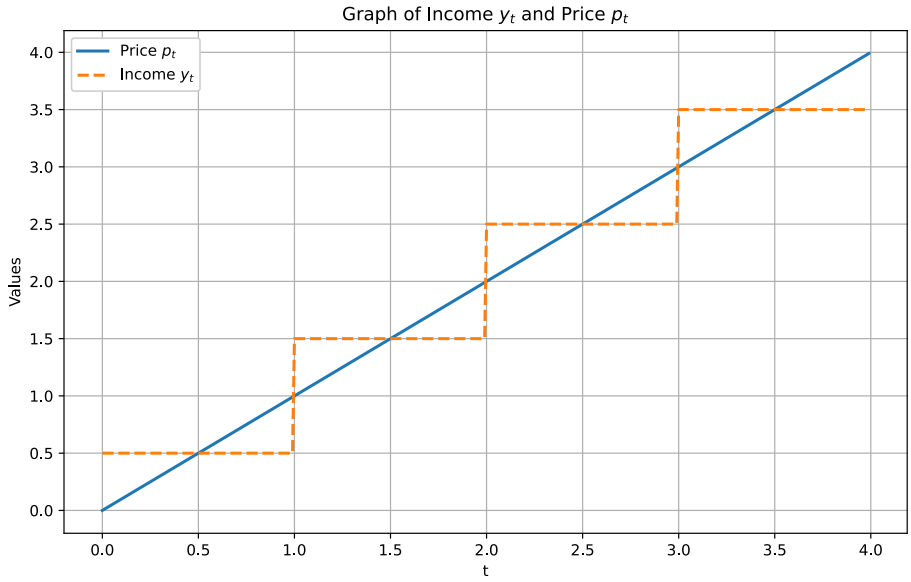


Figure 1: The price increase continuously, but the agents's nominal income increases only once a year (those quantities are in logs). That makes the planning process more complex, as the theory develops.

Hence, to the problem is added an extra complication: the π_{n+1} term. It captures the following: early in the year ($h \simeq 0$), agent should save: as he just got a raise, his real wage is high; later in the year ($h \simeq 1$), agent should dis-save: his real wage has been eroded by the inflation since the beginning of the year, so his real income is low.

The cost of inflation are the following. For simplicity, we suppose here that inflation is constant, at π.

Proposition 2. (Losses from inflation) *Suppose that nominal wages are readjusted only between periods of length τ. Call m the attention to inflation between wages increases, and assume first order complexity aversion. Then, the utility losses (expressed as a fraction of losses in permanent consumption) are:*

$$\ell^{CA} = \omega_\pi \left|\pi\right|, \qquad \omega_\pi = \frac{1}{4}\gamma\xi\tau\left(1 - m_\pi\right). \tag{12}$$

In particular, the complexity costs are higher than the variance costs iff $\pi < \frac{6\xi}{\tau} = \frac{6\times\frac{1}{2}}{1\text{yr}} = 300\%/\text{yr}$.

Because of first order complexity aversion, the loss is proportional to π, not π^2 as in the traditional model. So, for a consumer that typically doesn't pay attention to inflation ($m_\pi = 0$), then the loss is quite high, about ξ times $\tau\pi$. To calibrate ξ, let us use the simple

calibration in Section A.2. With $\xi = 0.5$ and $\gamma = 4$, and reset of wages every $\tau = 1$ year, $m_\pi = 0$, the losses are $\omega_\pi |\pi|$ with[13]

$$\omega_\pi = 0.5\text{yr}.$$

If annual inflation is $\pi = 4\%$, the loss is 2% — something sizable, while at $\pi = 20\%$,the loss is a large 10%. This is qualitatively consistent with empirical evidence that shows that people very averse to inflation even at low inflation rates in surveys, (Shiller (1997), Stantcheva (2024), Georgarakos et al. (2025)) and in elections (Di Tella et al. (2001)). This loss is much larger than the quadratic loss π^2 of the New Keynesian model, with both rational agents (Woodford (2003)), and behavioral but not complexity averse agents (Gabaix (2020)).

Complexity costs of income For consumption, we have likewise, as $\sum_t \beta^t \frac{\partial c_t}{\partial y_t} = 1$ (as income will need to be spent),

$$L = \omega_x |x_t|, \qquad \omega_x = \xi\gamma (1 - m_x)$$

Taking $m_x = 0.8$ (as people understand income better than inflation), we have $\omega_x = 0.5 \times 4 \times 0.2 = 0.4$. In summary, we provisionally calibrate:

$$\omega_\pi = 0.4\text{yr}, \qquad \omega_x = 0.4$$

and the welfare due to CA is

$$W_t^{CA} = -\omega_\pi |\pi_t| - \omega_x |x_t| \tag{13}$$

However, there is idiosyncratic risk so the extra penalty is actually, with $\bar{\varepsilon} \coloneqq \mathbb{E}[|\varepsilon|]$,

$$\mathbb{E}[|x_t + \varepsilon_{it}|] = \bar{\varepsilon} + f^\varepsilon(0) x_t^2$$

Hence, we get a quadratic extra term,

$$\Gamma_x = \Gamma_x^{\text{trad}} + \Gamma_x^{\text{CA}}, \qquad \Gamma_x^{\text{CA}} = 2 f^\varepsilon(0) \omega_x$$

For instance, for ε Gaussian distributed with standard deviation σ_ε,

$$\Gamma_x^{\text{CA}} = \sqrt{\frac{2}{\pi}} \frac{\omega_x}{\sigma_\varepsilon}$$

[13]Georgarakos et al. (2025) reports that people are on average willing to forgo 5% per year in annual consumption to eliminate inflation, which they perceive to be around 5%. This can be matched with a more aggressively behavioral calibration, say with $m = 0$, $\xi = 1$ and $\pi = 5\%$, which is the inflation often perceived by survey respondents. Then, the loss is $\ell^{\text{CA}} = 5\%$.

If we take an "inverse average" $\bar{\sigma}_\varepsilon \coloneqq \mathbb{E}\left[\frac{1}{\sigma_\varepsilon}\right]^{-1} = 3\%$, that gives an extra $\Gamma_x^{\text{CA}} = 0.8\frac{0.4}{3\%} = 10$. Given that the traditional $\Gamma_x^{\text{trad}} = \gamma + \phi \simeq 5$, we get $\Gamma_x = 15$. Hence, complexity aversion triples the quadratic losses from the business cycle.

4.2 Complexity-aware welfare in the New Keynesian model

4.2.1 Economic environment

I take a basic New Keynesian model (Woodford (2003), Galí (2015)). I use the behaviorally enriched version developed in Gabaix (2020), which allows for partially myopic agents, which is relevant for realism, and monetary-fiscal policy (e.g. McKay et al. (2016); Eichenbaum et al. (2026)). [14] It microfounds the following the behavior of the output gap x_t and inflation π_t:

$$x_t = M\mathbb{E}_t\left[x_{t+1}\right] - \sigma\left(i_t - \mathbb{E}_t\pi_{t+1} - r_t^n\right) \text{ (IS curve)}, \tag{14}$$

$$\pi_t = \beta M^f \mathbb{E}_t\left[\pi_{t+1}\right] + \kappa x_t + \nu_t \text{ (Phillips curve)}, \tag{15}$$

where $M, M^f \in [0,1]$ are the aggregate-level behavioral "cognitive discounting" parameters of consumers and firms, respectively. The rational case corresponds to $M = M^f = 1$.

Aggregate production is $Y_t = e^{\zeta_t} N_t$, where TFP level ζ_t, follows an AR(1), and N_t is labor supply. Equation (15) also allows for a "cost-push" shock ν_t, which could be microfounded as a markup shock (Galí (2015), Section 5.2), or an expectational error.

4.2.2 Welfare with complexity costs

In the traditional models (Woodford (2003)), the welfare function is:

$$W^{\text{trad}} = -K^{\text{trad}}\mathbb{E}_0\sum_{t=0}^{\infty}\frac{1}{2}\beta^t\left(\pi_t^2 + \vartheta^{\text{trad}}\left(x_t - x_*\right)^2\right) + W_-, \tag{16}$$

for positive coefficients K^{trad}, ϑ^{trad}.[15] The term x_* is zero is there are no distortions at the natural rate of interest (i.e., when removing monetary frictions). The losses come from

[14] This model is meant to be particularly tractable, and reasonably respectful of psychology. One could imagine also adding complexity costs of inflation, as modelled here, to other frameworks (e.g. McKay et al. (2016); Angeletos and Lian (2018b); Farhi and Werning (2019); Woodford (2019); Bilbiie (2020); Cochrane (2023)), and also enrich the expectation process (Gabaix and Laibson (2002); Mankiw and Reis (2002); Bordalo et al. (2020); Afrouzi et al. (2023); Mei and Wu (2024)). In addition, the model uses a Calvo friction, but one could use other pricing frictions (Nakamura et al. (2018)), e.g. a Calvo where adjustment frequency depends on inflation itself.

[15] Here $\vartheta^{\text{trad}} = \frac{\bar{\kappa}}{\varepsilon}$, $K^{\text{trad}} = u_c c\left(\gamma + \phi\right)\frac{\varepsilon}{\bar{\kappa}}$, and W_- is a constant, $\bar{\kappa}$ is the Phillips curve coefficient with rational firms, and ε is the elasticity of demand.

misallocation costs, due to inflation π_t (the relative price of firms is off), and output gap x_t (people don't work exactly the correct amount). In contrast, complexity aversion adds one more term.

Proposition 3. *In the New Keynesian model with complexity aversion costs, welfare is, to the leading order:*

$$W^{tot} = -\sum_t \beta^t C_t^{CA} + W^{trad}$$

where the complexity aversion costs at time t is:

$$C_t^{CA} = \omega_\pi |\pi_t| + \omega_x |x_t| . \tag{17}$$

for positive coefficients ω, and W^{trad} is the traditional welfare loss (16), with quadratic loses in inflation and output.

But for small inflations and output gap, first order costs dwarf second-order costs. Hence, the welfare is, to the leading order, only made of first order costs, $\sum_t \beta^t C_t^{CA}$.[16]

The planner wants to maximize $\sum_t \beta^t W_t$, where the loss at time t is

$$W_t = K\left(-\omega_\pi |\pi_t| + g\left(x_t\right)\right), \qquad g\left(x_t\right) = \gamma x_t - \frac{1}{2}\theta x_t^2 \tag{18}$$

where the first order costs come from complexity aversion, and $\theta = \frac{\vartheta^{\text{trad}}}{w_\pi} > 0$ is the relative weight on output gaps, and $\gamma = -x_* \frac{\vartheta^{\text{trad}}}{w_\pi}$. Empirically, almost everyone "feels" that booms are better than recessions, so $\gamma > 0$ seems the empirically relevant case.[17]

So now we have a someone different problem: optimal policy with a absolute value penalty on inflation, rather than a quadratic one. This apparently minor and technocratic difference leads to quite deep changes in the optimal dynamic monetary policy, as we next see.

This dynamic analysis has a partial static antecedent in Orphanides and Wilcox (2002), who propose an objective function $-\pi_t^2 - \theta |\hat{y}_t|$, so with first order loss in output (as a deviation from full employment), rather than in inflation. They see that as a reverse-engineered way to formalize the preferences of central bankers, who have a high distastes for output losses (and implicitly target unemployment). In contrast here, I provide complexity microfoundations (rather than first order losses because of unemployed labor), and analyze dynamic policies.

Before that, I provide some simple assumptions ensuring that the model boils down to (14)-(15). Those are more technical, and the reader may skip them. There is a continuum of agents i. Their nominal wage is reset every H periods, and a fraction $1/H$ of agents has

[16] For simplicity, only consumers are complexity averse, not firms.

[17] In the traditional NK model, this is dues to the fact that there positive markups at the steady state, which this microfoundation is not necessary

their wage reset each period. The agents don't fully appreciate inflation, as in Section 4.1. Then, we have the following reassuring result.

Lemma 1. *The aggregate behavioral model (14)-(15) still holds in the aggregate, despite wage staggering and partial nominal illusion by workers.*

4.3 What a non-zero inflation target should be, given aversion to nominal wage cuts, and the possibility of hitting the ZLB

What is the optimal target level of inflation, $\bar{\pi}$? So far, it would be simply zero. Sadly, the world is more complicated. The main traditional reason for to have inflation above zero is that workers exhibit an aversion to nominal (not just real) wage cuts, so a bit of inflation helps the avoiding nominal wage cuts, which would hurt "morale" in the labor market (Tobin (1972); Akerlof et al. (1996); Bewley (1999); Shafir et al. (1997); Kaur (2019); Grigsby et al. (2021)). I propose a simple model of this and calibrate it. Then, I derive the optimal target—providing a simple model for this, that perhaps calibrates better given the new complexity costs of inflation.

I then enrich the model by allowing for the possibility of hitting the zero lower bound on the nominal interest rate (ZLB). This provides an additional reason to target positive inflation—namely, to reduce the economic impact of the ZLB. However, based on a provisional calibration, I conclude that this is a relatively minor force that does not alter the main analysis. It raises the optimal inflation rate from about 2% to 2.2%.

4.3.1 Modelling workers' happiness from a higher nominal wage

So far, the optimal inflation target was zero. Now, the real world brings one more complication: workers have some nominal illusion, and in particular do not like nominal wage cuts. So, if the neoclassical real wage was to fall, with zero inflation, firms would not cut wages, and rather fire some workers. To avoid this, it is useful to have positive inflation.

To model this, let us say that workers morale affects raw log productivity ζ_t, to change it into effective log productivity:

$$\zeta_t^{\text{eff}} = \zeta_t + f_t$$

where the new term f_t captures the efficiency wage idea that disgruntled work less well, and indeed may sabotage a firm. We take the following formulation for their morale:

$$f_t = f\left(w_t^{\$} - w_{t-1}^{\$}\right) = f\left(w_t - w_{t-1} + \pi_t\right) \tag{19}$$

where $w_t^{\$}$ (resp. w_t) is the log nominal (resp. real) wage, for some increasing function f that we detail later. This is in the spirit previous formulations (e.g. Benigno and Antonio Ricci

(2011); Schmitt-Grohé and Uribe (2016)), which typically use a "hard" constraint such as a rule that the nominal wage can never fall. The formulation here likely would yield rather similar conclusion as those earlier models, but is more tractable.[18]

4.3.2 Optimal nonzero steady state inflation

Given we have inflation has both a first order cost (complexity) and a first order benefit (worker morale linked to the nominal wage), we can discuss an optimal steady state inflation. Total per period log consumption is

$$C_t = e^{\zeta_t + f_t} L_t$$

in this model, without complexity costs of inflation. With complexity costs, log utility becomes:

$$W_t = -\omega_\pi |\pi_t| + f_t$$

as in the steady state, there is no trend productivity growth, and we normalize average log productivity to be 0. So optimal steady state inflation balances the complexity costs of inflation with the efficiency inflation benefits: $\max_\pi W$, with:

$$W = -\omega_\pi |\pi| + f(\pi)$$

Numerical illustration To calibrate, I use a function

$$f(\pi) = \frac{b\pi\mathcal{L}}{b + |\pi|\mathcal{L}} \tag{20}$$

where $b, \mathcal{L}$ are positive parameters. Figure 2 plots it. So f has a sigmoid shape: morale effects affect productivity by an amount between $-b$ and b.[19] I calibrate $f(\infty) = b = 5\%$ to be the maximum effect (a 5% increase or decrease in steady state productivity), and $f'(0) = \mathcal{L} = 1\text{yr}$ so that the slope is 1: 1%/yr deflation would great a 1% productivity loss. Those rough orders of magnitudes.[20] Here, the moral function is a psychological primitive,

[18] Guerreiro et al. (2024) explore a related though different mechanism, where inflation requires conflict-inducing demands of nominal wage increases.

[19] A variant is an asymmetric function between gains and losses, so that $f(x) = \frac{b\pi T}{b+|\pi|T(1+\alpha\,\text{sign}(\pi))}$, for some $\alpha \in [0,1)$, so that the effect of is larger for wage decreases than wage increases: $-f(-\infty) = \frac{b}{1-\alpha} > f(+\infty) = \frac{b}{1+\alpha}$.

[20] The literature gives only indirect guidance to the magnitude. It tends to estimate the elasticity of productivity to a permanent change in the real wage, rather than what we need here, a temporary change in the nominal wage. Increasing the wage by 1% increases productivity by about 2% in Coviello et al. (2022) and 1% in Emanuel and Harrington (2025). Ehrlich and Montes (2024) find that nominal wage cuts lower productivity by 33%, which seems very high. On fact is clear: the cost of nominal wage cuts is high, say

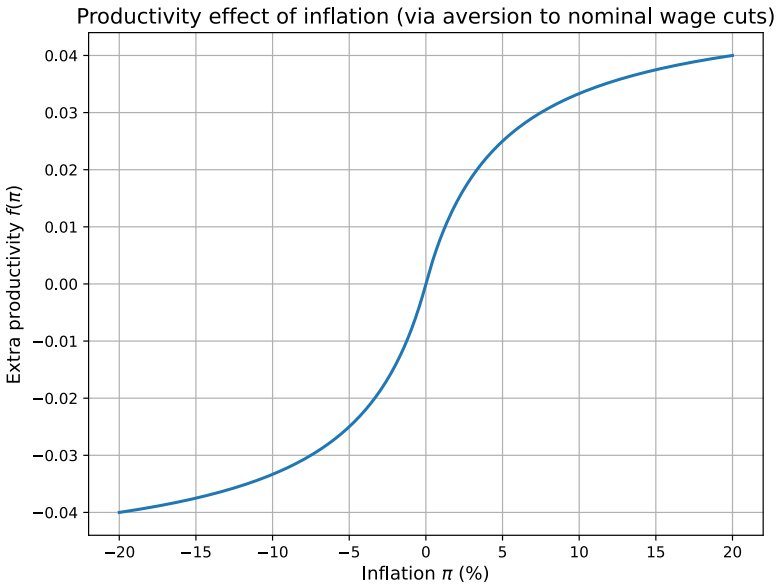


Figure 2: Productivity effect of inflation. This plots the hypothesized productivity effect of inflation: in particular, wage disinflation creates a loss in productivity.

though they might be microfounded (e.g. as a gift-exchange game with a reference at 0 nominal wage increases); and more importantly, they could be empirically calibrated.

Proposition 4. (Optimal steady-state inflation) *The optimal steady state inflation level trades off the complexity costs of inflation* $-\omega_\pi |\pi|$ *with the workers' morale benefits of nominal wage increases,* $\max_{\bar{\pi}} -\omega_\pi |\bar{\pi}| + f(\bar{\pi})$*, which gives, using the morale function (20):*

$$\bar{\pi} = \frac{b}{\mathcal{L}} \left(\sqrt{\frac{\mathcal{L}}{\omega_\pi}} - 1 \right)^+ = \frac{b}{4\omega_\pi} H\left(\frac{\mathcal{L}}{\omega_\pi} \right), \qquad H(\tau) = 4 \left(\frac{1}{\sqrt{\tau}} - \frac{1}{\tau} \right)^+ \tag{21}$$

The function $H(\tau)$ has values in $[0, 1]$, and reaches its maximum for $\tau = 4$, i.e $\mathcal{L} = 4\omega_\pi$. Hence, for $\mathcal{L}$ around that value, we can approximate the optimal inflation rate as:

$$\bar{\pi}^* = \frac{b}{4\omega_\pi} \tag{22}$$

Using the above calibration for the morale function f, and for the welfare losses of in-

first order or zero-th order (i.e. they might create a discrete loss in productivity), as shown that there are virtually no small negative nominal wage cuts: they're either no cuts (with a big mass at 0), or a drastic cuts (perhaps because the firm is on the verge of bankruptcy). Kim and Ruge-Murcia (2009) posit a very small (second order) costs of nominal wage changes, which naturally leads to a very small (about 0.3%) optimal long run inflation.

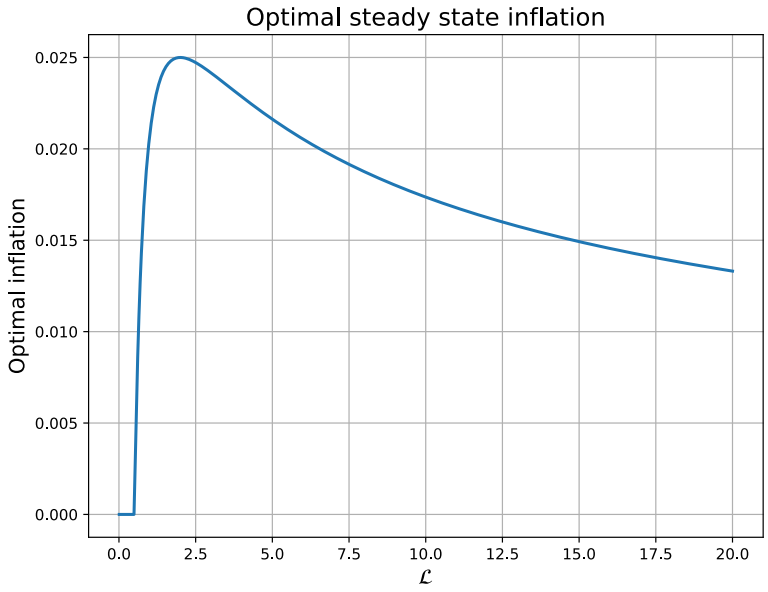


Figure 3: Optimal steady-state inflation. This figure plots the optimal inflation as a function of the slope $\mathcal{L}$ of the "morale function" (20). We see that, above a certain threshold, the optimal inflation is quite insensitive to that slope, yielding some robustness to the analysis.

flation (right after Proposition 2), optimal inflation is $\pi = \frac{5\%}{1}\left(\sqrt{\frac{1}{0.5}} - 1\right) = 2\%/\text{yr}$. The corresponding maximum is $\bar{\pi}^* = \frac{b}{4\omega_\pi} = \frac{5\%}{4\times 0.5} = 2.5/\text{yr}$. Of course, this is all quite a guesstimate, but at least, we have roughly plausible order of magnitude, and more importantly, a conceptual and calibration framework to do a more systematic investigation.

4.3.3 By how much does the ZLB increases optimal steady state inflation?

The existence of a ZLB makes higher inflation more useful, so must increases optimal steady state inflation. Let us see by how much – we shall see that the effect is actually quite small, for a reasonable calibration.

Notations: I use $M = M_f = e^{-\mu\Delta t}$, where μ is the continuous time value (I will soon takes the continuous time limit, $\Delta t \to 0$).

I take a stylized example. The natural rate of interest, r_t^n, can sometimes jump and take a very low value, $\underline{r}^n$ (this could be microfounded as a change in the pure rate of time preference; though of course, in reality, it comes from deeper forces, e.g. people want to save more decrease their leverage): this value will be so low that it can be called a "ZLB" state. Then, the economy remains there, until it goes back to normal, which happens with probability λ_z: so that the average duration of a ZLB is $\frac{1}{\lambda_z}$.

At the ZLB, the central bank just sets the interest rate to the optimal, which is 0. For

simplicity, I assume that it doesn't do forward guidance (e.g. because it cannot commit); so that after the ZLB, the output gap comes back to 0, and inflation to $\bar{\pi}$.

At the ZLB, the Phillips curve gives (with anchored expectations) $\pi_{ZLB} = \bar{\pi} + \hat{\pi}_{ZLB}$ with:

$$\hat{\pi}_{ZLB} = \frac{\kappa}{\tilde{\lambda}_z + r} x_{ZLB} \tag{23}$$

where I define

$$\tilde{\lambda}_z \coloneqq \mu + \lambda_z \tag{24}$$

The IS equation gives

$$x_{ZLB} = \frac{\sigma}{\tilde{\lambda}_z} \left(\underline{r}^n + \bar{\pi} + \hat{\pi}_{ZLB} \right) \tag{25}$$

Solving, we get

$$x_{ZLB} = \frac{\partial x_{ZLB}}{\partial \bar{\pi}} \left(\underline{r}^n + \bar{\pi} \right) \tag{26}$$

$$\frac{\partial x_{ZLB}}{\partial \bar{\pi}} = \frac{\sigma}{\tilde{\lambda}_z \left(1 - \sigma\kappa / \tilde{\lambda}_z^2 \right)} \tag{27}$$

Now, we can think of welfare. At the ZLB, the welfare change is simply:

$$W^{\text{zlb}} = -w_\pi \left| \hat{\pi}_{ZLB} \right| + a_x x_{ZLB}$$

so that

$$\frac{\partial W^{\text{zlb}}}{\partial \bar{\pi}} = \left(a_x + w_\pi \frac{\kappa}{\tilde{\lambda}_z + r} \right) \frac{\partial x_{ZLB}}{\partial \bar{\pi}}$$

and total welfare changes as

$$\Delta W = f^{\text{zlb}} W^{\text{zlb}}$$

where f^{zlb} is the fraction of time spent in ZLB episodes. Hence, the new problem for optimal inflation is

$$\max_{\bar{\pi}} -a_{\bar{\pi}} \left| \bar{\pi} \right| + f^{\text{Tobin}} \left(\bar{\pi} \right) + f^{\text{zlb}} W^{\text{zlb}}$$

The calibration gives

$$\frac{\partial x_{ZLB}}{\partial \bar{\pi}} \simeq \frac{\sigma}{\mu + \lambda_z} = 0.2$$

$$\frac{\partial W^{\text{zlb}}}{\partial \bar{\pi}} = \left(a_x + w_\pi \frac{\kappa}{\tilde{\lambda}_z + r} \right) \frac{\partial x_{ZLB}}{\partial \bar{\pi}} \left(0.4 + 0.2 \right) \frac{\partial x_{ZLB}}{\partial \bar{\pi}}$$

$$\partial_{\bar{\pi}} f^{\text{zlb}} W^{\text{zlb}} \simeq f^{\text{zlb}} a_x \frac{\partial x_{ZLB}}{\partial \bar{\pi}} \simeq 0.3 \times 0.6 \times 0.2 \simeq 0.04$$

This corrects $a_{\bar{\pi}} \simeq 0.5$. So this is a very small force, changing it by about 1/10. So the effective weight on inflation goes from $\omega_\pi = 0.5$ to $\omega_\pi = 0.5 - 0.04 = 0.46$: multiplied by 0.9.

Hence, by Proposition 4, the optimal steady state inflation is multiplied by 1.1: goes from 2% to 2.2%. This is a very small effect.

This section's main result is a simple framework for calculating the impact of the ZLB on optimal inflation. Under the proposed calibration, the impact appears to be fairly small: the optimal inflation rate is multiplied by about 1.1.

4.4 Optimal monetary policy with complexity costs

4.4.1 Optimal policy when the inflation target is 0%

The previous section analyzed the optimal steady state level of inflation. I now turn to the day-to-day conduct of monetary policy. I take the case with commitment, which is arguably the relevant case (the appendix deals with the no-commitment case). To simplify the analysis, I examine first the case with a zero inflation target.

When there are no cost-push shocks, the optimal monetary policy is very simple, and the same as the traditional New Keynesian model: ensure zero output gap and zero inflation, by adjusting the nominal rate to be the natural rate ($i_t = r_t^n$ in (14), so that $x_t = 0$ and $\pi_t = 0$ are solutions).[21]

However, when there is cost-push shock, the $\nu_t \neq 0$ in (15), one cannot get the first best: $x_t = \pi_t = 0$ is not consistent with (15). In the traditional model, a positive cost-push shocks is best handles by a bit of recession and a bit of inflation: both x_t and π_t are non-zero, generically.With complexity aversion, things change.

In what follows, the sign (π) is meant to be the generalized derivative of the function $|\pi|$. So it is 1 if $\pi > 0$, -1 if $\pi < 0$, and some number between -1 and 1 if $\pi = 0$.[22]

Proposition 5. (Optimal response to a cost-push shocks, with commitment: conditions to have 0 inflation as a state by state target). *With welfare (18), given a shock at $t = 0$ the optimal commitment policy is characterized by*

$$\kappa\omega_\pi \operatorname{sign}(\pi_t) = g'(x_t) - M^f g'(x_{t-1}) 1_{t>0} \tag{28}$$

Call $x_t^* := \frac{-\nu_t}{\kappa}$ the output gap that ensures zero inflation. Suppose that:

$$\kappa\omega_\pi \geq \left|\Gamma_x \left(x_t^* - M^f x_{t-1}^* 1_{t>0}\right)\right| \tag{29}$$

which is true when the cost-push shocks are small enough. Then the optimal central bank

[21] I omit here the issue of the zero-lower bound, i.e. the constraint $i_t \geq 0$. The first best solution is then use fiscal policy.

[22] The machinery of subgradients gives the general framework for this.

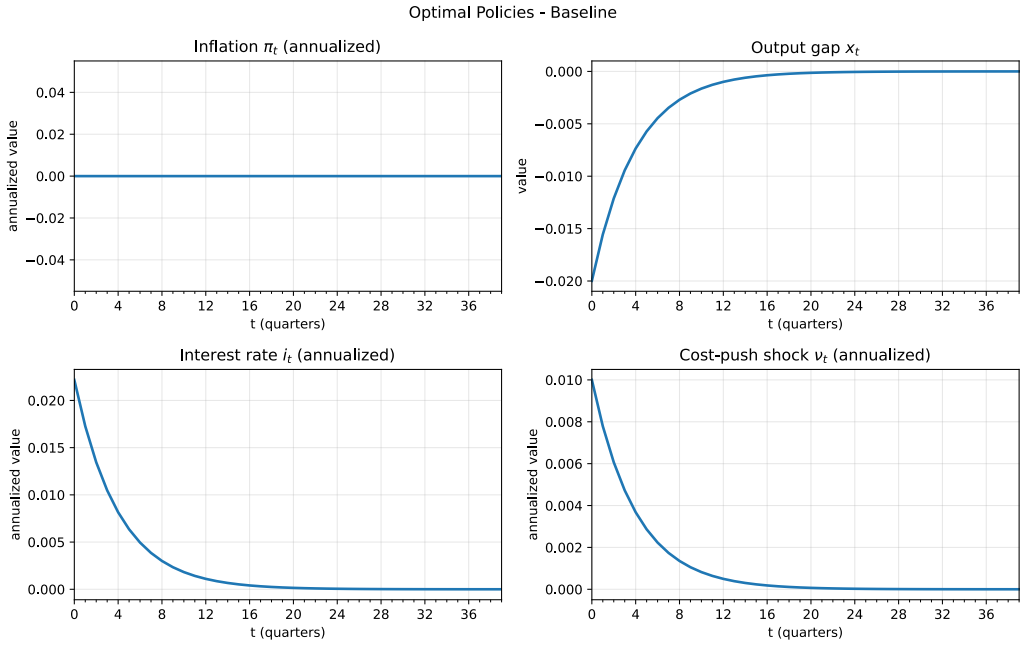


Figure 4: Impulse-response for a small to moderate cost-push shock. Given the exogenous cost-push shock shown in the last panel, the figure shows the resulting economy under the optimal policy. Inflation is pinned down at 0 at all periods.

policy under commitment is simply to keep inflation at 0 at all dates,

$$\pi_t = 0 \tag{30}$$

and adjust the output gap x_t so that $x_t = x_t^*$. In turn, the nominal interest rate i_t is set so that (14) holds.

Consider a small shock, so that 29 holds. Then, we have a starkly simple policy: you keep inflation at its long run target, which here is 0. This is the lasso intuition (Tibshirani (1996)): given the absolute value penalty in $-|\pi_t|$, in an open zone of parameters, one wants to have exactly zero inflation.

For a larger cost-push shock, then inflation will partially accommodate the shock, but quickly, inflation goes back to being exactly 0. Figures 4 and 5 show the behavior of the economy for a small vs a large cost-push shock.

We conclude by considering a consider a special, but illustrative case.

Proposition 6. (Optimal response to a large, persistent cost-push shocks, with commitment) *Suppose that we have a permanent cost-push shock $\nu_t = \nu$. Define $\nu_- := \frac{1}{\theta}\left(\gamma - \frac{\kappa}{1-M^f}\right)$, $\nu_+ := \frac{1}{\theta}\left(\gamma + \frac{\kappa}{1-M^f}\right)$, and $x_\pm := -\frac{\nu_\pm}{\kappa}$. Then, the steady state commitment optimal policy is the following. (i) For moderate cost-push shocks, $\nu \in [\nu_-, v_+]$, we set inflation at 0, $\pi = 0$*

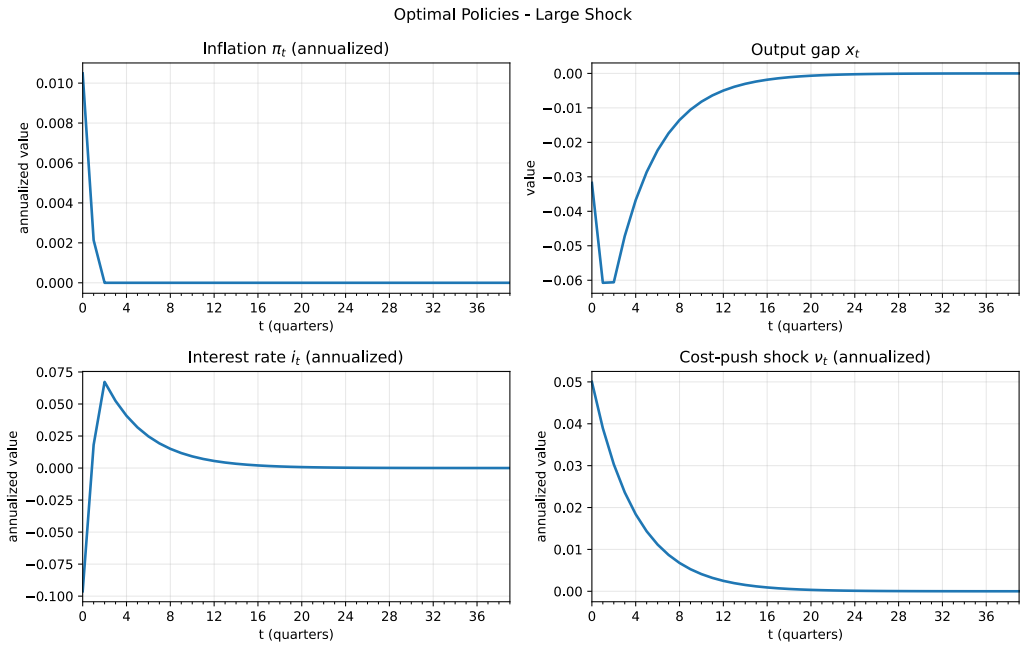


Figure 5: Impulse-response for a large cost-push shock. Given the exogenous cost-push shock shown in the last panel, the figure shows the resulting economy under the optimal policy. This shock is 5 times larger that in Figure 4. Inflation is now above baseline for a little while, then it is pinned at 0 again in all later periods.

and $x = -\frac{\nu}{\kappa}$. *(ii) For very inflationary pressures* ($\nu > \nu_+$)*, the central banks sets* $x = x_-$ *and* $\pi = \frac{\nu - \nu_+}{1-\beta M^f} > 0$*: the central bank sets a recession floor, and lets inflation absorb the brunt of the cost-push shock. (iii) Similarly, if deflationary pressures are very high* ($\nu < \nu_-$)*, the central bank sets* $x = x_+$ *(i.e. a maximum-allowed boom to fight deflations), and sets* $\pi = \frac{\nu - \nu_-}{1-\beta M^f} < 0$.

As in Proposition 5, moderate cost-push pressures call for exactly zero inflation, even at a cost of a moderate recession. However, if there is a very high cost-push shock ($\nu > \nu_+$), then there is a recession floor: $x = x_-$, while inflation bears the brunt of the cost-push shocks, $\pi = \frac{\nu - \nu_+}{1-\beta M^f} > 0$. This is because recessions creates a convex welfare loss ($\theta > 0$), so that they shouldn't be too deep.

4.4.2 Welfare and optimal policy with a non-zero inflation target

Finally, I derive optimal monetary policy with a non-zero inflation targets. The earlier conclusions broadly remain the same, replacing "inflation" by "inflation minus target".

Call $\bar{\pi}$ the target inflation, and $\hat{\pi}_t = \pi_t - \bar{\pi}$ the deviation from the target. I now use the enriched model in Gabaix (2020, section V.6), that allows for non-zero target inflation: then, the basic model (14)-(15) remains the same, replacing π_t by $\hat{\pi}_t$ in the Phillips curve

(15).

Proposition 7. *Suppose that we morale reasons to have positive steady inflation, which is set optimally. Then, up to second order terms in inflation, welfare is:*

$$W_t = \bar{W} - h\left(\hat{\pi}_t\right) + g\left(x_t\right) \tag{31}$$

with g as in (18), and

$$h\left(\hat{\pi}_t\right) = \omega_{\hat{\pi}}^{-} \left|\hat{\pi}_t\right| + \omega_{\pi}\hat{\pi}_t = \omega_{\hat{\pi}}^{-} \left|\hat{\pi}_t\right| 1_{\hat{\pi}_t<0} + \omega_{\hat{\pi}}^{+} \left|\hat{\pi}_t\right| 1_{\hat{\pi}_t>0} \tag{32}$$

where $0 \leq \omega_{\hat{\pi}}^{+} \leq \omega_{\hat{\pi}}^{-}$ (the values are in the proof) so that a temporary increase in inflation is less harmful than a temporary decrease.

Hence, temporary inflation $\hat{\pi}_t$ hurts welfare, much like $-\left|\hat{\pi}_t\right|$, but with a $\omega_{\hat{\pi}}^{+}$ weight for positive temporary inflation, and a higher weight $\omega_{\hat{\pi}}^{-}$ for negative temporary inflation. A negative deviation of inflation is worse than a positive one, because the complexity cost are the same in both directions, but the labor market gains are positive for positive inflation.

Intuitively, this implies the central bank policy is again to set $\hat{\pi}_t = 0$, in a broad set of circumstances, but with some asymmetry. We make this more precise.

Proposition 8. *Make the same assumption as Proposition 7. The first order condition for optimal monetary policy becomes*

$$\kappa h'\left(\hat{\pi}_t\right) = g'\left(x_t\right) - M^f g'\left(x_{t-1}\right) \tag{33}$$

So, it is "more acceptable" to have positive deviations (as $\left|h'\left(0^+\right)\right| \leq \left|h'\left(0^-\right)\right|$) than negative deviations of temporary inflation, as negative ones are costlier. This may be a reason why empirically inflation seems to overshoot its target more often than undershoot it.

Otherwise, the earlier analysis goes through, with the caveat that the previous $\text{sign}\left(\pi_t\right)$ changes into $h'\left(\hat{\pi}_t\right)$. For instance, the condition (29) to have always $\hat{\pi}_t = 0$ becomes:

$$g'\left(x_t\right) - M^f g'\left(x_{t-1}\right) \in \left[\kappa h'\left(0^-\right), \kappa h'\left(0^+\right)\right]. \tag{34}$$

We finished our tour of the complexity costs of inflation. It leads to much higher cost, a different monetary policy (targeting exactly zero inflation, or zero deviation from a target), and indeed allows to think about the optimal steady state level of inflation.

Numerous extensions are possible, e.g. refining the welfare function to allow for rich heterogeneity. More generally, thinking through complexity costs in welfare and then optimal policy should be productive.

5 Conclusion

This paper extended the static modelling of complexity aversion in Gabaix (2025), to dynamic contexts, in a way useful for macroeconomics. It appears to make macro more sensible. The business cycle and inflation have first order costs (as they make planning more difficult), rather than second order, hence minuscule, costs as in Lucas (1987) and Woodford (2003). Taking into account the plain fact that workers seem particularly averse to nominal wage cuts, this leads to a simple calibratable framework to calibrate the first order costs and benefits or inflation, hence the optimal inflation level. It also leads to an updated welfare function, hence a different monetary policy, which is typically more targeted towards a stable inflation, except in deep recessions. The model captures hitherto hard to model notions (such as the complexity cost of inflation), and is very tractable, so it might be useful to make macroeconomic analysis more realistic.

References

Abeler, J. and S. Jäger (2015). Complex tax incentives. *American Economic Journal: Economic Policy 7*(3), 1–28.

Afrouzi, H., S. Y. Kwon, A. Landier, Y. Ma, and D. Thesmar (2023). Overreaction in expectations: Evidence and theory. *The Quarterly Journal of Economics 138*(3), 1713–1764.

Akerlof, G. A., W. T. Dickens, and G. L. Perry (1996). The macroeconomics of low inflation. *Brookings papers on economic activity 1996*(1), 1–76.

Angeletos, G.-M. and C. Lian (2018b). Forward guidance without common knowledge. *American Economic Review 108*(9), 2477–2512.

Augenblick, N., B. K. Jack, S. Kaur, F. Masiye, and N. Swanson (2024). Retrieval failures and consumption smoothing: A field experiment on seasonal poverty. Technical report, Technical Report, University of California, Berkeley.

Bar-Gill, O. (2012). *Seduction by Contract: Law, Economics, and Psychology in Consumer Markets*. Oxford University Press.

Ben-David, I. and A. Chinco (2024). Expected EPS × Trailing P/E. Technical report, National Bureau of Economic Research.

Benigno, P. and L. Antonio Ricci (2011). The inflation-output trade-off with downward wage rigidities. *American Economic Review 101*(4), 1436–1466.

Bewley, T. F. (1999). *Why wages don't fall during a recession*. Harvard university press.

Bilbiie, F. O. (2020). The new keynesian cross. *Journal of Monetary Economics 114*, 90–108.

Binetti, A., F. Nuzzi, and S. Stantcheva (2024). People's understanding of inflation. Technical report, National Bureau of Economic Research.

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer (2020). Over-reaction in macroeconomic expectations. *American Economic Review*.

Bordalo, P., N. Gennaioli, and A. Shleifer (2013). Salience and consumer choice. *Journal of Political Economy 121*(5), 803–843.

Caplin, A., D. Csaba, J. Leahy, and O. Nov (2020). Rational inattention, competitive supply, and psychometrics. *The Quarterly Journal of Economics 135*(3), 1681–1724.

Carroll, G. (2015). Robustness and linear contracts. *American Economic Review 105*(2), 536–563.

Célérier, C. and B. Vallée (2017). Catering to investors through security design: Headline rate and complexity. *Quarterly Journal of Economics 132*(3), 1469–1508.

Cochrane, J. H. (2023). *The Fiscal Theory of the Price Level*. Princeton University Press.

Coviello, D., E. Deserranno, and N. G. Persico (2022). Minimum wage and individual worker productivity: Evidence from a large us retailer. *Journal of Political Economy 130*(9), 2315–2360.

Dari-Mattiacci, G. and F. Marotta-Wurgler (2022). Learning in standard-form contracts: Theory and evidence. *Journal of Legal Analysis 14*(1), 244–314.

Dean, M. and N. L. Neligh (2023). Experimental tests of rational inattention. *Journal of Political Economy*.

Di Tella, R., R. J. MacCulloch, and A. J. Oswald (2001). Preferences over inflation and unemployment: Evidence from surveys of happiness. *American economic review 91*(1), 335–341.

Ehrlich, G. and J. Montes (2024). Wage rigidity and employment outcomes: Evidence from administrative data. *American Economic Journal: Macroeconomics 16*(1), 147–206.

Eichenbaum, M. S., J. Guerreiro, and J. Obradovic (2026). Ricardian non-equivalence. Technical report, National Bureau of Economic Research.

Emanuel, N. and E. Harrington (2025). The payoffs of higher pay. Technical report, Working Paper.

Enke, B. and T. Graeber (2023). Cognitive Uncertainty. *Quarterly Journal of Economics 138*(4), 2021–2067.

Enke, B., T. Graeber, and R. Oprea (2023). Complexity and time. Working paper, National Bureau of Economic Research.

Enke, B., T. Graeber, R. Oprea, and J. Yang (2024). Behavioral attenuation. Working paper, National Bureau of Economic Research.

Enke, B. and C. Shubatt (2023). Quantifying lottery choice complexity. Technical report, National Bureau of Economic Research.

Farhi, E. and I. Werning (2019). Monetary Policy, Bounded Rationality, and Incomplete Markets. *American Economic Review 109*(11), 3887–3928.

Fudenberg, D. and I. Puri (2023). Evaluating and extending theories of choice under risk. Working paper, MIT Economics.

Gabaix, X. (2014). A sparsity-based model of bounded rationality. *Quarterly Journal of Economics 129*(4), 1661–1710.

Gabaix, X. (2019). Behavioral inattention. *Handbook of Behavioral Economics 2*, 261–344.

Gabaix, X. (2020). A behavioral new keynesian model. *American Economic Review 110*(8), 2271–2327.

Gabaix, X. (2023). Behavioral macroeconomics via sparse dynamic programming. *Journal of the European Economic Association 21*(6), 2327–2376.

Gabaix, X. (2025). A theory of complexity aversion. *Available at SSRN 5185671*.

Gabaix, X. and T. Graeber (2024). The complexity of economic decisions. Working Paper 33109, National Bureau of Economic Research.

Gabaix, X. and D. Laibson (2002). The 6D bias and the equity-premium puzzle. *NBER Macroeconomics Annual 16*, 257–312.

Gabaix, X., D. Laibson, G. Moloche, and S. Weinberg (2006, September). Costly information acquisition: Experimental analysis of a boundedly rational model. *American Economic Review 96*(4), 1043–1068.

Gagnon-Bartsch, T., M. Rabin, and J. Schwartzstein (2018). Channeled attention and stable errors. Technical report.

Galí, J. (2015). *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and its Applications*. Princeton University Press.

Georgarakos, D., K. H. Kim, O. Coibion, M. Shim, M. A. Lee, Y. Gorodnichenko, G. Kenny, S. Han, and M. Weber (2025). How costly are business cycle volatility and inflation? a vox populi approach. Technical report, National Bureau of Economic Research.

Gertler, M. and P. Karadi (2011). A model of unconventional monetary policy. *Journal of monetary Economics 58*(1), 17–34.

Gormsen, N. J. and K. Huber (forth.). Corporate discount rates. Technical report.

Graham, J. R. (2022). Presidential address: Corporate finance and reality. *The Journal of Finance 77*(4), 1975–2049.

Grigsby, J., E. Hurst, and A. Yildirmaz (2021, February). Aggregate nominal wage adjustments: New evidence from administrative payroll data. *American Economic Review 111*(2), 428–471.

Guerreiro, J., J. Hazell, C. Lian, and C. Patterson (2024). Why do workers dislike inflation? wage erosion and conflict costs. *NBER Working Paper* (w32956).

Gulati, M. and R. E. Scott (2012). *The three and a half minute transaction: Boilerplate and the limits of contract design*. University of Chicago Press.

Hart, O. (1995). *Firms, contracts, and financial structure*. Clarendon press.

Hartman-Glaser, B. and B. Hébert (2020). The insurance is the lemon: Failing to index contracts. *The Journal of Finance 75*(1), 463–506.

Huck, S. and G. Weizsäcker (1999). Risk, complexity, and deviations from expected-value maximization: Results of a lottery choice experiment. *Journal of Economic Psychology 20*(6), 699–715.

Iyengar, S. S. and E. Kamenica (2010). Choice proliferation, simplicity seeking, and asset allocation. *Journal of Public Economics 94*(7-8), 530–539.

Kahan, M. and M. Klausner (1997). Standardization and innovation in corporate contracting (or the economics of boilerplate). *Virginia Law Review 83*, 713.

Kaur, S. (2019). Nominal wage rigidity in village labor markets. *American Economic Review 109*(10), 3585–3616.

Khaw, M. W., Z. Li, and M. Woodford (2021). Cognitive imprecision and small-stakes risk aversion. *The Review of Economic Studies 88*(4), 1979–2013.

Kim, J. and F. J. Ruge-Murcia (2009). How much inflation is necessary to grease the wheels? *Journal of Monetary Economics 56*(3), 365–377.

Lucas, R. E. J. (1987). *Models of Business Cycles*. Oxford, UK; New York, USA: Basil Blackwell.

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

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

Matějka, F. and A. McKay (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review 105*(1), 272–298.

McKay, A., E. Nakamura, and J. Steinsson (2016). The Power of Forward Guidance Revisited. *American Economic Review 106*(10), 3133–3158.

Mei, P. and L. Wu (2024). Thinking about the economy, deep or shallow?

Nakamura, E., J. Steinsson, P. Sun, and D. Villar (2018). The elusive costs of inflation: Price dispersion during the us great inflation. *The Quarterly Journal of Economics 133*(4), 1933–1980.

Oprea, R. (2020). What makes a rule complex? *American Economic Review 110*(12), 3913–51.

Oprea, R. (2024b). Decisions under risk are decisions under complexity. *American Economic Review 114*(12), 3789–3811.

Orphanides, A. and D. W. Wilcox (2002). The opportunistic approach to disinflation. *International finance 5*(1), 47–71.

Puri, I. (2025). Simplicity and risk. *The Journal of Finance 80*(2), 1029–1080.

Rubinstein, A. (1998). *Modeling Bounded Rationality*. MIT press.

Schmitt-Grohé, S. and M. Uribe (2016). Downward nominal wage rigidity, currency pegs, and involuntary unemployment. *Journal of Political Economy 124*(5), 1466–1514.

Schwartzstein, J. and A. Sunderam (2021). Using models to persuade. *American Economic Review 111*(1), 276–323.

Shafir, E., P. Diamond, and A. Tversky (1997). Money illusion. *Quarterly Journal of Economics 112*(2), 341–374.

Shiller, R. J. (1997). Why do people dislike inflation? In *Reducing inflation: Motivation and strategy*, pp. 13–70. University of Chicago Press.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics 50*(3), 665–690.

Sonsino, D., U. Benzion, and G. Mador (2002). The complexity effects on choice with uncertainty–experimental evidence. *Economic Journal 112*(482), 936–965.

Spier, K. E. (1992). Incomplete contracts and signalling. *The RAND Journal of Economics*, 432–443.

Stantcheva, S. (2024). Why do We Dislike Inflation? *Brookings Papers on Economic Activity 55*(1).

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological) 58*(1), 267–288.

Tirole, J. (2009). Cognition and incomplete contracts. *The American Economic Review 99*(1), 265–294.

Tobin, J. (1972). Inflation and unemployment. *American Economic Review 62*(1), 1–18.

Woodford, M. (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.

Woodford, M. (2019). Monetary policy analysis when planning horizons are finite. *NBER Macroeconomics Annual 33*(1), 1–50.

Woodford, M. (2020). Modeling imprecision in perception, valuation, and choice. *Annual Review of Economics 12*, 579–601.

Young, H. P. and M. A. Burke (2001). Competition and custom in economic contracts: a case study of illinois agriculture. *American Economic Review 91*(3), 559–573.

A Theory complements

A.1 The complexity cost of business cycles

Here I complement the analysis of welfare in the main text, in particular Section 3.

A.1.1 The cost of business cycles with idiosyncratic risk

With idiosyncratic risk, there's a danger that it will swamp the impact of business cycle, and make business cycle risk minimal again. We show that, if there area enough agents with small idiosyncratic risk, we still get a large cost of business cycle.

Recall that the loss associated with complexity is $\ell^{\mathrm{CA}} = \omega_y \sigma_y$ with

$$\omega_y \coloneqq \gamma\xi\left(1-m\right)$$

With idiosyncratic risk, the risk is $y_t + \varepsilon_{it}$, where ε_{it} is idiosyncratic risk. This is

$$\ell^{\mathrm{CA},0} = \omega_y \sqrt{\sigma_{\varepsilon_i}^2 + \sigma_y^2}$$

Hence, the loss solely due to the business cycle is

$$\ell^{\mathrm{CA}} = \omega_y g\left(\sigma_y\right), \qquad g\left(\sigma_y\right) \coloneqq \mathbb{E}\left[\sqrt{\sigma_{\varepsilon_i}^2 + \sigma_y^2} - \sigma_{\varepsilon_i}\right]$$

where the expectation averages over the distribution of σ_{ε_i}: so $g\left(\sigma_y\right)$ is the specific contribution of aggregate risk to total complexity losses.

To get a feel for that term, suppose first that a fraction a_0 of agents have 0 idiosyncratic risk. Then, they contribute to a term

$$g\left(y\right) \sim a_0 \sigma_y \tag{35}$$

in the limit of small σ_y. Again, this is first order.

Suppose on the other hand that σ_ε is typically much larger than σ_y; then, by Taylor expansion

$$g\left(y\right) \sim \mathbb{E}\left[\frac{1}{\sigma_{\varepsilon_i}}\right]\frac{\sigma_y^2}{2} \tag{36}$$

which shows again how important is the quantity of agents with close to 0 idiosyncratic risk.

Now, suppose that have a distribution of idiosyncratic risk:

$$f\left(\sigma_\varepsilon\right) = a_0\delta\left(\sigma_\varepsilon\right) + b\sigma_\varepsilon^{-\alpha} + h\left(\sigma_\varepsilon\right),$$

with $\alpha \in (0,1)$, and $h(\sigma_\varepsilon)$ is smooth. The new term $b\sigma_\varepsilon^{-\alpha}$ means that there is a "lots" of agents with very small idiosyncratic risk. Then, the extra cost of a business cycle of size σ_y is:

$$g(\sigma_y) = a_0\sigma_y + bc_\alpha\sigma_y^{2-\alpha} + O\left(\sigma_y^2\right)$$

for small σ_y, with a positive constant c_α.[23] So we get terms much larger than σ_y^2. So even with idiosyncratic risk, we get "large" (proportional to σ_y^β with $\beta < 2$) losses from the business cycle.

A.1.2 The full cost of the business cycle in a distorted economy

Call T_p (resp. T_w) the average time between price (resp. wage) resets. The welfare cost is (when $rT_p \ll 1$),

$$W_t = W_t^{\text{trad}} + W_t^{\text{CA}} + W_t^{\text{Tobin}} \tag{37}$$

$$W_t^{\text{trad}} = -\frac{\varepsilon}{2}T_p^2\pi_t^2 - \frac{1}{2}(\gamma+\sigma)(x_t - x^{**})^2 \tag{38}$$

$$W_t^{\text{CA}} = -\xi^\pi T_w|\pi_t| - \xi^x|x_t| \tag{39}$$

and $W_t^{\text{Tobin}} = f(\pi_t)$ is the effect on inflation to help avoid nominal wage cuts.

We have the complexity losses from the output gap, $\xi^x|x_t|$, where $\xi^x = \xi(1-m_x)$ is the complexity penalty associated with x, and m_x the attention to x.

For the ideal output, it's the one corresponding to $\max_C u(C) - v(C)$, where productivity is normalized to 1. However, because of (i) markup $\mu = \frac{1}{1-1/\varepsilon}$ and taxes (assumed to be affine, with a marginal tax rate τ_{tax}, the FOC of the consumer is (indeed, $p = w\mu$,so even without taxes, we have $\frac{1}{\mu}u'(C) = V'(C)$)

$$\left(1 - \frac{1}{\varepsilon}\right)(1-\tau_{\text{tax}})u'(C) - V'(C) = 0$$

Call $\tau^{\text{eff}} = \frac{1}{\varepsilon} + \tau_{\text{tax}}$ the total efficiency wedge. In the limit of small markups and taxes, this gives a wedge $\delta C = \frac{\tau^{\text{eff}}u'}{u''+v''}$, i.e. in relative consumption, the output gap is

$$x^{**} = \frac{\tau^{\text{eff}}}{\gamma+\phi} \tag{40}$$

As a calibration, we have $\varepsilon \simeq 5$ and $\tau_{\text{tax}} \simeq 30\%$ so $\tau^{\text{eff}} \simeq 50\%$. With $\gamma = 4$ and $\phi = 1$,that gives $x^{**} = 10\%$: the economy produces 10% less than if all the wedges where removed.

[23] Indeed, $c_\alpha = \int_0^\infty \frac{v^{-\alpha}}{v+\sqrt{1+v^2}}dv = \frac{\Gamma\left(\frac{\alpha}{2}\right)\Gamma\left(\frac{1-\alpha}{2}\right)}{2\sqrt{\pi}(2-\alpha)}$. This implies for that $\alpha \to 1$, $c_\alpha \sim \frac{1}{1-\alpha}$.

So, the Taylor expansion is:

$$W_t^{\text{trad}} = \tau^{\text{eff}} x_t - \xi^x |x_t| - \frac{1}{2}(\gamma + \phi) x_t^2 + A_t \tag{41}$$

where the A_t terms are independent of x_t. How to calibrate that? First, it is clear that, around $x_t = 0$,people prefers a boom to a recession: this means that $\tau^{\text{eff}} > \xi^x$. With the calibration $\xi = 0.5$, $m_x = 0.8$ (as people understand reasonably well how to spend their income, so higher that $m_\pi = 0.5$), we get $\xi^x = \xi(1 - m_x) = 0.1$, $\tau^{\text{eff}} - \xi^x = 0.4$, so things qualitatively make sense.

Suppose that there is a boom, and \$10 billion of extra output is produced. When $x^{**} = 0$, this is a second order welfare loss: people work "too much". With a "naive" Keynesian view, this is \$10 billion of extra goods, so a welfare equivalent of 10 billion. But once we take into account that resources (e.g. labor) are used to produce is, with $\tau^{\text{eff}} = 0.5$, this is still worth \$5 billion in welfare (that's because we're closer to the optimum). With the complexity costs, this is \$4 billion – because this adds a bit of complexity to the situation.

A.2 A simple calibration of complexity aversion

Suppose a very simple gamble, with equal probabilities of payoffs $\mu+\sigma$, $\mu-\sigma$, with $\sigma > \mu > 0$. A prospect theoretic agent with loss aversion λ (usually calibrated to $\lambda = 2$) will have a value: $v = \mu + \sigma + \lambda(\mu - \sigma) = (1 + \lambda)\mu + (\lambda - 1)\sigma$. So, to be indifferent ($v = 0$), she will require a Share ratio: $S = \frac{\mu}{\sigma} = \frac{\lambda - 1}{\lambda + 1}$. With the typical calibration $\lambda \simeq 2$, this gives $\frac{\mu}{\sigma} = \frac{1}{3}$.

The complexity interpretation (in the spirit of Oprea (2024b)) may be that the certainty equivalent is $\mu - \xi\sigma(1 - m)$, so that the require a Sharpe ratio $S = \frac{\mu}{\sigma} = \xi(1 - m)$. So, we are led to

$$\xi(1 - m) = \frac{\lambda - 1}{\lambda + 1} \simeq \frac{1}{3}$$

With $m \simeq 0.5$, that leads to $\xi = \frac{2}{3}$. To avoid excess pretense of knowledge, and to be a bit conservative, we'll round this to $\xi \simeq \frac{1}{2}$.

Note that this theory predicts that there will be more "effective loss aversion" in complex gambles, and under cognitive load. This is qualitatively true in Iyengar and Kamenica (2010) (in their experiment, when there are more lotteries to chose from, the simple, lower expected value lottery is chosen more).

A variant is the the CA parameter ξ could be higher when in an adversarial relationship, i.e. when the agent knows that there is the possibility for an Akerlof lemons effect.

B Omitted proofs

Proof of Proposition 2 The complexity costs are, in general:

$$C^{\mathrm{CA}} = \xi \sum_{i,t} \mathbb{E} \left| \tilde{a}_t u_{a_t a_t} a_{t,x_i} x_i \right| (1 - m_i) \tag{42}$$

In our context, with $\tilde{a}_t = c_t$, this gives:

$$C^{\mathrm{CA}}/\xi = \sum_{t} \mathbb{E} \left| \gamma \beta^t c_t^{1-\gamma} \frac{\partial \ln c_t}{\partial x_i} \sigma_{x_i} \right| (1 - m_i)$$

In the limit $r\tau \to 0$,

$$\begin{aligned} C^{\mathrm{CA}}/\xi &= \gamma c_0^{1-\gamma} (1-m) \int_0^\infty e^{-rt} \left| \frac{\partial \ln c_t}{\partial x_i} \sigma_{x_i} \right| dt \\ &= \gamma c_0^{1-\gamma} (1-m) \sigma_{x_i} \sum_{n \geq 0} e^{-rn} \int_0^1 \left| \frac{1}{2} - h \right| dt \\ &= \gamma c_0^{1-\gamma} (1-m) \sigma_{x_i} \sum_{n \geq 0} e^{-rn} \frac{1}{4} \\ &= \frac{1}{r} c_0^{1-\gamma} \frac{1}{4} \xi \gamma (1-m) \sigma_{x_i} \end{aligned}$$

Proof of Proposition 5 Let us call $g(x) = \gamma x - \frac{1}{2}\theta x^2$ the welfare impact of output gap x. The Lagrangian is:

$$L = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[- |\pi_t| + g(x_t) - \Xi_t \left(-\pi_t + \beta M^f \pi_{t+1} + \kappa x_t + \nu_t \right) \right], \tag{43}$$

where Ξ_t are Lagrange multipliers. The first order conditions are: $L_{x_t} = 0$ and $L_{\pi_t} = 0$, i.e.:

$$\begin{aligned} g'(x_t) &= \Xi_t \kappa & (44) \\ \mathrm{sign}(\pi_t) &= \Xi_t - M^f \Xi_{t-1} 1_{t>0} & (45) \end{aligned}$$

hence

$$\mathrm{sign}(\pi_t) = g'(x_t) - M^f g'(x_{t-1}) 1_{t>0} \tag{46}$$

which gives (28).

We have solution with $\pi_t = 0$ at all dates iff (i) $\kappa x_t + \nu_t = 0$ from the Phillips curve, so

$x_t = x_t^* \coloneqq \frac{-\nu_t}{\kappa}$ and (ii) (28) holds with $|\text{sign}(\pi_t)| \leq 1$, i.e.

$$\kappa \geq \left|g'(x_t) - M^f g'(x_{t-1})\right| = \left|\gamma\left(1 - M^f\right) - \theta\left(x_t^* - M^f x_{t-1}^*\right)\right|$$

which is the announced expression (29).

If this inequality is violated, then the optimum is remains by (28).

Proof of Proposition 6 From the Phillips curve, the steady state values π, x satisfy

$$\pi = \frac{\kappa x + \nu}{1 - \beta M^f} \tag{47}$$

and the optimality condition (28) gives: $\kappa\,\text{sign}(\pi) = \left(1 - M^f\right) g'(x)$ i.e.

$$\kappa\,\text{sign}(\pi) = \left(1 - M^f\right)(\gamma - \theta x) \tag{48}$$

We have $\pi = 0$ iff $x = \frac{-\nu}{\kappa}$ and (48) holds, i.e. $\kappa \geq \left|\left(1 - M^f\right)(\gamma - \theta x)\right|$. Hence, the boundary condition is $\pm\kappa = \left(1 - M^f\right)(\gamma - \theta x)$, i.e. $x = \frac{\gamma}{\theta} - \frac{\pm\kappa}{\theta\left(1 - M^f\right)}$, as announced. The corresponding value of ν is $\nu_\pm \coloneqq -\kappa x_\pm$.

When $\nu > \nu_+$, (so inflationary pressures are very large), we have $\pi > 0$, so $\kappa = \left(1 - M^f\right)(\gamma - \theta x)$, and $x = x_+$. Then, (47) gives $\pi = \frac{\nu - \nu_+}{1 - \beta M^f}$ Likewise, when $\nu < \nu_-$, we have $\pi < 0$ and $x = x_-$.

Proof of Proposition 4 The optimal inflation satisfies $\max_\pi -\omega_\pi |\pi| + \frac{b\pi T}{b + |\pi| T}$. For $\pi > 0$, the FOC is $-\omega_\pi + \frac{b^2 T}{(b + |\pi| T)^2} = 0$, so $\pi = \frac{b}{T}\left(\sqrt{\frac{T}{\omega_\pi}} - 1\right)^+$.

Proof of Proposition 7 Welfare is:

$$W_t = -\omega_{\bar{\pi}} |\bar{\pi}| - \omega_{\hat{\pi}} |\hat{\pi}_t| + g(x_t) + f_t$$

where f_t is the "efficiency" from term. As temporarily inflation is less well perceived than steady state one, we have $\omega_{\hat{\pi}} \geq \omega_{\bar{\pi}} \geq 0$. Around the steady state, $f_t = f(\bar{\pi} + \hat{\pi}_t)$. Taking the Taylor expansion around $\hat{\pi}_t = 0$, we get:

$$f_t - f(\bar{\pi}) = f'(\bar{\pi})\,\hat{\pi}_t = \omega_{\bar{\pi}} \hat{\pi}_t$$

so $W_t = \bar{W} - \omega_{\hat{\pi}} |\hat{\pi}_t| + \omega_{\bar{\pi}} \hat{\pi}_t + g(x_t)$ i.e.

$$W_t = \bar{W} - h(\hat{\pi}_t) + g(x_t) \tag{49}$$

with

$$h\left(\hat{\pi}_t\right) = \omega_{\hat{\pi}} \left|\hat{\pi}_t\right| - \omega_{\bar{\pi}} \hat{\pi}_t = \omega_{\hat{\pi}}^{-} \left|\hat{\pi}_t\right| 1_{\hat{\pi}_t<0} + \omega_{\hat{\pi}}^{+} \left|\hat{\pi}_t\right| 1_{\hat{\pi}_t>0}$$

with $\omega_{\hat{\pi}}^{+} \coloneqq \omega_{\hat{\pi}} - \omega_{\bar{\pi}}, \omega_{\hat{\pi}}^{-} \coloneqq \omega_{\hat{\pi}} + \omega_{\bar{\pi}}$. so, we get an asymmetric kink in $\hat{\pi}_t$.