A Theory of Complexity Aversion*

Xavier Gabaix

February 2, 2026

Abstract

This paper proposes a tractable model of "complexity aversion." The key ingredient is first-order complexity aversion: when people know they are making a mistake because the situation is complex, they experience dread—a utility loss proportional to the expected size of the error in their action. This adds a new complexity-aversion term to the traditional utility function. This simple enrichment has a range of implications, which I illustrate: (i) Whereas the traditional model predicts that contracts should be indexed to aggregate factors (e.g., inflation), with sufficient complexity aversion contracts become non-indexed. (ii) If complexity aversion is high enough, the price of a good will be constant over time—even when marginal cost varies—to avoid burdening consumers with a complex pricing system. (iii) Traditional economics often recommends tax systems far more complex than what we observe. With complexity aversion, the optimal tax system is simpler—e.g. has uniform taxes or a few tax brackets. (iv) 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. (v) This allows to calculate an optimal 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. The zero lower bound does not change this optimal inflation much—it would multiply it by a factor of 1.1. (vi) These considerations in turn alter the optimal monetary policy, which should often ensure completely stable inflation, to the exclusion of other goals. Finally, I discuss how this model of complexity aversion can underpin a useful "behavioral mechanism design," leading to more realistic—and simpler—mechanisms.

*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 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

I propose a tractable notion of "complexity aversion" and explore its implications for "simple mechanisms" in several settings: simple price systems, simple taxes, simple (non-indexed) contracts, and the complexity cost of inflation and the business cycle, and the resulting optimal monetary policy.

Actual economic systems are typically much "simpler" than a frictionless rational model would predict. For instance, basic economic theory suggests that most contracts should be indexed to aggregate risks—such as inflation, GDP, or the stock market—because this reduces the risk faced by consumers. But that is typically not the case in practice. Intuitively, this is because doing so would be "too complicated." The model formalizes this and, sensibly, implies that contracts will generally not be indexed unless the risk-reduction effects are large enough.

Likewise, prices for Whole Foods delivery are the same across hours and days of the week, even though they "should" vary in a frictionless world, as the marginal cost of deliveries changes over the day.[1] A barber typically charges prices that do not depend on the time of day, even though the 5pm slot is much more popular than the 2pm slot, so (generically, at least) their prices should differ.

The core contribution is a broadly applicable model of "complexity aversion." I attend closely to the psychology while ensuring the model is portable: it scales appropriately as problem size changes, whereas more simplistic devices—such as a fixed cost—do not, because they fail to scale. I also show that other devices, such as information costs, or smooth ambiguity aversion, cannot generate "full simplicity": for example, they invariably predict some indexation rather than none. Complexity aversion (or a close variant) is necessary and sufficient to obtain full simplicity.

Two key properties of "first order complexity aversion" yield full simplicity. The model is essentially a half-equation: the complexity-aversion cost in (5), which is subtracted from the traditional utility function. At the same time, that half-equation has sharp consequences, forcing simplicity in economic outcomes.

This complexity aversion cost distills the following psychology. First, the consumer experiences a complexity "dread", the sense that she is making a mistake, which directly lowers her utility (above the traditional cost from lowering one's utility because of a mistake). Unlike say cognition or information cost, this "dread" cost is paid even when attention is 0 (whereas with zero attention, no cognition cost is paid). Hence, even for a very behavioral consumer who pays zero thinking costs, lower complexity (and the scope for mistake) will

[1] In the spring 2024, this was \$9.99 for 2-hour slots, e.g. 6-8am, 8-10am, etc. There are also 1-hour slots (6-7am, 7-8am etc.), that all cost \$14.99.

be useful to lower his dread. This is a sense in which cognitive costs, and complexity dread are distinct, though they interact with each other: cognitive costs generate mistakes, which in terms generate complexity aversion.

Second, this complexity aversion cost is (in the main specification) first order in the size of the mistake (in action space) rather than second order: if the expected size of the mistake is e (in absolute value), the dread cost is proportional to e, rather than e^2. This "first order complexity aversion" is a close cousin, or perhaps indeed an offspring, of the "first order risk aversion" found necessary to explain the behavior under risk (Kahneman and Tversky (1979), Rabin (2000)). It is Kahneman and Tversky (1979)'s first order risk aversion, but for the risk of making a mistake because of bounded rationality, rather than (as KT) for the exogenous risks in the outside world.

Traditional model ambiguity aversion would generate something akin to second order complexity aversion. But to robustly generate simplicity at the optimum, we need first order complexity aversion. Why? This is because making mechanism more "rationally efficient" by a small factor b will create extra allocative gains proportional to b, typically. With first order complexity aversion, the dread costs will also be proportional to b, and so they might overcome the "rational benefits" of order b. However, if the complexity aversion is proportional to b^2, then it will be much lower than the the first order rational allocative gains (of order b); so at the optimum, we shall want to have b away from 0: the mechanism will be a complex, as b is not 0. So second order complexity aversion does not (most of the time) generate fully simple mechanism. The logic is actually more rich, as the size of the mistakes, hence the dread, depends on attention which is endogenous, but this is the essence of the reasoning.

Complexity aversion, when it is high enough, leads to very simple mechanisms at the optimum *Contracts are not indexed.* I start the non-indexed contracts. Indexing allows to reduce risk, but also (if the indexed pays off at some point in the future) creates extra complexity, as a consumer should keep track of the value of the index. Hence, if complexity aversion is high enough, no indexation is better. To illustrate this, I take the simplest contract — the choice of buying a hedge, or not. This is just a solo decision problem (against nature), and this way I eschew all complications of strategic interactions between humans. But of course, the model applies also in those richer settings, which I explore next.

Price schemes are more rigid. When complexity is high enough, the price of a good (e.g. Whole Foods deliveries) will be constant across time-periods, even though the marginal cost might be variable. Intuitively, this is to avoid annoying the consumer with a very complex price system.

A price system is optimally simpler when complexity aversion is higher, when the com-

plexity of thinking about the price is higher, when the marginal cost is less variable, and when the price elasticity is lower (which often means that competition is lower). So Whole Foods deliveries will have a constant price (low elasticity of demand), while the price of strawberries will be more variable (as the marginal cost is more variable, and the price elasticity is higher). This result of a fully simple price system may seem intuitive, but it does require some care in the modeling of complexity aversion.

Optimal tax systems are much "simpler" than in the traditional model, e.g. with one or a few tax brackets, rather than a different tax for each good and state of the world.

This paper revisits the (Ramsey) theory of optimal goods taxation, with now complexity-averse agents. In the traditional Ramsey model, the optimal tax on each good is (generically) different for each good (as it is inversely proportional to the elasticity of demand for the good). This appears ridiculously complex—with ten thousand different goods, rational theory wants ten thousand different tax rates. This is probably one of the reasons why those strict recommendations of optimal taxation theory are scarcely followed in practice. With complexity-averse agents, things are much more sensible. Indeed, if complexity aversion is high enough, the optimal tax system is "simple", e.g. just features a uniform tax rate.

Complexity aversion creates "aversion to deciding". If an "active decision" (e.g., rebalance one's portfolio) entails complexity aversion, while a "passive course" (e.g. do nothing), when complexity aversion is high enough, and the decision is hard enough, people will stay passive.[2]

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 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

[2]Years ago, I described this to Robert Lucas, the great theorist of rational behavior. He told me that this what he did—he had never rebalanced his retirement portfolio, and just let it drift, like the behavioral agents in this present model.

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). In addition, 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), and the 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 2%, but really, the advantage is to have a calibratable framework in which those estimates can be made.

The advantage of a unified framework. The reader might be a bit bewildered by all those illustrations. Are they really needed? I found it useful to have several examples, for two reasons. First, to verify that the theory is expressive enough, and simple enough, that it can be used in a number of situations. Second, for a given situation, there could be a number of explanations: for instance, the rigidity of prices might come from fairness (Eyster et al. (2020)). But fairness, presumably, cannot explain why consumers do not wish to index contracts. Likewise, the lack of indexation might be due that the bank, say, knows more than the consumer (Spier (1992), Hartman-Glaser and Hébert (2020)), but that does explain why Whole Foods would have uniform delivery prices across time windows. So it is useful to show how complexity sheds light on outcomes across a domain of situations.

I view these results as potentially useful since they give the hope that one can explain existing institutions (and their degree of "simplicity"). The underlying model yields a principled way to do "behavioral mechanism design", i.e. design optimal institutions (i.e. price or tax systems) that aim both at high allocative efficiency and not too high complexity. This will augment the power of mechanism design, which is often hampered by the fact that it generates too complex recommendations to be used in practice.

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)).[3] 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

[3]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.

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 precisely, 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 (2024) 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)).[4] 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),

[4]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.

Dari-Mattiacci and Marotta-Wurgler (2022)).[5] 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 simple prices, the fact that people avoid deciding at all, simple tax systems, aversion to indexation, 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 Complexity aversion: the core model

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

[5] 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.

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:[6]

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

where $m_i \in [0,1]$ is the precision of the signal, and $\sigma^2_{\eta_i} = m_i \left(1 - m_i\right) \sigma^2_{x_i}$.[7] If all the shocks are jointly Gaussian, we have $\mathbb{E}\left[x_i | x_i^s\right] = x_i^s$. The rational case corresponds to $m_i = 1$, and the extremely inattentive case, to $m_i = 0$.[8] The s signifies "using subjective perceptions", 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 $a^r\left(x^s\right)$. Linearizing:

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

The traditional utility losses from an imperfect decision are then:

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

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 $C^{\text{Cognition}}(m) = \sum_i C_i^{\text{Cognition}}(m_i)$.[9] I normalize $C^{\text{Cognition}}(0) = 0$.

[6] This involves a choice of the "basis" or "frame" for the x_i's. Fortunately, this is recoverable from data on principle, see Section A.1.

[7] 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 A.2 gives the derivation.

[8] 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, 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 $\mathbb{E}\left[y_1 | y_1^s\right]$.

[9] Obvious variants include functions of that sum.

2.2 Complexity aversion: The innovation in this paper

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

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

where $a^s = a^r(x^s(x, m))$ 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 (which comes from (81)).

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

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

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 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

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

and the typical mistake size is[10]

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

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

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

possible,I set $\bar{\xi} = \xi\kappa_a^u$, where ξ is a unitless that can be the same across problems, and κ_a^u 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 $\kappa_a^u = |\sigma_a u_{aa}|$, 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 $C^{\text{CA}} = \bar{\xi} e = \xi\kappa_a^u \sigma_{a^\delta}$, hence:

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

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.[11] 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.[12] Hence, there is an incentive for the planner to create a "simple scheme", that will not create too 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 $\left|a_{x_i}^r\right|$. 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 $\left|\sigma_a u_{aa}\right| \sigma_{x_i}$ makes the penalty scaled correctly with incentives of the problem.

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

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

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

aversion is:

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

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

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

2.3 Discussion: Why this formulation?

Complexity aversion and mistake aversion Complexity lowers utility via two channels in (4). First, that it requires extra thinking (or in other language, more information). Second, if the agent makes a mistake (or can fear she's making a mistake), complexity leads to first order "mistake aversion cost".

One can have the first without the second: suppose that entering the decision requires each time filling a long form; this is a transaction cost akin to cognition cost. But, if the task is straightforward enough (though burdensome), there will be no mistake, and no mistake aversion cost.

One finding in this paper is that the "mistake aversion" term, is actually the central term to explain simplicity: otherwise, generically, we do not get full simplicity.

Why model complexity aversion as first order mistake aversion? We'll see that this "first order" is basically necessary and sufficient to create "simplicity".[13] More deeply, this first order mistake aversion for mistakes in the inner world of the decision maker, might be seen as the natural correlate of Kahneman and Tversky (1979)'s first order risk aversion for gambles in the outside world. Recall that in decision-making, people tend to reject a gamble yielding with equal probability +\$10 or −\$9. More generally, they reject a gamble with standard deviation σ when the risk premium is less than $\lambda\sigma$, for some λ around 0.2 to 0.6 (Kahneman and Tversky (1979)). This "first order risk aversion" (where the risk penalty is proportional to the standard deviation), which has been found many times to describe people better than rational model's prediction of second order risk aversion (where the risk penalty is proportional to the squared standard deviation), see e.g. Rabin (2000).

Other conceivable alternatives to CA, and why I didn't choose them *First order complexity aversion vs traditional regret theory.* A variant studied in the appendix (and then ultimately discard) is "second order complexity aversion", which which is proportional to

[13] This is true for integer functional forms. Strictly speaking, the analysis will show that any penalty e^α with $\alpha \in [0, 1]$ would deliver full simplicity.

the expected loss from mis-optimization:

$$C^{\mathrm{CA},2} = \xi \mathbb{E}\left[u\left(a^r, x\right) - u\left(a^s, x\right)\right] = \xi \sum_i \mathbb{E}\left[-\frac{1}{2} u_{aa}\left(a^r_{x_i} x_i\right)^2\right]\left(1 - m_i\right) \tag{9}$$

It is in the tradition of regret, e.g. Loomes and Sugden (1982), Sarver (2008). It is more natural given the tradition, but ultimately I will discard, as it typically does not generate full simplicity, e.g. rigid prices, or no indexation (see Proposition 3 below).

The key difference between "complexity aversion" (proposed here) and "regret aversion" (from the tradition) are:

1. CA is first order (the loss is proportional to the error e, not e^2: otherwise, second order losses will not robustly lead to full simplicity (e.g. no indexation); while uncertainty aversion is second order. This is crucial, as this will allow CA to generate "robust simplicity", i.e. exactly 0 indexation in contracts (if ξ is large enough), and in general, 0 reaction to a number of features. In contrast, uncertainty aversion will typically generate a some sensitivity—contradicting the "stylized fact" of pervasive "simplicity".

2. CA applies to the mistakes made by the agents in their mind, not in nature. It is not to the ex post regret because of nature's moves. Hence, CA does not change the utility of a rational agent (who makes no mistakes), unlike ambiguity aversion.

3. CA applies features by feature: if there are several features (e.g. the decision entails looking at a tax rate, and an expected return, and a price), each of them is subject to CA, so the loss is $\xi \sum_i w_i \left(1 - m_i\right)$ for some weight w_i, rather than in a "unified manner". To see why this matters, consider the alternative, e.g. having e.g. $CA' = \xi \sqrt{b_i^2 \sigma_{x_i}^2 \left(1 - m_i\right)}$. Then, the partial impact of each b_i in b_i^2,not $|b_i|$: so de facto we have second order complexity aversion at the feature level, and we will not get full simplicity.

Information cost. This approach doesn't work: we don't get full simplicity, as we shall see (e.g. Proposition 3). This also includes the *Fixed cost of thinking*, paid optionally, without complexity aversion, which is just variant of the information cost.

Maximin and taste for robustness. That has been a productive enterprise (Gilboa and Schmeidler (1989); Hansen and Sargent (2008)). But in many cases, this doesn't quite the effect: indeed it goes to "pessimism" rather than "taste for simplicity". For instance, in the indexation example below, if the agent knows that $cov(x, y)$ could be negative, then (without CA) the maximin value is not to index. But if he knows that this $\rho_{x,y} \geq \underline{\rho}$ for some $\underline{\rho} > 0$, then he still want to index. Typically, "maximin" demands very extreme beliefs to yield

full simplicity. Also, in the "Whole Foods" example, it's hard to use maximin: if the cost if higher at 6pm, should I be worried?

More general version of the complexity aversion term in this paper:

$$\tilde{C}^{CA} = \xi \left|\sigma_a u_{aa}\right| \sum_i \left(\frac{\left|a_{x_i}\right| \sigma_{x_i}}{\sigma_a}\right)^{\beta} (1 - m_i)^{\gamma} \tag{10}$$

In the main model, I took $\beta = \gamma = 1$. To get full simplicity, any $\beta \in [0, 1]$ and $\gamma \geq 0$ would work (so the CA cost is large enough to compensate for a first order gain).[14] I took $\beta = 1$, because it "scales linearly", so that the penalty is in line with Kahneman and Tversky (1979)'s first order risk aversion. But any $\beta \in [0, 1]$ would deliver "full simplicity". For instance, $\beta =$ means "any feature that I don't know how to react to me generate a fixed cost of dread, independently of the economic importance of the feature"; while when $\beta = 1$, the cost is proportional to the economic. importance of the feature.

I took $\gamma = 1$, for simplicity. One could image other values of $\gamma \geq 0$, using some experimental method. The limit $\gamma \searrow 0$ is particularly convenient. Then, we're in a dichotomous world where "obvious things" ($m_i = 1$) have no penalty, and not completely obvious things ($m_i < 1$) incur a full penalty. This leads to a somewhat simpler model, where one may take (for some features i) the limit $m_i = 1 - \varepsilon$ with $\varepsilon \searrow 0$, so that the CA penalty is simply (with $\beta = 1$)

$$C^{CA} = \xi \left|\sigma_a u_{aa}\right| \sum_i \left|a_{x_i}\right| \sigma_{x_i} 1_{i \in I^{\text{Complex}}} \tag{11}$$

where $I^{\text{Complex}} = \{i : m_i < 1\}$ is a set of "complex" features. Compare to (5), this is the same, except that $1 - m_i$ is replaced by an indicator function for "is complex, or not". For for the rest of the analysis, one can still use the case $m_i = 1$ (though m_i is an infinitesimal quantity shy of 1). I actually recommend this model as the "simplest" model of CA. Still, to make the analysis more orthodox, I took the case $\gamma = 1$, i.e. (5).

Loss aversion for risk in the world. I'm closest to it, as first order complexity aversion is "loss aversion for randomness in one's mind", rather than randomness in the objective world. First, it might work for some cases, e.g. the baguette example. However, it doesn't seem to apply well do the other examples, e.g. lack of indexation: to use loss aversion, one would need to segregate the two variables ($\hat{y}$ and x), and apply loss aversion to them. Likewise, for the cost of inflation, the loss aversion model (perhaps expressed in nominal term) would predict that people like deflation – after all, my purchasing power has increased.

[14] I keep the $\sigma_a^{2-\beta}$ scaling so that the expression is invariant when changing of the units of a (as well as each x_i).

3 Aversion to indexation

I present the application of the general model to a simple situation: the aversion to indexation. It provides the key economic lessons, valid in later applications.

When a risk-averse agents deals with a risk-neutral firm, in basic rational economic theory, contracts should be indexed, to remove some extraneous risk from the consumer: for instance, debt or mortgage contracts should be indexed on inflation, or other macroeconomic variables (e.g., Holmström (1979)). However, the stylized facts on indexation are arguably the following. (i) Most consumer contracts are not indexed, contradicting the rational prediction (ii) On the other hand, in extreme cases, e.g. hyperinflations, contracts become indexed again.

Intuitively, the reason for the lack of indexation is that an indexed contract would be too "complicated", compared to the traditional gains (lower risks). I formalize that now.

3.1 Basic analysis

There are three periods, $t = 0, 1, 2$. Consumption happens are times 1, 2, and utility is

$$V = v(c_1) + v(c_2), \qquad v(c) = \left(c - \frac{\Gamma}{2}c^2\right)/\Gamma \tag{12}$$

The consumer receives income at time 1 only, which is

$$Y_1 = 2(\bar{y} + \hat{y})$$

where $\hat{y}$ is a mean 0 random variable. The real interest rate is 0.

At time 0, the consumer can choose to "index", i.e. to enter into a contract adding $-2bx$, which is received at time 2, where b is the consumer's choice of indexation, and x is random variable, with mean 0 and variance 1, positively correlated with $\hat{y}$. So, his full intertemporal income is

$$Y(b) = Y_1 = 2(\bar{y} + \hat{y} - bx)$$

The action a is the time-1 consumption. So full utility is

$$u(a, b) = v(a) + v(Y(b) - a) \tag{13}$$

The value of $\hat{y}$ and x is announced at time 1. The rational decision at time 1 is $a^r = \frac{Y}{2}$, i.e.

$$a^r = \bar{y} + \hat{y} - bx \tag{14}$$

The expected utility is

$$u^r = \mathbb{E}_0 [u(a)] = 2v(\bar{y}) - var(\hat{y} - bx) \tag{15}$$

Hence, to maximize utility $\mathbb{E}[u(a,b)]$ utility from the point of view of time 0, a rational consumer would minimize the variance of income, i.e. choose the rationally optimal indexation level: $b^r = \text{argmin}_b\, var(\hat{y} - bx)$, i.e[15]

$$b^r = cov(\hat{y}, x) \geq 0 \tag{16}$$

But intuitively, a behavioral consumer will dislike the additional complexity of having to think through bx at time 1 (the values of $\hat{y}$, x are realized at time 1). To model this, let us say that at time 1, his time-1 income y_1 is visible at no cost; but anticipating the value of the hedge $(-2bx)$ is effortful and mistake-prone: it is knowable at time 1, but received at time 2, [16] The perceived value of the hedge at time 1 is $x^s = mx + \varepsilon$ (with ε uncorrelated with x and $\hat{y}$, and $\sigma_\varepsilon^2 = m(1-m)\sigma_x^2$, see (2)). Hence the consumer chooses the time-1 consumption:

$$a^s = \bar{y} + \hat{y} - bx^s$$

So, his utility is

$$V(b) = \mathbb{E}[u(a^s, b)] - C^{CA} \tag{17}$$

The utility is more explicitly as follows.

Proposition 1. (Indexation example: Consumer's utility with complexity aversion) *The expected utility of the consumer choosing indexation level b is:*

$$V(b) = V(0) + \underbrace{(2b^r b - b^2)}_{\substack{\textit{Allocative gains} \\ \textit{from indexation}}} - \underbrace{(1-m)b^2}_{\substack{\textit{2nd order losses} \\ \textit{from mistakes}}} - \underbrace{2\xi\sigma_a|b|(1-m)}_{\substack{\textit{1st order losses} \\ \textit{from CA}}} \tag{18}$$

where $V(0) = 2\mathbb{E}v(\bar{y} + \hat{y})$ is the utility without indexation, $b^r = cov(\hat{y}, x)$ is the optimal hedge for a rational consumer, and m is the consumer's attention to the future income from the hedge.

The term $2b^r b - b^2$ is the rational benefit of indexation for a rational consumer, and it is maximized for $b = b^r$: more generally, this represents the traditional allocative gains from having a more complex scheme. The term $-(1-m)b^2$ is the loss coming from the fact that

[15] To be complete in that first example: $var(\hat{y} - bx) = \sigma_y^2 - 2cov(\hat{y}, x)b + b^2\sigma_x^2$. Optimizing over b, using $\sigma_x = 1$, gives the announced result, $b^r = cov(\hat{y}, x)$.

[16] One could expand to the case where $\hat{y}$ also is hard to see, without changing the conclusions.

the consumer doesn't understand the hedge well at time 1, hence doesn't choose the optimal consumption. The term $2\xi\sigma_a |b| (1-m)$ is the complexity aversion penalty.

Proof. I derive it slowly, so that everything should be very clear. The full consumption utility is in (17), which can also be written as:

$$V(b) = u^r + \frac{1}{2}u_{aa}\mathbb{E}\left[(a_x x)^2\right](1-m) - C^{CA} \tag{19}$$

i.e. the full rational utility, minus the misoptimization costs due to limited attention m, minus the complexity aversion dread. We calculate them in turn. The rational decision (14) is $a^r = \bar{y} + \hat{y} - bx$.

The rational utility is in (15)

$$u^r = 2v(\bar{y}) - var(\hat{y} - bx) = 2v(\bar{y}) - \left(\sigma_y^2 - 2bcov(\hat{y}, x) + b^2\right) = V^0 + 2bb^r - b^2$$

where $V^0 \coloneqq 2\mathbb{E}v(\bar{y} + \hat{y})$ is the utility with no hedge. Here $2bb^r - b^2$ are the utility gains of the hedge, maximized at $b = b^r$.

The misallocation term is (given that $a_x = -b$, and $u_{aa} = 2v''(c) = -2$):

$$\frac{1}{2}u_{aa}\mathbb{E}\left[(a_x x)^2\right](1-m) = -b^2(1-m) \tag{20}$$

The complexity aversion term is (5) with just one variable x :

$$C^{CA} = \xi |\sigma_a u_{aa}| |a_x^r| \sigma_x (1-m) = 2\xi\sigma_a |b| (1-m) \tag{21}$$

Hence, the full utility obtains by added the three terms in (19). □

We move on to the optimal indexation level b.

Proposition 2. (How complexity dampens or annuls the incentive to index) *Suppose first order complexity aversion, and exogenous attention m. The optimal indexation level is*

$$b = \frac{\max(b^r - \xi\sigma_a(1-m), 0)}{2-m} \tag{22}$$

Hence, the consumer chooses not to index ($b = 0$) iff

$$b^r \leq \xi\sigma_a(1-m) \tag{23}$$

i.e. iff the traditional allocative benefits of indexation ($b^r = cov(\hat{y}, x)$) are less than the complexity aversion costs of indexation ($\xi\sigma_a(1-m)$).

Put simply, the consumer does not index iff complexity aversion ξ is large enough, and $m < 1$.

Proof. We use (18). The FOC is

$$0 = \frac{1}{2}V_b = b^r - (2-m)\, b - \xi\sigma_a \operatorname{sign}(b)\,(1-m)$$

where I use $\operatorname{sign}(b) = \frac{d}{db}|b|$, which is the generalized derivative of the function $|b|$. So it is 1 if $b > 0$, -1 if $b < 0$, and some number between -1 and 1 if $b = 0$.[17] Hence, if $b > 0$

$$b = \frac{b^r - \xi\sigma_a\,(1-m)}{2-m}$$

But this require $b^r > \xi\sigma_a\,(1-m)$. If that inequality fails, then $V_b < 0$ for $b > 0$. As $V_b > 0$ for $b < 0$, the optimal b is then $b = 0$. So the solution is (22). □

Proposition 2 is qualitatively consistent with the main facts, and indeed may explain why most contracts (e.g. debt contracts) are not indexed.. When the traditional allocative benefits of indexation ($cov\,(\hat{y}, x)$) are small, and complexity aversion ξ is large enough, there is no indexation. However, when stakes increase (e.g. in hyperinflations), then there is some indexation.

Why pure "information economics" doesn't yield full simplicity, and we need first order complexity aversion If there's no complexity aversion ($\xi = 0$), the optimal index is (via (22)):

$$b = \frac{b^r}{2-m} \in \left[\frac{b^r}{2}, b^r\right], \qquad b > 0 \tag{24}$$

So, importantly, there is always some indexation, $b > 0$. This means that a pure "information cost" model (with exogenous or endogenous m) cannot generate "full simplicity" (i.e. $b = 0$). It only creates dampening compared to the rational case, $b < b^r$: this is because the misperception of the index introduces noise in the decision, hence the index is less desirable. Likewise, if we had only second order complexity aversion, things would not work: more generally, with $C^{CA} = kb^\gamma$ with $\gamma > 1$ (with $\gamma = 2$ for second-order complexity aversion), then we would have $V'(b = 0) > 0$, so $b > 0$ is optimal. Hence, we proved the following.

Proposition 3. (Information cost, or second order losses, do not generate full simplicity). *Suppose that instead of C^{CA}, we have an information cost to acquire attention m, or a penalty proportional to $|b|^\gamma$ for $\gamma > 1$. Then, there is always a positive amount of indexation, $b > 0$.*

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

On the other hand, we see that any penalty $C^{CA} = kb^{\gamma}$ with $\gamma \in [0,1]$ (where $\gamma = 0$ is a fixed cost) would work, i.e. generate $b = 0$ when k is large enough. The reason why I advocate first order complexity aversion ($\gamma = 1$) is that is that this way, it scaled one for one with stakes, and it has indirect support via Kahneman and Tversky (1979)'s first order risk aversion.

This example made the point that complexity-averse will not want to index, under a wide range of parameters. More generally, they will not want to contract on some observables, which might be a way to generate incomplete contracts, in the spirit of Hart (1995) and Segal (1999), but with a more explicit cognitive structure, which would lend itself to measurement. I defer this exploration to a future iteration of this project.

3.2 Extension: Endogenous attention

We can extend the model to endogenous attention. Then, the problem is to maximize the same utility (17)-(18), minus cognition costs $C^{\text{cognition}}(m)$. So, the value function is (using 18).

$$\max_{b,m} V(b,m) = 2b^r b - (2-m)\, b^2 - 2\xi\sigma_a |b| (1-m) - C^{\text{cognition}}(m) \tag{25}$$

The following proposition gives a necessary and sufficient condition for $b = 0$ to be optimum, with endogenous attention.

Proposition 4. (Indexation problem, with endogenous attention). *The jointly optimal choice of indexation and attention m is to take $b(m)$, the value in (22), and m to maximize $\mathcal{V}(m) \coloneqq V(b(m), m)$, which is:*

$$\mathcal{V}(m) = b(m)^2 (2-m) - C^{cognition}(m) \tag{26}$$

Hence, if cognition costs are high enough, in the sense that for all m,

$$b(m)^2 (2-m) \leq C^{cognition}(m)$$

then the optimal indexation is 0.

This means that we have zero indexation, with endogenous attention, if the cognition costs large enough. Again, complexity aversion is crucial. Without it, when $m = 0$, some indexation is better than none.[18]

[18]To see why, consumer the completely naive consumer with $m = 0$. The index will not help him at time 1, as he does not take the index into account. But it will help him at time 2, the index will decrease the volatility of his final consumption.

The online appendix (Section A.3) complements the discussion. In particular, it shows that with linear costs $C^{\text{cognition}}(m) = \kappa m$, *the indexation chosen is either zero indexation or full indexation.*[19]

This indexation example illustrated the core economics, with a minimum of math. I now turn to a richer, market setting.

4 Simple pricing schemes

I apply the general model to a slightly more richer situation: the pricing of one good. The motivation is that many pricing schemes look "simple", compared to a frictionless benchmark. For instance, prices for Whole Foods delivery are the same for all hours and days of the week, even though they "should" vary in the frictionless world, as the marginal cost to Whole Foods must be different across the slots. Why does Whole Foods adopt this very simple pricing scheme? Most likely, this is because of some form of complexity aversion by the consumer.[20] This sort of situation is extremely prevalent: for instance, restaurant typically (not always) charge the same price for meals are different hours of the evening; bars do not implement "congestion pricing".

4.1 Setup and basic analysis

A competitive firm (maximizing social welfare) sells to a behavioral consumer, who can buy a continuous quantity of its good.[21] The marginal cost is an exogenous $c(x) = 1 + x$ where x is a random variable with mean 0 and variance σ_x^2. The endogenous price per unit is is $p(x) = 1 + bx$, where b is chosen by the benevolent firm, i.e. the planner. If consumers were rational, the efficient value would be to have the price equal marginal cost, i.e. $b = 1$, but with behavioral consumers the firm might choose another b. For instance, if $b = 0$, the price is constant over time, which "simplifies" the consumer's choice.

The consumer chooses to consume a quantity $C = 1 + a$, where a will be a deviation of consumption from a value of 1, and will have a mean 0. Her total utility, net of the cost of

[19] This is because then, the optimum attention is either 0 or 1.

[20] This could also logically be a form of complexity aversion by the producer – i.e. some employee at Whole Foods – who does not know all the marginal costs. The modelling would be similar as the one we proposed here, and might apply to the situations documented in DellaVigna and Gentzkow (2019). It couldn't plausibly be a monetary "menu cost" for Whole Foods – their sales are so enormous that any cost of changing the web site should be comparatively trivial.

[21] Here, I take the benchmark of firms forced to maximize social welfare by competitive forces, hence to be "benevolent". Section A.4 gives a microfoundation. This is the opposite tack of the literature where firms exploit the naive consumers' bounded rationality (Gabaix and Laibson (2006), Heidhues and Kőszegi (2018a)), because there is a "shrouded attribute" (e.g. the spare part to be bought later) in addition to a "visible base good" (the car). Considering that those issues are now well-understood, I focus complexity with a benevolent firms, and discuss the case of exploitation to Section A.6.

good, is:

$$u(a,x) = U(1+a) - p(x)(1+a) \tag{27}$$

where the gross utility from consuming a quantity $1+a$ of the good is:

$$U(1+a) = 1 + a - \frac{1}{2\psi}a^2 \tag{28}$$

where $\psi > 0$ is the elasticity of demand (when the price is centered at 1). The rational action is $a^r = \text{argmax}_a\, u(a,x)$ i.e.

$$a^r = -\psi b x \tag{29}$$

4.1.1 The behavioral consumer's utility

Imperfect cognition The consumer perceives the price as $p^s = 1 + bx^s$ with

$$x^s = mx + \varepsilon \tag{30}$$

with $\sigma_\varepsilon^2 = m(1-m)\sigma_x$ and $m \in [0,1]$ is the precision of the perception (exo- or endogenous). She takes the action $a = -\psi b x^s$, see (29). Her utility gains from taking the right action (i.e. $a^r = -\psi b x$) is:

$$\chi = -\frac{1}{2}u_{aa}\mathbb{E}\left[(a_x^r x)^2\right] = \frac{1}{2}\psi\sigma_x^2 b^2 \tag{31}$$

Traditional consumption utility is increasing in attention[22'23]:

$$U^{\text{cons}} = \mathbb{E}\left[u(a,x)\right] = \frac{1}{2}\psi\sigma_x^2 b^2 m.$$

Complexity aversion With first order complexity aversion (5), we have the dread cost, of complexity aversion cost:

$$C^{\text{CA}} = \xi\mathbb{E}\left[(\sigma_a u_{aa} a_x^r x)\right]^{1/2}(1-m) = \xi\sigma_a\frac{1}{\psi}\psi\left|b\right|\sigma_x(1-m)$$

[22]The derivation is:

$$U^{\text{cons}} \quad = \quad \mathbb{E}\left[u(a,x)\right] = \mathbb{E}\left[u(a^r,x)\right] + \frac{1}{2}u_{aa}\mathbb{E}\left[(a_x^r x)^2\right](1-m) = \chi - \chi(1-m) = \chi m$$

[23]It may be surprising that a more volatile price makes the consumer better off, even though the consumer is risk-averse (U^{cons} increases with $\sigma_x^2 b^2$). This is always true with quasi-linear utility $u(c_0,\ldots,c_n) = c_0 + U(c_1,\ldots,c_n)$, with U concave, which is the case here. This is not true for general utilities, as can be seen for the one-good case: the indirect utility is $u\left(\frac{w}{p}\right)$, which shows the tension between the convex $\frac{1}{p}$ and the concave u in determining whether indirect utility is locally convex or concave in p.

hence

$$C^{\mathrm{CA}} = \xi \sigma_a \sigma_x |b| (1 - m) \tag{32}$$

The crucial part is that, when $m = 0$ (when the agent doesn't think), the dread term $C^{\mathrm{CA},1}$ is strictly positive (even though cognition cost is zero, $C(0) = 0$). Specifically, it is proportional to the expected absolute mistake from not thinking.[24] Also, in (32), if the consumer makes no mistake, $m = 1$, the dread term is 0: there is no cognitive dread when all is well-understood.[25]

The consumer's full utility is, from (4):

$$V^{\mathrm{cons}}(m) = \chi m - C(m) - C^{\mathrm{CA}}(m). \tag{34}$$

In the case of endogenous attention, I posit that the consumer optimizes cognition, i.e. does

$$\max_m V^{\mathrm{cons}}(m). \tag{35}$$

This point is not central, but it simplifies the exposition.

4.1.2 Social welfare

Social welfare is:

$$W = V^{\mathrm{cons}}(m) + \mathbb{E}\left[(p(x) - c(x))(1 + a(x^s))\right]$$

i.e. the consumer's full utility (including dread and cognition cost), plus profits from consumption $1 + a(x^s)$, which also reflects that the social cost of the good is c, not p:

So, the social planner's (or benevolent firm's) wants to maximize total social surplus, while taking into account that the variability of prices (indexed by b) affects the endogenous attention m. Mathematically:

$$\max_b W(m(b), b) \tag{36}$$

In the exogenous attention case, $m(b)$ is just constant; with endogenous attention, $m = \mathrm{argmax}_m V^{\mathrm{cons}}(m, b)$, where $V^{\mathrm{cons}}(m, b)$ is (34). We next calculate that social welfare.

[24] To anticipate why this is important, consider a very cognitively limited consumer that exerts no cognition in equilibrium ($m = 0$): the dread terms increases the deadweight loss from the mechanism, as it makes the agent "feel anxious", in addition to making him lose utils from poor decision-making.

[25] With second order CA, the CA penalty (9) is the expected loss from mis-optimization, times the coefficient of complexity aversion ξ:

$$C^{\mathrm{CA},2} = \xi \frac{1}{2} \psi \sigma_x^2 b^2 (1 - m) \tag{33}$$

The interpretation is similar to that of the first order CA (32, but with a term proportional to b^2 rather than $|b|$.

Proposition 5. *Social welfare is*

$$W(m,b) = \underbrace{(b-\frac{b^2}{2})\psi\sigma_x^2 m}_{\substack{\text{Allocative gains from pricing} \\ \text{closer to marginal cost}}} \underbrace{-2\xi\sigma_a|b|(1-m)}_{\substack{\text{1st order losses} \\ \text{from CA}}} - C^{Cognition}(m) \tag{37}$$

We verify that when there is no cognition cost nor complexity aversion ($m = 1$, $\xi = 0$), social welfare is $W = \left(b - \frac{b^2}{2}\right)\psi\sigma_x^2$, so is maximized when $b = 1$, i.e. when the consumer pays the marginal cost.[26]

4.2 When is the optimal price system "simple"?

We now arrive at the meat of the analysis. We can derive the optimum price complexity, particularly focusing on when the price is "perfectly simple" and constant ($b = 0$).

First order complexity aversion reliably generates simplicity Let us start with a transparent case, with exogenous attention m.

Proposition 6. (Optimal price simplicity, first order complexity aversion) *Suppose that consumers have first order complexity aversion. Then, the optimal sensitivity of price to marginal cost is:*

$$b = \max\left(1 - \frac{\xi\sigma_a\sigma_x(1-m) - W_m\frac{dm}{db}}{\psi\sigma_x^2 m}, 0\right) \tag{38}$$

So, with exogenous attention ($\frac{dm}{db} = 0$), the optimum entails rigid prices if and only if:

$$\xi\sigma_a\sigma_x(1-m) \geq \psi\sigma_x^2 m \tag{39}$$

i.e. if marginal complexity aversion cost, $\xi\sigma_a\sigma_x(1-m)$, exceeds that the marginal increase in allocative benefits, $\psi\sigma_x^2 m$.

Proof. The total derivative of $W(m(b), b)$in (37) is:

$$\frac{d}{db}W = (1-b)\psi\sigma_x^2 m - \xi\sigma_a\sigma_x \operatorname{sign}(b)(1-m) + W_m\frac{dm}{db}$$

Solving for b gives the announced expression (38). □

Proposition 6, especially (39) also shows that contracts are more likely to be "simple" (lower b, including $b = 0$) when (i) complexity aversion ξ is high; (ii) when the elasticity of

[26] In particular, there is no motive for the firm to "insure" the consumer. The consumer has quasilinear utility, so is risk-neutral to income shocks; hence, optimality demands that price equal marginal cost, with no further consideration of insurance.

demand ψ is low – hence, if competition is low (in many models, higher competition leads to a higher elasticity of demand ψ, see e.g. Atkeson and Burstein (2008), Gabaix et al. (2016)); when (iii) volatility of the demand would be low ($\psi\sigma_x$ low) even under a frictionless model; (iv) when the degree of understanding is low (low m).

The important point is that we get a simple price system if the cost of cognition is large enough. We see why Whole Foods has a uniform pricing system for delivery, but variable prices between goods: the demand for delivery is reasonably inelastic, while the demand for an individual goods is much more elastic (e.g. raspberries vs blueberries, or raspberries at Whole Foods vs another store), and the marginal cost of perishable goods is very variable, depending on crops for instance. Likewise, in very competitive markets (e.g. airlines), the elasticity of demand is very high, and as a result the schemes are "complex".

The role of complexity aversion cost The next proposition shows that the CA is essential to obtain simple mechanism. Without CA, we do not get outright simplicity ($b = 0$), and instead we get $b = 1$ as the optimum.

Proposition 7. (We need complexity aversion to explain full simplicity) *Suppose that there is no complexity aversion ($\xi = 0$), and there is either exogenous or endogenous attention. Then social welfare is optimized for $b = 1$, so that prices are not rigid at the optimum.*

We conclude that, with first order CA, a high CA parameter ξ generates simplicity, $b = 0$. Let us next verify that this works with endogenous attention.

Deepening the role of the thinking cost with endogenous attention The following gives sufficient conditions for full simplicity.

Proposition 8. (Optimal price simplicity, first order complexity aversion, endogenous attention) *Suppose first order complexity aversion, and endogenous attention m. A perfectly simple contract ($b = 0$) is* locally *optimal (i.e. the social welfare $W(b)$ has a local optimum at 0) if and only if (39) holds. For a perfectly simple contract to be* globally *optimal, the following are sufficient conditions: $C^{cognition}(1) > \psi\sigma_x^2$ (reaching full precision is very costly), and the complexity aversion cost ξ is large enough: $\xi \geq \xi^*$, with:*

$$\xi^* := \max_m \frac{\psi\sigma_x^2 - C^{cognition}(m)}{\sigma_a\sigma_x(1-m)} \tag{40}$$

The intuition is the following: if it's mentally costly enough to fully understand the situation, people will have a limited understanding of it, so $m < 1$. But then, if $b > 0$ and the complexity aversion ξ is high enough, the "dread" term (32) is high enough to wipe out the allocative benefits, $\frac{1}{2}\psi\sigma_x^2$.

The cutoff ξ^* is higher when cognition is cheaper ($C^{\text{cognition}}$ is lower), and when the traditional benefits from a rational allocation ($\frac{1}{2}\psi\sigma_x^2$) are higher (when the rational elasticity of demand ψ and the volatility of prices σ_x are higher).

To complete the picture, let us consider a case where we need not have perfect simplicity: intuitively, this is the case when the marginal cost of thinking are low enough, compared to the social surplus gains of having a complex contract.

Proposition 9. (When full rationality is easy enough, contracts shouldn't be maximally simple*) Suppose that $C^{cognition\prime}(1) < \frac{1}{2}\psi\sigma_x^2$. Then, perfect simplicity ($b = 0$), is non-optimal (and indeed, less desirable than $b = 1$, the optimum slope with traditional agents), even with a high complexity aversion parameter ξ.*

The economic interpretation of the condition is that the (marginal) cost of thinking $C^{\text{cognition}\prime}(1)$ is less than the surplus generated at the optimum, which is $\frac{1}{2}\psi\sigma_x^2$. Hence, some complexity is important if the task is "easy", or a "very important" (it yields a high social surplus $\frac{1}{2}\psi\sigma_x^2$). Hence, we have simplicity when the cost of getting it exactly right is not too large.

This concludes the analysis of first order complexity aversion. We see that when it is large enough (ξ large enough), it generates full simplicity, at least if it is hard enough for the agents to completely understand the situation (Propositions 6 and 8). At the same time, complexity aversion is necessary to account for full simplicity (Proposition 7).

The appendix A.5 analyzes second order complexity aversion, and finds that it can sometimes generates simplicity, but in a non-robust way.

5 Simple tax systems: Ramsey problem with simple taxes at the optimum

The Ramsey (1927) taxation problem is the most basic taxation problem, and still leads to complex optimal taxes: generically, each good has its own specific tax, proportional to the inverse elasticity of demand. This is even true when agents are behavioral, in the sense that they misperceive the tax, but are not complexity averse (Farhi and Gabaix (2020), Rees-Jones and Taubinsky (2020)). This is true in richer tax problems, where generically the tax scheme depends on each good, and each state to the world.

Let us see how complexity aversion lead to an optimal tax system that is simple.

5.1 Traditional, rational version

We start with the traditional framework for goods taxation. There are $G+1$ goods indexed by g. The agent has utility $U(c) = \sum_{g=1}^{G} u^g(c_g) + c_0$ where $u^g(c_g) = c_g - \frac{1}{2\psi_g}(c_g - 1)^2$ and c_0 is the a residual good with price 1 (which is utility from residual income), untaxed. Hence the rational demand is $c_g = 1 - \psi_g(p_g - 1)$, so that ψ_g is a sensitivity of demand of good g to its price.

Each good g is taxed at a rate of τ_g, so the price paid by the consumer is $\tilde{p}_g \coloneqq p_g(1+\tau_g)$. The pre-tax price p_g is independent of taxes, e.g. fixed by a linear production technology, and it is normalized to be $p_g = 1$.

The government has a social welfare function:

$$W = U(c) + (1+\lambda)\sum_g \tau_g c_g \tag{41}$$

where $\lambda > 0$ is the excess marginal value of public funds, which are the collected taxes, $\sum_g \tau_g c_g$.

Following the tradition, we take the limit of a small λ, so that taxes are small. Taking a Taylor expansion up to second order terms, rational utility is

$$W = W^0 + \sum_g \left(-\frac{1}{2}\psi_g \tau_g^2 + \lambda \tau_g\right)$$

where W^0 is the utility with 0 taxes. With rational consumers, optimizing over τ_g gives the traditional Ramsey (1927) result, that the optimal tax is proportional to the inverse of demand elasticity.

$$\tau_g = \frac{\lambda}{\psi_g} \tag{42}$$

Each good has a different tax: the tax system is very "complex" in that sense.

5.2 The case with complexity averse agents

The government sets up n categories of taxes, with tax $\bar{\tau} + \hat{\tau}_{G_k}$ (for $k = 1 \dots n$). The number of categories will be endogenous. The tax in group k is $\tau_{G_k} = \bar{\tau} + \hat{\tau}_{G_k}$, where $\bar{\tau}$ is the average tax and $\hat{\tau}_{G_k}$ is a deviation.

I assume that the consumer sees a tax

$$\tau^s_{G_k} = \bar{\tau} + \hat{\tau}^s_{G_k} \tag{43}$$

where the subjective perception $\hat{\tau}^s_{G_k}$ is an inattentive version of the true $\hat{\tau}^s_{G_k}$: the consumer

sees average tax $\bar{\tau}$ perfectly , but sees the deviation $\hat{\tau}_g$ only partially. The consumer also sees for free the category $k(g)$ of a good g.

The limited capacity to react to taxes is intuitive. Chetty et al. (2009), Taubinsky and Rees-Jones (2017) and Rees-Jones and Taubinsky (2020) provide compelling experimental evidence that people do not react optimally to taxes, and indeed typically under-react to them, consistent with limited attention. Glaeser and Shleifer (2001) propose that quantities rather than prices are sometimes regulated (e.g., no alcohol should be sold on Sundays), when quantities are easier to observe, a form of simplicity. Aghion et al. (2024) find significant frictions to optimal reaction to taxes, consistent with complexity costs.

This model applies directly to the US, where it's customary for firms to show the pre-tax price p_g, but not the price inclusive of the tax, $\tilde{p}_g \coloneqq p_g(1+\tau_g)$. However, it also applies to other countries, where the price inclusive of taxes is shown: simply then, the cognitive difficulty is on the side of firms. Indeed, consider a world where what is shown is the price $\tilde{p}_g$ inclusive of the tax, as in European countries; and where, as in all sensible systems, intermediary production is not distorted, so that firms do not pay the VAT on intermediary inputs. Then, the model still applies, but the complexity is for firms, not consumers.[27] However, the economics are the same.

The consumer's demand is:

$$c_g = 1 - \psi_g \tau_g^s \tag{44}$$

Calling m_g the attention to the special tax $\hat{\tau}_g$ to good g, the (first order) complexity aversion cost associated with good g is as in (32):

$$C_{G_k}^{\mathrm{CA}} = \xi \sigma_{G_k} |\hat{\tau}_{G_k}| (1 - m_{G_k}) \tag{45}$$

where σ_{G_k} is the size modulus for the variation consumption of good g in group (as usual, it will play no role except giving the correct units).

Social welfare is:

$$W = \mathbb{E}\left[\sum_g u^g(c_g)\right] - C_g^{\mathrm{CA}} - C_g + (1+\lambda)\sum_g \tau_g c_g \tag{46}$$

i.e. the consumer's utility (inclusive of the complexity aversion, and the cognition cost C_g^{CA}, plus the government revenues (with a weight $1+\lambda$, where $\lambda \geq 0$).

[27]Suppose that firms see the price $\tilde{p}_g$ but the price they pay is $\frac{\tilde{p}_g}{1+\tau_g}$, as the VAT that a firm pays is rebated to it. Then, a firm h's problem is to maximize over the intermediary inputs X_g^h of good g, $\max_X p_h F^h\left(X_g^h\right) - \sum_g \frac{\tilde{p}_h}{1+\tau_g} X_g^h$. So, it's the firm that now faces the complexity cost of knowing each τ_g. As decision-makers in firms are people, then if prices inclusive of taxes are shown, those people will pay the complexity cost instead of the consumer.Their attention may be higher than for consumers, but (at least for non-giant firms), likely to be imperfect.

Lemma 1. (Welfare in the Ramsey problem with CA agents: Taylor expansion) *In the limit of small taste for government funds λ, hence small taxes, social welfare is, up to higher order terms:*

$$W = \sum_{k;g\in G_k} [\frac{-1}{2}\psi_g \mathbb{E}\left[\left(\bar{\tau}+\hat{\tau}^s_{G_k}\right)^2\right] - \xi\sigma_{c_g}\left|\hat{\tau}_{G_k}\right|\left(1-m_{G_k}\right) + \lambda\left(\bar{\tau}+\hat{\tau}_{G_k}\right)] - \sum_k C^{Cognition}_{G_k} + W^0 \tag{47}$$

where W^0 social when all taxes are zero. Also, $\mathbb{E}\left[\left(\bar{\tau}+\check{\tau}^s_G\right)^2|\check{\tau}_G\right] = \bar{\tau}^2 + m_G\left(2\bar{\tau}\check{\tau}_G + \check{\tau}^2_G\right)$.

The government problem is to maximize social welfare W, subject to the average tax being τ, i.e. $\sum_{G_k} s_{G_k}\hat{\tau}_{G_k} = 0$, so that $\bar{\tau}$ is truly the average tax.

If all consumers are rational ($m_{G_k} = 1$, so $\hat{\tau}^s_{G_k} = \hat{\tau}_{G_k}$), optimizing social welfare (47), we obtain $\tau_g = \frac{\lambda}{\psi_g}$, the traditional Ramsey inverse elasticity rule: the tax system is complex, as each good has a different price.

With complexity-averse consumers, we obtain the following very different result, Proposition 10.

For simplicity, I consider the case with continuum of goods, whose elasticities are uniformly distributed in $[\psi_{\min}, \psi_{\max}]$. I call $\bar{\psi} = \frac{\psi_{\min}+\psi_{\max}}{2}$ the average elasticity, and $v_\psi = \frac{\psi_{\max}-\bar{\psi}}{\bar{\psi}}$ the half-width of the range of taxes. I take a Taylor expansion in the limit of small v_ψ and $\bar{\xi}$, defined below.

We will sometimes consider the endogenous attention case. To simplify, I model it as the agent having costless but finite cognitive resource, so that the attention m to each of the n bracket is some function $m(n)$.

Proposition 10. (Optimal tax simplicity, with a concave cost of cognition) *Consider the Ramsey problem with a continuum of goods, and complexity averse consumers. Suppose that the government implements n tax brackets, chosen optimally and with optimal taxes. When moving from one brackets to n, welfare increases by:*

$$V = \frac{1}{3}\lambda\bar{\tau}v^2_\psi m \max\left(1 - z\xi\frac{1-m}{m}, 0\right)^3\left(1-\frac{1}{n^2}\right) \tag{48}$$

with $z := \frac{\sigma_{c_G}}{\lambda v_\psi}$. Hence:

If agents are complexity averse: (i) when m is small enough and ξ is large enough, the optimum is just to have a uniform tax taxes. (ii) With endogenous attention $m(n)$, the optimum number of brackets (weakly) decreases with the degree of complexity aversion.

If agents are not complexity averse: (iii) with exogenous attention m, there is an infinity of tax brackets (iv) with endogenous attention $m(n)$, there must be at least two tax brackets.

When complexity aversion ξ increases, there are fewer tax buckets; when ξ is high enough, there is just one uniform tax bracket. The situation is shown in Figure 1. With rational

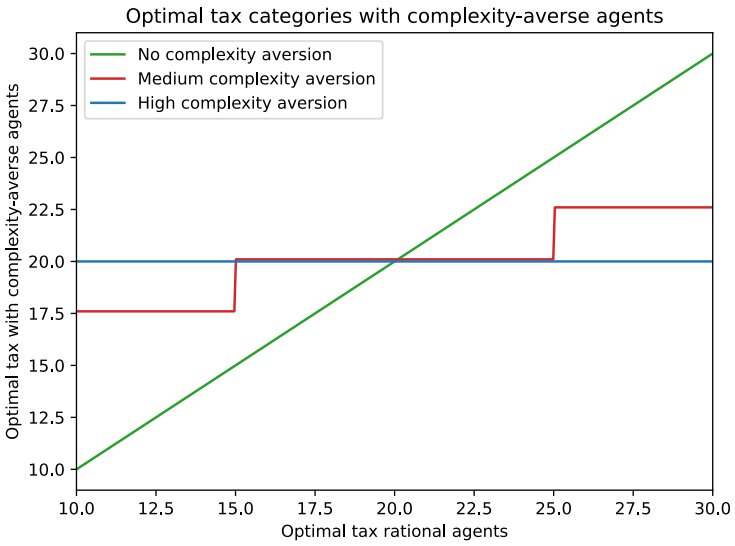


Figure 1: Optimal tax system for three levels of complexity aversion. With rational agents, the optimal tax is just the Ramsey tax τ^R that would be used with rational agents: there are generically as many tax categories as goods. With very high complexity aversion, there is a uniform tax: there is just one tax bucket. For an intermediate complexity aversion, there are three tax buckets.

agents, the optimal tax is just the Ramsey tax $\tau_g^R = \frac{\lambda}{\psi_g}$ that would be used with rational agents. With very high complexity aversion, there is a uniform tax: there is just one tax bucket. For an intermediate complexity aversion, there are three tax buckets in this example.

According to Proposition 10, the optimal number of tax brackets is decreasing of complexity aversion ξ, the cost of thinking κ, and increasing in attention m, the taste for public funds λ, the average optimal Ramsey tax $\bar{\tau} = \frac{\lambda}{\bar{\psi}}$, the dispersion of elasticities v_ψ. This makes intuitive sense: when cognition costs are very high, it's best to have a uniform tax.

In each of tax category, the tax is given by (128). In particular, the taxes are always more compressed (towards $\bar{\tau}$) than under a Ramsey model with the same tax brackets, but no complexity aversion. The proof of the proposition (in Section B) describes the shape of the tax categories. The central bracket, around the average elasticity, has a width $2b^* v_\psi$ with b^* given in (131); and otherwise the other tax brackets have equal size.

The biggest step here was to find finding the conditions (namely, first order complexity aversion) that generates full simplicity.[28] This is not trivial. Indeed, without complexity aversion, if there is a fringe of rational, zero-cost agents, then again we want (without

[28] We started Farhi and Gabaix (2020) hoping to obtain simple tax systems, but we didn't. Optimal taxes in the Ramsey problem were modulated by attention, but were not simple: we were using inattention and costs of cognition, which was not enough. But now the new ingredient of complexity aversion allows to obtain simple taxes.

complexity aversion) a different tax for every good. Furthermore, the proof of Proposition 10 gives intermediary results that are useful to compute numerically the optimum tax scheme, under more general assumptions, e.g. without uniformly distributed elasticities.

Other potential ways to derive a uniform tax, or a small amount of tax brackets. Let us discuss other ways to potentially derive the result.

A penalty $g(n)$ for psychological cost. The simplest is to say that the welfare function is

$$W = W^{\text{rat}} - g(n) \tag{49}$$

where W^{rat} is the rational utility, and $g(n)$ is some exogenous psychological cost of having n brackets; taken to be weakly increasing and convex in n, which implies

$$\lim_{n\to\infty} g(n) = \infty \tag{50}$$

This seems to get the job done. I reject this path, as not scaling correctly at all. Indeed, in life there are lots of small disturbances that change are actions, and the cost associated with them is very bounded. However, the above formulation leads to infinite complexity losses.

In contrast, the formulation with first order complexity aversion is such that even with an infinite amount of "small" costs of size s_n, the total discomfort is proportional to $\sum_n s_n$, which is typically finite. For instance, in the Ramsey problem above, with a continuum of goods, if the planner set up the Ramsey solution, where each good has its own tax, this would not be optimal; but the complexity aversion would still be finite.

If $g(n)$ is not convex, then conceivably $\lim_{n\to\infty} g(n) < \infty$: but then, the increment $g'(n)$ must go to 0. Hence, as the margin, in a sufficiently rich environment, the incremental complexity costs go to zero, and the impact of complexity is 0. This is a self-defeating modelling device.

Administrative costs. They would work, and would be a form of cost $g(n)$. I don't pursue them here, because the conceptually interesting reason here is limited processing power by the people, not by the government. In that perspective, even if administrative powers were infinite, a sage administration would keep the VAT simple and tidy.

Purely limited attention, without complexity aversion. This would not work. Take the exogenous attention case: Proposition 10 shows that then, the amount of tax brackets is actually infinite. Next, take the endogenous attention case: the number of brackets is always at least two. Why? with one tax bracket, attention to the average tax is full. When going away from one tax bracket, there is always some allocative gain, even if attention is very small.[29] Hence, some form of complexity aversion is necessary.

[29] If the attention to the average tax fell, that would be good: the planner's problem is helped by a lower

Let us conclude this analysis of the Ramsey problem. It is meant to be a prototype for muhc more. More generally, this approach yields a principled way to do "behavioral mechanism design", i.e. design optimal institutions (i.e. price or tax systems) that aim both at high allocative efficiency and not too high complexity (see Gonczarowski et al. (2023) for advances in behavioral mechanism design).

This is a fruitful research avenue. One is theoretical. Another is an issue of measurement: how to measure the complexity of rules? Though promising avenue is text analysis, as in Colliard and Georg (2025), perhaps enriched with economic reasoning about the importance of features, as in the present paper.

6 Inertia in decision-making coming from complexity aversion

Very often, people prefer not to decide, at all. For instance, they do not rebalance their portfolio. Complexity aversion offers a natural hypothesis for this: thinking is aversive, especially if one dreads taking the wrong decision.[30]

There could be other explanations. The thinking cost is paid now and the benefits are received later, so a hyperbolic discounting will fail to act (Laibson (1997)). Or, the act of choosing itself could be aversive (Bernheim et al. (2024)), or entail some adjustment cost (see the models reviewed in Gabaix (2019)). Or, choice could active loss aversion, as looking up information can reveal to bad news (Kőszegi and Rabin (2009), Andries and Haddad (2020)). Still, complexity aversion and the dread of mistakes seem like a plausible alternative route, so I briefly explore it here.

Suppose that the active decision is the traditional utility:

$$U^{\text{active}} = \mathbb{E}\left[u\left(a^s\left(x^s\right), x\right)\right] - C^{\text{cognition}}\left(m\right) - C^{\text{CA}}\left(m\right) \tag{51}$$

where a^s is the noisy decision given the imperfect perception x^s. But there is a "passive" decision a^d, evaluated as:[31]

$$U^{\text{passive}} = \mathbb{E}\left[u\left(a^d, x\right)\right] \tag{52}$$

Hence, we can hypothesize that people will stay passive, and do nothing, when $U^{\text{passive}} \geq U^{\text{active}}$. The main psychological difference is that the passive decision does not bear any complexity aversion.

average tax, as this makes taxes less distortionary.

[30] More generally, complexity aversion could potentially explain the phenomenon that people prefer "simple" options to more complex ones (e.g. Iyengar and Kamenica (2010), de Clippel et al. (2024)).

[31] An interesting variant would be: $U^{\text{passive}} = \mathbb{E}\left[u\left(a^d, x^d\right)\right]$, which gives an even higher utility.

To complete this simple psychology, we can add "motivated cognition" with[32]

$$\tilde{U}^{\text{active}} = U^{\text{active}} + \mu \sum_i u_{x_i} x_i^s$$

where $\mu > 0$ is weight on motivated attention. This means that reacting to "bad news" (news x_i that lower utility if $u_{x_i} x_i^s < 0$) is penalized, an "ostrich effect" (Karlsson et al. (2009); Sicherman et al. (2016)).[33]

To see the effects, take the model $\tilde{U}(a, x) = -\frac{1}{2}(a - x)^2 + vx$, so that the passive decision at default action a is $U^{\text{passive}} = -\frac{1}{2}\left[\left(a^d\right)^2 + \sigma_x^2\right]$. Take the case $C(m) = \kappa m$ for simplicity.

Proposition 11. (When the agent fails to act, and passively follows the default) *The agent takes passive course of action ($\tilde{U}^{passive} \geq \max_m U^{active}(m)$) if and only if:*

$$\min\left(\xi\sigma_a\sigma_x, \kappa - \frac{\sigma_x^2}{2} - \mu v x\right) \geq \frac{1}{2}\left(a^d\right)^2 \tag{53}$$

Condition (53) means that people are more likely to be passive if: complexity aversion ξ is high, cognition cost κ is high, the news is "bad news" ($vx < 0$ high), and if the passive allocation isn't too far off the optimum ($\left|a^d\right|$ small). If there is no complexity aversion, the whole effect disappears.[34]

Take the lack of rebalancing of one's portfolio. The active decision involves thinking about the correct determinants for portfolio choice, a complicated vector x, involving e.g. the expected returns of US vs International stocks, or long vs short term bonds. Hence, many people just choose not rebalance at all.[35] They will do so only when the portfolio is very off, and when they have a no complexity aversion.

More generally, this "concern for mistake by taking an active decision" will lead to rigid contracts. Most hedge funds have charged "2 and 20" (2% fixed, plus 20% of the upside). Likewise, sharecropping contracts tend to have a rigid split (often, 50%) between landlord and tenants. The initial split might be arbitrary, and dictated by historical accidents (Young and Burke (2001)). But the persistence then may well stem from first order complexity aversion.

[32]This modeling of motivating cognition is in Gabaix (2019), Section 4.3.2.

[33]Sensibly, real ostriches do not exhibit that effect—Pliny the Elder spread that myth.

[34]In this model with linear cost, if the agent is active, then $m = 1$, so that she fully optimizes. With costly attention, then even under "active" decision, the attention wouldn't be full.

[35]Calvet et al. (2009) finds some rebalancing at the level of individual securities. Gabaix et al. (2025) finds very low "active" rebalancing of a the level of aggregate risk — e.g. in the mix of stocks vs cash.

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

Here I pursue the theme that business cycle costs makes things more complicated. This is a straightforward application of the general framework, but useful 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

$$W = V^{\text{trad}} - \xi\bar{y}^2\,|u''|\,\sigma_y(1-m_y) \tag{54}$$

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

Proposition 12. (Losses from the business cycle) *The consumption loss equivalent is:*

$$\ell = \ell^{Lucas} + \ell^{CA}, \qquad \ell^{Lucas} = \frac{1}{2}\gamma\sigma_y^2, \qquad \ell^{CA} = \omega_y\sigma_y, \qquad \omega_y \coloneqq \gamma\xi(1-m) \tag{55}$$

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.[36]

In the online appendix (Section A.8), I detail some potential variants with idiosyncratic risks. With only large idiosyncratic risks, the loss becomes again quadratic in σ_y. But with 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 $\beta \in [1,2)$).

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

8.1 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

$$y^{\$}_{n+h} = p_n + \frac{1}{2}\pi_{n+1} \tag{56}$$

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.[37]

Figure 2 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

[36] 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.

[37] It would be easy to make inflation uncertain – then in the income process (56) we'd replace π_n by its expectation.

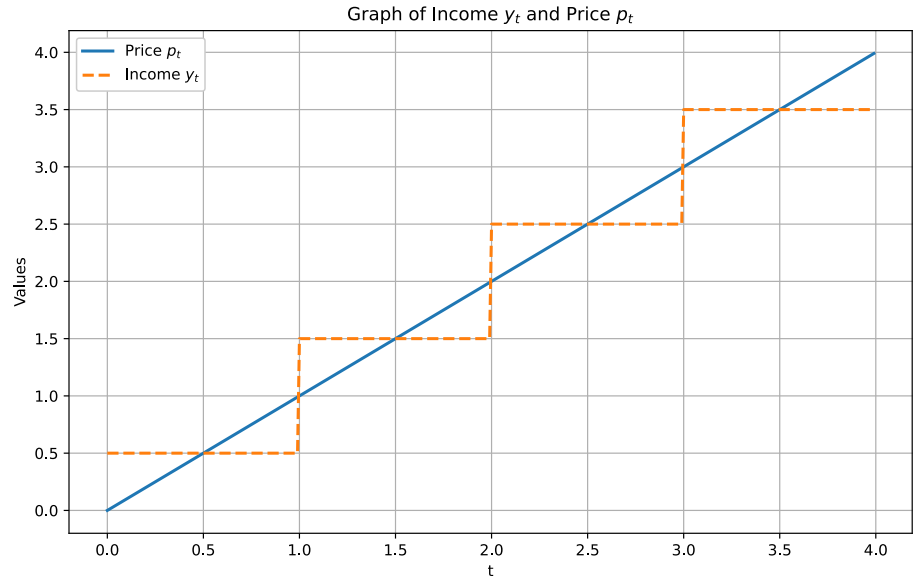


Figure 2: The price increase continuously, but the agents's nominal income increases only once a year (those quantities are in logs). That makes the planning process more complex, as the theory develops.

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)

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

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 13. (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:*

$$\ell^{CA} = \omega_\pi \left|\pi\right|, \qquad \omega_\pi = \frac{1}{4}\gamma\xi\tau\left(1 - m_\pi\right). \tag{57}$$

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 9. 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[38]

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

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

$$W_t^{CA} = -\omega_\pi |\pi_t| - \omega_x |x_t| \tag{58}$$

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

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

[38] 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 $\ell^{\text{CA}} = 5\%$.

Hence, we get a quadratic extra term,

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

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

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

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.

8.2 Complexity-aware welfare in the New Keynesian model

8.2.1 Economic environment

We take a simple New Keynesian model. We take the model in Gabaix (2020), which generalizes Woodford (2003) and Galí (2015), to allow for partially myopic agents, which is relevant for monetary policy, including this analysis.[39] It microfounds the following the behavior of the output gap x_t and inflation π_t:

$$x_t = M\mathbb{E}_t[x_{t+1}] - \sigma\left(i_t - \mathbb{E}_t\pi_{t+1} - r_t^n\right) \text{ (IS curve)}, \tag{59}$$

$$\pi_t = \beta M^f \mathbb{E}_t[\pi_{t+1}] + \kappa x_t + \nu_t \text{ (Phillips curve)}, \tag{60}$$

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 (60) 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.

[39] 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.

8.2.2 Welfare with complexity costs

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

$$W^{\text{trad}} = -K^{\text{trad}}\mathbb{E}_0 \sum_{t=0}^{\infty} \frac{1}{2}\beta^t \left(\pi_t^2 + \vartheta^{\text{trad}} \left(x_t - x_*\right)^2\right) + W_-, \tag{61}$$

for positive coefficients K^{trad}, ϑ^{trad}.[40] 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 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 14. *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:

$$C_t^{CA} = \omega_\pi \left|\pi_t\right| + \omega_x \left|x_t\right|. \tag{62}$$

for positive coefficients ω, and W^{trad} is the traditional welfare loss (61), 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, $\sum_t \beta^t C_t^{CA}$.[41]

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

$$W_t = K\left(-\omega_\pi \left|\pi_t\right| + g\left(x_t\right)\right), \qquad g\left(x_t\right) = \gamma x_t - \frac{1}{2}\theta x_t^2 \tag{63}$$

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.[42]

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.

[40] 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.

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

[42] 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

This dynamic analysis has a partial static antecedent in Orphanides and Wilcox (2002), who propose an objective function $-\pi_t^2 - \theta \left|\hat{y}_t\right|$, 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 (59)-(60). 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 their wage reset each period. The agents don't fully appreciate inflation, as in Section 8.1. Then, we have the following reassuring result.

Lemma 2. *The aggregate behavioral model (59)-(60) still holds in the aggregate, despite wage staggering and partial nominal illusion by workers.*

8.3 Optimal monetary policy with complexity costs

We can now analyze policy – here with commitment, which is arguably the relevant case (the appendix deals with the no-commitment case). 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 (59), so that $x_t = 0$ and $\pi_t = 0$ are solutions).[43]

However, when there is cost-push shock, the $\nu_t \neq 0$ in (60), one cannot get the first best: $x_t = \pi_t = 0$ is not consistent with (60). 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 $\pi = 0$.[44]

Proposition 15. (Optimal response to a cost-push shocks, with commitment: conditions to have 0 inflation as a state by state target). *With welfare (63), given a shock at $t = 0$ the optimal commitment policy is characterized by*

$$\kappa \omega_\pi \operatorname{sign}(\pi_t) = g'(x_t) - M^f g'(x_{t-1}) 1_{t>0} \tag{64}$$

[43] 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.

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

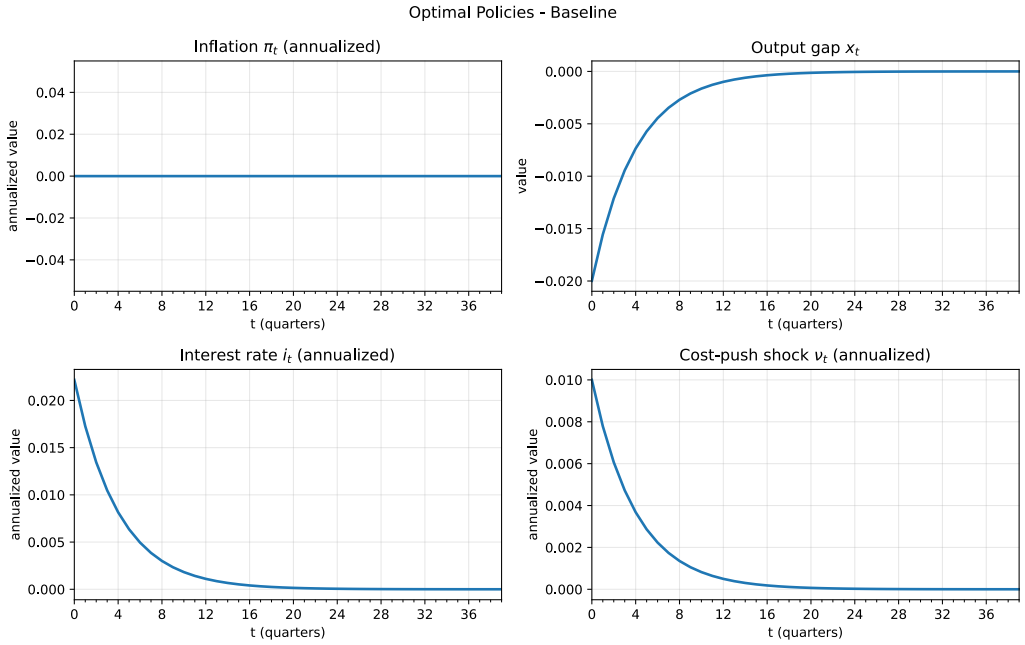


Figure 3: Impulse-response for a small to moderate cost-push shock. Given the exogenous cost-push shock shown in the last panel, the figure shows the resulting economy under the optimal policy. Inflation is pinned down at 0 at all periods.

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

$$\kappa\omega_\pi \geq \left|\Gamma_x \left(x_t^* - M^f x_{t-1}^* 1_{t>0}\right)\right| \tag{65}$$

which is true when the cost-push shocks are small enough. Then the optimal central bank policy under commitment is simply to keep inflation at 0 at all dates,

$$\pi_t = 0 \tag{66}$$

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 (59) holds.*

Consider a small shock, so that 65 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 3 and 4 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.

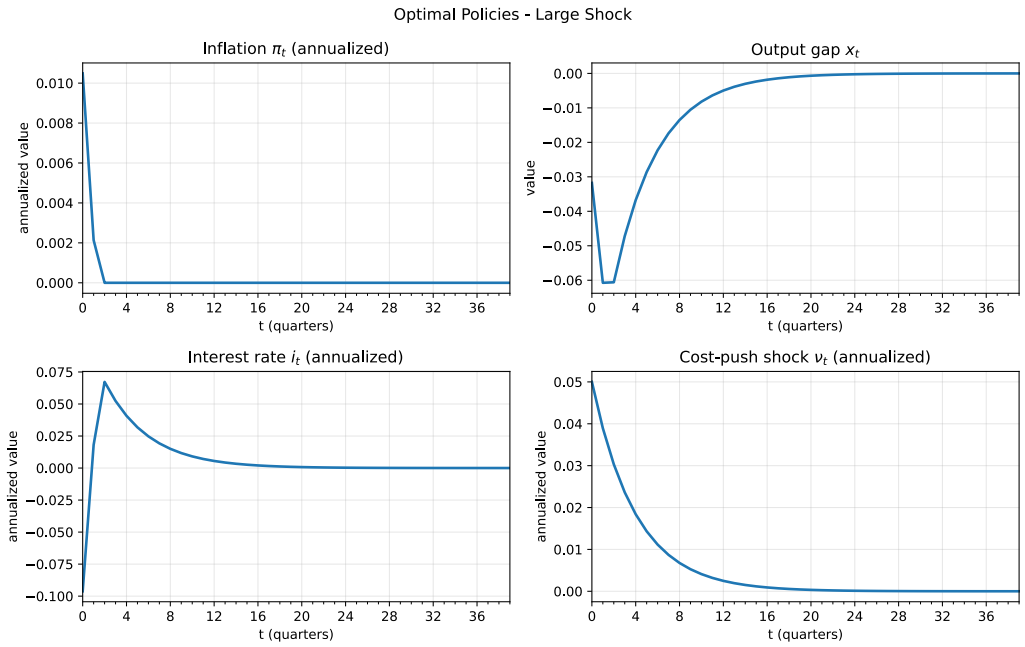


Figure 4: Impulse-response for a large cost-push shock. Given the exogenous cost-push shock shown in the last panel, the figure shows the resulting economy under the optimal policy. This shock is 5 times larger that in Figure 3. Inflation is now above baseline for a little while, then it is pinned at 0 again in all later periods.

Proposition 16. (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,* $\nu \in [\nu_-, \upsilon_+]$, *we set inflation at 0,* $\pi = 0$ *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 15, 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.

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

So far, the optimal inflation target was 0. I now study what happens with non-zero inflation target $\bar{\pi}$. 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)). 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%.

8.4.1 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:

$$f_t = f\left(w_t^{\$} - w_{t-1}^{\$}\right) = f\left(w_t - w_{t-1} + \pi_t\right) \tag{67}$$

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.[45]

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

8.4.2 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

$$f(\pi) = \frac{b\pi\mathcal{L}}{b + |\pi|\mathcal{L}} \tag{68}$$

where $b, \mathcal{L}$ are positive parameters. Figure 5 plots it. So f has a sigmoid shape: morale effects affect productivity by an amount between $-b$ and b.[46] 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.[47] Here, the moral function is a psychological primitive, 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.

[46] 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: $-f(-\infty) = \frac{b}{1-\alpha} > f(+\infty) = \frac{b}{1+\alpha}$.

[47] 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 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.

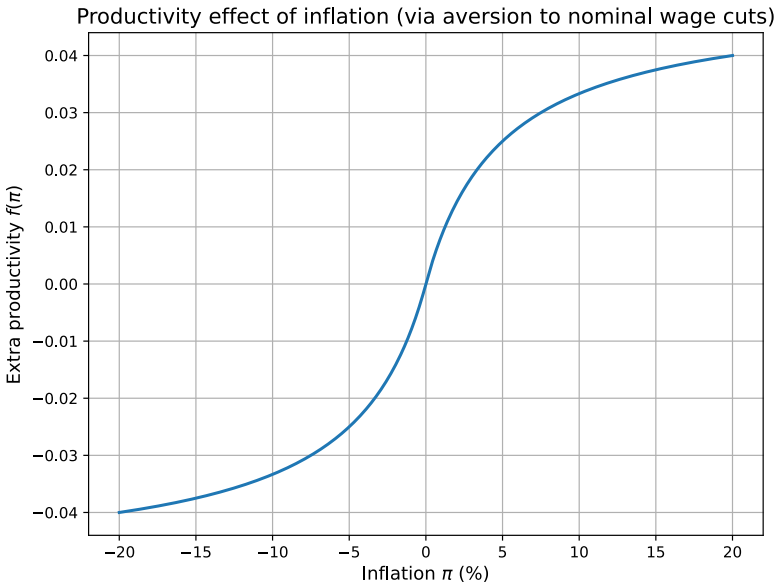


Figure 5: Productivity effect of inflation. This plots the hypothesized productivity effect of inflation: in particular, wage disinflation creates a loss in productivity.

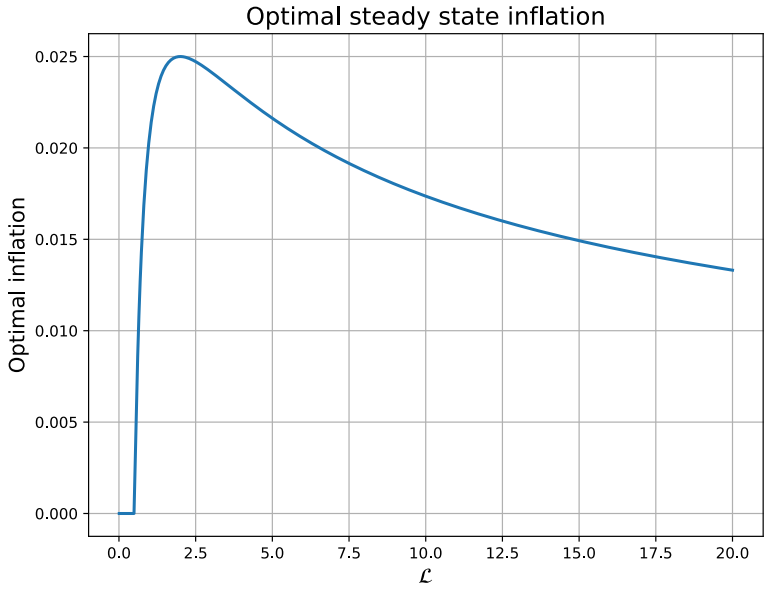


Figure 6: Optimal steady-state inflation. This figure plots the optimal inflation as a function of the slope $\mathcal{L}$ of the "morale function" (68). We see that, above a certain threshold, the optimal inflation is quite insensitive to that slope, yielding some robustness to the analysis.

Proposition 17. (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 (68):*

$$\bar\pi = \frac{b}{\mathcal{L}}\left(\sqrt{\frac{\mathcal{L}}{\omega_\pi}} - 1\right)^+ = \frac{b}{4\omega_\pi} H\left(\frac{\mathcal{L}}{\omega_\pi}\right), \qquad H(\tau) = 4\left(\frac{1}{\sqrt{\tau}} - \frac{1}{\tau}\right)^+ \tag{69}$$

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:

$$\bar\pi^* = \frac{b}{4\omega_\pi} \tag{70}$$

Using the above calibration for the morale function f, and for the welfare losses of inflation (right after Proposition 13), 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.

8.4.3 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:

$$\hat\pi_{ZLB} = \frac{\kappa}{\tilde\lambda_z + r} x_{ZLB} \tag{71}$$

where I define

$$\tilde{\lambda}_z \coloneqq \mu + \lambda_z \tag{72}$$

The IS equation gives

$$x_{ZLB} = \frac{\sigma}{\tilde{\lambda}_z} \left(\underline{r}^n + \bar{\pi} + \hat{\pi}_{ZLB} \right) \tag{73}$$

Solving, we get

$$x_{ZLB} = \frac{\partial x_{ZLB}}{\partial \bar{\pi}} \left(\underline{r}^n + \bar{\pi} \right) \tag{74}$$

$$\frac{\partial x_{ZLB}}{\partial \bar{\pi}} = \frac{\sigma}{\tilde{\lambda}_z \left(1 - \sigma\kappa / \tilde{\lambda}_z^2 \right)} \tag{75}$$

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 17, 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.

8.5 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 (59)-(60) remains the same, replacing π_t by $\hat{\pi}_t$ in the Phillips curve (60).

Proposition 18. *Suppose that we morale reasons to have positive steady inflation, which is set optimally. Then, up to second order terms in inflation, welfare is:*

$$W_t = \bar{W} - h\left(\hat{\pi}_t\right) + g\left(x_t\right) \tag{76}$$

with g as in (63), and

$$h\left(\hat{\pi}_t\right) = \omega_{\hat{\pi}}^{-} \left|\hat{\pi}_t\right| + \omega_{\pi} \hat{\pi}_t = \omega_{\hat{\pi}}^{-} \left|\hat{\pi}_t\right| 1_{\hat{\pi}_t<0} + \omega_{\hat{\pi}}^{+} \left|\hat{\pi}_t\right| 1_{\hat{\pi}_t>0} \tag{77}$$

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 19. *Make the same assumption as Proposition 18. The first order condition for optimal monetary policy becomes*

$$\kappa h'\left(\hat{\pi}_t\right) = g'\left(x_t\right) - M^f g'\left(x_{t-1}\right) \tag{78}$$

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 $\operatorname{sign}\left(\pi_t\right)$ changes into $h'\left(\hat{\pi}_t\right)$. For instance, the condition (65) to have always $\hat{\pi}_t = 0$ becomes:

$$g'\left(x_t\right) - M^f g'\left(x_{t-1}\right) \in \left[\kappa h'\left(0^-\right), \kappa h'\left(0^+\right)\right]. \tag{79}$$

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.

9 Discussion

This section discusses some remaining high-level points.

Why is it hard to handle tasks requiring to juggle with numbers, and relatively "easy" to drive – does the model see that? We conclude with a remark on complexity across domains (linked to Moravec's Paradox: what's hard for computer is often easy for people and vice-versa). *Why is it hard to handle tasks requiring to juggle with numbers, and relatively "easy" to drive* (at least, after two dozen hours of lessons) – *does the model see that?* Intuitively, this is because processing actual numbers is hard, while processing visual information is quite easier. In the model, this is because the cost of processing c_i in (81) is high when dealing with mathematical operations on numbers, and and low when processing visual information: e.g., there's a fork on the road, and I should turn right. Hence, inside the model, when driving, the imprecision $1 - m_i$ is very small, and the complexity dread term (5) is very low. Under a heavy rain, so that seeing the road is harder, imprecision is higher, and driving is harder.

Are people aware that they're making a mistake? In this model, people are completely aware that they're making a mistake. This is clearly just a benchmark, probably right for more decisions. However, it is useful to keep in mind that there are some problems where a "tempting" answer seems "obvious", and is actually wrong. A famous one is the Monty Hall three door problems[48], and behavioral economics has provided other (Frederick (2005); Enke and Zimmermann (2019)). Those people are not "aware" of the mistake. This matters for economics: for instance, firms can exploit consumers' naïveté (Gabaix and Laibson (2006), Heidhues and Kőszegi (2018a)), those consumers are not aware that they're making a mistakes. It would be nice to extend the model for that, with an "attention" to the mistake.

[48] Amazingly, the great mathematician Erdős got ot wrong initially, and he was convinced only after several days, see Vazsonyi (1999)

A very 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.

10 Conclusion

This paper proposed a tractable theory of complexity aversion. It explains fairly easily a number of "obvious" features of the real world (that contracts and prices are "simple", e.g. not indexed)—-obvious to common sense, but puzzling to the traditional economic model with no complexity costs.

Along the way, the model shed light on a number of issues of basic modelling. One, cognition costs or information cost per se are not enough to generate complexity: one needs to have a form of "complexity aversion" that is non-zero even when the agent pays no cognition cost. Another, we saw that first order complexity aversion is the robust way to generate simplicity, much more than second order complexity aversion (or regret from regret theory)

Many interesting issues seem now within reach. In particular, designing optimal mechanisms, taking into account the complexity costs by consumers, seem particularly useful from a practical point of view. The present paper's analysis of simple tax systems for a taxation of goods is a beginning, and much more can be done.

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.

Aghion, P., U. Akcigit, M. Gravoueille, M. Lequien, and S. Stantcheva (2024). Tax simplicity or simplicity of evasion? evidence from self-employment taxes in france. Working paper, National Bureau of Economic Research.

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.

Andries, M. and V. Haddad (2020). Information aversion. *Journal of Political Economy 128*(5), 1901–1939.

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

Atkeson, A. and A. Burstein (2008). Pricing-to-market, trade costs, and international relative prices. *American Economic Review 98*(5), 1998–2031.

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.

Bajari, P. and S. Tadelis (2001). Incentives versus transaction costs: A theory of procurement contracts. *Rand journal of Economics*, 387–407.

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.

Bernheim, B. D., K. Kim, and D. Taubinsky (2024). Welfare and the act of choosing. Technical report, National Bureau of Economic Research.

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.

Calvet, L. E., J. Y. Campbell, and P. Sodini (2009). Fight or flight? Portfolio rebalancing by individual investors. *Quarterly Journal of Economics 124*(1), 301–348.

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.

Chetty, R., A. Looney, and K. Kroft (2009). Salience and taxation: Theory and evidence. *American economic review 99*(4), 1145–1177.

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

Colliard, J.-E. and C.-P. Georg (2025). Measuring regulatory complexity. Working paper, HEC Paris Research Paper Series.

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.

de Clippel, G., P. Moscariello, P. Ortoleva, and K. Rozen (2024). Caution in the face of complexity. Technical report, Mimeo.

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

DellaVigna, S. and M. Gentzkow (2019). Uniform pricing in us retail chains. *The Quarterly Journal of Economics 134*(4), 2011–2084.

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.

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.

Enke, B. and F. Zimmermann (2019). Correlation neglect in belief formation. *The Review of Economic Studies 86*(1), 313–332.

Eyster, E., K. Madarász, and P. Michaillat (2020, 08). Pricing Under Fairness Concerns. *Journal of the European Economic Association 19*(3), 1853–1898.

Farhi, E. and X. Gabaix (2020). Optimal Taxation with Behavioral Agents. *American Economic Review 110*(1), 298–336.

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

Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic perspectives 19*(4), 25–42.

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

Gabaix, X., R. S. Koijen, F. Mainardi, S. S. Oh, and M. Yogo (2025). Limited risk transfer between investors: A new benchmark for macro-finance models. Technical report, 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. and D. Laibson (2006). Shrouded attributes, consumer myopia, and information suppression in competitive markets. *Quarterly Journal of Economics 121*(2), 505–540.

Gabaix, X., D. Laibson, D. Li, H. Li, S. Resnick, and C. G. de Vries (2016). The impact of competition on prices with numerous firms. *Journal of Economic Theory 165*, 1–24.

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.

Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with non-unique prior. *Journal of mathematical economics 18*(2), 141–153.

Glaeser, E. L. and A. Shleifer (2001). A reason for quantity regulation. *American Economic Review, Papers and Proceedings 91*(2), 431–435.

Gonczarowski, Y. A., O. Heffetz, and C. Thomas (2023). Strategyproofness-exposing mechanism descriptions. Technical report, National Bureau of Economic Research.

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.

Hansen, L. P. and T. J. Sargent (2008). *Robustness*. Princeton university 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.

Heidhues, P. and B. Kőszegi (2018a). Behavioral industrial organization. *Handbook of Behavioral Economics: Applications and Foundations 1 1*, 517–612.

Holmström, B. (1979). Moral hazard and observability. *The Bell journal of economics*, 74–91.

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.

Kahneman, D. and A. Tversky (1979). Prospect theory: An analysis of decision under risk. *Econometrica 47*(2), 263–291.

Karlsson, N., G. Loewenstein, and D. Seppi (2009). The ostrich effect: Selective attention to information. *Journal of Risk and uncertainty 38*(2), 95–115.

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.

Kőszegi, B. and M. Rabin (2009). Reference-dependent consumption plans. *The American Economic Review 99*(3), 909–36.

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics 112*(2), 443–478.

Loomes, G. and R. Sugden (1982, 12). Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty. *The Economic Journal 92*(368), 805–824.

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. (2024). Simplicity and risk. *Journal of Finance, Forthcoming*.

Rabin, M. (2000). Risk aversion and expected-utility theory: A calibration theorem. *Econometrica 68*(5), 1281–1292.

Ramsey, F. P. (1927). A Contribution to the Theory of Taxation. *Economic Journal 37*(145), 47–61.

Rees-Jones, A. and D. Taubinsky (2020). Measuring 'schmeduling'. *The Review of Economic Studies 87*(5), 2399–2438.

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

Sarver, T. (2008). Anticipating regret: Why fewer options may be better. *Econometrica 76*(2), 263–305.

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.

Segal, I. (1999). Complexity and renegotiation: A foundation for incomplete contracts. *The Review of Economic Studies 66*(1), 57–82.

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.

Sicherman, N., G. Loewenstein, D. J. Seppi, and S. P. Utkus (2016). Financial attention. *The Review of Financial Studies 29*(4), 863–897.

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).

Taubinsky, D. and A. Rees-Jones (2017). Attention variation and welfare: Theory and evidence from a tax salience experiment. *Review of Economic Studies 85*(4), 2462–2496.

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.

Vazsonyi, A. (1999). Which door has the cadillac? *Decision Line 30*(1), 17–19.

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 How to recover frames, attention and default from data

In this research, the "attention", the "default", and the "frame" are taken as knowable. Here, I discuss how to elicit them from data.[49]. In practice, there is often on "obvious" frame (e.g. prices of apples and bananas are processed separately), though sometimes the choice is a bit delicate (e.g. do people use nominal or inflation-adjusted prices).[50]

First, let us define terms. We consider problems $\max_a u(a, x)$, where $a \in \mathbb{R}^A$ and $x \in \mathbb{R}^N$. Linearizing, and normalizing the default a to be 0, the rational action is:

$$a^r = a_x^r x$$

where a_x^r is an $A \times N$ matrix. With inattention, we get for the average action, we get an attention matrix M, which we take to be diagonal, and the default vector x^d:

$$a^s = a_x^r x^s, \qquad x^s = Mx + (I - M) x^d$$

But one could imagine the people use a transformation $\tilde{x} = Fx$, where $F = \frac{\partial \tilde{x}}{\partial x} = \tilde{x}_x$ is an $N \times N$ matrix (we keep the transformation to be linear; or F could be the derivative for small x of a non-linear transformation). For instance, the frame could be nominal, or real, assuming the inflation is somewhere in x. We can call the choice of F (or equivalently, the representation $x \mapsto \tilde{x}$) the "frame". Under that frame

$$\tilde{x}^s = M\tilde{x} + (1 - M) \tilde{x}^d = MFx + (I - M) Fx^d$$

and the action is $a^s = a_{\tilde{x}}^r \tilde{x}^s = a_x^r x^s$ with $a_{\tilde{x}}^r = a_x^r F^{-1}$ and $x^s = F^{-1}\tilde{x}^s$, so

$$a^s = a_x^r x^s, \qquad x^s = F^{-1}MFx + \left(I - F^{-1}MF\right) x^d \tag{80}$$

Suppose that we have data of x, $a^s(x)$, where $a^s(x)$is the average choice of the behavioral agent (with s meaning subjective, or sparse), and perhaps $x^s(x)$ — where we elicit the perceived value of a variable (say, after the choice). Can we recover attention M and the frame F?Throughout, we suppose that we know a_x^r — perhaps, this is elicited under clear and distinct conditions, with the help of a guide.

Proposition 20. (Recover of attention, and the frame) *Suppose that M's eigenvalues are all different, and that we know a_x^r. Suppose that we know $D_x = G_a a_x^s + G_x x_x^s \in \mathbb{R}^{d_D \times N}$,*

[49] I already discussed how to elicit attention in Gabaix (2014) and Gabaix (2019).

[50] One could also endogenize them, e.g. by using the "most useful" frame for instance.

where selection matrices G_a and G_x have $d_D \times A$ and $d_D \times A$), and that D_x has rank least N. For instance we know x_x^s or we know a_x^s and $A \geq N$. Then, up to reordering and rescaling, we can recover the attention M and the frame F. That is, we can recover some M_0 and F_0 that account for the data; and all the other solutions have the form $M = \Pi M_0 \Pi^{-1}$ and $F = \Pi S F_0$, where Π is a permutation matrix, and S is a diagonal invertible "scaling" matrix. If we know the frame F, it is enough that $D_x \neq 0$: then we can recover attention M.

Proof. Let us start with a special case, where we have x_x^s from data. In the model, this is $x_x^s = F^{-1}MF$. Hence, given M is diagonal, x_x^s must be diagonalizable, and M is (up to permutation) the matrix of its eigenvalues. The rest is well-known diagonalization algebra: any solution for the diagonal has the form $M = \Pi M_0 \Pi^{-1}$, where Π is a permutation matrix (there are $N!$ of them). Also, a solution can be rescaled, as in $F = \Pi S F_0$ where S is a diagonal invertible "scaling" matrix. Indeed $F^{-1}MF = F_0^{-1}M_0F_0 = x_x^s$, and this describes fully the set of solutions.

In the more general case, we have the data $D_x = G_a a_x^s + G_x x_x^s = (G_a a_x^r + G_x) F^{-1}MF$, for some (selection) matrices G_a and G_x. The proposition assumes assume that $G_a a_x^r + G_x$ has rank at least N. Hence, we can recover $x_x^s = F^{-1}MF$ from data, and we apply the special case we start with.

If we know the frame F (for instance, it is the identity matrix, a sort of natural frame). Then, it is enough that D_x has rank 1. This gives $F^{-1}MF$, hence M. □

Economically, the "up to reordering and rescaling" makes sense: we do not care about the order of the perception vector (i.e., does the agent order the price vector as $(p_{\text{apple}}, p_{\text{banana}}, p_{\text{cherries}})$ or with another order), or scale (i.e. does the agent count in grams or kilograms). The hard data requirement is the we can recover the frame we enough data on actions, or on perceptions: we need N^2 pieces of independent data.[51]

It is comparatively easier to recover the default x^d, as we record here.

Proposition 21. (Recovery of the default) *Suppose the we observe $D^d = G_a a^d + G_x x^d$, some transform of the default action and default parameter, where $G_a a_x^r + G_x$ has rank N at least. Then, we can recover the default x^d.*

Proof. We observe $D^d = (G_a a_x^r + G_x) x^d$. Then, as $G_a a_x^r + G_x$ has rank N at least, we recover x^d. □

In conclusion, we can in principle recover all the essential elements: the attention, the default, and the frame.

[51] The special case where we only have a_x^s is also interesting. We can recover $a_x^s = a_x^r x_x^s = a_x^r F^{-1}MF$. This is an $A \times N$ matrix. However, the unknowns live in a manifold of dimension N^2: N for M; plus $N^2 - N$ for F, up to the diagonal scaling matrix S. Hence, we can recover attention iff $A \times N \geq N^2$, i.e. if $A \geq N$.

A.2 Some simple cognitive microfoundations

Derivation of (2) One derivation of (2), and the functional form for the noise, is the following. People receive noisy "raw" signals $\tilde{x}_i^s$ about x_i: $\tilde{y}_i^s = x_i + \tilde{\eta}_i$, where η_i is a mean-zero noise uncorrelated with x_i. Then, if noises are Gaussian, we have $\mathbb{E}\left[x_i|\tilde{y}_i^s\right] = x_i^s \coloneqq m_i\tilde{x}_i^s$, with $m_i = \frac{\sigma_{x_i}^2}{\sigma_{x_i}^2+\sigma_{\tilde{\eta}_i}^2}$. If noises are not Gaussian, then $x_i^s = m_i\tilde{x}_i^s$ is still the least-squared best estimator of x_i, $m_i = \text{argmin}_{m_i}\mathbb{E}\left[\left(x_i - m_i\tilde{x}_i^s\right)^2\right]$, which yields $m_i = \frac{\mathbb{E}\left[x_i\tilde{x}_i^s\right]}{\mathbb{E}\left[\left(\tilde{x}_i^s\right)^2\right]} = \frac{\sigma_{x_i}^2}{\sigma_{x_i}^2+\sigma_{\tilde{\eta}_i}^2}$. We then define the "processed" signal as $x_i^s \coloneqq m_i\tilde{x}_i^s$, and set $\eta_i = m_i\tilde{\eta}_i$, which is uncorrelated with x_i. This yields (2), with:

$$\sigma_{\eta_i}^2 = m_i^2\sigma_{\tilde{\eta}_i}^2 = m_i^2\sigma_{x_i}^2\left(\frac{1}{m_i} - 1\right) = \sigma_{x_i}^2 m_i\left(1 - m_i\right)$$

Production function of thought

Typically, in calculations, we will assume a production function of thought (or precision) as follows: an effort L_i leads to a precision

$$m_i\left(L_i\right) = \min\left(\left(\frac{L_i}{c_i}\right)^{1-\alpha}, 1\right), \qquad \alpha \in (0,1) \tag{81}$$

where $\alpha \in (0,1)$ and c_i parametrizes the basic "cost" of thinking, and is taken here as a primitive (it can be endogenized, as in Gabaix and Graeber (2024)). The cognitive cost $C_i\left(m_i\right)$ is then wL_i, where L_i is the effort needed to reach precision m_i according to (81), and w is a cost of effort, which can be taken as a constant in applications, so is

$$C_i\left(m_i\right) = wc_i m_i^{\frac{1}{1-\alpha}}, \qquad 0 < \alpha < 1 \tag{82}$$

We could also have the production function of thought with $\alpha > 1$:

$$m_i\left(L_i\right) = \max\left(1 - \left(\frac{L_i}{c_i} + 1\right)^{1-\alpha}, 0\right), \qquad \alpha > 1 \tag{83}$$

This model induces sparsity, i.e. when benefits are small enough, the optimal effort and attention are exactly 0. Then, the cost is

$$C_i\left(m_i\right) = wc_i\left(\left(1 - m_i\right)^{\frac{1}{1-\alpha}} - 1\right), \qquad \alpha > 1 \tag{84}$$

Yet another tractable format is a linear-quadratic cost of cognition. The cost $C\left(m\right)$ is convex in m, $C\left(0\right) = 0$, and we normalize $w = 1$.

A.3 Closed form for endogenous attention in the indexation problem

Solving for the optimum m is in general a bit ugly. One closed form obtains is

$$C^{\text{cognition}}(m) = \kappa \left(m - m^d\right)^+ \tag{85}$$

for some "default" attention $m^d \in [0, 1)$. This captures that there's a "free" minimum of attention m^d, but that paying higher attention is effortful.

Proposition 22. (Optimal attention and indexation with linear cognitive cost function). *Suppose the cognition cost (85). Then, the full utility function of the consumer (26) is maximized at either $m = m^d$ or $m = 1$. We have $m = 1$ iff $\mathcal{V}(1) \geq \mathcal{V}\left(m^d\right)$, i.e. iff:*

$$b(1)^2 - b\left(m^d\right)^2 \left(2 - m^d\right) \geq \kappa \left(1 - m^d\right) \tag{86}$$

and otherwise $m = m^d$ and $b = b\left(m^d\right)$.

In particular, when $b\left(m^d\right) = 0$, the optimum indexation level is either no indexation ($b = 0$), or full indexation ($b = b(1) = b^r$).

Proof. Using (25) and the envelope theorem), with $\mathcal{V}(m) \coloneqq V(b(m), m)$

$$\mathcal{V}'(m) = \frac{d}{dm} V(b(m), m) = \frac{\partial V(b, m)}{\partial m} = b^2 + 2\xi\sigma_a |b| - \kappa$$

Observe that $b(m)$ is increasing weakly. So, if for some $m_0 \in \left[m^d, 1\right]$, $\mathcal{V}'(m_0) \geq 0$, then for all $m \geq m_0$, we will have $\mathcal{V}'(m) \geq \mathcal{V}'(m_0) \geq 0$ hence $\mathcal{V}(1) \geq \mathcal{V}(m)$. Likewise, if $\mathcal{V}'(m_0) \leq 0$, $\mathcal{V}\left(m^d\right) \geq \mathcal{V}(m)$. Hence, the two candidate solutions are $m = m^d$ or $m = 1$. Hence, we have $m = 1$ iff $\mathcal{V}(1) \geq \mathcal{V}\left(m^d\right)$. This is true if the cognition costs κ are not too large. In the case $m^d = 0$, we have $b = 1$ iff $\mathcal{V}'(1) \geq \mathcal{V}'(0)$, i.e. iff $(b^r)^2 \geq \kappa$. □

This "extremal property" (that the optimal attention is either m^d or 1) works also for "nice" cost functions in a range $[m_d, 1]$. Indeed, observe that (calling $C^{\text{cognition}}(m)$ simply $C(m)$):

$$\mathcal{V}'(m) = \bar{c}(m) - C'(m), \qquad \bar{c}(m) \coloneqq b^2 + 2\xi\sigma_a |b| \tag{87}$$

So, if $C'(m) \leq \bar{c}(m)$ for all m, then $\mathcal{V}'(m) \geq 0$ for all m, and the optimum is $m = 1$.If $C'(m) \geq \bar{c}(m)$ for all m, then $\mathcal{V}'(m) \leq 0$ for all m, and the optimum is $m = m^d$.

A.4 A microfoundation for the fact that competitive firms want maximize social surplus

Section 4 assumes that the firm is "benevolent" – wants maximize social surplus. Here is a specific microfoundation. It is quite traditional: as in Adam Smith and Arrow-Debreu, competitive firms, by their actions, maximize social surplus (at least, when consumers appreciate enough their expected utility). Still, for completeness, it is useful to fill in the details.

There are at least 2 firms, and N consumers. On day t, the marginal costs is $1+x_t$, the same across firms (this is not essential). At time 0 (at the beginning of the year), each firm f proposes a pricing scheme $p_f(x_t) = \bar{p}_f + b_f x_t$ (where is chooses the average price $\bar{p}_f$ and the slope b_f) and also "membership fee" that the consumer has to pay, a lump-sum payment F_f (this is not essential either, but makes the analysis more general, and simpler).[52]

At the beginning of the year, each consumer sees the menus offered by the firms, and chooses which firm to sign up with—which firms to go to every day $t = 1, \ldots, T$ in that year. For instance, the consumer chooses which bakery to go to every day, in a way that maximizes her utility over the year. Formally, the firm gets an exclusive contract within the year (but has committed to the pricing scheme above).[53]

The consumer is fully conscious her expected utility, including and cognition and complexity aversion costs.[54] She chooses a firm f that maximizes his expected utility surplus, from this pricing scheme, i.e.

$$V_f = -F_f + TV_f^{\text{cons}} \tag{88}$$

The first part if the lump-sum payment. The second part is the expected utility of the consumer, net of the cost. The consumption is chosen as in the main body of the paper, as $1 + a_f(x^s) = \text{argmax}_{1+a} U(1+a) - p_f(x^s)(1+a)$, based on the perceived price $p_f(x^s)$.

The firm's profit is

$$\Pi_f = N_f \left(F_f + T\mathbb{E}\left[\left(p_f(x) - c(x)\right)\left(1 + a_f(x^s)\right)\right]\right)$$

where N_f is the number of consumers choosing the firms, and the last term is the average profit per consumer.

The timing is: each firm f propose simultaneously $(F_f, \bar{p}_f, b_f)$, to maximize its expected

[52] I assume that the firm can commit to the scheme. This is obviously easy to do when the outcome is a constant price $\bar{p}_f$.

[53] This is important, to prevent arbitrage: otherwise, take the case the outcome is a constant price equal to 1. Then, on a day with a negative x_t, a deviating firm could undercut its competitors (which offer a price of 1), by offering a price of $1 + \frac{1}{2}x_t < 1$, explain to consumers that they save money by coming to this deviating firm, and make a profit.

[54] This is where the paper deviates from the literature on shrouded attributes and firms exploiting naive consumers (see footnote 21): in that benchmark, the consumer is not conscious of all the future elements of the situations, for instance does not take into account "surprise" fees.

profit Π_f; then the each consumer decides which firms to choose. Firms play Nash.

Lemma 3. *In the competition game above, the only Nash equilibrium is: $\bar{p}_f = 1$, $F_f = 0$, and b_f as in the paper, maximizing total welfare. Firms' profits are zero.*

Proof. The reasoning is completely conventional. Because we allow for fixed payments F_f, firms have zero profits, and offer the maximum consumer's surplus given zero profits – hence, they maximize total social surplus. Given that, the average price right should be the average marginal cost, so $\bar{p}_f = 1$. Then, the derivation of the optimal b_f is the one that maximizes social surplus (which is the consumer's utility, as equilibrium profits are zero) as in the main body of this paper. □

A.5 Second order complexity aversion sometimes generates simplicity, but in a non-robust way

This appendix probes second-order complexity aversion, and finds that it generally fails to generate full simplicity, e.g. rigid prices or zero indexation. We explore this in the pricing problem of Section 4

Proposition 23. (Optimal price simplicity, with second order complexity aversion) *Suppose second order complexity aversion. Suppose that attention m is exogenous. Then the optimum slope is*

$$b = \frac{m}{m + (1+\xi)(1-m)} \in [0,1]$$

Hence, contracts are typically non-simple, as $b > 0$ whenever $m > 0$.

Proof. With second order CA, we use the analogue of (119)-(120), with $\bar{\xi} = 1 + \xi$

$$W = \left(b - \frac{1}{2}b^2\right)\psi\sigma_x^2 m - \bar{\xi}\frac{1}{2}\psi\sigma_x^2 b^2 (1-m) - C(m)$$

so

$$W_b = (1-b)\psi\sigma_x^2 m - \bar{\xi}\psi\sigma_x^2 b(1-m) + W_m\frac{dm}{db}$$

First, with exogenous attention, so $\frac{dm}{db} = 0$. Then, the optimum is $b = \frac{m}{m+\bar{\xi}(1-m)}$. In general, $b = \frac{m+\frac{1}{\psi\sigma_x^2}W_m\frac{dm}{db}}{m+\bar{\xi}(1-m)}$. □

So, if there is a "default" level of attention that's positive, the some positive slope $b > 0$ is optimum: contracts are not "perfectly simple". They are not the optimum contract with rational agents, which would be $b = 1$.

Let us next explore a different variant, to probe the robustness of the second order complexity aversion. I suppose that there is a non-zero fraction $\varepsilon \in (0,1)$ of rational consumers, with zero cognition costs, and that the same contract must be offered to all consumers.

Proposition 24. (Optimal contract when there is a small fraction of rational consumers) *Suppose that there is a fraction $\varepsilon \in (0,1)$ of rational consumers. Then, with first order complexity aversion: with either exogenous attention $m < 1$, or endogenous attention with the conditions of Proposition 8, then if ξ is large enough one still wants a simple contract ($b = 0$). However, with second order complexity aversion, a simple contract is never optimal.*

The moral is that first order complexity aversion is a much more robust way to get simplicity than second order CA.

Now, let us turn to endogenous attention, which gives another way to appreciate the same phenomenon.

Proposition 25. (Optimal price simplicity, with second order complexity aversion, endogenous attention) *Suppose second order complexity aversion, and endogenous attention. Then, we have the following: with $\alpha < \frac{2}{3}$ or $\alpha > 1$ (where α parametrizes the curvature of the production function of thought, (81)-(84)), then if complexity aversion $\xi > 0$ is large enough, then the optimum entails rigid prices ($b = 0$). However, if $\alpha \in \left(\frac{2}{3}, 1\right)$, then the optimum entails flexible prices ($b > 0$).*

We see that both first and second order complexity aversion generate simplicity. However, the first order complexity works with fewer condition. Both tools achieve complexity, but first order CA is very directly effective, like sword, whereas second order CA is more delicate, like a scalpel. It is useful to have both tools in the economist's kit, though the sword will be more powerful most of the time.

A.6 Complexity aversion as a partial force for consumer protection from firms' complexity

How does that modify firm's ability to manipulate consumer. Complexity (adding features) has new effect (i) lowering utility by making the product more "suspicious"; in addition to two effects already understood: (ii) a source of profit, the pitfall for naive consumers (ii) a source of inefficiency: some effort required for rational consumers (iii). To see how they balance, consider Gabaix and Laibson (2006).

There a rational and naive consumer; the fraction of naives is α. Each firm sells a "base good"(say, a printer), whose price p everyone sees. It also sells an add-on good (say, a color cartridge) $\hat{p}$, that only rational consumers see — it's opaque or "shrouded" to the naive

consumers. Rational consumer can avoid paying it at an effort cost e (e.g. by printer black and white). The marginal cost of both goods is 0. For legal constraints, maximum price on the shrouded attribute is $\bar{p}$ (firms cannot charge a $1 million for hidden fee).

Rational consumers have surplus:

$$V^{\text{rat}} = -p - \min(\hat{p}, e) \tag{89}$$

i.e. they pay the price p, and pay the effort cost e it's less than the add-on good price $\hat{p}$.

Naive consumers have perceived surplus (I take the case of full naivete, $m = 0$):[55]

$$V^{\text{naive},s} = -p - \xi M |\hat{p}| \tag{90}$$

where M is the "awareness" of the shrouded attribute. This means that they have some sense of the shrouded attribute, in a way modulated by ξM. This term $-\xi M |\hat{p}|$ is the new one, compared to GL.

There is a potential entry of firms offering marginal cost pricing, $p = \hat{p} = 0$, the "simple" scheme without shrouded prices. The question is: will consumers go to a "shrouded attribute" equilibrium with high prices for the shrouded attribute?

Proposition 26. (Complexity aversion makes a shrouded attribute equilibrium less likely) *We have a shrouded price equilibrium if the fraction of naive agent α is large enough, and they're not too complexity-aware:*

$$\alpha \geq \alpha^* = \max\left(\xi M, \frac{e}{\bar{p}}\right) \tag{91}$$

i.e. (i) $\alpha \geq \xi M$,i.e. if complexity aversion is not too strong; and (ii) the older condition in Gabaix and Laibson (2006) $\alpha \geq \frac{e}{\bar{p}}$, i.e. there

The GL analysis assumed $\xi M = 0$, no awareness-adjusted complexity aversion, so the condition was simply $\alpha \geq \frac{e}{\bar{p}}$. Hence, complexity aversion helps make it less likely that the market less shrouded, though this force is moderate. The reason is that the naives "feel bad" about the product (via complexity aversion), as they feel that there is a snag, and defect to the simple "scheme", without any shrouded attribute. For small complexity aversion ξ this force will make no difference in allocations.

Proof of Proposition 26. The proof is a variant of the one in GL (Gabaix and Laibson (2006)), which gives all the technical details, including the careful analysis of all the possible deviations. Instead, let us concentrate on the economic essence of the argument. Suppose

[55] The more general case is $V^{\text{naive},s} = -p - \min(m\hat{p}, e) - \xi(1-m)|\hat{p}|M$, and is similar.

that we have a shrouded equilibrium, a segment of firms serves both, with a fraction of naive consumers is α. In that equilibrium, it's optimal for the firms to offer the maximum allowable price, $\hat{p} = \bar{p}$. Then, the profit is $\Pi = p + \alpha\bar{p}$, and it should be 0 with perfect competition (this can be relaxed, as in GL), so $p = -\alpha\bar{p}$. With complexity aversion, perceived surplus of naives thus:

$$V^{\text{naive},s} = -p - \xi M \left|\hat{p}\right| = \alpha\bar{p} - \xi M\bar{p}$$

The surplus offered by a firm offering goods at marginal cost is 0. Hence, for a shrouded attributes equilibrium to hold, we need $V^{\text{naive}} \geq 0$, i.e. $\alpha \geq \xi M$.

For rational consumers, the analysis is as in GL: the surplus of rationals is $V^{\text{rat}} = \alpha\bar{p} - e$. For the shrouded equilibrium to hold, we need $V^{\text{rat}} \geq 0$, i.e. $\alpha \geq \frac{e}{\bar{p}}$.

A.7 More on simple taxes

I complement Section 5. I ask when we have a fully uniform tax system when there is no cognition costs, but only attention m_g to the goods.

Condition (93) means that we have a simple tax system if complexity aversion ξ is large enough compared to the relative dispersion of elasticities ($\frac{\max_g \psi_g - \min_g \psi_g}{\bar{\psi}}$), and if attention m is low enough. Condition (92) has the same message, in a more general way.[56]

Proposition 27. (Optimal tax simplicity, with a concave cost of cognition) *Suppose first order complexity aversion, and exogenous attention. Define the average elasticity* $\bar{\psi} = \frac{1}{G}\sum_g \psi_g$. *Suppose for all goods, the average attention to the tax is imperfect,* $m_g < 1$. *Then if complexity aversion* $\xi > 0$ *is large enough, then the optimum entails a simple tax system: the tax rate is the same for all goods,* $\tau_g = \bar{\tau} = \frac{\lambda}{\bar{\psi}}$. *More specifically, a necessary and sufficient for this simple tax system is that* ξ *is large enough to satisfy is:*

$$\max_g \lambda m_g \frac{\psi_g}{\bar{\psi}} - \xi\sigma_g\left(1 - m_g\right) \leq \min_g \lambda m_g \frac{\psi_g}{\bar{\psi}} + \xi\sigma_g\left(1 - m_g\right) \tag{92}$$

When m_g *and* σ_g *are the same across goods, this necessary and sufficient condition becomes:*

$$\lambda \frac{\max_g \psi_g - \min_g \psi_g}{\bar{\psi}} \leq 2\xi\sigma \frac{1 - m}{m} \tag{93}$$

Proof of Proposition 27 First, we observe that if the optimum all the $\check{\tau}_g = 0$,then the optimum $\bar{\tau} = \frac{\lambda}{\bar{\psi}}$.

Next, we examine the optimum $\check{\tau}_g$. We maximize W, subject to the constraint that the

[56]There is a value $\xi_* \geq 0$ such that (92) holds if and only if $\xi \geq \xi_*$. Indeed, observe that the difference between the right hand size of (92) and its left-hand size increases in ξ, and goes to ξ as $\xi \to \infty$.

average deviation $\check{\tau}_g$ is 0, so the Lagrangian is $L = W + \mu \sum_g \check{\tau}_g$. The FOC is for $\check{\tau}_g$ is:

$$0 = L_{\check{\tau}_g} = -\psi_g m_g \left(\bar{\tau} + \check{\tau}_g\right) - \xi \sigma_{c_g} \left(1 - m_g\right) \text{sign}\left(\check{\tau}_g\right) + \lambda + \mu$$

So we have $\check{\tau}_g = 0$ iff:

$$v_g^- \coloneqq \psi_g m_g \bar{\tau} - \xi \sigma_{c_g} \left(1 - m_g\right) \leq \lambda + \mu \leq \psi_g m_g \bar{\tau} + \xi \sigma_{c_g} \left(1 - m_g\right) =: v_g^+$$

So all $\check{\tau}_g = 0$ iff we can find a μ for which this inequality holds for all g, i.e. iff:

$$\max_g v_g^- \leq \min_g v_g^+$$

This is the expression in (92), as $\bar{\tau} = \frac{\lambda}{\psi}$.

A.8 The complexity cost of business cycles

Here I complement the analysis of welfare in the main text, in particular Section 7.

A.8.1 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 $\ell^{\text{CA}} = \omega_y \sigma_y$ 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

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

Hence, the loss solely due to the business cycle is

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

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

$$g(y) \sim a_0 \sigma_y \tag{94}$$

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

$$g(y) \sim \mathbb{E}\left[\frac{1}{\sigma_{\varepsilon_i}}\right] \frac{\sigma_y^2}{2} \tag{95}$$

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:

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

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:

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

for small σ_y, with a positive constant c_α.[57] 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.8.2 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 $rT_p \ll 1$),

$$W_t = W_t^{\text{trad}} + W_t^{\text{CA}} + W_t^{\text{Tobin}} \tag{96}$$

$$W_t^{\text{trad}} = -\frac{\varepsilon}{2} T_p^2 \pi_t^2 - \frac{1}{2}(\gamma + \sigma)(x_t - x^{**})^2 \tag{97}$$

$$W_t^{\text{CA}} = -\xi^\pi T_w |\pi_t| - \xi^x |x_t| \tag{98}$$

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

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

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

$$x^{**} = \frac{\tau^{\text{eff}}}{\gamma + \phi} \tag{99}$$

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.

So, the Taylor expansion is:

$$W_t^{\text{trad}} = \tau^{\text{eff}} x_t - \xi^x |x_t| - \frac{1}{2}(\gamma + \phi) x_t^2 + A_t \tag{100}$$

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.9 Fixed price vs Cost plus contract: why the slope in a contract is often exactly 0 or 1

To build a house, most contracts are fixed cost: the buyer pays a pre-determined price $\bar{p}$ to the seller; but otherwise, they tend to be "cost plus" (C+).[58] So the slope of the payment is either $b = 0$ or $b = 1$, rather than some intermediary solution. How to explain that? A key model (Bajari and Tadelis (2001)) models whether people will choose Fixed or Cost+, but they need to posit that those are the only to possibilities; the acknowledge that, logically,

[58]The "plus" refers to e.g. a fixed cost, or some extra incentive e.g. on-time performance, or an extra 10% "administrative" cost, which I fold in the total cost

intermediary values of $b \in (0,1)$ would be optimal, and they exclude them by fiat (adding some intuitive reasons why that might be possible).

Here I show how complexity aversion can be a natural force that can be added to their analysis, and make natural the fact that (often) people will use only $b = 0$ or $b = 1$. To show that, I take the "chemically pure" case where only complexity-sharing and risk-sharing are the only forces, without incentives or incomplete contracts.

Suppose that the cost is

$$c = \bar{c} + x$$

where x is random with mean 0. The buyer pays

$$p = \bar{p} + bx$$

to the seller, for some $\bar{p}$ and the slope b to be determined. The utility of the buyer and sellers are:

$$v^B = u^B(-p) \qquad v^S = u^S(c - p)$$

Assume linear-quadratic utility, absolute risk aversion Γ_B, Γ_S. The fully utilities are, removing the constants:

$$V_B = -\bar{p} - \frac{1}{2}\Gamma_B b^2 \sigma_x^2 - \xi_B |b| \sigma_x, \qquad V_S = \bar{p} - \bar{c} - \frac{1}{2}\Gamma_S (1-b)^2 \sigma_x^2 - \xi_S |1-b| \sigma_x$$

where $\xi_B = \xi_B \kappa (1-m)$ is the effective complexity aversion of the buyer, where κ is the unit converter. So, the optimal contract will maximize (assuming Nash bargaining would yield something very similar).

$$\max_b V, \qquad V = V_B + V_S$$

First, without complexity aversion (or first order risk aversion), we get an interior solution, $b = \frac{\Gamma_S}{\Gamma_B + \Gamma_S}$, so it is typically strictly interior — contradicting the empirical facts — except for a risk-neutral firm, which does lead to $b = 0$.

Next, with complexity aversion, we have the following result:[59]

Proposition 28. (Optimality of Fixed vs Cost+ contracts) *We have $b = 0$ iff $\xi_B - \xi_S \geq \Gamma_S \sigma_x$. We have $b = 1$ iff $\xi_B - \xi_S \leq -\Gamma_B \sigma_x$. In between, if $\xi_B - \xi_S \in (-\Gamma_B \sigma_x, \Gamma_S \sigma_x)$ we have $b = \frac{\Gamma_S - \frac{\xi_B - \xi_S}{\sigma_x}}{\Gamma_B + \Gamma_S} \in (0,1)$.*

When risk aversions become negligible ($\Gamma_B, \Gamma_S \to 0$), only $b = 0$ and $b = 1$ are possible. More generally, in an open zone of parameters, we have $b = 0$ or $b = 1$. In most cases with a

[59] The proof involves simply taking the derivative. E.g., taking the case $\sigma_x = 1$, $V_{b|b=0} = -\xi_B \operatorname{sign}(b) + \Gamma_S + \xi_S$, so $b = 0$ is optimal iff $\xi_B \geq \Gamma_S + \xi_S$.

"small" buyer and a large firm, $\xi_B \gg \xi_S$, so $b = 0$ and we have a fixed contract. In contrast, when the "government" is large, perhaps $\xi_B \ll \xi_S$, and $b = 1$.

So, we find that complexity aversion, or its elder cousin first order risk aversion, explain easily why contracts will have a slope of 0 or 1. This means that complexity aversion is an ingredient that can be added to previous analysis (e.g. Bajari and Tadelis (2001)), which helps guide the analysis towards simpler schemes.

We have "complexity sharing" – rather than "risk sharing". But very often, in "complexity sharing", just one party bears all the complexity.

A.10 Complexity aversion in discrete choice

Most of this paper is for a continuous action (perhaps multidimensional). Here I show how to handle discrete choice.

A.10.1 Prior results on discrete choice, without complexity aversion

I draw from Gabaix and Graeber (2024), which I first summarize here. We consider the discrete choice

$$\max_{k:k\in K\subseteq\mathbb{N}} v(k,x) \tag{101}$$

where the choice k is in a discrete set K, and yields rational utility $v(k,x)$, where for each k, $v(k,x)$ is twice continuously differentiable in x.

We define as $a(k)$ the probability of playing action k, so $\sum_a a(k) = 1$, $a \in \Delta(K)$. We then define the extended utility function, concave in a, as

$$u(a,x) \coloneqq \sum_k a(k)\, v(k,x) - \beta D\left(a, a^d\right), \tag{102}$$

where $D\left(a, a^d\right)$ is a divergence measure that is convex in a, $a^d \in \Delta(K)$ is a default action probability, and $\beta > 0$ is discussed later. We shall take the Kullback-Leibler divergence between a and a^d:

$$D\left(a, a^d\right) = 1 + \sum_k a(k)\left(\ln \frac{a(k)}{a^d(k)} - 1\right) \tag{103}$$

The main analytical results are the following. The solution of $\max_a u(a,x)$ is $a = a^F\left(a^d, v(x), \beta\right)$, where the probability of playing action k is:

$$a^{F,k}\left(a^d, v(x), \beta\right) = \frac{a^{k,d} e^{\frac{v^k}{\beta}}}{\sum_{k'} a^{k',d} e^{\frac{v^{k'}}{\beta}}} \tag{104}$$

A change in information dx_i creates a change in the probability of choosing action k equal

to

$$da^k = \frac{1}{\beta} a^k \sum_i \left(v_{x_i}^k - \bar{v}_{x_i} \right) dx_i \tag{105}$$

where $\bar{v}_{x_i}$ is the average value of valuation changes created by the change in variable x_i:

$$\bar{v}_{x_i} = \sum_k a^{k,d} v_{x_i}^k. \tag{106}$$

In addition, $u_{aa} = -\beta Diag\left(\frac{1}{a^k}\right)$.

Hence, the expected utility after acquiring precision m_i is, up to higher order terms in $\sigma_{x_i}^2$:

$$U(m) = U\left(0^+\right) + \sum_i V_i m_i, \qquad U\left(0^+\right) = \sum_k a^{d,k} \mu^k \tag{107}$$

where $V_i = \frac{V_i^0}{2\beta}$ is the value of information in dimension i, with:

$$V_i^0(a) = \sum_k a^k \left(v_{x_i}^k - \bar{v}_{x_i} \right)^2 \sigma_{x_i}^2 \tag{108}$$

If information in dimension i only affects an option $k(i)$ (i.e., if $v_{x_i}^{k'} = 0$ for $k' \neq k$), V_i^0 reduces to:

$$V_i^0(a) = a^{k(i)} \left(1 - a^{k(i)}\right) \left(v_{x_i}^{k(i)}\right)^2 \sigma_{x_i}^2 \tag{109}$$

Finally, We set a^d and β in the following manner. First, we define $a^{d,0}(k) \coloneqq \frac{1}{K}$ the "ignorance prior" on the actions, assigning equal probability to all actions; and compute an initial value of β, $\beta^0 \coloneqq \beta^F\left(a^{d,0}\right)$,

$$\beta^F\left(a^d\right) \coloneqq \bar{\beta}(\sum_i V_i^0\left(a^d\right))^{1/2}, \tag{110}$$

where $\bar{\beta}$ is a unitless parameter and V_i^0 is in (191). using (110). Then, we obtain the default action probability, $a^d = a^F\left(a^{d,0}, v\left(x^d\right), \beta^0\right)$, where a^F is the function in (104), and its associated $\beta = \beta^F\left(a^d\right)$.

A.10.2 Complexity aversion in discrete choice

The agent adds a complexity penalty to each option:

$$\tilde{v}(k, x) \coloneqq v(k, x) - C^{CA,k} \tag{111}$$

$$C^{CA,k} = \xi \sum_i \left| v_{x_i}^k \right| \sigma_{x_i} (1 - m_i) \tag{112}$$

Total complexity aversion is

$$C^{CA} = \sum_k a^{k,d} C^{CA,k} = \xi \sum_i \left(\sum_k a^{k,d} \left| v_{x_i}^k \right| \right) \sigma_{x_i} (1 - m_i) = \xi \sum_i D_{x_i}^u \sigma_{x_i} (1 - m_i) \quad (113)$$

$$D_{x_i}^u \equiv \sum_k a^{k,d} \left| v_{x_i}^k \right| \quad (114)$$

This $D_{x_i}^u$ is the discrete-action counter part of what we had in the basic continuous action case: $D_{x_i}^u = \left| \sigma_a u_{aa} \frac{\partial a}{\partial x_i} \right|$. They capture ow much x_i changes the action, times the utility consequences of that.

Otherwise, the model is the same. The optimization of attention does take into account the CA term. So, the full utility is:

$$U(m) = U^r - \sum_i V_i (1 - m_i) - C^{CA}$$

i.e.

$$U(m) = U^r - \sum_i \tilde{V}_i (1 - m_i), \qquad \tilde{V}_i \coloneqq V_i + \xi \left(\sum_k a^{k,d} \left| v_{x_i}^k \right| \right) \sigma_{x_i} \quad (115)$$

The benefits of attention are increased by the CA term, which is the average value of $\left| v_{x_i}^k \right|$, averaging over actions k.

XX or is it: complexity of each deviation from the average, and complexity of the average?

Ex "0" is clearly very simple: so the "absolute" $\left| v_{x_i}^k \right|$ (rather than the relative $\left| v_{x_i}^k - \bar{v}_{x_i} \right|$) matters).

But if all have a value $+x_n$, then it's edited away in the comparison: so, it's in V_i.

Implication Suppose that we have a "complex" and a simple option. The complex option will suffer the CA penalty, but the simple option won't. Hence, the simple option is more likely to be chosen. This accounts, qualitatively, for the findings in the "choice overload" literature (e.g. Iyengar and Kamenica (2010)).

A.10.3 Application to the Whole Foods example

The consumer's problem There are k slots for the delivery. The value to the consumer is $v^k = h^k - \beta p^k$, where h^k (as in hedonic value) is the convenience of slot k if the slot was costless, p^k is the price, and β is the exchange rate between the dollar price and utils. Trade off both dimensions is difficult. I capture that as saying that $\beta = \bar{\beta} + x$, where x is hard to

perceive. Hence, the perceived value is

$$v^k(x) = h^k - \left(\bar{\beta} + x\right) p^k$$

Hence,

$$\left|v_x^k\right| = \left|p^k\right|$$

I assume that a choice needs to be made. Define $\hat{p}^k \coloneqq p^k - \bar{p}$. So (XXX fill the blanks), we have

$$\left|v_x^k\right| = \left|\hat{p}^k\right|$$

Hence, using (113),

$$C^{CA} = \xi \sigma_x \sigma_p (1 - m_x), \qquad \sigma_p \coloneqq \sum_k a^{k,d} \left|\hat{p}^k\right|$$

where σ_p is the dispersion of prices.

There's also a "macro" complexity, of the agent can opt out not to buy the good (and delivery) at all, e.g. if there's just 1 slot. It's

$$C^{CA} = \xi \pi (1 - \pi) \sigma_x \left|\bar{p}\right| (1 - m_x)$$

where π is the prior probability of buying the good. TBC

Whole Foods' problem Because of CA, Whole Foods will set a uniform price p^k for all delivery slots.

Now, what price will that be? A priori, it could be any value. It's remarkable that the price is historically either 0 or \$9.99, rather than any value, like \$9.52. What's going on, of course, is the psychology of price editing. The consumer perceives a price $p = \$9.82$, as $p = \$9$; or more generally, as $p = 9 + m_d 0.12$,where m_d is the attention to the decimals. Then, the demand curve is discontinuous between \$9.99 and \$10, so the optimal price will be \$9.99, for an open set of parameters. That explains why the price might be \$9.99, or \$4.99, rather than \$9.12 or \$4.83.

But it doesn't quite explain why the price will be \$0. Well, it's clear that a price of \$0.82 is not perceived at 0.

For the price "\$0", it's tempting to assume (realistically) that "free" is simple, and very attractive, in a way that "cost 0.1" is not. The modelling is the following. A first pass filter finds "simple to interpret" messages (e.g. free, certain); when things are not "fully simple".

Price $p = x$ is perceived at

$$p^s = mx + \xi x (1 - m) = [m + \xi (1 - m)] x = x - (1 - m)(1 - \xi) x$$

B Omitted proofs

Proof of Proposition 4 Then, the value function has a term (with $b \geq 0$): $Ab - Bb^2$ with $A = 2b_r - 2\xi\sigma_a$ and $B = 2 - m$. The optimum (subject to $b \geq 0$) is $b = \frac{A}{2B}$, so the value function is $\frac{A^2}{4B} = Bb^2$. Hence,

$$V(m) = b(m)^2 (2 - m) - C(m) \tag{116}$$

Hence, if $C(m) \geq b(m)^2 (2 - m)$, then $V(m) \leq 0 = V(0)$. So the optimum has m and b equal to 0.

Proof of Proposition 5 We have:

$$\mathbb{E}[(p(x) - c(x)) a(x^s)] = \mathbb{E}[(b - 1) x (-\psi b x^s)] = (1 - b) b\psi \mathbb{E}[x x^s] = \left(b - b^2\right) \psi\sigma_x^2 m$$

and as $\mathbb{E}[(p(x) - c(x))] = 0$, the total producer surplus is:

$$\mathbb{E}[(p(x) - c(x))(1 + a(x^s))] = \left(b - b^2\right) \psi\sigma_x^2 m \tag{117}$$

This implies:

$$\begin{aligned} W &= V^{\text{cons}}(m) + \mathbb{E}[(p(x) - c(x))(1 + a(x^s))] \\ &= \frac{1}{2}\psi\sigma_x^2 b^2 - C^{\text{CA}} - C(m) + \psi\sigma_x^2 \left(b - b^2\right) m \\ &= \left(b - \frac{b^2}{2}\right) \psi\sigma_x^2 m - C^{\text{CA}} - C(m) \end{aligned}$$

Proof of Proposition 7 Recall that with $\chi = \frac{1}{2}\psi\sigma_x^2$, social welfare is:

$$W = \Pi + V^{\text{cons}}, \qquad \Pi = 2\left(b - b^2\right) m\chi, \qquad V^{\text{cons}} = b^2 m\chi - C(m)$$

where Π is the average profit (117), and V^{cons} is the consumer's expected utility. Hence,

$$W_b = \Pi_b + V_b^{\text{cons}} + \left(\Pi_m + V_m^{\text{cons}}\right) m_b$$

As the consumer optimizes m to maximize V^{Cons}, we have $V_m^{\text{Cons}} = 0$, so

$$W_b = 2(1-b)\, m\chi + 2\left(b - b^2\right) m_b \chi \tag{118}$$

For all $b \in [0,1)$, as $m_b \geq 0$, we have $W_b \geq 0$; and whenever $m > 0$, we have $W_b > 0$. Hence, $b = 1$ is the social optimum.

We next study when we have some $m > 0$. We observe that the first order condition for m is $C'(m) = b^2\chi$.

If $C'(0) < \chi$, when $b \in [0,1]$ is such that $b^2\chi > C'(0)$ we have $m > 0$, and $b = 1$ is a strict global optimum.

If $C'(0) \leq \chi$, then naive consumers always give $m = 0$ (they "give up" on paying attention), and any $b \in [0,1]$ gives the same social surplus. But with a minuscule fraction of rational consumers (as in Proposition 24) shifts the social optimum to $b = 1$, strictly.

Proof of Proposition 8 Suppose now that we have endogenous attention. Then, as $b = 0$, $W_m = 0$, so

$$\frac{d}{db} W_{|b=0} = \psi\sigma_x^2 m^d - \xi\sigma_a\sigma_x sign(b)\left(1 - m^d\right)$$

so locally, $\frac{d}{db} W_{|b=0} \leq 0$ iff the same condition (39) holds. We note that a simple way to think about this is that the allocative benefits are $\sim bm$, while CA costs are ξb, so that around $b = 0$, rigidity is better if ξ is large enough.

The consumer's attention is: $\max_m V^{\text{cons}}(m)$. It useful to define with

$$V(m) \coloneqq V^{\text{cons}}(m) - \frac{1}{2}\psi\sigma_x^2 b^2$$

where $\frac{1}{2}\psi\sigma_x^2 b^2$ is the surplus with a rational agent. Hence, $V(m) \leq 0$ and it is easier to reason with it. We have

$$V(m) \coloneqq V^{\text{cons}}(m) - \frac{1}{2}\psi\sigma_x^2 b^2 = \frac{1}{2}\psi\sigma_x^2 b^2 m - \xi\sigma_a |b| \sigma_x (1-m) - C(m) - \frac{1}{2}\psi\sigma_x^2 b^2$$

hence

$$V(m) = -B(1-m) - C(m), \qquad B \coloneqq \frac{1}{2}\psi\sigma_x^2 b^2 + \xi\sigma_a |b| \sigma_x, \tag{119}$$

which gives $m^* = \text{argmax}_m V(m)$. . By (117) social welfare is: $W = V^{\text{cons}}(m) + (b - b^2)\psi\sigma_x^2 m$, i.e.

$$W = V(m) + \left(b - \frac{b^2}{2}\right)\psi\sigma_x^2 m \tag{120}$$

Suppose that $C(1) > \psi\sigma_x^2$ and ξ is large enough. Let us show that with endogenous attention, a perfectly simple contract is optimal.

When $b = 0$, welfare W is 0. So, a sufficient condition is: for all $b \in (0, 1]$ (later we'll say "all b"), $W \leq 0$, i.e.

$$-V(m^*) \geq \left(b - \frac{b^2}{2}\right) \psi \sigma_x^2 m^*$$

A sufficient condition for that is: for all b

$$-V(m^*) \geq b \psi \sigma_x^2$$

i.e.

$$\min_m \left(\frac{1}{2}\psi \sigma_x^2 b^2 + \xi \sigma_a b \sigma_x\right)(1 - m) + C(m) \geq b \psi \sigma_x^2$$

i.e. (after division by b), for all b,

$$\min_m \left(\frac{1}{2}\psi \sigma_x^2 b + \xi \sigma_a \sigma_x\right)(1 - m) + \frac{C(m)}{b} \geq \psi \sigma_x^2$$

A sufficient condition for that is that

$$\xi \sigma_a \sigma_x (1 - m) + C(m) \geq \psi \sigma_x^2, \qquad \forall m \in [0, 1]$$

i.e. $\xi \geq \xi^* \coloneqq \max_m \frac{\psi \sigma_x^2 - C(m)}{\sigma_a \sigma_x (1-m)}$.

Proof of Proposition 9 If $b = 0$, social welfare is $W(0) = 0$. Now suppose that the firm sets $b = 1$. Then, optimal attention allocation follows $\max_m Bm - C(m)$ with $B = \frac{1}{2}\psi \sigma_x^2 + \xi \sigma_a \sigma_x$. So, by the proposition's assumption, we have $C'(1) < \frac{1}{2}\psi \sigma_x^2 \leq B$, so consumer attention is $m = 1$. Hence, welfare is $W(1) = \frac{1}{2}\psi \sigma_x^2 - C(1)$. As $C(1) \leq C'(1)$ by convexity, we have $C(1) < \frac{1}{2}\psi \sigma_x^2$ and $W(1) > 0 = W(0)$. Choosing $b = 1$ yields greater social welfare better than $b = 0$.

Proof of Proposition 10 [Don't read too closely: The proof is long, has some typos, and I'm in the process of streamlining it]

I call $\bar{\psi}$ the average elasticity (so $\bar{\psi} = \frac{\psi_{\min} + \psi_{\max}}{2}$); $v_\psi = \frac{\psi_{\max} - \bar{\psi}}{\bar{\psi}}$ the half-width of the range of taxes; and $\hat{\psi}_g \coloneqq \frac{\psi_g - \bar{\psi}}{\bar{\psi}}$, so that $\frac{\hat{\psi}_g}{v_\psi}$ is uniformly distributed in $[-1, 1]$. I normalize the mass of goods to be 1. I define the thresholding function

$$T(y, \kappa) = \text{sign}(y) \max(|y| - \kappa, 0) = \underset{x}{\text{argmin}} \frac{1}{2}(y - x)^2 + \kappa x \tag{121}$$

As said in the main text, I take the limit where v_ψ is "small", and of the same order of magnitude at $\bar{\xi}$ defined in (123). I look for an odd number of categories, $n = 2N + 1$ for an

integer N. This is without loss of generality, as the central category could have a mass of 0.

Lemma 4. *The tax categories, in $\frac{\hat{\psi}_g}{v_\psi}$ are of the form: a central one in $[-b, b]$ for some $b \in (0, 1]$; and otherwise, $2N$ equal-sized categories partitioning $[-1, 1] \setminus [-b, b]$ into $2N$ intervals of size $\delta_{\hat{\psi}} \coloneqq \frac{1-b}{N}$, so intervals $I_k = (b + (k-1)\,\delta_{\hat{\psi}}, b + k\delta_{\hat{\psi}}]$ for $k = 1, \ldots, N$ and symmetrically in the negatives.*

[Proof to be typed in]

Next, we determine the optimal tax in category G_k. We call $\hat{\psi}_G$ the average elasticity gap in category k, so

$$\frac{\hat{\psi}_{G_k}}{v_\psi} = b + \left(k - \frac{1}{2}\right)\delta. \tag{122}$$

We call $\hat{\tau}_G \coloneqq \frac{\tau_G - \bar{\tau}}{\bar{\tau}}$ the special tax for τ_G.

The average welfare coming from category k is:

$$V_{G_k} \coloneqq \frac{-1}{2}\bar{\psi}\left(1 + \hat{\psi}_{G_k}\right)\mathbb{E}\left[\bar{\tau}^2\left(1 + 2m\hat{\tau}_{G_k} + m\hat{\tau}_{G_k}^2\right)\right] - \xi\sigma_{c_G}\bar{\tau}\left|\hat{\tau}_{G_k}\right|\left(1 - m_{G_k}\right) + \lambda\bar{\tau}\left(1 + \hat{\tau}_{G_k}\right)$$

The planner's problem (given attentions, for now taken as exogenous) is (where s_{G_k} is the mass in each category $k \neq 0$)

$$\max_{\hat{\tau}_{G_k}} \sum_k s_{G_k} V_{G_k} \text{ s.t. } \sum_k s_{G_k}\hat{\tau}_{G_k} = 0$$

The last condition means that the average tax deviation is 0. We call $\mu - 1$ the multiplier for that last condition. We observe that $\bar{\tau} = \frac{\lambda}{\bar{\psi}}$,so the Lagrangian is

$$\frac{L}{\lambda\bar{\tau}} = \frac{1}{\lambda\bar{\tau}}\left(\sum_k s_{G_k} V_{G_k}\right) + (\mu - 1)\sum_k s_{G_k}\hat{\tau}_{G_k}$$

So, with

$$\bar{\xi} \coloneqq \xi\frac{\sigma_{c_G}}{\lambda m}\left(1 - m\right) \tag{123}$$

the Lagrangian is

$$\frac{L}{\lambda\bar{\tau}} = \sum_k \left[\frac{-1}{2}\left(1 + \hat{\psi}_{G_k}\right)\left(1 + 2m\hat{\tau}_{G_k} + m\hat{\tau}_{G_k}^2\right) - m\bar{\xi}\left|\hat{\tau}_{G_k}\right| + 1 + \mu\hat{\tau}_{G_k}\right] s_{G_k} \tag{124}$$

The FOC is

$$-\left(1 + \hat{\psi}_{G_k}\right) m\left(1 + \hat{\tau}_{G_k}\right) - m\bar{\xi}\,\mathrm{sign}\left(\hat{\tau}_{G_k}\right) + \mu = 0$$

This gives, assuming an interior solution with $\hat{\tau}_{G_k} \neq 0$,

$$1 + \hat{\tau}_{G_k} = \frac{\mu - m\bar{\xi}\,\mathrm{sign}\,(\hat{\tau}_{G_k})}{m\left(1 + \hat{\psi}_{G_k}\right)} \tag{125}$$

Now, we solve by taking the limit of small $\hat{\tau}_{G_k}$ and $\hat{\psi}_{G_k}$, neglecting the higher-order terms. This gives,

$$\hat{\tau}_{G_k} = \frac{1}{m}\left[\mu - m\bar{\xi}\,\mathrm{sign}\,(\hat{\tau}_{G_k})\right]\left(1 - \hat{\psi}_{G_k}\right) - 1.$$

Now, we use the condition that $\langle \hat{\tau}_{G_k} \rangle = 0$, which gives, as $\left\langle \hat{\psi}_{G_k} \right\rangle = 0$ (where the mean is taken over the non-central categories, $k \neq 0$), $\frac{1}{m}\left[\mu - m\bar{\xi}\left\langle \left|\hat{\psi}_{G_k}\right| \right\rangle\right] - 1 = 0$, i.e.

$$\mu = m + m\bar{\xi}\left\langle \left|\hat{\psi}_{G_k}\right| \right\rangle \tag{126}$$

so, keeping only first order terms,

$$\hat{\tau}_{G_k} = \frac{1}{m}\left[m + m\bar{\xi}\left\langle \left|\hat{\psi}_{G_k}\right| \right\rangle - m\bar{\xi}\,\mathrm{sign}\,(\hat{\tau}_{G_k})\right]\left(1 - \hat{\psi}_{G_k}\right) - 1 = -\hat{\psi}_{G_k} - \bar{\xi}\,\mathrm{sign}\,(\hat{\tau}_{G_k})$$

We record this:

Lemma 5. *The optimal tax deviation for a category k is:*

$$\hat{\tau}_{G_k} = T\left(\hat{\psi}_{G_k}, \bar{\xi}\right) \tag{127}$$

where T is the thresholding function (121) and $\bar{\xi}$ is in (123). So, when $\hat{\tau}_{G_k}$ is not zero, it is

$$\hat{\tau}_{G_k} = -\left(\hat{\psi}_{G_k} - \bar{\xi}\,\mathrm{sign}\left(\hat{\psi}_{G_k}\right)\right). \tag{128}$$

This is like the Ramsey rule (which would give $\hat{\tau}_{G_k} = -\hat{\psi}_{G_k}$): higher elasticities lead to lower taxes. Also, there is a dampening, via the term $\bar{\xi}$, which signifies that there is a Lasso-like shrinkage towards 0, to mitigate complexity aversion.

Let us next calculate the gain V_{G_k} from moving $\hat{\tau}_{G_k}$ from zero to its optimal value. Observe that if we optimize $\max_x -\frac{1}{2}Ax^2 + Bx$, the maximand is at $x = \frac{B}{A}$, and the value is $\frac{B^2}{2A} = \frac{A}{2}x^2$. Applying this to (124), with $x = \hat{\tau}_{G_k}$, $A = \left(1 + \hat{\psi}_{G_k}\right)m$, we get (to the leading order):

$$\frac{1}{\lambda\bar{\tau}}V_{G_k} = \frac{1}{2}m\left(\hat{\tau}_{G_k}\right)^2 s_{G_k} \tag{129}$$

This is

$$V' = \frac{1}{\lambda\bar{\tau}}V \quad = \quad 2\frac{1}{2}m\sum_{k=1}^{N}\left(\hat{\tau}_{G_k}\right)^2\delta = m\delta\sum_{k=1}^{N}\left(v_\psi\left(b+\left(k-\frac{1}{2}\right)\delta\right)-\bar{\xi}\right)^2$$

with $\delta = \frac{1-b}{N}$. Normalizing $v_\psi = 1$ we get:

$$\tilde{V} \coloneqq \frac{1}{\lambda\bar{\tau}m}V_N = \frac{(1-\bar{\xi})^3-(b-\bar{\xi})^3}{3} - \frac{(1-b)^3}{12N^2}$$

So

$$\frac{d\tilde{V}}{db} = -(b-\bar{\xi})^2 - \frac{(1-b)^2}{4N^2}$$

which is 0 for $b = \bar{\xi} \pm \frac{\delta}{2}$. When b goes from $-\infty$ to $+\infty$, $\frac{d\tilde{V}_N}{db}$ is negative, then positive in $\left(\bar{\xi}-\frac{\delta}{2}, \bar{\xi}+\frac{\delta}{2}\right)$, then negative. Hence, the maximum of $\tilde{V}_N$ is reached for:

$$b^* = \bar{\xi} + \frac{\delta}{2} \tag{130}$$

which is the value of b for which the optimal tax in the central band, $\hat{\tau}_0 = b - \frac{\delta}{2} - \bar{\xi}$, is just 0. Also, as $b^* = \bar{\xi} + \frac{1-b^*}{2N}$, we have more explicitly:

$$b^* \coloneqq \frac{\bar{\xi} + \frac{1}{2N}}{1+\frac{1}{2N}} \tag{131}$$

When $b = b^*$, this is

$$\frac{1}{\lambda\bar{\tau}m}V = \frac{(1-\bar{\xi})^3}{3}\left(1-\frac{1}{(2N+1)^2}\right) = \frac{1}{3}(1-\bar{\xi})^3\left(1-\frac{1}{n^2}\right) \tag{132}$$

I record total welfare here:

Lemma 6. *With n tax categories (with optimally chosen bounds), welfare is*

$$V_n = V_1 + \frac{1}{3}\lambda\bar{\tau}v_\psi^2 m(1-\frac{\bar{\xi}}{v_\psi})^3\left(1-\frac{1}{n^2}\right) \tag{133}$$

where $\bar{\xi}$ is in (123), and V_1 is the welfare with just a uniform tax.

Hence the situation is as follows.

First, if we have complexity aversion, and theres' an upper bound $M < 1$ on attention ($m \leq M < 1$), and ξ large enough, a uniform tax is best, $n^* = 1$. Also, with endogenous attention, then the optimal n can be any odd number, but it decreases with ξ. IIt is always

1 for ξ large enough. [60]

In contrast, take the case without complexity aversion ($\xi = 0$). With exogenous attention, we aways have $n^* = \infty$: we want a continuum of tax rates. ndeed, the optimal tax is the Ramsey tax, without modification.[61] With exogenous attention, the optimum must have $n \geq 3$: this is because the maximum of $m\left(1 - \frac{1}{n^2}\right)$ is always achieved for a $n > 1$ (as this is 0 at $n = 1$). One does not get a uniform tax rate as the optimum.

Hence, complexity aversion is the only way to get $n^* = 1$ as an optimum; and this is obtained in an open set of parameters.

For a fixed m, a higher number n of categories slightly increases welfare. But higher attention does increase welfare.[62] Hence, we have $V(m, N)$ with $V_m \geq 0$ and $V_N \leq 0$.

There are extensions within reach: e.g. Ramsey taxes with redistributive motives, Pigouvian taxes.

In the future, it's probably the case the numerical methods will be most useful. So, clearly write the problem for that.

Proof of Proposition 11. For this utility $\tilde{U}(a) = -\frac{1}{2}(a - x)^2 + vx$, we have

$$\tilde{U}^{\text{active}}(m) = -B(1 - m) - C(m), \qquad B = \frac{1}{2}\sigma_x^2 + \xi\sigma_a\sigma_x \tag{134}$$

while the passive decision is $U^{\text{passive}} = -\frac{1}{2}\left[\left(a^d\right)^2 + \sigma_x^2\right]$. We have

$$\begin{aligned}\tilde{U}^{\text{active}} &= -B(1 - m) - C(m) + \mu vx = -B(1 - m) - C(m) - \mu vx(1 - m) + \mu vx \\ &= -\tilde{B}(1 - m) - C(m) + \mu vx\end{aligned}$$

with $\tilde{B} = B + \mu vx$. We observe that

$$\min_{m \in [0,1]} \tilde{B}(1 - m) + \kappa m = \min\left(\tilde{B}, \kappa\right) \tag{135}$$

[60] Indeed, calling $v = \frac{1}{3}\ln(V - V_1)$, we have $v = \ln(1 - \xi a\frac{1-m}{m}) + g(m, n)$ with g independent of ξ, for a constant a. So,

$$v_\xi = \frac{-\frac{1-m}{m}}{1 - \xi a\frac{1-m}{m}}$$

so $v_{\xi m} \geq 0$, and as m decreases with n,we have $v_{\xi n} \leq 0$.

[61] This may seem surprising, as normally, a lower attention leads to a higher tax (as it makes the tax less distortionary), see Farhi and Gabaix (2020). The reason is that here, we assume that the average tax is correctly $\bar{\tau}$ perceived (put another way, the average tax deviation $\hat{\tau}_G$ has mean 0). Hence, limited attention does not lead the planner to increase taxes.

[62] Lower attention to the average tax $\bar{\tau}$ would increase welfare, as in Farhi and Gabaix (2020). But here we take the case of agents who not the average tax, and are just inattentive to the categorical deviations from it.

as indeed the optimum attention is $m = 1_{\kappa \leq \tilde{B}}$. So, optimizing on attention m, with $C(m) = \kappa m$, we get:

$$\tilde{U}^{\text{active}} = -\min\left(\tilde{B}, \kappa\right) + \mu v x = -\min\left(\tilde{B} - \mu v x, \kappa - \mu v x\right) = -\min\left(B, \kappa - \mu v x\right)$$

Hence,

$$\tilde{U}^{\text{active}} = -\min\left(B, \kappa - \mu v x\right)$$

Hence, we have an passive decision iff $\tilde{U}^{\text{active}} \leq U^{\text{passive}}$, i.e. iff

$$\min\left(B, \kappa - \mu v x\right) \geq \frac{1}{2}\left[\left(a^d\right)^2 + \sigma_x^2\right]$$

i.e.

$$\min\left(\xi \sigma_a \sigma_x, \kappa - \frac{\sigma_x^2}{2} - \mu v x\right) \geq \frac{1}{2}\left(a^d\right)^2$$

Proof of Proposition 13 The complexity costs are, in general:

$$C^{\text{CA}} = \xi \sum_{i,t} \mathbb{E}\left|\tilde{a}_t u_{a_t a_t} a_{t,x_i} x_i\right|\left(1 - m_i\right) \tag{136}$$

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

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

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

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

Proof of Proposition 15 Let us call $g(x) = \gamma x - \frac{1}{2}\theta x^2$ the welfare impact of output gap x. The Lagrangian is:

$$L = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[-\left|\pi_t\right| + g\left(x_t\right) - \Xi_t\left(-\pi_t + \beta M^f \pi_{t+1} + \kappa x_t + \nu_t\right)\right], \tag{137}$$

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

$$\begin{aligned} g'(x_t) &= \Xi_t \kappa && (138) \\ \operatorname{sign}(\pi_t) &= \Xi_t - M^f \Xi_{t-1} 1_{t>0} && (139) \end{aligned}$$

hence

$$\operatorname{sign}(\pi_t) = g'(x_t) - M^f g'(x_{t-1}) 1_{t>0} \tag{140}$$

which gives (64).

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) (64) holds with $|\operatorname{sign}(\pi_t)| \leq 1$, 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 (65).

If this inequality is violated, then the optimum is remains by (64).

Proof of Proposition 16 From the Phillips curve, the steady state values π, x satisfy

$$\pi = \frac{\kappa x + \nu}{1 - \beta M^f} \tag{141}$$

and the optimality condition (64) gives: $\kappa \operatorname{sign}(\pi) = \left(1 - M^f\right) g'(x)$ i.e.

$$\kappa \operatorname{sign}(\pi) = \left(1 - M^f\right)(\gamma - \theta x) \tag{142}$$

We have $\pi = 0$ iff $x = \frac{-\nu}{\kappa}$ and (142) 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, (141) gives $\pi = \frac{\nu - \nu_+}{1 - \beta M^f}$ Likewise, when $\nu < \nu_-$, we have $\pi < 0$ and $x = x_-$.

Proof of Proposition 17 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 18 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.

$$W_t = \bar{W} - h(\hat{\pi}_t) + g(x_t) \tag{143}$$

with

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

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 $\hat{\pi}_t$.

Proof of Proposition 24 With a fraction ε of rational consumers, the welfare function has now a term:

$$W = W^0(1-\varepsilon) + \varepsilon\frac{1}{2}\psi\sigma_x^2\left(b - \frac{b^2}{2}\right)$$

where W^0 is the social welfare with behavioral aversion, and $\frac{1}{2}\psi\sigma_x^2\left(b - \frac{b^2}{2}\right)$ is the social surplus linked to rational consumers.

With first order CA, we had (**??**). Now, this gives:

$$W \leq \underline{W}(b) \coloneqq -(1-\varepsilon)\min\left(kd^{1-\gamma}, d\right) + b.$$

So, the exact same reasoning as in the proof of Proof of Proposition 23 shows that, if ξ is large enough, $\underline{W}(b) < 0$ for $b > 0$. So, $b = 0$ is preferable.

With second order CA, (144) implies that $W^0 \geq -Eb^2$ for a positive constant E. So, we have

$$W \geq -(1-\varepsilon)Eb^2 + \varepsilon b$$

hence social surplus W is strictly positive for a small b. As a result, some $b > 0$ is preferable to $b = 0$.

Proof of Proposition 25 I build on the proof of Proposition 23. I use the attention function in (81). Let us analyze the incentive to go from $b = 0$ to a small $b > 0$. First, we have $\chi = \frac{1}{2}\psi\sigma_x^2 b^2$, $m = \left(\bar{\xi}\frac{\chi}{wc}\right)^{\frac{1}{\alpha}-1} = \bar{k}b^{\frac{2}{\alpha}-2}$, for a positive constant $\bar{k}$ independent of b, and one can show that the cost is: $C = (1-\alpha)m\bar{\xi}\chi = O(b^2)m$, so that, using (37), in the

neighborhood of $b = 0$, we have

$$\begin{aligned} W &= \left(b - \frac{1}{2}b^2\right)\psi\sigma_x^2 m + m\bar{\xi}\frac{1}{2}\psi\sigma_x^2 b^2 - C(m) - \bar{\xi}\frac{1}{2}\psi\sigma_x^2 b^2 \\ &= mb(D + o(1)) - Eb^2, \qquad \text{where } D, E \text{ are positive constants} \\ W &= \bar{k}b^{\frac{2}{\alpha}-1}(D + o(1)) - Eb^2 \end{aligned} \tag{144}$$

Hence, if $\alpha \in \left(\frac{2}{3}, 1\right)$ we have $\frac{2}{\alpha} - 1 < 2$, so $W(b) > 0$ for b in a neighborhood of 0: locally $b = 0$ is worse than $b > 0$.

If $\alpha < \frac{2}{3}$, then the reverse is true then, $W(b) < 0$ for small b's, and $b = 0$ is a local optimum. One can also show that, if ξ is large enough, this is a global optimum, as in the proof of Proposition 8.