Inattentive Economies*

George-Marios Angeletos† **Karthik A. Sastry**‡

March 8, 2021

Abstract

We study the efficiency of inattentive but otherwise frictionless economies by augmenting the Arrow-Debreu framework with a general form of rational inattention. An appropriate amendment of the First Welfare Theorem holds when the model of inattention satisfies an invariance condition roughly equal to assuming mutual-information costs. Away from this benchmark, a planner can help agents economize attention and reduce mistakes by regulating the stochastic properties of prices or replacing markets with other means of communication and coordination. Our main result therefore links Hayek's (1945) argument about the informational optimality of the price system to Sims's (2003) specification of attention costs and experimental tests thereof. Additional results address equilibrium existence, the Second Welfare Theorem, and whether agents pay attention to objects other than prices.

*An earlier version circulated in December 2018 under the title "Welfare Theorems for Inattentive Economies." We thank Jakub Steiner and John Leahy for discussing our paper at the 2019 and 2021 ASSA meetings, respectively; and Daron Acemoglu, Benjamin Hébert, Jennifer La'O, Stephen Morris, Alessandro Pavan, and Harald Uhlig for comments and stimulating discussions. Angeletos acknowledges the support of the National Science Foundation under Grant Number SES-1757198.

†MIT and NBER; angelet@mit.edu

‡MIT; ksastry@mit.edu

1 Introduction

People are inattentive, forgetful, impulsive, and otherwise cognitively constrained. In such circumstances, it is natural to question the functioning of the invisible hand. Sims (2010) warns that prices "cannot play their usual market-clearing role" in inattentive economies, or economies in which both buyers and sellers are inattentive; and review articles by Maćkowiak, Matějka, and Wiederholt (2018) and Gabaix (2019) claim that the Fundamental Welfare Theorems fail in the same context. The latter claim is of course valid if market outcomes are judged relative to non-market alternatives that remove the friction (as implicitly or explicitly assumed by the aforementioned authors). But can market outcomes be improved upon without undoing people's cognitive constraints?

A related question concerns Hayek's (1945) classic argument about the informational optimality of the price system. In his words:

> We must look at the price system as a mechanism for communicating information if we want to understand its real function. [...] The most significant fact about this system is the economy of knowledge with which it operates, or how little the individual participants need to know in order to be able to take the right action.

This argument, at least as formalized by Grossman (1981), presumes not only that markets are complete but also that prices are observed perfectly and costlessly. But if attention to prices is costly, what exactly is the "economy of knowledge" achieved by markets? And if people make mistakes when reacting to price changes, do markets still provide the best means of communication and coordination?

In this paper, we show how to reduce all these inherently general-equilibrium questions to a simpler, decision-theoretic question, whose answer can be informed by experimental evidence. Our main result, Theorem 1, states that an appropriate version of the First Welfare Theorem, expanded to require the efficiency in attention choices and the informational content of prices, holds when attention costs satisfy an *invariance condition* that is tightly connected to the mutual-information specification of Sims (1998, 2003), its axiomatization by Caplin et al. (2020), and experimental tests thereof.

This result provides an important benchmark in which inattention does not invite government intervention. But there is a caveat: when the invariance condition holds, the "economy of knowledge" afforded by markets is weak instead of strict, or the same allocation can be implemented with agents' paying attention to one another's fundamentals as opposed to prices. And when our invariance condition does not hold, markets can strictly economize attention costs, but then it is also generally possible to improve upon them by manipulating the manner in which prices encode information. Hayek's (1945) argument is thus turned on its head: markets are likely to be inefficient precisely when they provide a strict economy of knowledge.

Framework and main result. Towards these lessons, we augment the Arrow-Debreu framework with a general form of rational inattention. Following Sims (1998, 2003), Tirole (2015), Woodford (2019) and others, we treat "attention" and "cognition" as interchangeable notions and model them as the choice of a signal subject to a cost. For consumers, the cost is in terms of forgone utility; for firms, it is in terms of resources diverted away from production. We let these costs take an essentially arbitrary functional form, nesting not only mutual information but also any other specification found in the literature.[1]

Two elements of this framework are worth emphasizing. The first is that it gives a new meaning to complete markets, as consumers are assumed to be insured against any noise-driven variation in their marginal utility of wealth. While unrealistic, this assumption plays a familiar role: had we ruled out such insurance from *both* markets and planning alternatives, competitive equilibria would generally be inefficient for basically the same reason as that articulated in Geanakoplos and Polemarchakis (1986), the existence of pecuniary externalities. By muting this familiar source of inefficiency, we both give the best chance to Hayek's argument and isolate the role played by different assumptions on attention costs.

The second key element is that we let agents flexibly obtain information about the exogenous state of nature and/or equilibrium prices. Even though prices will be a function of the state of nature in equilibrium, treating them as a distinct object in the agents' attention problem is essential for formalizing the sense in which prices may or may not "economize knowledge." It also allows us to capture a number of behavioral phenomena that are well-characterized by imperfect perception of prices.[2]

In this context, equilibrium embeds a new function for prices, as their stochastic properties enter attention costs. This opens the door to a "cognitive" externality: one's actions affects others' attention costs, or cognition, by influencing the stochastic properties of prices (e.g., their volatility, contingency, or "sparsity"). The amended notion of Pareto optimality, on the other hand, maps to a planner who cannot undo the cognitive friction itself, but can internalize the aforementioned externality and can even replace prices with other "messages" that agents in turn may pay attention to. This allows us to formalize the question of whether markets are an optimal means for communication and coordination, in the spirit of Hayek (1945).

Clearly, equilibria can be efficient only if the aforementioned cognitive externality is muted. Theorem 1 shows that a sufficient condition for this to be true and, more generally, for Hayek's argument to be correct is that attention costs are invariant to "repackaging" of information in prices. We prove this result in two steps.

We first show that there is no way to improve upon *laissez-faire* outcomes if we restrict the planner's messages to package information in exactly the same way as equilibrium prices. This intermediate result follows from mapping the economy with fixed messages to a "twin" economy that can be readily nested to the standard Arrow-Debreu framework. Put differently, insofar as one abstracts from the aforementioned endogeneity and markets are complete in the sense highlighted earlier, rational inattention translates merely to a modification of preferences and technologies that may influence observable behavior in interesting ways but fits in the familiar conditions for the Welfare Theorems.

We next relax the aforementioned restriction on messages and check if it is beneficial to replace markets with other means of communication and coordination. The answer is no under our invariance condition. This condition, spelled out in Definition 8, allows attention costs to depend on the exogenous stochastic properties of the environment but mutes their dependence on how these properties are transformed by,

[1] Including the full class of posterior-separable costs described by Caplin and Dean (2015) and Denti (2018), the neighborhood-based costs of Hébert and Woodford (2020), and the costs proposed in Pomatto et al. (2018) and Denti et al. (2019).

[2] Chetty et al. (2009) provide evidence consistent with inattention toward prices and a supporting theoretical framework. Additional examples can be found in Section 3.2 of Sims (2010) and Section III of Gabaix (2014).

or "repackaged" in, prices, other endogenous market outcomes (e.g., average trades, industry output, GDP, etc.), or messages sent as part of non-market mechanisms. Under this property, our intermediate result directly translates to efficiency of the market mechanism.

Our invariance condition is closely related to the notion of invariance in information geometry (Amari and Nagaoka, 2000; Amari, 2016) and the axiom of invariance under compression introduced by Caplin et al. (2020). It is naturally satisfied by Sims's (2003) mutual-information specification, but not by different specifications that appear to find support in experimental work.[3]

The following two examples help shed light on what a failure of invariance, and of Theorem 1, may mean in a market economy. Suppose first that people struggle to discern small price changes. This is inconsistent with mutual-information costs, because mutual information is scale free, but seems consistent with experiments emphasizing "perceptual distance" and can be rationalized with the kind of attention costs proposed in Hébert and Woodford (2020). In these circumstances, welfare can be improved by inducing larger variation in prices, or by requiring that prices such as $99 are rounded up to $100. Suppose next that people struggle to track volatile or "complex" objects. This, too, can be rationalized by an appropriate departure from mutual-information costs. And it calls for the opposite kind of intervention: welfare is now improved by stabilizing or "simplifying" the price system, perhaps even by closing some markets.

Additional results and extensions. Our main result establishes the efficiency of equilibria presuming their existence. But as long as invariance holds, we can establish existence and a version of the Second Welfare Theorem by leveraging once more the aforementioned mapping to "twin" economies. In particular, we show that, while not needed for our main result, the conventional assumption that attention costs are posterior separable (Caplin and Dean, 2015; Denti, 2018) helps guarantee that the modified preferences and technologies of these "twin" economies are convex even even if the primitive preferences and technologies are not. Our existence result and our Second Welfare Theorem then follow from standard arguments.

We next combine invariance with another key property of the mutual-information specification, a form of monotonicity encoding cognitive savings from "coarsening" signals, to more sharply characterize the informational structure of efficient equilibria. We show, first, that equilibrium allocations are pinned down by fundamentals (preferences, endowments, technologies), so there is no room for sunspots or correlation devices; and, second, that every agent pays attention only to variables that *directly* enter their payoffs or budget constraint. This adds to our perspective on Hayek (1945) by clarifying the conditions under which prices are "sufficient statistics" in people's minds about the state of the economy.

Finally, we explain how our results readily extend to settings where agents track endogenous objects other than prices (e.g., GDP, industry output, the trades of others), or to a more powerful planner that can not only replace prices with other messages but also more directly manage people's attention and cognition (e.g., by precluding them from learning directly about fundamentals or by relabeling the state space). And we also discuss how to accommodate behavioral frictions not directly nested in our main framework, like narrow bracketing, bounded recall, general stochastic choice, and sparsity.

[3] We discuss this issue further in Section 6.3.

Related literature. The literature on rational inattention spurred by Sims (1998, 2003) is voluminous. Some works focus on single-agent behavior (Matějka, 2016; Matějka, Steiner, and Stewart, 2015); others study specific macroeconomic models (Maćkowiak and Wiederholt, 2009, 2015) or games (Colombo, Femminis, and Pavan, 2014; Myatt and Wallace, 2012). Our paper's added value is to adapt the analysis of rational inattention to the Arrow-Debreu framework, to develop the appropriate amendments of the Welfare Theorems, and to show how equilibrium efficiency hinges on whether attention costs take Sims's (2003) preferred specification or other specifications proposed by an emerging experimental and decision-theoretic literature.

Letting consumers and firms choose their attention to prices in a market environment is akin to letting players obtain information about others' actions in a large game. This link is explored in a recent, complementary paper by Hébert and La'O (2020). These authors establish efficiency of equilibria in their class of games under two conditions: a close cousin of our invariance condition; and a restriction on payoffs akin to the netting-out of pecuniary externalities in our setting. Hébert and La'O (2020) push the frontier further by showing that invariance is not only sufficient but also generically necessary for efficiency. They also provide a result on non-fundamental volatility which, unlike our Proposition 2, requires only monotonicity and not invariance, and therefore extends to inefficient equilibria.

Related are also Angeletos and La'O (2018), Colombo, Femminis, and Pavan (2014), Gul, Pesendorfer, and Strzalecki (2017), and Tirole (2015). These works focus on different applied questions but share the following feature: they study economies that trivially satisfy our invariance condition, because attention costs are specified as a function of the joint distribution of an agent's signal and the exogenous state of nature *alone*. This explains why the conditions for (in)efficiency found in these works relate exclusively to pecuniary or payoff externalities, as opposed to the kind of cognitive externalities identified here.

Closer to our question of how markets with inattention work, Ravid (2020) studies a bargaining game in which a buyer can flexibly but costly pay attention to a good's quality and a seller's take-it-or-leave-it offer. His model, like ours, includes inattention to the terms of trade. But by focusing on mutual-information costs, his model's equilibrium does not feature the cognitive externality that drives our kind of inefficiency. Instead, it produces inefficiency from the interaction of rational inattention with market power.

At a high level, our paper's cognitive externality is, of course, a specific form of information externality: the information an agent has about payoff-relevant objects (most notably prices) is a function of the choices of other agents. Along with our emphasis on markets, this brings to mind the literature on Noisy Rational Expectations Equilibria (Grossman and Stiglitz, 1980; Laffont, 1985; Vives, 2017). But there are subtle differences. While this literature emphasizes learning from prices about fundamentals, such learning is not essential for our purposes (although it is allowed). Instead, the essential friction is that consumers and firms are inattentive to prices themselves. Furthermore, in this literature the existence of information externalities hinge on missing markets (Grossman, 1976, 1981) and the ideal policy intervention would be to complete the markets. In our context, instead, cognitive externalities are possible despite complete markets, and uncompleting the markets can actually be welfare-improving insofar it helps "simplify" the price system,[4]

[4]One paper in the Grossman-Stiglitz tradition that comes closer to what we do is Vives and Yang (2018). These authors study an asset market in which traders incur a cost for "interpreting" prices (i.e., extracting information about the asset's return). We suspect

Outline. Section 2 illustrates the main ideas with an example. Section 3 sets up our general framework. Section 4 presents our main result regarding efficiency. Section 5 provides additional results about equilibrium existence, the Second Welfare Theorem, the role of prices as sufficient statistics, and efficiency under a looser planning concept. Section 6 uses our results to discuss the link between the price system's "economy of knowledge" and the mutual-information framework. Section 7 discusses possible extensions to other models of behavioral frictions. And Section 8 concludes.

2 Example

In this section, we study a simple, two-good, exchange economy, designed to nest in the linear-quadratic-Gaussian framework used in Sims (2003) and various other works. In this example, efficiency is guaranteed with mutual information costs but not with two different specifications. This finding foreshadows our main result of how efficiency relates to an invariance property at the core of mutual information.

2.1 Set-up

There is a continuum of agents, indexed by $i \in [0,1]$, and two goods, called "coconuts" and "money." Consumer i enjoys the following utility from coconut consumption $x_i \in \mathbb{R}$ and money consumption $y_i \in \mathbb{R}$:

$$u(x_i, y_i) = x_i - \tfrac{1}{2}x_i^2 + y_i \tag{1}$$

Each consumer's endowment of coconuts is $\xi \in \mathbb{R}$, their endowment of money is zero, and the price of coconuts in terms of money is $p \in \mathbb{R}$. The consumer's budget is therefore $px_i + y_i \le p\xi$.

The consumer chooses their demand of coconuts under imperfect perception of both ξ and p. Formally, we treat $z = (\xi, p) \in \mathcal{Z} = \mathbb{R}^2$ as a random variable, require that x_i be measurable in the realization $\omega_i \in \Omega = \mathbb{R}$ of a signal of z, and let the consumer design the signal's relationship with z subject to a cost. Assuming that all relevant distributions admit density functions, we denote the consumer's prior about z by $\pi \in \Pi$, where Π is a set of density functions over $\mathcal{Z}$, and write the density of ω conditional on each z as $\phi(\cdot|z)$. We then denote the collection of such densities by $\phi = (\phi(\cdot \mid z))_{z\in\mathcal{Z}}$; let Φ be the admissible set of such collections; and represent attention costs as $C(\phi,\pi)$, for some function $C : \Phi \times \Pi \to \mathbb{R}$.

Finally, we allow the consumption of money to adjust mechanically so as to meet the budget for all realizations of uncertainty.[5] Solving out for y_i and suppressing the i index, we can write the consumer's problem as the following choice of a consumption rule $x : \mathbb{R} \to \mathbb{R}$ and a signal structure $\phi \in \Phi$:

$$\max_{x,\phi} \int_{\mathcal{Z}}\int_{\Omega}\left(x(\omega) - \frac{x(\omega)^2}{2} + p(\xi - x(\omega))\right)\phi(\omega|z)\pi(z)\,\mathrm{d}\omega\,\mathrm{d}z - C(\phi,\pi) \tag{2}$$

their specific assumptions about signals and costs amount, under the lens of our analysis, to a joint violation of invariance and monotonicity. The first opens the door to inefficiency, the second to non-fundamental volatility. One of our examples, although differently motivated, has a very similar flavor (see Online Appendix E).

[5] The use of a "residual" good to meet budgets is common in applications but is not needed for our main results. See Section 3.

Remark 1. If we were to treat the consumer's prior about z as a fixed primitive, program (2) would be similar to those studied in a growing, decision-theoretic literature on unrestricted information acquisition (e.g., Caplin and Dean, 2015; Caplin et al., 2020; Denti, 2018). But, in our model, the aforementioned prior (and, specifically, the behavior of p) is determined in equilibrium.

Remark 2. Apart from adding with tractability, the quasi-linear specification has two substantial implications. The first is that the optimal consumption of coconuts is insensitive to wealth, which in turn means that consumers do not care to know the endowment ξ per se; they only care to know the price p.[6] This highlights that, unlike Grossman and Stiglitz (1980), the key issue in our context is not how much agents know about a fundamental but rather how attentive they are to prices. The second implication is that the marginal utility of wealth is equated across realizations of the idiosyncratic noise in ω. This is preserved in our general framework, thanks to a strong notion of complete markets, and allows us to abstract from the question of redistributing wealth from agents with "good" noise realizations to agents with "bad" noise realizations.

2.2 Equilibrium and Efficiency

Let (x, ϕ) be a solution to (2). We assume that a law of large numbers applies in the cross section of agents so that the idiosyncratic noise in their signals washes out at the aggregate level. We can therefore write market clearing as $X(\xi, p) = \xi$ for each ξ, where $X(\xi, p) \equiv \int_\Omega x(\omega)\, \phi(\omega \mid \xi, p)\, \mathrm{d}\omega$. This suggests that, although the decision problem in (2) may plausibly treat the pair $z = (\xi, p)$ as an arbitrary random variable, market clearing may impose that the second element of z is merely a transformation of its first element.

This claim is subject to the following qualifications. First, if the function X is non-monotonic, it may be possible to construct equilibria in which p varies with a sunspot variable in addition to ξ. Second, the existence of such sunspot-driven variation in prices may influence attention costs and thereby the entire equilibrium. Both of these possibilities relate to the question of how rich is the exogenous state of nature. Our general model accommodates arbitrary richness for it. Here, we simplify the exposition by restricting it to coincide with the fundamental ξ. We thus define an equilibrium in the following way:

Definition 1. A competitive equilibrium is a combination of a demand function $x : \mathbb{R} \to \mathbb{R}$, a signal choice $\phi \in \Phi$, a price function $P : \mathbb{R} \to \mathbb{R}$, a prior $\pi \in \Pi$ such that :

1. Agents optimize, or (x, ϕ) solves program (2).
2. The market for coconuts clears, or $\int_\Omega x(\omega)\, \phi(\omega \mid \xi, p)\, \mathrm{d}\omega = \xi$ for all (ξ, p).
3. The consumers' prior about z is consistent with the equilibrium price function, or $\pi(\xi, p) = \pi_\xi(\xi) \cdot D(P(\xi) = p)$ for all (ξ, p), where π_ξ is the density of ξ and $D(\cdot)$ is the Dirac delta function.

This is like a textbook definition but for two modifications. First, individual optimality is extended to incorporate the optimal signal choice. And second, the agents' prior about z, which matters only because

[6] We could have introduced an attentive supplier, who owned the coconuts but valued only money. Then, from the perspective of the inattentive consumer, ξ would be the random supply of coconuts, with no intrinsic interest independent from the price.

it enters attention costs, is generated by combining the exogenous prior about ξ with the equilibrium price function. This is where Rational Inattention (RI) meets Rational Expectations Equilibrium (REE).

It is useful to re-cast this definition in the familiar metaphor of a Walrasian auctioneer, to obtain a firmer sense of how, paraphrasing Sims (2010), prices *can* "play their usual market-clearing role" despite inattention to prices. Before ξ and p are realized, consumers use their knowledge of the price functional P to construct their prior over $z = (\xi, p)$ and to design their optimal signal. Then, nature chooses ξ and the auctioneer sets $p = P(\xi)$. Agents observe their noisy signals of p (and of ξ) and submit demands based on these signals.This yields at the aggregate level a demand for coconuts that is a deterministic function of p, thanks to our assumption of a continuum of agents (and an appropriate law of large numbers). Finally, the loop closes by requiring that P is such that this demand meets supply, for on any possible realization of ξ. Clearly, this is similar to how prices clear markets in standard, attentive economies—and this is precisely the point.

Now consider a benevolent planner that cannot eliminate the cognitive friction but can otherwise regulate consumers' behavior. To fix ideas, suppose that the planner can tax any consumer as a flexible function of the state, the consumer's demand for coconuts, and the consumer's choice of signal. Clearly, the planner can induce any combination of consumption, attention, and prices with such a panoply of tax instruments. This motivates the following notion of efficiency.[7]

Definition 2. A triplet (x^*, ϕ^*, M^*) is efficient if it solves the following program:

$$\begin{aligned} \max_{\phi\in\Phi, x:\Omega\to\mathbb{R}_+, M:\Theta\to\mathbb{R}^N} & \int_{\mathcal{Z}}\int_{\Omega}\left(x(\omega) - \frac{x(\omega)^2}{2}\right)\phi(\omega|z)\,\pi(z)\,\mathrm{d}\omega\,\mathrm{d}z - C(\phi,\pi) \\ \text{s.t.} & \int_{\Omega} x(\omega)\,\phi(\omega\mid\xi,m)\,\mathrm{d}\omega = \xi \text{ for all } (\xi,m)\in\mathcal{Z} \\ & \pi(\xi,m) = \pi_{\xi}(\xi)\cdot D[M(\xi)=m] \text{ for all } (\xi,m)\in\mathcal{Z} \end{aligned} \tag{3}$$

where $D(\cdot)$ is the Dirac delta function. An equilibrium (x, ϕ, P) is efficient if it solves the above with $M = P$. An economy is efficient if all equilibria are efficient.

The change of notation from the price p and the functional P to, respectively, the variable m and the functional M highlights the following point. As already mentioned, prices play a dual role in equilibrium: the traditional one of affecting incentives and clearing markets; and the novel one of being a component of the variable z that agents try to learn or understand. But in the planner's solution, incentive compatibility is not an issue and market clearing is replaced by resource feasibility. It follows that prices remain relevant only because of the second role: they are merely "messages" (hence, m and M) that enter agents' cognition.

This facilitates the following, literal interpretation of the problem defined above. First, a central planner (in place of the Walrasian auctioneer) observes ξ and announces the message m according to rule M. Next, the planner recommends to every agent a signal structure ϕ for learning about $z = (\xi, m)$ and consumption plan. And finally, all agents obediently follow these recommendations.[8]

[7] When writing the planner's problem in (3) and the consumer's problem in (2), we *presume* that the respective maximums exist. The technicalities are taken care of in our general analysis. Also, our general analysis poses efficiency in terms of Pareto dominance. The planner interpretation is used only to ease the exposition.

[8] As in textbook welfare analysis, our planner has no problem in either observing the state of nature or controlling agents' strate-

This mechanism nests "free markets" as a special case, but offers the planner two degrees of freedom: the regulation of the agents' consumption and attention strategies; and the use of an arbitrary message in lieu of the price. The question of interest thus boils down to this: when are these degrees of freedom redundant?

2.3 Specialization: Gaussian Structure

For the rest of the example, we assume that $\xi \sim N(0,1)$ and that agents can only choose signals ω_i that are jointly Gaussian with z, or else they face an infinite cost. This buys us the following simplifications. First, we can represent any feasible signal as

$$\omega_i = a_1 z_1 + a_2 z_2 + a_3 \eta_i \tag{4}$$

where z_1 stands for ξ, z_2 stands for p or m, $\eta_i \sim N(0,1)$ is idiosyncratic noise, and $a = (a_1, a_2, a_3) \in \mathbb{R}^3$ is a vector of coefficients. Second, we can infer that in any equilibrium p is linear in ξ, or else z would fail to be Gaussian, agents would be unable to comprehend z (i.e., choose any signal about it), and an equilibrium would not exist. We can therefore write $P(\xi) = \psi_0 - \psi_1 \xi$ for some scalars ψ_0 and ψ_1; and since we have normalized the mean and variance of ξ, we can also characterize the joint distribution of (ξ, p) by only these scalars. A similar point applies to the rule M and the joint distribution of (ξ, m) in the planner's problem.

Combining these observations, we infer that ϕ and π, the consumer's signal choice and prior, are parameterized by the vectors a and (ψ_0, ψ_1), respectively. For the rest of this section, we therefore write attention costs as $C(a;\psi_0,\psi_1)$. The dependence of C on (ψ_0,ψ_1) underscores how the "translation" of fundamentals via prices or other messages may influence attention costs. We next show how this effect is muted, and efficiency is obtained, with mutual-information costs but not with two other specifications.

2.4 An Efficient Economy

Let attention costs be a function K of the mutual information between ω and z, where $K : \mathbb{R}_+ \to \mathbb{R}$ is increasing, convex, and differentiable, with $K'(0) = 0$. Because p is a function of ξ, the mutual information between ω and $z = (\xi, p)$ is the same as that between ω and ξ. Attention costs thus reduce to

$$C(a;\psi_0,\psi_1) = c(\delta) \equiv K(-\log(1-\delta)), \quad \text{where} \quad \delta = d(a;\psi_1) \equiv \frac{(a_1 - \psi_1 a_2)^2}{a_3^2 + (a_1 - \psi_1 a_2)^2}. \tag{5}$$

See that δ is the correlation between ω and ξ (or, equivalently, between ω and p whenever $\psi_1 \neq 0$). Hence, in the present context, δ offers a simple measure of attention and $c(\delta)$ gives the corresponding cost.

To compute the benefits of attention, proceed as follows. Suppose that *other* agents, $j \neq i$, choose some $a^e = (a_1^e, a_2^e, a_3^e)$ and let $\delta^e = d(a^e;\psi_1)$. Suppose further that they choose their demands optimally, which means that $x_j(\omega) = 1 - \mathbb{E}[p \mid \omega_j] = 1 - \delta^e p$. It follows that the equilibrium price is given by $P(\xi) = 1 - \frac{\xi}{\delta^e}$. Using this fact along with the fact that i's own optimal demand satisfies $x_i(\omega) = 1 - \delta p$, we can calculate her *ex ante*

gies. Needless to say, this is far from real-world applications, where governments have access to a limited set of policy tools and contingencies.

utility from consumption as follows (detailed derivations can be found in Online Appendix C):

$$b(\delta, \delta^e) \equiv \max_x \mathbb{E}\left[x(\omega) - \frac{x(\omega)^2}{2} + P(\xi)(\xi - x(\omega))\right] = \frac{\delta - 2\delta^e}{2(\delta^e)^2} \tag{6}$$

This makes clear not only that attention is privately valuable ($b_1 > 0$) but also that the gains from attention are most significant when prices are more volatile or, equivalently, others are more *in*attentive ($b_{12} < 0$)

Combining the above expressions of the costs and benefits of attention, we conclude that any equilibrium boils down to a solution of the following fixed-point problem:

$$\delta^e \in \arg\max_\delta \left\{b(\delta, \delta^e) - c(\delta)\right\}. \tag{7}$$

Under the assumed conditions on K (or c), the above has a unique solution and this solution is characterized by the consumer's first-order condition, namely $b_1(\delta^e, \delta^e) = c'(\delta^e)$.

We can similarly reduce the planner's problem to the following:

$$\delta^* \in \arg\max_\delta \{b(\delta, \delta) - c(\delta)\}, \tag{8}$$

or equivalently $b_1(\delta^*, \delta^*) + b_2(\delta^*, \delta^*) = c'(\delta^*)$. Comparing this to the equilibrium counterpart, we see the following key properties. On the benefits side, the planner internalizes the pecuniary externality whereby increasing one agent's attention affects others' budgets and utility. On the costs side, the planner's trade-off is the same as the agents' despite the planner's flexibility to send an arbitrary message.

A direct calculation from (6) yields $b_2(\delta, \delta) = 0$ for any δ. This reflects the fact that pecuniary externalities are muted, thanks to the linearity of preferences to money. It follows that $\delta^e = \delta^*$, or that the economy's unique equilibrium coincides with the planner's solution. Summing up:

Proposition 1. *With mutual-information costs, an equilibrium exists, is unique, and is efficient.*

2.5 Two Inefficient Economies

The economy studied above was efficient because externalities were muted on both the benefits and the costs of attention. We already explained why this was true on the benefit side. We now show that the absence of an externality on the cost side depended critically on the mutual-information assumption. Away from this assumption, inefficiency is possible via what we call *cognitive externalities.*

Consider first a case in which agents can only obtain signals of the form $\omega = p + a_3\eta$ and must pay in proportion to the signal's precision about prices. That is,

$$C(a; \psi_0, \psi_1) = \begin{cases} a_3^{-2} & \text{if } a_1 = 0, a_2 = 1 \\ \infty & \text{otherwise} \end{cases} \tag{9}$$

We take this class of costs to represent imperfect observation/cognition in the physically relevant units of prices. In particular, such a cost implies that it is more difficult to perceive price changes from \$1.99 to \$2.00 than price changes from \$2.00 to \$2.10. As discussed in Section 6.2, such scale-dependent cognition relates loosely to the decision-theoretic work of Hébert and Woodford (2020) on "perceptual distance" and "neighborhood-based" costs of information acquisition.

We can still express the benefits of attention by $b(\delta, \delta^e)$, where b is defined as in (6) and (δ, δ^e) is the pair of the correlations between ξ and, respectively, one's own signal and others' signal. But once expressed in these units, the cost of attention is different than in the mutual-information benchmark. In particular, using $\delta = \psi_1^2/(a_3^2 + \psi_1^2)$, solving the latter for a_3, and replacing the outcome into (9), we get $C(a;\psi_0,\psi_1) = \frac{\delta}{1-\delta}\psi_1^{-2}$. Using the fact that equilibrium price function has $\psi_1 = 1/\delta^e$, we get

$$C(a;\psi_0,\psi_1) = c(\delta,\delta^e) \equiv \left(\frac{\delta}{1-\delta}\right)(\delta^e)^2$$

The cost of achieving a given signal-to-noise ratio regarding ξ, which is the relevant notion of signal "quality" in terms of expected utility gains, increases with *others'* signal-to-noise ratio. Why? When others are more attentive, and their demands are more elastic, equilibrium prices are less dispersed across states of nature. Under the present model, these less dispersed prices are harder to tell apart from one another.

As a second example, we consider an opposite scenario in which it is relatively *easier* for agents to track prices (or messages) when the latter are less dispersed. Agents are again restricted to obtain signals of the form $\omega = p + a_3\eta$ but now incur a cost in proportion to the signal's variance:

$$C(a;\psi_0,\psi_1) = \begin{cases} C_0 - (a_3^2 + \psi_1^2) & \text{if } a_1 = 0, a_2 = 1 \\ \infty & \text{otherwise} \end{cases} \tag{10}$$

for some constant $C_0 > 0$. One may interpret this as a continuous analogue for the idea put forth in Gabaix (2014) that agents prefer to keep track of fewer possible realizations of ω (or concentrate probability on a smaller measure of realizations). Following similar steps as above, these costs can be re-expressed as

$$C(a;\psi_0,\psi_1) = c(\delta,\delta^e) \equiv -\left(1 + \sqrt{\frac{1-\delta}{\delta}}\right)(\delta^e)^2$$

The cognitive externality is now the opposite: higher attention from others raises price volatility, which makes it *less* costly to obtain any given value for δ.

What does this mean for efficiency? In both of the above cases, the equilibrium is characterized by $b_1(\delta^e,\delta^e) = c_1(\delta^e,\delta^e)$. And while we still have $b_2(\delta,\delta) = 0$, meaning that pecuniary externalities are muted, we now have $c_2(\delta,\delta) \neq 0$, or there is a cognitive externality creating a wedge between equilibrium and efficient attention. What changes is the direction of this externality, as summarized below.

Proposition 2. *With the costs described in (9) or (10), the economy is inefficient. Attention is inefficiently high and prices are insufficiently volatile in the first case, and the opposite is true in the second case.*

2.6 Summary

The analysis in this section hints that the efficiency of inattentive economies hinges on whether the cognitive process, as represented by the attention cost specification, are suitably invariant to the stochastic properties of prices, or any other "messages" society could potentially use in their place. But it does not fully clarify the *precise* form of invariance needed for efficiency. In particular, the examples presented above seem to emphasize how steep or flat the mapping from ξ to p is. But what does this mean in a world where there are multiple prices and these are non-linear and possibly non-monotone transformations of multiple fundamental variables, or even of non-fundamental variables?

We address these questions in the rest of the paper. We shall show that the required property encodes invariance not only to rescaling but also to more complicated transformations, including "simplifying" prices or making them "sparser" by removing certain contingencies from them, or perhaps adding noise in them.[9] And we will draw a sharper link between the general-equilibrium questions of interest and the decision-theoretic literature on rational inattention.

3 General Model

This section introduces our general framework, which augments a standard Arrow-Debreu economy with rational inattention. We also define our paper's notions of invariance and monotonicity for attention costs.

3.1 Goods, State of Nature, and Cognition State

Let the *state of nature* be a random variable θ, drawn from a finite set Θ according to probability distribution $\pi_\theta \in \triangle(\Theta)$, with $\pi_\theta(\theta) > 0$ for all $\theta \in \Theta$. This encodes not only any shocks to fundamentals (endowments, preferences and technologies) but also any available correlation device (sunspots). Next, let there be a finite set of primitive, non-contingent goods, indexed by $n \in \{1, ..., N\}$. The corresponding price vector is denoted by $p \in \mathbb{R}^N_+$. This, too, is a random variable, whose distribution will be determined in equilibrium.[10]

To introduce rational inattention, we must first take a stand on what objects an agent can possibly collect signals on. We denote the collection of such objects by z and refer to it is that agent's *cognition state*.[11] In Section 5.3, we discuss how to accommodate a flexible specification of z, which allows us to capture the possibility that agents may pay attention not only to fundamentals, sunspots, and prices but also to taxes, macroeconomic statistics, blogs, or twits. For our main analysis, however, we confine attention to $z = (\theta, p)$. We then introduce rational inattention by (i) restricting agents to condition their behavior on a noisy signal of z and (ii) letting them design this signal subject to a cost.

[9] See Online Appendix E for an example in which the use of a sunspot to add noise in prices helps economize attention costs. In this example, the form of invariance needed for efficiency (Theorem 1) breaks in a subtle way, despite the *apparent* use of mutual information—and so does the form of monotonicity invoked in our result about fundamental equilibria (Theorem 2).

[10] The assumptions that the state space Θ is finite (made here) and that the signal space Ω is also finite (made shortly) are purely technical: they allow us to work with probability distributions instead of more complicated measures.

[11] It is straightforward to let z differ across j so as to capture heterogeneity in what agents can possibly learn or think about. Alternatively, such heterogeneity can be embedded in the attention cost functional below, while preserving the symmetry in z.

It may appear that z double counts θ, because in equilibrium p will have to be a function of θ. But this is an *equilibrium* property and, in the decision problem of the typical consumer or firm, θ and p appear as two conceptually distinct variables. In other words, the probability of rain and the price of umbrellas may loom differently in people's cognition, even if they end up being tightly connected in equilibrium. A similar point applies to the "messages" that replace prices in our upcoming notions of feasibility and efficiency.

We describe the stochastic properties of z as follows. We let π be a probability mass function over realizations of $z = (\theta, p)$ and require that it belongs in the set $\mathscr{P}$ defined by the compositions of π_θ, the prior about the state of nature, with arbitrary functions from Θ to $\mathbb{R}_+^N$:

$$\mathscr{P} \equiv \left\{\pi : \Theta \times \mathbb{R}_+^N \to [0,1] \text{ s.t. } \pi(\theta, p) = \pi_\theta(\theta)\,\mathbb{I}[f(\theta) = p], \text{ for some } f : \Theta \to \mathbb{R}_+^N\right\}. \tag{11}$$

This restriction is without loss of generality once we recognize that p is a function of θ in equilibrium and only serves a technical purpose: to ease the upcoming description of attention cost functionals. Along the same lines, we also define, for every $\pi \in \mathscr{P}$, the π-indexed set $Z_\pi \equiv \{(\theta, p) : \pi(\theta, p) > 0\} \subset \Theta \times \mathbb{R}_+^N$ and the π-indexed function $f_\pi : \theta \to \Theta \times \mathbb{R}_+^N$ with $f_\pi(\theta) = \{(\theta, p) : \pi(\theta, p) > 0\}$ for any $\theta \in \Theta$; the first gives the (necessarily finite) support of any $\pi \in \mathscr{P}$, the second reverse-engineers the mapping from (11) associated with π.

3.2 Inattentive Consumers

There is a unit-measure continuum of households, split into a finite number of types $j \in \{1, \dots, J\}$. The mass of type j is given by $\mu^j \in (0,1)$, with $\sum_{j=1}^J \mu^j = 1$. Each consumer, irrespective of type, can consume goods within a set $\mathscr{X} \subset \mathbb{R}^N$. For technical reasons, we assume $\mathscr{X}$ is a closed rectangle in $\mathbb{R}_+^N$, or $\mathscr{X} = \prod_{n=1}^N [0, x_n^{\max}]$ for a vector $(x_n^{\max})_{n=1}^N \gg 0$; but we make these bounds arbitrarily large so as to make their value irrelevant.[12]

Consumers have two income sources. The first is a state-dependent endowment, denoted by $e^j(\theta)$ for some fixed $e^j : \Theta \to \mathbb{R}_+^N$. The second is a fraction a^j of aggregate firm profits, with $\sum_{j=1}^J a^j = 1$. The profits are denoted by $\Pi(\theta)$ and, like prices, will be determined as part of the equilibrium.

Consumers choose two variables: a signal about z, and a plan for how much to consume of each good for each signal realization. Denote the realization of the signal by ω and let the *signal space* (i.e., the set of possible ω values) be a fixed, finite set Ω. The choice of a consumption plan is then represented by the choice of a mapping $x : \Omega \to \mathscr{X}$. The choice of a signal, on the other hand, is represented by the choice of a probability distribution over Ω for every $z \in Z_\pi$, where Z_π was previously defined as the set of all z which occur with positive probability under prior π. When there is no confusion, we refer to this collection of "likelihood distributions" $(\phi(\cdot|z))_{z \in Z_\pi}$ with the short-hand notation ϕ.

As long as $\pi \in \mathscr{P}$, ϕ is necessarily an element of $\Phi \equiv \Delta(\Omega)^{|\Theta|}$. We thus restrict attention to pairs (ϕ, π) belonging to the set $\Phi \times \mathscr{P}$, refer to any such pair as an *information structure*, and assign it a cost $C^j[\phi, \pi]$

[12] These bounds help guarantee existence of a solution to the consumer's problem. But since endowments are finite, the production set is compact, and $\mu^j > 0$ for all j, it is straightforward to let, for all n, $x_n^{\max}$ be greater than $1/\min_j \mu^j$ times the maximum conceivable amount of the respective good, so no consumer type can reach the bound in a feasible allocation.

measured in utils, for some function $C^j : \Phi \times \mathscr{P} \to \mathbb{R}_+$. This defines the cost of attention, or cognition.[13,14]

We then represent the consumer's preferences over the joint of consumption plans and attention choices by the expected utility of consumption net of this cost:

$$\mathscr{U}(x,\phi;\pi) \equiv \sum_{\omega\in\Omega, z\in Z_\pi} u^j(x(\omega),\theta)\,\phi(\omega \mid z)\,\pi(z) - C^j[\phi,\pi],$$

where $u^j : \mathscr{X} \times \Theta \to \mathbb{R}$ is the consumer's Bernoulli utility function over goods. We assume that, conditional on each state θ, this Bernoulli utility function is (i) continuous in x and (ii) represents *weakly monotone* preferences over goods, or satisfies $u^j(x',\theta) > u^j(x,\theta)$ for each $x, x' \in \mathscr{X}$ such that $x' \gg x$ and each $\theta \in \Theta$ (where we use $\gg$ to denote strict inequality for each element).

We close the consumer's problem by assuming complete markets over not only θ but also ω. This is analogous to our assumption of quasi-linear preferences in Section 2, insofar as it allows the consumer to equalize marginal utility over realizations of the signal ω. As mentioned in the Introduction, this makes sure that pecuniary externalities are muted. But it does not eliminate the mistakes due to rational inattention, nor does it equalize marginal utility across consumer types. For example, "less sophisticated" consumers (i.e., those with higher attention costs) will naturally fare worse in the absence of compensating transfers.

Taken together, the consumer's behavior can be expressed as the solution to the following program:[15]

$$\begin{aligned}
\max_{x,\phi} \quad & \sum_{\omega\in\Omega, z\in Z_\pi} u^j(x(\omega),\theta)\,\phi(\omega \mid z)\,\pi(z) - C^j[\phi,\pi] \\
\text{s.t.} \quad & \sum_{\omega\in\Omega, z\in Z_\pi} \Big(p\cdot x(\omega) - p\cdot e^j(\theta) - a^j\Pi(\theta)\Big)\,\phi(\omega \mid z)\,\pi(z) \le 0 \\
& x:\Omega \to \mathscr{X}; \quad \phi(\cdot \mid z) \in \Delta(\Omega),\ \forall\, z \in Z_\pi
\end{aligned} \tag{12}$$

This formulation accommodates a wide range of attention costs, like the ubiquitous posterior-separable class (as studied by Caplin and Dean, 2015; Denti, 2018). But for now, we do not need to impose posterior separability or any other substantial restriction. We only assume that C^j is continuous in its first argument conditional on the second, to help guarantee that the decision problem stated above is well posed.[16]

Finally, note that we have let attention costs enter linearly in preferences. This is standard in the literature but is not essential for our main result: we could interpret $c^j = C^j[\phi,\pi]$ as a "bad" and let it enter u^j in a non-linear way alongside the other goods. We follow such an approach in the firm's problem below.

[13] Let $K \equiv |\Theta|$. Our formulation presumes that C^j knows how to "read" from any pair $(\phi,\pi) \in \Phi \times \mathscr{P}$ which of the K components of ϕ corresponds to which of the K values in the support of z implied by π. To make sure this is the case, we proceed as follows. First, we order Θ in an arbitrary way, so that we can write $\Theta = \{\theta_1, ..., \theta_K\}$. Next, for any $\pi \in \mathscr{P}$, we order Z_π in a way that preserves the order over Θ, that is, we let $Z_\pi = \{z_1, ..., z_K\}$ with $z_k \equiv (\theta_k, f_\pi(\theta_k))$. We also write the typical element of Φ as $\phi = (\phi_1, ..., \phi_K)$, with $\phi_k \in \Delta(\Omega)$ for all $k \in \{1, ..., K\}$. And finally, we interpret ϕ_k as the distribution of ω conditional on z taking the value $z_k \equiv f_\pi(\theta_k)$. In simple words, for any π and ϕ, we always associate the k-th component of ϕ with the k-th value in the support of z implied by π.

[14] Our formulation leaves the "degenerate" likelihoods, $\phi(\cdot|z)$ for any $z \notin Z_\pi$, undefined. This is not a problem, because these likelihoods never appear in the decision problems we write below. But one may normalize $\phi(\cdot \mid z) = \frac{1}{|\Omega|}$ for all $z \notin Z_\pi$.

[15] In the sums appearing in the consumer's problem (12), z is of course the same variable as the pair (θ, p). The same applies to the firm's problem (13) below.

[16] Since our probability distributions are finite, it is sufficient to think of continuity in the vector space $[0,1]^{|\Omega|\times|\Theta|}$. If the signal space were continuous, we would need to define continuity with respect to the appropriate weak topology.

3.3 Inattentive Firms

There is a unit-measure continuum of identical firms.[17] Firms, like consumers, choose two objects. The first is a signal parameterized again by a collection $\phi = (\phi(\cdot \mid z))_{z \in Z_\pi} \in \Phi \equiv \Delta(\Omega)^{|\Theta|}$. The second is a ω-contingent production plan, $y : \Omega \to \mathscr{Y}$, where $\mathscr{Y} \subset \mathbb{R}^N$ is non-empty, is closed, and contains the zero vector (no production). Firms maximize expected profits subject to a technological constraint, which embeds attention costs. Formally, an output vector $y \in \mathscr{Y}$ is feasible for the firm in state θ if and only if it satisfies $H\left(y, c, \theta\right) \leq 0$, where $H : \mathbb{R}^{N+1} \times \Theta \to \mathbb{R}$ is continuous and increasing in its first $N+1$ elements and c measures the cost of cognition. The latter is specified as $c = C^F[\phi, \pi]$, where $C^F : \Phi \times \mathscr{P} \to \mathbb{R}$ is defined analogously to the consumers' cost function. Finally, we introduce normalizations such that the firm can always "shut down." That is, we let $C^F(\phi, \pi) = 0$ whenever ϕ is such that $\phi(\omega|z) = \frac{1}{|\Omega|}$ for all ω and all $z \in Z_\pi$ and assume $H(0, 0, \theta) < 0$.

Putting everything together, firm behavior is summarized in the following program:

$$\begin{aligned} \max_{y,\phi} \; & \sum_{\omega \in \Omega, z \in Z_\pi} (p \cdot y(\omega))\, \phi(\omega \mid z)\, \pi(z) \\ \text{s.t. } & H\left(y(\omega), C^F[\phi, \pi], \theta\right) \leq 0, \forall\, (\omega, \theta) : \phi(\omega \mid f_\pi(\theta)) > 0 \\ & y : \Omega \to \mathscr{Y}; \quad \phi(\cdot \mid z) \in \Delta(\Omega),\ \forall\, z \in Z_\pi \end{aligned} \tag{13}$$

This formulation treats firm cognition as a non-tradable, "in-house" production activity that diverts resources from production (insofar as it increases C^F and, thereby, reduces H), but also lets production plans respond more efficiently to demand and supply shocks (insofar as it allows ω to be more informative of z). Note that this readily nests the scenario in which attention costs emerge as a linear penalty in profits by letting $H(y, c, \theta) = \tilde{H}(y + \nu c, \theta)$ for some function $\tilde{H}$ and some constant $\nu \in \mathbb{R}_+$.

3.4 Equilibrium

Throughout, we focus on equilibria in which strategies are symmetric *within* types. This is without serious loss of generality, because we can always partition types into sub-types with the opportunity to make different decisions. We thus define equilibrium as follows:

Definition 3. An equilibrium is a profile of consumption and production strategies, $\left((x^j)_{j=1}^J, y\right)$, attention choices, $\left((\phi^j)_{j=1}^J, \phi^F\right)$, a price function $P : \Theta \to \mathbb{R}_+^N$, and a prior $\pi \in \mathscr{P}$ such that

1. Consumers and firms optimize, respectively solving programs (12) and (13), taking as given π.

2. Markets clear, or for all $\theta \in \Theta$,

$$\sum_{j=1}^J \mu^j \overline{x}^j(\theta) = \sum_{j=1}^J \mu^j e^j(\theta) + \overline{y}(\theta) \tag{14}$$

where

$$\overline{x}^j(\theta) \equiv \sum_{\omega \in \Omega} x^j(\omega)\, \phi^j(\omega \mid f_\pi(\theta)) \quad \text{and} \quad \overline{y}(\theta) \equiv \sum_{\omega \in \Omega} y(\omega)\, \phi^F(\omega \mid f_\pi(\theta)) \tag{15}$$

[17] It would be straightforward to add types of firms as well, but we abstract from this for simplicity.

are, respectively, the average demand of type-j consumers and the average supply of firms in state θ.

3. Profits are rebated to consumers, or $\Pi(\theta) = p(\theta) \cdot \overline{y}(\theta)$ for all $\theta \in \Theta$.

4. The prior π about z is consistent with the equilibrium price functional, or $f_\pi(\theta) = (\theta, P(\theta))$ for all $\theta \in \Theta$.

The following two properties carry over from the equilibrium definition in the example (Definition 1). First, the prior about z, which enters each agent's cognitive cost, is an endogenous object required to be consistent with the equilibrium price functional—this is, again, where "RI meets REE." And second, because of a law of large numbers applied to idiosyncratic realizations of ω, all aggregate quantities and prices are functions of θ in equilibrium—but now θ may contain a multitude of fundamentals as well as sunspots.

Remark 3. While the distinction between states of nature and "true," non-contingent goods is immaterial in the standard Arrow-Debreu framework, it was necessary in our context in order to introduce rational inattention. Still, this distinction disappears once we reach Definition 3, thanks to the following two basic properties: aggregate demands/supplies and market clearing are defined at the good-by-state level; and attention choices are subsumed in aggregate demands/supplies. This anticipates our mapping from inattentive economies to "twin" attentive economies developed in Section 4.

Remark 4. The above discussion also invites the following intuition for what rational inattention "means" in more standard GE language: "confusing θ and θ'" is tantamount to submitting similar demands for "apples-in-θ" and "apples-in-θ'." From this perspective, an complete-markets economy with rational inattention may look like one with "endogenously incomplete markets," in the sense that agents choose in equilibrium to ignore certain contingencies that are not valuable to learn about relative to attention costs.

3.5 Feasibility and Efficiency

We now define efficiency. This is the same as Pareto optimality under an amended version of feasibility that takes into account the cognitive friction. As in the example, this involves the replacement of the price variable p in the cognition state of the agents' cognitive process with a message m. We require that this message, like the equilibrium price, be representable as a function $M : \Theta \to \mathbb{R}^N_+$. The consumer and firm problems are adjusted to re-define the cognition state as $z = (\theta, m)$, which has a prior $\pi \in \mathscr{P}$ which must now be consistent with the message function M. We can then define feasibility and efficiency as follows:

Definition 4. A feasible arrangement is a profile of consumption and production strategies, $\left((x^j)_{j=1}^J, y\right)$, attention choices, $\left((\phi^j)_{j=1}^J, \phi^F\right)$, a message rule $M : \Theta \to \mathbb{R}^N_+$, and a prior $\pi \in \mathscr{P}$ such that

1. Consumption and production are informationally feasible, or $x^j : \Omega \to \mathscr{X}$ for all j and $y : \Omega \to \mathscr{Y}$.

2. Consumption and production are technologically feasible, or (14)-(15) hold along with

$$H\left(y(\omega), C^F[\phi^F, \pi], \theta\right) \leq 0, \forall\, (\omega, \theta) : \phi^F(\omega \mid f_\pi(\theta)) > 0 \tag{16}$$

3. The prior π about z is consistent with the message rule, or $f_\pi(\theta) = (\theta, M(\theta))$ for all $\theta \in \Theta$.

Definition 5. An equilibrium is efficient if there does not exist a feasible alternative such that (i) all agents are weakly better off and (ii) a positive measure of agents is strictly better off.

As discussed in Section 2.2, replacing prices with messages allows our question about the "economy of knowledge" to be well-posed. Such messages are herein restricted to live in the same space as prices. This simplifies the exposition by making sure that z itself remains in the same space as we move back and forth between equilibria and planning alternatives. In Section 5.3, however, we discuss how to accommodate two related extensions: to enrich the messages sent by a planner, and to let endogenous objects other than prices, such as aggregate trades or taxes, enter the cognitive process as distinct elements of z.

3.6 Invariance and Monotonicity of Costs

In this subsection we define two properties of attention costs that we invoke in our subsequent results. The first is the form of invariance alluded to at the end of Section 2.6. The second is a form of monotonicity that, while distinct in principle, is also satisfied in the leading example of mutual-information costs.

We start by fixing language and notation. Recall that we refer to any pair (ϕ, π), describing a signal choice and a prior about z, as an *information structure*; and that z is a random variable with domain $\Theta \times \mathbb{R}^N_+$. Let $\mathscr{G} \equiv \{g : (\Theta \times \mathbb{R}^N_+) \to (\Theta \times \mathbb{R}^N_+)\}$ and, for any $g \in \mathscr{G}$, consider the transformation $z \mapsto \tilde{z} = g(z)$. Informally, we can say that such transformations either "relabel" states if they are bijective ("one-to-one") or "relabel and merge" states if they are surjective ("one-to-many"). And we define the corresponding transformations of information structures in the following way.

Definition 6 (Transformations of Information Structures)**.** Consider two information structures (π, ϕ) and $(\tilde{\pi}, \tilde{\phi})$ and a function $g \in \mathscr{G}$. We say that $(\tilde{\phi}, \tilde{\pi})$ is the transformation of (ϕ, π) under g if

$$\tilde{\pi}(z) = \sum_{z'} \pi(z')\mathbb{I}[g(z') = z] \qquad \forall z \in Z_\pi \tag{17}$$

$$\tilde{\phi}(\omega|z) = \frac{\sum_{z' \in Z_\pi} \phi(\omega|z')\pi(z')\mathbb{I}[g(z') = z]}{\tilde{\pi}(z)} \qquad \forall \omega \in \Omega, z \in Z_{\tilde{\pi}} \tag{18}$$

Such transformations amount to a "change of variables" in the following sense. Consider a random variable z and replace it with the random variable $\tilde{z} = g(z)$, for some $g \in \mathscr{G}$. If the former's distribution is $\pi \in \mathscr{P}$, then the latter's is $\tilde{\pi} \in \mathscr{P}$ constructed as in (17). Furthermore, if the agent had chosen a signal ϕ for the original random variable and wishes to preserve the informational content of that original signal with respect to the new random variable $\tilde{z}$, then the new signal $\tilde{\phi}$ would be constructed according to (18).

We next define a *sufficiency* relationship between information structures, recasting the familiar definition of a *sufficient statistic* in our language:

Definition 7. Consider two information structures (ϕ, π) and $(\tilde{\phi}, \tilde{\pi})$ such that $(\tilde{\phi}, \tilde{\pi})$ is the transformation of (ϕ, π) under g for some $g \in \mathscr{G}$. We say that $\tilde{\pi}$ is sufficient for π with respect to ϕ if $\phi(\omega \mid z) = \tilde{\phi}(\omega \mid g(z))$ for all ω and all z such that $\pi(z) > 0$.

This definition is equivalent to saying, in slightly different language, that the random variable $\tilde{z} = g(z)$ is a sufficient statistic for the random variable z with respect to the signal ω, or that the distribution of z conditional on $\tilde{z}$ does not depend on ω.[18] More informally, under the sufficient statistic condition, the signal ω does not contain any information about z other than the one it contains about $\tilde{z}$. See that this is trivial if g is bijective, or merely "relabels" states, and a meaningful restriction when g is surjective, or both "relabels and combines" states. In particular, if the original signal structure allowed one to learn the relative likelihood of two states $z \neq z'$ such that $g(z) = g(z')$, then the sufficiency property does not hold.

With these definitions in hand, we now state the *invariance* and *monotonicity* properties of interest, both defined with respect to an arbitrary subset G of the universe $\mathscr{G}$ of possible transformations:

Definition 8. Fix a set $G \subseteq \mathscr{G}$. Consider any function $g \in G$ and any two information structures (ϕ, π) and $(\tilde{\phi}, \tilde{\pi})$ such that $(\tilde{\phi}, \tilde{\pi})$ is the transformation of (ϕ, π) under g. A cost functional C is

1. invariant with respect to G if $C[\phi, \pi] = C[\tilde{\phi}, \tilde{\pi}]$ whenever $\tilde{\pi}$ is sufficient for π with respect to ϕ.

2. monotone with respect to G if $C[\phi, \pi] > C[\tilde{\phi}, \tilde{\pi}]$ whenever $\tilde{\pi}$ is not sufficient for π with respect to ϕ

Invariance requires that rescaling or relabeling the states does not affect attention costs, while merging particular states has no effect provided the signal did not originally distinguish between those states. Monotonicity requires additionally that combining states that were originally distinguished from one another strictly decreases costs. The first property will alone be the key for efficiency (Theorem 1). The second property will guarantee that efficient equilibria are sunspot-free and, in addition, that agents pay attention only to prices and not to one another's fundamentals (Proposition 2).

In Section 6.2 and Appendix B, we connect our invariance and monotonicity notions to related notions from the economic literature on flexible information acquisition (Caplin et al., 2020) and the statistics literature on information geometry (Amari, 2016). We also clarify that mutual information costs are invariant and monotone with respect to the entire $\mathscr{G}$. Below, we use a simple example to illustrate the particular invariance property that is relevant in the proof of Theorem 1. This amounts to allowing arbitrary transformations of prices (or their replacement with arbitrary messages) but ruling out transformations of the state of nature. In Section 5.3, on the other hand, we explain why "full" invariance must be invoked once we consider plausible expansions of our planner's powers or the definition of z.

Remark. In earlier drafts, we had combined our invariance and monotonicity notions in a single condition (Assumption 3 in Angeletos and Sastry, 2019), thus also blending Theorems 1 and 2 together. The separation of these notions and of their implications in the present draft mirrors that in Hébert and La'O (2020).

Illustration. A consumer lives in a world with two states of nature with equal prior probabilities: $\theta \in \Theta = \{A, B\}$, with $\pi_\theta(A) = \pi_\theta(B) = \frac{1}{2}$. There is a single price, p, taking values in $[0, \infty)$. The cognition state is

[18] This equivalence follows from an elementary calculation combining the transformation of Definition 6 with Bayes' rule. This calculation is contained in the proof of Lemma 4. Also, the reason that in this draft, unlike an earlier one, we opt to use the language of distributions (e.g., π and $\tilde{\pi}$) instead of that of random variables (respectively, z and $\tilde{z}$) is merely that this facilitates a cleaner connection to the decision-theoretic literature on flexible information acquisition. This connection is spelled out in Section 6.2.

Figure 1: Visualization of a Non-Invariant Cost Functional

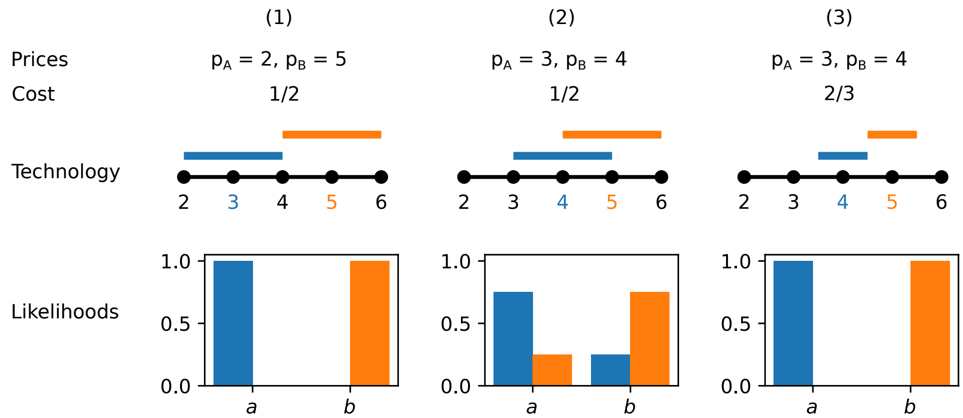

Notes: Blue corresponds with state A and orange with state B. The first row of diagrams visualizes the agent's signal "technology" and the second plots, in a two-color bar graph, their likelihood distributions.

$z = (\theta, p) \in \{A, B\} \times \mathbb{R}$. There are two possible price functionals, denoted by P and $\tilde{P}$. The first maps $\{A, B\}$ to $\{3, 5\}$, the second to $\{4, 5\}$. These mappings correspond to two possible distributions for z, denoted by π and $\tilde{\pi}$, with respective supports $Z_\pi = \{(A,3),(B,5)\}$ and $Z_{\tilde{\pi}} = \{(A,3),(B,5)\}$. The signal space is $\Omega = \{a, b\}$.

Our question is how costs do (or do not) change as we move from price functional P to price functional $\tilde{P}$, or equivalently from prior π and prior $\tilde{\pi}$. Note that this amounts to a transformation of the original z to a new $\tilde{z}$ by a function g that preserves the first element of z but changes the second. Namely, g maps $\{(A,3),(B,5)\}$ to $\{(A,4),(B,5)\}$. It is this kind of "partial" transformation of z that will be relevant for understanding the efficiency of the equilibrium below.

We first construct a cost that is neither monotone nor invariant and that captures the idea of having difficulty distinguishing nearby prices; this is like the cost assumed in equation (9) for one of our earlier examples of inefficiency. The agent has no way to directly learn θ; perhaps they do not even understand what "A" and "B" mean. But they can try to learn p. In particular, they can choose among a continuum of "experiments," indexed by $\eta \geq 0$.[19] Running experiment η costs $1/(\eta+1)$ and returns an outcome that is distributed uniformly on the interval $[p-\eta, p+\eta]$. The agent then applies the following (Bayesian) algorithm to map experiments to signals. If the outcome of the experiment is uniquely consistent with the price in A (respectively, the price in B), the agent assigns probability 1 to $\omega = a$ (respectively, $\omega = b$). Otherwise, the agent attaches probability $\frac{1}{2}$ to each value of ω. See that the choice of an experiment corresponds to the choice of a ϕ, or the distribution for ω conditional on z, and that we can define $C[\phi, \pi]$ in this context as the lowest cost experiment generating a given ϕ.

Imagine that, when the price mapping is P, or $(p_A, p_B) = (3,5)$, the consumer picks $\eta = 1$ at cost $C = 1/2$ and perfectly distinguishes the two states at the lowest possible cost. This is demonstrated in Column 1 of Figure 1, which visualizes the signal technology and plots the likelihood distributions. Next, see that when the price mapping is $\tilde{P}$, or $(p_A, p_B) = (3,4)$, the same experiment of $\eta = 1$ yields a different and strictly less

[19]We use "experiment" informally here, and not in the precise sense of Blackwell (1951).

informative signal (Column 2). Finally, see that the cheapest method of preserving the informational content of the original signal under the new price functional requires a lower η and therefore a strictly higher cost (Column 3). This cost functional is therefore neither monotone nor invariant in the sense of Definition 8.

On the other hand, any cost functional which depended only on the signal distributions visualized in the fourth row of Figure 1, without reference to the "names" or values of the states, would take an equal value in columns (1) and (3) and therefore be invariant with respect to this transformation. The mutual information of ω and (θ, p), or any increasing function thereof, is one such cost functional.[20] This connects with our observation in Section 2.6 that mutual information costs were somehow "invariant" to changes in the stochastic properties of prices.

4 Main Result: Efficiency of Inattentive Economies

We now state our main result:

Theorem 1 (First Welfare Theorem)**.** *Let $G^p \subset \mathcal{G}$ be the class of functions $(\Theta \times \mathbb{R}_+^N) \to (\Theta \times \mathbb{R}_+^N)$ that can be written as $g(\theta, p) = (\theta, h(p))$ for some $h : \mathbb{R}_+^N \to \mathbb{R}_+^N$; and assume that all consumers and firms have attention costs that are invariant with respect to G^p. Then, all equilibria are efficient.*

The stated assumption ("invariance within G^p") specializes the properties of Definition 8 in two dimensions. First, it requires only invariance, not monotonicity. Second, it requires invariance in only one particular dimension: that of manipulating the informational content of prices or replacing them with other messages. In the rest of this section, we sketch the proof of Theorem 1 in three steps, which help clarify both how our result leverages the standard argument underpinning the First Welfare Theorem and why exactly invariance is the key for extending this theorem to inattentive economies.[21]

4.1 Proof

While there are multiple proof strategies for Theorem 1, we will pursue one that most clearly maps to textbook treatments (e.g., Chapter 16 of Mas-Colell et al., 1995). Here, we describe the main logic up to a few supporting results proved in the Appendix.[22]

Step One: An Equivalent Economy

We first define a notion of preferences and technology at the level of type-specific aggregate quantities, $\overline{x}^j(\theta)$ and $\overline{y}(\theta)$. This representation subsumes the attention choice to the preferences and technology of a friction-

[20] It is simple to calculate that the mutual information in each case is 0.69, 0.13, and 0.69.

[21] The statement and proof of the standard First Welfare Theorem requires neither the existence of an equilibrium nor its comparison to the solution of a planning problem. Instead, it presumes the existence of an equilibrium and proceeds to rule out Pareto improvements. The same applies to Theorem 1 here and its proof below. The study of equilibrium existence and Pareto optima is therefore *not* needed for our main result and is postponed to Section 5.1.

[22] We followed a different proof strategy in an earlier draft (Angeletos and Sastry, 2019), by comparing "state-tracking economies" (in which agents received signals only of θ) and "price tracking economies" (in which agents received signals of both θ and p, as in the current draft). The efficiency of the former class of economies, stated in Theorem 1 of our earlier draft, follows from Proposition 3 below. The efficiency of the latter class, stated as Theorem 4 of the earlier draft, is the content of Theorem 1 here.

less “twin economy,” and clarifies what is and is not standard about our problem.[23]

Let $\overline{\boldsymbol{x}} \in \mathscr{X}^{|\Theta|}$ be a shorthand for $(\overline{x}(\theta))_{\theta\in\Theta}$. We define the following program for each consumer type j that solves for an optimal consumption plan and an optimal signal structure subject the constraint that the average consumption across realizations of signals should not exceed an available basket $\overline{\boldsymbol{x}}$; and returns type j's expected utility net of attention costs:

$$\begin{aligned}\overline{u}^j(\overline{\boldsymbol{x}},\pi) \equiv \max_{x,\phi} \sum_{\omega,\theta} u^j(x(\omega),\theta)\,\phi(\omega \mid f_\pi(\theta))\,\pi_\theta(\theta) - C^j[\phi,\pi] \\ \text{s.t.} \sum_{\omega} x(\omega)\,\phi(\omega\mid f_\pi(\theta)) \le \overline{x}(\theta),\ \forall\theta\in\Theta \\ x:\Omega\to\mathscr{X};\quad \phi(\cdot\mid f_\pi(\theta))\in\Delta(\Omega),\ \forall\theta\in\Theta\end{aligned} \tag{19}$$

See that, under the maintained assumptions that (i) Ω is finite, (ii) $\mathscr{X}$ is compact, and (iii) $u^j(\cdot,\theta)$ and $C^j[\cdot,\pi]$ are continuous, Weierstrauss' theorem guarantees that this program has a solution and therefore $\overline{u}^j((\overline{x}(\theta))_{\theta\in\Theta},\pi)$ is well-defined. We refer to the value function $\overline{u}^j(\overline{\boldsymbol{x}},\pi)$ as a *reduced preference* for the basket of state-contingent commodities $\overline{\boldsymbol{x}}$, which depends on the prior π as an auxiliary parameter.

The consumer program (12), owing to complete markets over ω, can now be re-written as a more standard consumer optimization with the altered preferences defined above:

$$\begin{aligned}\max_{\overline{\mathbf{x}}}\ \overline{u}^j((\overline{x}(\theta))_{\theta\in\Theta},\pi) \\ \text{s.t.} \sum_{\theta}(P(\theta)\cdot\overline{x}(\theta) - P(\theta)\cdot e^j(\theta) - a^j\Pi(\theta))\,\pi_\theta(\theta)\le 0\end{aligned} \tag{20}$$

Two specific features of this representation are important. First reduced preferences inherit monotonicity in the goods space from the primitive utility function u^j.[24] This is formalized as Lemma 5 stated and proven in the Appendix, using a simple constructive argument in Program 19. Second, prices affect preferences directly via the prior π (and, more fundamentally, their influence on the learning problem in Program 19). This feature embodies the cognitive externality on the consumer side of the economy

For firms, we analogously define, for each prior π, the following *reduced production set*:

$$\begin{aligned}\overline{\boldsymbol{F}}(\pi) \equiv \Big\{\overline{\boldsymbol{y}} = (\overline{y}(\theta))_{\theta\in\Theta}\in\mathscr{Y}^{|\Theta|} : \exists\,(y,\phi)\ \text{s.t.:}\ & \sum_{\omega} y(\omega)\,\phi(\omega\mid f_\pi(\theta)) \le \overline{y}(\theta),\forall\theta\in\Theta \\ & H\big(y(\omega),C^F[\phi,\pi],\theta\big)\le 0, \forall\,(\omega,\theta):\phi(\omega\mid f_\pi(\theta))>0 \\ & y:\Omega\to\mathscr{Y};\quad \phi(\cdot\mid f_\pi(\theta))\in\Delta(\Omega),\ \forall\,\theta\in\Theta\Big\}\end{aligned} \tag{21}$$

This set describes the combinations, across goods and states of nature, of the *aggregate* production levels that are both technologically and cognitively feasible: that is, they can be disaggregated with at least one

[23]Close readers may note that, while this twin economy is well-defined under our simplifying assumptions like a finite state and signal space and continuous payoffs, the offered representation could be bypassed in our formal argument. We follow this strategy because it provides the sharpest possible understanding of our main result and it is also useful for proving the additional results in Section 5.

[24]The use of such monotonicity in the proof of Proposition 3 could be relaxed for an appropriate form of local non-satiation, as in our earlier draft (Angeletos and Sastry, 2019).

attention strategy $\phi = (\phi(\cdot \mid z))_{z \in Z_\pi}$ and a production plan $y : \Omega \to \mathscr{Y}$. The representation further extends the metaphor from Section 3.3 about attention as an "in-house" productive activity.

The firms' profit maximization problem then reduces to the following program:

$$\max_{\overline{\mathbf{y}}} \sum (P(\theta) \cdot \overline{y}(\theta))\, \pi(\theta) \quad \text{s.t. } (\overline{y}(\theta))_{\theta \in \Theta} \in \overline{\boldsymbol{F}}(\pi) \tag{22}$$

This is like the profit-maximization problem of a standard, attentive firm, except that the production set is allowed to depended on the prior. This embodies the cognitive externality on the firm side of the economy.

Step Two: Efficiency Ignoring Messages

We now state a restricted version of the First Welfare Theorem that removes the planner's flexibility to replace prices with other messages, but also holds without *any* specification of attention costs:

Proposition 3. *For any equilibrium with price functional $p = P(\theta)$, there does not exist a Pareto dominating allocation that is feasible with message $m = P(\theta)$*

By restricting the messages to replicate the equilibrium prices, we fix π in the reduced preferences and production sets. The rest of the proof then reads much like the familiar, textbook argument for why competitive equilibria can not be Pareto dominated (e.g., in section 16.C of Mas-Colell et al., 1995). In this respect, Proposition 3 is quite trivial. But it contains an important lesson: barring the elimination of inattention itself, the *only* way to possibly improve upon the equilibrium of an inattentive but otherwise frictionless (i.e., competitive and complete-markets) economy is to manipulate the informational content of prices. This clarifies some of the confusion in the literature about the channels via which inattention may or may not open the door to inefficiency; it goes to the heart of Hayek's argument; and it paves the way to the next, and last, step of the proof of our main result.

Step Three: Using Invariance to Prove Messages Cannot Help

The translation from Proposition 3 to Theorem 1 uses informational invariance. We first establish the following result: if attention costs are invariant to transformations of the price component of the economic state, then preferences are production sets do not depend directly on the stochastic properties of prices. By the same token, the cognitive externality, or the dependence of preferences and technologies on the *endogenous* part of the prior π, is muted.

Lemma 1. *When consumers' attention costs are invariant with respect to G^p, there exist functions $\tilde{u}^j$ such that the reduced preferences satisfy $\overline{u}^j(\overline{\boldsymbol{x}}, \pi) = \hat{u}^j(\overline{\boldsymbol{x}}, \pi_\theta)$ for all $\overline{\boldsymbol{x}}$ and all π. Similarly, when firms' attention costs are invariant with respect to G^p, there exists a function $\hat{\boldsymbol{F}}$ such that the reduced production set satisfies $\overline{\boldsymbol{F}}(\pi) = \hat{\boldsymbol{F}}(\pi_\theta)$ for all π.*

The proof of this Lemma is quite simple: if costs would be the same for any transformation of the prices in z, then the distribution of p conditional on θ cannot matter for preferences or production sets. The proof of Theorem 1 is then completed by using this Lemma to establish that there cannot exist a Pareto dominating allocation even with the use of *arbitrary* messages. Imagine there were. Lemma 1 implies that each consumer's payoff as well as the feasibility constraints would be identical under the message $m = P(\theta)$. Therefore, the allocation must also be implementable with that message; but Proposition 3 implies that such an allocation cannot exist.

4.2 Remarks

We end this section by commenting on the role played by two key assumptions—the inclusion of p in z and the insurance over ω—and on whether invariance is necessary in addition to being sufficient for efficiency.

Restricting to signals of θ instead of z. Throughout we have allowed agents to pay attention to, or learn about, both the exogenous state of nature and the endogenous prices (or any messages used in their place). What if instead we had restricted agents to receive signals only about the state of nature, that is, what if we had specified attention costs as a function of the joint distribution of ω and θ only?

This restriction is inconsistent with our motivation, the decision-theoretic literature we relate to, and a large behavioral literature, all of which emphasize cognition about prices or other variables that may be exogenous to the individual decision maker but certainly not to the economy as a whole. But it has been common in applications (e.g., Angeletos and La'O, 2018; Colombo et al., 2014; Maćkowiak and Wiederholt, 2015; Tirole, 2015). Economies satisfying this restriction can be nested in our analysis by noting that the associated cost functionals are trivially members of G^p, the class of costs invariant to transforming the price in the cognition state, since the price does not appear in them.[25] Thus any claims about efficiency or inefficiency in such prior works have abstracted from the kind of cognitive externality we have emphasized here, and therefore did not identify informational invariance in any form as a key condition for efficiency.

Insurance over ω. Our proof leveraged heavily the assumption of complete markets over the idiosyncratic realizations of the signal in defining and using the "reduced" preferences. Without complete markets over ω, or over any other idiosyncratic uncertainty, the familiar argument applies that pecuniary externalities do *not* net out and a social planner may wish to manipulate prices in order to redistribute wealth across realizations of ω, or to mimic the missing transfers. That is, equilibrium may fail to be constrained inefficient in the sense of Geanakoplos and Polemarchakis (1986), wherein the social planner faces a comparable restriction on insurance as the market. [26]

[25] In an earlier draft we referred to such economies as "state-tracking" economies; see Section 5 of Angeletos and Sastry (2019).

[26] In Online Appendix F, we spell out how exactly this logic can be applied to our context. Markets are complete over θ but incomplete over ω. Budgets nevertheless clear for all realizations of ω thanks to the existence of an "adjustment good" like "money" in the example of Section 2. Constrained efficiency is defined similarly to Definitions 2 and 5, plus the new requirement that the planner cannot not make transfers directly contingent on ω (as in Geanakoplos and Polemarchakis, 1986). We then show that constrained efficiency is essentially impossible *unless* the adjustment good has constant marginal utility across realizations of ω, like in the example of Section 2. The quasilinear assumption in that example equates the marginal value of wealth across realizations of ω *in spite* of the lack of complete markets in this dimension, and hence is a substitute for complete markets in our main arguments.

From this perspective, efficiency seems unlikely *even if* our invariance condition holds. But a policy-maker may be excused if she can't tell in which direction prices must be manipulated in order to mimic the missing insurance over ω. Furthermore, if agents make the same choices repeatedly and the noise is independent across rounds of choices, the noise may wash out over their life cycle. Loosely speaking, we think of this as a situation in which the assumed law of large numbers over ω applies *within* agents, substituting for insurance *across* agents, and that such an assumption may be more plausible than true complete markets.

Markets vs games. Complete, competitive markets can be understood as a special class of games, in which players are infinitesimal and payoff externalities are muted on equilibrium. This suggests that our main insight, regarding the role of invariance for efficiency, may extend to such games, as indeed established by Hébert and La'O (2020). Furthermore, because our state of nature is allowed to contain correlation devices ("sunspots"), Theorem 1 also allows for rich, non-fundamental correlation in the agents' signals and, thereby, for non-fundamental volatility in aggregate allocations despite efficiency. We return to this issue 5.2. But it is also worth clarifying the following, subtler point. Our equilibrium and efficiency concepts specify the agents' prior about z and their associated attention choices on equilibrium but not off equilibrium. This is innocuous here because agents are infinitesimal and off-equilibrium beliefs are immaterial. But this is not true in settings with large players, where the threat not to pay attention to something off equilibrium could influence what happens on equilibrium (see, e.g., Ravid, 2020). The extension of our paper's logic to games with non-infinitesimal players remains an open question.

Sufficiency versus necessity. Theorem 1 establishes that invariance is sufficient for efficiency but not that it is necessary. It is of course easy to envision cost functionals that violate invariance but support efficient equilibria. A trivial example is one in which attention costs happen to be zero when $z = (\theta, P^*(\theta))$ and P^* is an equilibrium price function in the underlying *attentive* economy, and positive and sufficiently large whenever $z = (\theta, g(\theta))$ for any $g \neq P^*$.

Such situations, however, seem contrived. Indeed, the following construction shows how "small" departures from invariance lead to inefficiency. Suppose that C satisfies our invariance condition, pick an equilibrium, and let P be the associated price functional. Next, consider a change in C that leaves attention costs unchanged for *all* signal choices as long as $z = (\theta, P(\theta))$ but strictly reduces them by $\epsilon > 0$ whenever $z = (\theta, g(\theta))$ for $g \neq P$. Clearly, the original equilibrium remains an equilibrium but ceases to be efficient, no matter how small ϵ is.

This logic, while well short of a formal converse to Theorem 1, invites us to think of invariance and efficiency as "nearly" equal. Hébert and La'O (2020) corroborate this point by showing that, at least in their setting, invariance is not only sufficient but also generically necessary for efficiency (where "generically" is relative to utilities and attention costs).

5 Additional Results

In this section, we provide four additional results that further characterize efficient inattentive economies. The first provides conditions for equilibrium existence. The second is our version of the Second Welfare Theorem. The third investigates when equilibria are "fundamental" and "price tracking." The fourth establishes efficiency under a more "powerful" planning concept. These results are useful but not strictly needed for our perspectives on rational inattention and Hayek's (1945) economy of knowledge ; readers interested on this discussion can jump to Section 6.

5.1 Equilibrium Existence and the Second Welfare Theorem

Here we provide sufficient conditions for equilibria to exist and for Pareto optima to be implementable as competitive equilibria with transfers.[27] To this goal, we leverage heavily on the twin-economy representation developed in Section 4.1. In particular, we start by assuming that the twin economy satisfies appropriate convexity and continuity conditions; we next use this assumption along with our invariance condition to obtain our existence result and our version of the Second Welfare Theorem; and we finally discuss how the assumed convexity and continuity of the twin economy can be derived from first principles.

The aforementioned convexity and continuity conditions are defined below.

Definition 9. Reduced preferences are

1. *convex* if, for every π, for every pair $\overline{\boldsymbol{x}}, \overline{\boldsymbol{x}}' \in \mathscr{X}^{|\Theta|}$ such that $\overline{u}^j(\overline{\boldsymbol{x}}, \pi) < \overline{u}^j(\overline{\boldsymbol{x}}', \pi)$ and every $\alpha \in (0,1)$, we have $\overline{u}^j(\alpha\overline{\boldsymbol{x}} + (1-\alpha)\overline{\boldsymbol{x}}, \pi) > \alpha\overline{u}^j(\overline{\boldsymbol{x}}, \pi) + (1-\alpha)\overline{u}^j(\overline{\boldsymbol{x}}', \pi)$

2. *continuous* if, for every π, $\overline{u}^j(\overline{\boldsymbol{x}}, \pi)$ is continuous for every π.

Definition 10. Reduced production sets are

1. *convex* if, for every π, every pair $\overline{\boldsymbol{y}}, \overline{\boldsymbol{y}}' \in \overline{\boldsymbol{F}}(\pi)$, and every $\alpha \in (0,1)$, $\alpha\overline{\boldsymbol{y}} + (1-\alpha)\overline{\boldsymbol{y}}' \in \overline{\boldsymbol{F}}(\pi)$.

2. *closed* if, for every π and every convergent sequence $\{\overline{\boldsymbol{y}}_k\}_{k=1}^{\infty}$ with $\overline{\boldsymbol{y}}_k \in \overline{\boldsymbol{F}}(\pi)$ for all k, $\lim_{k\to\infty}\overline{\boldsymbol{y}}_k \in \overline{\boldsymbol{F}}(\pi)$.

We then have the following two results:

Proposition 4 (Equilibrium Existence)**.** *Assume that attention costs are invariant with respect to G^p; that reduced preferences are convex and continuous; and that reduced production sets are convex and closed. Then, an equilibrium exists.*

Proposition 5 (Second Welfare Theorem)**.** *Under the assumptions stated above, any Pareto optimum can be implemented as an equilibrium with transfers.*

[27] An equilibrium with transfers in our setting is a direct extension of Definition 3 in which each consumer has a state-dependent wealth $p(\theta)e^j(\theta) + T^j(\theta)$, where $T^j(\theta) \in \mathbb{R}$ is a transfer that nets out across agents, that is, $\sum_{j=1}^{J} T^j(\theta) = 0$ for all $\theta \in \Theta$.

The proof of these results, in the Appendix, replicate the arguments in Arrow and Debreu (1954) and Debreu (1954), respectively. The assumed properties for the *reduced* preferences and technologies play the familiar role but do not suffice in our context. They must be combined with our invariance condition in order to rule out dependence of these preferences and technologies on prices.[28]

Our version of the Second Welfare Theorem may be of interest for researchers looking to establish comparative statics in inattentive economies, which may be considerably easier to derive in the planner's problem than in the corresponding equilibrium fixed-point.[29] Our existence result, on the other hand, limits attention to efficient economies. But it verifies that our framework passes the usual "sanity test" and that our main result is not vacuous. Its proof strategy also complements the intuition from Section 2 about how prices can play their usual market-clearing role even if both sides of the market are inattentive to prices.

By mapping the inattentive economy to a twin attentive economy, and by equating existence in the former to existence in the latter, we make clear that rational inattention does not pose any *new* difficulty for existence than that familiar from standard general equilibrium theory, namely, the possibility of discontinuity in aggregate excess demands due to non-convexities in preferences and technologies. In fact, one may conjecture that, starting from an attentive but non-convex economy, the introduction of rational inattention may *aid* equilibrium existence—and thereby the Second Welfare Theorem, too—by smoothing out the associated discontinuities in aggregate excess demands.

To check this intuition and complete the analysis of this section, we now return to the question of how the assumptions on the reduced preferences and technologies invoked in Propositions 4 and 5 translate in terms of primitives. Lemma 2, stated below and proved in Online Appendix D, provides an answer based on the following logic. Continuity and closedness follow from the appropriate notions of continuity on the primitives and application of Berge's Theorem. Convexity, on the other hand, is related to the question of whether agents can *costlessly* randomize over consumption and production plans. To use language familiar from general equilibrium theory, the key question is whether the optimal design of a noisy signal can subsume the use of lotteries over bundles of goods. The Lemma provides sufficient conditions for this to be true.

Lemma 2. Suppose that the following properties hold:

1. the signal space is given by $\Omega = [0,1]$.

2. technology can be written as $H(y,c,\theta) = \tilde{H}(y+c\cdot v,\theta)$, for a vector $v \in \mathbb{R}_+^N$ and a function $\tilde{H}: \mathbb{R}^N \times \Theta \to \mathbb{R}$.

3. attention costs are posterior separable (in the sense of Caplin and Dean, 2015).

Then, reduced preferences are convex and continuous and reduced production sets are convex and closed.

[28]Naturally, invariance is as essential for our version of the Second Welfare Theorem as it is for our version of the First Welfare Theorem. It may be possible, thought, to prove existence without invariance (and without efficiency) by appropriately adapting techniques from the general equilibrium literature on analysis with price-dependent preferences. See, for instance, Section III in Sonnenschein (2017) for an overview of these techniques.

[29]For example, one may use this approach to investigate which properties of the Neoclassical Growth Model are robust to the introduction of rational inattention.

Property 1 is technical, allowing enough "richness" in the signal space to accommodate different randomizations at different points of primitive non-convexities. Property 2 lets the firms' attention costs enter as a linear penalty on its outputs or inputs. This assumption is not trivial but easy to motivate: it nests the special case in which attention is "paid for" in the unit of a specific good (the vector v has a single non-zero element) and therefore enters as an additive penalty in profits denominated in the price of that good. As noted earlier, this is the case invariably assumed in macroeconomic applications. Furthermore, this assumption is the mirror in the production side to a simplifying assumption already made in the consumer side of our environment, namely that attention costs subtract linearly from the expected utility of goods. Together, these assumptions allow us to accommodate arbitrary non-convexity in the *primitive* preferences and technologies, and to derive the needed convexity for the *reduced* preferences and technologies by invoking an appropriate convexity with respect to attention choices. Property 3 completes the picture by showing such convexity is implied by assuming that attention costs are posterior separable. This nests mutual information along with virtually any other specification considered in the related, decision-theoretic literature.

5.2 Fundamental Equilibria and Optimal Attention

The examples of inefficiency given in Section 2.5 emphasized a particular failure of invariance due to scale dependence. But invariance may fail also in terms of transformations that add or remove contingencies. To illustrate, revisit the simple economy of Section 2; expand the state of nature to include not only the coconut endowment ξ but also a sunspot v (so $\theta = (\xi, v)$ in the notation of our general model); and suppose that attention costs depend on the mutual information of ω with p *alone*, as opposed to that of ω with $z = (\theta, p)$. Then, the original equilibrium, in which the price p and the signal ω were uncorrelated with v, remains an equilibrium. But there are now other, Pareto superior, equilibria that let the price and agents' signals be correlated with v. This is because, under the aforementioned assumption about attention costs, the addition of "noise" in the price allows agents to pay less attention to prices and, nevertheless, make consumption choices that are better aligned with the underlying fundamental.[30]

This example highlights the following subtlety: mutual-information costs satisfy invariance and guarantee efficiency only if they are "holistic," in the sense that they depend on the mutual information of the signal with the *entire* z. Otherwise, there is a potential "free lunch" in exploiting the contingencies in z that are cost-free. This example also raises the question of what it takes for equilibrium allocations and attention choices to be not only efficient but also non-dependent on payoff-irrelevant objects. We address this question in Proposition below.

Towards this result, we first define notions of (i) the *economy-wide fundamental* and (ii) a *group-specific fundamental* for each consumer type or firm. Heuristically, these concepts coarsen the state at the level at which all or some agents' preferences, endowments, and technologies are identical. Our specific definition describes one such coarsening that works when Θ can be ordered (which is trivially true since Θ is finite). We then use these definitions to describe when equilibria are *fundamental* and *price-tracking*.

[30] See Online Appendix E for a detailed exposition of this example.

Definition 11. The *economy-wide fundamental* is a random variable $\theta^* \in \Theta$ that can be expressed as

$$\begin{aligned} \theta^* = Q(\theta) \equiv \min\Big\{ t \in \Theta \text{ s.t. } : & \; e^j(t) = e^j(\theta), \; \forall j \in \{1, \ldots, J\} \\ & u^j(x,t) = u^j(x,\theta), \; \forall j \in \{1,\ldots,J\}, x \in \mathscr{X} \\ & H(y,c,t) = H(y,c,\theta), \; \forall y \in \mathscr{Y}, c \in \mathbb{R}_+ \Big\} \end{aligned} \tag{23}$$

The *group-specific fundamental* for consumers of type j is a random variable $\theta^j \in \Theta$ such that $\theta^j = Q^j(\theta) = \min\{t \in \Theta : u^j(x,\theta) = u^j(x,\theta'),\ \forall x \in \mathscr{X}; e^j(\theta) = e^j(t)\}$, and the corresponding object for the firm is a random variable $\theta^F \in \Theta$ such that $\theta^F = Q^F(\theta) = \min\{t \in \Theta : H(y,c,\theta) = H(y,c,t),\ \forall y \in \mathscr{Y}, c \in \mathbb{R}_+\}$.

Definition 12. An equilibrium is *fundamental* if $\overline{x}^j(\theta) = \overline{x}^{j*}(\theta^*)$ and $\overline{y}(\theta) = \overline{y}^*(\theta^*)$, for some functions $(\overline{x}^{j*})_{j=1}^J$ and $\overline{y}^*$.

Definition 13. An equilibrium is *price-tracking* if (θ^j, p) is a sufficient statistic for (θ, p) with respect to to (ϕ^j, π), for all j, and similarly (θ^F, p) is a sufficient statistic for (θ, p) with respect to to (ϕ^F, π).[31]

In a fundamental equilibrium, allocations can not depend on sunspots; and in a price-tracking equilibrium, agents only pay attention to the objects that enter their decision problems, namely their own fundamentals and the price vector.[32] Our result is that full invariance plus monotonicity guarantees that every equilibrium is both fundamental and price-tracking:

Theorem 2. *Assume that all agents' attention costs are invariant and monotone with respect to $\mathscr{G}$. Then, all equilibria are efficient, fundamental, and price-tracking.*

The proof is provided in the Appendix but the basic argument can be summarized as follows. We first show that monotonicity alone suffices for any efficient allocation, whether this is an equilibrium or not, to be fundamental. Intuitively, because non-fundamental volatility in aggregate allocations is never optimal in standard, attentive economies, it can be optimal in *in*attentive economies only if helps economize attention costs. But for this to be the case, it has to be that costs fall when agents start pay attention to non-fundamental contingencies, which is a violation of monotonicity.[33] Along with Theorem 1, this shows that the combination of invariance and monotonicity suffice for equilibria to be both efficient and fundamental. To show the price-tracking property, we finally use monotonicity to argue that an agent's learning about any other feature of the cognition state incurs strictly positive attention costs while providing no benefit relative to a constructed alternative that "ignores" these realizations of the state.

This argument also invites the following intuition for why the invisible hand may optimally produce non-fundamental volatility when attention costs are invariant but not monotone. In such circumstances, consumers can enjoy lower attention costs by making their signals, and thereby also their consumption

[31] Here, we are using the language from the remarks immediately after Definition 7. This can be translated into the original definition of sufficiency by applying the function $(\theta, p) \mapsto (\theta^j, p)$ for each agent type and describing the associated transformed prior.

[32] In fact, it is straightforward to strengthen Theorem 2 below so as to show that, when an agent enjoys utility from only a subset of the available goods, she only pays attention to the prices of this particular subset as opposed to the entire price vector.

[33] We state and prove a stronger intermediate result of this form as Lemma 3 in the Appendix.

choices, co-vary with sunspots. This willingness to pay for sunspots is quite literally manifest in the demand functions of the "twin" economy (under this representation, the sunspots are effectively transformed to preference shocks), and profit-maximizing firms try to exploit it.

What if *both* invariance and monotonicity are violated? This is precisely the scenario captured by the example mentioned in the beginning of this section and explored further in Appendix E. In such circumstances, the invisible hand may fail to produce the efficient amount of non-fundamental volatility, and may even yield multiple, Pareto-ranked equilibria. This possibility represents a form of "cognitive trap" that, unlike that articulated in Tirole (2015), does not depend on payoff (or pecuniary) externalities, originating instead in the properties attention costs *themselves*.

5.3 Expanding the Planner's Options and the Cognition State z

Our notions of feasibility and efficiency in Definitions 4-5 allowed messages to replace prices but did not give the planner any of the following options: to "customize" messages, sending different messages to different agents; to relabel or merge the underlying states of nature; and to preclude all or some agents from learning directly about the state of nature. The last option could be interpreted quite literally as the power to restrict access to information.[34]The option to merge states, on the other hand, could represent the planner's choice to "uncomplete" the markets (remove certain contingencies in markets). We can readily show an extension of our main result to such an expansion of the planner's powers, provided a commensurate enlargement in the assumed invariance of attention costs.

Toward this result, let the planner now manipulate the entire z via a collection of functions $Z^j : \Theta \to \Theta \times \mathbb{R}_+^N$, one for each type of consumer and firms. Next, let $\mathscr{A}^{\text{all}}$ denote the set of *all* such collections of functions, and let the planner choose this collection of functions from a subset $\mathscr{A} \subseteq \mathscr{A}^{\text{all}}$, which embeds the precise ways in which the planner may or may not manipulate cognition. Our benchmark scenario is then nested by restricting $\mathscr{A} = \mathscr{A}^p$, where

$$\mathscr{A}^p \equiv \left\{ (Z^j)_{j \in \{1,\dots,J\} \cup \{F\}} \;\middle|\; \exists m : \theta \to \Theta \times \mathbb{R}_+^N \text{ s.t. } Z^j(\theta) = (\theta, m(\theta)) \; \forall \theta, j \right\},$$

And any possible expansion of the planner's options relative to this benchmark maps to some $\mathscr{A}$ such that $\mathscr{A}^p \subset \mathscr{A} \subseteq \mathscr{A}^{\text{all}}$.[35]

[34]Under this first interpretation, our upcoming result, showing conditions under which such restrictions are *not* optimal, relates also to a literature on information design with inattentive receivers (Lipnowski et al., 2019; Bloedel and Segal, 2018).

[35]For instance, suppose we give the planner a single option on top of replacing prices, that of precluding any *direct* learning about the state. This maps to $\mathscr{A} = \mathscr{A}^p \cup \mathscr{A}^0$, where

$$\mathscr{A}^0 \equiv \left\{ Z = (Z_j) \;\middle|\; \exists m : \Theta \times \mathbb{R}_+^N \text{ s.t. } Z^j(\theta) = (\theta^0, m(\theta)) \; \forall \theta, j \right\},$$

for some fixed element $\theta^0 \in \Theta$. As another example, we can capture a situation where the planner may preclude direct learning about any aspect of the state other than that pertaining to an agent's *own* fundamental by letting $\mathscr{A} = \mathscr{A}^p \cup \mathscr{A}^f$, where

$$\mathscr{A}^f \equiv \left\{ Z = (Z_j) \;\middle|\; \exists m : \Theta \times \mathbb{R}_+^N \text{ s.t. } Z^j(\theta) = (Q^j(\theta), m(\theta)) \; \forall \theta, j \right\}$$

and the function Q^j identifies j's fundamental (see Definition 11).

Our notions of equilibrium (Definition 3) and efficiency (Definition 5) are unchanged, modulo to the following two natural adjustments. First, different types of agents are now allowed to have different priors, henceforth denoted by π^j; and second, consistency of the prior requires $\pi^j(z) = \sum_{\theta'} \pi_\theta(\theta') \cdot \mathbb{I}[Z^j(\theta') = z]$ for all $z \in \mathrm{Im}[Z^j]$ and all j.[36] The following version of Theorem 1 then applies provided attention costs are invariant to the full set of transformations $\mathscr{G}$:

Corollary 1. *Assume that all consumers and firms have attention costs that are invariant with respect to $\mathscr{G}$. Then, equilibria are efficient with respect to the enlarged feasibility concept described above.*

The proof of this result is quite trivial given the machinery developed above. The only difference from the proof of Theorem 1 is the translation from Proposition 3 to the desired result. Because the planner can now manage attention not only by replacing prices but also by directly relabeling the states of nature or compressing them in fewer contingencies, we must invoke invariance with respect to the entire $\mathscr{G}$ as opposed to the subset G^p. Intuitively, while invariance with respect to G^p suffices for the planner to be unable to improve upon equilibria when her options are limited to $\mathscr{A}^p$, invariance with respect to $\mathscr{G}$ is needed to guarantee that this remains true *even* when her options are expanded to $\mathscr{A}^{\text{all}}$.[37]

One can push this argument in a slight different direction as follows. In our main analysis, we defined the cognition state z as the combination of the state of nature and the price vector. This was motivated by the sufficiency of (θ, p) to describe each agent's decision problems in "free markets." But if we think of more realistic market structures, augmented with taxation or regulation, policy tools may naturally enter as additional element in the agent's decision problem. This invites a redefinition of the cognition state from $z = (\theta, p)$ to $z = (\theta, p, \tau)$, where τ is the tax, with the goal of formalizing how agents may be differentially able to attend to prices and taxes.[38] Alternatively, one could imagine including in z other endogenous objects that may enter cognition even if they do not directly enter payoffs. Examples include various indicators of others' activity, such as the average trade $\overline{x}_j$ of each type, industry output, or GDP; or even messages emitted by agents other than a social planner (e.g., media).

In Online Appendix G, we build the machinery to study economies with such a flexible definition of the cognition state z and even the potential for the social planner to send messages in a different space. It is simple to see that our main result extends with the appropriately extended notion of invariance—as implied, for example, by measuring costs by the mutual information between ω and z regardless of the space of z.

[36] It is also necessary to redefine the domain of attention costs when priors over z do not have a marginal distribution π_θ on the first element. In particular, we define the set of mass functions

$$\mathscr{P}' \equiv \left\{ \pi : \Theta \times \mathbb{R}_+^N \to [0,1] \text{ s.t. } \pi(\theta, z) = \sum_{\theta' \in \Theta} \pi_\theta(\theta') \cdot \mathbb{I}[f(\theta') = z] \text{ for some } f : \Theta \to \Theta \times \mathbb{R}_+^N \right\} \quad (24)$$

and require that attention costs are well defined functions $C : \Delta(\Omega)^{|\Theta|} \times \mathscr{P}' \to \mathbb{R}_+$.

[37] Needless to say, invariance with respect to $\mathscr{G}$ is *more* than enough when $\mathscr{A}$ is a *strict* subset of $\mathscr{A}^{\text{all}}$: as already illustrated by Theorem 1, the invariance requirement is naturally commensurate to what the planner can do. But if the analyst doesn't want to take a specific stand on this issue, then "full" invariance must be invoked.

[38] The behavioral literature on "tax salience" (e.g., Chetty et al., 2009) provides direct evidence that agents behave as if they perceive the tax component of prices more noisily.

6 Discussion: Hayek (1945) Meets Sims (2003)

We now discuss how our results relate the general-equilibrium question of whether free markets offer the best way to economize attention to an ongoing debate in decision theory and experimental economics about the appropriate formalization of attention costs.

6.1 Revisiting Hayek's (1945) Economy of Knowledge

Let us circle back to our interpretation of Hayek (1945): Does the price system in complete, competitive markets economize knowledge? And does it offer the *best* means of doing so relative to other mechanisms?

Our results offer affirmative answers to both questions under appropriate conditions. With regard to the first question, Theorem 2 says that prices summarize everything that an agent needs to know about the rest of the economy—agents need to look no further, and doing the opposite would strictly increase attention costs and reduce welfare. And with regard to the second question. Theorem 1 says that, barring a change in the "technology of knowledge," there is no way to improve upon markets.

But this seemingly reassuring lesson comes with the following important caveat: under the invariance condition used to guarantee efficiency, a social planner could implement the *same* outcomes by merely announcing the state of nature and a completely degenerate message:

Corollary 2. *Under the conditions of Theorem 1, any equilibrium can be replicated with a feasible mechanism that uses the same consumption, production and attention strategies but replaces prices with an "uninformative" message rule, namely a rule M such that $M(\theta) = \overline{m}$ for all θ and for arbitrary $\overline{m} \in \mathbb{R}_+^N$.*

Formally, the result is immediate from the last step of our proof of Theorem 1. In the context of Hayek's economy of knowledge, the result says that the same conditions that guarantees the efficiency of markets also implies that there is no welfare loss from scrapping market signals altogether and, instead, having agents redirect all their attention to learning the underlying state of nature alone. It follows that the economization of knowledge embedded in Theorem 1 is a rather weak one: it does not rely on prices' coarsely representing the state of nature, or repackaging the state of nature in a cognitively-friendly manner, but instead on the agents' ability to costlessly generate an equivalent transformation of the state of nature in their minds.

It may be instructive to return to the illustration of Section 3.6. Imagine that the state $\theta \in \{A, B\}$ represented the primitive demand shifter for another agent in the economy. Our main example cost was motivated by the idea that learning about θ was prohibitively difficult, while learning about prices was feasible but scale-dependent. This mapped to a failure of our invariance condition. Mutual information costs, on the other hand, implied that there was no difference between learning "directly" about this taste (e.g., via direct research) versus learning about prices. These costs were, in light of Theorem 1, compatible with efficiency.

More generally, this presents the following paradox: the idea that markets *strictly* economize knowledge seems most meaningful in cases that open the door to inefficiency. Away from invariance, the invisible hand may naturally do better than a centralized mechanism with "flat" messages; but an *optimal* mechanism, or regulated markets, can do *even* better.[39] For example, a planner can contemplate the removal of unnecessarily confusing markets as strictly welfare-improving. Alternatively, the planner might internalize the value of "improving" price data so as to contribute to better decisions. The examples in Section 2.5 embodied this idea, with the planner preferring to adjust the variance of prices to make learning easier.

6.2 Characterizing Invariant and Monotone Cost Functionals; and Relation to Sims (2003)

Our notions of invariance and monotonicity are closely connected to related notions from the statistics literature on information geometry (Amari, 2016), which have recently been applied in information economics by Hébert and Woodford (2019, 2020), Hébert and La'O (2020), and Caplin et al. (2020). This connection is spelled out in Appendix B within the class of *posterior-separable costs*, studied by Caplin and Dean (2015) and Denti (2018) and otherwise ubiquitous in the literature. This specialization amounts to measuring the costs of any signal by the expected divergence (loosely speaking, the distance) between the prior and the posterior induced by this signal. Within this class, our definition of invariance and monotonicity with respect to the full set of transformations $\mathscr{G}$ maps directly to a notion of *invariant and monotone divergences* due to Amari and Nagaoka (2000). And because Sims's (2003) mutual-information specification corresponds to such a divergence (namely the the Kullback-Leibler divergence), the following is a direct implication of Theorems 1 and 2:

Corollary 3. *Suppose attention costs are given by the mutual information between ω and z. Then, equilibria are efficient, fundamental and price-tracking.*

More generally, our Definition 8 is effectively a natural extension of the aforementioned information-geometry notions outside the posterior-separable class. This in turn allows us to clarify which decision-theoretic properties drive different equilibrium properties. In particular, Theorems 1 and 2 clarify that posterior-separability is neither necessary nor sufficient for equilibria to be either efficient or fundamental. Propositions 4 and 5, on the other hand, show that posterior-separability in combination with other assumptions is sufficient to prove equilibrium existence and implementation of efficient allocations, by merit of being sufficient for appropriate continuity and convexity.

Our notion of invariance, and the discussion surrounding it, is closely related to a property of state-dependent stochastic choice data defined by Caplin et al. (2020), *invariance under compression*. Loosely speaking, this property requires that observed stochastic choice patterns be invariant to relabeling states of nature and/or merging those that correspond with identical payoffs. Theorem 3 of Caplin et al. (2020) establishes that mutual information is the only cost that is consistent with invariance under compression among uniformly posterior separable costs, a subset of the posterior-separable class.[40] This confines attention to a narrower class than the one allowed in our analysis, but suggests that one can loosely think of the provided condition for efficiency as synonymous to mutual-information costs. In this sense, we justify the following "if and only if" result:

[39] This conclusion is subject to the qualification made in Section 4.2 about necessity versus sufficiency: it is possible to construct examples, albeit contrived ones, where efficiency holds despite a violation of invariance.

[40] See Section 4.1 and Definition 3 in Caplin et al. (2020).

Corollary 4. *Consider the class of uniformly posterior separable costs, as defined in Caplin et al. (2020). Within this class, Hayek's (1945) argument about the informational optimality of the price system, as formulated herein, is valid if and only if Sims's (2003) proposal for how to measure attention costs is also valid.*

As for the broader class of posterior separable costs, Theorem 2 of Caplin et al. (2020) says that consistency of stochastic choice data with invariance under compression holds if and only if the cost functional is invariant in the sense of Amari and Nagaoka (2000). In the light of the previous discussion of how this notion in turn relates to ours, it seems a safe guess that Theorem of Caplin et al. (2020) extends outside the posterior-separable class with our expanded definition of invariance. We conclude that the combination of our results with those of Caplin et al. (2020) provide a pathway for testing the validity of of Hayek's (1945) argument on the basis of stochastic choice data. We expand on this idea next.

6.3 Experimental Tests

The best available evidence on the validity of invariance under compression comes from perceptual experiments, in which there is an objective "state of the world" (e.g., "51 of 100 balls are red"); participants in the laboratory observe some representation of that state (e.g., a picture of the 100 colored balls) and then subsequently make a decision the payoffs of which depend on the state (e.g., choosing a payment schedule that depends on the number of red balls). Dean and Neligh (2017) design such an experiment, with the aforementioned set-up, to test numerous axioms of state-dependent stochastic choice including invariance under compression. Their experimental data reject invariance under compression and the authors propose a variant of the Shannon mutual information model that can rationalize the results. In an earlier experiment, Shaw and Shaw (1977) had subjects try to recall the identity of a symbol (a letter E, T, or V) which was briefly displayed, and varied the location of the letter in the display. In their data, subjects are more able to distinguish letters if those letters consistently appear in the same locations. Woodford (2012) interprets this as a rejection of mutual information's implication that the probabilities of a decision-irrelevant state, the location of the symbol, are irrelevant for decisions.

A number of recent theoretical works studying choice under uncertainty (e.g., Pomatto et al., 2018; Morris and Yang, 2019; Hébert and Woodford, 2020) have explored the formal underpinnings for cost functionals that embody various dependencies on the physical attributes of the state space. These latter two works, in particular, prioritize the idea that distinguishing "closer" states in some metric may be more difficult, like in the example of Section 2.5.

But to our knowledge there is comparatively less research on the dependence of consumer or firm decisions on scale or other stochastic properties of prices, which is the exact setting of interest. While complete invariance seems unlikely, there are plausible economic arguments for departures in both directions as illustrated in Section 2.5. For the micro-economic literature, our results underscore the importance of testing the "scale-free" properties of mutual information costs in consumer and firm choice.[41]

[41] For applied macroeconomic modelers, we would argue that mutual information (and its associated equilibrium properties) may still be a reasonable approximation. A relevant analogy may be to assuming homothetic preferences or constant-returns-to-scale production, each of which is easily rejected in the data but may provide a tractable platform on which to study other issues.

6.4 Policies for a Better Economy of Knowledge

How exactly can a planner improve welfare when invariance is violated? Literally taken, our analysis invites the replacement of markets with mechanisms that *both* regulate agents' choices *and* replace prices with other messages. But Proposition 1 suggests that such double deviations are not strictly needed: because inefficiency is present only in one dimension (the informational content of prices), improvements may be possible even if the planner hands are tied by a limited set of policy instruments.

Suppose in particular that the planner has to work *inside* the market mechanism: she cannot outright replace prices with other messages and that she is limited to allocations implemented as competitive equilibria with taxes.[42] Suppose further that the planner can only tax consumption choices, not attention choices. By Proposition 1, we know that a small tax will only have a second-order welfare loss in terms of distorting consumption and attention choices. But as long as our invariance condition fail, the implied change in the stochastic properties of prices may well have a first-order welfare gain via attention costs.

This can be illustrated in the example of Section 2 by the planner imposing a flat, non-contingent tax τ on coconut expenditure. The consumer's demand is then given by $x = 1-(1+\tau)\mathbb{E}[p|\omega]$, and as a result the equilibrium price satisfies $p = 1-\frac{1}{(1+\tau)\delta}\xi$. For given δ, a positive tax makes demands more elastic and, hence, prices less dispersed across states of nature. A small tax therefore reduces attention costs when C has the "precision cost" form of Equation 9 and a small subsidy does the same under the "variance cost" form of Equation 10, at an arbitrarily small cost in terms of allocative efficiency.[43] And of course, as implied *a fortiori* by efficiency in our standard definition, no such taxes are desirable with mutual information costs.

7 Alternative Behavioral Frictions

In this section we discuss how two straightforward extensions of our main result can shed light on the efficiency properties of a disparate set of behavioral frictions that are not nested in our main model, including narrow bracketing, bounded recall, general stochastic choice, and sparsity.

7.1 Narrow Bracketing and Bounded Recall

Our main analysis assumed that an agent's consumption of all goods is conditioned on the same information. This assumption is in line with the Arrow-Debreu formulation, which has the consumption of all goods be conditioned on the same, and complete, set of prices. But it may be natural to relax this result to capture phenomena like narrow bracketing and bounded recall within the rational inattention framework.

To see what we have in mind, let the signal variable have N sub-components, indexed as $\omega = (\omega_n)_{n=1}^N$, and require that x_n is measurable in ω_n for all n. This defines ω_n as the information set upon which the

[42] It is straightforward to extend the formulations of the consumer and firm problems and the equilibrium definition to the presence of taxes, provided that ones maintains $z = (\theta, p)$. We could alternatively expand z to include taxes as in Section 5.3.

[43] Is this merely a theoretical possibility or a practical issue? Our framework is not geared to answer this question—it is just too abstract. But going back to our real-world example of how consumers may disproportionately associate $1.99 with $1 than with $2, it may not be too far fetched to envision a regulation that requires that posted prices are rounded up or down.

consumption of good n must be conditioned on, but does not by itself put any restriction on how correlated this information may be across n. Our main analysis can now be nested by assuming that there is neither a gain nor a loss, in terms of C, from making ω_n have the same information for all n. But if we let that information to differ across goods, we can nest the model of "narrow thinking" proposed by Lian (2018) and, by extension, the type of narrow bracketing captured therein. And if interpret the index of goods, n, as different time periods, we can capture bounded recall.

Observe that there is no obstacle to proving an extension of Theorem 1 that carries the restriction that x_n be measurable in ω_n.[44] Thus we can accommodate fairly broad notions of asymmetric cognitive constraints across different goods choices or different time periods without necessarily opening the door to government intervention. In fact, this is possible even in the "vanilla" model from our main analysis. This follows from combining the results of Kőszegi and Matějka (2020), which show how narrow bracketing can arise as the optimal solution to multi-dimensional tracking problems with mutual information costs, with our result that the invariance property of such costs guarantees efficiency.

On the other hand, it may be reasonable to let narrow bracketing arise *because* of a failure of invariance. In particular, the core assumption in Lian (2018) is that it is primitively cheaper for ω_n to contain information about p_n than about p_r for $r \neq n$. This assumption captures the natural idea that "the price of apples is less salient than the price of bananas when choosing how many bananas to buy, and vice versa." But under the lenses of our analysis, this assumption also amounts to a violation of invariance. Similarly, if bounded recall means it is less costly for ω_n to contain information about current prices than about the past ones, then it, too, amounts to a violation of our invariance condition. In such cases, our results suggest that it may be possible to improve welfare by manipulating the informational content of prices or other endogenous objects that attract people's attention. Put differently, if "What You See is All There Is" in the pithy phrasing of Kahneman (2011), a social planner may be motivated to make "What You See" especially informative.

7.2 Stochastic Choice, Default Points, and Sparsity

Models of rational inattention are nested within the broader class of models of *state-dependent stochastic choice*. The converse is not true: there are models of state-dependent stochastic choice which *cannot* be micro-founded as models of information acquisition. Such models can directly be motivated as models of costly control or trembling hands (see, e.g., Morris and Yang, 2019; Flynn and Sastry, 2020). We now illustrate how our results can be extended to these contexts.

For the sake of simplifying the argument, let the consumption space $\mathscr{X}$ be discrete.[45] Next, let $\psi(x|z) \in \Delta(\mathscr{X})$ denote the probability of consuming x when the cognition state is z. Keeping with previous notation, we use the shorthand notation $\psi = (\psi(\cdot \mid z))_{z \in Z_\pi} \in (\Delta(\mathscr{X}))^{|\Theta|}$. Let attention costs be given by some function

[44] This was actually the version of this result proved in an earlier draft (Angeletos and Sastry, 2019), because the specification with choice-specific signals described here was that paper's main model.

[45] Otherwise, we could extend our notation and continuity notion to handle distributions on a continuous commodity space, similar to what we do in Appendix D for a continuous signal space.

$K^j : (\Delta(\mathscr{X}))^{|\Theta|} \times \mathscr{P} \to \mathbb{R}$. The consumer's problem is as follows:

$$\begin{aligned} \max_{\psi} \; & \sum_{x,z} u^j(x,\theta)\, \psi(x \mid z)\, \pi(z) - K_X^j[\psi,\pi] \\ \text{s.t.} \; & \sum_{x,z} (p\cdot x - p\cdot e^j(\theta) - a^j \Pi(\theta))\, \psi(x\mid z)\, \pi(z) \leq 0 \end{aligned} \tag{25}$$

The definitions of equilibrium, efficiency and invariance are similarly adapted. The arguments in Lemma 1 and hence also Theorem 1 then follow from the same premises. What changes is only the interpretation of our invariance condition: invariance now refers to whether the costs of *random plans of action*, as opposed to signals, are sensitive or not to a specific labeling of the state space.

When cost functionals satisfy Infeasible Perfect Discrimination (IPD) as proposed by Morris and Yang (2019), which is loosely speaking a notion of "continuous stochastic choice," invariance is necessarily violated. Concretely, if the grocery store randomly switches which of two substitutable products are \$1.99 versus \$2.00, an IPD consumer struggles to shift consumption decisively from one to the other, and a planner would do better to move those prices further apart. Conversely, the likelihood-separable costs motivated by Flynn and Sastry (2020), to study trembling hands conditional on observing the state, are invariant in the required ways. Inattentive economies with such a friction would therefore be efficient as per our results.

The formulation developed above can also capture the role of "default points," as studied in a long tradition of behavioral economics (e.g., Tversky and Kahneman, 1991). To illustrate, suppose that, on top of or in place of any penalty for deterministic choice, $K_X^j[\psi,\pi]$ contains a penalty for picking x away from $x^{j,d}$, where $x^{j,d}$ is a type-specific default point.[46] The latter could be a deterministic variable that depends on π, or a random variable that depends on both z and p. The dependence on π allows the default point to depend on "average" properties of the environment, such as what is "usually" optimal. The dependence on z, on the other hand, allows the default point to vary across different realizations of the relevant prices, or the actions of others. If either dependence is flexible enough, invariance is almost certainly a lost cause.

Gabaix's (2014) model of "sparsity" naturally fits in this discussion. This model amounts to a set of assumptions about the default point and the aforementioned penalty. Importantly, the default point is connected to what it would have been optimal in the absence of sparsity, which is endogenous to prices. We therefore agree with the author that sparsity opens the door to government intervention, but with a different interpretation of what this means. Proposition 8 of Gabaix (2014), which claims that the First Welfare Theorem generally fails in sparse economies, uses the *standard* notion of efficiency. That is, Gabaix's result presumes that the planner could entirely eliminate the mistakes due to inattention and concludes that the planner should of course do that. But this result may be of limited value for guiding policy, unless we imagine a planner that literally takes over people's lives. If, instead, we think of markets augmented with taxes or regulation, the appropriate notion of efficiency may be the one developed here, which treats inattention as an inevitable fact of life. Our results then say that government intervention can improve welfare *only* by exploiting the endogeneity of default points.

[46]For example, let this penalty be $\sum_{x,z} \|x - x^{j,d}(z,\pi)\|\, \psi(x \mid z)\, \pi(z)$ where the norm is the standard Euclidean one. This penalty gives the agent an incentive to "anchor" toward $x^{j,d}(z,\pi)$.

8 Conclusion

Cognitive frictions cause people to make mistakes. And these mistakes can propagate in markets, causing other people to change their behavior and the economy as a whole to malfunction relative to the textbook scenario where agents are fully rational and fully attentive. But unless a social planner has the power to "cure" the cognitive friction itself, it is not obvious why the planner should try to regulate these mistakes or manipulate market outcomes.

Our main result formalized the conditions under which such intervention may or may not be desirable. If attention costs are invariant in the sense defined in this paper, then there is no way to improve upon market outcomes, outside the elimination or bypassing of the primitive friction itself. If, instead, invariance fails, there is room for policies that manipulate or "simplify" the stochastic properties of prices, or even shut down certain "confusing" markets. Additional results provided conditions for existence of equilibria, for implementation of Pareto optima, and for equilibrium attention to be concentrated on "fundamental" objects. Last but not least, we discussed how to map other behavioral notions of imperfect optimization, such as general stochastic choice and sparsity, into our framework.

All in all, our analysis drew a link between two seemingly disparate issues in information economics: the validity of Hayek's (1945) argument about the "economy of knowledge" afforded by the price system was shown to hinge on the appropriateness of Sims's (2003) mutual information specification for attention costs. This reinforces the value of an active decision-theoretic and experimental literature that departs from mutual-information costs within the rational-inattention framework. Such departures offer the promise of understanding jointly individual behavior (the focus of this literature) and equilibrium properties including efficiency (the focus of our paper), without either a violation of individual rationality or the presumption of a policy maker that can cure, bypass, or ignore people's cognitive constraints.

Finally, our analysis committed to the interpretation of signals as a representation of inattention, cognition, or stochastic choice. This interpretation was most suitable for the connections we built to the related advances in decision theory and experimental economics (e.g., Caplin et al., 2020; Dean and Neligh, 2017; Hébert and Woodford, 2020). A *literal* interpretation in terms of collecting and processing market data is also possible. Our results then apply to the extent that such information represents primarily a private good, as in the case of a monopolist learning about demand or costs. How reasonable invariance is in this context is an open question.

References

AMARI, S.-I. (2016): *Information geometry and its applications*, vol. 194, Springer.

AMARI, S.-I. AND H. NAGAOKA (2000): *Methods of Information Geometry*, vol. 191, American Mathematical Society and Oxford University Press.

ANGELETOS, G.-M. AND J. LA'O (2018): "Optimal Monetary Policy with Informational Frictions," *Journal of Political Economy*, forthcoming.

ANGELETOS, G.-M. AND K. SASTRY (2019): "Inattentive Economies," Working Paper 26413, National Bureau of Economic Research, December 2019 Revision.

ARROW, K. J. AND G. DEBREU (1954): "Existence of an equilibrium for a competitive economy," *Econometrica*, 265–290.

BLACKWELL, D. (1951): "Comparison of Experiments," in *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, The Regents of the University of California.

BLOEDEL, A. W. AND I. R. SEGAL (2018): "Persuasion with Rational Inattention," *Available at SSRN 3164033.*

CAPLIN, A. AND M. DEAN (2015): "Revealed Preference, Rational Inattention, and Costly Information Acquisition," *American Economic Review*, 105, 2183–2203.

CAPLIN, A., M. DEAN, AND J. LEAHY (2020): "Rationally Inattentive Behavior: Characterizing and Generalizing Shannon Entropy," miméo; earlier version available as nber working paper 23652., National Bureau of Economic Research.

CHETTY, R., A. LOONEY, AND K. KROFT (2009): "Salience and Taxation: Theory and Evidence," *American Economic Review*, 99, 1145–77.

COLOMBO, L., G. FEMMINIS, AND A. PAVAN (2014): "Information Acquisition and Welfare," *Review of Economic Studies*, 81, 1438–1483.

DEAN, M. AND N. L. NELIGH (2017): "Experimental Tests of Rational Inattention," .

DEBREU, G. (1954): "Valuation Equilibrium and Pareto Optimum," *Proceedings of the National Academy of Sciences of the United States of America*, 40, 588–592.

DENTI, T. (2018): "Posterior-Separable Cost of Information," Working paper, MIT.

DENTI, T., M. MARINACCI, A. RUSTICHINI, ET AL. (2019): "Experimental cost of information," Tech. rep.

FLYNN, J. P. AND K. SASTRY (2020): "Strategic Mistakes in Large Games," *Available at SSRN.*

GABAIX, X. (2014): "A Sparsity-Based Model of Bounded Rationality," *Quarterly Journal of Economics*, 129, 1661–1710.

——— (2019): "Behavioral Inattention," in *Handbook of Behavioral Economics*, ed. by B. D. Bernheim, S. DellaVigna, and D. Laibson, Elsevier, vol. 2.

GEANAKOPLOS, J. AND H. M. POLEMARCHAKIS (1986): "Existence, Regularity, and Constrained Suboptimality of Competitive Allocations When the Asset Market Is Incomplete," *Uncertainty, Information and Communication: Essays in Honor of K.J. Arrow*, 3, 65–96.

GROSSMAN, S. (1976): “ON THE EFFICIENCY OF COMPETITIVE STOCK MARKETS WHERE TRADES HAVE DIVERSE INFORMATION,” *The Journal of Finance*, 31, 573–585.

GROSSMAN, S. J. (1981): “An Introduction to the Theory of Rational Expectations Under Asymmetric Information,” *The Review of Economic Studies*, 48, 541–559.

GROSSMAN, S. J. AND J. E. STIGLITZ (1980): “On the Impossibility of Informationally Efficient Markets,” *American Economic Review*, 70, 393–408.

GUL, F., W. PESENDORFER, AND T. STRZALECKI (2017): “Coarse Competitive Equilibrium and Extreme Prices,” *American Economic Review*, 107, 109–37.

HAYEK, F. A. (1945): “The Use of Knowledge in Society,” *American Economic Review*, 519–530.

HÉBERT, B. AND J. LA’O (2020): “Information Acquisition, Efficiency, and Non-Fundamental Volatility,” *Stanford GSB and Columbia University mimeo.*

HÉBERT, B. M. AND M. WOODFORD (2019): “Rational inattention when decisions take time,” Tech. rep., National Bureau of Economic Research.

——— (2020): “Neighborhood-Based Information Costs,” Tech. rep., National Bureau of Economic Research.

HELLWIG, C. AND L. VELDKAMP (2009): “Knowing what Others Know: Coordination Motives in Information Acquisition,” *Review of Economic Studies*, 76, 223–251.

KAHNEMAN, D. (2011): *Thinking, fast and slow*, Macmillan.

KŐSZEGI, B. AND F. MATĚJKA (2020): “Choice simplification: A theory of mental budgeting and naive diversification,” *The Quarterly Journal of Economics*, 135, 1153–1207.

LAFFONT, J.-J. (1985): “On the Welfare Analysis of Rational Expectations Equilibria with Asymmetric Information,” *Econometrica*, 53, 1–29.

LIAN, C. (2018): “A Theory of Narrow Thinking,” Working paper, Massachusetts Institute of Technology.

LIPNOWSKI, E., L. MATHEVET, AND D. WEI (2019): “Attention Management,” *Available at SSRN 3161782.*

MAĆKOWIAK, B. AND M. WIEDERHOLT (2009): “Optimal Sticky Prices under Rational Inattention,” *American Economic Review*, 99, 769–803.

——— (2015): “Business Cycle Dynamics under Rational Inattention,” *Review of Economic Studies*, 82, 1502–1532.

MAĆKOWIAK, B. A., F. MATĚJKA, AND M. WIEDERHOLT (2018): “Survey: Rational Inattention, a Disciplined Behavioral Model,” CEPR Discussion Papers 13243, C.E.P.R. Discussion Papers.

MAS-COLELL, A., M. D. WHINSTON, AND J. R. GREEN (1995): *Microeconomic Theory*, vol. 1, Oxford University Press.

MATĚJKA, F. (2016): "Rationally Inattentive Seller: Sales and Discrete Pricing," *Review of Economic Studies*, 83, 1125–1155.

MATĚJKA, F., J. STEINER, AND C. STEWART (2015): "Rational Inattention Dynamics: Inertia and Delay in Decision-Making," *CERGE-EI mimeo*.

MORRIS, S. AND M. YANG (2019): "Coordination and Continuous Stochastic Choice," *SSRN Working Paper*.

MYATT, D. P. AND C. WALLACE (2012): "Endogenous Information Acquisition in Coordination Games," *Review of Economic Studies*, 79, 340–374.

POMATTO, L., P. STRACK, AND O. TAMUZ (2018): "The Cost of Information," Working paper, arXiV paper 1812.04211.

RAVID, D. (2020): "Ultimatum Bargaining with Rational Inattention," *American Economic Review*, 110, 2948–63.

SHAW, M. L. AND P. SHAW (1977): "Optimal allocation of cognitive resources to spatial locations," *Journal of Experimental Psychology: Human Perception and Performance*, 3, 201–211.

SIMS, C. A. (1998): "Stickiness," *Carnegie-Rochester Conference Series on Public Policy*, 49, 317–356.

——— (2003): "Implications of Rational Inattention," *Journal of Monetary Economics*, 50, 665–690.

——— (2010): "Rational Inattention and Monetary Economics," *Handbook of Monetary Economics*, 3, 155–181.

SONNENSCHEIN, H. F. (2017): "Chicago and the origins of modern general equilibrium," *Journal of Political Economy*, 125, 1728–1736.

TIROLE, J. (2015): "Cognitive Games and Cognitive Traps," *Toulouse School of Economicss mimeo*.

TVERSKY, A. AND D. KAHNEMAN (1991): "Loss aversion in riskless choice: A reference-dependent model," *The quarterly journal of economics*, 106, 1039–1061.

VIVES, X. (2017): "Endogenous Public Information and Welfare in Market Games," *Review of Economic Studies*, 84, 935–963.

VIVES, X. AND L. YANG (2018): "Costly Interpretation of Asset Prices," mimeo, IESE/University of Toronto mimeo.

WOODFORD, M. (2012): "Inattentive valuation and reference-dependent choice," .

——— (2019): "Modeling Imprecision in Perception, Valuation and Choice," *NBER Working Paper No. 26258*.

Appendix

A Proofs of Main Results

In this Appendix, we first prove Lemma 1, which along with Proposition 3 completes the proof of Theorem 1. We then provide the proofs of Theorem 2 and Corollary 1, which complete our perspective on Hayek (1945). The proofs of the remaining results are in Online Appendix C.

Proof of Lemma 1

Select an arbitrary element $\overline{p} \in \mathbb{R}^N_+$. See that $g_{\overline{p}} : (\theta, p) \mapsto (\theta, \overline{p})$, the transformation that "eliminates" the variation of prices from the cognition state, is a member of G^p. Thus attention costs are invariant to any such transformation. We will use this construction to prove parts of the claim.

Reduced Preferences. We first show that, for any $\pi \in \mathscr{P}$, $\overline{u}^j((\overline{x}(\theta))_{\theta\in\Theta}, \pi) = \overline{u}^j((\overline{x}(\theta))_{\theta\in\Theta}, \pi_\theta, \overline{p})$, where

$$\begin{aligned}\overline{u}^j((\overline{x}(\theta))_{\theta\in\Theta}, \pi_\theta, \overline{p}) \equiv \max_{x,\theta} & \sum_{\omega,\theta} u^j(x(\omega), \theta)\, \psi(\omega \mid \theta, \overline{p})\, \pi_\theta(\theta) - C^j[\psi, \pi_{\overline{p}}] \\ \text{s.t.} & \sum_{\omega,\theta} x(\omega)\, \psi(\omega \mid \theta) \leq \overline{x}(\theta), \ \forall \theta \in \Theta \\ & x : \Omega \to \mathscr{X}; \quad \psi(\cdot \mid \theta, \overline{p}) \in \Delta(\Omega), \ \forall \theta \in \Theta\end{aligned} \tag{26}$$

in which $\pi_{\overline{p}}(\theta, p) = \pi_\theta(\theta) \cdot \mathbb{I}\{p = \overline{p}\}$. Note that, by the same argument given in the main text for (19) (e.g., Weierstrauss' Theorem), a solution to this program exists.

We first show $\overline{u}^j((\overline{x}(\theta))_{\theta\in\Theta}, \pi_\theta, \overline{p}) \geq \overline{u}^j((\overline{x}(\theta))_{\theta\in\Theta}, \pi)$. Suppose the opposite is true. Let (x, ϕ) denote the arg max of the program defining $\overline{u}^j((\overline{x}(\theta))_{\theta\in\Theta}, \pi)$. Construct $\psi(\omega|\theta, \overline{p}) = \phi(\omega| f_\pi(\theta))$ via the standard change of variables formula. We now consider the bundle (x, ψ) in (26). See that expected utility is unchanged by "relabeling" states $f_\pi(\theta)$. Next, informational invariance with respect to G^p guarantees that $C^j[\psi, \pi_{\overline{p}}] = C^j[\phi, \pi]$. These facts together imply that the payoff from (x, ψ) in (26) equals $\overline{u}^j((\overline{x}(\theta))_{\theta\in\Theta}, \pi)$ and therefore exceeds $\overline{u}^j((\overline{x}(\theta))_{\theta\in\Theta}, \pi_\theta, \overline{p})$. But this contradicts the definition of the maximum implicit in the last.

To establish $\overline{u}^j((\overline{x}(\theta))_{\theta\in\Theta}, \pi_\theta, \overline{p}) \leq \overline{u}^j((\overline{x}(\theta))_{\theta\in\Theta}, \pi)$, we apply an essentially identical argument in reverse using the transformation $(\theta, \overline{p}) \mapsto f_\pi(\theta)$. We omit this for brevity. Combining the two arguments, we conclude that $\overline{u}^j((\overline{x}(\theta))_{\theta\in\Theta}, \pi) = \overline{u}^j((\overline{x}(\theta))_{\theta\in\Theta}, \pi_\theta, \overline{p})$.

The equivalent arguments above replicated for a different choice $\overline{p}'$ shows that $\overline{u}^j((\overline{x}(\theta))_{\theta\in\Theta}, \pi_\theta, \overline{p}) = \overline{u}^j((\overline{x}(\theta))_{\theta\in\Theta}, \pi_\theta, \overline{p}')$ for any pair $(\overline{p}, \overline{p}')$, which means that that "prices do not enter reduced preferences," that is, the representation $\overline{u}^j((\overline{x}(\theta))_{\theta\in\Theta}, \pi) = \hat{u}^j((\overline{x}(\theta))_{\theta\in\Theta}, \pi_\theta)$ is valid.

Reduced Production Sets. We use essentially the same argument for the firms' constraint. Let

$$\overline{\boldsymbol{F}}(\pi_\theta,\overline{p}) \equiv \Big\{\overline{y} = (\overline{y}(\theta))_{\theta\in\Theta} \in \mathscr{Y}^{|\Theta|} : \exists\, (y,\psi) \text{ s.t. : } \quad \sum_\omega y(\omega)\,\psi(\omega\mid\theta,\overline{p}) \leq \overline{y}(\theta), \forall \theta\in\Theta$$

$$H\big(y(\omega), C^F[\psi,\pi_{\overline{p}}],\theta\big) \leq 0, \forall\,(\omega,\theta) : \psi(\omega\mid\theta,\overline{p}) > 0 \tag{27}$$

$$y:\Omega\to\mathscr{Y}; \quad \psi(\cdot\mid\theta,\overline{p}) \in \Delta(\Omega),\ \forall\,\theta\in\Theta\Big\}$$

We first show $\overline{\boldsymbol{F}}(\pi) \subseteq \overline{\boldsymbol{F}}(\pi_\theta,\overline{p})$. Take any element $(\overline{y}(\theta))_{\theta\in\Theta}$ of the former and the associated $(y(\omega),\phi)$. Construct $\psi(\omega\mid\theta,\overline{p}) \equiv \phi(\omega\mid f_\pi(\theta))$ and consider the bundle $(y(\omega),\psi)$. See that

$$\sum_\omega y(\omega)\,\psi(\omega\mid\theta,\overline{p}) = \sum_\omega y(\omega)\,\phi(\omega\mid f_\pi(\theta)) \leq \pi_\theta(\theta)\overline{y}(\theta) \tag{28}$$

Next, see that $\phi(\omega\mid f_\pi(\theta)) > 0$ implies that $\psi(\omega\mid\theta,\overline{p}) > 0$. Moreover, informational invariance with respect to G^p guarantees that $C^F[\psi,\pi_{\overline{p}}] = C^F[\phi,\pi]$. Combining these observations shows that $H\big(y(\omega),,C^F[\psi,\pi_{\overline{p}}],\theta\big) \leq 0$ for all (ω,θ) such that $\psi(\omega\mid\theta,\overline{p}) > 0$. Therefore $(\overline{y}(\theta))_{\theta\in\Theta} \in \overline{\boldsymbol{F}}(\pi_\theta,\overline{p})$. We again omit the reverse argument for brevity as it is essentially identical, using the transformation $(\theta,\overline{p}) \mapsto f_\pi(\theta)$. We finally replicate the argument for any two distinct $\overline{p},\overline{p}' \in \mathbb{R}^N_+$ to drop the dependence on $\overline{p}$.

Proof of Theorem 2

The claim that the equilibrium is efficient follows from Theorem 1. To prove the remaining two claims, we first state and prove an intermediate Lemma showing the consumer's strict preference for, and the firm's ability to produce, bundles that average over states of nature irrelevant to preferences, endowments and technologies. This result is the core of both remaining parts of the proof.

Step 1: Intermediate Result

Define for each consumer type j the set of functions $W^j \subseteq \{w:\Theta\to\Theta\}$ that do not separate any two states corresponding to the same payoffs and do not alter prices. That is, $w(\theta) = w(\theta') \implies u^j(x,\theta) = u^j(x,\theta'),\ \forall x \in \mathscr{X}$. For firms we similarly define the transformations that keep intact the state-dependent component of the feasibility constraint: W^F includes all functions such that $w(\theta) = w(\theta') \implies H(y,c,\theta) = H(y,c,\theta'),\ \forall y\in\mathscr{Y}$. The intermediate result is the following:

Lemma 3. *Let attention costs be invariant and monotone with respect to $\mathscr{G}$. The following properties hold:*

1. *Fix a $(\overline{x}(\theta))_\theta$ and a $w\in W^j$, and construct $(\overline{x}'(\theta))_\theta$ such that*

$$\overline{x}'(\theta) = \frac{\sum_{\theta'} \overline{x}(\theta')\pi_\theta(\theta')\mathbb{I}[w(\theta')=\theta]}{\sum_{\theta'} \pi_\theta(\theta')\mathbb{I}[w(\theta')=\theta]} \tag{29}$$

for each $\theta\in\Theta$. Then $\overline{u}^j((\overline{x}'(\theta))_\theta,\pi) \geq \overline{u}^j((\overline{x}(\theta))_\theta,\pi)$ with equality if and only if $\overline{x}(\theta) = \overline{x}'(\theta)$ for all θ.

2. *Fix a* $(\overline{y}(\theta))_\theta \in \overline{\boldsymbol{F}}(\pi)$ *and a* $w \in W^F$*, and construct* $(\overline{y}'(\theta))_\theta$ *such that*

$$\overline{y}'(\theta) = \frac{\sum_{\theta'} \overline{y}(\theta')\pi_\theta(\theta')\mathbb{I}[w(\theta') = \theta]}{\sum_{\theta'} \pi_\theta(\theta')\mathbb{I}[w(\theta') = \theta]} \tag{30}$$

for each $\theta \in \Theta$*. Then* $(\overline{y}'(\theta))_\theta \in \overline{\boldsymbol{F}}(\pi)$*.*

Proof. We start with part 1. Let us denote by (x, ϕ) and (x', ϕ') any selection of the solutions to the program (19) corresponding respectively to the parameters $(\overline{x}(\theta))_\theta$ and $((\overline{x}'(\theta))_\theta$. See that $\overline{x}(\theta) \neq \overline{x}'(\theta)$ for at least one $\theta \in \Theta$ implies that $\sum_p \phi(\omega \mid \theta, p) \neq \sum_p \phi(\omega \mid \theta', p)$ for at least some pair some (θ, θ') such that $g(\theta) = g(\theta')$.

Let us now construct a lower bound for $\overline{u}^j((\overline{x}'(\theta))_\theta, \pi)$. To do this, we define a transformation as in Definition 6, for some $g = (\theta, p) \mapsto (w(\theta), \overline{p})$ for some $\overline{p} \in \mathbb{R}^N_+$, which defines ϕ'' and π''. We then define the signal structure function (ϕ''', π) such that, for all ω, z,

$$\phi'''(\omega \mid z) = \phi''(\omega \mid g(z)) \tag{31}$$

We then propose the allocation (x, ϕ''') in program (19) corresponding to the parameters $((\overline{x}'(\theta))_\theta$. See that expected utility for the agent is the same under the proposed allocation and under (x, ϕ):

$$\begin{aligned}
\sum_{\omega,\theta} u^j(x(\omega), \theta)\, \phi(\omega \mid f_\pi(\theta))\pi_\theta(\theta) &= \sum_{\omega,\theta} \sum_{\theta'} u^j(x(\omega), \theta')\, \phi(\omega \mid f_\pi(\theta'))\, \pi_\theta(\theta')\, \mathbb{I}[w(\theta') = \theta] \\
&= \sum_{\omega,\theta} u^j(x(\omega), \theta) \sum_{\theta'} \phi(\omega \mid f_\pi(\theta'))\, \pi_\theta(\theta')\, \mathbb{I}[w(\theta') = \theta] \\
&= \sum_{\omega,\theta} u^j(x(\omega), \theta)\, \phi'''(\omega \mid f_\pi(\theta))\, \pi_\theta(\theta)
\end{aligned}$$

where the first line re-writes the sum; the second uses the definition of W^j and in particular that $u(\cdot, \theta) = u(\cdot, \theta')$ whenever $w(\theta') = w(\theta)$; and the third substitutes the definition of ϕ'''.

We will next show the cognitive cost is strictly lower under the proposed signal technology than under the technology described by ϕ. See first that $C[\phi'', \pi''] < C[\phi, \pi]$, by monotonicity of the cost function, since $g \subset \mathcal{G}$; and, second, that $C[\phi''', \pi] = C[\phi'', \pi'']$ by informational invariance, as (31) immediately implies the sufficiency relationship. Therefore, $C[\phi''', \pi] < C[\phi, \pi]$. Putting this together with the previous observation,

$$\sum_{\omega,\theta} u^j(x(\omega), \theta)\, \phi'''(\omega \mid f_{\pi'''}(\theta))\, \pi_\theta(\theta) - C[\phi''', \pi] > \overline{u}^j((\overline{x}(\theta))_\theta, \pi) \tag{32}$$

Observe now that the allocation (x, ϕ''') is feasible in program (19) with parameter $((\overline{x}'(\theta))_\theta$ by the following direct calculation for each state $\theta \in \Theta$:

$$\begin{aligned}\sum_\omega x(\omega)\,\phi'''(\omega \mid f_\pi(\theta)) &= \sum_\omega x(\omega)\,\frac{\sum_{\theta'} \phi(\omega \mid f_\pi(\theta'))\cdot \pi_\theta(\theta')\cdot \mathbb{I}[w(\theta') = w(\theta)]}{\sum_{\theta''} \pi_\theta(\theta'')\cdot \mathbb{I}[w(\theta'') = w(\theta)]} \\ &= \sum_{\theta'} \frac{\pi_\theta(\theta)\mathbb{I}[w(\theta') = w(\theta)]}{\sum_{\theta''} \pi_\theta(\theta'')\mathbb{I}[w(\theta'') = \theta]} \sum_\omega x(\omega)\,\phi(\omega \mid f_\pi(\theta'))\,\pi_\theta(\theta') \\ &\le \sum_{\theta'} \overline{x}(\theta') \frac{\pi_\theta(\theta')\mathbb{I}[g(\theta') = \theta]}{\sum_{\theta''} \pi_\theta(\theta'')\mathbb{I}[g(\theta'') = \theta]} = \overline{x}'(\theta)\end{aligned} \tag{33}$$

where the first line uses the definition of ϕ'''; the second re-arranges terms; and the third uses the feasibility of the original strategy (x,ϕ) in each state of the world.

Since (x,ϕ'''), is feasible, and $\overline{u}^j((\overline{x}'(\theta))_\theta, \pi)$ is the maximized value of the program, we must have

$$\overline{u}^j((\overline{x}'(\theta))_\theta, \pi) \ge \sum_{\omega,\theta} u^j(x(\omega),\theta)\,\phi'''(\omega \mid f_\pi(\theta))\,\pi_\theta(\theta) - C[\phi''',\pi] \tag{34}$$

This combined with (32) gives $\overline{u}^j((\overline{x}'(\theta))_\theta, \pi) > \overline{u}^j((\overline{x}(\theta))_\theta, \pi)$ as desired.

We now establish the second part of the result, for firms. Let (y,ϕ) denote the original production plan that exists, satisfies the constraints in (21), and satisfies

$$\sum_\omega y(\omega)\,\phi(\omega \mid f_\pi(\theta)) = \overline{y}(\theta),\ \forall \theta \in \Theta \tag{35}$$

construct ϕ'', ϕ''' and π'' just as above, via the change of variables associated with $g = (\theta, p) \mapsto (w(\theta), \overline{p})$. By an analogue of the same argument used for consumers, $C[\phi''',\pi] < C[\phi,\pi]$. We now consider the production constraints. See that a necessary condition for $\phi'''(\omega \mid f_\pi(\theta)) > 0$ is that $\phi(\omega \mid f_\pi(\theta')) > 0$ for some θ' such that $w(\theta') = w(\theta)$. Feasibility of the original allocation plan implies that

$$H\big(y(\omega), C^F[\phi,\pi], \theta'\big) \le 0 \tag{36}$$

The definition of W^F implies that $H(y(\omega), c, \theta') = H(y(\omega), c, \theta)$. Moreover, monotonicity of the production possibilities function in its second argument means that $H(y(\omega), c', \theta') < H(y(\omega), c, \theta')$ for any $c' < c$. Putting this together with the strict inequality for costs gives

$$H(y(\omega), C^F[\phi''',\pi], \theta) < H(y(\omega), C^F[\phi,\pi], \theta) = H(y(\omega), C^F[\phi,\pi], \theta') \le 0 \tag{37}$$

so production satisfies the technological constraint. Finally a calculation identical to (33) shows that

$$\sum_\omega y(\omega)\,\phi'''(\omega \mid f_\pi(\theta)) = \frac{\sum_{\theta'} \overline{y}(\theta')\pi_\theta(\theta')\mathbb{I}[w(\theta') = \theta]}{\sum_{\theta'} \pi_\theta(\theta')\mathbb{I}[w(\theta') = \theta]} = \overline{y}'(\theta) \tag{38}$$

for all $\theta \in \Theta$. Therefore $(\overline{y}'(\theta))_{\theta\in\Theta} \in \overline{\boldsymbol{F}}[\pi]$ as desired. □

Step 2: Equilibrium is Price-tracking

We now continue with the proof of Theorem 2. We now verify the equilibrium is price-tracking.

Assume otherwise. Imagine first that there is at least one consumer type j such that $Q^j(\theta)$ is not a sufficient statistic for θ. Apply the construction in the proof of Lemma 3 with respect to $w(\theta) = Q^j(\theta)$ to generate a new information structure $(\phi^{j\prime}, \pi')$ and consumer bundle $(x^j, \phi^{j\prime})$.[47] See that, because we have assumed that $Q^j(\theta)$ is not a sufficient statistic for θ in the conditional density, this transformation produces a strict improvement in payoffs. Additionally, under the construction, we have $\phi^{j\prime}(\omega \mid \theta, p) = \phi^{j\prime}(\omega \mid Q(\theta), p)$ which implies that $(Q(\theta), p)$ is sufficient for (θ, p).[48]

See next that the new bundle is feasible. By the following elementary calculation, which uses the law of large numbers, the new bundle has the same cost as the original bundle

$$\begin{aligned}\sum_{\omega,z}(x^j(\omega)\cdot p)\,\phi^j(\omega \mid \theta,p)\,\pi(\theta,p) &= \sum_p p\cdot\left(\sum_\omega x^j(\omega)\sum_\theta \phi^j(\omega\mid\theta,p)\,\pi(\theta,p)\right)\\ &= \sum_p p\cdot\left(\sum_\omega x^j(\omega)\sum_\theta \phi^{j\prime}(\omega\mid\theta,p)\,\pi(\theta,p)\right)\\ &= \sum_{\omega,z}(x^j(\omega)\cdot p)\,\phi^{j\prime}(\omega\mid\theta,p)\,\pi(\theta,p)\end{aligned} \tag{39}$$

Therefore the feasibility of the original bundle implies the feasibility of the one. The existence of a feasible bundle with strictly higher payoffs contradicts consumer optimality, as we have found something that is higher payoff and feasible. So the proposed equilibrium cannot exist.

Imagine next that, for the firms, $Q^F(\theta)$ is not a sufficient statistic for θ. Apply the construction in the proof of Lemma 3 with respect to $w(\theta) = Q^F(\theta)$ to generate a new production bundle (y, ϕ'). Observe that, according to (37), there is strict slack in the production constraint or

$$H(y(\omega), C^F[\phi', \pi], \theta) < 0 \tag{40}$$

in all ω, θ such that $\phi'(\omega \mid f_\pi(\theta)) > 0$. Because H is continuous and strictly increasing in its first N arguments, there exists an $\epsilon > 0$ such that

$$H(y(\omega) + \epsilon \boldsymbol{e}, C^F[\phi', \pi], \theta) < 0 \tag{41}$$

where $\boldsymbol{e}$ is an $N \times 1$ vector of ones. Hence (in some abuse of notation) the production plan $(y + \epsilon\boldsymbol{e}, \phi')$ is also feasible.[49] Producing this plan increases profits by $\epsilon(p \cdot \boldsymbol{e}) > 0$. Hence the existence of this deviation contradicts profit maximization. So the proposed equilibrium cannot exist.

[47] Specifically, $\phi^{j\prime}(\omega \mid z) = \phi'''(\omega \mid z)$ where the right-hand-side is in the Lemma's terminology.

[48] To translate to Definition 7, we would state this as a property of π' defined by (17) with $g(\theta, p) = (Q(\theta), p)$, with respect to π.

[49] For simplicity, we ignore the issue of $y(\omega) + \epsilon\boldsymbol{e} \notin \mathscr{Y}$ for some ω. To "fix" this, we would apply the exact same reasoning as in the proof of Lemma 5: we can restrict attention away from allocations at a "corner" of $\mathscr{Y}$ by making that set sufficiently large, and otherwise construct our deviation to increase $y(\omega)$ only on the (necessarily positive measure subset) of ω in which it is possible.

Step 3: Equilibrium is Fundamental

Assume there exists a non-fundamental equilibrium. Since invariance and monotonicity with respect to $\mathscr{G}$ implies invariance with respect to G^p, the equilibrium is not Pareto dominated by any allocation supported by an arbitrary message (Theorem 1).

We now show a contradiction to Pareto optimality: the social planner could remove the non-fundamental contingency in the allocation to achieve a Pareto improvement. We start by showing that at least one consumer conditions their aggregate demand or production on the non-fundamental state in the equilibrium. Assume not. The market clearing condition is

$$\sum_{j=1}^{J} \mu^j \overline{x}^j(\theta) = \sum_{j=1}^{J} \mu^j e^j(\theta) + \overline{y}(\theta) \tag{42}$$

for all $\theta \in \Theta$. Take any two states θ, θ' such that $Q(\theta) = Q(\theta')$. Under the conjecture, $\sum_{j=1}^{J} \mu^j \overline{x}^j(\theta) - \sum \mu^j e^j(\theta) = \sum_{j=1}^{J} \mu^j \overline{x}^j(\theta') - \sum \mu^j e^j(\theta')$. But this violates the implication of market clearing that $\overline{y}(\theta) \neq \overline{y}(\theta')$. Therefore at least one consumer type conditions on non-fundamental volatility.

We now propose the following allocation. For each consumer, set the allocation constructed in the proof of Lemma 3 using $w(\theta) = Q(\theta)$, which is by construction in W^j for each type j. For each producer, also set the allocation constructed in the proof of Lemma 3 using $w(\theta) = Q(\theta)$, which is by construction in W^F. Lemma 3 guarantees this allocation is strictly preferred by at least one agent type, whose allocation has changed, and is weakly preferred by all others.

We finally show that it is resource feasible. The feasibility constraint for state θ is

$$\sum_{j=1}^{J} \mu^j \overline{x}^{j\prime}(\theta) \leq \sum_{j=1}^{J} \mu^j e^j(\theta) + \overline{y}'(\theta) \tag{43}$$

which is, using the law of large numbers,

$$\sum_{j=1}^{J} \mu^j \frac{\sum_{\theta'} \overline{x}^j(\theta') \cdot \pi_\theta(\theta') \cdot \mathbb{I}[Q(\theta') = Q(\theta)]}{\sum_{\theta'} \pi_\theta(\theta') \cdot \mathbb{I}[Q(\theta') = Q(\theta)]} \leq \sum_{j=1}^{J} \mu^j E^j(\theta) + \frac{\sum_{\theta'} \overline{y}(\theta') \cdot \pi_\theta(\theta') \cdot \mathbb{I}[Q(\theta') = Q(\theta)]}{\sum_{\theta'} \pi_\theta(\theta') \cdot \mathbb{I}[Q(\theta') = Q(\theta)]} \tag{44}$$

Next see that, because $e^j(\theta) = e^j(\theta')$ for all $\theta, \theta' : Q(\theta) = Q(\theta')$, we can write

$$e^j(\theta) = \frac{\sum_{\theta'} e^j(\theta') \cdot \pi_\theta(\theta') \cdot \mathbb{I}[Q(\theta') = Q(\theta)]}{\sum_{\theta'} \pi_\theta(\theta') \cdot \mathbb{I}[Q(\theta') = Q(\theta)]} \tag{45}$$

But we know from the feasibility of the original allocation that, for all θ',

$$\sum_{j=1}^{J} \mu^j \overline{x}^j(\theta') \leq \sum_{j=1}^{J} \mu^j e^j(\theta') + \overline{y}(\theta') \tag{46}$$

Multiplying both sides by $\mathbb{I}[Q(\theta') = Q(\theta)] \cdot \pi_\theta(\theta')$, then dividing by $\sum_{\theta'} \pi_\theta(\theta') \cdot \mathbb{I}[Q(\theta') = Q(\theta)]$, gives (44).

Therefore we have shown the existence of a feasible allocation that Pareto dominates the original allocation. This is a contradiction. Therefore there cannot exist a non-fundamental equilibrium.

Proof of Corollary 1

We proceed in the following two steps. Consider any possible Pareto-dominating allocation. If the proposed message $M(\theta)$ is invertible in θ, then a twin of the argument from the proof of Theorem 1 shows it must be feasible and payoff-equivalent with the message $(\theta, P(\theta))$ and therefore cannot Pareto dominate the equilibrium. In particular, we can apply a map $M(\theta) \mapsto (\theta, P(\theta))$ and use the full invariance condition to prove an equivalent of Lemma 1.

If the proposed message is not invertible in θ, we claim (and prove at the end of this section) that a payoff-equivalent arrangement is possible with the message $(\theta, P(\theta))$. Therefore, were a Pareto-dominating allocation to exist with message $M(\theta)$, our construction has provided a Pareto-dominating allocation with message $(\theta, P(\theta))$ and contradicted the result of Proposition 3.

We now prove the intermediate step. Denote the consumer choices in the proposed arrangement via x^j, ϕ^j for each type j; and the proposed prior distribution, compatible with the message M, by π_M. Let us now construct the signal structure with likelihood distributions

$$\phi_P^j(\omega \mid \theta, P(\theta)) = \phi^j(\omega \mid f_{\pi_M}(\theta)) \quad \forall \omega \in \Omega \tag{47}$$

and the prior distribution

$$\pi_P(\theta, P(\theta)) = \pi_\theta(\theta) \tag{48}$$

This new construction in particular "un-coarsens" and re-labels the state space, maintaining the assumption that agents do not distinguish between the previously coarsened states of nature . See that (ϕ^j, π_M) is a transformation of (ϕ_P^j, π_P) for a mapping $(\theta, P(\theta)) \mapsto M(\theta)$ which is in $\mathcal{G}$. Moreover, under this mapping, π_M is sufficient with respect to π_P with respect to ϕ_P^j in the sense of Definition 7. Therefore, by invariance, $C[\phi^j, \pi_M] = C[\phi_P^j, \pi_P]$. Next, see that the consumer payoffs from program (12) are trivially the same under both information structures when the consumption strategy is fixed at x^j; and aggregate demands $\overline{x}^j(\theta)$ conditional on any θ are also the same. We then can replicate this argument for each consumer type j.

We next turn to feasibility. We use a similar constructive argument to define a new attention strategy that is valid under π_P and leads to the same costs. The feasibility constraint is similarly unaltered. Thus, with this new arrangement, the planner can generate the same aggregate supply $\overline{y}(\theta)$.

Therefore, we have constructed a new feasible arrangement that results in the same payoffs for each agent in the economy. This ends the proof of the intermediate result.

B Invariance in the Space of Posteriors

In this Appendix, we formalize the relationship of our invariance condition with notions of invariance in the literature on information geometry. This allows more direct comparison of our analysis with those of Caplin et al. (2020) and Hébert and La'O (2020).

We restrict attention to the class of *posterior separable costs*, which is studied by Caplin and Dean (2015) and Denti (2018) and encompasses many specifications used in the literature.[50] With our notation, this class is defined as follows.

Definition 14. Attention costs are *posterior separable* if they admit the following representation:

$$C[\phi,\pi]=\sum_{\omega\in\Omega}\phi_\omega(\omega)\cdot T\left[\phi_{z|\omega}(\cdot|\omega);\pi\right]-T\left[\pi;\pi\right] \tag{49}$$

where $\phi_\omega(\omega)=\sum_z \phi(\omega\mid z)\,\pi(z)\ \forall\omega$, $\phi_{z|\omega}(z\mid\omega)=\frac{\phi(\omega|z)\,\pi(z)}{\phi_\omega(\omega)}\ \forall\omega:\phi_\omega(\omega)>0$ (and $\phi_{z|\omega}(z\mid\omega)=\pi(z)$ otherwise), and $T[\cdot;\pi]:\mathscr{P}\to\mathbb{R}$ is strictly convex for each $\pi\in\mathscr{P}$.

Loosely speaking, a posterior-separable cost functional corresponds the expected increase in a measure of the difference between the posterior to the prior. A signal structure, represented by its induced posterior distributions, costs more if it generally induces posteriors that differ from the prior. This makes the specification of T the key to understanding the economic properties of information costs. We therefore ask how our notions of invariance and monotonicity of C translate to properties of T.

Consider, in particular, the following notion of invariance drawn from the literature on information geometry, which has recently been applied in economics by Hébert and Woodford (2019, 2020), Hébert and La'O (2020), and Caplin et al. (2020).

Definition 15 (Invariance, from Amari and Nagaoka (2000) and Amari (2016))**.** Let $\pi\in\mathscr{P}$ and $\pi'\in\mathscr{P}$ be two distributions on $\mathcal{Z}$, fix a $g:\mathcal{Z}\to\mathcal{Z}$, and construct $\tilde{\pi}$ and $\tilde{\pi}'$ as in (17). The functional T satisfies invariance if $T[\pi',\pi]\geq T[\tilde{\pi}',\tilde{\pi}]$ for any g, with equality if and only if the distributions can be written as $\pi(z)=\tilde{\pi}(g(z))r(z)$ and $\pi'(z)=\tilde{\pi}'(g(z))r(z)$ for the same $r:\mathcal{Z}\to[0,1]$.

When applying this definition to T in the posterior-separable cost (49), π and π' are the prior and posterior about a random variable z, and $\tilde{\pi}$ and $\tilde{\pi}'$ are the implied prior and posterior about the random variable $\tilde{z}=g(z)$, for some $g\in\mathcal{G}$. T is invariant if the "difference" between prior and posterior weakly decreases with any such transformation, and remains the same if and only if the prior and posterior completely agree about the realizations of the state that have been relabeled and/or combined.

The next Lemma, proved in Online Appendix C, verifies that invariant and monotone T in a posterior separable model implies a cost functional that is invariant and monotone per our definition:

Lemma 4. *Suppose that the cost functional is posterior separable, that T is invariant in the sense of Definition 15, and that $T[\pi;\pi]=0$ for any $\pi\in\mathscr{P}$. Then, the cost functional is invariant and monotone with respect to $\mathcal{G}$ in the sense of Definition 8.*

[50]This class include, not only the mutual-information specification proposed by Sims (2003), but also the alternatives proposed by Pomatto et al. (2018) and Hébert and Woodford (2020).

We can use this method to construct invariant cost functionals. Theorem 3.1 in Amari (2016) (p. 54) demonstrates that the class of divergences defined by

$$T[\pi;\pi'] = \sum_z \pi(z) \cdot f\left(\frac{\pi'(z)}{\pi(z)}\right) \tag{50}$$

for a differentiable convex $f(\cdot)$ satisfying $f(1) = 0$ are invariant. When $f(u) = -\log u$, T is the Kullback-Leibler divergence, which verifies the following:

Corollary 5. *Mutual information costs are invariant and monotone in the sense Definition 8.*

Applying our Theorems 1 and 2, we then conclude that mutual information costs suffice for equilibria to be not only efficient but also fundamental and price taking, as stated in Corollary 3 in the main text. This is subject to the clarification made in the beginning of Section 5.2, namely that attention costs be measured by the mutual information of the signal ω with the *whole* cognition state $z = (\theta, p)$, as opposed to, say, only p.

Moving from mutual information costs to the broader class of posterior separable costs, Theorem 2 of Caplin et al. (2020) says that, within this class, consistency of choice data with invariance under compression holds if and only if the cost functional is invariant in the sense of Definition 15. Combined with our Lemma 4, this makes invariance under compression a sufficient condition for invariance and monotonicity in the sense of Definition 8, and suggests a path for testing our invariance condition in choice data, as discussed in the main text.

Finally, Hébert and La'O (2020) work in the class of posterior separable costs and define invariance as in Definition 15, which as already explained is basically the same as ours except for the fact that our extends outside the aforementioned class. This explains the proximity between their efficiency result and ours. At the same time, Hébert and La'O (2020) offer a more thorough investigation of the type of monotonicity needed to rule out non-fundamental volatility. Loosely speaking, what is needed is monotonicity only with respect to the subset of transformations that discard information about non-fundamental variables. And while our Theorem 2 restricts attention to efficient equilibria, their result about the absence of non-fundamental volatility extends to inefficient equilibria.

Online Appendices

C Additional Proofs

In this Appendix, we fill in the details of the proofs of Propositions 1 and 2 from the example in Section 2. We then provide the proofs of the results in the main body of the paper, except for those already included in Appendix A

Proof of Proposition 1

The main logic of the proof is given in the main text. Here, we fill in the detailed calculations.

Consumer Problem. We first derive the optimal coconut demand of an agent conditional on receiving a Gaussian signal of the price with signal-to-noise ratio δ (i.e., $\mathbb{E}[p \mid \omega] = \delta\omega + (1-\delta)\mathbb{E}[p]$ and $\mathbb{E}[\omega \mid p] = p$). See that the first-order condition in the agent's program

$$\max_{x:\Omega\to\mathbb{R}} \int_{\mathcal{Z}}\int_{\Omega}\left(x(\omega) - \frac{x(\omega)^2}{2} + p(\xi - x(\omega))\right)\phi(\omega|z)\,\pi(z)\,\mathrm{d}\omega\,\mathrm{d}z \tag{51}$$

is, for each ω,

$$1 - x(\omega) = \mathbb{E}[p \mid \omega] \tag{52}$$

where $\mathbb{E}[p \mid \omega]$ is defined in the standard way by applying Bayes' rule to the density $\phi(\omega|z)\,\pi(z)$. Throughout this proof, we will use such "expectations notation" to ease the notational burden.

Substituting in the conjectured information structure, see that $x(\omega) = 1-\delta\omega-(1-\delta)\mathbb{E}[p]$. We now derive aggregate demand in the market clearing condition. Under the conjecture that $p = P(\xi)$, and the stated assumption that $\mathbb{E}[\omega \mid p] = p$, see that

$$\int_{\Omega} x(\omega)\,\phi_{\omega|z}(\omega \mid z)\,\mathrm{d}\omega = 1-\delta p-(1-\delta)\mathbb{E}[p] \tag{53}$$

As derived in the main text, market clearing requires that $1-\delta p-(1-\delta)\mathbb{E}[p] = \xi$ for all ξ. By taking the expectation of both sides we derive $\mathbb{E}[p] = 1$. Then, by substituting in this finding and solving algebraically, we establish $p - \mathbb{E}[p] = -\xi/\delta$. Thus the unique market clearing price $p = P(\xi) = 1-\frac{\xi}{\delta}$.

We next formulate agents' reduced-form benefits as a function of the choice variable δ and the conjectured price $P(\xi) = 1-\frac{\xi}{\delta'}$. See in particular that $\mathbb{E}[x(\omega)] = 0$; $\mathbb{E}\left[\frac{-x(\omega)^2}{2}\right] - \frac{\delta}{2(\delta')^2}$; $\mathbb{E}[p\xi] = -\frac{1}{\delta'}$; and $\mathbb{E}[-px(\omega)] = \frac{\delta}{(\delta')^2}$. Adding together these terms defines the reduced-form benefits function

$$b(\delta,\delta') = \frac{\delta-2\delta'}{2(\delta')^2} \tag{54}$$

as desired.

We now return to the signal acquisition problem. As established in the main text, it is without loss to focus on signals of the form $\omega_i = p + a_3\eta_i$. The signal to noise ratio of this signal, as a function of δ and

$\psi_1 = -1/\delta'$, is

$$d(a;\delta') \equiv \frac{a_3^{-2}}{a_3^{-2} + (\delta')^2} \tag{55}$$

See that, for any $\delta' \in (0,1]$, we can set $a_3 = \left(\frac{\delta}{1-\delta}\right)^{-\frac{1}{2}} \frac{1}{\delta'}$ to achieve any $\delta \in (0,1)$ as claimed. This establishes the claim that δ' is irrelevant for attention costs written as a function of (δ,δ').

We finally provide a sharper characterization of the equilibrium fixed-point equation

$$\delta^e \in \arg\max_{\delta} \left\{ b(\delta,\delta^e) - c(\delta) \right\}. \tag{56}$$

Note that the first-order condition

$$b_1(\delta^e,\delta^e) = c'(\delta^e) \tag{57}$$

is necessary and sufficient for equilibrium because $b-c$ is differentiable and strictly concave in $(0,1)$, and the corners $\delta = 0$ or $\delta = 1$ are ruled out by, respectively, $c'(0) = 0$ and $\lim_{\delta\to 1} c'(\delta) = \infty$. See then that existence is guaranteed by the continuities of b_1 and c', and uniqueness by their monotonicity.

Social Planner's Problem. Let us now consider the social planner's problem, re-printed here:

$$\begin{aligned} \max_{x,(\phi_{\omega|z})_{z\in\mathcal{Z}},M} \; & \int_{\mathcal{Z}} \int_{\Omega} \left(x(\omega) - \frac{x(\omega)^2}{2} \right) \phi(\omega|z)\, \pi(z)\; \mathrm{d}\omega\, \mathrm{d}z - C(\phi,\pi) \\ \text{s.t.} \; & \int_{\Omega} x(\omega)\, \phi_{\omega|z}(\omega \mid \xi, m)\; \mathrm{d}\omega = \xi \text{ for all } (\xi,m) \in \mathcal{Z} \\ & \pi(\xi,m) = \pi_\theta(\theta)\cdot D[M(\xi) = m] \text{ for all } (\xi,m) \in \mathcal{Z} \end{aligned} \tag{58}$$

To complete the argument in the main text, we derive the social planner's objective conditional on a fixed signal structure of the form $\omega_i = \xi + a_3\eta_i$. See that the problem conditional on the signal choice and message is globally concave and characterized by first-order conditions. Let $\hat{\lambda}(\xi)$ denote the Lagrange multiplier for the first (continuum of) constraints and $\lambda(\xi) = \frac{\hat{\lambda}(\xi)}{\pi(\xi)}$ denote the Lagrange multiplier divided by the prior. The first-order condition for the choice of $x(\omega)$, for each ω, is

$$1 - x(\omega) = \frac{\int_{\mathcal{Z}} \int_{\Omega} \lambda(\xi)\, \phi(\omega|z)\, \pi(z)\; \mathrm{d}\omega\, \mathrm{d}z}{\int_{\mathcal{Z}} \phi(\omega|z)\, \pi(z)\; \mathrm{d}z} \tag{59}$$

We can re-write the right-hand-side more illustratively as $\mathbb{E}[\lambda(\xi) \mid \omega]$, where the expectation is taken over the conditional density associated with the joint density $\phi(\omega \mid z)\pi(z)$. We substitute this expression into the resource constraint to obtain the condition

$$1 - \mathbb{E}[\mathbb{E}[\lambda(\xi) \mid \omega] \mid \xi] = \xi \tag{60}$$

using the simpler conditional expectation notation for the integral in the constraint. The Gaussianity of the right-hand-side requires that the left-hand side is also Gaussian. This pins down that the normalized co-state $\lambda(\xi)$ must be Gaussian, which also restricts it to be linear in ξ. Let us then represent $\lambda(\xi) = \lambda_0 + \lambda_1\xi$ and use the assumed signal structure to write out the conditional expectation as

$$1-\lambda_0-\lambda_1\delta\xi=\xi \tag{61}$$

from which we recover $\lambda_0=1$ and $\lambda_1=1/\delta$. Next, see that optimal consumption maximizes the Lagrangian

$$\mathscr{L}=\int_{\mathcal{Z}}\int_{\Omega}\left(x(\omega)-\frac{x(\omega)^2}{2}\right)\phi(\omega|z)\,\pi(z)\,\mathrm{d}\omega\,\mathrm{d}z-\int(\lambda_0+\lambda_1\xi)\left(\xi-\int_{\Omega}x(\omega)\,\phi_{\omega|z}(\omega\mid\xi,m)\,\mathrm{d}\omega\right)\pi(\xi)\,\mathrm{d}\xi \tag{62}$$

which can be re-expressed using the expectation notation as the following:

$$\mathbb{E}\left[x(\omega)-\frac{x(\omega)^2}{2}+\lambda(\xi-x(\omega))\right] \tag{63}$$

Observe that, comparing with the consumer's program, $\lambda=\lambda(\xi)=P(\xi)$ and $x(\omega)=1-\mathbb{E}[\lambda(\xi)\mid\omega]=1-\mathbb{E}[P(\xi)\mid\omega]$. This is identical to the consumer's program, so the same calculations apply to derive the benefits function modulo the replacement of δ' with δ. From this point, the analysis follows from the analysis in the main text.

Proof of Proposition 2

In the first economy, see that by the same arguments in the proof of Proposition 1 the social planner's objective function can be written as the following function of signal quality δ and message slope ψ_1:

$$W(\delta,\psi_1)=b(\delta,\delta)-c(\delta;\psi_1)=b(\delta,\delta)-\frac{\delta}{1-\delta}\psi_1^{-2} \tag{64}$$

See that the slope of the welfare function in ψ_1 is

$$\frac{\partial}{\partial\psi_1}W(\delta,\psi_1)=-\frac{\partial}{\partial\psi_1}c(\delta;\psi_1)=2\frac{\delta}{1-\delta}\psi_1^{-3} \tag{65}$$

and, evaluated at the equilibrium, this is

$$\frac{\partial}{\partial\psi_1}W(\delta,\psi_1)|_{\delta=\delta^e,\psi_1=1/\delta^e}=2\frac{(\delta^e)^4}{1-\delta^e}>0 \tag{66}$$

The last part ($\neq 0$) follows provided $\delta^e\neq 0$, which is easily verified using the fixed-point equation characterizing equilibrium. Therefore starting from any equilibrium, there is a first-order benefit to deviating from the equilibrium outcome and the economy is not efficient. Relative to this first-order deviation, equilibrium prices are insufficiently volatile, which corresponds with attention being inefficiently high.

Similarly, for the second economy, the social planner's objective up to scale is

$$W(\delta,\psi_1)=b(\delta,\delta)-c(\delta;\psi_1)=b(\delta,\delta)-\left(1+\sqrt{\frac{1-\delta}{\delta}}\right)\psi_1^{-2} \tag{67}$$

which has the following derivative in ψ_1 evaluated at the equilibrium:

$$\frac{\partial}{\partial \psi_1} W(\delta, \psi_1)|_{\delta=\delta^e, \psi_1=1/\delta^e} = -2\left(1+\sqrt{\frac{1-\delta^e}{\delta^e}}\right)(\delta^e)^3 < 0 \tag{68}$$

provided $\delta^e \neq 0$. Thus the planner has a first-order benefit to deviate from implementing the equilibrium outcome and the economy is inefficient. Relative to this first-order deviation, equilibrium prices are excessively volatile, which corresponds with attention being inefficiently low.

Proof of Proposition 3

We prove this by contradiction, following closely the textbook proof in Chapter 16 of Mas-Colell et al. (1995).

Let the competitive equilibrium, which we assume to exist, be denoted by $((x^j, \phi^j)_{j=1}^J, (y, \phi^F), P)$, and aggregate demands and production by $((\overline{x}^j(\theta))_{j=1}^J, \overline{y}(\theta))$. Since we have assumed equilibrium exists, it must be the case that for each consumer (x^j, ϕ^j) solves program (19). Let the proposed variant allocation be denoted by $((x^{j\prime}, \phi^{j\prime})_{j=1}^J, (y', \phi^{F\prime}), P)$, and aggregate demands and production be denoted by $((\overline{x}^{j\prime}(\theta))_{j=1}^J, \overline{y}'(\theta))$. It is without loss to assume that $(x^{j\prime}, \phi)$ solves (19) for parameters $(\overline{x}^{j\prime}(\theta))_{j=1}^J$ and π. Otherwise we could construct another feasible allocation with weakly higher payoffs for each agent, and then apply the proof by contradiction to this variant allocation. Note finally that the message and prior are unchanged from the price and prior in the equilibrium.

Assume that, under the variant allocation, all agents are weakly better off and some positive mass of agents are strictly better off. Under our restriction to symmetry in allocations within types, this implies that

$$\overline{u}^j((\overline{x}^{j\prime}(\theta))_{\theta\in\Theta}, \pi) \geq \overline{u}^j((\overline{x}^j(\theta))_{\theta\in\Theta}, \pi) \tag{69}$$

for all types j, holding strictly for at least one type.

We now establish that $\sum_\theta P(\theta)\cdot\overline{x}^{j\prime}(\theta)\,\pi_\theta(\theta) \geq \sum_\theta (P(\theta)\cdot e^j(\theta) + a^j\Pi(\theta))\,\pi_\theta(\theta)$ for all agents j, with inequality for at least one type. Let us first establish the inequality. Assume instead that $\sum_\theta P(\theta)\overline{x}^{j\prime}(\theta)\,\pi_\theta(\theta) < \sum_\theta (P(\theta)e^j(\theta) + a^j\Pi(\theta))\,\pi_\theta(\theta)$. In this case, given that $P(\theta) \in \mathbb{R}_+^N$ for each θ, there exists an ϵ such that $\sum_\theta (P(\theta)\cdot\overline{x}^{j\prime}(\theta))\,\pi_\theta(\theta) + \epsilon\sum_\theta (P(\theta)\cdot\boldsymbol{e})\,\pi_\theta(\theta) < \sum_\theta (P(\theta)e^j(\theta) + a^j\Pi(\theta))\,\pi_\theta(\theta)$, where $\boldsymbol{e} \in \mathbb{R}^N$ is a vector of ones. Moreover,

$$\overline{u}^j((\overline{x}^{j\prime}(\theta)+\epsilon\boldsymbol{e})_{\theta\in\Theta}, \pi) > \overline{u}^j((\overline{x}^{j\prime}(\theta))_{\theta\in\Theta}, \pi) \geq \overline{u}^j((\overline{x}^j(\theta))_{\theta\in\Theta}, \pi) \tag{70}$$

where the first, strict inequality uses the monotonicity of preferences established in Lemma 5, stated and proven at the end of this proof.[51] Because this bundle is strictly preferred to (x^j, ϕ^j) and feasible given the same prices (and profits), its existence would contradict consumer optimality. We use a similar argument to establish that $\sum_\theta P(\theta)\overline{x}^{j\prime}(\theta)\,\pi_\theta(\theta) > \sum_\theta (P(\theta)e^j(\theta) + a^j\Pi(\theta))\,\pi_\theta(\theta)$ for agents experiencing a strict utility gain in the new allocation. If not, $\overline{x}^{j\prime}(\theta)$ would be feasible and preferred, contradicting consumer optimality.

[51] We also use the fact that the original equilibrium allocation was interior to $\mathscr{X}^\Theta$ to establish that the proposed deviation lies in $\mathscr{X}^\Theta$. This was without loss of generality from setting the boundaries of $\mathscr{X}$ sufficiently large (see footnote 12).

Adding up the previously established conditions gives

$$\sum_j \sum_\theta (P(\theta) \cdot \overline{x}^{j\prime}(\theta))\, \pi_\theta(\theta) > \sum_j \sum_\theta (P(\theta) \cdot e^j(\theta) + a^j \Pi(\theta))\, \pi_\theta(\theta) \tag{71}$$

We next use $\sum a^j = 1$ and $\Pi(\theta) = P(\theta) \cdot y(\theta)$ to write

$$\sum_j \sum_\theta (P(\theta)) \cdot \overline{x}^{j\prime}(\theta))\, \pi_\theta(\theta) > \sum_j \sum_\theta (P(\theta) \cdot e^j(\theta)) + \sum_\theta (P(\theta) \cdot y(\theta)) \tag{72}$$

Observe that profit maximization guarantees that $\sum_\theta (P(\theta) \cdot y(\theta)) \geq \sum_\theta (P(\theta) \cdot y''(\theta))$ for any $(y''(\theta))_{\theta \in \Theta}$ in $\overline{\boldsymbol{F}}(\pi)$. This includes the proposed aggregate production plan $(y'(\theta))_{\theta\in\Theta}$. Therefore, (72) can be re-written as

$$\sum_j \sum_\theta (P(\theta) \cdot \overline{x}^{j\prime}(\theta))\, \pi_\theta(\theta) > \sum_j \sum_\theta (P(\theta) \cdot e^j(\theta)) + \sum_\theta (P(\theta) \cdot y'(\theta)) \tag{73}$$

But if this is true for $(P(\theta))_{\theta\in\Theta} \gg 0$, then it is a contradiction of feasibility. Therefore the proposed Pareto dominating allocation cannot exist.

We now complete the argument by stating and proving the required monotonicity of preferences:

Lemma 5. *For each type j, each prior π, and each pair $\overline{\boldsymbol{x}}, \overline{\boldsymbol{x}}' \in \mathscr{X}^{|\Theta|}$, the following is true: if $\overline{\boldsymbol{x}}' \gg \overline{\boldsymbol{x}}$, then $\overline{u}^j(\overline{\boldsymbol{x}}', \pi) > \overline{u}^j(\overline{\boldsymbol{x}}, \pi)$.*

Proof. This proof is constructive. Let $\overline{\boldsymbol{x}}$ and $\overline{\boldsymbol{x}}' \gg \overline{\boldsymbol{x}}$ be two consumption vectors in $\mathscr{X}^{|\Theta|}$. Necessarily, $\overline{\boldsymbol{x}}$ is in the interior of $\mathscr{X}^{|\Theta|}$. Let $\overline{x}_n(\theta)$ denote consumption of the nth good. Define $d \in \mathbb{R}^N_+$ as the point-wise minimum increase across states in the aggregate consumption vector or $d \equiv (\min_\theta(\overline{x}'_n(\theta) - \overline{x}_n(\theta)))_{n=1}^N$. Now let us take the optimizers (x, ϕ) which we assume to exist for program (19) with parameters $\overline{\boldsymbol{x}}$ and π. Let $d^\omega = (d^\omega_n)_{n=1}^N = (x_n^{\max} - x_n(\omega))_{n=1}^N$ be the distance to the boundary of $\mathscr{X}$ in each dimension; and note that for any interior $\overline{\boldsymbol{x}}$, that the "capped" vector $\tilde{d} \equiv (\min\{d_n, d^\omega_n\})_{n=1}^N$ also has all strictly positive elements.

Construct $x'(\omega) = x(\omega) + \tilde{d}(\omega)$ for each ω and note that

$$\sum_\omega x'(\omega)\, \phi(\omega \mid f_\pi(\theta)) \leq d + \sum_\omega x(\omega)\, \phi(\omega \mid f_\pi(\theta)) \leq d + \overline{x}(\theta), \ \forall \theta \in \Theta \tag{74}$$

where the third statement uses the feasibility constraint in (19). Note that $d + \overline{x}(\theta) < \overline{x}'(\theta)$ for all θ by construction. Therefore $(x'(\omega), \phi)$ is feasible in the variant program with parameters $(\overline{x}'(\theta))_{\theta\in\Theta}$ and π.

Observe next that, for a positive measure of ω, $x'(\omega) > x(\omega)$.[52] Moreover, because $x'(\omega) \gg x(\omega)$ for a positive measure of ω and $u^j(\cdot, \theta)$ represents weakly monotone preferences for each θ, expected utility is also strictly ranked:

$$\sum_{\omega,\theta} u^j(x(\omega), \theta)\, \phi(\omega \mid f_\pi(\theta)) \pi_\theta(\theta) < \sum_{\omega,\theta} u^j(x'(\omega), \theta)\, \phi(\omega \mid f_\pi(\theta)) \pi_\theta(\theta) \tag{75}$$

[52] If not, then for measure 1 of ω, we have $x(\omega) = (x_n^{\max})_{n=1}^N$. But this implies $\overline{x}(\theta) = (x_n^{\max})_{n=1}^N$ for all θ and there cannot exist an $\overline{x}'(\theta)$ that is larger in every dimension.

Therefore,

$$\sum_{\omega,\theta} u^j(x'(\omega),\theta)\,\phi(\omega \mid f_\pi(\theta))\pi_\theta(\theta) - C[\phi,\pi] > \overline{u}^j(\overline{\boldsymbol{x}},\pi) \tag{76}$$

Since a maximum exists to program (19) we know also that

$$\overline{u}^j(\overline{\boldsymbol{x}}',\pi) \geq \sum_{\omega,\theta} u^j(x'(\omega),\theta)\,\phi(\omega \mid f_\pi(\theta))\pi_\theta(\theta) - C[\phi,\pi] \tag{77}$$

by the definition of the maximum. Therefore, $\overline{u}^j(\overline{\boldsymbol{x}}',\pi) > \overline{u}^j(\overline{\boldsymbol{x}},\pi)$, which completes the proof. □

Proof of Proposition 4

Since an equilibrium in the twin economy trivially maps to an equilibrium in the original economy, it is sufficient to prove to prove existence in the twin economy.

We first prove existence of equilibrium in an economy in which all preferences are production sets are conditioned on the "message" $M(\theta) = \overline{p}$ for some arbitrary $\overline{p} \in \mathbb{R}^N_+$. Let $\pi_{\overline{p}}$ denote the associated prior. We now map our analysis to the setting of Arrow and Debreu (1954) and verify conditions I-IV in that article when preferences for each type j are represented by $\overline{u}^j(\cdot,\pi_{\overline{p}})$ and the production set is $\overline{\boldsymbol{F}}(\pi_{\overline{p}})$. Condition Ia, which requires closed and convex production sets, is assumed in our setting. Ib and Ic, which restrict pathological outcomes like purely positive production plans, are implied by the assumed monotonicity of $H(\cdot)$ and the normalization for a shut-down production of 0 is possible in any state of the world, at any cognitive cost. Condition II, that the consumption set $\mathscr{X}$ is bounded from above and below, is implied by taking $\mathscr{X}$ is a closed rectangle in the first quadrant. Condition III requires first that the utility function is continuous (IIIa) and convex in the sense of Definition 9.2 (IIIc). As stated it requires also the lack of a satiation point, but it is immediate (and remarked upon by the authors) that this can be replaced by there not existing a satiation point that is consistent with feasibility; and this latter point is implied by our assumption in footnote 12. Assumption IV requires that consumers hold positive claims on the firms, summing to one, and endowments bounded above by some element of $\mathscr{X}$. These are part of our environment.

Theorem I in Arrow and Debreu (1954) guarantees the existence of a competitive equilibrium under the stated assumptions. In particular, this equilibrium is supported by some $P : \Theta \to \mathbb{R}^N_+$.[53]

Observe finally, from Lemma 3.6, that under invariance within G^p, preferences and production sets are unchanged conditional on any π of the form $\pi(z) = \pi_\theta(\theta) \cdot \mathbb{I}[p = f(\theta)]$ for any $f : \Theta \to \mathbb{R}^N_+$. More precisely, we can apply Lemma 1 to show $\overline{u}^j(\cdot,\pi_{\overline{p}}) = \overline{u}^j(\cdot,\pi_P)$, where π_P is the prior induced by the price functional $P(\cdot)$; and $\overline{\boldsymbol{F}}(\pi_{\overline{p}}) = \overline{\boldsymbol{F}}(\pi_P)$. It is trivial then to show that we have obtained an equilibrium of the economy in which preferences for each type j are represented by $\overline{u}^j(\cdot,\pi_P)$ and the production set is $\overline{\boldsymbol{F}}(\pi_P)$. This concludes the proof.

[53] As written, Theorem I in Arrow and Debreu (1954) guarantees the existence of a vector of state-contingent prices $p \in \mathbb{R}^{N|\Theta|}_+$, but this is readily transformed to the price functional under the maintained assumption that $\pi_\theta(\theta) > 0$ for all states.

Proof of Proposition 5

The implementation of an equilibrium with transfers, and the notion of a Pareto optimum, are trivially equivalent between the inattentive economy and its attentive twin. It is therefore sufficient to prove the result in the twin economy.

Consider, then, a Pareto optimum in the twin economy that can be written as $(\overline{x}^j(\theta))_{\theta\in\Theta}$, for each j, and implemented with message $m = M(\theta)$. Let π_M be the induced prior over $\theta \times \mathbb{R}^N_+$. We now use Theorem 2 in Debreu (1954) to verify the existence of a price vector that supports such an equilibrium, when preferences and production sets are conditioned on π_M. To do so we verify conditions (I) to (V) in that article. Condition I is that $\mathscr{X}$ is convex, which is assumed. Condition II is convexity of preferences as stated in the main text. Condition III is a form of continuity. In particular, as written, it asks for every $x, x', x'' \in \mathscr{X}^\Theta$ and agent j, that $\{\alpha : \overline{u}^j((1-\alpha)x' + \alpha x'', \pi) \geq \overline{u}^j(x,\pi)\}$ and $\{\alpha : \overline{u}^j((1-\alpha)x' + \alpha x'', \pi)] \leq \overline{u}^j(x,\pi)\}$ are closed subsets of $[0,1]$.[54] See that this is a trivial consequence of the assumed continuity in the utility function in $\mathscr{X}^\Theta$. Condition IV is the convexity of the production set, which is trivial in the endowment economy. And Condition V is that $\mathscr{X}^\Theta$ is finite dimensional, guaranteed by the finite state space.

Theorem 2 in Debreu (1954) guarantees the existence of a price function $P : \Theta \to \mathbb{R}^N_+$ such that[55]

$$\overline{u}^j((\overline{x}'(\theta))_{\theta\in\Theta}, \pi_M) \geq \overline{u}^j((\overline{x}^j(\theta))_{\theta\in\Theta}, \pi_M) \implies \sum_\theta \pi_\theta(\theta)\, P(\theta) \cdot \overline{x}'(\theta) \geq \sum_\theta \pi_\theta(\theta)\, P(\theta) \cdot \overline{x}^j(\theta) \tag{78}$$

This implies that $(\overline{x}^j(\theta))_{\theta\in\Theta}$ solves program 20 when the price functional is given by $P(\cdot)$, the utility function by $\overline{u}^j((\overline{x}'(\theta))_{\theta\in\Theta}, \pi_M)$, and income by $\sum_\theta \pi_\theta(\theta)\, P(\theta) \cdot \overline{x}^j(\theta)$. See that this equilibrium can be implemented with transfers $T^j(\theta) = P(\theta) \cdot (\overline{x}^j(\theta) - e^j(\theta))$ for each θ.

We finally argue that the given allocation is a price equilibrium with transfers. To do this, we apply invariance (Lemma 1) to show $\overline{u}^j(\cdot, \pi_M) = \hat{u}^j(\cdot, \pi_\theta) = \overline{u}^j(\cdot, \pi_P)$, where π_P is the prior induced by the price functional $P(\cdot)$; and $\overline{\boldsymbol{F}}(\pi_M) = \hat{\boldsymbol{F}}(\pi_\theta) = \overline{\boldsymbol{F}}(\pi_P)$. Thus we have shown how to implement the Pareto optimum as a price equilibrium with transfers as desired.

Proof of Lemma 4

The cost of the original distribution can be written using the definition of a posterior-separable cost (49) as

$$C[\phi, \pi] = \sum_{\omega\in\Omega} \phi_\omega(\omega) \cdot T\left[\phi_{z|\omega}; \pi\right] - T\left[\pi; \pi\right] \tag{79}$$

Now consider the transformation for any g. We use the invariance of the marginal on ω and the normalization $T[\pi,\pi] = 0$ to write

$$C[\phi, \pi] = \sum_{\omega\in\Omega} \tilde{\phi}_\omega(\omega) \cdot T\left[\phi_{z|\omega}; \pi\right] \tag{80}$$

[54] This already employs two simplifications relative to the article, using the compactness of $\mathscr{X}$ to use the whole interval $[0,1]$ inclusive of endpoints and the existence of a utility representation of preferences.

[55] As written, the theorem guarantees the existence of a vector of state-contingent prices $\overline{p} \in \mathbb{R}^{N|\Theta|}_+$, but this is readily transformed to the price functional under the maintained assumption that $\pi_\theta(\theta) > 0$ for all states.

and then note that, for each ω, $T\left[\phi_{z|\omega};\pi\right] \leq T\left[\tilde{\phi}_{z|\omega};\tilde{\pi}\right]$ with equality if and only if $g(z)$ is a sufficient statistic for z in each posterior distribution. Therefore,

$$C[\phi,\pi] = \sum_{\omega\in\Omega} \tilde{\phi}_{\omega}(\omega)\cdot T\left[\phi_{z|\omega};\pi\right] \leq \sum_{\omega\in\Omega} \tilde{\phi}_{\omega}(\omega)\cdot T\left[\tilde{\phi}_{z|\omega};\tilde{\pi}\right] = C[\tilde{\phi},\tilde{\pi}] \tag{81}$$

with strict inequality if and only if the sufficient statistic condition holds in each posterior. Therefore C is monotone and invariant in the sense of Definition 8, provided that the two notions of sufficiency coincide.

It remains to verify directly this last point. We outline the direct calculation for convenience. Using the construction of Definition 8 and Bayes' rule, the posterior distributions are

$$\tilde{\phi}_{z|\omega}(z \mid \omega) = \frac{\sum_{z'\in\mathcal{Z}} \phi_{\omega|z'}(\omega \mid z')\cdot\pi(z')\cdot\mathbb{I}[g(z') = z]}{\tilde{\phi}_{\omega}(\omega)} \tag{82}$$

The previous can be re-factored as

$$\tilde{\phi}_{z|\omega}(z \mid \omega) = \sum_{z'\in\mathcal{Z}} \frac{\phi_{\omega|z'}(\omega \mid z')\pi(z')}{\tilde{\phi}_{\omega}(\omega)}\cdot\mathbb{I}[g(z') = z] \tag{83}$$

See that the marginal distribution over ω is unchanged, or $\tilde{\phi}_{\omega}(\omega) = \phi_{\omega}(\omega)$. Thus the first term of (83) is $\phi_{z|\omega}(z \mid \omega)$, and so

$$\tilde{\phi}_{z|\omega}(z \mid \omega) = \sum_{z'\in\mathcal{Z}} \phi_{z|\omega}(z \mid \omega)\cdot\mathbb{I}[g(z') = z] \tag{84}$$

This expression and the construction for the prior in Definition 8 therefore both fit the construction of $\tilde{\pi}$ and $\tilde{\pi}'$ in Lemma 4, with $\pi = \phi_{z|\omega}$; $\pi' = \pi$; $\tilde{\pi} = \tilde{\phi}_{z|\omega}$; $\tilde{\pi}' = \tilde{\pi}$. We now show the equivalence of the sufficient statistic expressions. Starting with the expression in Definition 8, $\phi_{\omega|z}(\omega \mid z) = \tilde{\phi}_{\omega|z}(\omega \mid g(z))$ for all ω and z, we use the definition of the construction to write

$$\phi_{z|\omega}(z \mid \omega) = \frac{\tilde{\phi}_{\omega|z}(\omega \mid g(z))\pi(z)}{\phi_{\omega}(\omega)} \tag{85}$$

which we can factor further as

$$\phi_{z|\omega}(z \mid \omega) = \frac{\tilde{\phi}_{\omega|z}(\omega \mid g(z))\tilde{\pi}(g(z))}{\tilde{\phi}_{\omega}(\omega)}\cdot\frac{\pi(z)}{\pi(\tilde{g(z)})} = \tilde{\phi}_{z|\omega}(g(z) \mid \omega)\cdot\frac{\pi(z)}{\tilde{\pi}(g(z))} \tag{86}$$

Set $r(z) = \frac{\pi(z)}{\tilde{\pi}(g(z))}$. See that $\tilde{\pi}(g(z)) \geq \pi(z)$ and therefore $r(z) \in [0,1]$, with the natural extension that $r(z) = 0$ if $\tilde{\pi}(g(z)) = 0$. Moreover it is trivial that $\pi(z) = \tilde{\pi}(g(z))\cdot r(z)$. For the other direction, which is not strictly needed in our proof, see that the $r(z)$ defined above is the unique choice such that $\pi(z) = \tilde{\pi}(g(z))\cdot r(z)$ and the posteriors average to the prior. We then use the same argument in reverse. Therefore the two notions of sufficiency are the same.

D Convex and Continuous Inattentive Economies

Here, we describe in more detail a version of our model economy in which the *reduced* preferences and production sets defined in Section 4.1 are convex and continuous (or closed) in the ways required in Section 5.1, to guarantee equilibrium existence and the implementability of Pareto optima (Proposition 4).

To do this, we consider an enrichment of the model with a continuous state space. Without loss of generality, we set $\Omega = [0,1]$. Define $L_{\Omega,\mathscr{X}}(\ell)$ in our context as the set of Lipschitz-continuous functions from Ω to $\mathscr{X}$ with constant ℓ, equipped with the sup norm; and $L_{\Omega,\mathscr{Y}}(\ell)$ the equivalent subset of functions mapping Ω to $\mathscr{Y}$.[56] Next, let $P(\Omega)$ denote the set of all probability measures defined on Ω (equipped with the Borel σ-algebra), and equip $P(\Omega)$ with the standard weak topology.[57] We represent such probability measures by cumulative distribution functions $\Phi \in P(\Omega)$. We re-define the domain of the each agent's cost functional as $(P(\Omega))^{|\Theta|} \times \mathscr{P}$, where $\mathscr{P}$ still carries its definition from Section 3.1. The richness of the state space will assist in proving convexity of preferences, as we will see soon.

D.1 Preferences

We will first show that preferences are continuous, which requires no special features of the cost function or state space (and indeed could have been proven, though it was not used, in our baseline environment). We will then show how the additional assumption of *posterior separable* cost functionals can be used to show convexity of preferences.

D.1.1 Problem Statement and Continuity

Let us first consider the (appropriately amended) program that defines reduced preferences:

$$\begin{aligned}\overline{u}^j((\overline{x}(\theta))_{\theta\in\Theta},\pi) \equiv \max_{x,(\Phi(\cdot|z))_{z\in S[\pi]}} \ & \sum_\theta \int_\Omega u^j(x(\omega),\theta)\ \mathrm{d}\Phi(\omega \mid f_\pi(\theta))\ \pi_\theta(\theta) - C^j[\phi,\pi] \\ \text{s.t.} \ & \sum_\theta \int_\Omega x(\omega)\ \mathrm{d}\Phi(\omega \mid f_\pi(\theta)) \le \overline{x}(\theta),\ \forall \theta \in \Theta \\ & x \in L_{\Omega,\mathscr{X}}(\ell); \quad \Phi(\cdot \mid f_\pi(\theta)) \in P(\Omega),\ \forall \theta \in \Theta\end{aligned} \tag{87}$$

We assume now that, for fixed π, that $C^j[\cdot,\pi]$ is continuous in the collection of $(\Phi(\cdot \mid z))_{z\in S[\pi]}$, equipped with the weak topology.[58] Observe that the whole objective function is therefore continuous in the choice variables. Next, see that the constraint is continuous in $\overline{x}(\theta)$, in particular because the integral operation $\int_\Omega x(\omega)\ \mathrm{d}\Phi(\omega \mid f_\pi(\theta))$ converges for a combined sequence of bounded functions $x_n(\omega) \to x(\omega)$ and $\Phi_n(\omega \mid z) \to \Phi(\omega \mid z)$ by definition.[59] Finally, note that the respective domains for the goods choice and likelihood

[56] This restriction is technical and, we conjecture, not particularly restrictive given the richness of the signal space.

[57] Let each measure in $P(\Omega)$ be associated with a cumulative distribution function F. A sequence of measures corresponds to a limit measure if the corresponding sequence of cumulative distribution functions F_n converge at all points of continuity to the corresponding limit measure's cumulative distribution, F.

[58] In particular, extend the previous definition to hold for all measures in the (finite) collection.

[59] Fixing the measure Φ_n, this requires use of the dominated convergence theorem; and then across the sequence of measures, this follows from the definition of convergence in the weak topology.

choices are compact. We can then apply Berge's theorem to program (87) to assert the existence of a maximum and, moreover, that the value function is continuous in the parameter $(\overline{x}(\theta))_{\theta\in\Theta}$. Observe that a similar argument could be used to establish continuity of preferences in our main, discrete-signal-space case in program (19). As stated in the main text, it is trivial to translate this into the required continuity of preferences.

D.1.2 Convexity

We now discuss the convexity of the preferences defined in (87). We first demonstrate convexity in the posterior-separable case.[60] In particular, say that we can write

$$C[\phi,\pi] = \int_{\Omega} T\left[\phi_{z|\omega};\pi\right] \mathrm{d}\Phi_{\omega}(\omega) - T[\pi;\pi] \tag{88}$$

where $\Phi_{\omega}(\omega)$ is the marginal CDF over ω and $\phi_{z|\omega} \in \mathscr{P}$ is the posterior distribution over z given ω.

Now imagine we have two bundles $(\overline{x}(\theta))_{\theta\in\Theta}$ and $(\overline{x}'(\theta))_{\theta\in\Theta}$, such that the associated solutions of (87) are (x,ϕ) and (x',ϕ'). Our goal is to show that

$$\overline{u}^j((\alpha\overline{x}(\theta)+(1-\alpha)\overline{x}'(\theta))_{\theta\in\Theta},\pi) \geq \alpha\overline{u}^j((\overline{x}(\theta))_{\theta\in\Theta},\pi)+(1-\alpha)\overline{u}^j((\overline{x}'(\theta))_{\theta\in\Theta},\pi) \tag{89}$$

We do this by constructing a feasible consumption, in the program with constraints $(\alpha\overline{x}(\theta)+(1-\alpha)\overline{x}'(\theta))_{\theta\in\Theta}$, that achieves the payoff on the left of the previous expression. The plan is an "α lottery" for these two plans. In particular, we set

$$\Phi''(\omega \mid z) = \begin{cases} \frac{1}{2}\Phi(2\omega \mid z) & \text{if } \omega \in [0,1/2] \\ \frac{1}{2} + \Phi'\left(2\omega - \frac{1}{2}\right) & \text{if } \omega \in (1/2,1] \end{cases} \tag{90}$$

for all z. See that the countably infinite state space allows us to accommodate such a lottery. For consumption, similarly define x'' such that[61]

$$x''(\omega) = \begin{cases} x\left(\frac{\omega}{\alpha}\right) & \text{if } \omega \in [0,\alpha] \\ x'\left(\frac{\omega-\alpha}{1-\alpha}\right) & \text{if } \omega \in (\alpha,1] \end{cases} \tag{91}$$

The combination of ϕ'', x'' replicates the strategy x,ϕ with probability α and the strategy x',ϕ' with probability $1-\alpha$.

First, see that the cost decomposes in the following calculation:

[60] Since our problem is specified in terms of signals and decision rules separately, our argument is much simpler (but also less strong) than the result in Denti (2018) showing that the stochastic choice programs associated with posterior-separable cost functionals are convex.

[61] If this construction results in a Lipschitz discontinuous $x''(\omega)$, it can trivially be amended (alongside the construction of Φ'') to leave a small "gap" between the functions which is interpolated continuously.

$$
\begin{aligned}
C[\phi,\pi] &= \int_\Omega T\left[\phi''_{z|\omega};\pi\right] \mathrm{d}\Phi''_\omega(\omega) - T[\pi;\pi] \\
&= \int_0^\alpha T\left[\phi''_{z|\omega};\pi\right] \mathrm{d}\Phi''_\omega(\omega) + \int_\alpha^1 T\left[\phi''_{z|\omega};\pi\right] \mathrm{d}\Phi''_\omega(\omega) - T[\pi;\pi] \\
&= \alpha \int_\Omega T\left[\phi_{z|\omega};\pi\right] \mathrm{d}\Phi_\omega(\omega) + (1-\alpha)\int_\Omega T\left[\phi'_{z|\omega};\pi\right] \mathrm{d}\Phi'_\omega(\omega) - T[\pi;\pi] \\
&= \alpha C[\phi,\pi] + (1-\alpha)C[\phi',\pi]
\end{aligned} \tag{92}
$$

This leverages the fact that the posterior-separable cost is linear in posteriors and hence in linear in the operation of "combining" signals in the current way.[62] Next, see that by almost the same logic the expected utility term is linear:

$$
\begin{aligned}
\sum_\theta \int_\Omega u^j(x''(\omega),\theta)\ \mathrm{d}\Phi''(\omega \mid f_\pi(\theta))\ \pi(f_\pi(\theta)) = \alpha \sum_\theta \int_\Omega u^j(x(\omega),\theta)\ \mathrm{d}\Phi(\omega \mid f_\pi(\theta))\ \pi(f_\pi(\theta)) + \\
(1-\alpha)\sum_\theta \int_\Omega u^j(x'(\omega),\theta)\ \mathrm{d}\Phi'(\omega \mid f_\pi(\theta))\ \pi(f_\pi(\theta))
\end{aligned} \tag{93}
$$

and finally that the constraints are linear:

$$
\int_\Omega x''(\omega)\ \mathrm{d}\Phi''(\omega \mid f_\pi(\theta)) \leq (\alpha \overline{x}(\theta) + (1-\alpha)\overline{x}'(\theta)) \tag{94}
$$

This allows us to establish a lower bound for $\overline{u}^j((\alpha\overline{x}(\theta) + (1-\alpha)\overline{x}'(\theta))_{\theta\in\Theta},\pi)$, and therefore show (89) as intended.

D.2 Convex and Closed Production Sets

We now describe a similar extension for firms. We will first show that production sets are closed, which like the previous argument of continuous utility functions requires no specific properties of the cost function or state space (and could have been shown in our baseline environment).

To establish convexity, we will require the combination of posterior separability of costs with a particular linearity restriction on H. In particular, for some vector $\nu \in \mathbb{R}^N_+$, we require the representation

$$
H\big(y(\omega), C^F[\phi,\pi],\theta\big) = \tilde{H}\big(y(\omega) + \nu C^F[\phi,\pi],\theta\big) \tag{95}
$$

in which we maintain $\tilde{H}(0,\theta) = 0$ for all θ and that H is increasing. We also require the regularity condition that $\tilde{H}(\cdot,\theta)$ is continuous for any value of θ.

See that this formulation can capture the "replacement" of any element y_n of y with $y_n + \nu_n C^F[\phi,\pi]$ in the relevant part of a production function. This allows attention costs to capture requirements of additional inputs and/or destroyed outputs. In addition, building on the discussion in footnote 12, it allows for atten-

[62]An additional feature of posterior-separable cost functionals, which we do *not* need to use in this construction to achieve weak convexity, is that costs would be strictly reduced if a positive measure of signals corresponding to the same posteriors (or optimal decisions) were combined.

tion costs to be specified in terms of a "dummy input" with a normalized price (e.g., of 1), so firms effectively maximize profits net of attention costs.

D.2.1 Problem Statement and Closed Production Sets

We carry over the regularity assumptions stated above for production plans and feasible distributions. We redefine the feasible production set as

$$\begin{aligned}\overline{\boldsymbol{F}}(\pi) \equiv \Big\{(\overline{y}(\theta))_{\theta\in\Theta} : \exists\ (y(\omega),\phi) \text{ s.t.} & \\ & \int_\Omega y(\omega)\ \mathrm{d}\Phi(\omega \mid f_\pi(\theta)) \le \overline{y}(\theta), \forall \theta\in\Theta \\ & \tilde{H}\left(y(\omega)+\vec{v}C^F[\phi,\pi],\theta\right) \le 0, \forall(\omega,\theta)\ \phi \text{ a.e.} \\ & y \in L_{\Omega,\mathscr{Y}}(\ell); \quad \phi(\cdot\mid f_\pi(\theta)) \in P(\Omega),\ \forall\,\theta\in\Theta\Big\}\end{aligned} \tag{96}$$

where "ϕ a.e." denotes that the statement holds for any subset of $\Omega\times\Theta$ that has positive measure under the signal distribution.

We first argue that this set is closed. Assume it is not. Then there exists some point $(\overline{y}(\theta))_{\theta\in\Theta} \notin \overline{\boldsymbol{F}}(\pi)$, and a sequence $(\overline{y}_n(\theta))_{\theta\in\Theta} \to (\overline{y}(\theta))_{\theta\in\Theta}$ (in the Euclidean distance of $\mathbb{R}^{N|\Theta|}$) such that $(\overline{y}_n(\theta))_{\theta\in\Theta} \in \overline{\boldsymbol{F}}(\pi)$ for each n. In particular, for any ϵ, there exists some $(\overline{y}_K(\theta))_{\theta\in\Theta}$ such that $\|(\overline{y}_K(\theta))_{\theta\in\Theta} - (\overline{y}(\theta))_{\theta\in\Theta}\| < \epsilon$. Next, consider the program

$$\begin{aligned}\max_{y,(\Phi(\cdot|z))_{z\in S[\pi]}} \quad & -\|(\overline{y}'(\theta))_{\theta\in\Theta} - (\overline{y}(\theta))_{\theta\in\Theta}\| \\ \text{s.t.} \quad & \int_\Omega y(\omega)\ \mathrm{d}\Phi(\omega \mid f_\pi(\theta)) \le \overline{y}'(\theta), \forall \theta\in\Theta \\ & \tilde{H}\left(y(\omega)+\vec{v}C^F[\phi,\pi],\theta\right) \le 0, \forall(\omega,\theta)\ \phi \text{ a.e.} \\ & y \in L_{\Omega,\mathscr{Y}}(\ell); \quad \phi(\cdot\mid f_\pi(\theta)) \in P(\Omega),\ \forall\,\theta\in\Theta\end{aligned} \tag{97}$$

See that all the constraints define a compact set. Boundedness is by assumption, given the spaces in which y and ϕ lie. To see closedness, we first establish in the first constraint that for any sequence $y_n \to y$ and $\Phi_n \to \Phi$, $\overline{y}'_n(\theta) = \int_\Omega y_n(\omega)\ \mathrm{d}\Phi_n(\omega \mid f_\pi(\theta))$ is a convergent sequence (using the definition of the weak topology for probability measures); and if each $\overline{y}'_n(\theta) \le \overline{y}'(\theta)$, then also $\lim_{n\to\infty} \overline{y}'_n(\theta) \le \overline{y}'(\theta)$. Next, for the second constraint, see that for any subset of $\Omega\times\Theta$ which has positive measure under the limit signal distribution represented by Φ, there must exist a subsequence of Φ_n for which this set has a positive measure; index this subsequence by k. Along the subsequence, we argue

$$H\left(y_n(\omega)+\vec{v}C^F[\phi_n,\pi],\theta\right) \to H\left(y(\omega)+\vec{v}C^F[\phi,\pi],\theta\right) \tag{98}$$

using the continuity of H and C^F; and then by a similar argue to the above, using the fact that

$$H\left(y_n(\omega)+\vec{v}C^F[\phi_n,\pi],\theta\right) \le 0 \tag{99}$$

everywhere along the subsequence, argue that $H\left(y(\omega)+\vec{v}C^F[\phi,\pi],\theta\right)\leq 0$.

Having established points, (97) is a maximization of a continuous function on a compact set and by Weierstrauss' theorem must admit a solution. Denote the maximized value of the program as V. Since we have assumed that $(\overline{y}(\theta))_{\theta\in\Theta}\notin\overline{\boldsymbol{F}}(\pi)$, it must be that the value function V of this program satisfies $V<0$. But then, as argued above using the fact that $(\overline{y}(\theta))_{\theta\in\Theta}\in \mathrm{cl}\overline{\boldsymbol{F}}(\pi)$, there exists some implementable solution of (97) that achieves value $-V/2>V$. This is a contradiction. Therefore the set $\overline{\boldsymbol{F}}(\pi)$ must be closed.

Like with the proof of continuous preferences, see that a much less technical version of the same argument could be used to prove the closedness of $\overline{\boldsymbol{F}}(\pi)$ (and the existence of a solution to our profit maximization problem) when the state space was discrete, as in the main model of Section 3.

D.2.2 Convexity

We now show that the aggregate production set is convex under the restriction in (89). To do this, we show that for any $(\overline{y}(\theta))_{\theta\in\Theta},(\overline{y}'(\theta))_{\theta\in\Theta}\in\overline{\boldsymbol{F}}(\pi)$, we have also $(\alpha\overline{y}(\theta)+(1-\alpha)\overline{y}'(\theta))_{\theta\in\Theta}\in\overline{\boldsymbol{F}}(\pi)$. We show this with the following construction which mirrors the construction for convex preferences. Let y,ϕ and y',ϕ' be the production plan and attention choice that exist and satisfy the conditions in (96) with respect to constraints $(\overline{y}(\theta))_{\theta\in\Theta},(\overline{y}'(\theta))_{\theta\in\Theta}$. Let us now construct a variant plan y'',ϕ'' in which attention is given by the α lottery as in (90):

$$\Phi''(\omega\mid z)=\begin{cases}\frac{1}{2}\Phi(2\omega\mid z) & \text{if } \omega\in[0,1/2]\\ \frac{1}{2}+\Phi'\left(2\omega-\frac{1}{2}\right) & \text{if } \omega\in(1/2,1]\end{cases}\tag{100}$$

for all z. For the same argument given in the last subsection, $C^F[\phi'',\pi]=\alpha C^F[\phi,\pi]+(1-\alpha)C^F[\phi',\pi]$ on account of the posterior separability. The production plan is more complex than the analogue with consumer demand and is given by the following:[63]

$$y''(\omega)=\begin{cases}y\left(\frac{\omega}{\alpha}\right)+\vec{v}(1-\alpha)(C^F[\phi,\pi]-C^F[\phi',\pi]) & \text{if } \omega\in[0,\alpha]\\ y'\left(\frac{\omega-\alpha}{1-\alpha}\right)+\vec{v}\alpha(C^F[\phi',\pi]-C^F[\phi,\pi]) & \text{if } \omega\in(\alpha,1]\end{cases}\tag{101}$$

See that, since attention costs are bounded, the domain of $\mathscr{Y}$ can be specified such that $y''(\omega)\in\mathscr{Y}$ without upsetting compactness of $\mathscr{Y}$.

[63] As stated previously: if this construction results in a Lipschitz discontinuous $y''(\omega)$, it can trivially be amended (alongside the construction of Φ'') to leave a small "gap" between the functions which is interpolated continuously.

See that this satisfies the capacity constraint as

$$\begin{aligned}\int_{\Omega} y''(\omega)\, \mathrm{d}\Phi''(\omega \mid f_\pi(\theta)) &= \alpha \int_{\Omega} (y(\omega) + \vec{v}(1-\alpha)(C^F[\phi,\pi] - C^F[\phi',\pi]))\, \mathrm{d}\Phi(\omega \mid f_\pi(\theta)) \\ &\quad + (1-\alpha)\alpha \int_{\Omega} (y'(\omega) + \vec{v}\alpha(C^F[\phi',\pi] - C^F[\phi,\pi]))\, \mathrm{d}\Phi'(\omega \mid f_\pi(\theta)) \\ &\le \pi(\theta)(\alpha\overline{y}(\theta) + (1-\alpha)\overline{y}'(\theta)) + \alpha(1-\alpha)\vec{v}(C^F[\phi,\pi] - C^F[\phi',\pi]) \\ &\quad + \alpha(1-\alpha)\vec{v}(C^F[\phi',\pi] - C^F[\phi,\pi]) \\ &= \alpha\overline{y}(\theta) + (1-\alpha)\overline{y}'(\theta)\end{aligned} \tag{102}$$

Next, we check feasibility. See first that, for $\omega \in [0,\alpha]$,

$$\begin{aligned}y''(\omega) + \vec{v}\cdot C^F[\phi'',\pi] &= y''(\omega) + \vec{v}\cdot(\alpha C[\phi,\pi] + (1-\alpha)C[\phi',\pi]) \\ &= y\left(\frac{\omega}{\alpha}\right) + \vec{v}(1-\alpha)(C^F[\phi,\pi] - C^F[\phi',\pi]) + \vec{v}\cdot(\alpha C[\phi,\pi] + (1-\alpha)C[\phi',\pi]) \\ &= y\left(\frac{\omega}{\alpha}\right) + \vec{v}((\alpha+1-\alpha)C^F[\phi,\pi] + (1-\alpha-(1-\alpha))C[\phi',\pi]) \\ &= y\left(\frac{\omega}{\alpha}\right) + \vec{v}(C^F[\phi,\pi])\end{aligned} \tag{103}$$

An essentially identical argument shows $y''(\omega) + \vec{v}\cdot C^F[\phi'',\pi] = y\left(2\omega - \frac{1}{2}\right) + \vec{v}(C^F[\phi',\pi])$ when $\omega \in (\alpha,1]$. Since we know

$$\begin{aligned}H\left(y(\omega) + \vec{v}C^F[\phi,\pi],\theta\right) &\le 0, \forall(\omega,\theta)\ \phi \text{ a.e.} \\ H\left(y'(\omega) + \vec{v}C^F[\phi',\pi],\theta\right) &\le 0, \forall(\omega,\theta)\ \ \phi' \text{ a.e.}\end{aligned} \tag{104}$$

we therefore know the corresponding feasibility condition for y'',ϕ''. Thus all conditions in (96) are satisfied and $(\alpha\overline{y}(\theta) + (1-\alpha)\overline{y}'(\theta))_{\theta\in\Theta} \in \overline{\boldsymbol{F}}(\pi)$.

E An Example with Non-fundamental Volatility

In this Appendix, we extend the coconuts example of Section 2 to include a non-fundamental component in the state of nature. We use this to illustrate that both inefficiency and non-fundamental volatility may be possible under a specification of attention costs that violates the invariance and monotonicity properties invoked in, respectively, 1 and 2, even thought it may *appear* to be mutual information. On the way, we also draw a connection to the much-studied issue of coordination in information choice (e.g., Hellwig and Veldkamp, 2009; Myatt and Wallace, 2012) and in particular to Tirole (2015) and Vives and Yang (2018)

The state of nature is now given by $\theta = (\xi, \nu)$, where ξ is still the endowment of coconuts while $\nu \sim \mathcal{N}(0,1)$ is independent of ξ and plays no fundamental role in the economy. The assumption that ν is a "sunspot" is particularly stark, though it could just as well be a variable that is fundamental for another group of agents that do not interact with the coconut-consuming agents. Signals are restricted to be jointly Gaussian with the cognition state z, where $z = (\theta, p) = (\xi, \nu, p)$. We show first that, when attention costs are an increasing and convex function of the mutual information of ω with z, that neither the equilibrium nor planner conditions on non-fundamental volatility and that the economy remains efficient. On the other hand, when costs are given by the mutual information of ω and p alone, there exists multiple, Pareto-ranked equilibria, all but the *worst* one featuring non-fundamental volatility. These possibilities extend to situations where attention costs separately put some weight on the mutual information of ω with the three components of z, as opposed to "holistically" measuring the mutual information of ω with z.

E.1 Efficient Case: Mutual Information with (ξ, ν, p)

We first describe an efficient case of the model in which cognitive costs are given by the mutual information of the agent's signal ω with the entire vector (ξ, ν, p).

We start by characterizing equilibrium. Conjecture that the price has the form $p = P(\xi, \nu) = 1 + \psi_1 \xi + \psi_2 \nu$. Say the agent gets a signal of the form $\omega_i = a_1 \xi + a_2 p + a_3 \eta_i + a_4 \nu$. Define $V_\omega = (a_1 + a_2 \psi_1)^2 + (a_4 + a_2 \psi_2)^2 + a_3^2$ as the variance of the signal and see that

$$\mathbb{E}[p \mid \omega] = 1 + \beta(\omega - a_2) \tag{105}$$

with

$$\beta \equiv \frac{(a_1 + a_2 \psi_1)\psi_1 + (a_4 + a_2 \psi_2)\psi_2}{V_\omega} \tag{106}$$

Furthermore, the agent's utility is given by

$$\max_{x:\Omega \to \mathbb{R}} \mathbb{E}\left[x(\omega) - \frac{x(\omega)^2}{2} + p(\xi - x(\omega))\right] = \frac{\beta^2 V_\omega}{2} - \psi_1 \tag{107}$$

We now restrict attention, without loss, to signals that set $a_2 = 0$ and calculate

$$\beta^2 V_\omega = \frac{(a_1 \psi_1 + a_4 \psi_2)^2}{a_1^2 + a_4^2 + a_3^3} \tag{108}$$

We now consider attention costs. Se that the mutual information of ω with (ξ, ν, p), maintaining $a_2 = 0$, is proportional to

$$\log\left(\frac{a_3^2 + a_1^2 + a_4^2}{a_3^2}\right) \tag{109}$$

and hence the cost can be written as

$$K\left(\log\left(\frac{a_3^2 + a_1^2 + a_4^2}{a_3^2}\right)\right) \tag{110}$$

for an increasing and convex K satisfying $K'(0) = 0$.

Finally, market clearing requires that

$$-\beta(a_1 \xi + a_4 \nu) = \xi \tag{111}$$

for all realizations of ξ and ν. This pins down that, if other agents play an attention strategy (a_1', a_3', a_4'), equilibrium prices satisfy

$$\psi_1 = -\frac{(a_1')^2 + (a_3')^2}{(a_1')^2} \tag{112}$$

while ψ_2 is undetermined. Note also that, in any fixed point, we would also require $a_4' = 0$ lest markets not clear.

We show by contradiction that, in equilibrium, $\psi_2 = 0$. Imagine not. Then, see that the marginal benefit of increasing or decreasing a_4 around $a_4 = 0$ is necessarily non-zero, evaluated at the fixed point, while the marginal cost of moving in either direction is 0. Therefore $a_4 \neq 0$. But this contradicts market clearing.

Having established this, the fixed point problem reduces to

$$(a_1^e, a_3^e) \in \arg\max_{a_1, a_3} \left(-\frac{1}{2} \frac{a_1^2}{a_1^2 + a_3^3} \left(\frac{(a_1')^2 + (a_3')^3}{(a_1')^2} \right)^2 + \left(\frac{(a_1')^2 + (a_3')^2}{(a_1')^2} \right) - K \left(\log \left(\frac{a_3^2 + a_1^2}{a_3^2} \right) \right) \right) \tag{113}$$

which is the same fixed-point problem solved in the previously studied efficient economy with mutual information costs.

Let us now consider the social planner's problem. Observe again the irrelevance of the message, since (ξ, ν) is necessarily a sufficient statistic for any $m = M(\xi, \nu)$. Hence it is without loss to consider the same mutual information cost described above by (110). If the co-state variable is $\lambda(\xi) = \lambda_0 + \lambda_1 \xi$, we replicate the last subsection's argument that the planner's optimal consumption plan conditional on the information structure is $x(\omega) = 1 - \mathbb{E}[\lambda \mid \omega]$; and plugging into the constraint establishes that $a_4 = 0$ and $\lambda_1 = -\frac{a_1^2 + a_3^2}{a_1^2}$. The planner's problem is therefore

$$(a_1^*, a_3^*) \in \arg\max_{a_1, a_3} \left(\frac{1}{2} \left(\frac{(a_1')^2 + (a_3')^3}{(a_1')^2} \right) - K \left(\log \left(\frac{a_3^2 + a_1^2}{a_3^2} \right) \right) \right) \tag{114}$$

which is again the planner's problem from the previous subsection. Thus the competitive equilibrium is efficient, for exactly the same reasons outlined in the main efficient example.

E.2 Inefficient Case: Mutual Information with p

Let us now assume that cognitive costs are given by some transformation of the mutual information of the signal ω with *only* the price p.

We start with equilibrium. We continue to assume agents get a signal of the form $\omega_i = a_1 \xi + a_2 p + a_3 \eta_i + a_4 \nu$ but now normalize, necessarily without loss when p is linear in (ξ, ν), that $a_2 = 1$. We further restrict attention to signals such that $\mathbb{E}[p(\omega_i - p)] = 0$, which is also without loss given our normalization (as any part of ω_i that projected onto p could be re-normalized out). The substantial restriction implied by this assumption, under the conjecture $p = P(\xi, \nu) = 1 + \psi_1 \xi + \psi_2 \nu$, is that

$$\psi_1 a_1 + \psi_2 a_4 = 0 \tag{115}$$

See that the signal-to-noise ratio can now be written as

$$\delta_p = \frac{\psi_1^2 + \psi_2^2}{\psi_1^2 + \psi_2^2 + (a_1^2 + a_3^2 + a_4^2)} \tag{116}$$

Moreover, $\mathbb{E}_i[p] = \delta_p \omega_i + (1-\delta_p)$ and the cognitive cost is $K(-\log(1-\delta_p))$, by direct analogues of the arguments used in Proposition 1.

Market clearing can be re-arranged to

$$-\delta_p((a_1+\psi_1)\xi + (a_4+\psi_2)\nu) = \xi, \tag{117}$$

which in turn gives the following coefficient restrictions:

$$\begin{aligned} a_4 + \psi_2 &= 0 \\ -\delta_p(a_1+\psi_1) &= 1 \end{aligned} \tag{118}$$

The first, in particular, implies that the signal ω_i does not co-vary with ν.

We now solve for equilibrium information structures that are consistent with equilibrium. See that Equations 116 and 118, re-arranged, imply the following restrictions for (ψ_1, ψ_2, a_3)

$$\begin{aligned} \psi_2^2 &= -\psi_1(\delta_p^{-1} + \psi_1) \\ a_3^2 \delta_p &= \psi_1^2 + \psi_2^2 - \frac{1}{\delta_p} \end{aligned} \tag{119}$$

The signs of (a_3, ψ_2) are unsurprisingly indeterminate so we normalize them to be positive for now. Modulo this issue, Equation 119 is a system of two non-colinear equations with three unknowns and admits a continuum of solutions. In particular, these solutions are indexed by $\psi_1 \in [-\delta_p^{-1}, -1]$ and have

$$\begin{aligned} \psi_2 &= \sqrt{-\psi_1(\delta_p^{-1} + \psi_1)} \\ a_3 &= \delta_p^{-1}\sqrt{-1-\psi_1} \end{aligned} \tag{120}$$

Note that $\psi_1 = -1$ has $a_3 = 0$ and $\psi_2 = \sqrt{\delta_p^{-1} - 1}$. In this equilibrium, $\omega \propto \xi$ and the agent can obtain the first-best. This works with respect to cognitive constraints because the price is sufficiently contaminated with noise that "precise" observation of ξ corresponds with "imprecise" observation of p. This can be understood also as a cognitive externality whereby changing the dependence of prices on ν affects the cost of obtaining a fixed posterior about the fundamental ξ. And it can be considered a violation of the spirit of our later invariance and monotonicity conditions (Definition 8) along the same lines.

Let us now characterize equilibrium. The "costs and benefits" representation from the main example remains valid up to the indeterminacy of the slope $\frac{\partial p}{\partial \xi} = \psi_1$. In particular, a straightforward extension of the calculation in the proof of Proposition 1 reveals, as long as $\psi_1 \neq 0$, the benefits of attention can be expressed as

$$b(\delta_p;\psi_1) = \frac{\psi_1(\delta_p\psi_1 + 2)}{2} \tag{121}$$

which depends on others' actions only through the slope of the demand curve. Thus any pair (δ, ψ_1) that solves

$$\delta \in \arg\max_{\delta_p} \left[\frac{\psi_1(\delta_p\psi_1 + 2)}{2} - K(\log(1-\delta_p)) \right] \qquad \text{and} \qquad \psi_1 \in [-1, -\delta^{-1}] \tag{122}$$

evaluated at the restriction $\delta_p = \delta$. As two particular examples, there exist the maximally non-fundamental equilibria such that $\psi_1 = -1$, and the idiosyncratic noise equilibrium in which $\psi_1 = -1/\delta$, as studied in Proposition 1.

See that these equilibria are Pareto-ranked: welfare (the consumer's ex ante utility) is strictly higher in the equilibria with "more public noise," or lower $|\psi_1|$, for a fixed δ_p. We can thus interpret equilibria with less noise in prices as "cognitive traps," where agents fail to correlate their information/inattention in a welfare-improving manner. This reminds the cognitive traps articulated in Tirole (2015). But whereas that particular form of cognitives traps depends on pecuniary or payoff externalities (equivalently, on some inefficiency in the underlying, attentive economy), ours does not: it originates exclusively in the specification of attention costs and, in particular, on the joint violation of invariance and monotonicity (as established by Theorems 1 and 2).

Finally, see that any equilibrium in which $\psi_1 \neq 0$ (i.e., p varies even slightly with ξ) is dominated by the social planner's allocation in which $m = M(\xi, \nu) = \nu$, $\omega_i = \xi$, and $x_i = \omega_i$. This yields the first-best allocation without any costs of learning, as the mutual information of ω_i with m is 0. This embodies the most extreme possible exploitation of the endogeneity of the price or message and the associated cognitive externality.

F Incomplete Insurance over Noise

In this Appendix, we sketch a variant of our model which disallows complete markets over the noise in the agent's signal ω. Our specification of "partially complete markets" and a parallel efficiency concept is an adaptation of the framework of Geanakoplos and Polemarchakis (1986). We observe that the appropriate notion of constrained efficiency, which keeps symmetric restriction on markets' and the social planner's ability to insurance across the realizations of ω, is unlikely to hold outside the case of quasi-linear utility. The economic intuition, as in Geanakoplos and Polemarchakis (1986), is that a social planner can manipulate prices to partially simulate insurance over ω.

For the example, we restrict the environment in a number of ways to simplify analysis. First, we study a single-type ($J = 1$) endowment economy. Second, we assume throughout that a "first-order-condition" approach is necessary and sufficient to characterize the optimum for consumers and the social planner. This presumes differentiability and concavity of the utility function. Finally, we allow the good N to be a residual good or "money" which has no (binding) domain constraint and automatically adjusts to meet the budget constraint of the agent given each realization of the signal ω and the physical state θ.

We depart from our conventional notation in the following way. We write the consumption vector as $x = (x_{-N}, x_N)$, where x_{-N} is an $N-1$ length vector of the other goods' consumption and x_N is a scalar representing money consumption. We similarly write the utility function as $u(x_{-N}, x_N, \theta)$, with partial derivatives u_1 and u_2 in each sub-component of consumption. We write the price vector as $P(\theta) = (P_{-N}(\theta), P_N(\theta))$, where $P_{-N}(\theta)$ is an $N-1$ length vector of prices relative to the price of money and $P_N(\theta)$ is a scalar corresponding to the price of money. Finally we let the endowment be written as $e(\theta) = (e_{-N}(\theta), e_N(\theta))$ and the consumption domain as $\mathscr{X} = (\mathscr{X}_{-N}, \mathscr{X}_N)$ with a similar interpretation.

F.1 Equilibrium

We begin by characterizing necessary conditions for equilibrium. The consumer chooses her consumption of goods and money as a function of the cognitive state; a money balance $b(\theta)$ in each state θ; and a signal structure. Their optimal choices solve the following problem:

$$\begin{aligned}
&\max_{x_{-N}, x_N, b, \phi} \sum_\theta \sum_\omega u\left(x_{-N}(\omega), x_N(\omega,\theta), \theta\right) \phi(\omega|\theta)\pi(\theta) - C[\phi,\pi] \\
&\text{s.t. } x_N(\omega,\theta) = b(\theta) + P_{-N}(\theta)\cdot(e_{-N}(\theta) - x_{-N}(\omega)) \qquad \forall \omega,\theta \\
&\sum_\theta P_N(\theta)(b(\theta) - e_N(\theta)) \leq 0
\end{aligned} \tag{123}$$

with the domain constraints $x_{-N} : \Omega \to \mathscr{X}_{-N}$, $x_N : \Omega \times \Theta \to \mathscr{X}_N$, $b : \Theta \to \mathbb{R}$, and $\phi \in \Phi$. The combination of these constraints capture how the consumer can freely transfer money across θ but not ω.

Equilibrium is characterized by choices that solve (123) taking as given a price functional $P : \Theta \to \mathbb{R}_+^N$ and a prior consistent with that price functional, as well as the market clearing condition

$$\sum_\omega x_{-N}(\omega)\phi(\omega|\theta) = e_{-N}(\theta) \qquad \forall \theta \tag{124}$$

The market clearing condition for money is redundant from Walras' law.

We now provide a set of necessary conditions for equilibrium, coming from consumer optimization. Let the Lagrange multipliers for the first and the second constraint of (123) be, respectively, $\chi(\omega,\theta)\phi(\omega|\theta)\pi(\theta)$ and η. The first-order-conditions characterizing optimal the optimal choice of money consumption $x_N(\omega,\theta)$ is

$$u_2\left(x(\omega), x_N(\omega,\theta), \theta\right) = \chi(\omega,\theta) \qquad \forall \omega,\theta \tag{125}$$

where u_2 is the derivative with respect to the second argument. In words, the marginal value of wealth must equal the marginal utility from consuming money. The first-order conditions for choosing consumption are

$$\sum_\theta \left[u_1\left(x_{-N}(\omega), x_N(\omega,\theta), \theta\right) - P_{-N}(\theta)\chi(\omega,\theta)\right]\phi(\theta|\omega) = 0 \qquad \forall \omega \tag{126}$$

where u_1 is the derivative with respect to the first argument and $\phi(\theta \mid \omega)$ is constructed in the standard way via Bayes' rule. Finally, the first-order condition for choosing money balances is

$$\sum_{\omega} \chi(\omega,\theta)\phi(\theta|\omega)\pi(\theta) = P_N(\theta)\eta \qquad \forall\theta \tag{127}$$

which equates the price of money with its average value across realizations of ω.

Substituting in the marginal value of wealth, equal to the marginal consumption value of money, one derives the following two necessary conditions for consumer optimality and hence necessary conditions for equilibrium:

$$\begin{aligned} \sum_{\theta} \left[P_{-N}(\theta)\, u_2\left(x_{-N}(\omega), x_N(\omega,\theta),\theta\right)\right] \phi(\theta|\omega) &= \sum_{\theta} \left[u_1\left(x_{-N}(\omega), x_N(\omega,\theta),\theta\right)\right] \phi(\theta|\omega) \\ P_N(\theta)\eta &= \sum_{\omega} u_2(x_{-N}(\omega), x_N(\omega,\theta),\theta)\phi(\omega|\theta)\pi(\theta) \end{aligned} \tag{128}$$

where the first holds for all ω and the second for all θ.

F.2 Efficient Allocations

We now define the problem of a social planner who faces the same "asset-spanning" restriction. Loosely speaking, the planner solves the consumer's program (123) but with control over prices. More specifically, they choose the tuple $(x_{-N}, x_N, b, P, \phi)$ to maximize the following expected utility objective

$$\sum_{\theta}\sum_{\omega} u\left(x_{-N}(\omega), x_N(\omega,\theta)\right)\phi(\omega|\theta)\pi(\theta) - C[\phi,\pi] \tag{129}$$

subject to the following feasibility constraints. First, consumption of all goods, including money, is physically feasible in each state θ:

$$\begin{aligned} e_{-N}(\theta) &\geq \sum_{\omega} x_{-N}(\omega)\phi(\omega|\theta) \\ e_N(\theta) &\geq \sum_{\omega} x_N(\omega,\theta)\phi(\omega \mid \theta) \end{aligned} \tag{130}$$

Second, money consumption is residual-spending:

$$x_N(\omega,\theta) = b(\theta) + P_{-N}(\theta)\cdot\left(e_{-N}(\theta) - x_{-N}(\omega)\right) \qquad \forall\omega,\theta$$

Finally, the prior π is consistent with the price functional, or $f_\pi(\theta) = (\theta, P(\theta))$ for each θ. We define an efficient allocation as one that solves the above problem, which is without loss in our economy with one type and symmetric choices. Observe that this planner's problem is not perfectly parallel with the one in our main analysis, as the prices or messages have an instrumental role in transferring resources across states. We will show that even this "weaker" planner, with a tighter implementability constraint, can improve upon equilibrium allocations.

We now use a first-order approach to derive necessary conditions for an efficient allocation, again presuming such an allocation exists and it is characterized by first-order conditions. Let the Lagrange multipliers on the three constraints above respectively be $\lambda(\theta)\pi(\theta)$ (non-money goods feasibility), $\mu(\theta)\pi(\theta)$ (money feasibility), and $\tau(\theta,\omega)\phi(\omega|\theta)\pi(\theta)$ (residual spending). The first-order conditions for x_N, x_{-N}, and b are, respectively,

$$\begin{aligned} 0 &= \mu(\theta) + \tau(\omega,\theta) - (u_2(x_{-N}(\omega), x_N(\omega,\theta), \theta)) \quad \forall \omega, \theta \\ 0 &= \sum_\theta [u_1(x_{-N}(\omega), x_N(\omega,\theta),\theta) - \lambda(\theta) - P_{-N}(\theta)\tau(\omega,\theta)]\,\phi(\theta|\omega) \;\; \forall \omega \\ 0 &= \sum_\omega \tau(\omega,\theta)\phi(\omega|\theta) \quad \forall \theta \end{aligned} \tag{131}$$

See that these conditions resemble the ones in the consumer problem up to the introduction of shadow values $\mu(\theta)$ and $\lambda(\theta)$ for money and non-money consumption goods, respectively. Next, the first-order condition for the goods prices $P_{-N}(\theta)$, which show up only in the residual spending constraint, is[64]

$$\sum_\omega \tau(\omega,\theta)\,[e_{-N}(\theta) - x_{-N}(\omega)]\,\phi(\omega|\theta) = 0 \qquad \forall \theta \tag{132}$$

This condition bears special comment, as it is drives the key wedge between efficiency and equilibrium. By adjusting the price vector $P_{-N}(\theta)$ in any state θ, the planner affects marginal money consumption conditional on each (ω,θ) in proportion to the net endowment $e_{-N}(\theta) - x_{-N}(\omega)$. This has marginal value $\tau(\omega,\theta)$ to the consumer. Such adjustments are exactly what the invisible hand will not do in our environment.

We now develop the above argument mathematically. Like in the consumer's problem, we solve out for $\tau(\omega,\theta)$ and re-write the first-order conditions of (131) in the following way:

$$\begin{aligned} \sum_\theta [u_0(x_{-N}(\omega), x_N(\omega,\theta),\theta)]\,\phi(\theta|\omega) &= \sum_\theta \left[P_{-N}(\theta)\left(u_2(x_{-N}(\omega), x_N(\omega,\theta),\theta) - \mu(\theta)\right) + \lambda(\theta)\right]\phi(\theta|\omega) \\ \mu(\theta) &= \sum_\omega u_2(x_{-N}(\omega), x_N(\omega,\theta),\theta)\phi(\omega|\theta) \end{aligned} \tag{133}$$

which can be directly compared with the consumer's equilibrium first-order-conditions (128). More specifically, see that these exactly correspond if $\mu(\theta) \equiv P_N(\theta)\eta$ and $\lambda(\theta) = P_{-N}(\theta)$.

The condition for choosing prices is

$$\sum_\omega \left[u_2(x_{-N}(\omega), x_N(\omega,\theta),\theta) - \mu(\theta)\right][e_{-N}(\theta) - x_{-N}(\omega)]\,\phi(\omega|\theta) = 0 \qquad \forall \theta \tag{134}$$

where, in continuation of the discussion above of using prices for redistribution, $u_2(x_{-N}(\omega), x_N(\omega,\theta),\theta) - \mu(\theta)$ gives the marginal consumption value of money (net of the social cost), and the planner changes prices until the average effect on payoffs via residual money consumption is zero. See that this condition is generically incompatible with the previous conjecture for μ when evaluated at the equilibrium consumption levels, as there is no analogue for this condition in the equilibrium conditions. This is the sense in which one should not "expect" that an efficient allocation, which necessarily satisfies (134), should also correspond to an equi-

[64] The *dollar* prices $P_N(\theta)$ show up only in the prior and the cognitive cost, and therefore function like the purely informational messages of the main analysis.

librium allocation.[65] Note that the argument makes no reference to the structure of C apart from assuming there is *some* randomness over signal realizations. Thus the logic above applies in a number of simplified settings including exogenously-incomplete-information economies, which underscores our general point that the pathway of inefficiency via incomplete insurance is not a direct consequence of learning or rational inattention.

An important exception is a quasi-linear case, in which u_2 is constant as a function of x_{-N} and x_N. In particular, if we can write $u_2(x_{-N}(\omega), x_N(\omega,\theta),\theta) = f(\theta)$ for some function f, then the planner's second first-order condition in (133) reduces to $\mu(\theta) = f(\theta)$ and the condition for choosing prices simplifies via

$$\begin{aligned} 0 &= \sum_{\omega} \left[u_2(x_{-N}(\omega), x_N(\omega,\theta),\theta) - \mu(\theta) \right] [e_{-N}(\theta) - x_{-N}(\omega)]\, \phi(\omega|\theta) \\ &= \sum_{\omega} \left[f(\theta) - \mu(\theta) \right] [e_{-N}(\theta) - x_{-N}(\omega)]\, \phi(\omega|\theta) \\ &= \sum_{\omega} [0]\, [e_{-N}(\theta) - x_{-N}(\omega)]\, \phi(\omega|\theta) = 0 \end{aligned}$$

where the second and third line respectively use the simplifications above. This represents mathematically the logic that, with quasi-linear utility (accommodating even θ-dependence in the "slope"), there is no first-order benefit to distributing goods across realizations of ω. And it verifies that the precise role of the quasi-linear utility assumption in the example of Section 2 was to substitute for our main model's assumption of insurance over ω.

G Economies with a Broader Cognition State

In this Appendix, we describe how to model economies in which the cognition state includes variables other than the state of nature and prices. We then sketch how an appropriately extended notion of invariance delivers a straightforward extension of Theorem 1 and Corollary 1.

G.1 Environment

We first define an expanded notion of an inattentive market economy, in which the cognitive process (as captured by the definition of z) takes as inputs not only on the exogenous state of nature and the price vector but also the following additional objects: *transfers*, *goods taxes*, *aggregate trades*, and *exogenous signals* (or, more informally, "media").

Equilibrium is defined as in Definition 3 modulo the following changes. As in Section 5.1, we allow each consumer's wealth to include a state-dependent transfer, given by $t^j = T^j(\theta)$, for some exogenously specified and type-specific transfer rule $T^j : \Theta \to \mathbb{R}$. We next have consumers and firms take as given after-tax prices $p + \tau$, where the goods taxes are state-dependent, too, or $\tau = T(\theta)$ for some exogenously specified tax rule

[65] A more formal argument of "generic non-efficiency," as pursued in Geanakoplos and Polemarchakis (1986), would show that even if necessary condition (134) were satisfied evaluated at the equilibrium, then in a small perturbation of the economic environment (appropriately defined) it could be made not to hold.

$T:\Theta\to\mathbb{R}$ The collection $(T,(T^j)_{j=1}^F)$ is restricted to be such that the government budget balances in each state of nature, or $\sum_{j=1}^J(T^j(\theta)+T(\theta)\cdot\overline{x}^j(\theta))+T(\theta)\cdot\overline{y}(\theta)=0$ for all $\theta\in\Theta$. Next, we incorporate as part of the equilibrium a set of exogenous signals $s\in\mathbb{R}^Q$, for some $Q>0$, defined by the exogenous mapping $s=S(\theta)$ for some $S:\Theta\to\mathbb{R}$. And finally, we allow the cognition state in equilibrium to be

$$z=\left(\theta,p,\tau,(t^j)_{j=1}^J,(\overline{x}^j)_{j=1}^J,s\right)$$

We write the domain of this object as $z\in\Theta\times\mathscr{B}_0$ where

$$\mathscr{B}_0=\mathbb{R}_+^N\times\mathbb{R}^N\times\mathbb{R}^J\times\mathscr{X}^J\times\mathbb{R}^Q \tag{135}$$

is the composition of all the domains for the objects after the state of nature. We define a $\mathscr{B}\supseteq\mathscr{B}_0$ as an even larger possible domain for this state, which the social planner may take advantage of. We correspondingly re-define the admissible set of priors over z corresponding to any particular $B\subseteq\mathscr{B}$, including $B=\mathscr{B}_0$, as

$$\mathscr{P}_B\equiv\left\{\pi:\Theta\times\mathbb{R}_+^N\to[0,1]\ \text{ s.t. }\ \pi(\theta,z)=\pi_\theta(\theta)\,\mathbb{I}[f(\theta)=z],\ \text{for some } f:\Theta\to B\right\}. \tag{136}$$

Along the same lines, we also re-define, for every $\pi\in\mathscr{P}$ and domain B, the set $Z_{\pi,B}\equiv\{(\theta,b):\pi(\theta,b)>0\}\subset\Theta\times B$ and the function $f_{\pi,B}:\theta\to\Theta\times b$ with $f_{\pi,B}(\theta)=\{(\theta,b):\pi(\theta,b)>0\}\in\Theta\times B$ for any $\theta\in\Theta$.

Toward our efficiency concept, we start with the feasibility concept in 4 along with the following expanded notion of messages. The definition of an arrangement now includes an image set $B\subseteq\mathscr{B}$, the message is sent via a rule $M:\Theta\to B$.

G.2 Extending Invariance and Monotonicity

For the arguments above to be well-defined, we require an expanded notion of the cost functional. We assume, in particular, that for each $B\subseteq\mathscr{B}$, each agent (and the firm) has a well-defined cost indexed by B:

$$C_B^j:(\Delta(\Omega))^{|\Theta|}\times\mathscr{P}_B \tag{137}$$

maintaining the discussion from footnote 13 of the main text of how to order the elements of the support.

In this context, we re-define our notions of transformations and invariance. Let $\mathscr{H}\equiv\{h:(\Theta\times B)\to(\Theta\times B');B,B'\subseteq\mathscr{B}\}$ be the dictionary of transformations from any $\Theta\times B$ to any $\Theta\times B'$. We next define transformations of information structures based on this larger set of functions:

Definition 16 (Transformations of Information Structures, Revisited)**.** Consider two information structures (π,ϕ) and $(\tilde{\pi},\tilde{\phi})$ and a function $h\in\mathscr{H}$ mapping $\Theta\times B$ to $\Theta\times B'$. We say that $(\tilde{\phi},\tilde{\pi})$ is the transformation of

(ϕ, π) under h if

$$\tilde{\pi}(z) = \sum_{z'} \pi(z') \mathbb{I}[h(z') = z] \qquad \forall z \in Z_{\pi,B} \tag{138}$$

$$\tilde{\phi}(\omega|z) = \frac{\sum_{z' \in Z_{\pi,B}} \phi(\omega|z') \pi(z') \mathbb{I}[g(z') = z]}{\tilde{\pi}(z)} \qquad \forall \omega \in \Omega, z \in Z_{\tilde{\pi},B'} \tag{139}$$

Sufficiency is similarly extended:

Definition 17. Consider two information structures (ϕ, π) and $(\tilde{\phi}, \tilde{\pi})$ such that $(\tilde{\phi}, \tilde{\pi})$ is the transformation of (ϕ, π) under $h \in \mathscr{H}$ mapping $\Theta \times B$ to $\Theta \times B'$. We say that $\tilde{\pi}$ is sufficient for π with respect to ϕ if $\phi(\omega \mid z) = \tilde{\phi}(\omega \mid h(z))$ for all ω and all z such that $\pi(z) > 0$.

Finally, invariance and monotonicity are written as follows:

Definition 18. Fix a set $H \subseteq \mathscr{H}$. Consider any function $h \in H$ and any two information structures (ϕ, π) and $(\tilde{\phi}, \tilde{\pi})$ such that $(\tilde{\phi}, \tilde{\pi})$ is the transformation of (ϕ, π) under g. A cost functional C is

1. invariant with respect to H if $C[\phi, \pi] = C[\tilde{\phi}, \tilde{\pi}]$ whenever $\tilde{\pi}$ is sufficient for π with respect to ϕ.

2. monotone with respect to H if $C[\phi, \pi] > C[\tilde{\phi}, \tilde{\pi}]$ whenever $\tilde{\pi}$ is not sufficient for π with respect to ϕ

While the wording of these definitions is rather similar, the definition has changed considerably. In particular, invariance and monotonicity can hold along the natural "extension" of a cost functional to different domains or state spaces. What does this extension mean? For posterior-separable cost functionals, as defined in Appendix B, this relies on an appropriate extension of the divergence T. And if those divergences are f-divergences as defined in Equation 50, re-printed here:

$$T[\pi; \pi'] = \sum_z \pi(z) \cdot f\left(\frac{\pi'(z)}{\pi(z)}\right), \tag{140}$$

this is immediate as the support of the probability distribution always remains finite with a length of at most $|\Theta|$. To use an example, it is reasonable to say that "mutual information costs are invariant and monotone as per Definition 18" *provided* that one presumes the natural extension of "always taking the mutual information of the signal and the state," whatever the (finite-support) state space is.

G.3 Possible Equilibrium Results

Given the above adaptation, it seems safe to conjecture, although we do not formally state and prove, that all of our equilibrium results including efficiency (Theorem 1), existence (Proposition 4), implementability of Pareto optima (Proposition 5), and fundamentalness of equilibrium (Theorem 2) readily extend to our environment under invariance and monotonicity for all of $\mathscr{H}$ (e.g., as in our natural extension of mutual information). The content of these statements is that, under fully invariant and monotone costs like mutual information,

- Tax instruments are never optimal for a purely "informational" role, for instance to nudge agents toward or away from learning about certain contingencies.

- Media, expert opinions, blogs, and tweets, as captured in the exogenous signals $s = S(\theta)$, have no instrumental effect on equilibrium, as agents can always construct equivalent signals via unrestricted information acquisition.

- Noise in publicly available signals, which otherwise has no instrumental role, is optimally ignored in equilibrium.

- "Market data" in the form of trades or aggregate consumption statistics cannot be designed to have better informational content.

These lessons follow (we conjecture) from essentially the same premises articulated in our main analysis, but would be made more concrete in this extended environment.