



Institutional Members: CEPR, NBER and Università Bocconi

WORKING PAPER SERIES

A Framework for the Analysis of Self-Confirming Policies

*P. Battigalli, S. Cerreia-Vioglio, F. Maccheroni,
M. Marinacci, T. Sargent*

Working Paper n. 573

This Version: February 11, 2022

IGIER – Università Bocconi, Via Guglielmo Röntgen 1, 20136 Milano –Italy
<http://www.igier.unibocconi.it>

The opinions expressed in the working papers are those of the authors alone, and not those of the Institute, which takes non institutional policy position, nor those of CEPR, NBER or Università Bocconi.

A framework for the analysis of self-confirming policies*

P. Battigalli,^a S. Cerreia-Vioglio,^a F. Maccheroni,^a M. Marinacci,^a T. Sargent^b

^aUniversità Bocconi and IGIER

^bNew York University and Hoover Institution

February 11, 2022

Abstract

This paper provides a general framework for analyzing self-confirming policies. We study self-confirming equilibria in recurrent decision problems with incomplete information about the true stochastic model. We characterize stationary monetary policies in a linear-quadratic setting.

KEY WORDS: Self-confirming equilibrium, partial identification, monetary policy.

1 Introduction

Perspective Policies often persist. Absent switching costs, the reason must be that the goals and beliefs of the policy maker also persist, which is possible only if long-run data coincide with what the policy maker expected. This belief-confirmation property does *not* imply that a persistent policy is justified by *correct* beliefs. Indeed, even if the stochastic consequences of the persistent policy are observable in the long run, a policy maker's expectations about the consequences of alternative policies may be incorrect. We call *self-confirming* a policy justified by beliefs consistent with the long-run data affected by the policy itself. This paper provides a framework for the analysis of such self-confirming, stationary policies. We first develop a general analysis of self-confirming policies in recurrent decision problems with incomplete information about the true stochastic model. Next we apply and illustrate the theory with a characterization of stationary monetary policies in a linear-quadratic setting.

Consider an either moderately patient or impatient agent (she) who makes recurrent decisions under uncertainty. In each period she takes an action a that, via a feedback function f , delivers an observable outcome, or message, $m = f(a, s)$ that depends on an unobservable state of nature s . A fixed, unknown stochastic model σ^* (that is, a probability measure over the states) determines an i.i.d. process of state realizations. The agent knows the feedback function f , but not the stochastic model σ^* . Note that, for some action a , the same outcome

*This is an extended version of the article forthcoming in *Theory and Decision*. We thank Manuel Arnese, Roberto Corrao, Carlo Cusumano, Enrico De Magistris, Nicolò Generoso, Giacomo Lanzani, Julien Manili, Fabrizio Panebianco, Elia Sartori, and the participants to conferences and seminars at Bocconi University, CISEI-2015 (Capri), D-TEA 2015 (HEC, Paris), LUISS (Rome), New York University, Scuola Normale Superiore (Pisa), and the University of Vienna for useful comments. Pierpaolo Battigalli and Massimo Marinacci gratefully acknowledge financial support from ERC grants 324219 and 670337.

m may result from multiple states, i.e., $f(a, \cdot)$ need not be injective; in this case, the realized outcome does not reveal the realized underlying state as exemplified below. There are no structural links between periods, but the agent observes the realized outcome in each period t and therefore updates her subjective belief μ_t about the fixed unknown model σ^* . Over time, given a true model σ^* and a prior belief, the intertemporal subjective expected utility maximizing strategy yields a convergent active learning process, i.e., a stochastic process of actions and updated beliefs (\mathbf{a}_t, μ_t) that converges almost surely.¹ The realization (a^*, μ^*) of the stochastic limit satisfies almost surely the following two properties:

- **Confirmed beliefs:** μ^* assigns probability 1 to the set of models σ that are observationally equivalent to the true model σ^* given action a^* ;²
- **Subjective best reply:** action a^* maximizes the agent’s one-period subjective expected utility given belief μ^* .

We take “confirmed beliefs” and “subjective best reply” to be the characterizing properties of stationary actions and beliefs. We call *self-confirming equilibrium* an action-belief pair (a^*, μ^*) with these properties. Indeed, conceptually this is a special case of the self-confirming equilibrium idea of Battigalli (1987) and Fudenberg and Levine (1993a), applied to one-person games with incomplete information about the probabilities of states.³ The key observation is that the confirmed belief μ^* need not assign probability 1 to the true stochastic model σ^* and, therefore, action a^* may differ from the objective best reply to σ^* . In other words, although equilibrium beliefs are disciplined by long-run empirical frequencies of observations, they do not necessarily concentrate on the true model σ^* , so the long-run action a^* may be objectively sub-optimal. This can happen even if the decision maker is quite patient: in the learning phase a positive discount factor can induce experimentation with actions that do not maximize one-period subjective expected utility, but the option value of experimentation vanishes in the limit.⁴

Consider the following heuristic example. A decision maker is asked to repeatedly bet on the color of a ball that will be drawn from an urn that contains 90 black, green, or yellow balls. After the draw, she is told whether she won (in which case she receives 1 euro) or not (in which case he receives 0 euros). Thus, there are three states, $S = \{B, G, Y\}$, three actions $A = \{b, g, y\}$, and two monetary outcomes $M = \{0, 1\}$. The feedback function attains value 1 when the action matches the state ($f(B, b) = f(G, g) = f(Y, y) = 1$) and value 0 otherwise. Thus, winning reveals the realized state, but losing only rules out one state out of three. Suppose the urn contains 45 black balls, 35 green balls, and 10 yellow balls, i.e., $\sigma^*(B) = \frac{1}{2}$, $\sigma^*(G) = \frac{7}{18}$, and $\sigma^*(Y) = \frac{1}{9}$. The objective best reply is to bet on B , but the decision maker does not know this. Suppose she keeps betting on G . With high objective probability, she is going to win more than $\frac{1}{3}$ of the times and she may come to

¹See, e.g., Easley and Kiefer (1988) and the references therein. We use boldface letters to distinguish random variables from their realizations. The working paper version of Battigalli et al. (2019) describes the exact relationship between our framework and the stochastic control framework of Easley and Kiefer, showing that they are essentially equivalent.

²That is, to models that induce the same distribution of outcomes as σ^* when a^* is chosen.

³The one-person assumption matters only for two related results of our abstract analysis. We expand on this in the literature review.

⁴See Section 4.2, where we also consider a kind of converse result.

deem very likely that the urn contains more green than black or yellow balls. Indeed, betting on G infinitely many times she will almost certainly observe that the winning frequency is $\frac{7}{18}$ and in the long run her limit belief μ_∞ will assign probability 1 to the set models $\{\sigma : \sigma(G) = \frac{7}{18}\}$. As long as her limit belief μ_∞ also assigns a sufficiently high probability to the set of models $\{\sigma : \sigma(G) = \frac{7}{18} > \max\{\sigma(B), \sigma(Y)\}\}$, she will find it optimal to keep betting on G , which is an objectively sub-optimal choice. Indeed, betting on G with such beliefs is a self-confirming equilibrium given the true model σ^* . Even if she initially experiments betting on B , sufficiently many unlucky outcomes will dissuade her from doing it again. In other words, the vagaries of the active learning process may well drive her in the trap of choosing forever the “satisficing” action g that wins more than $\frac{1}{3}$ of the times rather than to experiment with b long enough to realize that it yields an even higher winning frequency.

Thus, in a self-confirming equilibrium, the decision maker may be best-replying to empirically confirmed but wrong views about the actual data generating model. She may thus get trapped in self-confirming behavior that differs substantially from the objectively optimal behavior postulated by rational expectations models.⁵ This trap and the resulting welfare loss is, at the same time, especially relevant and disturbing for policy making. It is relevant when policy makers cannot obtain enough reliable evidence before choosing (e.g., with externally valid lab or field experiments), but instead have to rely on evidence that is a by-product of their actual policies; it is disturbing because welfare in self-confirming equilibria can be lower than in rational expectations equilibria. The main contribution of the present paper is to provide a formal steady-state framework in which this important policy issue can be rigorously studied. We then use this framework and illustrate the macroeconomic relevance of our analysis in the context of a 70’s U.S. policy debate about whether there is a trade-off between inflation and unemployment that can be systematically exploited by a benevolent policy maker.

Illustrative Application We consider a stylized model economy in which a policy maker chooses average inflation a and observes an unemployment/inflation outcome $(u, \pi) = f(a, s)$ that depends on the unobservable random state s of the economy. This model economy can be interpreted as reflecting an aggregate response function of a continuum of market agents. Assuming a quadratic loss function, we completely characterize the self-confirming equilibrium map that associates each conceivable model economy with a corresponding set of self-confirming beliefs and monetary policies. Given a fixed policy, the monetary authority infers from long-run data the first and second moments of the joint distribution of u and π , and hence the slope of the Phillips curve; but it cannot infer the true policy multiplier. We show that observing (in the long run) the distribution of (u, π) leaves the monetary authority with a residual one-dimensional uncertainty about the model economy, parameterized by the direct impact of policy on unemployment (i.e., neglecting the impact on u through π).

For example, even if the true model is a rational expectations augmented Phillips curve,

⁵In order to remove pervasive inconsistencies of pre-rational-expectations models, rational-expectations models often assume that decision makers know the true data generating process, thus making decisions objectively optimal. The traditional Nash equilibrium concept shares this objective best-reply feature. The Bayesian Nash equilibrium concept of Harsanyi (1967) instead allows for subjective and incorrect beliefs about parameters, without imposing a confirmed-beliefs condition. Therefore, Bayesian Nash equilibrium does not refine self-confirming equilibrium.

in equilibrium the monetary authority may believe that its policy does not shift the Phillips curve and hence that there is an exploitable trade-off given by the slope of the Phillips regression; the (Keynesian) monetary policy is optimal given a (falsely) conjectured trade-off, the subjectively expected unemployment rate coincides with the natural rate, and average inflation is (objectively) excessive. But we do not take a stand on what the true model is and so also consider self-confirming equilibria where the monetary authority pushes average inflation to zero falsely believing that there is no exploitable trade-off. Whatever the case, our analysis shows how partial identification may trap policy makers in inferior, yet self-confirming, policies that result in significant losses compared to the objectively optimal policies.

Manifesto Partial identification pervades economic policy debates: despite the use of sophisticated econometric techniques, economists disagree about how the economy works. Therefore, at least some economists must be wrong, but all of them should hold beliefs consistent with the data, which indeed only partially identify the relevant unknowns. The agents that inhabit our models—in particular, policy makers—are in a similar position, but their partial identification problem is exacerbated because what they can infer about the relevant unknowns depends on their own behavior, so it is endogenous. Thus, different policies justified by different beliefs—so, ultimately, by different (possibly conflicting) economic views—may be self-confirming. Such beliefs may even be dogmatic, for example because they assign probability one to a parameter vector resulting from observed long-run frequencies and untested, possibly false, identifying hypotheses.

To escape the partial identification trap more experimentation may be advisable. But we do not see an easy way out: large-scale social experiments can have huge costs, captured in our framework by the opportunity cost of not using a one-period subjective best reply, while small-scale ones may have little external validity.

Roadmap As anticipated, the first part of this paper (Sections 2-4) develops an abstract analysis of self-confirming choices. The general contribution of this part is to provide a theoretical framework that is:

- *broad enough* to include the finite one-period setting in which self-confirming analysis was originally developed within game theory as well as the infinite setup relevant for many economic applications, including macroeconomic policy analysis;
- *specific enough* to provide welfare implications for relevant policy questions with the backdrop of a neat learning foundation.

The main issues that we address at this abstract level concern the scope of partial identification (Section 3), equilibrium values, and the resulting effects on the decision maker's welfare (Section 4.3). The most novel results of the abstract part concern the latter topic. We show that if among two policies justified by different self-confirming beliefs one allows better identification of the true model, even if only partially, then this policy yields higher welfare. Similarly, self-confirming equilibria justified by sharper beliefs yield higher welfare. The theoretical concepts of this first part are illustrated and clarified by a running example of a monopolist facing an uncertain demand.

The second part of the paper (Sections 5 and 6) builds on the abstract analysis to gain a better perspective on economic policy. Section 5 frames the classical debate on the possibility of systematically exploiting unemployment/inflation trade-offs and provides novel results. In particular, the scope of partial identification is characterized in Section 5.2, while equilibria, their values, and the welfare effects of model uncertainty are analyzed in Section 5.3. Sections 5.4 and 5.5 illustrate the analysis by considering two important special cases. Section 7 offers some concluding remarks. Finally, Section 6 sketches a more general analysis of self-confirming economic policies.

Appendix A collects some more technical material and all the formal proofs.

Related literature Our analysis contributes to and provides a bridge between two strands of literature, one in game theory and the other in macroeconomics, that are concerned with related issues, but have so far proceeded with limited cross fertilization and very different languages.

In the game-theoretic literature, a strategy profile that satisfies the properties of confirmed beliefs and subjective best reply has been called “conjectural equilibrium” (Battigalli, 1987, Battigalli and Guaitoli, 1988), “self-confirming equilibrium” (Fudenberg and Levine, 1993a,b) and “subjective equilibrium” (Kalai and Lehrer, 1993, 1995). Here we adopt the more self-explanatory terminology of Fudenberg and Levine. We refer the reader to Battigalli et al. (2015) for an up-to-date discussion of this literature. Here we point out that, although we focus on one-person decision problems with uncertainty, our abstract analysis extends seamlessly to n -person games except for the aforementioned comparative results about equilibrium values. To our knowledge, papers in the extant literature either consider finite (one-period) games, or games with no randomness. We extend the analysis of self-confirming equilibria to settings with inherent randomness and possibly infinite spaces of strategies and states of nature. Technically, this extension is not straightforward, it requires mathematical precision and care. We also point out that the learning foundation of the equilibrium concept is more solid in the one-person case analyzed here: while self-confirming equilibria of multi-person games represent the steady states of learning dynamics, convergence to a steady state is not guaranteed under general conditions, as is instead the case when the model can be effectively reduced to a one-person decision problem (see Section 4.2). Finally, note that Battigalli et al. (2015) is focused on the interaction between ambiguity aversion and self-confirming equilibria in games. Here instead we consider a decision maker who maximizes her subjective expected utility (i.e., she is ambiguity neutral). This simplifies the general analysis without affecting the illustrative examples and the application. Indeed, they feature conditions under which the degree of ambiguity aversion does not affect the set of self-confirming equilibria (see Section 7 and Battigalli et al. 2021), although it may well affect learning dynamics and the likelihood to be trapped in the long run in self-confirming equilibria with an objectively suboptimal choice (see Battigalli et al. 2019).

The macroeconomic literature focuses on policy making and learning dynamics. Sargent (1999) explains the rise and fall of US inflation assuming that the monetary authority sequentially estimates a Phillips curve, ignoring its impact on expectations, and best replies to updated beliefs. Standard OLS estimation leads to a Keynesian self-confirming equilibrium, but if instead recent observations are given more weight, because the monetary policy maker’s decisions make the Phillips curve slowly shift and rotate over time, the process first approaches a neighborhood of this equilibrium, but then recurrently abandons it when the

Phillips curve looks “more vertical” leading the monetary policy maker to lower inflation.⁶ Cho et al. (2002) and Sargent and Williams (2005) sharpen the theoretical analysis of such learning dynamics.⁷ Cho and Kasa (2015) note that the low inflation outcome at the end of Sargent’s (1999) narrative —according to the postulated learning model— cannot persist either; therefore, they consider an alternative stochastic learning dynamic in which the policy maker best responds to the current estimate of an aggregate supply model, out of a set of conceivable functional forms, as long as the model passes a statistical test; when the model is rejected, a new model is selected at random and the process is restarted. Also, in their model the Keynesian self-confirming equilibrium cannot persist, because, in the very long-run, the monetary authority adopts a vertical Phillips curve model.⁸ In our paper, we focus only on the set of possible limit points of learning dynamics. Furthermore, in our monetary policy application, we follow Sargent (2008) and assume that the monetary authority may believe in the exploitability of a trade-off between unemployment and inflation. Unlike the papers we have mentioned, we do not take a stand on a true model economy. Thus, instead of assuming that the true model economy features a rational-expectations augmented Phillips curve, we characterize the self-confirming equilibria and values for many conceivable models.

Other papers in the literature focus, like ours, mainly on self-confirming equilibrium policies rather than learning dynamics. In particular, Battigalli and Guaitoli (1988) analyze the self-confirming equilibria with rationalizable beliefs of a stylized policy game with incomplete information, showing that there are equilibria with Keynesian features and equilibria with new-classical features. Fudenberg and Levine (2009) discuss the Lucas critique through the analysis of refined self-confirming equilibria in some insightful illustrative examples; they emphasize the role of rationalizable beliefs and of robustness to experimentation. Unlike the foregoing papers, we formally analyze a one-agent framework, which makes the issue of the rationalizability of beliefs mute. According to the application, when the one-agent framework is interpreted as a reduced form of a multi-person game, the shape of outcome/feedback function f may implicitly represent such rationalizability constraints, e.g., the decision maker is a leader and the outcome function captures the best-reply behavior of followers; this is clarified by our monopoly example. As for the monetary policy application, only a genuinely game-theoretic model of the economy would allow a thorough analysis of the rationalizability of self-confirming beliefs, but tackling such difficult issue is beyond the scope of this article. In a series of papers, Saint Paul (e.g., 2013, 2018) considers an expert who knows the true model and advises the policy maker while pursuing her own policy agenda; the policy maker and the agents in the market fully trust the expert as long as the data are consistent with her advice. With this, the expert manipulates the policy maker and market agents under a

⁶See also, Cogley and Sargent (2005), Sargent et al. (2006), and Cogley et al. (2007).

⁷The phrase “escaping Nash inflation” in the title of Cho et al. (2002) deserves an explanation. When the decision model is interpreted as a game between the monetary authority and a representative agent, a self-confirming equilibrium outcome is also a (possibly subgame imperfect) Nash equilibrium outcome. Battigalli (1987) and Fudenberg and Levine (1993) provide sufficient conditions for the realization-equivalence between Nash and self-confirming equilibrium. Such conditions are satisfied in the model of Cho et al.

⁸In his work on rational belief equilibria, Kurz (1994a,b) analyzes stochastic dynamics where agents’ beliefs may be incorrect, but are eventually consistent with the long-run frequencies of observables, which is in the spirit of self-confirming equilibrium. The most important difference with the literature on the latter is that, although Kurz analyzes multi-agent systems, he does not use a game theoretic framework. Specifically, unlike game models, there is no function specifying how agents’ actions and—possibly—exogenous variables determine outcomes and observables.

self-confirmation constraint. Finally, Gaballo and Marimon (2021) analyze a directed search model of the credit market where lenders post excessively high interest rates because of confirmed pessimistic beliefs about returns on investments, but the monetary authority can break the spell by easing credit. The main difference with our monetary policy application is that we study the self-confirming actions and beliefs of the monetary authority, not of the agents in the market.⁹

To the best of our knowledge, besides the novelty of several results, our paper is unique in integrating an abstract analysis of self-confirming policies with an economic application. There are under-appreciated complementarities between abstract theory and applications. The former allows to focus on key concepts and properties uncluttered by specific modeling features, the latter helps to better understand the abstract theory and points to its relevance. Here we consider a monetary policy application, but the scope of our analysis goes well beyond that. For example, the difficulty of thorough experimentation and its consequences for welfare naturally arise in the context of environmental policies.

2 Preliminaries

2.1 Mathematics

Differently from Battigalli et al. (2015), the Phillips curve exploitation model that motivates and illustrates this paper features infinite action, state, and consequence spaces as well as unbounded payoff functions. The necessary adaptation is conceptually natural, but technically nontrivial. In particular, it requires that the analysis be carried out within a *standard Borel space* (X, \mathcal{X}) , where X is a completely metrizable and separable topological space and \mathcal{X} is its Borel sigma algebra. The Borel sets $B \in \mathcal{X}$ are themselves standard Borel spaces under the relative sigma algebra $\mathcal{X} \cap B$.¹⁰ When X is countable (i.e., finite or denumerable), standardness requires \mathcal{X} to be the power set of X (see Appendix A.1).

We denote by $\Delta(X)$ the collection of all probability measures on \mathcal{X} , endowed with the natural sigma algebra,¹¹ which in turn makes $\Delta(X)$ a standard Borel space too. With this, the Borel subsets Σ of $\Delta(X)$ with their relative sigma algebras are standard Borel spaces themselves. The meaning of $\Delta(\Sigma)$ is then obvious. Finally, $\delta : X \rightarrow \Delta(X)$ denotes the canonical Dirac embedding of X into $\Delta(X)$, that is, $\delta(x)$ is the probability measure on \mathcal{X} which assigns probability 1 to each Borel set containing $x \in X$.¹²

Let (Y, \mathcal{Y}) be another standard Borel space. The Cartesian product $X \times Y$ is a standard Borel space with respect to the product sigma algebra. Moreover, each measurable function $\varphi : X \rightarrow Y$ induces a measurable distribution map $\hat{\varphi} : \Delta(X) \rightarrow \Delta(Y)$ defined by

$$\hat{\varphi}(\xi) = \xi \circ \varphi^{-1}$$

for each probability measure $\xi \in \Delta(X)$. That is, $\hat{\varphi}(\xi)(B) = \xi(\varphi^{-1}(B))$ for all sets B in \mathcal{Y} .¹³

⁹The model is not explicitly represented as a game. Therefore the connection to the traditional self-confirming equilibrium concept is not immediate.

¹⁰See Kechris (2012) for the properties of standard Borel spaces that we use.

¹¹That is, the sigma algebra generated by the evaluation maps $\xi \mapsto \xi(B)$ for all $B \in \mathcal{X}$.

¹²The usual notation for the Dirac measure concentrated on x is δ_x .

¹³See Appendix A.1. In the applied probability literature, $\hat{\varphi}(\xi)(B)$ is sometimes denoted by $\hat{\varphi}(B | \xi)$, interpreted as the probability of observing a realization in B given ξ with “measurement” φ .

Lemma 1 *Let $\varphi : X \rightarrow Y$ be a measurable function. The following conditions are equivalent:*

- (i) φ is one-to-one,
- (ii) $\hat{\varphi}$ is one-to-one,
- (iii) $\varphi^{-1}(\mathcal{Y}) = \mathcal{X}$.

Interpreting x as a state and $y = \varphi(x)$ as an observable outcome, we can phrase these equivalent conditions as follows:

- (i) φ reveals the state x in X ,
- (ii) $\hat{\varphi}$ reveals the distribution ξ in $\Delta(X)$,
- (iii) φ generates the sigma-algebra \mathcal{X} .

Finally, we say that X and Y are *isomorphic*, written $X \cong Y$, if there is a bimeasurable bijection $\varphi : X \rightarrow Y$, that is, φ is measurable and $\varphi^{-1} : Y \rightarrow X$ is a well defined measurable function.

2.2 Classical subjective expected utility

Let S be a space of *states* of nature, A a space of *actions* available to the decision maker, C a space of *consequences*, and $\rho : A \times S \rightarrow C$ a measurable *consequence function* that associates a consequence $\rho(a, s) \in C$ with each pair $(a, s) \in A \times S$ of action and state. When consequences are monetary, C is a (Borel) subset of the real line.

The quartet

$$(A, S, C, \rho) \tag{1}$$

is the basic structure of the decision problem. The inherent randomness characterizing the realization of states —often called physical uncertainty— is described by probability models $\sigma \in \Delta(S)$ that can be regarded as possible generative mechanisms. For each probability model σ , actions a are evaluated through their expected utility

$$\int_S v(\rho(a, s)) d\sigma(s)$$

where $v : C \rightarrow \mathbb{R}$ is a measurable and bounded above von Neumann-Morgenstern utility function. It is often convenient to write the criterion in the expected-payoff form

$$R(a, \sigma) = \int_S r(a, s) d\sigma(s)$$

where $r : A \times S \rightarrow \mathbb{R}$ is the *payoff* (or *reward*) *function* $r = v \circ \rho$. Also the payoff function is easily seen to be measurable and bounded above. All our integrals are thus well defined, but may take value $-\infty$.

The decision maker may not know the true probability model σ^* but is able to posit a (measurable) collection $\Sigma \subseteq \Delta(S)$ of probability models that contains the true one; that is, $\sigma^* \in \Sigma$. We thus abstract from misspecification issues. We call *structural* the kind of

information that allows the decision maker to posit the collection Σ . For example, if the problem is to bet on the color, white or black, of a ball drawn from a two-color urn, and it is only known that the urn contains n balls, then Σ has $n + 1$ elements and is isomorphic to the set $\{0, \frac{1}{n}, \dots, \frac{n-1}{n}, 1\}$ of possible fractions of white balls. When Σ is a singleton, i.e., the true model is known, the decision maker confronts only *risk*. Otherwise, she faces *model uncertainty*.¹⁴ We can also give Σ a somewhat different interpretation: it represents a *backdrop theory* accepted by the decision maker, which happens to be correct (i.e., such that $\sigma^* \in \Sigma$). In particular, as we assume that the same decision problem is faced infinitely often, representing uncertainty with Σ rests on the assumption that the process of states is i.i.d.¹⁵

The decision maker ranks actions according to the *classical subjective expected utility* (SEU) criterion:¹⁶

$$V(a, \mu) = \int_{\Sigma} R(a, \sigma) d\mu(\sigma), \quad (2)$$

where $\mu \in \Delta(\Sigma)$ is a subjective *prior probability* over models that reflects personal beliefs about models that the decision maker may have, in addition to the structural information behind Σ .¹⁷ This representation admits the reduced form

$$\int_{\Sigma} R(a, \sigma) d\mu(\sigma) = \int_S r(a, s) d\sigma_{\mu}(s) = R(a, \sigma_{\mu}),$$

where $\sigma_{\mu} \in \Delta(S)$ is the subjective *predictive probability* defined by $\sigma_{\mu}(E) = \int_{\Sigma} \sigma(E) d\mu(\sigma)$ for each $E \in \mathcal{S}$. This reduced form is the original representation of Savage (1954), who derived σ_{μ} from preferences over bets.

The *decision problem* can then be summarized by the sextet

$$D = (A, S, C, \rho, \Sigma, v) \quad (3)$$

that combines the basic structure (1) with the information and taste traits Σ and v . A few special cases are noteworthy.

- (i) When the support of μ is a singleton $\{\sigma\}$, that is, $\mu = \delta(\sigma)$, the decision maker believes (maybe wrongly) that σ is the true model. The predictive probability trivially coincides with σ and criterion (2) reduces to the Savage expected payoff criterion $R(a, \sigma)$. Being a predictive probability, σ here is a subjective probability measure, albeit one derived from a dogmatic belief.
- (ii) When Σ is a singleton $\{\sigma^*\}$, the decision maker has maximal structural information and, as a result, knows that σ^* is the true model. In this case, there is only physical uncertainty, quantified by σ^* , without any model uncertainty. Criterion (2) again reduces

¹⁴Model uncertainty is also called model ambiguity (see Hansen and Marinacci, 2016).

¹⁵If, instead, the process of states were assumed to be Markovian, probability models would be kernel functions (with finitely many states, transition probability matrices) rather than elements of $\Delta(S)$.

¹⁶The integral is well defined, for each Σ and each $\mu \in \Delta(\Sigma)$, because the expected payoff function $R(a, \cdot) : \Delta(S) \rightarrow [-\infty, \infty)$ is, for every action a , measurable and bounded above on $\Delta(S)$ and hence on Σ .

¹⁷See Marinacci (2015) for a discussion of this setup. Classical SEU is proposed by Cerreia-Vioglio et al. (2013), where “classical” refers to the fact that a posited set of probability models is a basic feature of classical statistics.

to the expected payoff criterion $R(a, \sigma^*)$, but now interpreted as a von Neumann-Morgenstern criterion. For instance, if the decision maker either observed infinitely many draws from a given urn or were just able to count the balls of each color, she would learn/know the urn composition and Σ would be a singleton.

- (iii) When $\Sigma \subseteq \{\delta(s) : s \in S\}$, there is no physical uncertainty, but only model uncertainty, quantified by μ . We can identify prior and predictive probabilities: with a slight abuse of notation, we can write $\mu \in \Delta(S)$ and so (2) takes the form $R(a, \mu \circ \delta)$.¹⁸

Throughout this part (Sections 2-4), we illustrate the abstract theoretical concepts with a stripped-down monopoly example.

Example 1 (Monopoly: Unknown demand) *A monopolist choosing output $a \geq \underline{a}$ faces an imperfectly known (state-dependent) inverse demand function $a \mapsto P(a, s)$. We interpret the lower bound $\underline{a} \geq 0$, when strictly positive, as a pre-commitment to a minimum level of production. If $\underline{a} = 0$ there is no pre-commitment. The firm knows the slope, but not the intercept, which has a permanent component θ modified by an additive noise ε :*

$$P(a, s) = \max\{0, s - a\} = \max\{0, \theta + \varepsilon - a\},$$

where $s \in S = [\underline{s}, \bar{s}] \subseteq \mathbb{R}_{++}$, $\theta \in [\underline{\theta}, \bar{\theta}] = [\underline{s} + \bar{\varepsilon}, \bar{s} - \bar{\varepsilon}]$, and $\varepsilon \in [-\bar{\varepsilon}, \bar{\varepsilon}]$ is the realization of a random variable ε with known distribution η and 0 mean.¹⁹ The firm has a known linear cost function, with average and marginal cost $c > 0$. To further simplify the analysis, we assume that price is certainly strictly positive on the relevant range of outputs, i.e., also for the largest subjective best reply across all possible beliefs. This is the case if²⁰

$$\underline{\theta} - \frac{3}{2}\bar{\varepsilon} - \frac{\bar{\theta} - c}{2} > 0.$$

With this, we can ignore the 0-price floor, and the relevant inverse demand map becomes $a \mapsto (\theta + \varepsilon - a)$. For each $\theta \in [\underline{\theta}, \bar{\theta}]$, let $T_\theta : [-\bar{\varepsilon}, \bar{\varepsilon}] \rightarrow S$ denote the translation map $\varepsilon \mapsto (\theta + \varepsilon)$. We can parameterize Σ as follows:²¹

$$\Sigma = \{\sigma \in \Delta(S) : \exists \theta \in [\underline{\theta}, \bar{\theta}], \sigma = \eta \circ T_\theta^{-1}\} \cong [\underline{\theta}, \bar{\theta}].$$

The consequence function (again, in the relevant range of outputs) is the profit function

$$\rho(a, s) = a(s - a - c).$$

Under risk neutrality, v is the identity on the range of ρ ; thus, $r = \rho$. Given the parameterization of Σ , the objective expected payoff and subjective expected utility can be written as

$$R(a, \theta) = \mathbb{E}(a(\theta + \varepsilon - a - c)) = a(\theta - a - c)$$

¹⁸See Corollary 4 in Appendix A.1.

¹⁹We write random variables in boldface font and their realization in normal font.

²⁰Ignoring the 0-price floor, $(\bar{\theta} - c)/2$ is the best reply to the most optimistic belief. The condition guarantees that, at this largest output, price is strictly positive even with the lowest inverse demand function. This implies that the standard first-order conditions identify a global subjective optimum.

²¹Recall that $\Sigma \cong [\underline{\theta}, \bar{\theta}]$ means that Σ and $[\underline{\theta}, \bar{\theta}]$ are isomorphic.

and

$$V(a, \mu) = \int_{\underline{\theta}}^{\bar{\theta}} a(\theta - a - c) d\mu(\theta) = (\mathbb{E}_{\mu}(\theta) - c)a - a^2.$$

We obtain case (i) if the firm is certain of θ , case (ii) if it knows θ^* , and case (iii) if there is no noise, i.e., $\bar{\varepsilon} = 0$ and $S = [\underline{\theta}, \bar{\theta}]$.

The example clarifies that decision problem D could be the reduced form of a multi-agent model where the unknown state s represents the behavior of other agents, such as buyers. Such behavior is unaffected by choice a , either literally, or because it represents a profile of strategies (decision functions) rather than actual actions. In the quantity-setting monopoly, s may be determined by a distribution of private valuations, with output a sold in a multi-unit uniform price auction and with unit-demand buyers bidding their valuations as their dominant bid. For a price-setting monopolist valuations determine individual demand functions as optimal reactions to the set price. In these cases, the map $a \mapsto \rho_s(a)$ is determined by rational behavior of the un-modeled agents. In an alternative interpretation, the firm is a monopolistic competitor of negligible size and s represents general market conditions.

3 Partial identification

3.1 Feedback

The decision maker faces decision problem D recurrently in a stationary environment with an i.i.d. process of states determined by unknown probability model σ^* . To determine what actions and beliefs can be stable given σ^* , we have to specify the information obtained ex post by the decision maker for each action a and state s . We model such information through a (measurable) *feedback function*

$$f : A \times S \rightarrow M$$

where M is a space of messages. By selecting an action $a \in A$, the decision maker receives a *message*

$$m = f_a(s)$$

when s occurs.²² The decision maker's (ex post) information about the state is thus endogenous. When M is finite, such endogenous information is represented by the partition $\{f_a^{-1}(m) : m \in M\}$ of the state space S that the messages induce, which depends on the choice of action a . This partition generates the algebra of events whose probability can be inferred from the long-run frequencies of messages. When M is infinite, it may be the case that this collection of events cannot be recovered from the partition. Hence, it is technically convenient to represent information with the sigma algebra

$$\mathcal{F}_a = \{f_a^{-1}(B) : B \in \mathcal{M}\}.$$

A *decision problem with feedback* is described by the octet

$$(A, S, C, \rho, \Sigma, v, M, f) \tag{4}$$

²²Here $f_a : S \rightarrow M$ denotes the section $f(a, \cdot)$ of f at a .

where a feedback function and a message space are added to the decision problem (3).

When information does not depend on action a , we say that there is *own-action independence* of feedback about the state; formally, $\mathcal{F}_a = \mathcal{F}_{a'}$ for all $a, a' \in A$. The most important instance of own-action independence is *perfect feedback*, which occurs when each section f_a of the feedback function f generates \mathcal{S} —that is, in view of Lemma 1, when f_a is one-to-one for each $a \in A$. In this case, messages reveal to the decision maker which state obtained, regardless of the chosen action. When this is not the case, feedback about the state is *imperfect*, maximally so when each section f_a is constant, so that $\mathcal{F}_a = \{\emptyset, \mathcal{S}\}$ and all states return the same message.

An action a is *fully revealing* if f_a is one-to-one, that is, if it allows the decision maker to learn which state obtained. Under perfect feedback, all actions are fully revealing. The existence of fully revealing actions is a weak form of “endogenous” perfect feedback.

We assume throughout that *consequences are observable*. Formally, this amounts to assuming that, for each action $a \in A$, the section ρ_a of the consequence function ρ is \mathcal{F}_a -measurable. The next result, which will play an important role in our analysis, characterizes this assumption within a decision problem with feedback (4).

Proposition 1 *Consequences are observable if and only if, for each action $a \in A$, there exists a measurable function $g_a : M \rightarrow C$ such that*

$$\rho_a = g_a \circ f_a.$$

In this case, the payoff $r_a = v \circ \rho_a$ of each action a is \mathcal{F}_a -measurable.

In words, messages encode consequences and so payoffs. In particular, when the consequences of the actions are the only observed messages, we have $C = M$ and $f = \rho$. This is the most common and important case of feedback, which is also featured by our macroeconomic application.

Example 2 (Monopoly: Feedback) *A natural assumption about feedback for the quantity-setting monopoly of Example 1 is that the firm observes the realized market price, that is (for the relevant range of outputs)*

$$f(a, s) = P(a, s) = s - a$$

and g_a is the affine map $p \mapsto a(p - c)$ from market price to profit. Note, however, that if the firm has a zero lower bound $\underline{a} = 0$, assuming that a realized price can be observed even with 0 output to sell is contrived; indeed, the most plausible assumption is that nothing is observed. The same observability pattern occurs with an alternative assumption about feedback: the firm only observes its revenue, $f(a, s) = \rho(a, s) + ca$, e.g., because it is the grower of a unique variety of weed sold to a dealer who returns the proceeds from an auction. With 0 production, nothing can be observed, with positive production $a > 0$, the unit price and realized state can be backed out: $p = \rho/a + c$ (per-unit revenue) and $s = p + a$. To sum up, both feedback functions satisfy observability of consequences, and each interior choice ($a > 0$) is revealing. Thus, own-action independence of feedback about the state holds if there is a strictly positive lower bound on production $\underline{a} > 0$. Absent this constraint ($\underline{a} = 0$), own-action independence of feedback about the state fails because 0-output reveals nothing, while positive output is revealing.

3.2 Partial identification correspondence

In our steady state setting, a message distribution $\nu \in \Delta(M)$ can be interpreted as a long-run empirical frequency of messages received by the decision maker. Specifically, for each Borel set $B \in \mathcal{M}$, $\nu(B)$ is the long-run empirical frequency with which messages m belong to B . For any action $a \in A$, consider the distribution map $\hat{f}_a : \Sigma \rightarrow \Delta(M)$ defined by $\hat{f}_a(\sigma) = \sigma \circ f_a^{-1}$. That is,

$$\hat{f}_a(\sigma)(B) = \sigma(s \in S : f_a(s) \in B)$$

for each $B \in \mathcal{M}$.²³ Then $\hat{f}_a(\sigma)(B)$ is the long-run empirical frequency with which the decision maker receives messages m in B , when action a is chosen and σ is the true model. The inverse correspondence \hat{f}_a^{-1} from $\Delta(M)$ to $\Delta(S)$ partitions the latter set into classes

$$\hat{f}_a^{-1}(\nu) = \left\{ \sigma \in \Delta(S) : \hat{f}_a(\sigma) = \nu \right\}$$

of models that are observationally equivalent given that action a is chosen infinitely often and that the frequency distribution of messages ν is observed in the long-run conditional on a . In other words, $\hat{f}_a^{-1}(\nu)$ is the collection of all probability models that may have generated ν given a .

If action a is fully revealing, then \hat{f}_a is one-to-one and so $\hat{f}_a^{-1}(\nu)$ is at most a singleton for every ν . In this case the decision problem is identified under a since different models generate different message distributions, which thus uniquely pin down models. Otherwise, when $\hat{f}_a^{-1}(\nu)$ is nonsingleton for some ν , we have *partial identification* under action a . In the extreme case when \hat{f}_a is constant—that is, when all models generate the same message distribution—the decision problem is completely unidentified under action a . Interestingly, \hat{f}_a is constant if and only if f_a is constant, that is, all states generate the same message (see Lemma 7 in Appendix A.1).

Now recall that the decision maker posits a set of models Σ determined by structural information or a backdrop theory. Upon observing ν , one can conclude that the data generating model belongs to

$$\underbrace{\hat{f}_a^{-1}(\nu)}_{\text{data}} \cap \underbrace{\Sigma}_{\text{theory}}.$$

For this reason, models σ and σ' in Σ such that $\hat{f}_a(\sigma) = \hat{f}_a(\sigma') = \nu$ are observationally equivalent under action a . Formally, given an action a , two models $\sigma, \sigma' \in \Sigma$ are *observationally equivalent* if

$$\hat{f}_a(\sigma) = \hat{f}_a(\sigma').$$

We denote the class of models observationally equivalent to σ given a by

$$\hat{\Sigma}_a(\sigma) = \left\{ \sigma' \in \Sigma : \hat{f}_a(\sigma') = \hat{f}_a(\sigma) \right\} = \hat{f}_a^{-1}(\hat{f}_a(\sigma)) \cap \Sigma. \quad (5)$$

In other words, $\hat{\Sigma}_a(\sigma)$ is the *partially identified set* of models given action a .²⁴ We can thus regard the map $\hat{\Sigma}_a(\cdot) : \Sigma \rightrightarrows \Sigma$, which associates to each element of Σ its observational equivalence class, as the *partial identification correspondence* determined by action a .

²³As previously observed, this is sometimes denoted $\hat{f}_a(B | \sigma)$.

²⁴We can also write $\hat{\Sigma}_a(\sigma) = \{\sigma' \in \Sigma : \sigma'_{|\mathcal{F}_a} = \sigma_{|\mathcal{F}_a}\}$, i.e., partial identification is determined by the information sigma algebra \mathcal{F}_a .

It is easy to see that $\hat{\Sigma}_a$ has convex values if the collection Σ is convex. Moreover, if \hat{f}_a is one-to-one, then $\hat{\Sigma}_a$ is the identity: $\hat{\Sigma}_a(\sigma) = \{\sigma\}$ for all $\sigma \in \Sigma$. In this case, message distributions identify the true model. In contrast, when $\hat{\Sigma}_a(\sigma)$ is nonsingleton there is genuine partial identification.

Summing up, the collection $\{\hat{\Sigma}_a(\sigma)\}_{\sigma \in \Sigma}$ is a measurable partition of Σ and its cells consist of probability models that are observationally equivalent under action a . Clearly, the dependence on a is lost under own-action independence of feedback about the state.

Example 3 (Monopoly: Partial Identification) *Consider the quantity-setting monopoly of Example 1. If feedback is the realized price and output $a > 0$ is chosen infinitely often, the firm observes in the long-run the average price $\mathbb{E}(\mathbf{p}_a) = \theta - a$, and θ is identified: $\theta = \mathbb{E}(\mathbf{p}_a) + a$. With a 0-lower bound ($\underline{a} = 0$), producing 0 instead reveals nothing. Thus, given the parameterization of Σ , the partial identification correspondence is*

$$\hat{\Sigma}_a(\sigma(\theta)) \cong \hat{\Theta}_a(\theta) = \begin{cases} \{\theta\} & \text{if } a > 0, \\ [\underline{\theta}, \theta] & \text{if } a = 0. \end{cases}$$

3.3 Comparative statics

The extent of partial identification depends, intuitively, on how informative is the underlying feedback function. To formalize this intuition, we need to compare feedback functions according to their informativeness. To this end, we say that a feedback function f' is *coarser* (or less *fine*) than a feedback function f if, for each $a \in A$, f'_a is \mathcal{F}_a -measurable or, equivalently, if there exists a measurable function $h_a : M \rightarrow M'$ such that

$$f'_a = h_a \circ f_a.$$

In the monopoly example, with a 0-lower bound ($\underline{a} = 0$) and under the assumption that a notional realized price could be observed even at 0 output, realized revenue/profit $f'_a(s) = \rho_a(s)$ is a coarser feedback than realized price $f_a(s) = P_a(s)$, with $h_a(p) = ap$.

A coarser feedback function is less informative. Using this comparative notion, we show that a less informative feedback function aggravates the decision maker's partial identification problem, thus formalizing the previous intuition. Given feedback functions f and f' , we let $\hat{\Sigma}_a(\cdot)$ and $\hat{\Sigma}'_a(\cdot)$ denote the identification correspondences respectively derived from f and f' .

Proposition 2 *Fix feedback functions f and f' . If f' is coarser than f , then $\hat{\Sigma}_a(\sigma) \subseteq \hat{\Sigma}'_a(\sigma)$ for all $a \in A$ and $\sigma \in \Sigma$.*

Coarser feedback functions thus determine, for each action, coarser observational equivalence relations: a worse information translates into a lower degree of statistical identification. In particular, the assumption that consequences are observable makes the consequence function ρ the coarsest possible feedback. Perfect feedback is, instead, the finest.

4 Self-confirming actions and beliefs

4.1 Definition

Throughout this section, we fix a decision problem

$$(A, S, C, \rho, \Sigma, v, M, f)$$

with feedback and observable consequences, where Σ contains the true model σ^* that generates the states.

With this, we introduce a concept that is at the heart of our analysis and is motivated by the partial identification issues discussed in the previous section.

Definition 1 *An action-belief pair $(a^*, \mu^*) \in A \times \Delta(\Sigma)$ is a self-confirming equilibrium given σ^* if*

$$\forall a \in A, \quad V(a^*, \mu^*) \geq V(a, \mu^*) \quad (6)$$

and

$$\mu^* \in \Delta(\hat{\Sigma}_{a^*}(\sigma^*)). \quad (7)$$

The definition relies on two pillars: the optimality condition (6) that ensures that action a^* is subjectively optimal under belief μ^* , and the belief confirmation condition (7) that guarantees that belief μ^* is consistent with the data that action a^* reveals in the long run.²⁵ In fact, given model σ^* , action a^* determines the message distribution $\nu^* = \hat{f}_{a^*}(\sigma^*)$, which is the long-run evidence that disciplines the subjective belief μ^* . In this respect, note that

$$\hat{\Sigma}_{a^*}(\sigma^*) = \hat{f}_{a^*}^{-1}(\hat{f}_{a^*}(\sigma^*)) \cap \Sigma = \hat{f}_{a^*}^{-1}(\nu^*) \cap \Sigma.$$

Therefore, $\hat{\Sigma}_{a^*}(\sigma^*)$ depends only on the induced message distribution ν^* .

Note also that condition (7) makes self-confirming equilibrium for decision problems with feedback a *genuine* equilibrium concept. Indeed, we already mentioned in the Introduction that it characterizes the steady states of learning dynamics in stochastic control problems. Relatedly, it is a fixed-point concept: suppose for simplicity that there is a unique best reply $B(\mu)$ for each belief μ ; then, a self-confirming belief is a *fixed point* of the correspondence

$$\mu \mapsto \Delta(\hat{\Sigma}_{B(\mu)}(\sigma^*)).$$

Finally, it is worth noting that a self-confirming belief may exclude the true model.²⁶ We can indeed formulate the data confirmation condition (7) as follows:

$$\underbrace{\mu^*}_{\text{belief}} \left(\underbrace{\hat{f}_{a^*}^{-1}(\nu^*) \cap \Sigma}_{\text{data}} \right) = 1. \quad (8)$$

The equilibrium belief must thus exclude everything which is not consistent with either observations or structural information/theory, that is,

$$\underbrace{\mu^*}_{\text{belief}} \left(\underbrace{\hat{f}_{a^*}^{-1}(\nu^*)^c}_{\text{contra data}} \cup \underbrace{\Sigma^c}_{\text{contra theory}} \right) = 0, \quad (9)$$

but it may exclude other models as well, including the true one.

Under own-action independence of feedback about the state, the data confirmation condition (7) becomes $\mu^* \in \Delta(\hat{\Sigma}(\sigma^*))$. We thus return to a traditional optimization notion with

²⁵Here, since $\hat{\Sigma}_{a^*}(\sigma^*)$ is a measurable subset of Σ , the set $\Delta(\hat{\Sigma}_{a^*}(\sigma^*))$ is identified with the family of elements of $\Delta(\Sigma)$ that assign probability 1 to $\hat{\Sigma}_{a^*}(\sigma^*)$.

²⁶See the example in the introduction.

a purely exogenous data confirmation condition. In particular, under perfect feedback —and so full identification— the optimality condition (6) becomes

$$\forall a \in A, \quad R(a^*, \sigma^*) \geq R(a, \sigma^*), \quad (10)$$

since condition (7) requires $\mu^* = \delta(\sigma^*)$. In this case, common in the rational expectations literature, the decision maker has a correct belief about the true model and confronts only risk.

We say that an action $a^* \in A$ is *objectively optimal* given σ^* if it satisfies the optimality condition (10). Objectively optimal actions are the ones that the decision maker would select if she knew the true model, that is, under full identification. As such, they provide an important benchmark to assess alternative courses of action, as the next welfare analysis will show.

That said, observe that a “rational-expectations” pair $(a^*, \delta(\sigma^*))$, where action a^* is objectively optimal and belief $\delta(\sigma^*)$ is concentrated on the true model, is a self-confirming equilibrium. Indeed, $\sigma^* \in \hat{\Sigma}_{a^*}(\sigma^*)$ and so $\delta(\sigma^*)(\hat{\Sigma}_{a^*}(\sigma^*)) = 1$. Traditional rational-expectations analysis can thus be seen as the special case of ours that arises when the decision maker confronts only risk.

We close the section with a useful equivalence result. The optimality condition (6) can be written in predictive form as $R(a^*, \sigma_{\mu^*}) \geq R(a, \sigma_{\mu^*})$ for each $a \in A$. Relatedly, the data confirmation condition (7) implies that the predictive probability σ_{μ^*} belongs to $\hat{\Sigma}_{a^*}(\sigma^*)$ if it belongs to Σ .²⁷ In this case, $(a^*, \delta(\sigma_{\mu^*}))$ is a self-confirming equilibrium too. Hence we have the following *dogmatic equivalence principle*.

Proposition 3 *Let (a^*, μ^*) be a self-confirming equilibrium given σ^* . If $\sigma_{\mu^*} \in \Sigma$, then $(a^*, \delta(\sigma_{\mu^*}))$ is a self-confirming equilibrium as well, with $V(a^*, \mu^*) = V(a^*, \delta(\sigma_{\mu^*}))$.*

4.2 Steady state interpretation

To give perspective on the self-confirming equilibrium concept, we discuss in more detail its steady state interpretation by adapting to our framework some results in stochastic control, as found in the classic work of Easley and Kiefer (1988). The working paper version of Battigalli et al. (2019)²⁸ describes the exact relationship between our framework and the stochastic control framework of Easley and Kiefer, showing that they are essentially equivalent, moreover, they put forward an analysis of learning dynamics for the case of an ambiguity averse decision maker. To ease matters, throughout this subsection we assume that the sets M , S , and Σ are *finite*.

The distribution $\hat{f}_a(\sigma_\mu) \in \Delta(M)$ identifies, by assigning them a positive probability, the messages that the decision maker deems possible to receive when μ is his prior and a is the action he selected. If $m \in \text{supp } \hat{f}_a(\sigma_\mu)$, it is possible to compute the posterior probability

$$\forall \sigma \in \Sigma, \quad \mu(\sigma | a, m) = \mu(\sigma) \frac{\sigma(f_a = m)}{\sigma_\mu(f_a = m)} = \mu(\sigma) \frac{\hat{f}_a(\sigma)(m)}{\hat{f}_a(\sigma_\mu)(m)}.$$

²⁷The conjectural equilibrium conditions, stated for games by Battigalli (1987), are written in predictive form.

²⁸IGIER w.p. 588, Bocconi University.

The next result (cf. Lemma 2 in Easley and Kiefer, 1988) shows that priors and posteriors are equal under the data confirmation condition (7). Updating thus no longer operates when beliefs satisfy this condition, a property that clarifies the steady state rationale of the self-confirming equilibrium notion as a rest point of a learning process.

Lemma 2 *Let $a \in A$ and $\text{supp } \mu \cap \hat{\Sigma}_a(\sigma^*) \neq \emptyset$. Then $\mu \in \Delta(\hat{\Sigma}_a(\sigma^*))$ if and only if $\mu(\cdot | a, m) = \mu(\cdot)$ for all $m \in \text{supp } \hat{f}_a(\sigma_\mu)$.*

Consider a decision maker who faces a problem with feedback (D, f) for infinitely many periods. The state process is i.i.d. with unknown marginal measure $\sigma \in \Sigma$; the corresponding product measure on S^∞ is denoted σ^∞ . The decision maker starts with a prior belief $\mu_0 \in \Delta(\Sigma)$ at the beginning of period $t = 1$ and holds updated posterior belief $\mu_t \in \Delta(\Sigma)$ at the end of each period $t \geq 1$ and the beginning of period $t + 1$. If action a_t is chosen in period t and message m_t is observed at the end of the same period, then

$$\forall \sigma \in \Sigma, \quad \mu_t(\sigma | a_t, m_t) = \mu_{t-1}(\sigma) \frac{\hat{f}_{a_t}(\sigma)(m_t)}{\hat{f}_{a_t}(\sigma_\mu)(m_t)} \quad (11)$$

is the Bayes update of μ_{t-1} given (a_t, m_t) . The decision maker carries out a stationary strategy²⁹

$$\begin{aligned} \alpha^* &: \Delta(\Sigma) \rightarrow A \\ \mu_{t-1} &\mapsto a_t = \alpha^*(\mu_{t-1}) \end{aligned}$$

that maximizes the discounted expected utility

$$\int_{\Sigma} \left(\sum_{t=1}^{\infty} \beta^{t-1} R(a_t, \sigma) \right) d\mu_0(\sigma),$$

where $\beta \in (0, 1)$ is the discount factor. We call such α^* an *optimal stationary strategy* for the repeated decision problem (D, f, β, μ_0) .

Every stationary strategy α^* yields a corresponding stochastic process of actions and beliefs $(\alpha^*(\mu_{t-1}^*), \mu_{t-1}^*)_{t \geq 1}$ defined on probability space $(S^\infty, \sigma^\infty)$ and adapted to the filtration on S^∞ induced by feedback f , where $\mu_0^* = \mu_0$ and $\mu_t^*(s^\infty)$ is the Bayes update of $\mu_{t-1}^*(s^\infty)$ given action $\alpha^*(\mu_{t-1}^*(s^\infty))$ and message $f(\alpha^*(\mu_{t-1}^*(s^\infty)), s_t)$. Next we provide simple sufficient conditions for such an *active learning process* to converge to a self-confirming equilibrium. Of course, since the realized history of states $(s_t)_{t \geq 1}$ is random, the limit of the process is, in general, also random. With this, we say that a random pair $(\mathbf{a}^*, \boldsymbol{\mu}^*) : S^\infty \rightarrow A \times \Delta(\Sigma)$ is a *random self-confirming equilibrium* given model σ^* if, $\sigma^{*, \infty}$ -almost surely, the pair $(\mathbf{a}^*, \boldsymbol{\mu}^*)$ is a self-confirming equilibrium given σ^* .

Proposition 4 *Let (D, f, β, μ_0) satisfy the following assumptions:*

- (i) *A is a compact and convex subset of an Euclidean space;*

²⁹A strategy is stationary if it depends only on updated beliefs about the model. As the set of optimal strategies always contains a stationary strategy, the stationarity restriction is without substantial loss of generality.

(ii) $a \mapsto \hat{f}_a$ is continuous, and $\hat{f}_a(\sigma)(m) > 0$ for every $(a, \sigma, m) \in A \times \Sigma \times M$;

(iii) $r : A \times S \rightarrow \mathbb{R}$ is continuous and strictly concave on A ;

(iv) $\mu_0(\sigma) > 0$ for every $\sigma \in \Sigma$.

Then an optimal stationary strategy α^* exists and, for every $\sigma^* \in \Sigma$, the induced active learning process $(\alpha^*(\mu_{t-1}^*), \mu_{t-1}^*)_{t \geq 1}$ converges, $\sigma^{*,\infty}$ -almost surely, to a random self-confirming equilibrium $(\alpha^*(\mu_\infty^*), \mu_\infty^*)$ given model σ^* .

The proposition is, essentially, a self-confirming interpretation of stochastic control results of Easley and Kiefer (1988) and so we omit its proof.³⁰ Heuristically, the existence of an optimal stationary strategy follows from standard results in dynamic programming. The martingale convergence theorem implies that there exists a random belief $\mu_\infty^* : S^\infty \rightarrow \Delta(\Sigma)$ such that, $\sigma^{*,\infty}$ -almost surely, $\mu_t^* \rightarrow \mu_\infty^*$ (cf. Theorem 4 in Easley and Kiefer, 1988). The Bayesian updating function, which via (11) associates μ_t to μ_{t-1} given actions and messages (a_t, m_t) , is continuous on $\Delta(\Sigma)$. Hence, for every sample path $s^\infty = (s_t)_{t \geq 1}$ such that $\mu_t^*(s^\infty) \rightarrow \mu_\infty^*(s^\infty)$ (hence, by $\sigma^{*,\infty}$ -almost sure convergence, every s^∞ such that $\prod_{t=1}^T \sigma^*(s_t) > 0$ for all T) and for every action a_∞ in the set of limit points of $(\alpha^*(\mu_t^*(s^\infty)))_{t \geq 1}$ (which is not empty, by compactness of A), it must be the case that $\mu_\infty^*(s^\infty)$ is a_∞ -invariant, that is, $\mu_\infty^*(s^\infty)(\cdot | a_\infty, m) = \mu_\infty^*(s^\infty)(\cdot)$ for every possible m . In view of Lemma 2, this is equivalent to the confirmed beliefs condition (7). Since the value of experimentation vanishes in the limit, every such action a_∞ maximizes the one-period subjective expected value $V(\cdot, \mu_\infty^*(s^\infty))$ (cf. Lemma 4 in Easley and Kiefer, 1988). By assumption (iii), $V(\cdot, \mu_\infty^*(s^\infty))$ has a unique maximizer, therefore $\lim_{t \rightarrow \infty} (\alpha^*(\mu_t^*(s^\infty)), \mu_t^*(s^\infty)) = (\arg \max_{a \in A} V(a, \mu_\infty^*(s^\infty)), \mu_\infty^*(s^\infty))$, which is a self-confirming equilibrium (cf. Theorem 5 in Easley and Kiefer, 1988).

Can any self-confirming equilibrium action be interpreted as a limit point of an active learning process? We can provide a simple, but partial answer by recalling that if (a^*, μ^*) is self-confirming and $\sigma_{\mu^*} \in \Sigma$, then also $(a^*, \delta_{\sigma_{\mu^*}})$ is self-confirming. With this, if the initial belief is $\mu_0 = \delta_{\sigma_{\mu^*}}$, then the resulting active learning process is constant at $(a^*, \delta_{\sigma_{\mu^*}})$. Also, if the decision maker is impatient ($\beta = 0$) and (a^*, μ^*) is self-confirming, then the learning process starting at $\mu_0 = \mu^*$ is constant at (a^*, μ^*) .

4.3 Value and welfare

We now turn to a “welfare analysis,” that is, we compare equilibrium values with objective expected payoffs, including the maximum expected payoff that could be attained by the decision maker if she knew the true model σ^* . We start with an important preliminary result: since we assume that consequences are observable, it follows that, for each action, observationally equivalent models yield the same objective expected payoff.³¹

Lemma 3 *Let $(a, \sigma) \in A \times \Delta(S)$. If model $\sigma' \in \Delta(S)$ is observationally equivalent to model σ under a , then $R(a, \sigma') = R(a, \sigma)$.*

³⁰For more details on this result, also in the case of an ambiguity averse agent, see Battigalli et al. (2019).

³¹This lemma extends Lemma 1 of Battigalli et al. (2015) to the present setup with possibly infinite spaces of actions and states.

This has a noteworthy consequence for self-confirming equilibrium values.

Proposition 5 *If (a^*, μ^*) is a self-confirming equilibrium given σ^* , then*

$$V(a^*, \mu^*) = R(a^*, \sigma^*).$$

Thus, the value of any self-confirming equilibrium (a^*, μ^*) coincides with the true expected payoff of a^* , irrespective of the supporting belief μ^* . As a result, because of the data confirmation condition, the optimality condition (6) amounts to assuming that the “true value” of the self-confirming equilibrium action is higher than the subjective value, under the equilibrium belief, of all alternative actions. This interplay of objective and subjective features shows the substantial bite of the data confirmation condition.

Lemma 3 has interesting comparative welfare implications. For our welfare analysis, it is convenient to focus on actions that are part of some self-confirming equilibrium, thus neglecting the supporting confirmed beliefs.

Definition 2 *Action a^* is a self-confirming (equilibrium) action given σ^* if there exists a belief $\mu^* \in \Delta(\Sigma)$ such that (a^*, μ^*) is a self-confirming equilibrium.*

Since in this case $V(a^*, \mu^*) = R(a^*, \sigma^*) \leq \sup_{a \in A} R(a, \sigma^*)$, the decision maker incurs a welfare loss

$$\ell(a^*, \sigma^*) = \sup_{a \in A} R(a, \sigma^*) - R(a^*, \sigma^*) \geq 0$$

when she selects the self-confirming action a^* . In particular, $\ell(a^*, \sigma^*) = 0$ if and only if a^* is objectively optimal; the loss is caused by the decision maker’s ignorance, which makes it possible to assign positive subjective probability to (neighborhoods of) models different from the true one. Our next result shows that self-confirming equilibria with sharper basic subjective assessments yield higher welfare (lower loss). Formally, μ^* is *absolutely continuous* with respect to (i.e., “sharper than”) ν^* if, for every Borel set $B \subseteq \Sigma$, $\nu^*(B) = 0$ implies $\mu^*(B) = 0$ (equivalently, $\mu^*(B) > 0$ implies $\nu^*(B) > 0$). This means that μ^* rules out more models than ν^* ; in particular, if Σ is finite, it means that $\text{supp}\mu^* \subseteq \text{supp}\nu^*$.

Proposition 6 *Let (a^*, μ^*) and (b^*, ν^*) be self-confirming equilibria such that μ^* is absolutely continuous with respect to ν^* , then $\ell(a^*, \sigma^*) \leq \ell(b^*, \sigma^*)$.*

Consider the self-confirming equilibria (a^*, μ^*) and (b^*, ν^*) such that (i) a^* yields better identification than b^* (i.e., $\hat{\Sigma}_{a^*}(\sigma^*) \subseteq \hat{\Sigma}_{b^*}(\sigma^*)$), and (ii) μ^* and ν^* do not rule out any model consistent with the statistical evidence given a^* and b^* respectively. Then, we obtain a special case of Proposition 6 and we can conclude that $\ell(a^*, \sigma^*) \leq \ell(b^*, \sigma^*)$. The following result shows that we can dispense with condition (ii): Independently of their justifying confirmed beliefs, self-confirming actions with better identification properties exhibit lower losses.

Proposition 7 *Let a^* and b^* be self-confirming actions given σ^* . If $\hat{\Sigma}_{a^*}(\sigma^*) \subseteq \hat{\Sigma}_{b^*}(\sigma^*)$, then $\ell(a^*, \sigma^*) \leq \ell(b^*, \sigma^*)$.*

Propositions 6 and 7 are the only results in our analysis that depend on the one-person assumption in an essential way. In a multi-person game they hold only for the comparison

of equilibria where the strategies of all players but one are the same and the focus is on the welfare of the only agent playing a different strategy.

The next related result shows that an action with the best identification properties — thus, optimal from a purely statistical viewpoint — is self-confirming only when objectively optimal. Truth is ancillary to the decision maker’s pursuit of her goals (and so of her happiness).

Proposition 8 *An action $a \in A$ such that $\hat{\Sigma}_a(\sigma^*) \subseteq \hat{\Sigma}_{a'}(\sigma^*)$ for each $a' \in A$ is self-confirming given σ^* if and only if it is objectively optimal.*

Under own-action independence of feedback about the state, $\hat{\Sigma}_a(\sigma^*)$ is independent of a . Therefore, Proposition 8 yields the following noteworthy implication.

Corollary 1 *Under own-action independence of feedback about the state, every self-confirming action is objectively optimal.*

Example 4 (Monopoly: self-confirming equilibrium) *Under the assumptions of Example 1, certainty equivalence holds and the subjective best reply function of the monopolist is*

$$B(\mu) = B(\mathbb{E}_\mu(\theta)) = \max \left\{ \underline{a}, \frac{\mathbb{E}_\mu(\theta) - c}{2} \right\},$$

where θ parameterizes models according to the average intercept of the inverse demand function. Since any positive output is revealing (see Example 3), if the firm is pre-committed to a positive minimum output ($\underline{a} > 0$) own-action independence of feedback holds and the only self-confirming output is the objective best reply $\max\{\underline{a}, (\theta^* - c)/2\}$. Next, suppose that $\underline{a} = 0$, and furthermore $\underline{\theta} < c$ and $\theta^* > c$. Then own-action independence of feedback does not hold and there are two self-confirming actions: (i) the fully revealing action $a^* = (\theta^* - c)/2 > 0$ is the objective best reply, thus illustrating Proposition 8, and (ii) $b^* = 0$ is justified by any “pessimistic” belief μ such that $\mathbb{E}_\mu(\theta) < c$, which is trivially consistent with long-run evidence because $b^* = 0$ is fully un-revealing. The comparison of self-confirming actions a^* and b^* illustrates Proposition 7: indeed, a^* is more revealing than b^* and $\ell(a^*, \theta^*) = 0 < \left(\frac{\theta^* - c}{2}\right)^2 = \ell(b^*, \theta^*)$.

The example prompts the following question. We explained in Section 4.2 that self-confirming equilibria are limit steady states of active learning processes. Suppose that the monopolist believes it is optimal in the short run to produce 0, but deems it possible that the objective best reply is positive, i.e., that $\theta > c$. Should she not experiment with a positive output? This depends on several elements: her subjective belief, her degree of patience (discount factor), and the amount of noise. If the subjective probability $\mu(\theta > c)$ is relatively small and price is noisy, it is dynamically optimal not to experiment even if the decision maker is moderately patient. In particular, noise is important: only repeated experimentation with positive output can provide reliable evidence, and this has a high subjective opportunity cost.³² In sum, the decision maker is not just a statistician: she is

³²On the other hand, without noise ($\bar{\varepsilon} = 0$) a one-off experimentation with positive output would identify the true model, making the no-experimentation region in the belief-discount factor space very small.

not interested in discovering the true model per se, unless the action (played in the long run) that allows the discovery is subjectively optimal.

In this first part we expressed and analyzed the self-confirming equilibrium concept in an abstract framework amenable to policy applications. This requires to allow for an infinite action space (e.g., to use calculus) and for an infinite state space, and to posit an objective probability model characterizing the data generating process. Technically, the latter calls for the use of standard Borel spaces. Many of the themes analyzed within the framework of the first part are illustrated in the second part by an application to monetary policy.

5 Phillips curve exploitation model

We now illustrate our machinery in the context of a 1970's U.S. policy debate about whether a trade-off between inflation and unemployment can be systematically exploited by a benevolent policy maker. We extend a formulation of Sargent (1999, 2008), who presents a self-confirming equilibrium in which a policy maker believes in a model asserting an exploitable trade-off between unemployment and inflation while the truth is that the trade-off is not exploitable.³³

5.1 Steady state model economies

We study a class Θ of model economies θ at a (stochastic) steady state. We assume that unemployment u and inflation π , beyond depending on the unknown θ , are affected by random shocks w and ε with zero mean, and by a monetary policy variable a . Specifically, unemployment and inflation outcomes (u, π) are connected to the state of the economy $s = (w, \varepsilon, \theta)$ and the government action a according to

$$u = \theta_0 + \theta_{1\pi}\pi + \theta_{1a}a + \theta_2w, \quad (12)$$

$$\pi = a + \theta_3\varepsilon. \quad (13)$$

The vector parameter $\theta = (\theta_0, \theta_{1\pi}, \theta_{1a}, \theta_2, \theta_3) \in \mathbb{R}^5$, that is, the last component of the state vector, specifies the structural coefficients of an aggregate supply equation (12) and an inflation determination equation (13). Coefficients $\theta_{1\pi}$ and θ_{1a} are slope responses of unemployment to actual and planned inflation,³⁴ while the coefficients θ_2 and θ_3 quantify shock volatilities (see Sargent, 2008, p. 18). Finally, the intercept θ_0 is the baseline rate of unemployment that would (systematically) prevail at a zero planned inflation policy $a = 0$.

Throughout the section we maintain the following assumption about structural coefficients.

Assumption 1 $\theta_0 > 0$, $\theta_{1\pi} < 0$, $\theta_2 > 0$ and $\theta_3 > 0$.

In words, we posit a strictly positive baseline rate of unemployment, as well as strictly positive shock coefficients (nontrivial, possibly asymmetric, shocks thus affect both the inflation and the unemployment equations, their unknown values form the first component (w, ε)

³³Section 6 of the working paper version contains a more general analysis of self-confirming economic policies.

³⁴The economic interpretation is that planned inflation a affects agents' expectations to an extent parameterized by θ_{1a} .

of the state vector). Finally, we assume that —other things being equal— more inflation reduces unemployment.

The reduced form of each model economy is

$$u = \theta_0 + (\theta_{1\pi} + \theta_{1a})a + \theta_{1\pi}\theta_3\varepsilon + \theta_2w, \quad (14)$$

$$\pi = a + \theta_3\varepsilon. \quad (15)$$

The coefficients of the reduced form are $\xi = (\theta_0, \theta_{1\pi} + \theta_{1a}, \theta_{1\pi}\theta_3, \theta_2, \theta_3) \in \mathbb{R}^5$. Since $\theta_3 \neq 0$ (Assumption 1), it is easy to check that different structural parameter vectors $\theta \in \Theta$ correspond to different reduced form parameter vectors ξ , that is, $\theta \neq \theta'$ implies $\xi \neq \xi'$.

We assume that only realized unemployment and inflation are observable by the monetary authority. Thus, the reduced form above will give us the feedback function $(u, \pi) = f(a, s)$ of the previous sections. Specifically, rewriting (14) and (15) as

$$\begin{aligned} \mathbf{u}(a, w, \varepsilon, \theta) &= \theta_0 + (\theta_{1\pi} + \theta_{1a})a + \theta_{1\pi}\theta_3\varepsilon + \theta_2w, \\ \boldsymbol{\pi}(a, w, \varepsilon, \theta) &= a + \theta_3\varepsilon \end{aligned}$$

makes the dependence of observables (u, π) on action a and (unobservable) realized states (w, ε, θ) explicit, which allows us to study the present policy problem within our general framework. Formally, the message space $M = \mathbb{R}^2$ now consists of unemployment/inflation pairs, and the feedback function is $f = (\mathbf{u}, \boldsymbol{\pi}) : A \times (\mathbb{R}^2 \times \Theta) \rightarrow \mathbb{R}^2$.

The policy multiplier $\xi_2 = \theta_{1\pi} + \theta_{1a} = \theta_{1a} - |\theta_{1\pi}|$ quantifies the impact of planned inflation on unemployment. It is the sum of the direct and indirect impact of planned inflation on unemployment quantified, respectively, by θ_{1a} and $\theta_{1\pi}$. There is a systematic trade-off between unemployment and inflation when the multiplier is strictly negative, that is, $\xi_2 < 0$. If so, the model economy is *Keynesian*; otherwise, it is *new-classical*. In the rest of the section we make the following hypothesis on the multiplier.

Assumption 2 $\xi_2 \leq 0$.

Thus, we assume that an increase in planned inflation never increases unemployment. A possible interpretation of the model is that $\theta_{1a}/|\theta_{1\pi}|$ is the constant fraction of experienced/sophisticated agents in the economy who factor planned inflation into their expectations, and $\xi_2/\theta_{1\pi}$ is the fraction of inexperienced/naive agents.

To sum up, the set of parameters is

$$\Theta = \{\theta \in \mathbb{R}^5 : \theta_0 > 0, \theta_{1a} \leq -\theta_{1\pi}, \theta_{1\pi} < 0, \theta_2 > 0, \theta_3 > 0\}.$$

To clarify our language, we note that we keep using “model” in the same sense as in the previous sections, that is, a probability measure over states, or a specific parameter value that determines such measure. Thus, a set of parameterized equations like (12)-(13) corresponds to a class of models. We will therefore refer subclasses of models satisfying some restrictions as “kinds”. With this, our analysis will pay special attention to the following two competing kinds of model economies.

5.1.1 Lucas-Sargent models

The first kind of model economy, based on Lucas (1972) and Sargent (1973), is

$$\begin{aligned} u &= \theta_0 + \beta(\pi - a) + \theta_2 w = \theta_0 + \beta\theta_3 \varepsilon + \theta_2 w, \\ \pi &= a + \theta_3 \varepsilon, \end{aligned}$$

where $\beta \equiv \theta_{1\pi} = -\theta_{1a}$, and so $\theta = (\theta_0, \beta, -\beta, \theta_2, \theta_3)$ and $\xi = (\theta_0, 0, \beta\theta_3, \theta_2, \theta_3)$. In such new-classical models the policy multiplier ξ_2 is zero, and so the systematic part of inflation a has no effect on unemployment; only the unsystematic part $\theta_3 \varepsilon$ does.

5.1.2 Samuelson-Solow models

A second kind of model economy, based on Samuelson and Solow (1960), is

$$\begin{aligned} u &= \theta_0 + \theta_{1\pi} \pi + \theta_2 w = \theta_0 + \theta_{1\pi} a + \theta_{1\pi} \theta_3 \varepsilon + \theta_2 w, \\ \pi &= a + \theta_3 \varepsilon, \end{aligned}$$

where $\theta_{1a} = 0$ and so $\theta = (\theta_0, \theta_{1\pi}, 0, \theta_2, \theta_3)$ and $\xi = (\theta_0, \theta_{1\pi}, \theta_{1\pi} \theta_3, \theta_2, \theta_3)$. In such Keynesian models, the policy multiplier $\xi_2 = \theta_{1\pi}$ is strictly negative: monetary policies affect, at steady state, unemployment rates.

5.2 The policy problem: setup and identification

5.2.1 Setup

The monetary authority chooses policy a . As anticipated, the state space is the Cartesian product $S = W \times E \times \Theta$, which expresses that the monetary authority is uncertain about both shocks and permanent features of the economy, or models. The consequence space C consists of unemployment and inflation pairs $c = (u, \pi)$, so we set $C = U \times \Pi \subseteq \mathbb{R}^2$. The consequence function $\rho : A \times (W \times E \times \Theta) \rightarrow C$ is

$$\rho(a, w, \varepsilon, \theta) = (\mathbf{u}(a, w, \varepsilon, \theta), \boldsymbol{\pi}(a, w, \varepsilon, \theta)),$$

which is the unemployment/inflation pair (u, π) determined by policy a and state (w, ε, θ) , with matrix representation

$$\rho(a, w, \varepsilon, \theta) = \begin{bmatrix} \theta_0 \\ 0 \end{bmatrix} + a \begin{bmatrix} \theta_{1\pi} + \theta_{1a} \\ 1 \end{bmatrix} + \begin{bmatrix} \theta_2 & \theta_{1\pi} \theta_3 \\ 0 & \theta_3 \end{bmatrix} \begin{bmatrix} w \\ \varepsilon \end{bmatrix}. \quad (16)$$

5.2.2 Factorization

As anticipated, we assume that the messages received by the monetary authority are the policy outcomes. Hence, a message $m = (u, \pi)$ consists of an unemployment and inflation pair, and the feedback function

$$f = \rho = (\mathbf{u}, \boldsymbol{\pi})$$

corresponds to the reduced form of the model economy. When the monetary authority chooses policy a and in the long run observes a distribution over (u, π) pairs, it can partially

infer the underlying stochastic model σ . For example, if σ has finite support, the induced probability of outcome (u, π) is³⁵

$$\hat{f}_a(\sigma)(u, \pi) = \sigma(\{(w, \varepsilon, \theta) : (\mathbf{u}(a, w, \varepsilon, \theta), \boldsymbol{\pi}(a, w, \varepsilon, \theta)) = (u, \pi)\}). \quad (17)$$

The partially identified set $\hat{\Sigma}_a(\sigma)$ of stochastic models indistinguishable from σ is the set of σ' that induce the same joint distribution on unemployment/inflation outcomes given a .

At this point, it is convenient to add structure to this setup to provide a sharp characterization of the partially identified set corresponding to each policy a and model σ . Within a state $s = (w, \varepsilon, \theta)$, the pair (w, ε) represents random shocks and θ parameterizes a model economy. This suggests factorizing the probability models $\sigma \in \Sigma \subseteq \Delta(W \times E \times \Theta)$ as

$$\sigma = q \times \delta(\theta), \quad (18)$$

where the true marginal distribution of shocks $q \in \Delta(W \times E)$ is assumed to be *known* and $\delta(\theta) \in \Delta(\Theta)$ is a Dirac probability measure concentrated on a given economic model $\theta \in \Theta$, a permanent feature of the environment. We thus parameterize probability models with θ and write σ_θ .

The simplifying assumption that, at a steady state, the distribution q of shocks is known is common in the rational expectations literature since Lucas and Prescott (1971) and Lucas (1972). The resulting factorization (18) has two modeling consequences: (i) it establishes a one-to-one correspondence between model economies and probability models (in particular, a true economic model θ^* corresponds to a true probability model σ_{θ^*}); (ii) since q is known, it allows us to identify Σ with Θ via the relation

$$\Sigma = \{q \times \delta(\theta) \in \Delta(S) : \theta \in \Theta\},$$

and so to define the prior μ on Θ .³⁶

A first dividend of the factorization is that the objective function (2) takes the simpler form

$$V(a, \mu) = \int_{\Theta} \left(\int_{W \times E} r(a, w, \varepsilon, \theta) dq(w, \varepsilon) \right) d\mu(\theta), \quad (19)$$

where $r(a, w, \varepsilon, \theta) = v(\rho(a, w, \varepsilon, \theta))$ is the utility of outcome/message $(u, \pi) = \rho(a, w, \varepsilon, \theta)$.

In the rest of the section we maintain the following assumption on the known shock distributions.³⁷

Assumption 3 $\mathbb{E}_q(\varepsilon) = \mathbb{E}_q(\mathbf{w}) = \mathbb{E}_q(\varepsilon \mathbf{w}) = 0$ and $\mathbb{E}_q(\varepsilon^2) = \mathbb{E}_q(\mathbf{w}^2) = 1$.

In words, shocks are uncorrelated and normalized.

³⁵In the general case, for any measurable set of outcomes $O \subseteq U \times \Pi$,

$$\hat{f}_a(\sigma)(O) = (\sigma \circ f_a^{-1})(O) = \sigma(\{(w, \varepsilon, \theta) : (\mathbf{u}(a, w, \varepsilon, \theta), \boldsymbol{\pi}(a, w, \varepsilon, \theta)) \in O\}).$$

³⁶The map $\theta \mapsto q \times \delta_\theta$ is bijective and measurable. See Corollary 3 in the appendix.

³⁷Whenever convenient, in what follows we will use the shorthand notation \mathbb{E} for integrals, for example $\mathbb{E}_q(\varepsilon) = \int_{W \times E} \varepsilon dq(w, \varepsilon)$.

5.2.3 Identification

In this “factorized” setup, we can shift our focus from observationally equivalent probability models σ to observationally equivalent model economies θ . The partially identified set becomes:

$$\forall \theta \in \Theta, \hat{\Sigma}_a(\theta) = \left\{ \theta' \in \Theta : \hat{f}_a(\sigma_{\theta'}) = \hat{f}_a(\sigma_{\theta}) \right\}.$$

With this, a sharp identification result holds.

Proposition 9 *The partial identification correspondence $\hat{\Sigma}_a : \Theta \rightarrow 2^{\Theta}$ is*

$$\hat{\Sigma}_a(\theta) = \left\{ \theta' \in \Theta : \theta'_0 + \theta'_{1\mathbf{a}}a = \theta_0 + \theta_{1\mathbf{a}}a, \theta'_{1\pi} = \theta_{1\pi}, \theta'_2 = \theta_2, \theta'_3 = \theta_3 \right\}. \quad (20)$$

Given the true model θ , the shock coefficients θ_2 and θ_3 are thus identified, along with the slope $\theta_{1\pi}$ of the Phillips curve, independently of the chosen policy a . As we discuss below, the intercept of the curve is also identified, but it depends on the maintained policy a through the unidentified parameter $\theta_{1\mathbf{a}}$. This important identification result is made possible by some moment conditions, formally spelled out in the proof. We can, however, heuristically describe them via the bivariate random variable $(\mathbf{u}_a, \boldsymbol{\pi}_a) : W \times E \times \Theta \rightarrow U \times \Pi$ that, for a given policy a , represents the unemployment and inflation rates determined by the state (w, ε, θ) .³⁸ The monetary authority infers the following moments from the long-run distribution of outcomes:

- $\mathbb{E}_{\theta}(\mathbf{u}_a) = \theta_0 + (\theta_{1\pi} + \theta_{1\mathbf{a}})a$,
- $\mathbb{E}_{\theta}(\boldsymbol{\pi}_a) = a$,
- $\text{Var}_{\theta}(\mathbf{u}_a) = \theta_{1\pi}^2\theta_3^2 + \theta_2^2$,
- $\text{Var}_{\theta}(\boldsymbol{\pi}_a) = \theta_3^2$,
- $\text{Cov}_{\theta}(\mathbf{u}_a, \boldsymbol{\pi}_a) = \theta_{1\pi}\theta_3^2$.

Therefore,

$$\theta_{1\pi} = \frac{\text{Cov}_{\theta}(\mathbf{u}_a, \boldsymbol{\pi}_a)}{\text{Var}_{\theta}(\boldsymbol{\pi}_a)} \quad (21)$$

is the beta coefficient of the Phillips regression of unemployment on inflation,³⁹

$$\theta_2^2 = (1 - \text{Corr}_{\theta}^2(\mathbf{u}_a, \boldsymbol{\pi}_a)) \text{Var}_{\theta}(\mathbf{u}_a)$$

is the residual variance of \mathbf{u}_a (unexplained by the regression), and θ_3 is the standard deviation of inflation.

Finally, though the two structural coefficients θ_0 and $\theta_{1\mathbf{a}}$ remain unidentified even in the long-run, they satisfy

$$\theta_0 + \theta_{1\mathbf{a}}a = \mathbb{E}_{\theta}(\mathbf{u}_a) - \frac{\text{Cov}_{\theta}(\mathbf{u}_a, \boldsymbol{\pi}_a)}{\text{Var}_{\theta}(\boldsymbol{\pi}_a)} \mathbb{E}_{\theta}(\boldsymbol{\pi}_a), \quad (22)$$

where the right side is the alpha coefficient of the Phillips regression. In the long-run, the alpha coefficient is observed by the monetary authority, but what is observed depends on the policy a that the authority chooses.

³⁸Formally, \mathbf{u}_a and $\boldsymbol{\pi}_a$ are the sections $\mathbf{u}(a, \cdot)$ and $\boldsymbol{\pi}(a, \cdot)$ at policy a of the random variables \mathbf{u} and $\boldsymbol{\pi}$, respectively.

³⁹The Phillips regression $u = \alpha + \beta\pi$ is run by the monetary authority using long run data.

5.2.4 Estimated model economy

As an approximation of a situation in which the dataset is large and the sample variance is small, we take the idealized perspective of a monetary authority (or its econometrician) who can rely on an infinite dataset and therefore can perfectly estimate the identifiable parameters by observing some moments of the true distribution as specified in Proposition 9.

The moments that identify the three coefficients $\theta_{1\pi}$, θ_2 , and θ_3 do not depend on the chosen policy a , but only on the true model θ . To emphasize this key feature, we denote by $\hat{\beta}$ the beta regression coefficient that identifies $\theta_{1\pi}$,⁴⁰ by $\hat{\sigma}_{u|\pi}$ the residual standard deviation that identifies θ_2 , and by $\hat{\sigma}_\pi$ the standard deviation of inflation that identifies θ_3 . In contrast, the alpha regression coefficient that identifies the sum $\theta_0 + \theta_{1a}a$ depends on policy a ; we denote it by $\hat{\alpha}(a)$.

With this, we can write

$$\hat{\Sigma}_a(\theta) = \left\{ \theta' \in \Theta : \theta'_0 + \theta'_{1a}a = \hat{\alpha}(a), \theta'_{1\pi} = \hat{\beta}, \theta'_2 = \hat{\sigma}_{u|\pi}, \theta'_3 = \hat{\sigma}_\pi \right\}.$$

As a result, the long-run estimated version of the model economy (12)-(13) that the monetary authority considers is

$$u = \hat{\alpha}(a) + \hat{\beta}\pi + \hat{\sigma}_{u|\pi}w, \quad (23)$$

$$\pi = a + \hat{\sigma}_\pi\varepsilon, \quad (24)$$

$$\hat{\alpha}(a) = \theta_0 + \theta_{1a}a. \quad (25)$$

In particular, (23) is the estimated aggregate supply equation and (24) is the estimated inflation equation. The intercept of the former equation depends on policy a via eq. (25), which only partly identifies the two coefficients θ_0 and θ_{1a} . In turn, this makes the policy multiplier $\xi_2 = \hat{\beta} + \theta_{1a}$ unidentified. We will momentarily address this key partial identification issue.

5.2.5 Partial identification line

The monetary authority cannot identify—even in the long-run—the two structural coefficients θ_0 and θ_{1a} . The former is the average unemployment at zero planned inflation, $\theta_0 = \mathbb{E}_\theta(\mathbf{u}_0)$; the latter is the “direct” impact of policy on unemployment.

The parameter space of the estimated model economy (23)-(25) reduces to $\Theta = \tilde{\Theta} \times \{(\hat{\beta}, \hat{\sigma}_{u|\pi}, \hat{\sigma}_\pi)\}$, where $\tilde{\Theta} = \mathbb{R}_{++} \times (-\infty, -\hat{\beta}]$ is the collection of all possible values (θ_0, θ_{1a}) of the two remaining unidentified coefficients and $\{(\hat{\beta}, \hat{\sigma}_{u|\pi}, \hat{\sigma}_\pi)\}$ is the singleton containing the identified vector $(\theta_{1\pi}, \theta_2, \theta_3)$. To ease notation, in what follows we will consider directly $\tilde{\Theta}$ as the parameter space. As a result, the parameter space is now a subset of the plane. By (20), the partial identification correspondence $\hat{\Sigma}_a : \tilde{\Theta} \rightarrow 2^{\tilde{\Theta}}$ becomes

$$\hat{\Sigma}_a(\theta) = \left\{ (\theta'_0, \theta'_{1a}) \in \tilde{\Theta} : \theta'_0 = -\theta'_{1a}a + \theta_0 + \theta_{1a}a \right\}. \quad (26)$$

In words, $\hat{\Sigma}_a(\theta)$ is a straight line in the plane, with slope $-a$ and intercept $\theta_0 + \theta_{1a}a$ (determined by the policy a and by the true economic model θ). We thus have a partial

⁴⁰By Assumption 2, the beta coefficient of the Phillips regression is negative, that is, $\hat{\beta} < 0$. This negative sign will be tacitly assumed when interpreting our findings.

identification line that defines a linear relationship between the two unidentified coefficients, given the true model. In other words, partial identification is unidimensional.

Given true model $\theta = (\theta_0, \theta_{1a})$, the collection $\{\hat{\Sigma}_a(\theta) : a \in A\}$ of partial identification lines is the family of all straight lines in the plane that pass through the true model (θ_0, θ_{1a}) and have slope $-1/a$. In each such line there is a unique Lucas-Sargent model, characterized by $\theta'_{1a} = -\hat{\beta}$, as well as a unique Samuelson-Solow model, characterized by $\theta'_{1a} = 0$. In other words, partial identification lines feature a unique specimen of each kind of models.




Figure 5.2.5: Partial identification line.

Figure 5.2.5 illustrates the previous analysis. In particular, LS stands for Lucas-Sargent model and SS for Samuelson-Solow model, while the red (resp., blue) line is the partial identification line that correspond to policy $a = 0$ (resp., $a > 0$).

5.3 The policy problem: value, equilibria and welfare

5.3.1 Value and equilibrium

As much of the literature, we assume a quadratic von Neumann-Morgenstern utility function $v : C \rightarrow \mathbb{R}$ given by $v(u, \pi) = -u^2 - \pi^2$, so that the reward function $r : A \times S \rightarrow \mathbb{R}$ becomes:

$$r(a, w, \varepsilon, \theta) = -\mathbf{u}^2(a, w, \varepsilon, \theta) - \boldsymbol{\pi}^2(a, w, \varepsilon, \theta).$$

The linear model economy and quadratic utility together form a classic linear quadratic policy framework.

Lemma 4 *For every $(\theta, a) \in \tilde{\Theta} \times A$, we have $R(a, \theta) = v(\mathbb{E}_\theta(\mathbf{u}_a), \mathbb{E}_\theta(\boldsymbol{\pi}_a)) + \text{const}$.*

The linear quadratic framework thus allows us to express the expected reward as the utility of expectations. As a result, the objective function (19) becomes

$$V(a, \mu) = \int_{\tilde{\Theta}} v(\mathbb{E}_\theta(\mathbf{u}_a), \mathbb{E}_\theta(\boldsymbol{\pi}_a)) d\mu(\theta) + \text{const}. \quad (27)$$

As for self-confirming equilibria, we begin with a piece of notation: throughout the rest of this section we fix a true model economy θ^* (rather than θ) in $\tilde{\Theta}$, while θ (rather than θ') denotes a generic element of $\tilde{\Theta}$. With this notation, the partial identification line is

$$\hat{\Sigma}_a(\theta^*) = \left\{ (\theta_0, \theta_{1\mathbf{a}}) \in \tilde{\Theta} : \theta_0 = \theta_0^* + (\theta_{1\mathbf{a}}^* - \theta_{1\mathbf{a}}) a \right\}.$$

Hence, a policy and belief pair $(a^*, \mu^*) \in A \times \Delta(\tilde{\Theta})$ is self-confirming if and only if

$$a^* \in \arg \max_{a \in A} V(a, \mu^*)$$

and

$$\mu^* \left(\hat{\Sigma}_{a^*}(\theta^*) \right) = 1.$$

Next we characterize self-confirming equilibria of the estimated model economy (23)-(25). In both equilibrium conditions, the true multiplier $\xi_2^* = \hat{\beta}^* + \theta_{1\mathbf{a}}^*$ and its conjectured value $\mathbb{E}_{\mu^*}(\xi_2) = \hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})$ play a key role.⁴¹

Proposition 10 *A policy and belief pair $(a^*, \mu^*) \in A \times \Delta(\tilde{\Theta})$ is a self-confirming equilibrium given θ^* if and only if*

$$a^* = - \frac{\theta_0^* \left(\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) \right)}{1 + \left(\hat{\beta}^* + \theta_{1\mathbf{a}}^* \right) \left(\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) \right)} \quad (28)$$

and

$$\mu^* \left(\left\{ (\theta_0, \theta_{1\mathbf{a}}) \in \tilde{\Theta} : \theta_0 = \theta_0^* - \frac{\theta_0^* \left(\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) \right)}{1 + \left(\hat{\beta}^* + \theta_{1\mathbf{a}}^* \right) \left(\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) \right)} (\theta_{1\mathbf{a}}^* - \theta_{1\mathbf{a}}) \right\} \right) = 1. \quad (29)$$

The result can be heuristically derived in the special case of *dogmatic beliefs*, when μ^* is concentrated on a single parameter vector $\bar{\theta} = (\bar{\theta}_0, \bar{\theta}_{1\mathbf{a}}) \in \tilde{\Theta}$, that is, $\mu^* = \delta(\bar{\theta})$. By (27), up to a constant the monetary authority's value function is

$$V(a, \mu^*) = -\mathbb{E}_{\bar{\theta}}^2(\mathbf{u}_a) - \mathbb{E}_{\bar{\theta}}^2(\boldsymbol{\pi}_a).$$

The conjectured multiplier is $\bar{\xi}_2 = \bar{\beta}^* + \bar{\theta}_{1\mathbf{a}}$. For instance, a new-classical authority that believes that there is no systematically exploitable trade-off between inflation and unemployment assumes $\bar{\theta}_{1\mathbf{a}} = -\bar{\beta}^*$ (and so the conjectured multiplier is zero). In contrast, a Keynesian authority that believes in a trade-off may assume, for instance, $\bar{\theta}_{1\mathbf{a}} = 0$ (the conjectured multiplier is then $\bar{\beta}^*$, and so strictly negative).

Based on the estimated model economy (23)-(25), a dogmatic authority conjectures that, according to the chosen policy a , the expected values of inflation and unemployment are constrained by the equation

$$\mathbb{E}_{\bar{\theta}}(\mathbf{u}_a) = \bar{\theta}_0 + \left(\bar{\theta}_{1\mathbf{a}} + \bar{\beta}^* \right) \mathbb{E}_{\bar{\theta}}(\boldsymbol{\pi}_a).$$

⁴¹Recall that $\hat{\beta}^*$ is the beta regression coefficient of unemployment over inflation (given the true model θ^*).

This conjectured constraint is the version of the estimated aggregate supply equation (23) that the authority expects to face systematically given its dogmatic belief. So the authority's decision problem is

$$\begin{aligned} \min_{a \in A} \mathbb{E}_{\bar{\theta}}^2(\mathbf{u}_a) + \mathbb{E}_{\bar{\theta}}^2(\boldsymbol{\pi}_a), \\ \text{sub } \mathbb{E}_{\bar{\theta}}(\mathbf{u}_a) = \bar{\theta}_0 + \left(\bar{\theta}_{1\mathbf{a}} + \hat{\beta}^*\right) \mathbb{E}_{\bar{\theta}}(\boldsymbol{\pi}_a). \end{aligned}$$

With this, the Lagrangian is

$$\mathbb{E}_{\bar{\theta}}^2(\mathbf{u}_a) + \mathbb{E}_{\bar{\theta}}^2(\boldsymbol{\pi}_a) + \lambda \left(\mathbb{E}_{\bar{\theta}}(\mathbf{u}_a) - \left(\bar{\theta}_0 + \left(\bar{\theta}_{1\mathbf{a}} + \hat{\beta}^* \right) \mathbb{E}_{\bar{\theta}}(\boldsymbol{\pi}_a) \right) \right)$$

and the first-order conditions are

$$2\mathbb{E}_{\bar{\theta}}(\mathbf{u}_a) = -\lambda \quad 2\mathbb{E}_{\bar{\theta}}(\boldsymbol{\pi}_a) = \lambda \left(\bar{\theta}_{1\mathbf{a}} + \hat{\beta}^* \right) \quad \mathbb{E}_{\bar{\theta}}(\mathbf{u}_a) = \bar{\theta}_0 + \left(\bar{\theta}_{1\mathbf{a}} + \hat{\beta}^* \right) \mathbb{E}_{\bar{\theta}}(\boldsymbol{\pi}_a).$$

By solving them we get

$$\mathbb{E}_{\bar{\theta}}(\boldsymbol{\pi}_a) = B(\bar{\theta}) \equiv -\frac{\bar{\theta}_0 \left(\hat{\beta}^* + \bar{\theta}_{1\mathbf{a}} \right)}{1 + \left(\hat{\beta}^* + \bar{\theta}_{1\mathbf{a}} \right)^2}.$$

Since $\mathbb{E}_{\bar{\theta}}(\boldsymbol{\pi}_a) = a$, the monetary authority's best reply is thus the policy $a = B(\bar{\theta})$. As a result, a policy and belief pair $(a^*, \delta(\bar{\theta}))$ is a self-confirming equilibrium if and only if

$$a^* = B(\bar{\theta}) \quad (\text{subjective best reply}) \quad (30)$$

and

$$\bar{\theta}_0 = \theta_0^* + (\theta_{1\mathbf{a}}^* - \bar{\theta}_{1\mathbf{a}}) a^* \quad (\text{confirmed beliefs}). \quad (31)$$

Simple algebra shows that this is the case if and only if

$$a^* = -\frac{\theta_0^* \left(\hat{\beta}^* + \bar{\theta}_{1\mathbf{a}} \right)}{1 + \left(\hat{\beta}^* + \theta_{1\mathbf{a}}^* \right) \left(\hat{\beta}^* + \bar{\theta}_{1\mathbf{a}} \right)} \quad (32)$$

and

$$\bar{\theta}_0 = \theta_0^* - \frac{\theta_0^* \left(\hat{\beta}^* + \bar{\theta}_{1\mathbf{a}} \right)}{1 + \left(\hat{\beta}^* + \theta_{1\mathbf{a}}^* \right) \left(\hat{\beta}^* + \bar{\theta}_{1\mathbf{a}} \right)} \left(\theta_{1\mathbf{a}}^* - \bar{\theta}_{1\mathbf{a}} \right), \quad (33)$$

which are the equilibrium relations (28) and (29) in the case of dogmatic beliefs.⁴²

Figure 5.3.1.a illustrates the previous heuristic argument when the true model is of Lucas-Sargent kind, so that θ_0^* is the natural rate of unemployment and $\theta_{1\mathbf{a}}^* = -\hat{\beta}^*$ (and so the true policy multiplier ξ_2^* is zero). Under this true model, policy a induces average unemployment

⁴²Note that, with the dogmatic value $\bar{\theta}_{1\mathbf{a}}$ of $\theta_{1\mathbf{a}}$ in place of its expectation $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})$, the dogmatic equilibrium relations are identical to the general ones. This is a consequence of the certainty equivalence principle stated in Proposition 3.

$\mathbb{E}_{\theta^*}(\mathbf{u}_a) = \theta_0^*$ and average inflation $\mathbb{E}_{\theta^*}(\pi_a) = a$. But a monetary authority with dogmatic belief $\delta(\bar{\theta})$ expects to observe the pair of long-run averages $(\mathbb{E}_{\bar{\theta}}(\mathbf{u}_a), a)$. This dogmatic belief is confirmed, and so condition (31) is satisfied, if $\mathbb{E}_{\bar{\theta}}(\mathbf{u}_a) = \theta_0^*$, that is, if the pair of average unemployment and average inflation lies on the vertical partial identification line with abscissa θ_0^* . The subjective best reply condition (30) is represented by the tangency between the (red) indifference curve and the (green) conjectured constraint, according to which an increase Δa in average inflation yields a $-\bar{\xi}_2 \Delta a$ decrease in average unemployment, where $\bar{\xi}_2 = \hat{\beta}^* + \bar{\theta}_{1a}$ is the conjectured multiplier.




Figure 5.3.1.a: Self-confirming equilibrium in a new-classical world.

When the dogmatic belief is such that $\bar{\theta}_{1a} = 0$ so that $\bar{\xi}_2 = \hat{\beta}^*$ becomes the conjectured multiplier, the monetary authority is “orthodox” Keynesian. See Figure 5.3.1.b.




Figure 5.3.1.b: Equilibrium with “orthodox” Keynesian beliefs.

The conjectured constraint is $\mathbb{E}_{\bar{\theta}}(\mathbf{u}_a) = \bar{\theta}_0 + \hat{\beta}^* \mathbb{E}_{\bar{\theta}}(\boldsymbol{\pi}_a)$. Its slope is the beta coefficient of the Phillips regression, which represents the trade-off between inflation and unemployment that the Keynesian authority believes to be systematically exploitable.

5.3.2 Policy activism and welfare

To complete our equilibrium analysis we need to compare the self-confirming equilibrium action with the objectively optimal one and to compute the resulting welfare loss.

To this end we need to consider the estimated policy multiplier $\xi_2 = \hat{\beta} + \theta_{1\mathbf{a}}$. The authority underestimates the multiplier when $\mathbb{E}_{\mu^*}(\xi_2) > \xi_2^*$ and overestimates it when $\mathbb{E}_{\mu^*}(\xi_2) < \xi_2^*$.⁴³ In structural terms, $\mathbb{E}_{\mu^*}(\xi_2) \geq \xi_2^*$ if and only if $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) \geq \theta_{1\mathbf{a}}^*$. For instance, when $\theta_{1\mathbf{a}}^*$ and $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})$ are positive this means that the multiplier is under/overestimated if and only if the direct impact of planned inflation on unemployment is over/underestimated.

The objectively optimal policy is

$$a^o = -\frac{\theta_0^* (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)}{1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)^2}. \quad (34)$$

It is immediate to see that $a^* = a^o$ if and only if $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) = \theta_{1\mathbf{a}}^*$, (and so $\mathbb{E}_{\mu^*}(\xi_2) = \xi_2^*$). The equilibrium action is objectively optimal when the monetary authority has a correct expected value of the estimated policy multiplier ξ_2 . More generally, next we show that policy hyperactivism characterizes authorities that overestimate the policy multiplier, while hypoactivism characterizes authorities that underestimate it.⁴⁴

Proposition 11 *Given a true model θ^* , for every self-confirming equilibrium (a^*, μ^*) ,*

- (i) $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) < \theta_{1\mathbf{a}}^*$ if and only if policy a^* is hyperactive, i.e., $a^* > a^o$;
- (ii) $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) = \theta_{1\mathbf{a}}^*$ if and only if policy a^* is objectively optimal, i.e., $a^* = a^o$;
- (iii) $\theta_{1\mathbf{a}}^* < \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) < -\hat{\beta}^*$ if and only if policy a^* is hypoactive, i.e., $0 < a^* < a^o$;
- (iv) $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) = -\hat{\beta}^*$ if and only if policy a^* is zero-target-inflation, i.e., $a^* = 0$.

For the monetary authority, both kinds of deviations from objective optimality, hyperactivism and hypoactivism, cause the same welfare loss. Indeed:

Proposition 12 *The welfare loss is $\ell(a^*, \theta^*) = (1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)^2)(a^* - a^o)^2$.*

In the next section we will illustrate this result with a few examples.

5.4 Policy dogmatism and its welfare consequences

5.4.1 Equilibria

Assume that the monetary authority has dogmatic equilibrium beliefs $\mu^* = \delta(\bar{\theta})$. A pair $(a^*, \delta(\bar{\theta})) \in A \times \Delta(\bar{\Theta})$ is self-confirming if and only if it satisfies relations (32) and (33). Two special cases are noteworthy.

⁴³Both ξ_2^* and $\mathbb{E}_{\mu^*}(\xi_2)$ are negative (Assumption 2), and so $\mathbb{E}_{\mu^*}(\xi_2) \geq \xi_2^*$ if and only if $|\mathbb{E}_{\mu^*}(\xi_2)| \leq |\xi_2^*|$.

⁴⁴Since $\xi_2^* \leq 0$ (Assumption 2), the cases considered in the proposition exhaust all possibilities. Also note that, since $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) \leq -\hat{\beta}^*$, $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) = -\hat{\beta}^*$ is equivalent to $\mu^*(\theta_{1\mathbf{a}} = -\hat{\beta}^*) = 1$.

New-classical authority Suppose the monetary authority believes that the policy multiplier is zero, i.e., $\bar{\theta}_{1\mathbf{a}} = -\bar{\theta}_{1\boldsymbol{\pi}}$. Since in equilibrium $\theta_{1\boldsymbol{\pi}}$ is identified by the slope of the Phillips regression, we have $\bar{\theta}_{1\mathbf{a}} = -\hat{\beta}^*$. Here the conjectured constraint is vertical at the baseline unemployment rate θ_0^* : the new-classical authority does not believe in any systematically exploitable trade-off between inflation and unemployment. A zero-target-inflation equilibrium policy results (Proposition 11-(iv)).

Keynesian authority Suppose the monetary authority believes that there is a fully exploitable trade-off between inflation and unemployment, i.e., $\bar{\theta}_{1\mathbf{a}} = 0$. Then, in equilibrium, the conjectured policy multiplier $\bar{\xi}_2^* = \hat{\beta}^*$ is strictly negative. A positive-target-inflation equilibrium policy results:

$$a^* = -\frac{\theta_0^* \hat{\beta}^*}{1 + \hat{\beta}^* (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)} \quad \text{and} \quad \bar{\theta} = \left(\theta_0^* \left(\frac{1 + \hat{\beta}^{*2}}{1 + \hat{\beta}^* (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)} \right), 0 \right). \quad (35)$$

By Proposition 11, such a policy is hyperactive if $\theta_{1\mathbf{a}}^* > 0$, hypoactive if $\theta_{1\mathbf{a}}^* < 0$, and objectively optimal if $\theta_{1\mathbf{a}}^* = 0$.

To sum up, the two equilibria feature new-classical nonintervention a la Friedman-Hayek and Keynesian activism, respectively. Regardless of the true model economy, such policy prescriptions emerge through suitable dogmatic beliefs.

5.4.2 A new-classical world

So far we did not fix a specific economic model. Now, by way of example, assume that a Lucas-Sargent model economy $\theta^* = (\theta_0^*, -\hat{\beta}^*) \in \tilde{\Theta}$ is the true model, with no systematically exploitable trade-off between inflation and unemployment. Then, the pair $(a^*, \delta(\bar{\theta}))$ is a self-confirming equilibrium if and only if $a^* = -\theta_0^* (\hat{\beta}^* + \bar{\theta}_{1\mathbf{a}})$ and $\bar{\theta}_0 = \theta_0^* (1 - (\hat{\beta}^* + \bar{\theta}_{1\mathbf{a}})^2)$. Hence, the policy and belief pair

$$\left(-\theta_0^* (\hat{\beta}^* + \bar{\theta}_{1\mathbf{a}}), \delta \left(\theta_0^* \left(1 - (\hat{\beta}^* + \bar{\theta}_{1\mathbf{a}})^2 \right), \bar{\theta}_{1\mathbf{a}} \right) \right)$$

is the dogmatic self-confirming equilibrium in a Lucas-Sargent model economy. By Proposition 11, policy a^* is hyperactive when $\bar{\theta}_{1\mathbf{a}} < \theta_{1\mathbf{a}}^*$ and objectively optimal when $\bar{\theta}_{1\mathbf{a}} = \theta_{1\mathbf{a}}^*$. The welfare loss is $\ell(a^*, \theta^*) = \theta_0^{*2} (\hat{\beta}^* + \bar{\theta}_{1\mathbf{a}})^2$.

Next we consider two different equilibria in this new-classical world according to the monetary authority's dogmatic beliefs.

New-classical authority Suppose the monetary authority correctly believes that there is no exploitable trade-off between inflation and unemployment, that is, $\mu^* = \delta(\bar{\theta}_0, -\hat{\beta}^*)$. The pair $(a^*, \delta(\bar{\theta}_0, -\hat{\beta}^*))$ is a self-confirming equilibrium if and only if $a^* = 0$ and $\bar{\theta}_0 = \theta_0^*$. As a result, the policy and belief pair

$$\left(0, \delta \left(\theta_0^*, -\hat{\beta}^* \right) \right) \quad (36)$$

is the new-classical self-confirming equilibrium. It features a zero-target-inflation policy, which is the objectively optimal policy (so, there is no welfare loss) as well as the fully revealing one that allows the authority to learn, in the long-run, the true coefficient θ_0^* .

Keynesian authority Suppose the monetary authority wrongly believes that there is a fully exploitable trade-off between inflation and unemployment, with say $\mu^* = \delta(\bar{\theta}_0, 0)$. The pair $(a^*, \delta(\bar{\theta}_0, 0))$ is a self-confirming equilibrium if and only if $a^* = -\theta_0^* \hat{\beta}^{*2}$ and $\bar{\theta}_0 = \theta_0^*(1 - \hat{\beta}^{*2})$. The policy and belief pair

$$\left(-\theta_0^* \hat{\beta}^{*2}, \delta\left(\theta_0^*(1 - \hat{\beta}^{*2}), 0\right)\right) \quad (37)$$

is thus a Keynesian self-confirming equilibrium. It features an hyperactive positive-target-inflation policy. Since it is not the objectively optimal policy, the monetary authority suffers a welfare loss $\ell(a^*, \theta^*) = (\theta_0^* \hat{\beta}^{*2})^2$.

5.4.3 A Keynesian world

What we noted above can be reversed as we consider the case of a Keynesian model economy where the policy multiplier ξ_2 is different from zero, i.e., the monetary authority may systematically reduce average unemployment. To consider a stark (although implausible) example, suppose that $\theta^* = (\theta_0^*, 0) \in \tilde{\Theta}$ is the true model, that is, there is a full systematically exploitable trade-off between inflation and unemployment because monetary policy does not affect expectations ($\theta_{1a}^* = 0$). A Keynesian authority makes the objectively optimal positive-inflation choice in equilibrium. A new-classical authority chooses zero inflation, an inferior outcome.

5.4.4 Welfare consequences

What are the welfare implications of incorrect beliefs under dogmatism? By way of example, we consider a new-classical authority in a Keynesian economy, as well as a Keynesian authority in a new-classical economy. The loss of a new-classical zero inflation policy in a Keynesian economy, with $\theta_{1a}^* = 0$, is $(\theta_0^* \hat{\beta}^{*2})^2$. It is the same loss of a Keynesian nonzero inflation policy (37) in a new-classical economy: a mistaken new-classical authority has the same lower welfare as a mistaken Keynesian one.

5.5 Policy agnosticism and a curious interplay

5.5.1 Equilibria

Suppose that the monetary authority is not dogmatic, but has instead a two-model belief.⁴⁵ Specifically, she is uncertain whether the true model is of the Lucas-Sargent or Samuelson-Solow kind and her self-confirming subjective belief μ^* assigns positive probability mass to just one specimen of each kind, so that the (subjective) support consists of two points: a

⁴⁵Cogley and Sargent (2005) and Cogley et al. (2007) study dynamic Bayesian policy problems where beliefs assign positive probability to three model economies (dynamic specifications of the two models we consider here and a third related model).

Lucas-Sargent model $(\theta_0^{ls}(\mu^*), -\hat{\beta}^*)$ and a Samuelson-Solow (Keynesian) model $(\theta_0^{ss}(\mu^*), 0)$. Denoting by $\mu_k^* \in [0, 1]$ the subjective weight of the latter model, we can write belief μ^* as

$$\mu^* = (1 - \mu_k^*) \delta \left(\theta_0^{ls}(\mu^*), -\hat{\beta}^* \right) + \mu_k^* \delta \left(\theta_0^{ss}(\mu^*), 0 \right). \quad (38)$$

Since $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) = -(1 - \mu_k^*) \hat{\beta}^*$, the expected multiplier is $\mathbb{E}_{\mu^*}(\xi_2) = \mu_k^* \hat{\beta}^*$ and the pair (a^*, μ^*) is a self-confirming equilibrium if and only if

$$a^* = -\frac{\theta_0^* \hat{\beta}^* \mu_k^*}{1 + \hat{\beta}^* \mu_k^* (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)} \quad (39)$$

and

$$\theta_0^{ls}(\mu^*) = \frac{\theta_0^*}{1 + \hat{\beta}^* (\hat{\beta}^* + \theta_{1\mathbf{a}}^*) \mu_k^*}, \quad \theta_0^{ss}(\mu^*) = \frac{\theta_0^* (1 + \hat{\beta}^{*2} \mu_k^*)}{1 + \hat{\beta}^* (\hat{\beta}^* + \theta_{1\mathbf{a}}^*) \mu_k^*}. \quad (40)$$

As a result, in this case, a pair of the form

$$\left(-\frac{\theta_0^* \hat{\beta}^* \mu_k^*}{1 + \hat{\beta}^* \mu_k^* (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)}, (1 - \mu_k^*) \delta \left(\frac{\theta_0^*}{1 + \hat{\beta}^* (\hat{\beta}^* + \theta_{1\mathbf{a}}^*) \mu_k^*}, -\hat{\beta}^* \right) + \mu_k^* \delta \left(\frac{\theta_0^* (1 + \hat{\beta}^{*2} \mu_k^*)}{1 + \hat{\beta}^* (\hat{\beta}^* + \theta_{1\mathbf{a}}^*) \mu_k^*}, 0 \right) \right)$$

is a self-confirming equilibrium for every $\mu_k^* \in [0, 1]$. We thus have a continuum of equilibria parameterized by the subjective weight μ_k^* of the model of the Samuelson-Solow kind (and so by the expected multiplier $\mu_k^* \hat{\beta}^*$). In particular, the equilibrium policy a^* is increasing in μ_k^* : the higher the weight of the Keynesian model, the higher the planned inflation. If $\mu_k^* = 0$ we get back to the dogmatic new-classical equilibrium, while if $\mu_k^* = 1$ we get back to the dogmatic Keynesian equilibrium (Section 5.4.1).

In equilibrium, the coefficients (40) of the models of the Lucas-Sargent and Samuelson-Solow kind depend on the authority's subjective weight μ_k^* : different weights correspond to different Lucas-Sargent and Samuelson-Solow equilibrium specifications. Though the support of the equilibrium belief (38) always contains a specimen of both classes of model economies, that specimen changes as the weight μ_k^* changes. Finally, the welfare loss is

$$\ell(a^*, \theta^*) = \frac{\theta_0^{*2} (\hat{\beta}^* \mu_k^* + \hat{\beta}^* + \theta_{1\mathbf{a}}^*)^2}{\left(1 + \hat{\beta}^* \mu_k^* (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)\right)^2 \left(1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)^2\right)}. \quad (41)$$

This curious interplay between the models deemed possible and the weight on each kind of model is our main finding for the two-model self-confirming belief; therefore, it will be further clarified in a prominent special case.

5.5.2 A new-classical world

Assume that a Lucas-Sargent model economy $\theta^* = (\theta_0^*, -\hat{\beta}^*)$ is the true model. If so, by (39) and (40) the pair (a^*, μ^*) is a self-confirming equilibrium if and only if $a^* = -\theta_0^* \hat{\beta}^* \mu_k^*$,

$\theta_0^{ls}(\mu^*) = \theta_0^*$ and $\theta_0^{ss}(\mu^*) = \theta_0^*(1 + \hat{\beta}^{*2} \mu_k^*)$. Hence, in this case, the pair

$$\left(-\theta_0^* \hat{\beta}^* \mu_k^*, (1 - \mu_k^*) \delta(\theta_0^*, -\hat{\beta}^*) + \mu_k^* \delta(\theta_0^*(1 + \beta^{*2} \mu_k^*), 0)\right)$$

is a self-confirming equilibrium for every subjective weight $\mu_k^* \in [0, 1]$. The welfare loss is $\ell(a^*, \theta^*) = (\theta_0^* \hat{\beta}^* \mu_k^*)^2$.

As implied by the analysis of Section 5.5.1, we have a continuum of equilibria parameterized by the weight μ_k^* of the model of the Keynesian (Samuelson-Solow) kind: if $\mu_k^* > 0$ the equilibrium policy is hyperactive, if $\mu_k^* = 0$ we get the dogmatic new-classical equilibrium (36). Moreover, if $\mu_k^* = 1$ we get back to the dogmatic Keynesian equilibrium (37). Now, however, the equilibrium coefficient $\theta_0^{ls}(\mu^*)$ is pinned down by the true natural rate of unemployment θ_0^* : the monetary authority understands that, if the true model were of the Lucas-Sargent kind, average unemployment and baseline unemployment at 0-planned inflation would coincide; furthermore, in the case under consideration the average rate of unemployment must be the natural rate. In contrast, the subjective equilibrium coefficient $\theta_0^{ss}(\mu^*) = \theta_0^*(1 + \hat{\beta}^{*2} \mu_k^*)$ still depends on weight μ_k^* : a higher subjective weight of the Samuelson-Solow specification corresponds to a higher planned inflation in equilibrium, hence, to a higher Phillips regression line, whose horizontal intercept is $\theta_0^{ss}(\mu^*)$. Thus, the support of the equilibrium belief always contains a specimen of the Samuelson-Solow model; it, however, changes as μ_k^* changes. More generally, a two-model belief is determined by its (subjective) support and the relative likelihoods of the two models in the support. The self-confirming equilibrium conditions jointly discipline these two aspects of the belief.

Figure 5.5.2 illustrates. The monetary authority is uncertain about the true economic constraint, the vertical line at the natural rate of unemployment or the Phillips regression line. Since the true model is of the Lucas-Sargent kind, at a self-confirming equilibrium the average unemployment expected by the monetary authority must be the natural rate θ_0^* ; the subjective best reply condition is expressed by the tangency between the (red) indifference curve and a (green) line describing the expected constraint, the slope of which is intermediate between the vertical line at the natural rate θ_0^* and the Phillips regression line (which, in turn, depends on weight μ_k^* via the equilibrium relation $\theta_0 = \theta_0^*(1 + \hat{\beta}^{*2} \mu_k^*)$). Comparing the correct-belief equilibrium $(a, \mu) = (0, \delta(\theta^*))$ with the represented self-confirming equilibrium determined by $\mu_k^* > 0$, one can see that the latter features higher planned inflation $a^* > 0$

and higher horizontal intercept $\theta_0^{SS}(\mu^*) > \theta_0^*$.

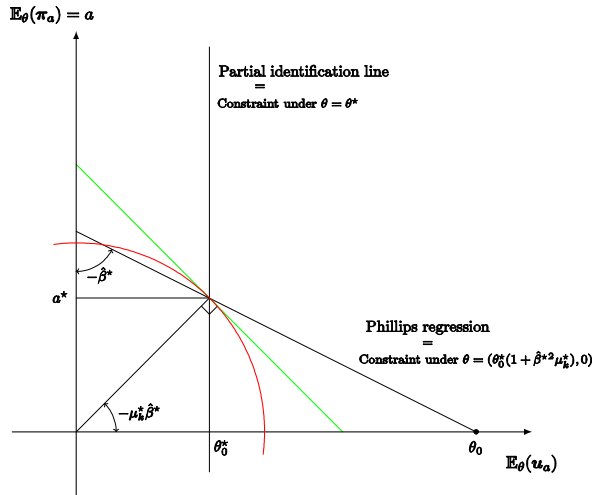


Figure 5.5.2.a: Two-Model self-confirming equilibrium.

Figure 5.5.2.b gives an alternative geometrical representation. Fix the true model θ^* and an alternative model θ . Every policy a induces a pair of objective expected rewards, the reward under model θ^* , $R(a, \theta^*)$, and the reward under model θ , $R(a, \theta)$. By changing a one obtains the locus of possible pairs of rewards. If $R(a, \theta^*) \neq R(a, \theta)$, the monetary authority can infer which of the two models is true from the observed long-run average payoff. Therefore, the partial identification condition is $R(a, \theta^*) = R(a, \theta)$. At a self-confirming equilibrium (a^*, μ^*) with $\text{supp } \mu^* = \{\theta^*, \theta\}$, this belief-confirmation condition must hold; therefore, the equilibrium point $(R(a^*, \theta^*), R(a^*, \theta))$ is at the intersection of the main diagonal in the $(R(\cdot, \theta^*), R(\cdot, \theta))$ -space, the “partial identification line,” with the locus of feasible pairs $\{(R(a, \theta^*), R(a, \theta)) : a \in A\}$, the constraint. At this intersection point, the constraint

curve must be tangent to the constant-SEU line with slope $(1 - \mu_k^*)/\mu_k^*$.




Figure 5.5.2.b: Self-confirming equilibrium in the expected rewards space.

Recall that $B(\cdot)$ denotes the best reply function. With this, $R(B(\theta^*), \theta^*) = V(B(\theta^*), \delta(\theta^*)) > V(B(\mu^*), \mu^*)$. Indeed, $V(B(\theta^*), \delta(\theta^*)) > V(B(\mu^*), \delta(\theta^*)) = R(B(\mu^*), \theta^*)$ because $B(\mu^*)$ is not a best reply to θ^* . On the other hand, $R(B(\mu^*), \theta^*) = V(B(\mu^*), \mu^*)$ because $R(B(\mu^*), \cdot)$ is constant on the support of self-confirming belief μ^* (see Lemma 3 and Proposition 5). The correct-belief equilibrium $(B(\theta^*), \delta(\theta^*))$ features sharper beliefs than $(B(\mu^*), \theta^*)$. Therefore, this is an instance of Proposition 6: self-confirming equilibria with sharper beliefs yield higher values and lower losses.

6 Steady state policies

The analysis of the Phillips curve example suggests a general form of the policy problem that extends to our steady state setup the classic policy Tinbergen-Theil framework. This is the subject matter of this final section that builds upon, and sums up, what we did so far in the paper.

6.1 Setup

A steady state model economy is described via the *structural form* relation

$$\varphi(x, y, \eta, \theta) = 0 \tag{42}$$

where $x \in X$ is a (specified) exogenous variable, $y \in Y$ is an endogenous variable, $\eta \in H$ is a shock variable (that is, an unspecified exogenous variable), and $\theta \in \Theta$ is a structural parameter that indexes the model economy. All variables can be multidimensional (more generally, they can live in vector spaces).

The endogenous variable is determined within the model, given the values that the other variables take on (outside the model); that is, the endogenous variable y solves equation (42). To ease matters, suppose that the solutions of this equation are always unique, so that we can write the solution function

$$y = \psi(x, \eta, \xi)$$

which is called the *reduced form* of the model economy, where $\xi = g(\theta)$ is a reduced form parameter.⁴⁶

Example 5 A simple model economy consists of a system of linear equations $\theta_1 x + \theta_2 y = \eta$. The entries of the (square, for simplicity) matrices θ_1 and θ_2 are the parameters, and $\varphi(x, y, \theta) = \theta_1 x + \theta_2 y - \eta$. If matrix θ_2 is invertible, the reduced form is $y = \psi(x, \eta, \xi) = \xi_1 x + \xi_2 \eta$, where $\xi = g(\theta_1, \theta_2) = (-\theta_2^{-1} \theta_1, \theta_2^{-1})$. \blacktriangle

Suppose that the exogenous variable can be decomposed as $x = (a, \zeta) \in A \times Z$, where the variables a and ζ are, respectively, under and outside the control of the policy maker. The policy maker chooses policy a in order to affect the value of the endogenous variable, which we assume to be payoff relevant, and so we denote it by c . The reduced form becomes

$$c = \psi(a, \zeta, \eta, \xi)$$

The state space is $S = Z \times H \times \Theta$, the consequence space is C , and $\rho : A \times S \rightarrow C$ is the consequence function, with $\rho(a, \zeta, \theta) = \psi(a, \zeta, g(\theta))$.⁴⁷ Policy multipliers correspond to the derivative of ψ with respect to a (partial derivatives if a is multidimensional).

Example 6 In the Phillips curve example the structural form is (12)-(13), with $c = (u, \pi)$, $\eta = (\varepsilon, w)$, and $\theta = (\theta_0, \theta_{1\pi}, \theta_{1a}, \theta_2, \theta_3)$. The function φ is linear, given by

$$\varphi(u, \pi, a, w, \varepsilon, \theta) = - \begin{bmatrix} 1 & -\theta_{1\pi} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} u \\ \pi \end{bmatrix} + \begin{bmatrix} \theta_0 \\ 0 \end{bmatrix} + a \begin{bmatrix} \theta_{1a} \\ 1 \end{bmatrix} + \begin{bmatrix} \theta_2 & 0 \\ 0 & \theta_3 \end{bmatrix} \begin{bmatrix} w \\ \varepsilon \end{bmatrix}$$

There are no exogenous variables (e.g., government expenditures). The reduced form is (14)-(15), with ψ linear and given by

$$\psi(a, w, \varepsilon, \xi) = \begin{bmatrix} \xi_1 \\ 0 \end{bmatrix} + a \begin{bmatrix} \xi_2 \\ 1 \end{bmatrix} + \begin{bmatrix} \xi_4 & \xi_3 \\ 0 & \xi_5 \end{bmatrix} \begin{bmatrix} w \\ \varepsilon \end{bmatrix}$$

Finally, $g(\theta) = (\theta_0, \theta_{1\pi} + \theta_{1a}, \theta_{1\pi} \theta_3, \theta_2, \theta_3)$, ρ is given by (16), and ξ_2 is the policy multiplier.⁴⁸ \blacktriangle

The rest of the decision problem specification is as the one presented in Subsection 2.2, we report here the essential elements. The physical uncertainty about the states' realizations

⁴⁶We do not discuss here the conditions that, along with the uniqueness of solutions, ensure the existence of the reduced form representation (see, e.g., Rothenberg, 1971). The economic relevance of multiple solutions is discussed by Jovanovic (1989); in this case, we have a solution correspondence, and so a reduce form correspondence.

⁴⁷Here we tacitly assume that ψ and g are such that ρ satisfies Savage's Consequentialism.

⁴⁸Assumption 1 of the example here amounts to requiring that the function $\theta \mapsto g(\theta)$ be one-to-one, so that distinct structural parameters correspond to distinct reduced form parameters.

is described by probability models $\sigma \in \Delta(S)$. Policy makers posit a collection $\Sigma \subseteq \Delta(S)$ of these distributions and, given any $\sigma \in \Sigma$, evaluate policy a via its expected utility $R(a, \sigma) = \int_S r(a, s) d\sigma$, where $r = v \circ \rho$ is a reward function based on a von Neumann-Morgenstern utility function $v : C \rightarrow \mathbb{R}$. Given his subjective prior μ over models σ , the policy maker then ranks actions according to the criterion $V(a, \mu) = \int_\Sigma R(a, \sigma) d\mu(\sigma)$.

6.2 Economic and stochastic model uncertainties

In the Phillips curve example we considered a state $s = (w, \varepsilon, \theta)$ in which we distinguish a shock pair (w, ε) and a model economy θ . The key factorization (18) builds on that distinction. This distinction also applies to the general policy problem, which features two types of model uncertainties. First, there is uncertainty about the economics of the phenomenon under consideration, and so about the economic models that explain it. Second, there is uncertainty about the statistical performance of such economic models, due to the errors that affect measurements and shocks; the latter represent the unexplained variation caused by, possibly many, minor explanatory variables that the policy maker is “unable and unwilling to specify”, as Marschak (1953) p. 12 remarks.⁴⁹

We can call, respectively, economic model uncertainty and stochastic model uncertainty the two types of uncertainty. The former is more fundamental than the latter since it reflects the economic views of policy makers.⁵⁰ This is why, by assuming q known, in the Phillips curve example we focused on economic model uncertainty.

In general, the state spaces relevant for policy problems can be represented as Cartesian products of factors that represent the two types of uncertainty. In this case, versions of the factorization (18) would apply with a twist: uncertainty about the probability model that describes the exogenous variable, absent in the Phillips curve example, must also be considered. We thus factor the probability models $\sigma \in \Sigma \subseteq \Delta(Z \times H \times \Theta)$ as $\sigma = q \times \delta_{\bar{\theta}}$, where – to focus on economic model uncertainty – we assume that the joint probability model $q \in \Delta(Z \times H)$ for the exogenous and shock variables is known and that $\delta_{\bar{\theta}} \in \Delta(\Theta)$ is a Dirac probability measure concentrated on a given economic model $\bar{\theta} \in \Theta$. As in the example, we can identify Σ with Θ without loss of generality. In particular, $R(a, \theta) = \int r(a, \zeta, \eta, \theta) dq(\zeta, \eta)$.

6.3 Partial identification and equilibrium

The analysis of partial identification and the notion of self-confirming equilibrium presents no any novelties for the generalized policy problem. But for the sake of completeness, here we briefly discuss them.

We assume that the messages that the policy maker receives are the values of the endogenous variable c , that is, $f = \rho$.⁵¹ Since we identify Σ with Θ , the distribution map

⁴⁹In a similar vein, Koopmans (1947) p. 169 writes that “... stochastic ... in that the behavior of any group of individuals, and the outcome of any production process, is determined in part by many minor factors, further scrutiny of which is either impossible or unrewarding.”

⁵⁰A main instance of what Denzau and North (1994) and, more recently, Rodrik (2014) call “ideas”. Note that economic model uncertainty subsumes the parametric uncertainty considered by Friedman (1953) and then Brainard (1967) in policy problems.

⁵¹Inter alia, this assumption implies that own-action independence of feedback cannot hold in nontrivial policy problems (recall our discussion in Section 3.2).

and the partial identification correspondence are given here by $\hat{f}_a(\theta)(E) = q(\rho_{a,\theta}^{-1}(E))$ and $\hat{\Sigma}_a(\theta) = \{\theta' \in \Sigma : \hat{f}_a(\theta') = \hat{f}_a(\theta)\}$.⁵² In particular, $\hat{\Sigma}_a(\theta)$ is the collection of all model economies that are observationally equivalent when θ is the true model and a is the chosen policy.

Moment conditions based on long-run observed values of the endogenous variable may allow the policy maker partially to identify the structural parameter θ , regardless of the chosen action. For instance, in the Phillips curve example the three coefficients $\theta_{1\pi}$, θ_2 and θ_3 were identified via suitable moment conditions (see Section 5.2.3). In this case, the parameter space can be factored as $\Theta = \tilde{\Theta} \times \{\hat{\theta}\}$, where $\hat{\theta}$ is the sub-vector estimated through long-run observations, independently of the chosen action (i.e., any learning dynamics should ensure the knowledge of $\hat{\theta}$). A lower dimensional parameter space $\tilde{\Theta}$ results and, consequently, the partial identification correspondence can be reduced to $\hat{\Sigma}_a : \tilde{\Theta} \rightarrow 2^{\tilde{\Theta}}$. In the Phillips curve example $\hat{\theta} = \{\hat{\beta}, \hat{\sigma}_{u|\pi}, \hat{\sigma}_u\}$ and $\tilde{\Theta} = \mathbb{R}_{++} \times (-\infty, -\hat{\beta}]$.

In view of all this, the notion of self-confirming equilibrium is easily stated for the policy problem: a policy and belief pair $(a^*, \mu^*) \in A \times \Delta(\tilde{\Theta})$ is self-confirming if and only if $a^* \in \arg \max_{a \in A} V(a, \mu^*)$ and $\mu^*(\hat{\Sigma}_{a^*}(\theta^*)) = 1$.

6.4 Steady state policy rules

A natural question is whether a notion of a steady state policy rule emerges in our setup. When Σ is a singleton, that is, when structural information allows the policy maker to know the true model, such a rule is given by the best reply correspondence $B(\theta) = \arg \max_{a \in A} R(a, \theta)$ for each $\theta \in \tilde{\Theta}$. A policy maker who knows the true model economy best replies to such knowledge by taking the objectively optimal policy. Since q is known, this corresponds in our steady state setup to the case originally studied by Tinbergen (1952) and Theil (1961).⁵³

The interplay between policies and information substantially complicates matters when Σ is nonsingleton. For simplicity we study the case, often considered in applications, when best replies to beliefs are unique. If so, self-confirming equilibria can be stated as equilibria in beliefs. Formally, say that a policy problem is *nice* if the best reply correspondence is a function $B : \Delta(\Sigma) \rightarrow A$, with $B(\mu) = \arg \max_{a \in A} V(a, \mu)$. In words, there is a unique best reply to each belief, as standard concavity conditions on reward functions would ensure. For instance, linear quadratic problems (like the Phillips curve example) are nice.

In a nice policy problem, given a true model $\theta^* \in \Sigma$, a pair $(a^*, \mu^*) \in A \times \Delta(\Sigma)$ of actions and beliefs is a self-confirming equilibrium if and only if $a^* = B(\mu^*)$ and

$$\mu^* \left(\hat{\Sigma}_{B(\mu^*)}(\theta^*) \right) = 1. \quad (43)$$

Hence, there is a unique equilibrium condition (43), which is cast in terms of beliefs (actions being pinned down by the best reply function). For this reason in nice policy problems we can view self-confirming equilibria as fixed points in the space of beliefs. To this end, define the correspondence $T_{\theta^*} : \Delta(\Sigma) \rightarrow 2^{\Delta(\Sigma)}$ by

$$T_{\theta^*}(\mu) = \Delta \left(\hat{\Sigma}_{B(\mu)}(\theta^*) \right).$$

⁵² $\rho_{a,\theta} : Z \times H \rightarrow C$ is the section of ρ at (a, θ) and $E \subseteq C$ is any measurable subset.

⁵³ In the terminology of Brainard (1967), a singleton Σ means a nonrandom parameter θ (see also Blinder, 1998, p. 11). A nonrandom θ is assumed, for example, also by Poole (1970) in his classic IS-LM policy analysis under uncertainty (see, e.g., Poole, 1970, p. 215).

Any belief $\nu \in \Delta(\hat{\Sigma}_{B(\mu)}(\theta^*))$ is consistent with the long-run frequency distribution of observed values of the endogenous variable. This suggests the following notion.

Definition 3 *Given a true model $\theta^* \in \Sigma$, a belief $\mu \in \Delta(\Sigma)$ is self-confirming if $\mu \in T_{\theta^*}(\mu)$.*

In words, self-confirming beliefs are consistent with the long-run data – that is, the observed values of the endogenous variable – gathered through the optimal policy that they induce.⁵⁴

Let $M(\theta^*) = \{\mu \in \Delta(\Sigma) : \mu \in T_{\theta^*}(\mu)\}$ be the collection of all self-confirming beliefs when θ^* is the true model. The *steady state policy rule* $h_{\theta^*} : M(\theta^*) \rightarrow A$ is the restriction of the best reply function on $M(\theta^*)$, that is, $h_{\theta^*}(\mu) = B(\mu)$ for all $\mu \in M(\theta^*)$. According to this rule, policy makers best reply to beliefs that, in the long run, are consistent with the data that they collect through their policies.

6.5 Brainard conservatism?

When the monetary authority knows the true model θ^* , the policy rule is

$$B(\theta^*) = -\frac{\theta_0^* (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)}{1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)^2}. \quad (44)$$

For each θ^* the rule prescribes the objectively optimal policy a^o . When, instead, the monetary authority does *not* know the true model θ^* , by Proposition 10 the set $M(\theta^*)$ consists of the beliefs $\mu \in \Delta(\Sigma)$ such that

$$\mu \left(\left\{ (\theta_0, \theta_{1\mathbf{a}}) \in \mathbb{R}^2 : \theta_0 = \theta_0^* - \frac{\theta_0^* (\hat{\beta}^* + \mathbb{E}_\mu(\theta_{1\mathbf{a}}))}{1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*) (\hat{\beta}^* + \mathbb{E}_\mu(\theta_{1\mathbf{a}}))} (\theta_{1\mathbf{a}}^* - \theta_{1\mathbf{a}}) \right\} \right) = 1.$$

This leads to the following characterization of steady state policy rules.

Proposition 13 *The steady state policy rule $h_{\theta^*} : M(\theta^*) \rightarrow \mathbb{R}$ satisfies the certainty equivalent principle, that is,*

$$h_{\theta^*}(\mu) = -\frac{\mathbb{E}_\mu(\theta_0) (\hat{\beta}^* + \mathbb{E}_\mu(\theta_{1\mathbf{a}}))}{1 + (\hat{\beta}^* + \mathbb{E}_\mu(\theta_{1\mathbf{a}}))^2}. \quad (45)$$

If μ is not dogmatic, then

$$h_{\theta^*}(\mu) = \begin{cases} \sqrt{\frac{\text{Var}_\mu(\theta_0)}{\text{Var}_\mu(\theta_{1\mathbf{a}})}} & \text{if } \text{Cov}_\mu(\theta_0, \theta_{1\mathbf{a}}) \neq 0 \\ 0 & \text{if } \text{Cov}_\mu(\theta_0, \theta_{1\mathbf{a}}) = 0 \end{cases}.$$

⁵⁴This fixed point characterization is easily seen to hold in full generality in our setup (beyond the policy case considered here) when best replies are unique.

Our steady state policies thus do not feature the “Brainard conservatism principle,” that is, they are not more prudent than what the certainty equivalent principle would prescribe.⁵⁵ The partial identification caused by the authority’s limited structural information does not translate into a policy conservatism a la Brainard: but the long-run data confirmation condition that characterizes our policy rules puts the certainty equivalent principle back in business.

The policy rule depends on the ratio of standard deviations $\sqrt{\text{Var}_\mu(\theta_0)/\text{Var}_\mu(\theta_{1\mathbf{a}})}$ when beliefs are not dogmatic (and so such deviations are not zero). Higher conjectured variability in θ_0 , the (systematic) rate of unemployment in absence of policy interventions, leads to higher policy activism. The opposite is true when it is the conjectured variability in $\theta_{1\mathbf{a}}$, and so in the policy multiplier, that becomes higher. However, within the certainty equivalent principle, in our steady state setup different types of parametric uncertainty determine different policy responses, some more aggressive and some more conservative. In particular, higher multiplier uncertainty makes policies more conservative; in this different sense, a form of Brainard conservatism does hold in steady state.

7 Concluding remarks

While applied theorists and economists more generally can benefit from seeing the self-confirming equilibrium concept in action, we think it is important to frame such applications within the context of an abstract analysis. Indeed, this allows to better understand key essential concepts like partial identification given the equilibrium choice, endogeneity of feedback about the state, and the role of observability of consequences. In this paper we put forward an abstract framework for the analysis of self-confirming policies amenable to economic applications, hence featuring both intrinsic randomness and (possibly) infinite spaces of actions and states. All the concepts and techniques can be extended to n -person games, but we focus on decision problems with uncertainty (i.e., one-person games with incomplete information) for several reasons. First, the analysis is simpler and it suffices for our examples and the economic policy applications. Second, it clarifies that self-confirming equilibrium is a *genuine* equilibrium concept also in a one-person setting, because equilibrium beliefs are disciplined by choice-dependent evidence. This should be contrasted with Harsanyi’s (1967) Bayesian equilibrium whereby subjective beliefs about unknown parameters are not disciplined by evidence; thus, in one-person settings Bayesian equilibrium just requires that the decision maker best reply to her subjective belief. Finally, we are not aware of simple and interesting n -person generalizations of our new comparative welfare results, Propositions 6 and 8. Our economic policy applications illustrate the abstract framework and extend previous work in several ways. In particular, they take a more neutral perspective on the true model economy and they consider general beliefs rather than dogmatic ones. Besides the n -person case, several other extensions of the selfconfirming equilibrium idea are conceivable. Here we consider a few that we find worth exploring.

⁵⁵Blinder (1997) names this principle after the classic finding of Brainard (1967) that, even in linear quadratic problems, parametric uncertainty may make policies more prudent in this sense. See also Soderstrom (2002) and Bernanke (2007) for a discussion of the principle.

Ambiguity aversion It is possible to allow for non-neutral attitudes toward perceived ambiguity,⁵⁶ e.g., by considering the smooth ambiguity model of Klibanoff et al. (2005), or the maxmin model of Gilboa and Schmeidler (1989). This is done in a companion paper (Battigalli et al. 2021). Here we give a hint of why such extension is immaterial in the examples and application of this paper. Go back to Figure 5.5.2.b. Choices are represented as vectors of objective expected rewards. The best-reply condition requires that the set of feasible vectors is separated by the upper-contour set of vectors preferred to the chosen one, which under ambiguity aversion is convex. The key observation is that, in our examples and application, every undominated feasible vector is on the “efficient” boundary of the convex hull of the feasible vectors, i.e., it is not dominated by convex combinations of feasible vectors. By an intuitive application of the separating hyperplane theorem, this means that if a feasible vector is a best reply under ambiguity aversion, then it is also a best reply under ambiguity neutrality (subjective expected utility maximization), with the upper half-space delimited by the separating hyperplane as upper-contour set.

Prospect Theory It would be also natural to extend the selfconfirming equilibrium idea and its applications to prospect theory models *à la* Kahneman and Tversky (1979) – see Wakker (2010) for an extensive treatment. The exercise is natural, but also challenging. On one hand, the equilibrium payoff is a natural (endogenous) reference point for the prospect theory analysis of selfconfirming equilibria. On the other hand, including the long-run empirical information represented by the partially identified set $\hat{\Sigma}_a(\sigma)$ in a prospect theory model is less immediate than doing it in a smooth ambiguity or in a maxmin model. A possibility is to require the distortion functions featured by prospect theory (for gains and losses) to affect a specific model in $\hat{\Sigma}_a(\sigma)$. An alternative route is to consider “smooth ambiguity like” versions of prospect theory *à la* Vinogradov (2013), and require the equilibrium prior to be supported on $\hat{\Sigma}_a(\sigma)$ (this yields the previous approach if the decision maker’s prior is a Dirac measure at some point in $\hat{\Sigma}_a(\sigma)$). In any case, the problem definitely deserves more attention, and presents an avenue for future research. The works of Peter Wakker on prospect theory provide a starting point for this intriguing companion quest.

Motivated beliefs Kunda (1990) wrote an influential paper on how motivation influences reasoning. Since then, “motivated beliefs” became an important topic in psychology and also in economics, as exemplified in the Introduction by Epley and Gilovich to an interesting symposium on this topic in the 2016 summer issue of the *Journal of Economic Perspectives*. In their contribution, Benabou and Tirole (2016) cite reports of how agents neglect negative information, distort it, or choose not to obtain important information at little or no cost. Such behavior is explained in economic models where agents’ utility directly depends on their posterior beliefs and agents take this into account in forming their (action-dependent) beliefs. The self-confirming equilibrium (SCE) idea instead posits agents who take information more seriously and exploit all the information they obtain, given their choices. Despite such clear differences, SCE can be combined with belief-dependent motivations to explain an important stylized fact studied by the motivated-beliefs literature, i.e., the reluctance to acquire materially useful and cheap information (see Mannahan 2021). Consider a decision maker (DM) with a prior belief μ over probabilistic models parameterized by an unknown

⁵⁶That is, lack of certainty about the objective probabilities of consequences of choices.

personal trait $\theta \in \Theta \subset \mathbb{R}$, such as her intelligence, general ability, or health. Let μ' denote her realized posterior belief, conditional on the received message (material outcome) given her action. To fix ideas, let the DM's "psychological utility" (Battigalli and Dufwenberg, forthcoming) be the sum of a standard utility function and an ego-utility component that depends on a posterior estimate of the unknown trait: $\bar{v}(m, a, \mu') = v(m, a) + e(\mathbb{E}_{\mu'}(\theta))$, where $e(\cdot)$ is an increasing function. The decision maker can either choose a status-quo action a^* that yields a known (or learned) distribution of material outcomes, or an alternative action a^t (e.g., taking a test) that yields a θ -dependent distribution of material outcomes, and that would teach her about her trait. It may well be the case that, absent the ego-utility component, the best choice (possibly the dominant one) would be to take the test. But if function $e(\cdot)$ is concave, the expected variability of the posterior estimate $\mathbb{E}_{\mu'}(\theta)$ may make taking the test too "ego-risky" for the DM. In this case, the status-quo action a^* would be an SCE action. This is somewhat similar to the preference for the status quo in an SCE under smooth-ambiguity aversion (Battigalli et al. 2015), with concavity of the "second-order utility" replaced by concavity of ego-utility.

A Appendix

A.1 Additional mathematical preliminaries

For the theory of standard Borel spaces we refer to Mackey (1957) and Kechris (2012). First of all, note that if X is a standard Borel space, then \mathcal{X} contains all the singletons. To prove it, it suffices to take a metric d on X that generates the Borel sigma algebra: $\mathcal{X} = \mathcal{B}(\mathcal{T}_d)$. For each $x \in X$ and all $n \in \mathbb{N}$, each open ball $B_{1/n}(x) = \{x' \in X : d(x, x') < 1/n\}$ belongs to the topology: $B_{1/n}(x) \in \mathcal{T}_d \subseteq \mathcal{X}$. Then

$$\{x\} = \bigcap_{n \in \mathbb{N}} B_{1/n}(x) \in \mathcal{X}.$$

In particular, if X is countable, then \mathcal{X} must be its power set.

Since several sigma algebras may be involved in the proofs, given a standard Borel space (X, \mathcal{X}) , we sometimes write \mathcal{B}_X instead of \mathcal{X} to denote its sigma algebra.

Lemma 5 *Let X and Y be measurable spaces and let $\varphi : X \rightarrow Y$ be measurable. Then $\hat{\varphi} : \Delta(X) \rightarrow \Delta(Y)$ is measurable with respect to the natural sigma algebra.*

Proof The natural sigma algebra $\mathcal{B}_{\Delta(Y)}$ of $\Delta(Y)$ is generated by the sets of the form

$$\{\nu \in \Delta(Y) : \nu(C) \leq c\}$$

for all $C \in \mathcal{B}_Y$ and $c \in \mathbb{R}$. Now, for all such sets

$$\begin{aligned} \hat{\varphi}^{-1}(\{\nu \in \Delta(Y) : \nu(C) \leq c\}) &= \{\xi \in \Delta(X) : \hat{\varphi}(\xi) \in \{\nu \in \Delta(Y) : \nu(C) \leq c\}\} \\ &= \{\xi \in \Delta(X) : (\xi \circ \varphi^{-1})(C) \leq c\} \\ &= \{\xi \in \Delta(X) : \xi(\varphi^{-1}(C)) \leq c\}, \end{aligned}$$

which belongs to $\mathcal{B}_{\Delta(X)}$ because the measurability of φ guarantees that $\varphi^{-1}(C) \in \mathcal{B}_X$, and $\mathcal{B}_{\Delta(X)}$ is generated by the sets of the form $\{\xi \in \Delta(X) : \xi(B) \leq b\}$ for all $B \in \mathcal{B}_X$ and $b \in \mathbb{R}$. Therefore, $\hat{\varphi}$ is measurable. \blacksquare

Note that this lemma does not require the measurable spaces (X, \mathcal{X}) or (Y, \mathcal{Y}) to be standard Borel, but rather hinges on the choice of the natural sigma algebras on $\Delta(X)$ and $\Delta(Y)$.

Lemma 6 *Let X and Y be standard Borel spaces and let $\varphi : X \rightarrow Y$ be measurable. If φ is one-to-one, then $\hat{\varphi}$ is one-to-one. In this case:*

- $\varphi(X) \in \mathcal{Y}$, hence $(\varphi(X), \mathcal{Y} \cap \varphi(X))$ is a standard Borel space;
- $\varphi : X \rightarrow \varphi(X)$ is a measurable isomorphism;
- $\mathcal{X} = \varphi^{-1}(\mathcal{Y})$, that is, φ generates \mathcal{X} ;
- $\hat{\varphi} : \Delta(X) \rightarrow \Delta(\varphi(X))$ is a measurable isomorphism (under the identification of $\hat{\varphi}(\xi)$ on \mathcal{Y} with its restriction to $\mathcal{Y} \cap \varphi(X)$).

In particular, the following three statements are equivalent:

- (i) $\varphi : X \rightarrow Y$ is one-to-one;
- (ii) $\hat{\varphi} : \Delta(X) \rightarrow \Delta(Y)$ is one-to-one;
- (iii) φ generates \mathcal{X} .

Proof The proof is based on Mackey's Monomorphism Theorem (henceforth MMT, see, Mackey, 1957, Theorem 3.2, and Kechris, 1994, Corollary 15.2).

MMT *If $\varphi : X \rightarrow Y$ is an injective and measurable map between two standard Borel spaces X and Y , then $\varphi(X) \in \mathcal{B}_Y$ and φ is a measurable isomorphism between (X, \mathcal{B}_X) and $(\varphi(X), \mathcal{B}_Y \cap \varphi(X))$. In particular, $\varphi(B) \in \mathcal{B}_Y$ for all $B \in \mathcal{B}_X$.*

By MMT, for each $B \in \mathcal{B}_X$, since $\varphi(B) \in \mathcal{B}_Y$, then $B = \varphi^{-1}(\varphi(B)) \in \varphi^{-1}(\mathcal{B}_Y)$, whence $\mathcal{B}_X \subseteq \varphi^{-1}(\mathcal{B}_Y)$; and the converse inclusion follows from the measurability of $\varphi : X \rightarrow Y$. Thus φ generates \mathcal{B}_X . This proves the first three points of the statement.

If $\xi, \xi' \in \Delta(X)$, then $\hat{\varphi}(\xi) = \hat{\varphi}(\xi')$ if and only if $\xi(\varphi^{-1}(C)) = \xi'(\varphi^{-1}(C))$ for all $C \in \mathcal{B}_Y$, that is, if and only if ξ and ξ' coincide on the sigma algebra $\varphi^{-1}(\mathcal{B}_Y)$ generated by φ . But $\varphi^{-1}(\mathcal{B}_Y) = \mathcal{B}_X$, then $\xi(B) = \xi'(B)$ for all $B \in \mathcal{B}_X$, thus $\xi = \xi'$. Therefore, $\hat{\varphi}$ is one-to-one.

If $\hat{\varphi} : \Delta(X) \rightarrow \Delta(Y)$ is one-to-one, since it is measurable and the spaces are standard Borel, by MMT, it follows that $\hat{\varphi}(\Delta(X))$ is a Borel subset of $\Delta(Y)$ and $\hat{\varphi}$ is a measurable isomorphism between $(\Delta(X), \mathcal{B}_{\Delta(X)})$ and $(\hat{\varphi}(\Delta(X)), \mathcal{B}_{\Delta(Y)} \cap \hat{\varphi}(\Delta(X)))$.

With this, every element $\nu = \xi \circ \varphi^{-1}$ of $\hat{\varphi}(\Delta(X))$ is a probability measure on \mathcal{B}_Y , $\varphi(X) \in \mathcal{B}_Y$, and $\nu(\varphi(X)) = \xi(\varphi^{-1}(\varphi(X))) = \xi(X) = 1$. Thus, the restriction $\iota(\nu)$ of an element $\nu \in \hat{\varphi}(\Delta(X))$ to the sigma algebra $\mathcal{B}_Y \cap \varphi(X)$ of $\varphi(X)$ is an element of $\Delta(\varphi(X))$. That is, the map

$$\begin{array}{ccc} \iota : \hat{\varphi}(\Delta(X)) & \rightarrow & \Delta(\varphi(X)) \\ \nu & \mapsto & \iota(\nu) \end{array}$$

that associates to each $\nu \in \hat{\varphi}(\Delta(X))$ with its restriction $\iota(\nu)$ to $\mathcal{B}_Y \cap \varphi(X)$ is well defined.

We want to show that, ι is a measurable isomorphism between $(\hat{\varphi}(\Delta(X)), \mathcal{B}_{\Delta(Y)} \cap \hat{\varphi}(\Delta(X)))$ and $(\Delta(\varphi(X)), \mathcal{B}_{\Delta(\varphi(X))})$.⁵⁷

By MMT again, it is sufficient to prove that it is bijective and measurable (since both spaces are standard Borel).

Measurability: First note that $\mathcal{B}_{\Delta(\varphi(X))}$ is generated by the sets of the form $\{\lambda \in \Delta(\varphi(X)) : \lambda(D) \leq d\}$ for all $D \in \mathcal{B}_{\varphi(X)} = \mathcal{B}_Y \cap \varphi(X) (\subseteq \mathcal{B}_Y)$ and $d \in \mathbb{R}$. Now, for all such sets

$$\begin{aligned} \iota^{-1}(\{\lambda \in \Delta(\varphi(X)) : \lambda(D) \leq d\}) &= \{\nu \in \hat{\varphi}(\Delta(X)) : \iota(\nu) \in \{\lambda \in \Delta(\varphi(X)) : \lambda(D) \leq d\}\} \\ &= \{\nu \in \hat{\varphi}(\Delta(X)) : \nu(D) \leq d\} \\ (\text{since } D \in \mathcal{B}_{\varphi(X)} = \mathcal{B}_Y \cap \varphi(X)) &= \{\nu \in \hat{\varphi}(\Delta(X)) : \nu(D) \leq d\} \\ &= \{\nu \in \Delta(Y) : \nu(D) \leq d\} \cap \hat{\varphi}(\Delta(X)) \in \mathcal{B}_{\Delta(Y)} \cap \hat{\varphi}(\Delta(X)) \end{aligned}$$

that is, ι is measurable.

Injectivity: Let $\nu, \nu' \in \hat{\varphi}(\Delta(X))$ be such that $\iota(\nu) = \iota(\nu')$. For all $C \in \mathcal{B}_Y$,

$$\nu(C \cap \varphi(X)) \leq \nu(C) = \nu(C \cap \varphi(X)) + \nu(C \cap \varphi(X)^c) \leq \nu(C \cap \varphi(X)) + \nu(\varphi(X)^c) = \nu(C \cap \varphi(X)),$$

that is, $\nu(C) = \nu(C \cap \varphi(X))$ and $C \cap \varphi(X) \in \mathcal{B}_Y \cap \varphi(X) = \mathcal{B}_{\varphi(X)}$, then $\nu(C) = \iota(\nu)(C \cap \varphi(X))$. Since the same considerations apply to ν' , it follows that

$$\nu(C) = \iota(\nu)(C \cap \varphi(X)) = \iota(\nu')(C \cap \varphi(X)) = \nu'(C),$$

and so $\nu = \nu'$.

Surjectivity: Next we show that, for every $\lambda \in \Delta(\varphi(X))$, the set function $\xi_\lambda(B) = \lambda(\varphi(B))$ for all $B \in \mathcal{B}_X$ belongs to $\Delta(X)$ and $\iota(\hat{\varphi}(\xi_\lambda)) = \lambda$. This is sufficient for surjectivity of ι because then $\lambda = \iota(\hat{\varphi}(\xi_\lambda)) \in \iota(\hat{\varphi}(\Delta(X)))$. First observe that $\xi_\lambda : \mathcal{B}_X \rightarrow [0, 1]$ is well defined because, by MMT, for all $B \in \mathcal{B}_X$, we have $\varphi(B) \in \mathcal{B}_Y \cap \varphi(X)$, and $\lambda : (\mathcal{B}_Y \cap \varphi(X)) \rightarrow [0, 1]$. Moreover, denoting by $\psi : \varphi(X) \rightarrow X$ the inverse isomorphism of φ , for every $B \in \mathcal{B}_X$, $\varphi(B) = \psi^{-1}(B)$.⁵⁸

Thus $\xi_\lambda(B) = \lambda(\varphi(B)) = \lambda(\psi^{-1}(B))$ is a probability measure on X . Finally, for every $D \in \mathcal{B}_Y \cap \varphi(X)$,

$$\iota(\hat{\varphi}(\xi_\lambda))(D) = \hat{\varphi}(\xi_\lambda)(D) = \xi_\lambda(\varphi^{-1}(D)) = \lambda(\varphi(\varphi^{-1}(D))) = \lambda(D)$$

as wanted.

So far, we have shown that, if φ is injective, then the map

$$\begin{aligned} \tilde{\varphi} = \iota \circ \hat{\varphi} : \Delta(X) &\rightarrow \Delta(\varphi(X)) \\ \xi &\mapsto (\xi \circ \varphi^{-1})|_{\mathcal{B}_Y \cap \varphi(X)} \end{aligned}$$

⁵⁷When this is done,

$$\iota \circ \hat{\varphi} : \Delta(X) \rightarrow \hat{\varphi}(\Delta(X)) \rightarrow \Delta(\varphi(X))$$

that associates to each $\xi \in \Delta(X)$ the restriction of $\hat{\varphi}(\xi)$ to $\mathcal{B}_Y \cap \varphi(X)$ is the isomorphism we were looking for.

⁵⁸If $y \in \varphi(B)$, then $y = \varphi(x)$ for some $x \in B$, then $\psi(y) = x \in B$ and $y \in \psi^{-1}(B)$. Conversely, if $y \in \varphi(X)$ is such that $y \in \psi^{-1}(B)$, then $y = \varphi(x)$ for some $x \in X$ and $\psi(y) \in B$, but $\psi(y) = x$, hence $x \in B$ and $y = \varphi(x) \in \varphi(B)$.

is an isomorphism of standard Borel spaces; and that for all $\lambda \in \Delta(\varphi(X))$, the set function $\xi_\lambda = \lambda \circ \varphi$, defined by $(\lambda \circ \varphi)(B) = \lambda(\varphi(B))$ for all $B \in \mathcal{B}_X$, belongs to $\Delta(X)$ and $\lambda = \iota(\hat{\varphi}(\xi_\lambda)) = \tilde{\varphi}(\lambda \circ \varphi)$; that is,

$$\tilde{\varphi}^{-1}(\lambda) = \lambda \circ \varphi$$

for all $\lambda \in \Delta(\varphi(X))$.

This proves the fourth point of the statement.

Finally, we already proved that if φ is injective, then φ generates \mathcal{B}_X , and that if φ generates \mathcal{B}_X , then $\hat{\varphi}$ is injective. The proof is concluded by showing that, if $\hat{\varphi}$ is injective, so is φ . We will actually prove the contrapositive statement.

Recall that $\delta : X \rightarrow \Delta(X)$ is the Dirac embedding of X into $\Delta(X)$. For each $x \in X$, we have $\hat{\varphi}(\delta(x)) = \delta(\varphi(x))$. In fact, for all $C \in \mathcal{B}_Y$,

$$\hat{\varphi}(\delta(x))(C) = \delta(x)(\varphi^{-1}(C)) = \begin{cases} 1 & \text{if } x \in \varphi^{-1}(C), \text{ i.e. } \varphi(x) \in C \\ 0 & \text{if } x \notin \varphi^{-1}(C), \text{ i.e. } \varphi(x) \notin C \end{cases} = \delta(\varphi(x))(C). \quad (46)$$

Therefore, if φ is not one-to-one, $\hat{\varphi}$ is not one-to-one. In fact, if there are $x \neq z$ in X such that $\varphi(x) = \varphi(z)$, then

$$\hat{\varphi}(\delta(x)) = \delta(\varphi(x)) = \delta(\varphi(z)) = \hat{\varphi}(\delta(z)),$$

but $\delta(x)$ and $\delta(z)$ are different probability measures on \mathcal{X} because the latter contains all singletons. ■

At the opposite side of the spectrum we have the case of constant functions.

Lemma 7 *Let X and Y be standard Borel spaces and let $\varphi : X \rightarrow Y$ be measurable. Then φ is constant if and only if $\hat{\varphi}$ is constant.*

Proof If $\varphi \equiv \bar{y}$ is constant, then given any $\sigma \in \Delta(X)$ and any $C \in \mathcal{Y}$,

$$\begin{aligned} \hat{\varphi}(\sigma)(C) &= \sigma(\varphi^{-1}(C)) = \sigma(\{x \in X : \varphi(x) \in C\}) \\ &= \sigma(\{x \in X : \bar{y} \in C\}) = \delta(\bar{y})(C) \end{aligned}$$

irrespective of σ .

The converse is proved by contraposition. If φ is not constant, then there exist $x \neq y$ such that $\varphi(x) \neq \varphi(y)$, but then

$$\hat{\varphi}(\delta(x)) = \delta(\varphi(x)) \neq \delta(\varphi(y)) = \hat{\varphi}(\delta(y))$$

where the two external equalities follow from (46) and the internal inequality by the fact that $\varphi(x) \neq \varphi(y)$ and singletons are measurable in C . ■

Let (X, \mathcal{X}) , (Y, \mathcal{Y}) , and (Z, \mathcal{Z}) be measurable spaces, $f : X \rightarrow Y$ and $g : Y \rightarrow Z$, and $h : X \rightarrow Z$. It is convenient to denote by $\mathcal{F} = f^{-1}(\mathcal{Y}) \subseteq \mathcal{X}$, $\mathcal{G} = g^{-1}(\mathcal{Z}) \subseteq \mathcal{Y}$, and $\mathcal{H} = h^{-1}(\mathcal{Z}) \subseteq \mathcal{X}$ the sigma algebras generated by f, g , and h , respectively. Moreover, we will say that h is *f-measurable* if it is \mathcal{F} - \mathcal{Z} -measurable.

Lemma 8 Let (X, \mathcal{X}) be a measurable space, and let (Y, \mathcal{Y}) and (Z, \mathcal{Z}) be standard Borel spaces. Then the following conditions are equivalent for two measurable functions $f : X \rightarrow Y$ and $h : X \rightarrow Z$:

- (i) h is f -measurable —that is, $h^{-1}(\mathcal{Z}) \subseteq f^{-1}(\mathcal{Y})$;
- (ii) there exists a measurable function $g : Y \rightarrow Z$ such that $h = g \circ f$.

Proof See, e.g., Chow and Teicher, 1997, Theorem 1.4.4. ■

The next result completes Lemma 1.

Corollary 2 Let $(\Theta, \mathcal{B}_\Theta)$ and (S, \mathcal{B}_S) be Borel spaces and fix a map $p : \Theta \rightarrow \Delta(S)$. Then p is measurable if and only if

$$\forall B \in \mathcal{B}_S, \forall b \in \mathbb{R}, \quad \{\theta \in \Theta : p(\theta)(B) \leq b\} \in \mathcal{B}_\Theta,$$

that is, $\theta \mapsto p(\theta)(B)$ is measurable for all $B \in \mathcal{B}_S$. If, moreover, p is one-to-one, then:

- $p(\Theta) \in \mathcal{B}_{\Delta(S)}$;
- $p : \Theta \rightarrow p(\Theta)$ is a measurable isomorphism;
- $\tilde{p} : \Delta(\Theta) \rightarrow \Delta(p(\Theta))$ defined by $\tilde{p}(\mu)(B) = \mu(p^{-1}(B))$ ($B \in \mathcal{B}_{p(\Theta)}$) is a measurable isomorphism and, for every $\lambda \in \Delta(p(\Theta))$, the inverse image of λ through \tilde{p} is $\lambda \circ p$.

Proof Since $\mathcal{B}_{\Delta(S)}$ is the sigma algebra generated by the functions $\phi_B : \Delta(S) \rightarrow \mathbb{R}$ defined by $\phi_B(\xi) = \xi(B)$ for all $B \in \mathcal{B}_S$, a map $p : \Theta \rightarrow \Delta(S)$ is measurable if and only if $\phi_B \circ p$ is measurable for all $B \in \mathcal{B}_S$ (see, e.g., Berberian, 1997, Proposition 1.3.8). But, given any $B \in \mathcal{B}_S$, $p(B | \theta) = p(\theta)(B) = (\phi_B \circ p)(\theta)$ for all $\theta \in \Theta$, thus $p(B | \cdot) = \phi_B \circ p$, proving the first part of the statement.⁵⁹ The rest follows from the statement of Lemma 6 setting $X = \Theta$, $Y = \Delta(S)$, and $\varphi = p$, with the exception of the explicit expression $\tilde{p}^{-1}(\lambda) = \lambda \circ p$, for which the last paragraph of the proof of Lemma 6 has to be inspected. ■

Corollary 3 Let $(\Theta, \mathcal{B}_\Theta)$ and (T, \mathcal{B}_T) be Borel spaces and fix $q \in \Delta(T)$. Then

$$\begin{aligned} p : \Theta &\rightarrow \Delta(T \times \Theta) \\ \theta &\mapsto q \times \delta(\theta) = p(\theta) \end{aligned}$$

is measurable and one-to-one.

Proof Injectivity is obvious, so we have only to show that

$$\forall B \in \mathcal{B}_{T \times \Theta}, \forall b \in \mathbb{R}, \quad \{\theta \in \Theta : p(\theta)(B) \leq b\} \in \mathcal{B}_\Theta,$$

that is, $\theta \mapsto q \times \delta(\theta)(B)$ is measurable for all $B \in \mathcal{B}_T \times \mathcal{B}_\Theta$. Now for each $\theta \in \Theta$,

$$q \times \delta(\theta)(B) = \int_{\Theta} q(B^\eta) d\delta(\theta)(\eta) = q(B^\theta),$$

where $B^\theta = \{t \in T : (t, \theta) \in B\}$, and a crucial step in the proof of the Fubini-Tonelli Theorem (see, e.g., Billingsley, 2012, p. 246) consists precisely in showing that the map $\theta \mapsto q(B^\theta)$ is measurable for all $B \in \mathcal{B}_T \times \mathcal{B}_\Theta$. ■

⁵⁹Notice that this part does not rely on the fact that the measurable spaces $(\Theta, \mathcal{B}_\Theta)$ or (S, \mathcal{B}_S) are Borel, but rather on the choice of the natural sigma algebra on $\Delta(S)$.

Corollary 4 *Let (S, \mathcal{B}_S) be a Borel space. Then: $\delta(S) \in \mathcal{B}_{\Delta(S)}$, $\delta : S \rightarrow \delta(S)$ is a measurable isomorphism, and $\lambda \mapsto \lambda \circ \delta$ is a measurable isomorphism between $\Delta(\delta(S))$ and $\Delta(S)$.*

Proof In order to apply the previous Corollary 2 with $\Theta = S$ and $p = \delta$, we have only to verify that $\{s \in S : \delta(s)(B) \leq b\} \in \mathcal{B}_S$ for all $B \in \mathcal{B}_S$ and $b \in \mathbb{R}$; but this follows from the fact that $\{s \in S : \delta(s)(B) \leq b\} = \{s \in S : 1_B(s) \leq b\}$ and indicators of measurable sets are measurable functions. ■

A.2 Feedback and identification

First recall that, for each $a \in A$, $f_a : S \rightarrow M$ is measurable and so is $\hat{f}_a : \Delta(S) \rightarrow \Delta(M)$. Since $\Sigma \in \mathcal{B}_{\Delta(S)}$, and points are measurable in standard Borel spaces, then for every $\nu \in \Delta(M)$ the set

$$\left\{ \sigma' \in \Sigma : \hat{f}_a(\sigma') = \nu \right\} = \left\{ \sigma' \in \Delta(S) : \hat{f}_a(\sigma') = \nu \right\} \cap \Sigma \in \mathcal{B}_{\Delta(S)} \cap \Sigma = \mathcal{B}_\Sigma,$$

and so $\hat{\Sigma}_a(\sigma) = \left\{ \sigma' \in \Sigma : \hat{f}_a(\sigma') = \hat{f}_a(\sigma) \right\}$ is a measurable subset of both Σ and $\Delta(S)$ for all $\sigma \in \Sigma$.

For the next result, recall that $\rho : A \times S \rightarrow C$ is the consequence function and that the feedback functions considered here satisfy observability of consequences, which is a maintained assumption.

Lemma 9 *Let f and f' be feedback functions for a decision problem D . Then:*

- (i) ρ is coarser than f ;
- (ii) if f_a is one-to-one for every $a \in A$, then f' is coarser than f ;
- (iii) if f' is coarser than f , then $\hat{\Sigma}_a(\sigma) \subseteq \hat{\Sigma}'_a(\sigma)$ for all $(a, \sigma) \in A \times \Sigma$.

Proof (i) Since consequences are observable, for each action $a \in A$ there exists a measurable function $g_a : M \rightarrow C$ such that $\rho_a(s) = g_a(f_a(s))$ for all $s \in S$. (ii) By assumption, for each $a \in A$, $f_a : S \rightarrow M$ is Borel measurable and one-to-one. By Lemma 6, $f_a(S)$ is a Borel subset of M and $f_a : S \rightarrow f_a(S)$ is a Borel isomorphism. Then, the inverse function $f_a^{-1} : f_a(S) \rightarrow S$ is Borel measurable.⁶⁰ Arbitrarily choose $\bar{s} \in S$ and set

$$k_a(m) \equiv \begin{cases} f_a^{-1}(m) & m \in f_a(S), \\ \bar{s} & m \notin f_a(S). \end{cases}$$

It is easy to see that k_a defines a Borel measurable map from M to S such that, for every $s \in S$,

$$f'_a(s) = f'_a(f_a^{-1}(f_a(s))) = f'_a(k_a(f_a(s))) = (f'_a \circ k_a)(f_a(s)).$$

⁶⁰ *Caveat:* In the proof of Lemma 6, the inverse isomorphism $f_a^{-1} : f_a(S) \rightarrow S$ is denoted f'_a , but here f'_a is a section of the feedback function f' that is not an inverse of f .

Setting $h_a = f'_a \circ k_a : M \rightarrow M'$ yields the desired result. (iii) Let $(a, \sigma) \in A \times \Sigma$. For every $\sigma' \in \hat{\Sigma}_a(\sigma)$, $\sigma'(f_a^{-1}(B_M)) = \sigma(f_a^{-1}(B_M))$ for all $B_M \in \mathcal{B}_M$. But $h_a^{-1}(B_{M'}) \in \mathcal{B}_M$ for all $B_{M'} \in \mathcal{B}_{M'}$, then

$$\begin{aligned} \sigma' \left((f'_a)^{-1}(B_{M'}) \right) &= \sigma' \left((h_a \circ f_a)^{-1}(B_{M'}) \right) = \sigma' \left(f_a^{-1}(h_a^{-1}(B_{M'})) \right) \\ &= \sigma \left(f_a^{-1}(h_a^{-1}(B_{M'})) \right) = \sigma \left((h_a \circ f_a)^{-1}(B_{M'}) \right) = \sigma \left((f'_a)^{-1}(B_{M'}) \right) \end{aligned}$$

and $\sigma' \in \hat{\Sigma}'_a(\sigma)$. ■

A.3 Additional definitions

The *self-confirming (equilibrium) correspondence*

$$\Gamma : \Sigma \rightarrow 2^{A \times \Delta(\Sigma)}$$

associates to each possible true model σ^* the collection $\Gamma(\sigma^*)$ of its self-confirming equilibria (a^*, μ^*) . It is also convenient to consider the *(equilibrium) action correspondence*

$$\gamma : \Sigma \rightarrow 2^A$$

that associates each possible true model σ^* with the collection $\gamma(\sigma^*)$ of its *self-confirming (equilibrium) actions*, that is, actions a^* such that $(a^*, \mu^*) \in \Gamma(\sigma^*)$ for some belief μ^* .

A.4 Model uncertainty

Let $\mu^* \ll \nu^*$ denote the μ^* is absolutely continuous with respect to (informally, “sharper than”) ν^* . We show that self-confirming equilibria with sharper basic subjective assessments have higher values, hence —by Proposition 5— lower losses.

Proposition 14 *If $(a^*, \mu^*), (b^*, \nu^*) \in \Gamma(\sigma^*)$ and $\mu^* \ll \nu^*$, then $R(a^*, \sigma^*) = V(a^*, \mu^*) \geq V(b^*, \nu^*) = R(b^*, \sigma^*)$.*

Proof Since $\mu^* \left(\hat{\Sigma}_{a^*}(\sigma^*) \right) = 1$ and $\nu^* \left(\hat{\Sigma}_{b^*}(\sigma^*) \right) = 1$, then $\mu^* \ll \nu^*$ implies $\mu^* \left(\hat{\Sigma}_{b^*}(\sigma^*) \right) = 1$ and so $\mu^* \left(\hat{\Sigma}_{b^*}(\sigma^*) \cap \hat{\Sigma}_{a^*}(\sigma^*) \right) = 1$. The optimality condition (6) for a^* and Proposition 5 deliver

$$R(a^*, \sigma^*) = V(a^*, \mu^*) \geq \int_{\hat{\Sigma}_{a^*}(\sigma^*)} R(b^*, \sigma) d\mu^*(\sigma) = \int_{\hat{\Sigma}_{a^*}(\sigma^*) \cap \hat{\Sigma}_{b^*}(\sigma^*)} R(b^*, \sigma) d\mu^*(\sigma).$$

Since, by Lemma 3, $R(b^*, \sigma) = R(b^*, \sigma^*)$ for all $\sigma \in \hat{\Sigma}_{b^*}(\sigma^*)$, it follows that $V(a^*, \mu^*) \geq R(b^*, \sigma^*) = V(b^*, \nu^*)$, where the last equality follows from Proposition 5. ■

Priors μ^* and ν^* that are mutually absolutely continuous are called *equivalent*, denoted $\mu^* \sim \nu^*$; this means that they agree on what models are possible (or impossible). By the previous result, if $\mu^* \sim \nu^*$ then $V(a^*, \mu^*) = V(b^*, \nu^*)$ for all pairs of self-confirming equilibria $(a^*, \mu^*), (b^*, \nu^*) \in \Gamma(\sigma^*)$. The value of self-confirming equilibria is thus pinned down by what the decision maker deems possible, whereas the specific shape of the prior is value-irrelevant. But more is actually true: actions can be exchanged across such self-confirming equilibria.

Proposition 15 *If $(a^*, \mu^*), (b^*, \nu^*) \in \Gamma(\sigma^*)$ and $\mu^* \sim \nu^*$, then $(a^*, \nu^*), (b^*, \mu^*) \in \Gamma(\sigma^*)$.*

Proof As observed, $R(a^*, \sigma^*) = V(a^*, \mu^*) = V(b^*, \nu^*) = R(b^*, \sigma^*)$, but then

- $R(b^*, \sigma^*) = V(a^*, \mu^*) \geq V(a, \mu^*)$ for all $a \in A$ and $\mu^* \left(\hat{\Sigma}_{b^*}(\sigma^*) \right) = 1$ since $\nu^* \left(\hat{\Sigma}_{b^*}(\sigma^*) \right) = 1$;
- $R(a^*, \sigma^*) = V(b^*, \nu^*) \geq V(b, \nu^*)$ for all $b \in A$ and $\nu^* \left(\hat{\Sigma}_{a^*}(\sigma^*) \right) = 1$ since $\mu^* \left(\hat{\Sigma}_{a^*}(\sigma^*) \right) = 1$.

■

The results on the value that we just established for subjective model uncertainty extend to objective model uncertainty. In particular, self-confirming (equilibrium) actions with better identification properties have higher values, regardless of which confirmed beliefs support them.

Proposition 16 *If $(a^*, \mu^*), (b^*, \nu^*) \in \Gamma(\sigma^*)$ and $\hat{\Sigma}_{a^*}(\sigma^*) \subseteq \hat{\Sigma}_{b^*}(\sigma^*)$, then $V(a^*, \mu^*) \geq V(b^*, \nu^*)$.*

Proof The optimality condition (6) for a^* and Proposition 5 deliver

$$R(a^*, \sigma^*) = V(a^*, \mu^*) \geq \int_{\hat{\Sigma}_{a^*}(\sigma^*)} R(b^*, \sigma) d\mu^*(\sigma) = \int_{\hat{\Sigma}_{b^*}(\sigma^*)} R(b^*, \sigma) d\mu^*(\sigma)$$

but, by Lemma 3, $R(b^*, \sigma) = R(b^*, \sigma^*)$ for all $\sigma \in \hat{\Sigma}_{b^*}(\sigma^*)$, it follows that

$$V(a^*, \mu^*) \geq R(b^*, \sigma^*) = V(b^*, \nu^*)$$

where the last equality follows from Proposition 5. ■

Also observe that, if $\hat{\Sigma}_{a^*}(\sigma^*) = \hat{\Sigma}_{b^*}(\sigma^*)$, then $R(a^*, \sigma^*) = V(a^*, \mu^*) = V(b^*, \nu^*) = R(b^*, \sigma^*)$. With this, if $(a^*, \mu^*), (b^*, \nu^*) \in \Gamma(\sigma^*)$ then

- $R(b^*, \sigma^*) = V(a^*, \mu^*) \geq V(a, \mu^*)$ for all $a \in A$ and $\mu^* \left(\hat{\Sigma}_{b^*}(\sigma^*) \right) = 1$ since $\mu^* \left(\hat{\Sigma}_{a^*}(\sigma^*) \right) = 1$;
- $R(a^*, \sigma^*) = V(b^*, \nu^*) \geq V(b, \nu^*)$ for all $b \in A$ and $\nu^* \left(\hat{\Sigma}_{a^*}(\sigma^*) \right) = 1$ since $\nu^* \left(\hat{\Sigma}_{b^*}(\sigma^*) \right) = 1$.

Finally, the following results relate self-confirming equilibrium actions to objectively optimal actions:

Corollary 5 *A fully revealing action is self-confirming if and only if it is objectively optimal.*

Under own-action independence of feedback about the state, we have a stronger result. Eq. (5) and Lemma 3 imply:

Corollary 6 *Under own-action independence of feedback about the state, an action is self-confirming if and only if it is objectively optimal.*

Thus, from a decision perspective, own-action independence of feedback is equivalent to perfect feedback.

A.5 Other proofs

Proof of Lemma 1 See Lemma 6. ■

Proof of Proposition 1 See Lemma 8. ■

Proof of Proposition 2 See Lemma 9. ■

Proof of Lemma 3 Fix $a \in A$. Observability of consequences implies that $\rho_a(s) = g_a(f_a(s))$ for each $s \in S$, where $g_a : M \rightarrow C$ is $\mathcal{B}_M - \mathcal{B}_C$ -measurable; as $f_a : S \rightarrow M$ is $\mathcal{F}_a - \mathcal{B}_M$ -measurable, then $\rho_a : S \rightarrow C$ is $\mathcal{F}_a - \mathcal{B}_C$ -measurable. Moreover, $v : C \rightarrow \mathbb{R}$ is $\mathcal{B}_C - \mathcal{B}_{\mathbb{R}}$ -measurable and bounded above, and so $r_a = v \circ \rho_a : S \rightarrow \mathbb{R}$ is $\mathcal{F}_a - \mathcal{B}_{\mathbb{R}}$ -measurable and bounded above. Thus,

$$\forall \sigma \in \Delta(S), \quad R_a(\sigma) = \int_S r_a d\sigma = \int_S r_a d\sigma|_{\mathcal{F}_a}. \quad (47)$$

In particular, if $\sigma \in \Sigma$ and $\sigma' \in \hat{\Sigma}_a(\sigma)$, then $R_a(\sigma) = \int_S r_a d\sigma|_{\mathcal{F}_a} = \int_S r_a d\sigma'|_{\mathcal{F}_a} = R_a(\sigma')$. ■

Proof of Proposition 5 If $(a^*, \mu^*) \in A \times \Delta(\Sigma)$ and $\mu^*(\hat{\Sigma}_{a^*}(\sigma^*)) = 1$, then

$$V(a^*, \mu^*) = \int_{\Sigma} R(a^*, \sigma) d\mu^*(\sigma) = \int_{\hat{\Sigma}_{a^*}(\sigma^*)} R(a^*, \sigma) d\mu^*(\sigma) = R(a^*, \sigma^*),$$

because, by Lemma 3, $R(a^*, \sigma) = R(a^*, \sigma^*)$ for all $\sigma \in \hat{\Sigma}_{a^*}(\sigma^*)$. ■

Proof of Proposition 6 It follows immediately from Proposition 5 and Proposition 14. ■

Proof of Proposition 7 It follows immediately from Proposition 5 and Proposition 16. ■

Proof of Proposition 8 Suppose that $\hat{\Sigma}_a(\sigma^*) \subseteq \hat{\Sigma}_b(\sigma^*)$ for each $b \in A$. We already observed that if a is objectively optimal, then $(a, \delta(\sigma^*)) \in \Gamma(\sigma^*)$ and $a \in \gamma(\sigma^*)$. As for the converse, let $\mu^* \in \Delta(\Sigma)$ be such that $(a, \mu^*) \in \Gamma(\sigma^*)$. Since $\hat{\Sigma}_a(\sigma^*) \subseteq \hat{\Sigma}_b(\sigma^*)$ for each $b \in A$ and, by Lemma 3, for each b it is true that $R(b, \sigma) = R(b, \sigma^*)$ when $\sigma \in \hat{\Sigma}_b(\sigma^*)$, then $R(a, \sigma^*) \geq \int_{\hat{\Sigma}_a(\sigma^*)} R(b, \sigma) d\mu^*(\sigma) = R(b, \sigma^*)$, as wanted. ■

Proof of Corollary 5 Given a true model $\sigma^* \in \Sigma$, the result follows from Proposition 8 since if a is fully revealing, then $\hat{\Sigma}_a(\sigma^*) = \{\sigma^*\} \subseteq \hat{\Sigma}_{a'}(\sigma^*)$ for every $a' \in A$. ■

Proof of Corollary 6 Given a true model $\sigma^* \in \Sigma$, the result follows from Proposition 8 since own-action independence of feedback implies $\hat{\Sigma}_a(\sigma^*) = \hat{\Sigma}_{a'}(\sigma^*)$ for every $a, a' \in A$. Hence, $\gamma(\sigma^*) = \arg \max_{a \in A} R(a, \sigma^*)$. ■

Proof of Proposition 9 Recall that a is fixed. We first prove the inclusion \subseteq . If $\theta' \in \hat{\Sigma}_a(\theta)$, then $\hat{\rho}_a(q \times \delta(\theta')) = \hat{\rho}_a(q \times \delta(\theta))$. In particular,

$$\int_S h(\rho_a) d(q \times \delta(\theta')) = \int_S h(\rho_a) d(q \times \delta(\theta)), \quad (48)$$

for all $h : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ for which the integral is defined. Next observe that:

1. For $h(u, \pi) = \pi$ and $\theta'' \in \Theta$, we have that $\int_S \boldsymbol{\pi} d(q \times \delta_{\theta''}) = a$.
2. For $h(u, \pi) = \pi^2$ and $\theta'' \in \Theta$, we have that $\int_S \boldsymbol{\pi}^2 d(q \times \delta(\theta'')) = a^2 + (\theta_3'')^2$.
3. For $h(u, \pi) = u$ and $\theta'' \in \Theta$, we have that $\int_S \mathbf{u} d(q \times \delta(\theta'')) = \theta_0'' + (\theta_{1\boldsymbol{\pi}}'' + \theta_{1\mathbf{a}}'') a$.
4. For $h(u, \pi) = u^2$ and $\theta'' \in \Theta$, we have that $\int_S \mathbf{u}^2 d(q \times \delta(\theta'')) = (\theta_0'' + (\theta_{1\boldsymbol{\pi}}'' + \theta_{1\mathbf{a}}'') a)^2 + (\theta_{1\boldsymbol{\pi}}'')^2 (\theta_3'')^2 + (\theta_2'')^2$.
5. For $h(u, \pi) = u\pi$ and $\theta'' \in \Theta$, we have that $\int_S \mathbf{u}\boldsymbol{\pi} d(q \times \delta(\theta'')) = a(\theta_0'' + (\theta_{1\boldsymbol{\pi}}'' + \theta_{1\mathbf{a}}'') a) + \theta_{1\boldsymbol{\pi}}'' (\theta_3'')^2$.

Given (48), note that point 2 gives $\theta_3' = \theta_3$, then points 3 and 5 give $\theta_{1\boldsymbol{\pi}}' = \theta_{1\boldsymbol{\pi}}$. With this, point 3 again yields $\theta_0' + \theta_{1\mathbf{a}}' a = \theta_0 + \theta_{1\mathbf{a}} a$. Then point 4 gives

$$(\theta_0' + (\theta_{1\boldsymbol{\pi}}' + \theta_{1\mathbf{a}}') a)^2 + (\theta_{1\boldsymbol{\pi}}')^2 (\theta_3')^2 + (\theta_2')^2 = (\theta_0 + (\theta_{1\boldsymbol{\pi}} + \theta_{1\mathbf{a}}) a)^2 + (\theta_{1\boldsymbol{\pi}})^2 (\theta_3)^2 + (\theta_2)^2;$$

point 3 says that the first summands on both sides coincide, and we already established $(\theta_{1\boldsymbol{\pi}}')^2 (\theta_3')^2 = (\theta_{1\boldsymbol{\pi}})^2 (\theta_3)^2$; therefore, $\theta_2' = \theta_2$. This concludes the proof of the first set inclusion and formalizes the moments heuristics described in the main text.

In order to obtain the opposite inclusion, note that some simple algebra delivers, for each $\theta'' \in \Theta$ and each $(u, \pi) \in \mathbb{R}^2$,

$$\hat{\rho}_a(q \times \delta(\theta''))((-\infty, u] \times (-\infty, \pi]) = q(Q_{\theta''}),$$

$$\text{where } Q_{\theta''} = \left\{ (w, \varepsilon) \in W \times E : \begin{array}{l} w\theta_2'' + \varepsilon\theta_{1\boldsymbol{\pi}}''\theta_3'' \leq u - (\theta_0'' + a\theta_{1\mathbf{a}}'') - a\theta_{1\boldsymbol{\pi}}'' \\ \theta_3''\varepsilon \leq \pi - a \end{array} \right\}.$$

Now, consider $\theta' \in \Theta$ such that $\theta_0' + \theta_{1\mathbf{a}}' a = \theta_0 + \theta_{1\mathbf{a}} a$, $\theta_{1\boldsymbol{\pi}}' = \theta_{1\boldsymbol{\pi}}$, $\theta_2' = \theta_2$, $\theta_3' = \theta_3$, then $Q_{\theta'} = Q_{\theta}$, hence $\hat{\rho}_a(q \times \delta(\theta'))((-\infty, u] \times (-\infty, \pi]) = \hat{\rho}_a(q \times \delta(\theta))((-\infty, u] \times (-\infty, \pi])$ implying that $\hat{\rho}_a(q \times \delta(\theta')) = \hat{\rho}_a(q \times \delta(\theta))$ and $\theta' \in \hat{\Sigma}_a(\theta)$. \blacksquare

Proof of Lemma 4 Some simple algebra shows that

$$\begin{aligned} R(a, \theta) &= - \int_{W \times E} \mathbf{u}^2(a, w, \varepsilon, \theta) dq(w, \varepsilon) - \int_{W \times E} \boldsymbol{\pi}^2(a, w, \varepsilon, \theta) dq(w, \varepsilon) \\ &= -(\theta_0 + (\theta_{1\boldsymbol{\pi}} + \theta_{1\mathbf{a}}) a)^2 - a^2 - \theta_2^2 - \theta_3^2 \theta_{1\boldsymbol{\pi}}^2 - \theta_3^2 \\ &= -\mathbb{E}_{\theta}^2(\mathbf{u}_a) - \mathbb{E}_{\theta}^2(\boldsymbol{\pi}_a) - \theta_2^2 - \theta_3^2 \theta_{1\boldsymbol{\pi}}^2 - \theta_3^2 \\ &= v(\mathbb{E}_{\theta}(\mathbf{u}_a), \mathbb{E}_{\theta}(\boldsymbol{\pi}_a)) + \kappa, \end{aligned}$$

where, being $\tilde{\Theta} = \{(\theta_0, \theta_{1\mathbf{a}})\} = \mathbb{R}^2$, we set $\kappa = -\theta_2^2 - \theta_3^2 \theta_{1\boldsymbol{\pi}}^2 - \theta_3^2$ since this polynomial can be regarded as a constant term. \blacksquare

Proof of Proposition 10 It holds

$$R(a, \theta) = - \left((\theta_{1\boldsymbol{\pi}} + \theta_{1\mathbf{a}})^2 + 1 \right) a^2 - 2\theta_0 (\theta_{1\boldsymbol{\pi}} + \theta_{1\mathbf{a}}) a + \kappa.$$

Thus, $V(a, \mu^*)$ is —up to a constant— equal to

$$\begin{aligned}
& - \int_{\widehat{\Sigma}_{a^*}(\theta^*)} \left(\left((\hat{\beta}^* + \theta_{1\mathbf{a}})^2 + 1 \right) a^2 + 2\theta_0 (\hat{\beta}^* + \theta_{1\mathbf{a}}) a \right) d\mu^*(\theta) \\
&= - \int_{\mathbb{R}} \left(\left((\hat{\beta}^* + \theta_{1\mathbf{a}})^2 + 1 \right) a^2 + 2(\theta_0^* + (\theta_{1\mathbf{a}}^* - \theta_{1\mathbf{a}}) a^*) (\hat{\beta}^* + \theta_{1\mathbf{a}}) a \right) d\mu^*(\theta_{1\mathbf{a}}) \\
&= - \int_{\mathbb{R}} \left((\hat{\beta}^{*2} + \theta_{1\mathbf{a}}^2 + 2\hat{\beta}^* \theta_{1\mathbf{a}} + 1) a^2 + 2\theta_0^* (\hat{\beta}^* + \theta_{1\mathbf{a}}) a + 2a^* (\theta_{1\mathbf{a}}^* - \theta_{1\mathbf{a}}) (\hat{\beta}^* + \theta_{1\mathbf{a}}) a \right) d\mu^*(\theta_{1\mathbf{a}}) \\
&= - \int_{\mathbb{R}} \left((\hat{\beta}^{*2} + \theta_{1\mathbf{a}}^2 + 2\hat{\beta}^* \theta_{1\mathbf{a}} + 1) a^2 + 2\theta_0^* (\hat{\beta}^* + \theta_{1\mathbf{a}}) a + 2a^* (\theta_{1\mathbf{a}}^* \hat{\beta}^* + \theta_{1\mathbf{a}}^* \theta_{1\mathbf{a}} - \theta_{1\mathbf{a}} \hat{\beta}^* - \theta_{1\mathbf{a}}^2) a \right) d\mu^*(\theta_{1\mathbf{a}}) \\
&= - \left(\hat{\beta}^{*2} + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}^2) + 2\hat{\beta}^* \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) + 1 \right) a^2 - 2\theta_0^* (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) a \\
&\quad - 2a^* \left(\theta_{1\mathbf{a}}^* \hat{\beta}^* + \theta_{1\mathbf{a}}^* \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) \hat{\beta}^* - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}^2) \right) a.
\end{aligned}$$

The first order condition $\partial V(a, \mu^*) / \partial a = 0$ implies

$$\begin{aligned}
& a \left(\hat{\beta}^{*2} + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}^2) + 2\hat{\beta}^* \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) + 1 \right) + a^* \left(\theta_{1\mathbf{a}}^* \hat{\beta}^* + \theta_{1\mathbf{a}}^* \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) \hat{\beta}^* - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}^2) \right) \\
&= -\theta_0^* (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})).
\end{aligned}$$

Putting $a = a^*$ we get

$$a^* \left(\hat{\beta}^{*2} + \hat{\beta}^* \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) + 1 + \theta_{1\mathbf{a}}^* \hat{\beta}^* + \theta_{1\mathbf{a}}^* \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) \right) = -\theta_0^* (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}))$$

and so

$$\begin{aligned}
a^* &= \frac{-\theta_0^* (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}))}{\hat{\beta}^{*2} + \hat{\beta}^* \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) + 1 + \theta_{1\mathbf{a}}^* \hat{\beta}^* + \theta_{1\mathbf{a}}^* \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})} \\
&= -\frac{\theta_0^* (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}))}{1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*) (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}))}.
\end{aligned}$$

As a result, $\widehat{\Sigma}_{a^*}(\theta^*)$ is equal to

$$\left\{ (\theta_0, \theta_{1\mathbf{a}}) \in \mathbb{R}^2 : \theta_0 = \theta_0^* - \frac{\theta_0^* (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}))}{1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*) (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}))} (\theta_{1\mathbf{a}}^* - \theta_{1\mathbf{a}}) \right\},$$

as desired. ■

Proof of Proposition 11 It holds

$$\begin{aligned}
a^* - a^o &= \frac{\theta_0^* (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)}{1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)^2} - \frac{\theta_0^* (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}))}{1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*) (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}))} \\
&= \frac{\theta_0^* \left((\hat{\beta}^* + \theta_{1\mathbf{a}}^*) \left(1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*) (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) \right) - (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) \left(1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)^2 \right) \right)}{\left(1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)^2 \right) \left(1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*) (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) \right)} \\
&= \frac{\theta_0^* \left(\theta_{1\mathbf{a}}^* + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)^2 (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) - (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)^2 \right)}{\left(1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)^2 \right) \left(1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*) (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) \right)} \\
&= \frac{\theta_0^* (\theta_{1\mathbf{a}}^* - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}))}{\left(1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)^2 \right) \left(1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*) (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) \right)}.
\end{aligned}$$

Hence, if $a^* \neq 0$, it holds

$$a^* - a^o = - \frac{a^*}{1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)^2} \frac{\theta_{1\mathbf{a}}^* - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})}{\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})}$$

and so

$$a^* \geq a^o \iff a^* \frac{\theta_{1\mathbf{a}}^* - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})}{\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})} \leq 0. \quad (49)$$

Having established this relation, we can now prove points (i) and (iii) (points (ii) and (iv) being obvious).

(i) Suppose $a^* > a^o$. By (28) $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) \neq -\hat{\beta}^*$ and so $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) < -\hat{\beta}^*$. By (49), $(\theta_{1\mathbf{a}}^* - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) / (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) < 0$, which in turn implies $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) < \theta_{1\mathbf{a}}^*$. Conversely, suppose $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) < \theta_{1\mathbf{a}}^*$. Since $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) \leq -\hat{\beta}^*$, by (28) it follows $a^* > 0$. Moreover, being $(\theta_{1\mathbf{a}}^* - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) / (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) < 0$, by (49) it holds $a^* > a^o$. (iii) Suppose $0 < a^* < a^o$. By (49), $(\theta_{1\mathbf{a}}^* - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) / (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) > 0$, that is, $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) \in (\theta_{1\mathbf{a}}^*, -\hat{\beta}^*)$. Conversely, suppose $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) \in (\theta_{1\mathbf{a}}^*, -\hat{\beta}^*)$. By (28), $a^* > 0$. Moreover, being $(\theta_{1\mathbf{a}}^* - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) / (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) > 0$, by (49) it holds $a^* < a^o$. ■

The loss function can be defined in terms of beliefs by setting $\ell(\mu, \sigma) = \ell(B(\mu), \sigma)$. For instance, next we show that for the Phillips curve example it holds

$$\ell(\mu^*, \theta^*) = \frac{\theta_0^{*2} (\theta_{1\mathbf{a}}^* - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}))^2}{\left(1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)^2 \right) \left(1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*) (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) \right)^2}. \quad (50)$$

There is a zero welfare loss if and only if $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) = \theta_{1\mathbf{a}}^*$, that is, if and only if the monetary authority's expected value of the coefficient $\theta_{1\mathbf{a}}$ is correct. Otherwise, the loss is nonzero, as (50) shows.

Proof of Proposition 12 and eq. (50) First note that

$$R(a^\circ, \theta^*) = -\theta_0^{*2} - \left(\hat{\beta}^* + \theta_{1\mathbf{a}}^*\right)^2 a^{\circ 2} - \left(\hat{\beta}^* \theta_3^*\right)^2 - \theta_2^{*2} - 2\theta_0^* \left(\hat{\beta}^* + \theta_{1\mathbf{a}}^*\right) a^\circ - a^{\circ 2} - \theta_3^{*2},$$

and

$$R(a^*, \theta^*) = -\theta_0^{*2} - \left(\hat{\beta}^* + \theta_{1\mathbf{a}}^*\right)^2 (a^*)^2 - \left(\hat{\beta}^* \theta_3^*\right)^2 - \theta_2^{*2} - 2\theta_0^* \left(\hat{\beta}^* + \theta_{1\mathbf{a}}^*\right) a^* - (a^*)^2 - \theta_3^{*2}.$$

Hence,

$$\begin{aligned} \ell(a^*, \theta^*) &= \max_{a \in A} R(a, \theta^*) - R(a^*, \theta^*) = R(a^\circ, \theta^*) - R(a^*, \theta^*) \\ &= -\left(\hat{\beta}^* + \theta_{1\mathbf{a}}^*\right)^2 (a^{\circ 2} - a^{*2}) - 2\theta_0^* \left(\hat{\beta}^* + \theta_{1\mathbf{a}}^*\right) (a^\circ - a^*) - (a^{\circ 2} - a^{*2}). \end{aligned}$$

Suppose $a^\circ = 0$, that is, $\theta_0^* \left(\hat{\beta}^* + \theta_{1\mathbf{a}}^*\right) = 0$. Then

$$\ell(a^*, \theta^*) = \left(1 + \left(\hat{\beta}^* + \theta_{1\mathbf{a}}^*\right)^2\right) a^{*2} = -\left(1 + \left(\hat{\beta}^* + \theta_{1\mathbf{a}}^*\right)^2\right) \frac{\theta_0^{*2} \left(\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})\right)^2}{\left(1 + \left(\hat{\beta}^* + \theta_{1\mathbf{a}}^*\right) \left(\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})\right)\right)^2}.$$

If $\theta_0^* \neq 0$, then $\hat{\beta}^* + \theta_{1\mathbf{a}}^* = 0$, thus,

$$\ell(a^*, \theta^*) = \theta_0^{*2} \left(\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})\right)^2 = \theta_0^{*2} \left(\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) - \theta_{1\mathbf{a}}^*\right)^2. \quad (51)$$

If $\hat{\beta}^* + \theta_{1\mathbf{a}}^* \neq 0$, then $\theta_0^* = 0$, thus,

$$\ell(a^*, \theta^*) = 0. \quad (52)$$

Next suppose $a^\circ \neq 0$. It holds $-2a^\circ \left(1 + \left(\hat{\beta}^* + \theta_{1\mathbf{a}}^*\right)^2\right) = 2\theta_0^* \left(\hat{\beta}^* + \theta_{1\mathbf{a}}^*\right)$, thus,

$1 + \left(\hat{\beta}^* + \theta_{1\mathbf{a}}^*\right)^2 = -\theta_0^* \left(\hat{\beta}^* + \theta_{1\mathbf{a}}^*\right) / a^\circ$. Hence,

$$\begin{aligned} \ell(a^*, \theta^*) &= -\left(\hat{\beta}^* + \theta_{1\mathbf{a}}^*\right)^2 (a^{\circ 2} - a^{*2}) - 2\theta_0^* \left(\hat{\beta}^* + \theta_{1\mathbf{a}}^*\right) (a^\circ - a^*) - (a^{\circ 2} - a^{*2}) \\ &= -(a^\circ - a^*) \left[\left(\hat{\beta}^* + \theta_{1\mathbf{a}}^*\right)^2 (a^\circ + a^*) + 2\theta_0^* \left(\hat{\beta}^* + \theta_{1\mathbf{a}}^*\right) + a^\circ + a^* \right] \\ &= -(a^\circ - a^*) \left[\left(\hat{\beta}^* + \theta_{1\mathbf{a}}^*\right)^2 (a^\circ + a^*) - 2a^\circ \left(1 + \left(\hat{\beta}^* + \theta_{1\mathbf{a}}^*\right)^2\right) + a^\circ + a^* \right] \\ &= -(a^\circ - a^*) \left[\left(\hat{\beta}^* + \theta_{1\mathbf{a}}^*\right)^2 (a^* - a^\circ) + a^* - a^\circ \right] = -(a^\circ - a^*) \left[\left(\hat{\beta}^* + \theta_{1\mathbf{a}}^*\right)^2 + 1 \right] (a^* - a^\circ) \\ &= (a^\circ - a^*)^2 \left[\left(\hat{\beta}^* + \theta_{1\mathbf{a}}^*\right)^2 + 1 \right] = (a^* - a^\circ)^2 \left[\left(\hat{\beta}^* + \theta_{1\mathbf{a}}^*\right)^2 + 1 \right] \\ &= -(a^* - a^\circ)^2 \frac{\theta_0^* \left(\hat{\beta}^* + \theta_{1\mathbf{a}}^*\right)}{a^\circ} = -\theta_0^* \left(\hat{\beta}^* + \theta_{1\mathbf{a}}^*\right) \frac{(a^* - a^\circ)^2}{a^\circ} \\ &= \left(1 + \left(\hat{\beta}^* + \theta_{1\mathbf{a}}^*\right)^2\right) \left(a^* + \frac{\theta_0^* \left(\hat{\beta}^* + \theta_{1\mathbf{a}}^*\right)}{1 + \left(\hat{\beta}^* + \theta_{1\mathbf{a}}^*\right)^2} \right)^2. \end{aligned}$$

In the previous proof we showed that

$$a^* - a^o = \frac{\theta_0^* (\theta_{1\mathbf{a}}^* - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}))}{\left(1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)^2\right) \left(1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*) (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}))\right)}.$$

Hence,

$$\begin{aligned} \ell(\mu^*, \theta^*) &= -\theta_0^* (\hat{\beta}^* + \theta_{1\mathbf{a}}^*) \frac{(a^* - a^o)^2}{a^o} \\ &= \theta_0^* (\hat{\beta}^* + \theta_{1\mathbf{a}}^*) \frac{\theta_0^{*2} (\theta_{1\mathbf{a}}^* - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}))^2}{\left(1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)^2\right)^2 \left(1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*) (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}))\right)^2} \frac{1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)^2}{\theta_0^* (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)} \\ &= \frac{\theta_0^{*2} (\theta_{1\mathbf{a}}^* - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}))^2}{\left(1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)^2\right) \left(1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*) (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}))\right)^2}. \end{aligned}$$

It is easy to check that, along with (51) and (52), this completes the proof. ■

Proof of eq. (41) It holds

$$\begin{aligned} \ell(a^*, \theta^*) &= \left(1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)^2\right) \left(-\frac{\theta_0^* \hat{\beta}^* \mu_k^*}{1 + \hat{\beta}^* \mu_k^* (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)} + \frac{\theta_0^* (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)}{1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)^2} \right)^2 \\ &= \left(1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)^2\right) \left(\frac{-\theta_0^* \hat{\beta}^* \mu_k^* \left(1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)^2\right) + \theta_0^* (\hat{\beta}^* + \theta_{1\mathbf{a}}^*) \left(1 + \hat{\beta}^* \mu_k^* (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)\right)}{\left(1 + \hat{\beta}^* \mu_k^* (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)\right) \left(1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)^2\right)} \right)^2 \\ &= \frac{\theta_0^{*2} (\hat{\beta}^* \mu_k^* + \hat{\beta}^* + \theta_{1\mathbf{a}}^*)^2}{\left(1 + \hat{\beta}^* \mu_k^* (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)\right)^2 \left(1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)^2\right)}, \end{aligned}$$

as desired. ■

References

- [1] Battigalli, P., *Comportamento razionale ed equilibrio nei giochi e nelle situazioni sociali*, unpublished thesis, Università Bocconi, 1987.
- [2] Battigalli, P., and D. Guaitoli, Conjectural equilibria and rationalizability in a macro-economic game with incomplete information, *Quaderni di Ricerca*, Università Bocconi, 1988 (published in *Decisions, Games and Markets*, Kluwer, Dordrecht, 97-124, 1997).
- [3] Battigalli, P., and M. Dufwenberg, Belief-Dependent Motivations and Psychological Game Theory, *Journal of Economic Literature*, forthcoming.

- [4] Battigalli, P., S. Cerreia-Vioglio, F. Maccheroni, and M. Marinacci, Self-confirming equilibrium and model uncertainty, *American Economic Review*, 105, 646-677, 2015.
- [5] Battigalli, P., S. Cerreia-Vioglio, F. Maccheroni, and M. Marinacci, Self-confirming policies and ambiguity. mimeo, Bocconi University, 2021.
- [6] Battigalli, P., A. Francetich, G. Lanzani, and M. Marinacci, Learning and self-confirming long-run biases, *Journal of Economic Theory*, 183, 740-785, 2019.
- [7] Benabou, R., and J. Tirole, Mindful economics: The production, consumption, and value of beliefs, *Journal of Economic Perspectives* 30: 141-64, 2016.
- [8] Berberian, S. K., *Borel Spaces*, mimeo, 1997.
- [9] Bernanke, B. S., Monetary policy under uncertainty, *Annual Economic Policy Conference*, Federal Reserve Bank of St. Louis, 2007.
- [10] Billingsley, P., *Probability and Measure*, anniversary ed., Wiley, New York, 2012.
- [11] Blinder, A. S., What central bankers could learn from academics - and vice versa, *Journal of Economic Perspectives*, 11, 3-20, 1997.
- [12] Blinder, A. S., *Central banking in theory and practice*, MIT Press, Cambridge, 1998.
- [13] Brainard, W., Uncertainty and the effectiveness of policy, *American Economic Review P&P*, 57, 411-425, 1967.
- [14] Cerreia-Vioglio, S., F. Maccheroni, M. Marinacci, and L. Montrucchio, Classical subjective expected utility, *Proceedings of the National Academy of Sciences*, 110, 6754-6759, 2013.
- [15] Cho, I.-K., and K. Kasa, Learning and model validation, *Review of Economic Studies*, 82, 45-82, 2015.
- [16] Cho, I.-K., N. Williams, and T. Sargent, Escaping Nash inflation, *Review of Economic Studies*, 69, 1-40, 2002.
- [17] Chow, Y. S., and H. Teicher, *Probability theory: independence, interchangeability, martingales*, Springer, New York, 1997.
- [18] Cogley, T., R. Colacito, and T. J. Sargent, Benefits from US monetary policy experimentation in the days of Samuelson and Solow and Lucas, *Journal of Money, Credit and Banking*, 39, 67-99, 2007.
- [19] Cogley, T., and T. J. Sargent, The conquest of US inflation: learning and robustness to model uncertainty, *Review of Economic Dynamics*, 8, 528-563, 2005.
- [20] Denzau, A. T. and D. C. North, Shared mental models: ideologies and institutions, *Kyklos*, 47, 3-31, 1994.
- [21] Easley, D., and N. M. Kiefer, Controlling a stochastic process with unknown parameters, *Econometrica*, 5, 1045-1064, 1988.

- [22] Epley, N., and T. Gilovich, The mechanics of motivated reasoning, *Journal of Economic Perspectives* 30: 133-40, 2016.
- [23] Friedman, M., The effects of full employment policy on economic stability: a formal analysis, in *Essays in Positive Economics* (M. Friedman, ed.) 117-132, University of Chicago Press, Chicago, 1953.
- [24] Fudenberg, D., and D. K. Levine, Self-confirming equilibrium, *Econometrica*, 61, 523-545, 1993a.
- [25] Fudenberg, D., and D. K. Levine, Steady state learning and Nash equilibrium, *Econometrica*, 61, 547-573, 1993b.
- [26] Fudenberg, D., and D. K. Levine, Self-confirming equilibrium and the Lucas critique, *Journal of Economic Theory*, 144, 2354–2371, 2009.
- [27] Gaballo, G., and R. Marimon, Breaking the spell with credit easing: Self-confirming credit crises in competitive search economies, *Journal of Monetary Economics*, 119, 1-20, 2021.
- [28] Gilboa, I. and D. Schmeidler, Maxmin expected utility with a non-unique prior, *Journal of Mathematical Economics*, 18, 141-153, 1989.
- [29] Hansen, L. P., and M. Marinacci, Ambiguity aversion and model misspecification: An economic perspective, *Statistical Science*, 31, 511-515, 2016.
- [30] Harsanyi, J. (1967-68): Games of incomplete information played by Bayesian players. Part I, *Management Science*, 14, 159-182, 1967.
- [31] Jovanovic, B., Observable implications of models with multiple equilibria, *Econometrica*, 57, 1431-1437, 1989.
- [32] Kahneman, D. and A. Tversky, Prospect theory: An analysis of decision under risk, *Econometrica*, 47, 363-391, 1979.
- [33] Kalai, E., and E. Lehrer, Subjective equilibrium in repeated games, *Econometrica*, 61, 1231-1240, 1993.
- [34] Kalai, E., and E. Lehrer, Subjective games and equilibria, *Games and Economic Behavior*, 8, 123-163, 1995.
- [35] Kechris A., *Classical descriptive set theory*, Springer, New York, 2012.
- [36] Klibanoff, P., M. Marinacci and S. Mukerji, A smooth model of decision making under ambiguity, *Econometrica*, 73, 1849-1892, 2005.
- [37] Koopmans, T. C., Measurement without theory, *Review of Economic Statistics*, 29, 161-172, 1947.
- [38] Kunda, Z., The Case for Motivated Reasoning, *Psychological Bulletin*, 108, 480-498, 1990.

- [39] Kurz, M., On rational belief equilibria, *Economic Theory*, 4, 859-876, 1994a.
- [40] Kurz, M., On the structure and diversity of Rational Beliefs, *Economic Theory*, 4, 877-900, 1994b.
- [41] Lucas, R. E. Jr., Expectations and the neutrality of money, *Journal of Economic Theory*, 4, 103-124, 1972.
- [42] Lucas, R. E. Jr., and E. C. Prescott, Investment under uncertainty, *Econometrica*, 39, 659-681, 1971.
- [43] Mackey, G. W., Borel structure in groups and their duals, *Transactions of the American Mathematical Society*, 134-165, 1957.
- [44] Mannahan, R., Self-Esteem and rational self-handicapping, mimeo, University of Arizona, 2021.
- [45] Marinacci, M., Model uncertainty, *Journal of the European Economic Association*, 13, 998-1076, 2015.
- [46] Marschak, J., Economic measurements for policy and prediction, in *Studies in Econometric Method* (W. Hood and T. J. Koopmans, eds.), 1-26, Wiley, New York, 1953.
- [47] Poole, W., Optimal choice of monetary policy instruments in a simple stochastic macro model, *Quarterly Journal of Economics*, 84, 197-216, 1970.
- [48] Rodrik, D., When ideas trump interests: preferences, worldviews, and policy innovations, *Journal of Economic Perspectives*, 28, 189-208, 2014.
- [49] Rothenberg, T. J., Identification in parametric models, *Econometrica*, 39, 577-591, 1971.
- [50] Saint Paul, G., Economic science and political influence, *Journal of the European Economic Association*, 11, 1004-1031, 2013.
- [51] Saint Paul, G., The possibility of ideological bias in structural macroeconomic models, *American Economic Journal: Macroeconomics*, 10, 216-241, 2018.
- [52] Samuelson, P. A., and R. M. Solow, Analytical aspects of anti-inflation policy, *American Economic Review P&P*, 50, 177-184, 1960.
- [53] Sargent, T. J., Rational expectations, the real rate of interest, and the natural rate of unemployment, *Brookings Papers on Economic Activity*, 429-472, 1973.
- [54] Sargent, T. J., *The conquest of American inflation*, Princeton University Press, Princeton, 1999.
- [55] Sargent, T. J., Evolution and intelligent design, *American Economic Review*, 98, 3-37, 2008.
- [56] Sargent, T. J., and N. Williams, Impacts of priors on convergence and escapes from Nash inflation, *Review of Economic Dynamics*, 8, 360-391, 2005.

- [57] Sargent, T. J., N. Williams, and T. Zha, Shocks and government beliefs: the rise and fall of the American inflation, *American Economic Review*, 96, 1193-1224, 2006.
- [58] Savage, L. J., *The Foundations of Statistics*, Wiley, New York, 1954.
- [59] Soderstrom, U., Monetary policy with uncertain parameters, *Scandinavian Journal of Economics*, 104, 125-145, 2002.
- [60] Theil, H., *Economic forecasts and policy*, North Holland, Amsterdam, 1961.
- [61] Tinbergen, J., *On the theory of economic policy*, North Holland, Amsterdam, 1952.
- [62] Vinogradov, D., Weighting probabilities in ambiguity, mimeo, University of Glasgow, 2013.
- [63] Wakker, P. P., *Prospect Theory: For Risk and Ambiguity*, Cambridge University Press, Cambridge, 2010.