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]

[a]Università Bocconi and IGIER

[b]New York University and Hoover Institution

August 9, 2017

Abstract

This paper provides a general framework for the analysis of self-confirming policies. We first study self-confirming equilibria in recurrent decision problems with incomplete information about the true stochastic model. Next we illustrate the theory with a characterization of 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: a policy maker's expectations about the consequences of *alternative* policies might be incorrect. This paper provides a framework for the analysis of such self-confirming policies. We first develop a general analysis of self-confirming equilibria in recurrent decision problems with incomplete information about the true stochastic model. Next we illustrate the theory with a characterization of stationary monetary policies in a linear-quadratic setting.

Consider an agent who makes recurrent decisions under uncertainty. In each period he takes an action a that, via a feedback function f, delivers an observable outcome $m = f(a, s)$ depending upon an unobservable state of nature s. A fixed, unknown stochastic model σ^* (that is, a probability measure over the states) determines state realizations. The agent knows the feedback function f, but not the stochastic model σ^*. There are no structural links between periods, but the agent observes the realized outcome in each period t and therefore updates his 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

*We thank Roberto Corrao, Enrico De Magistris, Giacomo Lanzani, 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.

strategy yields a convergent active learning process, i.e., a stochastic process of actions and updated beliefs (a_t, μ_t) that converges almost surely.[1] The realized stochastic limit (a^*, μ^*) almost surely satisfies 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. The key observation is that the confirmed belief μ^* need not assign probability one 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.

In a self-confirming equilibrium, decision makers might well be best replying to empirically confirmed, but wrong views about the actual data generating model. They may thus get trapped in self-confirming behavior that differs substantially from the objectively optimal behavior postulated by rational expectations models.[3] 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 experiment thoroughly 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 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.[4]

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 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 standard quadratic loss function, we completely characterize the self-confirming equilibrium map that associates each conceivable model economy with a corresponding set of

[1] See, e.g., Easley and Kiefer (1988) and the references therein. Battigalli et al. (2016) describe the exact relationship between our framework and the stochastic control framework of Easley and Kiefer, showing that they are essentially equivalent.

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

[3] In 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. Hence, rational-expectations models feature a Nash-type notion of equilibrium, where equilibrium choices are best replies to correct beliefs.

[4] But the scope of our analysis goes well beyond that. For example, similar considerations of the impossibility of thorough experimentation and of its consequences in terms of welfare naturally also apply to environmental policies.

self-confirming beliefs and monetary policies. Given a fixed policy, the monetary authority infers from long-run data the 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 the moments of 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, 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 natural rate, and average inflation is (objectively) excessive. But we do not take a stand on what the true model is and we 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 welfare 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 hypothesis.

To escape the partial identification trap more experimentation may be advisable. But we do not see an easy way out: large-scale social experiments may have huge costs, while small-scale ones may have little external validity.

Be that as it may, this paper provides an analysis of these situations that is at the same time general enough to be portable to many policy decisions problems and sufficiently specific to yield useful characterizations and welfare-relevant implications.

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 setting in which self-confirming analysis was originally developed within game theory as well as the infinite setup relevant for macroeconomic policy analysis;
- *specific enough* to provide welfare implications for relevant policy questions with the backdrop of a neat learning foundation.

The main questions that we address at this abstract level concern the scope of partial identification (Section 3), equilibrium values, and the resulting effects on welfare (Sections 4.2 and 4.3).

The second part of this paper (Section 5) builds on the abstract analysis to gain a better perspective and novel results on the classical debate on the possibility of systematically exploiting unemployment/inflation trade-offs. 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.

The Appendix collects some more technical material and all the formal proofs.[5]

Related literature Our analysis 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) 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. The latter paper is focused on the interaction between ambiguity aversion and self-confirming equilibria in games. Here instead we consider a decision maker who maximizes his subjective expected utility (i.e., he is ambiguity neutral). This simplifies the general analysis analysis without affecting the illustrative application.

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 approach a neighborhood of this equilibrium, but then abandons it, as the Phillips curve looks "more vertical" and after some time inflation is lowered.[6] Cho et al. (2002) and Sargent and Williams (2005) sharpen the theoretical analysis of such learning dynamics.[7] Cho and Kasa (2015) notice 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-

[5] Some further topics are analyzed in the working paper version.

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

[7] The 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.

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 allows for the possibility of a direct impact of target inflation on unemployment. Unlike the cited papers, we do not take a stand on the true model economy, i.e., we characterize the self-confirming equilibrium set for every conceivable model, instead of necessarily assuming that the true model economy features a rational expectations augmented Phillips curve.

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 rationalizable self-confirming equilibria 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. In a series of papers, Saint Paul (e.g., 2012, 2013) considers an expert who knows the true model and advises the policy maker while pursuing his own policy agenda; the policy maker and the market agents fully trust the expert as long as the data are consistent with his advice. With this, the expert manipulates the policy maker and market agents under a self-confirmation constraint. Finally, Gaballo and Marimon (2015) 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.[9]

2 Preliminaries

2.1 Mathematics

Finite setup To fix ideas, the reader can assume that the spaces considered in the paper are finite. In particular, if X is a finite set with n elements, the collection $\Delta(X)$ of all probability measures on X can be identified with the simplex $\left\{\xi \in \mathbb{R}^n_+ : \sum_{i=1}^n \xi_i = 1\right\}$ of $\mathbb{R}^n$. Moreover, integrals reduce to sums, that is, $\int_X f(x)\, d\xi(x) = \sum_{x\in X} f(x)\,\xi(x)$ for all $\xi \in \Delta(X)$ and $f : X \to \mathbb{R}$.

General setup In general, however, we need to consider both finite and infinite spaces.[10] A pair $(X, \mathcal{X})$ is a (standard) Borel space if there exists a metric that makes the space X

[8] In 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. The most important difference with the self-confirming equilibrium literature is that, although Kurz analyzes multi-agent systems, he does not use a game theoretic framework. This makes it difficult to compare rational-belief equilibrium with self-confirming equilibrium.

[9] The model is not explicitly represented as a game. Therefore the connection to the traditional self-confirming equilibrium concept is not immediate. Furthermore, the self-confirming policy analyzed in this paper is the one of creditors (banks), not of the monetary authority.

[10] For instance, action spaces are often assumed to be infinite in order to solve optimization problems through differential methods. The Phillips curve exploitation example studied later in the paper features, indeed, an infinite action space.

complete and separable (that is, a Polish space), and $\mathcal{X}$ is its Borel sigma algebra. The elements B of $\mathcal{X}$ are called Borel sets; they are Borel spaces themselves, with the relative sigma algebra. When X is finite, $\mathcal{X}$ is its power set 2^X.[11]

We denote by $\Delta(X)$ the collection of all (Borel) probability measures on X. We endow $\Delta(X)$ with the natural sigma algebra,[12] which in turn makes it a Borel space. It is then natural to endow any Borel subset Σ of $\Delta(X)$ with the relative sigma algebra; we denote by $\Delta(\Sigma)$ the collection of all probability measures on Σ. Finally, we denote by $\delta_x \in \Delta(X)$ the Dirac measure concentrated on $x \in X$, that is, $\delta_x(B) = 1$ if $x \in B$ and $\delta_x(B) = 0$ if $x \notin B$.

Given any two Borel spaces X and Y, their product $X \times Y$ is a Borel space with respect to the product sigma algebra.

Each measurable function $\varphi : X \to Y$ induces a measurable distribution map $\hat{\varphi} : \Delta(X) \to \Delta(Y)$ 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 the Borel sigma algebra $\mathcal{Y}$ of Y.

The following routine lemma describes a key feature of $\hat{\varphi}$.

Lemma 1 *If X and Y are Borel spaces and $\varphi : X \to Y$ is measurable, then $\hat{\varphi}$ is one-to-one if and only if φ is one-to-one.*

Unless otherwise stated, throughout the paper spaces are assumed to be Borel spaces, as usual in stochastic optimization (see, e.g., Puterman, 2014).

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 $\rho : A \times S \to C$ a measurable *consequence function* that associates a consequence $\rho(a, s) \in C$ to each pair $(a, s) \in A \times S$ of action and state; in particular, C becomes a subset of the real line when consequences are monetary.

The quartet (A, S, C, ρ) is the basic structure of the decision problem.

The inherent randomness that characterizes states' realizations – 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, each action a is evaluated through its expected utility $\int_S v(\rho(a, s))\, d\sigma(s)$, where $v : C \to \mathbb{R}$ is a von Neumann-Morgenstern (measurable and bounded above) 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 $r : A \times S \to \mathbb{R}$ is the *payoff* (or *reward*) *function* $r = v \circ \rho$. The payoff function is measurable and bounded above since the utility function has these properties; so, all our integrals are well defined. For every action $a \in A$, the section $R(a, \cdot) : \Delta(S) \to [-\infty, \infty)$ is measurable and bounded above too.

[11] The power set 2^X, which is the collection of all subsets of X, is the Borel sigma algebra of X under the discrete metric.

[12] That is, the sigma algebra generated by the functions $\phi_B : \Delta(X) \to \mathbb{R}$ defined by $\phi_B(\xi) = \xi(B)$ for all $B \in \mathcal{X}$ (cf. Theorem 2.3 of Gaudard and Hadwin, 1989).

We assume, a la Neyman-Pearson-Wald, that decision makers do not know the true probability model, but that they know a (measurable) collection $\Sigma \subseteq \Delta(S)$ of probability models that contains the true one (we thus abstract from misspecification issues). We call *structural* the kind of information that allows decision makers to posit the collection Σ. When Σ is not a singleton, decision makers face *model uncertainty*. They rank actions according to the *classical subjective expected utility* (SEU) criterion:

$$V(a,\mu) = \int_{\Sigma} R(a,\sigma)\, d\mu(\sigma) \tag{1}$$

where $\mu \in \Delta(\Sigma)$ is a subjective *prior probability* over models in Σ that reflects personal beliefs about models that decision makers may have, in addition to the structural information behind Σ.[13]

Representation (1) admits the reduced form $\int_{\Sigma} R(a,\sigma)\, d\mu(\sigma) = \int_S r(a,s)\, d\sigma_\mu = R(a,\sigma_\mu)$, 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 elicited σ_μ from betting behavior.

The decision problem can be summarized by the sextet

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

that combines the basic structure (A, S, C, ρ) 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 (1) 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 result, knows that σ is the true model. In this case, there is only physical uncertainty, quantified by σ, without any model uncertainty. Criterion (1) again reduces to the expected payoff criterion $R(a,\sigma)$, now interpreted as a von Neumann-Morgenstern criterion. For instance, if the decision maker either observed infinitely many drawings from a given urn, or if he just were able to count the balls of each color, he would learn/know the urn composition, so Σ would be a singleton.

(iii) When $\Sigma \subseteq \{\delta_s : s \in S\}$, there is no physical uncertainty; there is 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 (1) takes the form $R(a,\mu)$.[14]

[13] See Marinacci (2015) for a discussion of this setup; classical SEU is proposed by Cerreia-Vioglio et al. (2013).

[14] See Corollary 4 in Appendix A.1.

3 Partial identification

3.1 Feedback

We assume that the decision maker faces problem D recurrently in a stationary environment with an i.i.d. process of states determined by an 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* $f : A \times S \to M$, where M is a space of messages. By selecting action $a \in A$ the decision maker receives a *message*

$$m = f_a(s)$$

when s occurs.[15] The decision maker's information about the state is thus endogenous: if M is finite, the information is represented by the partition $\left\{f_a^{-1}(m) : m \in M\right\}$ of the state space S that the messages induce, which depends on the choice of action a; if M is infinite, this partition is replaced by the sigma algebra

$$\mathcal{F}_a = \left\{f_a^{-1}(B) : B \in \mathcal{M}\right\} \subseteq \mathcal{S}$$

When information does not depend on 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 is one-to-one. 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 imperfect when each section f_a is constant).

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 thus a weak form of "endogenous" perfect feedback.

We assume throughout that consequences are observable. Formally, we require that for each action $a \in A$, there exists a measurable function $g_a : M \to C$ such that

$$\rho_a(s) = g_a(f_a(s)) \qquad \forall s \in S$$

In words, messages encode consequences. In particular, when the consequences of the actions are the only observed messages, we have

$$C = M \text{ and } f = \rho$$

This is the most common and important case, and it will be the case in our macroeconomic application.

Example 1 Consider a decision maker who is asked to 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, he is

[15] Here $f_a : S \to M$ denotes the section $f(a, \cdot)$ of f at a.

told whether he won (in which case he receives 1 euro) or not (in which case he receives 0 euros). We have $S = \{B, G, Y\}$, $A = \{b, g, y\}$, and $C = M = \{0, 1\}$. Moreover,

$$\rho(b, B) = \rho(y, Y) = \rho(g, G) = 1$$

and

$$\rho(b, Y) = \rho(b, G) = \rho(y, B) = \rho(y, G) = \rho(g, B) = \rho(g, Y) = 0$$

The feedback function coincides with the consequence function, that is, $f = \rho$. Thus, the decision maker observes the realized color if he wins, but not if he loses. In particular, if he chooses action b, then

$$f_b^{-1}(1) = \{B\}, f_b^{-1}(0) = \{Y, G\}$$

that is, betting on b yields the algebra

$$\mathcal{F}_b = \{\emptyset, S, \{B\}, \{Y, G\}\}$$

of S. Similarly,

$$\mathcal{F}_y = \{\emptyset, S, \{Y\}, \{B, G\}\}$$

and

$$\mathcal{F}_g = \{\emptyset, S, \{G\}, \{B, Y\}\}$$

Therefore, own-action independence of feedback about the state (color) does not hold. ▲

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. If M is finite, $\nu(m)$ is the empirical frequency of message m. Given an action $a \in A$, consider the distribution map $\hat{f}_a : \Sigma \to \Delta(M)$ given, for each $\sigma \in \Sigma$, 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 $B \in \mathcal{M}$.[16] If M is finite, $\hat{f}_a(\sigma)(m)$ is the empirical frequency with which the decision maker receives message m when he chooses action a and σ is the true model.[17] The inverse correspondence $\hat{f}_a^{-1}$ partitions Σ into classes

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

of models that are observationally equivalent given that action a is chosen and that the frequency distribution of messages ν is observed in the long-run. 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 (Lemma 1), and so $\hat{f}_a^{-1}(\nu)$ is either a singleton or empty 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

[16] In the literature $\hat{f}_a(\sigma)(B)$ is sometimes denoted by $\hat{f}_a(B \mid \sigma)$, interpreted as the frequency of B given σ. Also note that here the distribution map $\hat{f}_a$ is restricted from $\Delta(S)$ to Σ.

[17] In the working paper version we make rigorous the long-run interpretation which we rely upon.

models. Otherwise, $\hat{f}_a^{-1}(\nu)$ is nonsingleton for some ν, so we have *partial identification* under action a. In the extreme case when $\hat{f}_a$ is constant – that is, all models generate the same message distribution – the decision problem is completely unidentified under action a.

For each action $a \in A$, consider the correspondence

$$\hat{\Sigma}_a = \hat{f}_a^{-1} \circ \hat{f}_a : \Sigma \to 2^{\Sigma}$$

For any fixed $\sigma \in \Sigma$, its image

$$\hat{\Sigma}_a(\sigma) = \left\{\sigma' \in \Sigma : \hat{f}_a\left(\sigma'\right) = \hat{f}_a(\sigma)\right\} \tag{2}$$

is the collection of models that are observationally equivalent given the long-run frequency distribution $\nu = \hat{f}_a(\sigma)$ of messages that action a generates along with model σ. In other words, $\hat{\Sigma}_a(\sigma)$ is the partially identified set of models given action a.[18]

Remark 1 The partially identified set can be written as $\hat{\Sigma}_a(\sigma) = \{\sigma' \in \Sigma : \sigma'_{|\mathcal{F}_a} = \sigma_{|\mathcal{F}_a}\}$, that is, partial identification is determined by the information sigma algebra $\mathcal{F}_a$. Therefore, own-action independence of feedback also implies that the partial identification correspondence is action independent: $\hat{\Sigma}_a(\sigma) = \hat{\Sigma}_{a'}(\sigma)$ for all $(a, a', \sigma) \in A \times A \times \Sigma$.

We can regard $\hat{\Sigma}_a$ as the *partial identification correspondence* determined by action 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 $\{\hat{\Sigma}_a(\sigma)\}_{\sigma\in\Sigma}$ of images is a measurable partition of Σ and each element of this partition consists of probability models that are observationally equivalent under action a.

Example 2 Consider the decision problem with feedback of Example 1. If the decision maker bets on Blue, his action prevents him from obtaining any evidence on the frequency of G and Y. In particular,

$$\hat{f}_b(\sigma)(1) = \sigma(B) \text{ and } \hat{f}_b(\sigma)(0) = 1 - \sigma(B) \qquad \forall \sigma \in \Sigma$$

where $\Sigma \subseteq \Delta(\{B, G, Y\})$ is the (finite) set of possible urn compositions that he posits. Hence, the evidence gathered through bet b only partially identifies the true model:

$$\hat{\Sigma}_b(\sigma) = \left\{\sigma' \in \Sigma : \hat{f}_b\left(\sigma'\right) = \hat{f}_b(\sigma)\right\} = \left\{\sigma' \in \Sigma : \sigma'(B) = \sigma(B)\right\}$$

For instance, if the true model σ is uniform, then

$$\hat{\Sigma}_b(\sigma) = \left\{\sigma' \in \Sigma : \sigma'(B) = \frac{1}{3}\right\}$$

[18] We show in the next section that coarser feedback implies a coarser partition of Σ, that is, a lower degree of identification.

that is, all probability models σ' that assign probability $1/3$ to B are observationally equivalent. More generally, if we denote by σ_n any model that assigns probability $n/90$ to B, then the partition $\{\hat{\Sigma}_b(\sigma)\}_{\sigma\in\Sigma} = \{\hat{\Sigma}_b(\sigma_n)\}_{n=0,\ldots,90}$ has 91 elements, each consisting of the probability models that assign probability $n/90$ to B. All models in the same cell $\hat{\Sigma}_b(\sigma_n)$ are observationally equivalent. ▲

Example 3 Suppose now that the decision maker observes *ex post* the color of the ball:

$$f(b,s) = f(g,s) = f(y,s) = s \qquad \forall s \in \{B,G,Y\}$$

Then there is perfect feedback and $\hat{\Sigma}_b(\sigma) = \hat{\Sigma}_g(\sigma) = \hat{\Sigma}_y(\sigma) = \{\sigma\}$ for each $\sigma \in \Sigma$. Regardless of the chosen action, the true model is identified. ▲

3.3 Comparative statics

We now show that the extent of model identification naturally depends on the underlying feedback function. To this end, given any two feedback functions f and f', say that f' is *coarser* than f if, for each $a \in A$, there exists a measurable function $h_a : M \to M'$ such that

$$f'_a(s) = h_a(f_a(s)) \qquad \forall s \in S$$

Our assumption that consequences are observable implies that ρ is the coarsest possible feedback, while perfect feedback is the least coarse (finest).

Lemma 2 *If f' is coarser than f, then, for each $a \in A$, $\hat{\Sigma}_a(\sigma) \subseteq \hat{\Sigma}'_a(\sigma)$ for all $\sigma \in \Sigma$.*

Coarser feedback functions thus determine, for each action, coarser partial identification correspondences: worse information translates into a lower degree of identification.

4 Self-confirming actions and beliefs

4.1 Definition

A decision problem with feedback can be described by the pair (D, f). The partial identification issues discussed in the previous section motivate the following definition.

Definition 1 *A pair $(a^*, \mu^*) \in A \times \Delta(\Sigma)$ of actions and beliefs is a* self-confirming equilibrium given $\sigma^* \in \Sigma$ *if*

$$V(a^*, \mu^*) \geq V(a, \mu^*) \qquad \forall a \in A \tag{3}$$

and

$$\mu^* \in \Delta(\hat{\Sigma}_{a^*}(\sigma^*)) \tag{4}$$

The definition relies on two pillars: the optimality condition (3) that ensures that action a^* is subjectively optimal under belief μ^*, and the belief confirmation condition (4) that guarantees that belief μ^* is consistent with the data that action a^* reveals.[19] Thus, the pair

[19] Here, 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 $\hat{\Sigma}_{a^*}(\sigma^*)$.

(a^*, μ^*) of actions and beliefs determines the message distribution $\nu^* = \hat{f}_{a^*}(\sigma^*)$ that is the evidence that disciplines the subjective belief μ^*.[20]

Example 4 In the urn setting of Example 1, suppose that $\Sigma = \{\sigma^*, \sigma_1, \sigma_2\} \subseteq \Delta(\{B, G, Y\})$, where the three possible models are described in the table below:

	B	G	Y
σ^*	$\frac{1}{3}$	0	$\frac{2}{3}$
σ_1	$\frac{1}{3}$	$\frac{2}{3}$	0
σ_2	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$

Impose the normalization $u(0) = 0$ and $u(1) = 1$, so that

$$R(b, \sigma^*) = R(b, \sigma_2) = R(b, \sigma_1) = R(y, \sigma_2) = R(g, \sigma_2) = \frac{1}{3} \tag{5}$$
$$R(y, \sigma^*) = R(g, \sigma_1) = \frac{2}{3} \quad ; \quad R(y, \sigma_1) = R(g, \sigma^*) = 0$$

(i) Consider a uniform belief μ^* on Σ: $\mu^*(\sigma^*) = \mu^*(\sigma_1) = \mu^*(\sigma_2) = 1/3$. The pair (b, μ^*) is self-confirming. Since $V(b, \mu^*) = V(g, \mu^*) = V(y, \mu^*) = 1/3$, the optimality condition (3) is satisfied. It is easy to check (cf. Example 1) that $\hat{\Sigma}_b(\sigma^*) = \{\sigma \in \Sigma : \sigma(B) = 1/3\} = \Sigma$. Hence $\mu^* \in \Delta(\hat{\Sigma}_b(\sigma^*))$, and so the data confirmation condition (4) is also satisfied. The self-confirming equilibrium (b, μ^*) generates the message distribution $\nu^* \in \Delta(\{0, 1\})$ with $\nu^*(1) = 1/3$, that is, a one-third frequency of wins.

(ii) Consider the belief $\mu^* = \delta_{\sigma^*}$ concentrated on the true model. The action and belief pair (y, μ^*) is self-confirming. The optimality condition (3) is easily seen to be satisfied, while the data confirmation condition (4) holds since $\hat{\Sigma}_y(\sigma^*) = \{\sigma \in \Sigma : \sigma(Y) = 2/3\} = \{\sigma^*\}$. The self-confirming equilibrium (y, μ^*) generates the message distribution $\nu^* \in \Delta(\{0, 1\})$ with $\nu^*(1) = 2/3$, that is, a two-thirds frequency of wins.

(iii) Since $\hat{\Sigma}_g(\sigma^*) = \{\sigma \in \Sigma : \sigma(G) = 0\} = \{\sigma^*\}$, action g is not part of any self-confirming equilibrium.

Actions b and y can be thus part of self-confirming equilibria. Since $\hat{\Sigma}_b(\sigma^*) \neq \hat{\Sigma}_y(\sigma^*)$, they differ in their identification properties. In particular, $\hat{\Sigma}_b(\sigma^*) \cap \hat{\Sigma}_y(\sigma^*) = \hat{\Sigma}_y(\sigma^*) = \{\sigma^*\}$. They also have different values: $R(b, \sigma^*) = 1/3 \neq 2/3 = R(y, \sigma^*)$. ▲

The next simple variation on the previous example shows the importance of structural information.

Example 5 If, in the previous example, we suppose that only actions b and g are available, i.e., $A = \{b, g\}$, then action g is still not part of any self-confirming equilibrium. But if we further suppose that the all-yellow urn belongs to the posited set of models, i.e., $\delta_Y \in \Sigma$ (as, for example, when there is no information about the urn), then this is no

[20] As mentioned in the Introduction, our steady-state interpretation of self-confirming equilibrium is justified by convergence results a la Easley and Kiefer (1988). See the working paper version of this paper for details. Battigalli el al. (2016) analyze the case of an ambiguity averse decision maker.

longer the case: the pair (g, δ_{δ_Y}) is self-confirming. In fact, $V(b, \delta_{\delta_Y}) = V(g, \delta_{\delta_Y}) = 0$ and $\delta_Y \in \hat{\Sigma}_g(\sigma^*) = \{\sigma \in \Delta(\{B, G, Y\}) : \sigma(G) = 0\}$. In words, the decision maker believes that he cannot win with either b or g, and he happens to choose the bet that truly gives him no chance. ▲

Under own-action independence of feedback (that is, actions do not affect information gathering), the data confirmation condition (4) becomes $\mu^* \in \Delta(\hat{\Sigma}(\sigma^*))$, where $\hat{\Sigma}(\sigma^*)$ is exogenously posited. We thus return to a traditional optimization notion with a purely exogenous data confirmation condition. In particular, under perfect feedback – and so full identification – the optimality condition (3) becomes

$$R(a^*, \sigma^*) \geq R(a, \sigma^*) \qquad \forall a \in A \tag{6}$$

since condition (4) 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 physical uncertainty (that is, risk).[21]

Action a^* is *objectively optimal* if it satisfies condition (6). Objectively optimal actions are the ones that the decision maker would select if he 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 section will show.

Remark 2 In Example 4, bet y is the objectively optimal action.

Of course, each "rational expectations" pair (a^*, δ_{σ^*}), where action a^* is objectively optimal and belief δ_{σ^*} is concentrated on the true model, is a self-confirming equilibrium.

The optimality condition (3) 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 (4) implies that the predictive probability σ_{μ^*} belongs to $\hat{\Sigma}_{a^*}(\sigma^*)$ if it belongs to Σ.[22] In this case, $(a^*, \delta_{\sigma_{\mu^*}})$ is a self-confirming equilibrium too. Hence we have the following *certainty equivalence principle*:

Proposition 1 *Given a true model $\sigma^* \in \Sigma$, if (a^*, μ^*) is a self-confirming equilibrium and $\sigma_{\mu^*} \in \Sigma$, then $(a^*, \delta_{\sigma_{\mu^*}})$ is a self-confirming equilibrium as well, with $V(a^*, \mu^*) = V(a^*, \delta_{\sigma_{\mu^*}})$.*

4.2 Value

Given a true model σ^*, the values of all self-confirming equilibria (a^*, μ^*) sharing the same equilibrium action a^* coincide with the expected payoff of a^* with respect to σ^*.[23] Formally:

Proposition 2 *If $(a^*, \mu^*) \in A \times \Delta(\Sigma)$ is a self-confirming equilibrium given $\sigma^* \in \Sigma$, then $V(a^*, \mu^*) = R(a^*, \sigma^*)$.*

[21] For condition (3) to reduce to (6) it is actually enough that the equilibrium action a^* be fully revealing, a weaker property than perfect feedback (see Corollary 5 below).

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

[23] Irrespective of the supporting belief.

The result is based on the following lemma of independent interest, in turn based on Battigalli et al. (2015), which shows that observationally equivalent models share the same expected utility.

Lemma 3 *Let $a \in A$ and $\sigma \in \Sigma$. We have $R(a, \sigma') = R(a, \sigma)$ for every $\sigma' \in \hat{\Sigma}_a(\sigma)$.*

In turn, the last two results easily imply the following characterization of self-confirming equilibria.

Corollary 1 *A pair $(a^*, \mu^*) \in A \times \Delta(\Sigma)$ is a self-confirming equilibrium given $\sigma^* \in \Sigma$ if and only if $R(a^*, \sigma^*) \geq V(a, \mu^*)$ for every $a \in A$ and the belief confirmation condition (4) holds.*

Given the data confirmation condition, the optimality condition (3) 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. In Appendix A.4, we show that self-confirming equilibria with sharper basic subjective assessments (smaller supports) have higher values.

4.3 Welfare

We say that an action a^* is a *self-confirming (equilibrium) action* given $\sigma^* \in \Sigma$, if there exists a belief $\mu^* \in \Delta(\Sigma)$ such that (a^*, μ^*) is a self-confirming equilibrium given σ^*. Since in this case $V(a^*, \mu^*) = R(a^*, \sigma^*) \leq \max_{a \in A} R(a, \sigma^*)$, the decision maker incurs a welfare loss

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

when he selects the self-confirming action a^*. In particular, $\ell(a^*, \sigma^*) = 0$ if and only if a^* is objectively optimal.

The loss is indeed caused by the decision maker's ignorance.

Proposition 3 *If a^* and b^* in A are self-confirming actions given $\sigma^* \in \Sigma$ and $\hat{\Sigma}_{a^*}(\sigma^*) \subseteq \hat{\Sigma}_{b^*}(\sigma^*)$, then $\ell(a^*, \sigma^*) \leq \ell(b^*, \sigma^*)$.*

That is, self-confirming actions with better identification properties exhibit lower losses. In this regard, the next result shows that for an action a with the best identification properties – thus, optimal from a purely statistical viewpoint – to be self-confirming amounts to being objectively optimal. Truth is ancillary to the decision maker's pursuit of his goals (and so of his happiness).

Proposition 4 *Given a true model $\sigma^* \in \Sigma$, suppose there is an action a such that $\hat{\Sigma}_a(\sigma^*) \subseteq \hat{\Sigma}_{a'}(\sigma^*)$ for each $a' \in A$. Then action a is self-confirming if and only if it is objectively optimal.*

In sum, the decision maker is not purely a statistician: he is not interested *per se* in discovering the true model unless the action that allows the discovery is subjectively optimal. In this sense, there is no separation between statistical estimation and decision in the present setup.

Example 6 Consider the decision problem

$A\backslash S$	s_1	s_2
a_1	1	1
a_2	0	2

(7)

where $r = \rho = f$ and $\Sigma = \Delta(S)$. Here the decision maker is risk neutral, has no structural information, and (monetary) consequences are the messages that he receives. Given any $\sigma \in \Sigma$, for the constant-payoff action a_1 we have $\hat{\Sigma}_{a_1}(\sigma) = \Sigma$, while for the risky action a_2 we have $\hat{\Sigma}_{a_2}(\sigma) = \{\sigma\}$. Action a_1 has no information value, and so no statistical interest, but it is not affected by state uncertainty. The opposite is true for action a_2, which is fully revealing but subject to uncertainty.

It holds that $R(a_1, \sigma) = 1$ and $R(a_2, \sigma) = 2\sigma(s_2)$ for all $\sigma \in \Sigma$. Hence, a_2 is a self-confirming equilibrium action when σ is the true model, i.e., $a_2 \in \gamma(\sigma)$, if and only if $\sigma(s_2) \geq 1/2$, that is, if and only if a_2 is objectively optimal (in accordance with the previous result). On the other hand, $a_1 \in \gamma(\sigma)$ for all $\sigma \in \Sigma$. In fact, a_1 is a best reply to every belief μ such that $\mu(\sigma(s_2) \leq 1/2)$, e.g., the belief $\delta_{s_1} \in \Sigma$ is concentrated on state s_1.

In sum, the constant-payoff action is always self-confirming, independently of the true model, while the risky one is self-confirming only when the true model makes it objectively optimal. ▲

5 Phillips curve exploitation example

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 a model asserting an exploitable trade-off between unemployment and inflation while the truth is that the trade-off is not exploitable.[24]

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 and by a monetary policy variable a. Specifically, unemployment and inflation outcomes (u, π) are connected to the state of the economy (w, ε, θ) and the government action a according to

$$u = \theta_0 + \theta_{1\boldsymbol{\pi}}\pi + \theta_{1\mathbf{a}}a + \theta_2 w \tag{8}$$

$$\pi = a + \theta_3\varepsilon \tag{9}$$

The vector parameter $\theta = (\theta_0, \theta_{1\boldsymbol{\pi}}, \theta_{1\mathbf{a}}, \theta_2, \theta_3) \in \mathbb{R}^5$, that is, the last component of the state vector, specifies the structural coefficients of an aggregate supply equation (8) and an inflation determination equation (9). Coefficients $\theta_{1\boldsymbol{\pi}}$ and $\theta_{1\mathbf{a}}$ are slope responses of unemployment

[24] The working paper version contains a more general analysis of self-confirming economic policies.

to actual and planned inflation,[25] 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 inflation policy $a = 0$.

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

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

In words, we posit a strictly positive intercept, 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) of the state vector). Finally, we assume that inflation and unemployment are inversely related.

The reduced form of each model economy is

$$u = \theta_0 + (\theta_{1\boldsymbol{\pi}} + \theta_{1\mathbf{a}})\, a + \theta_{1\boldsymbol{\pi}}\theta_3\varepsilon + \theta_2 w \tag{10}$$

$$\pi = a + \theta_3\varepsilon \tag{11}$$

The coefficients of the reduced form are $\xi = (\theta_0, \theta_{1\boldsymbol{\pi}} + \theta_{1\mathbf{a}}, \theta_{1\boldsymbol{\pi}}\theta_3, \theta_2, \theta_3) \in \mathbb{R}^5$. 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 $\xi \neq \xi'$.

Since momentarily we will assume that only realized unemployment and inflation are observable by the monetary authority, the reduced form above, will give us the feedback function $(u, \pi) = f(a, s)$ of the previous sections. Specifically, rewriting (10) and (11) as

$$\boldsymbol{u}(a, w, \varepsilon, \theta) = \theta_0 + (\theta_{1\boldsymbol{\pi}} + \theta_{1\mathbf{a}})\, a + \theta_{1\boldsymbol{\pi}}\theta_3\varepsilon + \theta_2 w$$

$$\boldsymbol{\pi}(a, w, \varepsilon, \theta) = a + \theta_3\varepsilon$$

will make the dependence of observables (u, π) from chosen actions a and (unobservable) realized states (w, ε, θ) explicit and it will allow us to study the present policy problem within our framework. Formally, the message space $M = \mathbb{R}^2$ now consists of unenployment/inflation pairs, and the feedback function is $f = (\boldsymbol{u}, \boldsymbol{\pi}) : A \times (\mathbb{R}^2 \times \Theta) \to \mathbb{R}^2$.

The policy multiplier $\xi_2 = \theta_{1\boldsymbol{\pi}} + \theta_{1\mathbf{a}} = \theta_{1\mathbf{a}} - |\theta_{1\boldsymbol{\pi}}|$ 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 $\theta_{1\mathbf{a}}$ and $\theta_{1\boldsymbol{\pi}}$. 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.

To sum up, the set of parameters is

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

Our analysis will pay special attention to the following two competing model economies.

[25] The economic interpretation is that planned inflation a affects agents' expectations to an extent parametrized by $\theta_{1\mathbf{a}}$.

5.1.1 The Lucas-Sargent model

The first model, 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 $\beta \equiv \theta_{1\boldsymbol{\pi}} = -\theta_{1\mathbf{a}}$, and so $\theta = (\theta_0, \beta, -\beta, \theta_2, \theta_3)$ and $\xi = (\theta_0, 0, \beta\theta_3, \theta_2, \theta_3)$. In this new-classical model 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 The Samuelson-Solow model

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

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

where $\theta_{1\mathbf{a}} = 0$ and so $\theta = (\theta_0, \theta_{1\boldsymbol{\pi}}, 0, \theta_2, \theta_3)$ and $\xi = (\theta_0, \theta_{1\boldsymbol{\pi}}, \theta_{1\boldsymbol{\pi}}\theta_3, \theta_2, \theta_3)$. In this Keynesian model, the policy multiplier $\xi_2 = \theta_{1\boldsymbol{\pi}}$ 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 models. The consequence space C consists of unemployment and inflation pairs (u, π), so we set $C = U \times \Pi$. The consequence function $\rho : A \times (W \times E \times \Theta) \to C$ is

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

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

$$\rho(a, w, \varepsilon, \theta) = \begin{bmatrix} \theta_0 \\ 0 \end{bmatrix} + a \begin{bmatrix} \theta_{1\boldsymbol{\pi}} + \theta_{1\mathbf{a}} \\ 1 \end{bmatrix} + \begin{bmatrix} \theta_2 & \theta_{1\boldsymbol{\pi}}\theta_3 \\ 0 & \theta_3 \end{bmatrix} \begin{bmatrix} w \\ \varepsilon \end{bmatrix} \tag{12}$$

5.2.2 Factorization

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

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

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) is[26]

$$\hat{f}_a(\sigma)(u,\pi) = \sigma\left(\{(w,\varepsilon,\theta) : (\boldsymbol{u}(a,w,\varepsilon,\theta), \boldsymbol{\pi}(a,w,\varepsilon,\theta)) = (u,\pi)\}\right) \tag{13}$$

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.

At this point, it is convenient to enrich 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 parametrizes a model economy. This suggests factorizing the probability models $\sigma \in \Sigma \subseteq \Delta(W \times E \times \Theta)$ as

$$\sigma = q \times \delta_\theta \tag{14}$$

where the marginal distribution of shocks $q \in \Delta(W \times E)$ is 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 parametrize 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 (14) has two modelling 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.[27]

A first dividend of the factorization is that the objective function (1) takes the simpler form

$$V(a,\mu) = \int_\Theta \left(\int_{W\times E} r(a,w,\varepsilon,\theta)\, dq(w,\varepsilon) \right) d\mu(\theta) \tag{15}$$

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.[28]

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

In words, shocks are uncorrelated and normalized.

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

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

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

[28] Whenever 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:

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

With this, a sharp identification result holds.

Proposition 5 *The partial identification correspondence* $\hat{\Sigma}_a : \Theta \to 2^\Theta$ *is*

$$\hat{\Sigma}_a(\theta) = \left\{\theta' \in \Theta : \theta'_0 + \theta'_{1\mathbf{a}} a = \theta_0 + \theta_{1\mathbf{a}} a,\ \theta'_{1\boldsymbol{\pi}} = \theta_{1\boldsymbol{\pi}},\ \theta'_2 = \theta_2,\ \theta'_3 = \theta_3\right\} \tag{16}$$

Given the true model θ, the shock coefficients θ_2 and θ_3 are thus identified, along with the slope $\theta_{1\boldsymbol{\pi}}$ 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 $(\boldsymbol{u}_a, \boldsymbol{\pi}_a) : W \times E \times \Theta \to U \times \Pi$ that, for a given policy a, represents the unemployment and inflation rates determined by the state (w, ε, θ).[29] The monetary authority observes in the long-run the following moments:

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

Therefore,

$$\theta_{1\boldsymbol{\pi}} = \frac{\mathrm{Cov}_\theta(\boldsymbol{u}_a, \boldsymbol{\pi}_a)}{\mathrm{Var}_\theta(\boldsymbol{\pi}_a)} \tag{17}$$

is the beta coefficient of the Phillips regression of unemployment on inflation,[30]

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

is the residual variance of $\boldsymbol{u}_a$ (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

$$\theta_0 + \theta_{1\mathbf{a}} a = \mathbb{E}_\theta(\boldsymbol{u}_a) - \frac{\mathrm{Cov}_\theta(\boldsymbol{u}_a, \boldsymbol{\pi}_a)}{\mathrm{Var}_\theta(\boldsymbol{\pi}_a)} \mathbb{E}_\theta(\boldsymbol{\pi}_a) \tag{18}$$

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 chose. What the authority learns depends on what it does in ways that it does not appreciate.

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

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

5.2.4 Estimated model economy

The moments that identify the three coefficients $\theta_{1\boldsymbol{\pi}}$, 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 $\theta_{1\boldsymbol{\pi}}$,[31] by $\hat{\sigma}_{\boldsymbol{u}|\boldsymbol{\pi}}$ the residual standard deviation that identifies θ_2, and by $\hat{\sigma}_{\boldsymbol{\pi}}$ the standard deviation of inflation that identifies θ_3. In contrast, the alpha regression coefficient that identifies the sum $\theta_0 + \theta_{1\mathbf{a}}a$ depends on policy a; we denote it by $\hat{\alpha}(a)$.

With this, we can write

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

As a result, the long-run estimated version of the model economy (8)-(9) that the monetary authority considers is

$$u = \hat{\alpha}(a) + \hat{\beta}\pi + \hat{\sigma}_{\boldsymbol{u}|\boldsymbol{\pi}}w \tag{19}$$

$$\pi = a + \hat{\sigma}_{\boldsymbol{\pi}}\varepsilon \tag{20}$$

$$\hat{\alpha}(a) = \theta_0 + \theta_{1\mathbf{a}}a \tag{21}$$

In particular, (19) is the estimated aggregate supply equation and (20) is the estimated inflation equation. The intercept of the former equation depends on the policy a via the equality (21), which only partly identifies the two coefficients θ_0 and $\theta_{1\mathbf{a}}$. In turn, this makes the policy multiplier $\xi_2 = \hat{\beta} + \theta_{1\mathbf{a}}$ 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 $\theta_{1\mathbf{a}}$. The former is the average unemployment at zero planned inflation, $\theta_0 = \mathbb{E}_\theta(\boldsymbol{u}_0)$; the latter is the "direct" impact of policy on unemployment.

The parameter space of the estimated model economy (19)-(21) reduces to $\Theta = \tilde{\Theta} \times \{(\hat{\beta},\, \hat{\sigma}_{\boldsymbol{u}|\boldsymbol{\pi}},\, \hat{\sigma}_{\boldsymbol{u}})\}$, where $\tilde{\Theta} = \mathbb{R}_{++} \times (-\infty, -\hat{\beta}]$ is the collection of all possible values $(\theta_0, \theta_{1\mathbf{a}})$ of the two remaining unidentified coefficients and $\{(\hat{\beta},\, \hat{\sigma}_{\boldsymbol{u}|\boldsymbol{\pi}},\, \hat{\sigma}_{\boldsymbol{u}})\}$ is the singleton containing the identified vector $(\theta_{1\boldsymbol{\pi}}, \theta_2, \theta_3)$. 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 (16), the partial identification correspondence $\hat{\Sigma}_a : \tilde{\Theta} \to 2^{\tilde{\Theta}}$ becomes

$$\hat{\Sigma}_a(\theta) = \left\{\left(\theta'_0, \theta'_{1\mathbf{a}}\right) \in \tilde{\Theta} : \theta'_0 = -\theta'_{1\mathbf{a}}a + \theta_0 + \theta_{1\mathbf{a}}a\right\} \tag{22}$$

In words, $\hat{\Sigma}_a(\theta)$ is a straight line in the plane, with slope $-a$ and intercept $\theta_0 + \theta_{1\mathbf{a}}a$ (determined by the policy a and by the true economic model θ). We thus have a partial identification line that defines a linear relationship between the two unidentified coefficients, given a true model. In other words, partial identification is unidimensional.

[31] By 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.

Given a true model $\theta = (\theta_0, \theta_{1\mathbf{a}})$, 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 $(\theta_0, \theta_{1\mathbf{a}})$ and have slope $-1/a$. In each such line there is a unique Lucas-Sargent model, characterized by $\theta'_{1\mathbf{a}} = -\hat{\beta}$, as well as a unique Samuelson-Solow model, characterized by $\theta'_{1\mathbf{a}} = 0$. In other words, partial identification lines feature a unique specimen of each class of models.

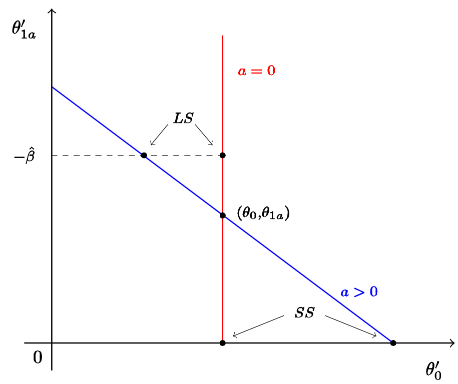

The figure illustrates the previous analysis. In particular, LS stands for Lucas-Sargent model and SS for Samuelson-Solow model, while the red (resp., blue) line is the partial identification line that correspond to policy $a = 0$ (resp., $a > 0$).

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 $v : C \to \mathbb{R}$ given by $v(u, \pi) = -u^2 - \pi^2$, so that the reward function $r : A \times S \to \mathbb{R}$ becomes:

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

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 $R(a, \theta) = v(\mathbb{E}_\theta(\boldsymbol{u}_a), \mathbb{E}_\theta(\boldsymbol{\pi}_a)) + const.$*

The linear quadratic framework thus allows us to express the expected reward as the utility of expectations. As a result, the objective function (15) becomes

$$V(a, \mu) = \int_{\tilde{\Theta}} v(\mathbb{E}_\theta(\boldsymbol{u}_a)), \mathbb{E}_\theta(\boldsymbol{\pi}_a))\, d\mu(\theta) + const. \tag{23}$$

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 (19)-(21). In both equilibrium conditions, the true multiplier $\xi_2^* = \hat{\beta}^* + \theta_{\mathbf{1a}}^*$ and its conjectured value $\mathbb{E}_{\mu^*}(\xi_2) = \hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}})$ play a key role.[32]

Proposition 6 *A policy and belief pair $(a^*, \mu^*) \in A \times \Delta(\tilde{\Theta})$ is a self-confirming equilibrium given θ^* if and only if*

$$a^* = -\frac{\theta_0^* \left(\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}}) \right)}{1 + \left(\hat{\beta}^* + \theta_{\mathbf{1a}}^* \right) \left(\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}}) \right)} \tag{24}$$

and

$$\mu^* \left(\left\{ (\theta_0, \theta_{\mathbf{1a}}) \in \tilde{\Theta} : \theta_0 = \theta_0^* - \frac{\theta_0^* \left(\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}}) \right)}{1 + \left(\hat{\beta}^* + \theta_{\mathbf{1a}}^* \right) \left(\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}}) \right)} (\theta_{\mathbf{1a}}^* - \theta_{\mathbf{1a}}) \right\} \right) = 1 \tag{25}$$

The result can be heuristically derived in the special case of *dogmatic beliefs*, when μ^* is concentrated on a single parameter vector $\bar{\theta} = \left(\bar{\theta}_0, \bar{\theta}_{\mathbf{1a}} \right) \in \tilde{\Theta}$, that is, $\mu^* = \delta_{\bar{\theta}}$. By (23), up to a constant the monetary authority's value function is

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

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

Based on the estimated model economy (19)-(21), a dogmatic authority conjectures that, according to the chosen policy a, the expected values of inflation and unemployment are constrained by the equation

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

This conjectured constraint is the version of the estimated aggregate supply equation (19) that the authority expects to face systematically given its dogmatic belief. So the authority's decision problem is

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

With this, the Lagrangian is

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

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

and the first-order conditions are

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

By solving them we get

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

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

$$a^*=B\left(\bar{\theta}\right) \qquad \text{(subjective best reply)} \tag{26}$$

and

$$\bar{\theta}_0=\theta_0^*+\left(\theta_{\mathbf{1a}}^*-\bar{\theta}_{\mathbf{1a}}\right)a^* \qquad \text{(confirmed beliefs)} \tag{27}$$

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

$$a^*=-\frac{\theta_0^*\left(\hat{\beta}^*+\bar{\theta}_{\mathbf{1a}}\right)}{1+\left(\hat{\beta}^*+\theta_{\mathbf{1a}}^*\right)\left(\hat{\beta}^*+\bar{\theta}_{\mathbf{1a}}\right)} \tag{28}$$

and

$$\bar{\theta}_0=\theta_0^*-\frac{\theta_0^*\left(\hat{\beta}^*+\bar{\theta}_{\mathbf{1a}}\right)}{1+\left(\hat{\beta}^*+\theta_{\mathbf{1a}}^*\right)\left(\hat{\beta}^*+\bar{\theta}_{\mathbf{1a}}\right)}\left(\theta_{\mathbf{1a}}^*-\bar{\theta}_{\mathbf{1a}}\right) \tag{29}$$

which are the equilibrium relations (24) and (25) in the case of dogmatic beliefs.[33]

The following figure illustrates the previous heuristic argument when the Lucas-Sargent model is true, so that θ_0^* is the natural rate of unemployment and $\theta_{\mathbf{1a}}^*=-\hat{\beta}^*$ (and so the true policy multiplier ξ_2^* is zero). Under this true model, policy a induces average unemployment $\mathbb{E}_{\theta^*}\left(\boldsymbol{u}_a\right)=\theta_0^*$ and average inflation $\mathbb{E}_{\theta^*}\left(\boldsymbol{\pi}_a\right)=a$. But a monetary authority with dogmatic belief $\delta_{\bar{\theta}}$ expects to observe the pair of long-run averages $\left(\mathbb{E}_{\bar{\theta}}\left(\boldsymbol{u}_a\right),a\right)$. This dogmatic belief is confirmed, and so condition (27) is satisfied, if $\mathbb{E}_{\bar{\theta}}\left(\boldsymbol{u}_a\right)=\theta_0^*$, 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 (26) 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,

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

where $\bar{\xi}_2 = \hat{\beta}^* + \bar{\theta}_{\mathbf{1a}}$ is the conjectured multiplier.

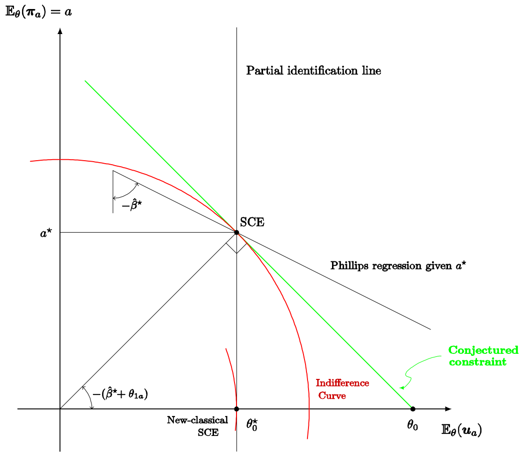

When the dogmatic belief is such that $\bar{\theta}_{\mathbf{1a}} = 0$ so that $\bar{\xi}_2 = \hat{\beta}^*$ becomes the conjectured multiplier, the monetary authority is "orthodox" Keynesian and the figure becomes:

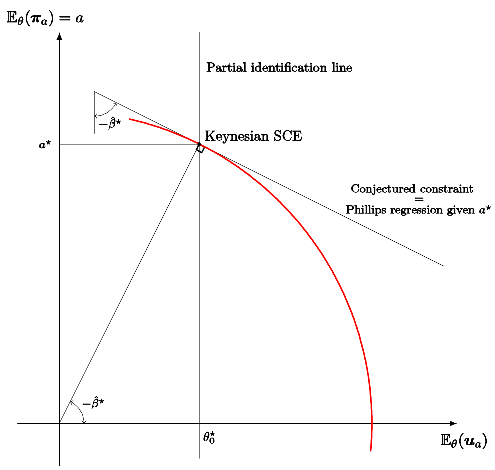

The conjectured constraint is $\mathbb{E}_{\bar{\theta}}(\boldsymbol{u}_a) = \bar{\theta}_0 + \hat{\beta}^* \mathbb{E}_{\bar{\theta}}(\boldsymbol{\pi}_a)$. Its slope is the beta coefficient of the Phillips regression, which represents the trade-off between inflation and unemployment that the Keynesian 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 $\xi_2 = \hat{\beta} + \theta_{\mathbf{1a}}$. The authority underestimates the multiplier when $\mathbb{E}_{\mu^*}(\xi_2) > \xi_2^*$ and overestimates it when $\mathbb{E}_{\mu^*}(\xi_2) < \xi_2^*$.[34] In structural terms, $\mathbb{E}_{\mu^*}(\xi_2) \gtrless \xi_2^*$ if and only if $\mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}}) \gtrless \theta_{\mathbf{1a}}^*$. For instance, when $\theta_{\mathbf{1a}}^*$ and $\mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}})$ 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.

[34] Both ξ_2^* and $\mathbb{E}_{\mu^*}(\xi_2)$ are negative (Assumption 2), and so $\mathbb{E}_{\mu^*}(\xi_2) \gtrless \xi_2^*$ if and only $|\mathbb{E}_{\mu^*}(\xi_2)| \lessgtr |\xi_2^*|$.

The objectively optimal policy is

$$a^o = -\frac{\theta_0^* \left(\hat{\beta}^* + \theta_{\mathbf{1a}}^*\right)}{1 + \left(\hat{\beta}^* + \theta_{\mathbf{1a}}^*\right)^2} \tag{30}$$

It is immediate to see that $a^* = a^o$ if and only if $\mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}}) = \theta_{\mathbf{1a}}^*$, (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.[35]

Proposition 7 *Given a true model θ^*, for every self-confirming equilibrium (a^*, μ^*),*

(i) $\mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}}) < \theta_{\mathbf{1a}}^*$ *if and only if policy a^* is hyperactive, i.e., $a^* > a^o > 0$;*

(ii) $\mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}}) = \theta_{\mathbf{1a}}^*$ *if and only if policy a^* is objectively optimal, i.e., $a^* = a^o$;*

(iii) $\theta_{\mathbf{1a}}^* < \mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}}) < -\hat{\beta}^*$ *if and only if policy a^* is hypoactive, i.e., $0 < a^* < a^o$;*

(iv) $\mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}}) = -\hat{\beta}^*$ *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 8 *The welfare loss is $\ell(a^*, \theta^*) = (1 + (\hat{\beta}^* + \theta_{\mathbf{1a}}^*)^2)(a^* - a^o)^2$.*

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 $(a^*, \delta_{\bar{\theta}}) \in A \times \Delta(\tilde{\Theta})$ is self-confirming if and only if it satisfies relations (28) and (29). Two special cases are noteworthy.

New-classical authority Suppose the monetary authority believes that the policy multiplier is zero, i.e., $\bar{\theta}_{\mathbf{1a}} = -\bar{\theta}_{\mathbf{1}\boldsymbol{\pi}}$. Since in equilibrium $\theta_{\mathbf{1}\boldsymbol{\pi}}$ is identified by the slope of the Phillips regression, we have $\bar{\theta}_{\mathbf{1a}} = -\hat{\beta}^*$. Here the conjectured constraint is vertical at the natural 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 7-(iv)).

[35] Since $\xi_2^* \leq 0$ (Assumption 2), the four cases considered exhaust all possibilities. Also note that, since $\mathbb{E}_{\mu^*}(\theta_{1a}) \leq -\hat{\beta}^*$, in (iv) it holds $\mathbb{E}_{\mu^*}(\theta_{1a}) = -\hat{\beta}^*$ if and only if $\mu^*(\theta_{1a} = -\hat{\beta}^*) = 1$.

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

$$a^* = -\frac{\theta_0^* \hat{\beta}^*}{1 + \hat{\beta}^* \left(\hat{\beta}^* + \theta_{\mathbf{1a}}^*\right)} \quad \text{and} \quad \bar{\theta} = \left(\theta_0^* \left(\frac{1 + \hat{\beta}^{*2}}{1 + \hat{\beta}^* \left(\hat{\beta}^* + \theta_{\mathbf{1a}}^*\right)}\right), 0\right) \tag{31}$$

By Proposition 7, such a policy is hyperactive if $\theta_{\mathbf{1a}}^* > 0$, hypoactive if $\theta_{\mathbf{1a}}^* < 0$, and objectively optimal if $\theta_{\mathbf{1a}}^* = 0$.

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 $a^* = -\theta_0^*(\hat{\beta}^* + \bar{\theta}_{\mathbf{1a}})$ and $\bar{\theta}_0 = \theta_0^*(1 - (\hat{\beta}^* + \bar{\theta}_{\mathbf{1a}})^2)$. Hence, the policy and belief pair

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

is the dogmatic self-confirming equilibrium in a Lucas-Sargent model economy. By Proposition 7, policy a^* is hyperactive when $\bar{\theta}_{\mathbf{1a}} < \theta_{\mathbf{1a}}^*$ and objectively optimal when $\bar{\theta}_{\mathbf{1a}} = \theta_{\mathbf{1a}}^*$. The welfare loss is $\ell(a^*, \theta^*) = \theta_0^{*2}(\hat{\beta}^* + \bar{\theta}_{\mathbf{1a}})^2$.

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

$$\left(0, \delta_{(\theta_0^*, -\hat{\beta}^*)}\right) \tag{32}$$

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 $a^* = -\theta_0^* \hat{\beta}^*$ and $\bar{\theta}_0 = \theta_0^*(1 - \hat{\beta}^{*2})$. The policy and belief pair

$$\left(-\theta_0^* \hat{\beta}^*, \delta_{\left(\theta_0^*\left(1-\hat{\beta}^{*2}\right), 0\right)}\right) \tag{33}$$

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 $\ell\left(a^*, \theta^*\right) = (\theta_0^* \hat{\beta}^*)^2$.

5.4.3 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 $\theta_{\mathbf{1a}}^* = 0$, is $(\theta_0^* \hat{\beta}^*)^2$. It is the same loss of a Keynesian nonzero inflation policy (33) in a new-classical economy: a mistaken new-classical authority has the same lower welfare as a mistaken Keynesian one.

5.5 Policy secularism and a curious interplay

5.5.1 Equilibria

Suppose that the monetary authority is not dogmatic, but has instead a two-models belief.[36] Specifically, it is uncertain whether the true model is Lucas-Sargent or Samuelson-Solow, so that the belief support consists of two points: a Lucas-Sargent model $(\theta_0^{ls}, -\hat{\beta}^*)$ and a Samuelson-Solow model $(\theta_0^{ss}, 0)$. If $\mu_k^* \in [0, 1]$ is the belief in the latter model, we can write belief μ^* as

$$\mu^* = (1 - \mu_k^*)\, \delta_{\left(\theta_0^{ls}, -\hat{\beta}^*\right)} + \mu_k^* \delta_{(\theta_0^{ss}, 0)} \tag{34}$$

Since $\mathbb{E}_{\mu^*}\left(\theta_{\mathbf{1a}}\right) = -\left(1 - \mu_k^*\right) \hat{\beta}^*$, the conjectured 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

$$a^* = -\frac{\theta_0^* \hat{\beta}^* \mu_k^*}{1 + \hat{\beta}^* \mu_k^* \left(\hat{\beta}^* + \theta_{\mathbf{1a}}^*\right)} \tag{35}$$

and

$$\theta_0^{ls} = \frac{\theta_0^*}{1 + \hat{\beta}^* \left(\hat{\beta}^* + \theta_{\mathbf{1a}}^*\right) \mu_k^*} \qquad \theta_0^{ss} = \frac{\theta_0^* \left(1 + \hat{\beta}^{*2} \mu_k^*\right)}{1 + \hat{\beta}^* \left(\hat{\beta}^* + \theta_{\mathbf{1a}}^*\right) \mu_k^*} \tag{36}$$

[36] Cogley 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).

As a result, in this case the pair

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

is a self-confirming equilibrium for every $\mu_k^* \in [0,1]$. We thus have a continuum of equilibria parametrized by the belief μ_k^* in the Samuelson-Solow model (and so by the conjectured multiplier $\mu_k^*\hat{\beta}^*$). In particular, the equilibrium policy a^* is increasing in μ_k^*: the higher the belief in a 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 (36) of the Lucas-Sargent and Samuelson-Solow models depend on the authority's belief μ_k^*: different such beliefs correspond to different Lucas-Sargent and Samuelson-Solow equilibrium specifications. Though the support of the equilibrium belief (34) always contains a specimen of both classes of model economies, that specimen changes as the belief μ_k^* changes. This curious interplay between self-confirming equilibrium models and beliefs is our main finding for the two-models belief case.

Finally, the welfare loss is

$$\ell\left(a^*,\theta^*\right)=\frac{\theta_0^{*2}\left(\hat{\beta}^*\mu_k^*+\hat{\beta}^*+\theta_{\mathbf{1a}}^*\right)^2}{\left(1+\hat{\beta}^*\mu_k^*\left(\hat{\beta}^*+\theta_{\mathbf{1a}}^*\right)\right)^2\left(1+\left(\hat{\beta}^*+\theta_{\mathbf{1a}}^*\right)^2\right)} \tag{37}$$

5.5.2 A new-classical world

To study two-models beliefs further, let us posit a true model. As we did in our study of dogmatism, assume that a Lucas-Sargent model economy $\theta^* = (\theta_0^*, -\hat{\beta}^*)$ is the true model. If so, by (35) and (36) the pair (a^*, μ^*) is a self-confirming equilibrium if and only if $a^* = -\theta_0^*\hat{\beta}^*\mu_k^*$, $\theta_0^{ls} = \theta_0^*$ and $\theta_0^{ss} = \theta_0^*(1+\hat{\beta}^{*2}\mu_k^*)$. Hence, in this case, the pair

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

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

As implied by the analysis of Section 5.5.1, we have a continuum of equilibria parametrized by the belief μ_k^* in the Samuelson-Solow model: if $\mu_k^* > 0$ the equilibrium policy is hyperactive, if $\mu_k^* = 0$ we get the dogmatic new-classical equilibrium (32). Moreover, if $\mu_k^* = 1$ we get back to the dogmatic Keynesian equilibrium (33). Now, however, the equilibrium coefficient θ_0^{ls} is pinned down by the true natural rate of unemployment θ_0^*. In contrast, the equilibrium coefficient $\theta_0^{ss} = \theta_0^*(1-\hat{\beta}^{*2}\mu_k^*)$ still depends on belief μ_k^*: such different beliefs correspond to different Samuelson-Solow equilibrium specifications. In other words, the support of the equilibrium belief always contains a specimen of the Samuelson-Solow model; it, however, changes as belief μ_k^* changes.

The following figures illustrate. The monetary authority is uncertain about the true economic constraint (the vertical line at the natural rate of unemployment) or the Phillips regression line. 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 conjectured 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 the belief μ_k^* via the equilibrium relation $\theta_0 = \theta_0^*(1 - \hat{\beta}^{*2}\mu_k^*)$).

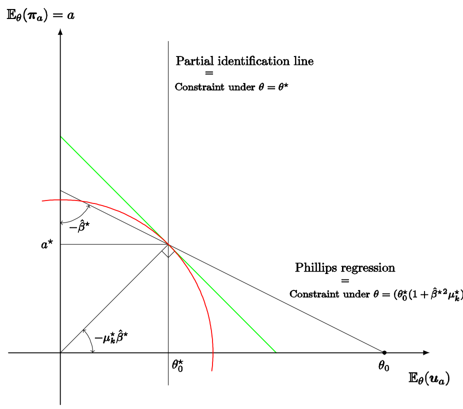

The second figure gives an alternative geometrical representation. Every policy a induces a pair of (objectively) expected rewards, the reward under model θ^*, $R(a, \theta^*)$, and the reward under model θ, $R(a, \theta)$. By changing a one obtains the locus of possible pairs of rewards. If $R(a, \theta^*) \neq R(a, \theta)$, the monetary authority can infer which of the two models is true by looking at its long-run average payoff. Therefore, the partial identification condition is that $R(a, \theta^*) = R(a, \theta)$. At a self-confirming equilibrium (a^*, μ^*) with $\operatorname{supp} \mu^* = \{\theta^*, \theta\}$, 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 pairs $\{(R(a^*, \theta^*), R(a^*, \theta)) : a \in A\}$, the constraint. At this intersection point, the constraint curve must be tangent to the indifference, constant-SEU line with slope $(1 - \mu_k^*)/\mu_k^*$.

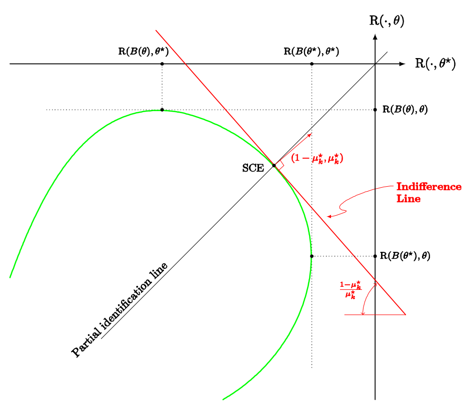

Note that $R(B(\theta^*), \theta^*) = V(B(\theta^*), \delta_{\theta^*}) > V(B(\mu^*), \mu^*)$. Indeed, $V(B(\theta^*), \delta_{\theta^*}) > V(B(\mu^*), \delta_{\theta^*}) = R(B(\mu^*), \theta^*)$ because $B(\mu^*)$ is not a best reply to θ^*. On the other hand, $R(B(\mu^*), \theta^*) = V(B(\mu^*), \mu^*)$ because $R(B(\mu^*), \cdot)$ is constant on the support of self-confirming belief μ^* (see Lemma 3). This is an instance of a more general result relating the values of equilibria with sharper or coarser self-confirming beliefs (see Proposition 9 in Appendix A.4).

A Appendix

A.1 Additional mathematical preliminaries

Since several sigma algebras may be involved in the proofs, given a Borel space $(S, \mathcal{S})$, we sometimes write $\mathcal{B}_S$ instead of $\mathcal{S}$ to denote its Borel sigma algebra. For every Borel set $B \in \mathcal{B}_S$, we let $\mathcal{B}_S \cap B = \{B' \in \mathcal{B}_S : B' \subseteq B\}$ denote the relative sigma algebra on B determined by $\mathcal{B}_S$. Berberian (1997) reviews the properties of Borel spaces.[37]

The next result completes Lemma 1.

Lemma 5 *Let X and Y be Borel spaces and $\varphi : X \to Y$ be measurable. Then:*

(i) $\hat{\varphi} : \Delta(X) \to \Delta(Y)$ is measurable;

(ii) $\hat{\varphi}$ is one-to-one if and only if φ is one-to-one; in this case:

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

Proof The Borel 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 $\varphi^{-1}(C) \in \mathcal{B}_X$ 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 $b \in \mathbb{R}$, that is, $\hat{\varphi}$ is measurable.[38]

If φ is one-to-one, then $\varphi(X) \in \mathcal{B}_Y$ and φ is a measurable isomorphism between $(X, \mathcal{B}_X)$ and $(\varphi(X), \mathcal{B}_Y \cap \varphi(X))$ (Mackey, 1957, Theorem 3.2, see also Berberian, 1997, Theorem 3.2.7). Denote by $\varphi' : \varphi(X) \to X$ the inverse isomorphism. For every $B \in \mathcal{B}_X$, $\varphi(B) \in \mathcal{B}_Y \cap \varphi(X) \subseteq \mathcal{B}_Y$ and so $B = \varphi'(\varphi(B)) = \varphi^{-1}(\varphi(B)) \in \varphi^{-1}(\mathcal{B}_Y)$, then $\mathcal{B}_X \subseteq \varphi^{-1}(\mathcal{B}_Y)$ and the converse inclusion follows from the measurability of φ. Now $\hat{\varphi}(\xi) = \hat{\varphi}(\xi')$ if and only if

[37] A terminological caveat: Berberian (1997) and other authors use Borel space as synonymous with measurable space; they always specify the adjective standard when they assume Polish metrizability (as we do here).

[38] Notice that this part of the statement does not rely on the fact that the measurable spaces $(X, \mathcal{B}_X)$ or $(Y, \mathcal{B}_Y)$ are Borel, but rather on the choice of the natural sigma algebras on $\Delta(X)$ and $\Delta(Y)$.

$\xi\left(\varphi^{-1}(C)\right) = \xi'\left(\varphi^{-1}(C)\right)$ for all $C \in \mathcal{B}_Y$, which implies $\xi(B) = \xi'(B)$ for all $B \in \mathcal{B}_X$, thus $\hat{\varphi}$ is one-to-one. Conversely, for each $x \in X$, we have $\hat{\varphi}(\delta_x)(C) = \delta_x\left(\varphi^{-1}(C)\right) = \delta_{\varphi(x)}(C)$ for all $C \in \mathcal{B}_Y$. Therefore, if φ is not one-to-one, $\hat{\varphi}$ is not one-to-one.

Finally, if $\hat{\varphi}$ is one-to-one, since it is measurable, then $\hat{\varphi}(\Delta(X))$ is a Borel subset of $\Delta(Y)$ and $\hat{\varphi}$ is a measurable isomorphism between $\left(\Delta(X), \mathcal{B}_{\Delta(X)}\right)$ and $\left(\hat{\varphi}(\Delta(X)), \mathcal{B}_{\Delta(Y)} \cap \hat{\varphi}(\Delta(X))\right)$. Now 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 $\nu(\varphi(X)) = \xi\left(\varphi^{-1}(\varphi(X))\right) = \xi(X) = 1$. Thus, when the standard Borel space $(\varphi(X), \mathcal{B}_Y \cap \varphi(X))$ is considered, the restriction of ν to $\mathcal{B}_Y \cap \varphi(X)$ is an element of $\Delta(\varphi(X))$ denoted $\nu_{\varphi(X)}$, that is,

$$\begin{array}{rccc} \iota: & \hat{\varphi}(\Delta(X)) & \to & \Delta(\varphi(X)) \\ & \nu & \mapsto & \nu_{\varphi(X)} \end{array}$$

is a well defined map (which coincides with the inclusion when φ is onto). We want to show that, indeed, ι is a measurable isomorphism between $\left(\hat{\varphi}(\Delta(X)), \mathcal{B}_{\Delta(Y)} \cap \hat{\varphi}(\Delta(X))\right)$ and $\left(\Delta(\varphi(X)), \mathcal{B}_{\Delta(\varphi(X))}\right)$. It is sufficient to prove that it is bijective and measurable (since both spaces are Borel).

First notice 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 $D \in \mathcal{B}_Y \cap \varphi(X)$ $(\subseteq \mathcal{B}_Y)$ and $d \in \mathbb{R}$. Now, for all such sets

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

that is, ι is measurable.

Now assume that $\nu, \nu' \in \hat{\varphi}(\Delta(X))$ and $\nu_{\varphi(X)} = \nu'_{\varphi(X)}$, then 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 $C \cap \varphi(X) \in \mathcal{B}_Y \cap \varphi(X) = \mathcal{B}_{\varphi(X)}$. It follows that

$$\nu(C) = \nu(C \cap \varphi(X)) = \nu_{\varphi(X)}(C \cap \varphi(X)) = \nu'_{\varphi(X)}(C \cap \varphi(X)) = \nu'(C \cap \varphi(X)) = \nu'(C)$$

and so ι is one-to-one.

In order to prove surjectivity of ι, next we show that, for every $\lambda \in \Delta(\varphi(X))$, the set function defined by $\xi(B) = \lambda(\varphi(B))$ for all $B \in \mathcal{B}_X$ belongs to $\Delta(X)$ and $\iota(\hat{\varphi}(\xi)) = \lambda$. First observe that $\xi : \mathcal{B}_X \to [0,1]$ is well defined because $\varphi : (X, \mathcal{B}_X) \to (\varphi(X), \mathcal{B}_Y \cap \varphi(X))$ is a measurable isomorphism. Moreover, for every $B \in \mathcal{B}_X$, $\varphi(B) = (\varphi')'(B) = (\varphi')^{-1}(B)$, thus $\xi(B) = \lambda(\varphi(B)) = \lambda\left((\varphi')^{-1}(B)\right)$ is a probability measure on X. Finally, for every $D \in \mathcal{B}_Y \cap \varphi(X)$,

$$\iota(\hat{\varphi}(\xi))(D) = \hat{\varphi}(\xi)_{\varphi(X)}(D) = \hat{\varphi}(\xi)(D) = \xi\left(\varphi^{-1}(D)\right) = \lambda\left(\varphi\left(\varphi^{-1}(D)\right)\right) = \lambda\left(\varphi\left(\varphi'(D)\right)\right) = \lambda(D)$$

as wanted. ■

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

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

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

- $\{p_\theta\}_{\theta\in\Theta} = p(\Theta) \in \mathcal{B}_{\Delta(S)}$;
- $p : \Theta \to \{p_\theta\}_{\theta\in\Theta}$ *is a measurable isomorphism;*
- $\tilde{p} : \Delta(\Theta) \to \Delta\left(\{p_\theta\}_{\theta\in\Theta}\right)$ *defined by* $\tilde{p}(\mu) = \left(\mu \circ p^{-1}\right)_{p(\Theta)}$ *is a measurable isomorphism and, for every* $\lambda \in \Delta\left(\{p_\theta\}_{\theta\in\Theta}\right)$, *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 $\phi_B : \Delta(S) \to \mathbb{R}$ defined by $\phi_B(\xi) = \xi(B)$ for all $B \in \mathcal{B}_S$, a map $p : \Theta \to \Delta(S)$ 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 $B \in \mathcal{B}_S$, $p(B \mid \theta) = p_\theta(B) = (\phi_B \circ p)(\theta)$ for all $\theta \in \Theta$, thus $p(B \mid \cdot) = \phi_B \circ p$, proving the first part of the statement.[39] The rest follows from the statement of Lemma 5 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 5 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*

$$\begin{array}{rccc} p: & \Theta & \to & \Delta(T \times \Theta) \\ & \theta & \mapsto & q \times \delta_\theta = p_\theta \end{array}$$

is measurable and one-to-one.

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

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

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 $\theta \in \Theta$,

$$q \times \delta_\theta(B) = \int_\Theta q(B^\eta)\, d\delta_\theta(\eta) = q\left(B^\theta\right)$$

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 $\theta \mapsto q\left(B^\theta\right)$ is measurable for all $B \in \mathcal{B}_T \times \mathcal{B}_\Theta$. ■

Corollary 4 *Let* $(S, \mathcal{B}_S)$ *be a Borel space and* $\delta : S \to \Delta(S)$ *the embedding* $s \mapsto \delta_s$. *Then:* $\{\delta_s\}_{s\in S} \in \mathcal{B}_{\Delta(S)}$, $\delta : S \to \{\delta_s\}_{s\in S}$ *is a measurable isomorphism, and* $\lambda \mapsto \lambda \circ \delta$ *is a measurable isomorphism between* $\Delta\left(\{\delta_s\}_{s\in S}\right)$ *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. ■

[39] Notice 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 $\Delta(S)$.

A.2 Feedback and identification

First recall that, for each $a \in A$, $f_a : S \to M$ is measurable and so is $\hat{f}_a : \Delta(S) \to \Delta(M)$. 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 $\sigma \in \Sigma$.

Lemma 6 *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) Recall that we assume that consequences are observable, thus for each action $a \in A$, there exists a measurable function $g_a : M \to C$ such that $\rho_a(s) = g_a(f_a(s))$ for all $s \in S$. (ii) For each $a \in A$, $f_a : S \to M$ is Borel measurable and one-to-one, by Lemma 5, $f_a(S)$ is a Borel subset of M and $f_a : S \to f_a(S)$ is a Borel isomorphism. Then the inverse function $f_a^{-1} : f_a(S) \to S$ is Borel measurable.[40] 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 $s \in S$

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

Setting $h_a = f'_a \circ k_a : M \to M'$ yields the desired result. (iii) Let $(a, \sigma) \in A \times \Sigma$. For every $\sigma' \in \hat{\Sigma}_a(\sigma)$, $\sigma'\left(f_a^{-1}(B_M)\right) = \sigma\left(f_a^{-1}(B_M)\right)$ 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

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

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

[40] *Caveat:* In the proof of Lemma 5, the inverse isomorphism $f_a^{-1} : f_a(S) \to S$ is denoted f'_a, but here f'_a is a section of the feedback function f' that is not an inverse of f.

A.3 Additional definitions

The *self-confirming (equilibrium) correspondence*

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

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*

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

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

We show that self-confirming equilibria with sharper basic subjective assessments have higher values. Formally, μ^* is absolutely continuous with respect to ν^*, denoted $\mu^* \ll \nu^*$, if and only if, for every Borel set $B \subseteq \Sigma$, $\mu^*(B) > 0$ implies $\nu^*(B) > 0$. This means that μ^* rules out more models than ν^*. In particular, if Σ is finite, $\mu^* \ll \nu^*$ is equivalent to $\operatorname{supp}\mu^* \subseteq \operatorname{supp}\nu^*$.

Proposition 9 *If $(a^*, \mu^*), (b^*, \nu^*) \in \Gamma(\sigma^*)$ and $\mu^* \ll \nu^*$, then $V(a^*, \mu^*) \geq V(b^*, \nu^*)$.*

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 (3) for a^* and Proposition 2 deliver

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

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 2. ■

Priors μ^* and ν^* that are mutually absolutely continuous are called *equivalent*, denoted $\mu^* \sim \nu^*$; they share the same possible and impossible models. 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 10 *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 prior captures the decision maker's subjective model uncertainty, which in a self-confirming equilibrium must be consistent with the objective model uncertainty $\hat{\Sigma}_{a^*}(\sigma^*)$ via the relation $\mu^*(\hat{\Sigma}_{a^*}(\sigma^*)) = 1$.[41] 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 11 *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 (3) for a^* and Proposition 2 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 2. ■

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

In words, this result implies that $V(a^*, \mu^*) = V(b^*, \nu^*)$ whenever $\hat{\Sigma}_{a^*}(\sigma^*) = \hat{\Sigma}_{b^*}(\sigma^*)$; that is, two self-confirming actions have the same value when they determine the same collection of probability models that are observationally equivalent with the true model. In this case, differences between the confirmed beliefs justifying the two actions are immaterial; the reason is that in each equilibrium the decision maker correctly predicts the distribution of consequences of *both* actions, which implies that they must yield the same objective expected reward, hence the same value (Proposition 2), otherwise at least one of them would not be a subjective best reply. In particular, the "exchangeability" thesis of Proposition 10 continues to hold even under the belief-free hypothesis $\hat{\Sigma}_{a^*}(\sigma^*) = \hat{\Sigma}_{b^*}(\sigma^*)$: if a^* and b^* are self-confirming equilibrium actions that identify the same set of models, then the sets of confirmed beliefs supporting a^* and b^* coincide.

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.*

[41] Which implies that all possible sets of models are essentially contained in $\hat{\Sigma}_{a^*}(\sigma^*)$.

Under own-action independence of feedback about the state, we have a stronger result. Remark 1 and Lemma 3 imply:

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

Self-confirming actions are thus always objectively optimal when information does not depend on choice. The reason is that, given our structural assumption of observability of consequences, in equilibrium the decision maker correctly predicts the distribution of consequences implied be each action, even if the true model is not identified. In this case partial identification becomes welfare irrelevant and so the analysis of feedback, the main topic of the paper, loses much of its interest. From a decision perspective, own action independence amounts to perfect feedback.

A.5 Other proofs

Proof of Lemma 1 See Lemma 5. ■

Proof of Lemma 2 See Lemma 6. ■

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 $g_a : M \to C$ is $\mathcal{B}_M - \mathcal{B}_C$-measurable; as $f_a : S \to M$ is $\mathcal{F}_a - \mathcal{B}_M$-measurable, then $\rho_a : S \to C$ is $\mathcal{F}_a - \mathcal{B}_C$-measurable. Moreover, $v : C \to \mathbb{R}$ is $\mathcal{B}_C - \mathcal{B}_\mathbb{R}$-measurable and bounded above, and so $r_a = v \circ \rho_a : S \to \mathbb{R}$ is $\mathcal{F}_a - \mathcal{B}_\mathbb{R}$-measurable and bounded above. Thus

$$R_a(\sigma) = \int_S r_a d\sigma = \int_S r_a d\sigma_{|\mathcal{F}_a} \qquad \forall \sigma \in \Delta(S) \tag{38}$$

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 2 If $(a^*, \mu^*) \in A \times \Delta(\Sigma)$ and $\mu^*\left(\hat{\Sigma}_{a^*}(\sigma^*)\right) = 1$, 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 3 It follows immediately from Proposition 2 and Proposition 11. ■

Proof of Proposition 4 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 4 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 4 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 5 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,

$$\int_S h(\rho_a)\, d(q \times \delta_{\theta'}) = \int_S h(\rho_a)\, d(q \times \delta_\theta) \tag{39}$$

for all $h : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ 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 $\int_S \boldsymbol{u} d(q \times \delta_{\theta''}) = \theta_0'' + (\theta_{1\boldsymbol{\pi}}'' + \theta_{1\mathbf{a}}'')\, a$.
4. For $h(u, \pi) = u^2$ and $\theta'' \in \Theta$, we have that $\int_S \boldsymbol{u}^2 d(q \times \delta_{\theta''}) = \left(\theta_0'' + (\theta_{1\boldsymbol{\pi}}'' + \theta_{1\mathbf{a}}'')\, a\right)^2 + (\theta_{1\boldsymbol{\pi}}'')^2 (\theta_3'')^2 + (\theta_2'')^2$.
5. For $h(u, \pi) = u\pi$ and $\theta'' \in \Theta$, we have that $\int_S \boldsymbol{u}\boldsymbol{\pi} d(q \times \delta_{\theta''}) = a\left(\theta_0'' + (\theta_{1\boldsymbol{\pi}}'' + \theta_{1\mathbf{a}}'')\, a\right) + \theta_{1\boldsymbol{\pi}}'' (\theta_3'')^2$.

Given (39), 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

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

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 $(u, \pi) \in \mathbb{R}^2$,

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

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

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

Proof of Lemma 4 Some simple algebra shows that

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

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

Proof of Proposition 6 It holds

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

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

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

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

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

Putting $a = a^*$ we get

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

and so

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

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

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

as desired. ■

Proof of Proposition 7 It holds

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

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

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

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

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

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

and so

$$a^* \gtreqless a^o \Longleftrightarrow a^* \frac{\theta_{\mathbf{1a}}^* - \mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}})}{\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}})} \lesseqgtr 0 \tag{40}$$

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

(i) Suppose $a^* > a^o > 0$. By (24) $\mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}}) \neq -\hat{\beta}^*$ and so $\mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}}) < -\hat{\beta}^*$. By (40), $\left(\theta_{\mathbf{1a}}^* - \mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}})\right)/(\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}})) < 0$, which in turn implies $\mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}}) < \theta_{\mathbf{1a}}^*$. Conversely, suppose $\mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}}) < \theta_{\mathbf{1a}}^*$. Since $\mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}}) \leq -\hat{\beta}^*$, by (24) it follows $a^* > 0$. Moreover, being $\left(\theta_{\mathbf{1a}}^* - \mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}})\right)/(\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}})) < 0$, by (40) it holds $a^* > a^o$. (iii) Suppose $0 < a^* < a^o$. By (40), $\left(\theta_{\mathbf{1a}}^* - \mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}})\right)/(\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}})) > 0$, that is, $\mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}}) \in (\theta_{\mathbf{1a}}^*, -\hat{\beta}^*)$. Conversely, suppose $\mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}}) \in (\theta_{\mathbf{1a}}^*, -\hat{\beta}^*)$. By (24), $a^* > 0$. Moreover, being $\left(\theta_{\mathbf{1a}}^* - \mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}})\right)/(\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}})) > 0$, by (40) it holds $a^* < a^o$. ■

In nice problems 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

$$\ell(\mu^*, \theta^*) = \frac{\theta_0^{*2} \left(\theta_{\mathbf{1a}}^* - \mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}})\right)^2}{\left(1 + \left(\hat{\beta}^* + \theta_{\mathbf{1a}}^*\right)^2\right)\left(1 + \left(\hat{\beta}^* + \theta_{\mathbf{1a}}^*\right)\left(\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}})\right)\right)^2} \tag{41}$$

There is a zero welfare loss if and only if $\mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}}) = \theta_{\mathbf{1a}}^*$, that is, if and only if the monetary authority's expected value of the coefficient $\theta_{\mathbf{1a}}$ is correct. Otherwise, the loss is nonzero, as (41) shows.

Proof of Proposition 8 and eq. (41) First note that

$$R(a^o,\theta^*) = -\theta_0^{*2} - \left(\hat{\beta}^* + \theta_{\mathbf{1a}}^*\right)^2 a^{o2} - \left(\hat{\beta}^*\theta_3^*\right)^2 - \theta_2^{*2} - 2\theta_0^*\left(\hat{\beta}^* + \theta_{\mathbf{1a}}^*\right)a^o - a^{o2} - \theta_3^{*2}$$

and

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

Hence,

$$\begin{aligned}\ell(a^*,\theta^*) &= \max_{a\in A} R(a,\theta^*) - R(a^*,\theta^*) = R(a^o,\theta^*) - R(a^*,\theta^*)\\ &= -\left(\hat{\beta}^* + \theta_{\mathbf{1a}}^*\right)^2\left(a^{o2} - a^{*2}\right) - 2\theta_0^*\left(\hat{\beta}^* + \theta_{\mathbf{1a}}^*\right)(a^o - a^*) - \left(a^{o2} - a^{*2}\right)\end{aligned}$$

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

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

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

$$\ell(a^*,\theta^*) = \theta_0^{*2}\left(\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}})\right)^2 = \theta_0^{*2}\left(\mathbb{E}_{\mu^*}(\theta_{\mathbf{1a}}) - \theta_{\mathbf{1a}}^*\right)^2 \tag{42}$$

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

$$\ell(a^*,\theta^*) = 0 \tag{43}$$

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

$$\begin{aligned}\ell(a^*,\theta^*) &= -\left(\hat{\beta}^* + \theta_{\mathbf{1a}}^*\right)^2\left(a^{o2} - a^{*2}\right) - 2\theta_0^*\left(\hat{\beta}^* + \theta_{\mathbf{1a}}^*\right)(a^o - a^*) - \left(a^{o2} - a^{*2}\right)\\ &= -(a^o - a^*)\left[\left(\hat{\beta}^* + \theta_{\mathbf{1a}}^*\right)^2(a^o + a^*) + 2\theta_0^*\left(\hat{\beta}^* + \theta_{\mathbf{1a}}^*\right) + a^o + a^*\right]\\ &= -(a^o - a^*)\left[\left(\hat{\beta}^* + \theta_{\mathbf{1a}}^*\right)^2(a^o + a^*) - 2a^o\left(1 + \left(\hat{\beta}^* + \theta_{\mathbf{1a}}^*\right)^2\right) + a^o + a^*\right]\\ &= -(a^o - a^*)\left[\left(\hat{\beta}^* + \theta_{\mathbf{1a}}^*\right)^2(a^* - a^o) + a^* - a^o\right] = -(a^o - a^*)\left[\left(\hat{\beta}^* + \theta_{\mathbf{1a}}^*\right)^2 + 1\right](a^* - a^o)\\ &= (a^o - a^*)^2\left[\left(\hat{\beta}^* + \theta_{\mathbf{1a}}^*\right)^2 + 1\right] = (a^* - a^o)^2\left[\left(\hat{\beta}^* + \theta_{\mathbf{1a}}^*\right)^2 + 1\right]\\ &= -(a^* - a^o)^2\frac{\theta_0^*\left(\hat{\beta}^* + \theta_{\mathbf{1a}}^*\right)}{a^o} = -\theta_0^*\left(\hat{\beta}^* + \theta_{\mathbf{1a}}^*\right)\frac{(a^* - a^o)^2}{a^o}\\ &= \left(1 + \left(\hat{\beta}^* + \theta_{\mathbf{1a}}^*\right)^2\right)\left(a^* + \frac{\theta_0^*\left(\hat{\beta}^* + \theta_{\mathbf{1a}}^*\right)}{1 + \left(\hat{\beta}^* + \theta_{\mathbf{1a}}^*\right)^2}\right)^2\end{aligned}$$

In the previous proof we showed that

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

Hence,

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

It is easy to check that, along with (42) and (43), this completes the proof. ■

Proof of eq. (37) It holds

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

as desired. ■

References

[1] Arrow, K. J., Alternative approaches to the theory of choice in risk-taking situations, *Econometrica*, 19, 404-437, 1951.

[2] Battigalli, P., *Comportamento razionale ed equilibrio nei giochi e nelle situazioni sociali*, unpublished thesis, Università Bocconi, 1987.

[3] Battigalli, P., and D. Guaitoli, Conjectural equilibria and rationalizability in a macroeconomic game with incomplete information, Quaderni di Ricerca, Università Bocconi, 1988 (published in *Decisions, Games and Markets*, Kluwer, Dordrecht, 97-124, 1997).

[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., A. Francetich, G. Lanzani, and M. Marinacci, Learning and self-confirming long-run biases, IGIER working paper 588, 2016.

[6] Berberian, S. K., *Borel Spaces*, mimeo, 1997.

[7] Billingsley, P., *Convergence of Probability Measures*, 2$^{\text{nd}}$ ed., Wiley, New York, 1999.

[8] Billingsley, P., *Probability and Measure*, anniversary ed., Wiley, New York, 2012.

[9] 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.

[10] Cho, I.-K., and K. Kasa, Learning and model validation, *Review of Economic Studies*, 82, 45-82, 2015.

[11] Cho, I.-K., N. Williams, and T. Sargent, Escaping Nash inflation, *Review of Economic Studies*, 69, 1-40, 2002.

[12] 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.

[13] 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.

[14] Easley, D., and N. M. Kiefer, Controlling a stochastic process with unknown parameters, *Econometrica*, 5, 1045-1064, 1988.

[15] Fudenberg, D., and D. Kreps, Learning in extensive games, I: self-confirming equilibria, *Games and Economic Behavior*, 8, 20-55, 1995.

[16] Fudenberg, D., and D. K. Levine, Self-confirming equilibrium, *Econometrica*, 61, 523-545, 1993a.

[17] Fudenberg, D., and D. K. Levine, Steady state learning and Nash equilibrium, *Econometrica*, 61, 547-573, 1993b.

[18] Fudenberg, D., and D. K. Levine, Self-confirming equilibrium and the Lucas critique, *Journal of Economic Theory,* 144, 2354–2371, 2009.

[19] Gaballo, G., and R. Marimon, Breaking the spell with credit easing, mimeo, European University Institute, 2015.

[20] Gaudard, M., and D. Hadwin, Sigma-algebras on spaces of probability measures, *Scandinavian Journal of Statistics*, 16, 169-175, 1989.

[21] Kalai, E., and E. Lehrer, Subjective equilibrium in repeated games, *Econometrica*, 61, 1231-1240, 1993.

[22] Kalai, E., and E. Lehrer, Subjective games and equilibria, *Games and Economic Behavior*, 8, 123-163, 1995.

[23] Kurz, M., On rational belief equilibria, *Economic Theory*, 4, 859-876, 1994a.

[24] Kurz, M., On the structure and diversity of Rational Beliefs, *Economic Theory*, 4, 877-900, 1994b.

[25] Lucas, R. E. Jr., Expectations and the neutrality of money, *Journal of Economic Theory,* 4, 103-124, 1972.

[26] Lucas, R. E. Jr., and E. C. Prescott, Investment under uncertainty, *Econometrica*, 39, 659-681, 1971.

[27] Mackey, G. W., Borel structure in groups and their duals, *Transactions of the American Mathematical Society*, 134-165, 1957.

[28] Marinacci, M., Model uncertainty, *Journal of the European Economic Association*, 13, 998-1076, 2015.

[29] Puterman, M. L., *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, Wiley, New York, 2014.

[30] Saint Paul, G., Toward a political economy of macroeconomic thinking, in J. Frenkel and C.Pissarides (eds) *International Seminar on Macroeconomics 2011*, 249-284, University of Chicago Press, Chicago, 2012.

[31] Saint Paul, G., Economic science and political influence, *Journal of the European Economic Association*, 11, 1004–1031, 2013.

[32] Samuelson, P. A., and R. M. Solow, Analytical aspects of anti-inflation policy, *American Economic Review P&P,* 50, 177-184, 1960.

[33] Sargent, T. J., Rational expectations, the real rate of interest, and the natural rate of unemployment, *Brookings Papers on Economic Activity*, 429-472, 1973.

[34] Sargent, T. J., *The conquest of American inflation*, Princeton University Press, Princeton, 1999.

[35] Sargent, T. J., Evolution and intelligent design, *American Economic Review,* 98, 3-37, 2008.

[36] Sargent, T. J., and N. Williams, Impacts of priors on convergence and escapes from Nash inflation, *Review of Economic Dynamics*, 8, 360-391, 2005.

[37] 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.

[38] Savage, L. J., *The Foundations of Statistics*, Wiley, New York, 1954.