

A Framework for the Analysis of Self-Confirming Policies*

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

^aUniversità Bocconi and IGIER

^bNew York University and Hoover Institution

March 2016

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. Finally we provide a more general discussion of self-confirming policies.

KEY WORDS: Self-confirming equilibrium, partial identification, law of large numbers, Keynesian, new classical.

1 Introduction

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 a true stochastic model. Next we illustrate the theory with a characterization of stationary monetary policies in a linear-quadratic setting and with a more general discussion of self-confirming economic policies.

Consider an agent who makes recurrent decisions under uncertainty. In each period he takes an action $a \in A$; given a , the observable outcome $m \in M$ – e.g., a payoff – depends on a random state of nature $s \in S$ according to a feedback function $f_a : S \rightarrow M$. A fixed, unknown stochastic model $\sigma \in \Delta(S)$ – e.g., an urn composition – determines the objective lottery $\hat{f}_a(\sigma) = \sigma \circ f_a^{-1} \in \Delta(M)$ corresponding to each action a ; the agent knows the maps $\sigma \mapsto \hat{f}_a(\sigma)$, but he does not know the stochastic model σ . There are no structural links

*We thank Fabrizio Panebianco, Elia Sartori, and the participants to conferences and seminar presentations at Bocconi University, CISEI-2015 (Capri), D-TEA 2015 (HEC, Paris), LUISS (Rome), New York University, and Scuola Normale Superiore (Pisa) for useful comments. Pierpaolo Battigalli and Massimo Marinacci gratefully acknowledge financial support from ERC grants 324219 and 670337.

between periods, but the agent observes the realized outcome $f_{a_t}(s_t)$ in each period t and therefore updates his subjective beliefs about the fixed unknown model σ . This dynamic setting has been analyzed in the literature on stochastic control (see, e.g., Easley and Kiefer, 1988, and the references therein): given the true model σ^* and the prior belief μ_0 , the intertemporal subjective expected utility maximizing strategy yields a converging active learning process, i.e., a stochastic process of actions and updated beliefs that converges almost surely. The realizations (a^*, μ^*) of the stochastic limit almost surely satisfy 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^* , i.e., $\mu^*(\{\sigma : \hat{f}_{a^*}(\sigma) = \hat{f}_{a^*}(\sigma^*)\}) = 1$; therefore, updated beliefs are constant in the limit;
- (*subjective best reply*) even if the agent cares about the future, action a^* maximizes his one-period subjective expected utility given belief μ^* .

A kind of converse also holds: for every (a^*, μ^*) satisfying the above properties, there is an active learning process so that in the long-run limit a^* is chosen and the agent assigns probability 1 to the set of observationally equivalent models $\{\sigma : \hat{f}_{a^*}(\sigma) = \hat{f}_{a^*}(\sigma^*)\}$. With this, 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 such 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^* . Thus, although equilibrium beliefs are disciplined by long-run empirical frequencies of observations, they do not necessarily coincide with the objective probabilities implied by the true model σ^* . The long-run action a^* may thus be objectively sub-optimal.

Thus, in a self-confirming equilibrium, decision makers might well be best replying to empirically confirmed, but wrong, even illusory, 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.¹ 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 where this important policy issue can be rigorously studied. We 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.

Specifically, we consider a stylized model economy in which a policy maker chooses average inflation, observes an inflation-unemployment outcome (π, u) , and has a standard quadratic loss function. The model economy can be interpreted as reflecting an aggregate

¹In order to remove pervasive inconsistencies of pre-rational-expectations models, rational-expectations models often assume that decision makers know the true data generating model, 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.

response function of a continuum of market agents. With this, we completely characterize the self-confirming equilibrium map that associates each conceivable model economy with a corresponding set of self-confirming beliefs and monetary policies. Given a fixed policy, the monetary authority infers from long-run data the moments of the joint distribution of π and u , 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 (π, u) leaves the monetary authority with a residual one-dimensional uncertainty about the model economy, parametrized 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. We then extend the analysis to more general policy problems.

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

In the game-theoretic literature, a strategy profile that satisfies the properties of confirmed beliefs and subjective best replies 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.² The above mentioned analysis of stochastic control problems applies to games with nature played recurrently with incomplete information about chance probabilities. In games played recurrently against other strategic agents the dynamics are more complex, because the environment of each agent is affected by the behavior of other learning agents, hence, it cannot be assumed to be stationary. Thus, convergence of learning dynamics is not guaranteed, although it is still true that the steady states are self-confirming equilibria, not necessarily Nash equilibria or refinements thereof (e.g., Fudenberg and Kreps, 1995).

Here we focus on self-confirming equilibria of recurrent decision problems, rather than games, for several reasons. First, as explained above, this framework allows for a compelling, active-learning foundation of the equilibrium concept. Second, it allows us to focus better on the interplay between choice and feedback, captured by the long-run frequencies map $a \mapsto \hat{f}_a(\sigma^*)$, given the true stochastic model σ^* . Thus, we can more easily analyze the long-run informativeness of actions, its relation with the equilibrium conditions, and the welfare loss implied by the deviation of self-confirming actions from the objective best reply to the true model. Finally, the true model σ^* can be interpreted as the constant strategy distribution of a large population of individually negligible agents (cf. Fudenberg and Levine, 1993b). As in much of the stochastic control literature, we assume that the decision maker observes

²Kalai and Lehrer’s subjective equilibrium in repeated games is similar to conjectural and self-confirming equilibrium if either players are impatient, or they believe that current choices do not affect the future behavior of opponents. The name “selfconfirming equilibrium” was initially used only for steady states in the recurrent play of sequential games, under the assumption that agents’ feedback in each period is the path of play (terminal node). Then the meaning was extended to encompass more general assumptions about feedback. See the discussion in Battigalli et al. (2015).

the relevant consequences resulting from actions and state realizations. This observability assumption features prominently in the game-theoretic work of Battigalli et al. (2015) from which we draw some preliminary results. While they focus on the impact of non-neutral attitudes toward model uncertainty, such as ambiguity aversion, in this paper we assume that the decision maker is a standard subjective expected utility maximizer. For the (linear-quadratic) applications considered here, this assumption is without loss of generality because we can show that, although ambiguity attitudes may affect the active learning dynamics, they do not affect the set of limit points of such dynamics, that is, the self-confirming equilibrium set.

Our work is related to a strand of literature 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.³ Cho et al. (2002) and Sargent and Williams (2005) sharpen the theoretical analysis of such learning dynamics.⁴ 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-confirming equilibrium cannot persist in the long run, but – unlike Sargent (1999) – they show that low inflation is a stable equilibrium outcome, as in the very long run, for most of the time, the monetary authority adopts a vertical Phillips curve model.⁵ In our paper we only focus 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. Also, 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

³See also, Cogley and Sargent (2005), Sargent, Williams, and Zha (2006), and Cogley et al. (2007).

⁴The phrase “escaping Nash inflation” in the title of this article 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 this model.

⁵In 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.

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. Finally, Gaballo and Marimont (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.⁶

Structure of the paper The rest of the paper is organized as follows. The first part (Sections 2-4) develops an abstract analysis of self-confirming choices in decision problems with model uncertainty, the second part (Sections 5-6) illustrates the abstract analysis with applications to economic policies. More in detail, Section 2 provides the preliminary elements of decision problems with model uncertainty; Section 3 specifies the key partial-identification problem; Section 4 analyzes the self-confirming equilibrium correspondence; Section 5 analyzes self-confirming monetary policies; Section 6 sketches a more general analysis of self-confirming economic policies. Section 7 is an appendix collecting some more technical material and all the formal proofs.

2 Preliminaries

2.1 Mathematics

A pair (X, \mathcal{X}) is a (standard) Borel space if there exists a metric that makes the space X 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.

We denote by $\Delta(X)$ the collection of all probability measures on X . We endow $\Delta(X)$ with the natural sigma algebra,⁷ 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 $\varphi : X \rightarrow Y$ induces a measurable distribution map $\hat{\varphi} : \Delta(X) \rightarrow \Delta(Y)$ defined by $\hat{\varphi}(\xi) = \xi \circ \varphi^{-1}$ for each probability measure $\xi \in \Delta(X)$; that is, $\hat{\varphi}(\xi)(B) = \xi(\varphi^{-1}(B))$ for all Borel sets $B \in \mathcal{Y}$. The following routine lemma describes a key feature of $\hat{\varphi}$.

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

⁶The 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.

⁷It is generated by the functions $\phi_B : \Delta(X) \rightarrow \mathbb{R}$ defined by $\phi_B(\xi) = \xi(B)$ for all $B \in \mathcal{X}$ (cf. Theorem 2.3 of Gaudard and Hadwin, 1989, and Theorem 15.15 of Aliprantis and Border, 2006).

Unless otherwise stated, throughout the paper spaces are assumed to be Borel spaces, as usual in stochastic optimization.⁸ That said, to fix ideas the reader can think of X as a finite space of cardinality n , with \mathcal{X} being its power set 2^X .⁹ In this case $\Delta(X)$ can be identified with the simplex of \mathbb{R}^n , that is, $\Delta(X) = \{\xi \in \mathbb{R}_+^n : \sum_{i=1}^n \xi_i = 1\}$, and 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 $f : X \rightarrow \mathbb{R}$.

2.2 Classical subjective expected utility

Let S be a space of *states* of nature, A a space of *actions* available to the decision maker, C a space of consequences, and $\rho : A \times S \rightarrow C$ a measurable *consequence function* that associates a consequence $\rho(a, s) \in C$ 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. As in Savage (1954), we assume that decision makers are indifferent among actions that share the same consequences in every state – i.e., $\rho(a', s) = \rho(a, s)$ for all $s \in S$. This is a form of consequentialism.

The inherent randomness that characterizes states’ realizations – often called physical uncertainty – is described by probability models $\sigma \in \Delta(S)$, which 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 $v : C \rightarrow \mathbb{R}$ 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 $r : A \times S \rightarrow \mathbb{R}$ 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 will be well defined. For every action $a \in A$, the section $R(a, \cdot) : \Delta(S) \rightarrow [-\infty, \infty)$ is measurable and bounded above too.

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 “physical” information behind Σ .¹⁰

⁸See, e.g., Bertsekas and Shreve (1978), Dynkin and Yushkevich (1979), and Puterman (2014).

⁹The power set 2^X is the collection of all subsets of X ; it is the Borel σ -algebra of X under the discrete metric.

¹⁰See Marinacci (2015) for a discussion of this setup; classical SEU is proposed by Cerreia-Vioglio et al. (2013).

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* $\sigma_{\mu}(E) = \int_{\Sigma} \sigma(E) d\mu(\sigma)$ 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, $\mu = \delta_{\sigma}$, decision makers believe (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 $\{\sigma\}$, decision makers have maximal structural information and, as result, know that σ is the true model. 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 decision makers either observed infinitely many drawings from a given urn or if they just were able to count the balls of each color, they 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 $R(a, \mu)$.¹¹

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* $f : A \times S \rightarrow M$, 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.¹² The decision maker's information about the state is thus endogenous: if M is finite, the information is represented by the partition $\{f_a^{-1}(m) : m \in M\}$ of the state space S that the messages induce, which depends on the choice of action a ; if M is infinite, this partition is replaced by the sigma algebra

$$\mathcal{F}_a = \{f_a^{-1}(B) : B \in \mathcal{M}\} \subseteq \mathcal{S}. \quad (2)$$

¹¹See Corollary 6 in Appendix 7.1.

¹²Here $f_a : S \rightarrow M$ denotes the section $f(a, \cdot)$ of f at a .

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.¹³ When this is not the case, information is *imperfect* (maximally imperfect when 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 $g_a : M \rightarrow C$ such that

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

In words, messages encode consequences. In particular, when the consequences of the actions are the only observed messages, we have $C = M$ and $f = \rho$.

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 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. \blacktriangle

¹³In general decision problems (or games), perfect feedback means that the terminal node of the game/decision tree is observed ex post. This implies that f_a is one-to-one if the choices of nature and the decision maker are essentially simultaneous, as assumed here. But the implication does not hold in general. Suppose, for example, that an outcome $y \in Y$ is selected after action a according to an unknown strategy $s \in S = Y^A$, with consequences determined through a known function $\gamma : A \times Y \rightarrow C$. Then $\rho(a, s) = \gamma(a, s(a))$, whereas perfect feedback implies that $f_a^{-1}(f_a(s)) = \{s' \in S : s'(a) = s(a)\}$. Thus, (if A and Y have at least two elements) s is not observed ex post and, furthermore, own-action independence of feedback 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, so that $\nu(m)$ is the empirical frequency of message m . Given an action $a \in A$, consider the distribution map $\hat{f}_a : \Sigma \rightarrow \Delta(M)$ given, for each $\sigma \in \Sigma$, by

$$\hat{f}_a(\sigma) = \sigma \circ f_a^{-1};$$

that is, $\hat{f}_a(\sigma)(B) = \sigma(s : f_a(s) \in B)$ for each $B \in \mathcal{M}$.¹⁴ In words, $\hat{f}_a(\sigma)(m)$ is the empirical frequency with which the decision maker receives message m when he chooses action a and σ is the true model.¹⁵ 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. In this case the decision problem is identified under a since different models generate different message distributions, which thus uniquely pin down models. Otherwise, $\hat{f}_a^{-1}(\nu)$ is nonsingleton and 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

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

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

$$\hat{\Sigma}_a(\sigma) = \left\{ \sigma' \in \Sigma : \hat{f}_a(\sigma') = \hat{f}_a(\sigma) \right\} \tag{3}$$

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 and messages.

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 $(a, a', \sigma) \in A \times A \times \Sigma$.

¹⁴In the literature $\hat{f}_a(\sigma)(B)$ is sometimes denoted by $\hat{f}_a(B | \sigma)$, 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 .

¹⁵Appendix 7.2 makes rigorous the long run interpretation which we rely upon.

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 function, with $\hat{\Sigma}_a(\sigma) = \{\sigma\}$ 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,

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

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:

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

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

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

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 $\{\hat{\Sigma}_b(\sigma)\}_{\sigma \in \Sigma} = \{\hat{\Sigma}_b(\sigma_n)\}_{n=0, \dots, 90}$ 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. \blacktriangle

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

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

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. \blacktriangle

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 $h_a : M \rightarrow M'$ such that

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

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 $a \in A$,*

$$\hat{\Sigma}_a(\sigma) \subseteq \hat{\Sigma}'_a(\sigma) \quad \forall \sigma \in \Sigma.$$

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

$$V(a^*, \mu^*) \geq V(a, \mu^*) \quad \forall a \in A \quad (4)$$

and

$$\mu^* \in \Delta(\hat{\Sigma}_{a^*}(\sigma^*)) \quad (5)$$

The definition relies on two pillars: the optimality condition (4) that ensures that action a^* is optimal under belief μ^* ; the data confirmation condition (5) that guarantees that belief μ^* is consistent with the data that action a^* reveals.¹⁶

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

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

$$\begin{aligned} R(b, \sigma^*) = R(b, \sigma_2) = R(b, \sigma_1) = R(y, \sigma_2) = R(g, \sigma_2) &= \frac{1}{3} \\ R(y, \sigma^*) = R(g, \sigma_1) = \frac{2}{3} \quad ; \quad R(y, \sigma_1) = R(g, \sigma^*) &= 0 \end{aligned} \quad (6)$$

- (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 (4) 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 (5) 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 (4) is easily seen to be satisfied, while the data confirmation condition (5) 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.

¹⁶Here, 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 $\hat{\Sigma}_{a^*}(\sigma^*)$.

(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: $R(b, \sigma^*) = 1/3 \neq 2/3 = R(y, \sigma^*)$. \blacktriangle

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 an all-Yellow urn is among the possible models, i.e. $\delta_{\delta_Y} \in \Sigma$ (as is the case when there is no information about the urn), then this is no 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 $\delta_Y \in \hat{\Sigma}_g(\sigma^*) = \{\sigma \in \Delta_q(\{B, G, Y\}) : \sigma(G) = 0\}$. 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. \blacktriangle

Under own-action independence of feedback (that is, actions do not affect information gathering), the data confirmation condition (5) 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 (4) becomes

$$R(a^*, \sigma^*) \geq R(a, \sigma^*) \quad \forall a \in A \quad (7)$$

since condition (5) 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).¹⁷

Action a^* is *objectively optimal* if it satisfies condition (7). 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.

Example 6 In Example 4, bet y is the objectively optimal action. \blacktriangle

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 (4) 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 (5) implies that the predictive probability σ_{μ^*} belongs to $\hat{\Sigma}_{a^*}(\sigma^*)$ if it belongs to Σ .¹⁸ In this case, $(a^*, \delta_{\sigma_{\mu^*}})$ is a self-confirming equilibrium too. Hence we have the following *certainty equivalence principle*:

¹⁷For condition (5) to reduce to (7) it is actually enough that the equilibrium action a^* be fully revealing, a weaker property than perfect feedback (see Corollary 2 below).

¹⁸The conjectural equilibrium conditions, stated for games by Battigalli (1987), are written in predictive form.

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 $V(a^*, \mu^*) = V(a^*, \delta_{\sigma_{\mu^*}})$.*

More generally, even if $\sigma_{\mu^*} \notin \Sigma$, we have

$$\hat{f}_{a^*}(\sigma_{\mu^*}) = \hat{f}_{a^*}(\sigma^*) \quad (8)$$

that is, the predictive probability and the true model both assign the same probability to messages. Condition (8) is an alternative data confirmation condition in terms of predictive distributions, weaker than (5).

Finally, consider

$$\sigma_{\mu^*} = \sigma^* \quad (9)$$

a much stronger data confirmation condition than (8). If (a^*, μ^*) is a self-confirming pair in which μ^* satisfies condition (9), then a^* is easily seen to be objectively optimal. Any belief that satisfies condition (9) is thus equivalent, in terms of actions, to the belief $\mu^* = \delta_{\sigma^*}$ concentrated on the true model. It does not matter how “accurate” is the belief, as long as its predictive probability coincides with the true model. If so, objectively optimal actions arise.

4.2 Steady state interpretation

To give perspective on the self-confirming equilibrium concept, we discuss in more detail its steady state interpretation by adapting to our framework some results in stochastic control, as found in the classic work of Easley and Kiefer (1988). To ease matters, throughout this subsection we assume that the sets M , S , and Σ are *finite*.

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

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

The next result (cf. Lemma 2 in Easley and Kiefer, 1988) shows that priors and posteriors are equal under the data confirmation condition (5). Updating thus no longer operates when beliefs satisfy this condition, a property that clarifies the steady state rationale of the self-confirming equilibrium notion as a rest point of a learning process (about which our steady state analysis is silent).

Lemma 3 *Let $a \in A$ and $\text{supp } \mu \cap \hat{\Sigma}_a(\sigma^*) \neq \emptyset$. Then $\mu \in \Delta(\hat{\Sigma}_a(\sigma^*))$ if and only if $\mu(\cdot | a, m) = \mu(\cdot)$ for all $m \in \text{supp } \hat{f}_a(\sigma_\mu)$.*

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

end of each period $t \geq 1$ and the beginning of period $t + 1$. If action a_t is chosen in period t and message m_t is observed at the end of the same period, then

$$\mu_t(\sigma | a_t, m_t) = \mu_{t-1}(\sigma) \frac{\hat{f}_{a_t}(\sigma)(m_t)}{\hat{f}_{a_t}(\sigma_\mu)(m_t)} \quad \forall \sigma \in \Sigma \quad (10)$$

is the Bayes update of μ_{t-1} given (a_t, m_t) . The decision maker carries out a stationary strategy¹⁹

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

that maximizes the discounted expected utility

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

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

Every stationary strategy α^* yields a corresponding stochastic process of actions and beliefs $(\alpha^*(\boldsymbol{\mu}_{t-1}^*), \boldsymbol{\mu}_{t-1}^*)_{t \geq 1}$ defined on probability space $(S^\infty, \sigma^\infty)$ and adapted to the filtration on S^∞ induced by feedback f , where $\boldsymbol{\mu}_0^* = \mu_0$ and $\boldsymbol{\mu}_t^*(s^\infty)$ is the Bayes update of $\boldsymbol{\mu}_{t-1}^*(s^\infty)$ given action $\alpha^*(\boldsymbol{\mu}_{t-1}^*(s^\infty))$ and message $f(\alpha^*(\boldsymbol{\mu}_{t-1}^*(s^\infty)), s_t)$. Next we provide simple sufficient conditions for such an *active learning process* to converge to a self-confirming equilibrium. Of course, since the realized history of states $(s_t)_{t \geq 1}$ is random, the limit of the process is, in general, also random. With this, we say that a random pair $(\mathbf{a}^*, \boldsymbol{\mu}^*) : S^\infty \rightarrow A \times \Delta(\Sigma)$ is a *random self-confirming equilibrium* given model σ^* if, $\sigma^{*, \infty}$ -almost surely, the pair $(\mathbf{a}^*, \boldsymbol{\mu}^*)$ is a self-confirming equilibrium given σ^* .

Proposition 2 *Let (D, f, β, μ_0) satisfy the following assumptions:*

- (i) *A is a compact and convex subset of an Euclidean space;*
- (ii) *$a \mapsto \hat{f}_a$ is continuous, and $\hat{f}_a(\sigma)(m) > 0$ for every $(a, \sigma, m) \in A \times \Sigma \times M$;*
- (iii) *$r : A \times S \rightarrow \mathbb{R}$ is continuous and strictly concave on A;*
- (iv) *$\mu_0(\sigma) > 0$ for every $\sigma \in \Sigma$.*

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

¹⁹A strategy is stationary if it depends only on updated beliefs about the model. As the set of optimal strategies always contains a stationary strategy, the stationarity restriction is without substantial loss of generality.

The proposition is, essentially, a self-confirming interpretation of stochastic control results of Easley and Kiefer (1988) and so we omit its proof. Heuristically, the existence of an optimal stationary strategy follows from standard results in dynamic programming. The martingale convergence theorem implies that there exists a random belief $\mu_\infty^* : S^\infty \rightarrow \Delta(\Sigma)$ such that, $\sigma^{*,\infty}$ -almost surely, $\mu_t^* \rightarrow \mu_\infty^*$ (cf. Theorem 4 in Easley and Kiefer, 1988). The Bayesian updating function, which via (10) associates μ_t to μ_{t-1} given actions and messages (a_t, m_t) , is continuous on $\Delta(\Sigma)$. Hence, for every sample path $s^\infty = (s_t)_{t \geq 1}$ such that $\mu_t^*(s^\infty) \rightarrow \mu_\infty^*(s^\infty)$ (hence, by $\sigma^{*,\infty}$ -almost sure convergence, every s^∞ such that $\prod_{t=1}^T \sigma^*(s_t) > 0$ for all T) and for every action a_∞ in the set of limit points of $(\alpha^*(\mu_t^*(s^\infty)))_{t \geq 1}$ (which is not empty, by compactness of A), it must be the case that $\mu_\infty^*(s^\infty)$ is a_∞ -invariant, that is, $\mu_\infty^*(s^\infty)(\cdot | a_\infty, m) = \mu_\infty^*(s^\infty)(\cdot)$ for every possible m . In view of Lemma 3, this is equivalent to the confirmed beliefs condition (5). Since the value of experimentation vanishes in the limit, every such action a_∞ maximizes the one-period subjective expected value $V(\cdot, \mu_\infty^*(s^\infty))$ (cf. Lemma 4 in Easley and Kiefer, 1988). By assumption (iii), $V(\cdot, \mu_\infty^*(s^\infty))$ has a unique maximizer, therefore $\lim_{t \rightarrow \infty} (\alpha^*(\mu_t^*(s^\infty)), \mu_t^*(s^\infty)) = (\arg \max_{a \in A} V(a, \mu_\infty^*(s^\infty)), \mu_\infty^*(s^\infty))$, which is a self-confirming equilibrium (cf. Theorem 5 in Easley and Kiefer, 1988).

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

4.3 Value

4.3.1 Equilibria

The true model anchors the decision problem (D, f) . The *self-confirming (equilibrium) correspondence*

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

associates to each possible true model σ^* the collection $\Gamma(\sigma^*)$ of its self-confirming equilibria (a^*, μ^*) . Through this correspondence we can characterize the value of self-confirming equilibria.

Proposition 3 $V(a^*, \mu^*) = R(a^*, \sigma^*)$ for each $(a^*, \mu^*) \in \Gamma(\sigma^*)$.

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

Lemma 4 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 $(a^*, \mu^*) \in A \times \Delta(\Sigma)$ belongs to $\Gamma(\sigma^*)$ if and only if $R(a^*, \sigma^*) \geq V(a, \mu^*)$ for every $a \in A$ and the data confirmation condition (5) holds.

Under the data confirmation condition, the optimality condition (4) thus 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 that the substantial bite of the data confirmation condition.

4.3.2 Model uncertainty

A basic subjective assessment of the decision maker is about which models in Σ he deems actually possible; if Σ is finite, these are the models with strictly positive mass according to his prior.²⁰ Next 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^* .

Proposition 4 *If $(a^*, \mu^*), (b^*, \nu^*) \in \Gamma(\sigma^*)$ and $\mu^* \ll \nu^*$, then $V(a^*, \mu^*) \geq V(b^*, \nu^*)$.*

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

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 $\mu^*(\hat{\Sigma}_{a^*}(\sigma^*)) = 1$.²¹ 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 6 *If $(a^*, \mu^*), (b^*, \nu^*) \in \Gamma(\sigma^*)$ and $\hat{\Sigma}_{a^*}(\sigma^*) \subseteq \hat{\Sigma}_{b^*}(\sigma^*)$, then $V(a^*, \mu^*) \geq V(b^*, \nu^*)$.*

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 3), otherwise at least one of them would not be a subjective best reply. In particular, the “exchangeability” thesis of Proposition 5 continues to

²⁰ A (well known) measure theoretic caveat: if S is infinite, not single models but sets of them have to be considered.

²¹ Which implies that all possible sets of models are essentially contained in $\hat{\Sigma}_{a^*}(\sigma^*)$.

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.

4.4 Welfare

In order to do welfare analysis it is convenient to consider the (*equilibrium*) *action correspondence*

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

that associates to each possible true model σ^* 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^* .

Let σ^* be the true model. Since $V(a^*, \mu^*) = R(a^*, \sigma^*) \leq \max_{a \in A} R(a, \sigma^*)$, the decision maker incurs a welfare loss

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

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 caused by the decision maker's ignorance. By Proposition 6,

$$\hat{\Sigma}_{a^*}(\sigma^*) \subseteq \hat{\Sigma}_{b^*}(\sigma^*) \implies \ell(a^*, \sigma^*) \leq \ell(b^*, \sigma^*) \quad \forall a^*, b^* \in \gamma(\sigma^*).$$

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 (one that is optimal from a purely statistical viewpoint) to be self-confirming amounts to being objectively optimal. An action that allows the decision maker to know the truth (or to get as close as possible to it) has to be optimal in light of the decision maker's objective: Truth is ancillary to the decision maker pursuit of his goals (and so of his happiness).

Proposition 7 *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 this action 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 optimal in the decision problem. In this sense, there is no separation between estimation and decision in the present setup.

Example 7 Consider the decision problem

$A \setminus S$	s_1	s_2	
a_1	1	1	(11)
a_2	0	2	

where $r = \rho = f$ and $\Sigma = \Delta(S)$. Here the decision maker is risk neutral, has no structural information, and the (monetary) consequences are the messages that he receives. Given any $\sigma \in \Sigma$, for the constant action a_1 we have $\hat{\Sigma}_{a_1}(\sigma) = \Sigma$, while for the non-constant 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 action is always self-confirming, independently of the true model, while the non-constant one is self-confirming only when the true model makes it objectively optimal. \blacktriangle

The next result further illustrates the point by showing that fully revealing actions are self-confirming if and only if they are objectively optimal.

Corollary 2 *A fully revealing action is self-confirming if and only if it is objectively optimal.*

Under own-action independence of feedback, we have a stronger result. Remark 1 and Lemma 4 imply (cf. Battigalli et al. 2015, Proposition 1):

Corollary 3 *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 by 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.

We close with some simple comparative statics. The extent of the partial identification loss that arises when σ^* is the true model is described by the set $L(\sigma^*) = \{\ell(a^*, \sigma^*) : a^* \in \gamma(\sigma^*)\}$. It is a singleton if and only if every self-confirming action is objectively optimal; in this case $L(\sigma^*) = \{0\}$. Otherwise it is a set that accounts for different losses that different self-confirming actions in $\gamma(\sigma^*)$ may cause.

In a decision problem with feedback (D, f) , we are interested in carrying out comparative statics results in f , that is, in information. The next result, a simple consequence of Lemma 2, shows that the set of self-confirming equilibria increases as feedback becomes coarser. The same behavior is featured by the partial identification loss.

Proposition 8 *Given decision problems with feedback (D, f) and (D, f') , if f' is coarser than f , then $\Gamma'(\sigma) \supseteq \Gamma(\sigma)$ and $L'(\sigma) \supseteq L(\sigma)$ for all $\sigma \in \Sigma$.*

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 adopt and 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.

5.1 Steady state model economies

We study a class Θ of model economies θ at a (stochastic) steady state in which unemployment u and inflation π are affected by random shocks ε and w and by a monetary policy variable a . Unemployment and inflation outcomes are connected to shocks and the government policy according to

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

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

The vector parameter $\theta = (\theta_0, \theta_{1\pi}, \theta_{1a}, \theta_2, \theta_3) \in \mathbb{R}^5$ specifies the structural coefficients of an aggregate supply equation (12). Coefficients $\theta_{1\pi}$ and θ_{1a} are slope responses of unemployment to actual and planned inflation,²² while the coefficients θ_2 and θ_3 quantify shock volatilities (see Sargent, 2008, p. 18). Finally, the intercept θ_0 is the 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 $\theta_0 > 0$, $\theta_{1\pi} < 0$, $\theta_2 > 0$ and $\theta_3 > 0$.

In words, we posit a strictly positive intercept, as well as strictly positive shock coefficients (nontrivial, possibly asymmetric, shocks affect both the inflation and unemployment equations). Finally, we assume that inflation and unemployment are inversely related.

The reduced form of each model economy is

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

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

The coefficients of the reduced form are $\xi = (\theta_0, \theta_{1\pi} + \theta_{1a}, \theta_{1\pi}\theta_3, \theta_2, \theta_3) \in \mathbb{R}^5$. It is convenient to rewrite the reduced form through a bivariate random variable $(\mathbf{u}, \boldsymbol{\pi})$ defined by²³

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

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

²²The economic interpretation is that planned inflation a affects agents' expectations to an extent parameterized by θ_{1a} .

²³Formally, $(\mathbf{u}, \boldsymbol{\pi}) : A \times W \times E \times \Theta \rightarrow U \times \Pi$ where $A = [0, +\infty)$ and $E = W = U = \Pi = \mathbb{R}$.

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 $\xi \neq \xi'$.

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

Assumption 2 $\xi_2 \leq 0$.

In words, we assume that an increase in planned inflation never increases unemployment. To ease matters we will tacitly maintain this assumption on the multiplier, though our analysis can be extended beyond it (Section 7.5).

Be that as it may, in what follows

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

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

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 $\beta \equiv \theta_{1\pi} = -\theta_{1a}$, and so $\theta = (\theta_0, \beta, -\beta, \theta_2, \theta_3)$ and $\xi = (\theta_0, 0, \beta\theta_3, \theta_2, \theta_3)$. In 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

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

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

5.2 The policy problem: setup and identification

5.2.1 Setup

The monetary authority chooses policy a . 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.²⁴ The consequence space C consists of unemployment and inflation pairs (u, π) , so we set $C = U \times \Pi$. The outcome function $\rho : A \times (W \times E \times \Theta) \rightarrow C$ is

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

which is the unemployment/inflation pair (u, π) determined by policy a and state (w, ε, θ) . Note that ρ is the reduced form of the model economy, with matrix representation

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

5.2.2 Factorization

We assume that the messages received by the monetary authority are their policies' consequences, that is, $f = \rho$. Hence, a message $m = (u, \pi)$ consists of an unemployment and inflation pair. When it chooses policy a and receives message (u, π) , the monetary authority can infer a set of possible states (w, ε, θ) through the inverse correspondence $\rho_a^{-1} : C \rightarrow 2^{W \times E \times \Theta}$. In particular,

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

and

$$\hat{\Sigma}_a(\sigma) = \left\{ \sigma' \in \Sigma : \hat{f}_a(\sigma') = \hat{f}_a(\sigma) \right\} \quad \forall \sigma \in \Sigma \quad (18)$$

At this point, it is convenient to enrich this setup. Within a state $s = (w, \varepsilon, \theta)$ the pair (w, ε) represents random shocks, while θ parametrizes a model economy. This suggests factoring the probability models $\sigma \in \Sigma \subseteq \Delta(W \times E \times \Theta)$ as

$$\sigma = q \times \delta_\theta \quad (19)$$

where $q \in \Delta(W \times E)$ is known and $\delta_\theta \in \Delta(\Theta)$ is a Dirac probability measure concentrated on a given economic model $\theta \in \Theta$. We thus parametrize probability models with θ and write σ_θ .

The 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 (19) 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

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

²⁴Section 6.2 further discusses the Cartesian structure of the state space.

and so to define the prior μ on θ .²⁵

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

$$V(a, \mu) = \int_{\Theta} \left(\int_{W \times E} r(a, w, \varepsilon, \theta) dq(w, \varepsilon) \right) d\mu(\theta), \quad (20)$$

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 tacitly maintain the following assumption on the known shock distributions.²⁶

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 suitably normalized.

5.2.3 Identification

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

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

With this, a sharp identification result holds.

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

$$\hat{\Sigma}_a(\theta) = \left\{ \theta' \in \Theta : \theta'_0 + \theta'_{1\mathbf{a}}a = \theta_0 + \theta_{1\mathbf{a}}a, \theta'_{1\pi} = \theta_{1\pi}, \theta'_2 = \theta_2, \theta'_3 = \theta_3 \right\} \quad (21)$$

Given the true model θ , the shock coefficients θ_2 and θ_3 are thus identified, along with the slope $\theta_{1\pi}$ of the Phillips curve, independently of the chosen policy a . As we discuss below, the intercept of the curve is also identified, but it depends on the maintained policy a through the unidentified parameter $\theta_{1\mathbf{a}}$. This important identification result is made possible by some moment conditions, formally spelled out in the proof. We can, however, heuristically describe them via the bivariate random variable $(\mathbf{u}_a, \boldsymbol{\pi}_a) : W \times E \times \Theta \rightarrow C$ that, for a given policy a , represents the unemployment and inflation rates determined by the state (w, ε, θ) .²⁷ The monetary authority observes in the long run the following moments:

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

²⁵The map $\theta \mapsto q \times \delta_{\theta}$ is bijective and measurable. See Corollary 5 in the appendix.

²⁶Whenever convenient, in what follows we will use the shorthand notation \mathbb{E} for integrals, for example $\mathbb{E}_q(\varepsilon) = \int_{W \times E} \varepsilon dq(w, \varepsilon)$.

²⁷Formally, \mathbf{u}_a and $\boldsymbol{\pi}_a$ are the sections $\mathbf{u}(a, \cdot)$ and $\boldsymbol{\pi}(a, \cdot)$ at policy a of the random variables \mathbf{u} and $\boldsymbol{\pi}$, respectively.

Therefore,

$$\theta_{1\pi} = \frac{\text{Cov}_\theta(\mathbf{u}_a, \boldsymbol{\pi}_a)}{\text{Var}_\theta(\boldsymbol{\pi}_a)} \quad (22)$$

is the beta coefficient of the Phillips regression of unemployment over inflation,²⁸

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

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

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

$$\theta_0 + \theta_{1\mathbf{a}}a = \mathbb{E}_\theta(\mathbf{u}_a) - \frac{\text{Cov}_\theta(\mathbf{u}_a, \boldsymbol{\pi}_a)}{\text{Var}_\theta(\boldsymbol{\pi}_a)} \mathbb{E}_\theta(\boldsymbol{\pi}_a) \quad (23)$$

where the right hand 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 doesn't appreciate.

5.2.4 Estimated model economy

The moments that identify the three coefficients $\theta_{1\pi}$, θ_2 , and θ_3 do not depend on the chosen policy a but only on the true model θ . To emphasize this key feature, we denote by $\hat{\beta}$ the beta regression coefficient that identifies $\theta_{1\pi}$,²⁹ by $\hat{\sigma}_{\mathbf{u}|\boldsymbol{\pi}}$ 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 $\hat{\alpha}(a)$.

In terms of this notation we can write

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

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

$$u = \hat{\alpha}(a) + \hat{\beta}\pi + \hat{\sigma}_{\mathbf{u}|\boldsymbol{\pi}}w \quad (24)$$

$$\pi = a + \hat{\sigma}_\boldsymbol{\pi}\varepsilon \quad (25)$$

$$\hat{\alpha}(a) = \theta_0 + \theta_{1\mathbf{a}}a \quad (26)$$

In particular, (24) is the estimated aggregate supply equation and (25) is the estimated inflation equation. The intercept of the former equation depends on the policy a via the equality (26), 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.

²⁸The Phillips regression $u = \alpha + \beta\pi$ is run by the monetary authority using long run data.

²⁹By 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.

5.2.5 Partial identification line


The monetary authority is uncertain, also in the long run, about the two structural coefficients θ_0 and θ_{1a} . The former is the average unemployment at zero planned inflation, $\theta_0 = \mathbb{E}_\theta(\mathbf{u}_0)$; the latter is the “direct” impact of policy on unemployment.

The parameter space of the estimated model economy (24)-(26) reduces to $\Theta = \tilde{\Theta} \times \{(-\hat{\beta}, \hat{\sigma}_{u|\pi}, \hat{\sigma}_u)\}$, where $\tilde{\Theta} = \mathbb{R}_{++} \times (-\infty, -\hat{\beta}]$ is the collection of all possible values (θ_0, θ_{1a}) of the two remaining unidentified coefficients. 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 (21), the partial identification correspondence $\hat{\Sigma}_a : \tilde{\Theta} \rightarrow 2^{\tilde{\Theta}}$ becomes

$$\hat{\Sigma}_a(\theta) = \left\{ (\theta'_0, \theta'_{1a}) \in \tilde{\Theta} : \theta'_0 = -\theta'_{1a}a + \theta_0 + \theta_{1a}a \right\} \quad (27)$$

In words, $\hat{\Sigma}_a(\theta)$ is a straight line in the plane, with slope $-a$ and intercept $\theta_0 + \theta_{1a}a$ (determined by the policy a taken 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.

Given a true model $\theta = (\theta_0, \theta_{1a})$, the set $\{\hat{\Sigma}_a(\theta) : a \in A\}$ of partial identification lines is the collection of all straight lines in the plane that pass through the true model (θ_0, θ_{1a}) and have slope $-1/a$. In each such line there is a unique Lucas-Sargent model, characterized by $\theta'_{1a} = -\hat{\beta}$, as well as a unique Samuelson-Solow model, characterized by $\theta'_0 = 0$. 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 $v : C \rightarrow \mathbb{R}$ given by $v(u, \pi) = -u^2 - \pi^2$, so that the reward function $r : A \times S \rightarrow \mathbb{R}$ becomes:

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

The linear model economy and quadratic utility together form a classic linear quadratic policy framework a la Tinbergen (1952) and Theil (1961).

Lemma 5 *For every $(\theta, a) \in \tilde{\Theta} \times A$, we have $R(a, \theta) = v(\mathbb{E}_\theta(\mathbf{u}_a), \mathbb{E}_\theta(\boldsymbol{\pi}_a)) + \text{cost}$.*

The linear quadratic framework thus allows us to express the expected reward as the utility of expectations. As a result, the objective function (20) becomes

$$V(a, \mu) = \int_{\tilde{\Theta}} v(\mathbb{E}_\theta(\mathbf{u}_a), \mathbb{E}_\theta(\boldsymbol{\pi}_a)) d\mu(\theta) + \text{cost}. \quad (28)$$

As for equilibria, we begin with a piece of notation: throughout the rest of the section we fix a true model economy $\theta^* \in \tilde{\Theta}$, while 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 (24)-(26). In both equilibrium conditions the true multiplier $\xi_2^* = \hat{\beta}^* + \theta_{1\mathbf{a}}^*$ and its conjectured value $\mathbb{E}_{\mu^*}(\xi_2) = \hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})$ play a key role.³⁰

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

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

and

$$\mu^* \left(\left\{ (\theta_0, \theta_{1\mathbf{a}}) \in \tilde{\Theta} : \theta_0 = \theta_0^* - \frac{\theta_0^* \left(\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) \right)}{1 + \left(\hat{\beta}^* + \theta_{1\mathbf{a}}^* \right) \left(\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) \right)} (\theta_{1\mathbf{a}}^* - \theta_{1\mathbf{a}}) \right\} \right) = 1. \quad (30)$$

³⁰Recall that $\hat{\beta}^*$ is the beta regression coefficient of unemployment over inflation (given the true model θ^*).

The result can be heuristically derived in the special case of dogmatic beliefs, that is, when μ is concentrated on a single parameter $\bar{\theta}$. By (28), up to a constant the monetary authority's value function is

$$V(a, \mu) = - \int_{\bar{\Theta}} v(\mathbb{E}_{\theta}(\mathbf{u}_a), \mathbb{E}_{\theta}(\boldsymbol{\pi}_a)) d\mu(\theta),$$

which depends only on the expected values of inflation and unemployment under the authority's beliefs. Suppose, by way of example, that such beliefs are dogmatic: there is some model economy $\bar{\theta} = (\bar{\theta}_0, \bar{\theta}_{1\mathbf{a}}) \in \bar{\Theta}$ such that $\mu = \delta_{\bar{\theta}}$.³¹ The conjectured multiplier is therefore $\bar{\xi}_2 = \hat{\beta}^* + \bar{\theta}_{1\mathbf{a}}$. For instance, a new-classical authority that believes that there is no systematically exploitable trade-off between inflation and unemployment assumes $\bar{\theta}_{1\mathbf{a}} = -\hat{\beta}^*$ (and so the conjectured multiplier is zero). In contrast, a Keynesian authority that believes in a trade-off may assume, for instance, $\bar{\theta}_{1\mathbf{a}} = 0$ (the conjectured multiplier is then $\hat{\beta}^*$, and so strictly negative).

Based on the estimated model economy (24)-(26), a dogmatic authority conjectures that, according to the chosen policy a , the expected values of inflation and unemployment are constrained by the equation

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

This conjectured constraint is the version of the estimated aggregate supply equation (24) that the authority expects to face systematically given its dogmatic belief. So the authority's decision problem is

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

With this, the Lagrangian is

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

and so the first-order conditions are

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

By solving them we get

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

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

$$a^* = B(\bar{\theta}) \quad (\text{subjective best reply}) \quad (31)$$

³¹Recall that $\delta_{\bar{\theta}}(B) = 1$ if and only if $B \ni \bar{\theta}$ for every Borel subset B of $\bar{\Theta}$.

and

$$\bar{\theta}_0 = \theta_0^* + (\theta_{1a}^* - \bar{\theta}_{1a}) a^* \quad (\text{confirmed beliefs}). \quad (32)$$

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


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

and

$$\bar{\theta}_0 = \theta_0^* - \frac{\theta_0^* (\hat{\beta}^* + \bar{\theta}_{1a})}{1 + (\hat{\beta}^* + \theta_{1a}^*) (\hat{\beta}^* + \bar{\theta}_{1a})} (\theta_{1a}^* - \bar{\theta}_{1a}), \quad (34)$$

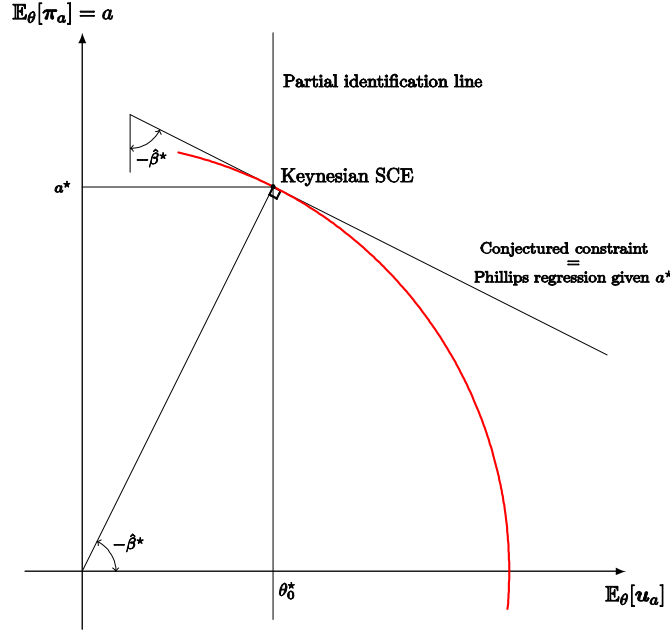
which are the equilibrium relations (29) and (30) in the case of dogmatic beliefs.³²

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 $\theta_{1a}^* = -\hat{\beta}^*$ (and so the true policy multiplier ξ_2^* is zero). Under this true model, policy a induces average unemployment $\mathbb{E}_{\theta^*}(\mathbf{u}_a) = \theta_0^*$ and average inflation $\mathbb{E}_{\theta^*}(\boldsymbol{\pi}_a) = a$. But a monetary authority with dogmatic belief $\delta_{\bar{\theta}}$ expects to observe the pair of long-run averages $(\mathbb{E}_{\bar{\theta}}(\mathbf{u}_a), a)$. This dogmatic belief is confirmed, and so condition (32) is satisfied, if $\mathbb{E}_{\bar{\theta}}(\mathbf{u}_a) = \theta_0^*$, that is, if the pair of average unemployment and average inflation lies on the vertical partial identification line with abscissa θ_0^* . The subjective best reply condition (31) is represented by the tangency between the (red) indifference curve and the (green) conjectured constraint, according to which an increase Δa in average inflation yields a $-\bar{\xi}_2 \Delta a$ decrease in average unemployment, where $\bar{\xi}_2 = \hat{\beta}^* + \bar{\theta}_{1a}$ is the conjectured multiplier.



³²Note that, with the dogmatic value $\bar{\theta}_{1a}$ of θ_{1a} in place of its expectation $\mathbb{E}_{\mu^*}(\theta_{1a})$, the dogmatic equilibrium relations are identical with the general ones. This is a consequence of the certainty equivalence principle stated in Proposition 1.

When the dogmatic belief is such that $\bar{\theta}_{1\mathbf{a}} = 0$ 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 $\mathbb{E}_{\bar{\theta}}(\mathbf{u}_a) = \bar{\theta}_0 + \hat{\beta}^* \mathbb{E}_{\bar{\theta}}(\boldsymbol{\pi}_a)$. Its slope is the beta coefficient of the Phillips regression, which represents the trade-off between inflation and unemployment that the Keynesian authority believes to be systematically exploitable.

5.3.2 Policy activism and welfare

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

To this end we need to consider the estimated policy multiplier $\xi_2 = \hat{\beta} + \theta_{1\mathbf{a}}$. The authority underestimates the multiplier when $\mathbb{E}_{\mu^*}(\xi_2) > \xi_2^*$ and overestimates it when $\mathbb{E}_{\mu^*}(\xi_2) < \xi_2^*$.³³ In structural terms, $\mathbb{E}_{\mu^*}(\xi_2) \geq \xi_2^*$ if and only if $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) \geq \theta_{1\mathbf{a}}^*$. For instance, when $\theta_{1\mathbf{a}}^*$ and $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})$ are positive this means that the multiplier is under/overestimated if and only if the direct impact of planned inflation on unemployment is over/underestimated.

The objectively optimal policy is

$$a^o = -\frac{\theta_0^* (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)}{1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)^2}. \quad (35)$$

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

³³ Both ξ_2^* and $\mathbb{E}_{\mu^*}(\xi_2)$ are negative (Assumption 2), and so $\mathbb{E}_{\mu^*}(\xi_2) \geq \xi_2^*$ if and only if $|\mathbb{E}_{\mu^*}(\xi_2)| \leq |\xi_2^*|$.

policy hyperactivism characterizes authorities that overestimate the policy multiplier, while hypoactivism characterizes authorities that underestimate it.³⁴

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

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

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

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

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

5.4 Policy dogmatism and its welfare consequences

5.4.1 Equilibria

Assume that the monetary authority has dogmatic equilibrium beliefs $\mu^* = \delta_{\bar{\theta}}$. A pair $(a^*, \delta_{\bar{\theta}}) \in A \times \Delta(\tilde{\Theta})$ is self-confirming if and only if it satisfies relations (33) and (34). Two special cases are noteworthy.

- (i) When $\bar{\theta}_{1\mathbf{a}} = -\hat{\beta}^*$ so that the conjectured policy multiplier is zero, we have the self-confirming equilibrium $a^* = 0$ and $\bar{\theta} = (\theta_0^*, -\hat{\beta}^*)$ of a new-classical authority (Proposition 11-(iv)). 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.
- (ii) When $\bar{\theta}_{1\mathbf{a}} = 0$, and so the conjectured policy multiplier $\bar{\xi}_2^* = \hat{\beta}^*$ is strictly negative, we obtain the self-confirming equilibrium

$$a^* = -\frac{\theta_0^* \hat{\beta}^*}{1 + \hat{\beta}^* (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)} \quad \text{and} \quad \bar{\theta} = \left(\theta_0^* \left(\frac{1 + \hat{\beta}^{*2}}{1 + \hat{\beta}^* (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)} \right), 0 \right) \quad (36)$$

of a Keynesian authority that believes that there is a systematically exploitable trade-off between inflation and unemployment (given by the beta coefficient $\hat{\beta}^*$ of the Phillips regression). Now a positive-target-inflation equilibrium policy results. By Proposition 11, such a policy is hyperactive if $\theta_{1\mathbf{a}}^* > 0$, hypoactive if $\theta_{1\mathbf{a}}^* < 0$, and objectively optimal if $\theta_{1\mathbf{a}}^* = 0$.

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

³⁴Being $\xi_2^* \leq 0$ (Assumption 2), the four cases it considers exhaust all possibilities. Appendix 7.5 will generalize the result to the case when Assumption 2 does not hold. Note that, being $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) \leq -\hat{\beta}^*$, in (iv) it holds $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) = -\hat{\beta}^*$ if and only if $\mu^*(\theta_{1\mathbf{a}} = -\hat{\beta}^*) = 1$, i.e., $\mu^* = \delta_{(\theta_0, -\hat{\beta}^*)}$.

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 economy, without any systematically exploitable trade-off between inflation and unemployment. The pair $(a^*, \delta_{\bar{\theta}})$ is a self-confirming equilibrium if and only if $a^* = -\theta_0^*(\hat{\beta}^* + \bar{\theta}_{1\mathbf{a}})$ and $\bar{\theta}_0 = \theta_0^*(1 - (\hat{\beta}^* + \bar{\theta}_{1\mathbf{a}})^2)$. Hence, the policy and belief pair

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

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

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

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

$$(0, \delta_{(\theta_0^*, -\hat{\beta}^*)}) \tag{37}$$

is the new-classical self-confirming equilibrium. It features a zero-target-inflation policy, which is the objectively optimal policy (and 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 an 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

$$\left(-\theta_0^*\hat{\beta}^*, \delta_{(\theta_0^*(1 - \hat{\beta}^{*2}), 0)} \right) \tag{38}$$

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

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 $\theta_{1\mathbf{a}}^* = 0$, is $\theta_0^{*2}\beta^{*2}$. It is the same loss of a Keynesian nonzero inflation policy (36) in a new-classical economy: A mistaken new-classical authority has the same lower welfare as a mistaken Keynesian one. Both mistakes result in the same welfare loss.

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. 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 $(\theta_0^{ss}, 0)$.³⁵ If $\mu_k^* \in [0, 1]$ is the belief in the latter model, we can write belief μ^* as

$$\mu^* = (1 - \mu_k^*) \delta_{(\theta_0^{ls}, -\hat{\beta}^*)} + \mu_k^* \delta_{(\theta_0^{ss}, 0)} \quad (39)$$

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

$$a^* = -\frac{\theta_0^* \hat{\beta}^* \mu_k^*}{1 + \hat{\beta}^* \mu_k^* (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)} \quad (40)$$

and

$$\theta_0^{ls} = \frac{\theta_0^*}{1 + \hat{\beta}^* (\hat{\beta}^* + \theta_{1\mathbf{a}}^*) \mu_k^*} \quad ; \quad \theta_0^{ss} = \frac{\theta_0^* (1 + \hat{\beta}^{*2} \mu_k^*)}{1 + \hat{\beta}^* (\hat{\beta}^* + \theta_{1\mathbf{a}}^*) \mu_k^*}. \quad (41)$$

As a result, in this case the pair

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

is a self-confirming equilibrium for every $\mu_k^* \in [0, 1]$. We thus have a continuum of equilibria 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 (41) 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 (39) 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 the main finding of the two-models belief case. Finally, the welfare loss is

$$\ell(a^*, \theta^*) = \frac{\theta_0^{*2} (\hat{\beta}^* \mu_k^* + \hat{\beta}^* + \theta_{1\mathbf{a}}^*)^2}{(1 + \hat{\beta}^* \mu_k^* (\hat{\beta}^* + \theta_{1\mathbf{a}}^*))^2 (1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)^2)} \quad (42)$$

³⁵Cogley and Sargent (2005) and Cogley, Colacito, and Sargent (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).

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 (40) and (41) 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


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

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

As implied by the analysis of Section 5.5.1, we have a continuum of equilibria 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 (37). Moreover, if $\mu_k^* = 1$ we get back to the dogmatic Keynesian equilibrium (38). 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^* : different such 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.

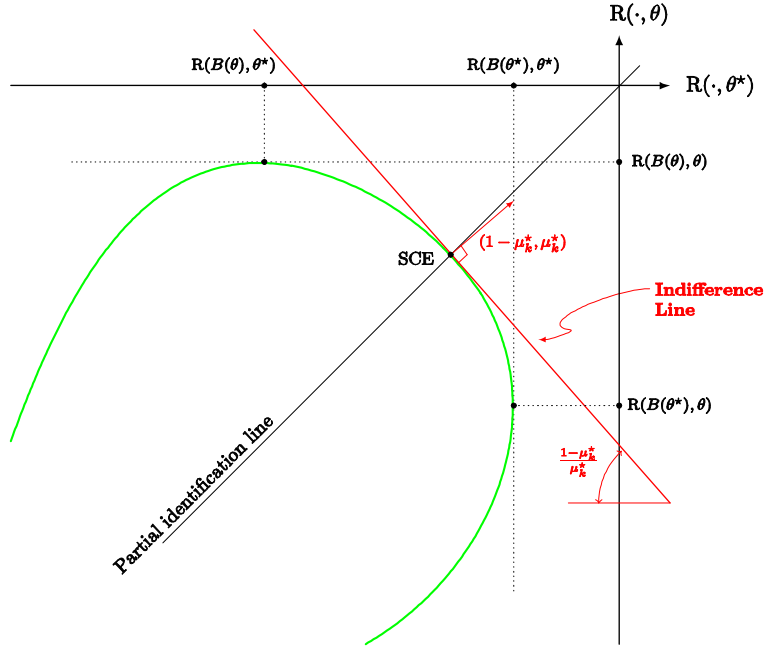
The following figures illustrate. The monetary authority is uncertain about the true economic constraint (the vertical line at the natural rate of unemployment) or the Phillips regression line. At a self-confirming equilibrium, the average unemployment expected by the monetary authority must be the natural 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 pair 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 $\text{supp } \mu^* = \{\theta^*, \theta\}$, this belief-confirmation condition must hold; therefore, the equilibrium point $(R(a^*, \theta^*), R(a^*, \theta))$ is at the intersection of the main diagonal in the $(R(\cdot, \theta^*), R(\cdot, \theta))$ -space, the “partial identification line,” with the locus of 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 $R(B(\theta^*), \theta^*) = V(B(\theta^*), \delta_{\theta^*}) > V(B(\mu^*), \mu^*)$, which illustrates Proposition 4: since the support of self-confirming belief δ_{θ^*} is included in the support of self-confirming belief μ^* , then the value of the former equilibrium is larger than the value of the latter.

6 Steady state policies

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

6.1 Setup

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

$$\varphi(x, y, \eta, \theta) = 0 \quad (43)$$

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

The endogenous variable is determined within the model, given the values that the other variables take on (outside the model); that is, the endogenous variable y solves equation (43). To ease matters, suppose that the solutions of this equation are always unique, so that we can write the solution function

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

which is called the *reduced form* of the model economy, where $\xi = g(\theta)$ is a reduced form parameter.³⁶

Example 8 A simple model economy consists of a system of linear equations $\theta_1 x + \theta_2 y = \eta$. The entries of the (square, for simplicity) matrices θ_1 and θ_2 are the parameters, and $\varphi(x, y, \theta) = \theta_1 x + \theta_2 y - \eta$. If matrix θ_2 is invertible, the reduced form is $y = \psi(x, \eta, \xi) = \xi_1 x + \xi_2 \eta$, where $\xi = g(\theta_1, \theta_2) = (-\theta_2^{-1} \theta_1, \theta_2^{-1})$. \blacktriangle

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

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

The state space is $S = Z \times H \times \Theta$, the consequence space is C , and $\rho : A \times S \rightarrow C$ is the consequence function, with $\rho(a, \zeta, \theta) = \psi(a, \zeta, g(\theta))$.³⁷ Policy multipliers correspond to the derivative of ψ with respect to a (partial derivatives if a is multidimensional).

Example 9 In the Phillips curve example the structural form is (12)-(13), with $c = (u, \pi)$, $\eta = (\varepsilon, w)$, and $\theta = (\theta_0, \theta_{1\pi}, \theta_{1a}, \theta_2, \theta_3)$. The function φ is linear, given by

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

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

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

Finally, $g(\theta) = (\theta_0, \theta_{1\pi} + \theta_{1a}, \theta_{1\pi} \theta_3, \theta_2, \theta_3)$, ρ is given by (16), and ξ_2 is the policy multiplier.³⁸ \blacktriangle

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

³⁶We do not discuss here the conditions that, along with the uniqueness of solutions, ensure the existence of the reduced form representation (see, e.g., Rothenberg, 1971). The economic relevance of multiple solutions is discussed by Jovanovic (1989); in this case, we have a solution correspondence, and so a reduce form correspondence.

³⁷Here we tacitly assume that ψ and g are such that ρ satisfies Savage's Consequentialism.

³⁸Assumption 1 of the example here amounts to requiring that the function $g : \Theta \rightarrow \Xi$ be one-to-one, so that distinct structural parameters correspond to distinct reduced form parameters.

6.2 Economic and stochastic model uncertainties

In the Phillips curve example we considered a state $s = (w, \varepsilon, \theta)$ in which we distinguish a shock pair (w, ε) and a model economy θ . The key factorization (19) builds on that distinction. This distinction also applies to the general policy problem, which features two types of model uncertainties. First, there is uncertainty about the economics of the phenomenon under consideration, and so about the economic models that explain it. Second, there is uncertainty about the statistical performance of such economic models, due to the errors that affect measurements and shocks; the latter represent the unexplained variation caused by, possibly many, minor explanatory variables that the policy maker is “unable and unwilling to specify”, as Marschak (1953) p. 12 remarks.³⁹

We can call, respectively, economic model uncertainty and stochastic model uncertainty the two types of uncertainty. The former is more fundamental than the latter since it reflects the economic views of policy makers.⁴⁰ This is why, by assuming q known, in the Phillips curve example we focused on economic model uncertainty.

In general, the state spaces relevant for policy problems can be represented as Cartesian products of factors that represent the two types of uncertainty. In this case, versions of the factorization (19) would apply with a twist: uncertainty about the probability model that describes the exogenous variable, absent in the Phillips curve example, must also be considered. We thus factor the probability models $\sigma \in \Sigma \subseteq \Delta(Z \times H \times \Theta)$ as $\sigma = q \times \delta_{\bar{\theta}}$, where $-$ to focus on economic model uncertainty – we assume that the joint probability model $q \in \Delta(Z \times H)$ for the exogenous and shock variables is known and that $\delta_{\bar{\theta}} \in \Delta(\Theta)$ is a Dirac probability measure concentrated on a given economic model $\bar{\theta} \in \Theta$. As in the example, we can identify Σ with Θ without loss of generality. In particular, $R(a, \theta) = \int r(a, \zeta, \eta, \theta) dq(\zeta, \eta)$.

6.3 Partial identification and equilibrium

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

We assume that the messages that the policy maker receives are the values of the endogenous variable c , that is, $f = \rho$.⁴¹ Since we identify Σ with Θ , the distribution map and the partial identification correspondence are given here by $\hat{f}_a(\theta)(E) = q(\rho_{a,\theta}^{-1}(E))$ and $\hat{\Sigma}_a(\theta) = \{\theta' \in \Sigma : \hat{f}_a(\theta') = \hat{f}_a(\theta)\}$.⁴² In particular, $\hat{\Sigma}_a(\theta)$ is the collection of all model economies that are observationally equivalent when θ is the true model and a is the chosen policy.

Moment conditions based on lung-run observed values of the endogenous variable may

³⁹In a similar vein, Koopmans (1947) p. 169 writes that “... stochastic ... in that the behavior of any group of individuals, and the outcome of any production process, is determined in part by many minor factors, further scrutiny of which is either impossible or unrewarding.”

⁴⁰A main instance of what Denzau and North (1994) and, more recently, Rodrik (2014) call “ideas”. Note that economic model uncertainty subsumes the parametric uncertainty considered by Friedman (1953) and then Brainard (1967) in policy problems.

⁴¹Inter alia, this assumption implies that own-action independence of feedback cannot hold in nontrivial policy problems (recall our discussion at the end of Section 3.2).

⁴² $\rho_{a,\theta} : Z \times H \rightarrow C$ is the section of ρ at (a, θ) and $E \subseteq C$ is any measurable subset.

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

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

6.4 Steady state policy rules

A natural question is whether a notion of a steady state policy rule emerges in our setup. When Σ is a singleton, that is, when structural information allows the policy maker to know the true model, such a rule is given by the best reply correspondence $B(\theta) = \arg \max_{a \in A} R(a, \theta)$ for each $\theta \in \tilde{\Theta}$. A policy maker who knows the true model economy best replies to such knowledge by taking the objectively optimal policy. Since q is known, this corresponds in our steady state setup to the case originally studied by Tinbergen (1952) and Theil (1961).⁴³

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

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

$$\mu^* \left(\hat{\Sigma}_{B(\mu^*)}(\theta^*) \right) = 1. \quad (44)$$

Hence, there is a unique equilibrium condition (44), which is cast in terms of beliefs (actions being pinned down by the best reply function). For this reason in nice policy problems we can view self-confirming equilibria as fixed points in the space of beliefs. To this end, define the correspondence $T_{\theta^*} : \Delta(\Sigma) \rightarrow 2^{\Delta(\Sigma)}$ by

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

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

Definition 2 *Given a true model $\theta^* \in \Sigma$, a belief $\mu \in \Delta(\Sigma)$ is self-confirming if $\mu \in T_{\theta^*}(\mu)$.*

⁴³In the terminology of Brainard (1967), a singleton Σ means a nonrandom parameter θ (see also Blinder, 1998, p. 11). A nonrandom θ is assumed, for example, also by Poole (1970) in his classic IS-LM policy analysis under uncertainty (see, e.g., Poole, 1970, p. 215).

In words, self-confirming beliefs are consistent with the long-run data – that is, the observed values of the endogenous variable – gathered through the optimal policy that they induce.⁴⁴

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

6.5 Brainard conservatism?

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

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

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

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

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

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

$$h_{\theta^*}(\mu) = -\frac{\mathbb{E}_\mu(\theta_0) (\hat{\beta}^* + \mathbb{E}_\mu(\theta_{1\mathbf{a}}))}{1 + (\hat{\beta}^* + \mathbb{E}_\mu(\theta_{1\mathbf{a}}))^2}. \quad (46)$$

If μ is not dogmatic, then

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

Our steady state policies thus do not feature the “Brainard conservatism principle,” that is, they are not more prudent than what the certainty equivalent principle would prescribe.⁴⁵ The partial identification caused by the authority’s limited structural information does not

⁴⁴This fixed point characterization is easily seen to hold in full generality in our setup (beyond the policy case considered here) when best replies are unique.

⁴⁵Blinder (1997) names this principle after the classic finding of Brainard (1967) that, even in linear quadratic problems, parametric uncertainty may make policies more prudent in this sense. See also Soderstrom (2002) and Bernanke (2007) for a discussion of the principle.

translate into a policy conservatism a la Brainard: but the long-run data confirmation condition that characterizes our policy rules puts the certainty equivalent principle back in business.

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

7 Appendix

7.1 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.⁴⁶

The next result completes Lemma 1.

Lemma 6 *Let X and Y be Borel spaces and $\varphi : X \rightarrow Y$ be measurable. Then:*

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

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

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

⁴⁶A terminological caveat: Berberian (1997) and other authors use Borel space as a synonymous of measurable space; they always specify the adjective standard when they assume Polish metrizable (as we do here).

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.⁴⁷

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 $\varphi' : \varphi(X) \rightarrow X$ 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 $\xi(\varphi^{-1}(C)) = \xi'(\varphi^{-1}(C))$ 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 $\hat{\varphi}(\delta_x)(C) = \delta_x(\varphi^{-1}(C)) = \delta_{\varphi(x)}(C)$ 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 $(\Delta(X), \mathcal{B}_{\Delta(X)})$ and $(\hat{\varphi}(\Delta(X)), \mathcal{B}_{\Delta(Y)} \cap \hat{\varphi}(\Delta(X)))$. 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 $\nu(\varphi(X)) = \xi(\varphi^{-1}(\varphi(X))) = \xi(X) = 1$. 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,

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

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 $(\hat{\varphi}(\Delta(X)), \mathcal{B}_{\Delta(Y)} \cap \hat{\varphi}(\Delta(X)))$ and $(\Delta(\varphi(X)), \mathcal{B}_{\Delta(\varphi(X))})$. 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 $\xi : \mathcal{B}_X \rightarrow [0, 1]$ is well defined because $\varphi : (X, \mathcal{B}_X) \rightarrow (\varphi(X), \mathcal{B}_Y \cap \varphi(X))$ is a measurable isomorphism. Moreover, for every $B \in \mathcal{B}_X$, $\varphi(B) = (\varphi')'(B) = (\varphi')^{-1}(B)$, thus $\xi(B) = \lambda(\varphi(B)) = \lambda((\varphi')^{-1}(B))$ is a probability measure on X . Finally, for every $D \in \mathcal{B}_Y \cap \varphi(X)$,

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

⁴⁷Notice 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 $\Delta(Y)$.

as wanted. ■

Corollary 4 *Let $(\Theta, \mathcal{B}_\Theta)$ and (S, \mathcal{B}_S) be Borel spaces and fix $p : \Theta \rightarrow \Delta(S)$. Then p is measurable if and only if*

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

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

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

Proof Since $\mathcal{B}_{\Delta(S)}$ is the sigma algebra generated by the functions $\phi_B : \Delta(X) \rightarrow \mathbb{R}$ defined by $\phi_B(\xi) = \xi(B)$ for all $B \in \mathcal{B}_S$, a map $p : \Theta \rightarrow \Delta(S)$ is measurable if and only if $\phi_B \circ p$ is measurable for all $B \in \mathcal{B}_S$ (see, e.g., Berberian, 1997, Proposition 1.3.8). But, given any $B \in \mathcal{B}_S$, $p(B | \theta) = p_\theta(B) = (\phi_B \circ p)(\theta)$ for all $\theta \in \Theta$, thus $p(B | \cdot) = \phi_B \circ p$, proving the first part of the statement.⁴⁸ The rest follows from the statement of Lemma 6 setting $X = \Theta$, $Y = \Delta(S)$, and $\varphi = p$, with the exception of the explicit expression $\tilde{p}^{-1}(\lambda) = \lambda \circ p$, for which the last paragraph of the proof of Lemma 6 has to be inspected. ■

Corollary 5 *Let $(\Theta, \mathcal{B}_\Theta)$ and (T, \mathcal{B}_T) be Borel spaces and fix $q \in \Delta(T)$. Then*

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

is measurable and one-to-one.

Proof Injectivity is obvious, we only have to show that

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

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 $\theta \in \Theta$,

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

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

Corollary 6 *Let (S, \mathcal{B}_S) be a Borel space and $\delta : S \rightarrow \Delta(S)$ the embedding $s \mapsto \delta_s$. Then: $\{\delta_s\}_{s \in S} \in \mathcal{B}_{\Delta(S)}$, $\delta : S \rightarrow \{\delta_s\}_{s \in S}$ is a measurable isomorphism, and $\lambda \mapsto \lambda \circ \delta$ is a measurable isomorphism between $\Delta(\{\delta_s\}_{s \in S})$ and $\Delta(S)$.*

Proof In order to apply the previous Corollary 4 with $\Theta = S$ and $p = \delta$, we only have 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. ■

⁴⁸Notice 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 $\Delta(S)$.

7.2 On the long-run interpretation

As we remarked in (2), if the decision maker chooses action a and observes a message in $B \in \mathcal{M}$, he can infer that the (unobservable) realized state s^* belongs to $f_a^{-1}(B) \in \mathcal{F}_a$. Assume now the decision maker chooses action a in n identical copies of the decision problem. If the (unobservable) realized states are $(s_1^*, s_2^*, \dots, s_n^*)$, he observes the messages $(f_a(s_1^*), f_a(s_2^*), \dots, f_a(s_n^*))$ and he can compute the empirical frequency

$$\frac{1}{n} \sum_{t=1}^n \delta_{f_a(s_t^*)}(B) \quad (47)$$

of every $B \in \mathcal{M}$, which corresponds to the empirical frequency of $f_a^{-1}(B) \in \mathcal{F}_a \subseteq \mathcal{S}$.⁴⁹

Next we address the question of the convergence of the sequence $(1/n) \sum_{t=1}^n \delta_{f_a(s_t^*)} \in \Delta(M)$ of empirical frequencies. In this perspective, it is convenient to consider:

- on M any metric d such that (M, d) is separable and \mathcal{M} is its Borel sigma algebra (for example, if M is a Borel subset of a Polish space (P, d_P) , then the restriction d_M of d_P to M has these properties, even if (M, d_M) is not necessarily complete);
- on $\Delta(M)$ the weak convergence of probability measures, according to which a sequence $(\mu_n)_{n \geq 1}$ converges to μ if and only if $\int_M \psi d\mu_n \rightarrow \int_M \psi d\mu$ for every continuous and bounded function $\psi : M \rightarrow \mathbb{R}$.

For each model $\sigma \in \Delta(S)$, denote by $(Z, \mathcal{Z}, \zeta_\sigma)$ the “i.i.d.” product space $(S^\infty, \mathcal{S}^\infty, \sigma^\infty)$ with generic element $z = (s_t)_{t \geq 1}$. The coordinate random variables $\mathbf{z}_t : Z \rightarrow S$, given by $\mathbf{z}_t(z) = s_t$, are independent with law $\zeta_\sigma \circ \mathbf{z}_t^{-1} = \sigma$ for all $t \geq 1$. By setting $\mathbf{x}_t = f_a \circ \mathbf{z}_t : Z \rightarrow M$ for all $t \geq 1$, (47) becomes

$$\mu_n(B)(z^*) \equiv \frac{1}{n} \sum_{t=1}^n \delta_{\mathbf{x}_t(z^*)}(B) \quad \forall B \in \mathcal{M}. \quad (48)$$

The random variables $\mathbf{x}_1, \mathbf{x}_2, \dots$ take values in M and are independent, with law $\mu_\sigma \equiv \zeta_\sigma \circ (f_a \circ \mathbf{z}_t)^{-1} = \sigma \circ f_a^{-1}$ for all $t \geq 1$. By a version of the Glivenko-Cantelli Theorem due to Varadarajan (Dudley, 2002, Theorem 11.4.1) applied to the sequence⁵⁰

$$\begin{aligned} \mu_n : Z &\rightarrow \Delta(M) \\ z &\mapsto \mu_n(z) \end{aligned}$$

we have that $\zeta_\sigma(\{z \in Z : \mu_n(z) \rightarrow \mu_\sigma\}) = 1$. That is,

$$\zeta_\sigma \left(\left\{ (s_t)_{t \geq 1} \in S^\infty : \frac{1}{n} \sum_{t=1}^n \delta_{f_a(s_t)} \rightarrow \sigma \circ f_a^{-1} \right\} \right) = 1. \quad (49)$$

⁴⁹For every $t = 1, \dots, n$, we have $\delta_{f_a(s_t^*)}(B) = 1 \Leftrightarrow f_a(s_t^*) \in B \Leftrightarrow s_t^* \in f_a^{-1}(B) \Leftrightarrow \delta_{s_t^*}(f_a^{-1}(B)) = 1$, i.e., $\delta_{f_a(s_t^*)}(B) = \delta_{s_t^*}(f_a^{-1}(B))$.

⁵⁰Note that, for all $n \geq 1$, and for all $B \in \mathcal{M}$, $\mu_n(B | z) = (1/n) \sum_{t=1}^n \delta_{\mathbf{x}_t(z)}(B) = (1/n) \sum_{t=1}^n 1_B(\mathbf{x}_t(z))$ defines a measurable function from Z to \mathbb{R} ; Corollary 4 implies that $\mu_n : Z \rightarrow \Delta(M)$ is measurable too.

If $\sigma^* \in \Delta(S)$ is the true model, a decision maker that has been choosing action a “many times” is able to estimate the messages’ distribution $\lim_{n \rightarrow \infty} (1/n) \sum_{t=1}^n \delta_{f_a(s_t^*)}$ so that he ζ_{σ^*} -almost surely identifies $\sigma^* \circ f_a^{-1}$.⁵¹

In the paper we assume this is the case, i.e., $\lim_{n \rightarrow \infty} (1/n) \sum_{t=1}^n \delta_{f_a(s_t^*)} = \sigma^* \circ f_a^{-1}$. The decision maker, conditional on choosing action a , can thus infer the true distribution $\sigma^* \circ f_a^{-1}$ of f_a , as well as that the true model belongs to the set

$$\hat{\Sigma}_a(\sigma^*) = \{\sigma \in \Sigma : \sigma \circ f_a^{-1} = \sigma^* \circ f_a^{-1}\}$$

under the structural assumption $\sigma^* \in \Sigma$.

7.3 Feedback and identification

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

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

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

Lemma 7 *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 coarse 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 $g_a : M \rightarrow C$ such that $\rho_a(s) = g_a(f_a(s))$ for all $s \in S$. (ii) For each $a \in A$, $f_a : S \rightarrow M$ is Borel measurable and one-to-one, by Lemma 6, $f_a(S)$ is a Borel subset of M and $f_a : S \rightarrow f_a(S)$ is a Borel isomorphism. Then the inverse function $f_a^{-1} : f_a(S) \rightarrow S$ is Borel measurable.⁵² Arbitrarily choose $\bar{s} \in S$ and set

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

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

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

⁵¹Specifically, he correctly identifies $\sigma^* \circ f_a^{-1}$ if and only if $(1/n) \sum_{t=1}^n \delta_{f_a(s_t^*)} \rightarrow \sigma^* \circ f_a^{-1}$ which happens with ζ_{σ^*} probability 1.

⁵²*Caveat:* In the proof of Lemma 6, the inverse isomorphism $f_a^{-1} : f_a(S) \rightarrow S$ is denoted f'_a , here f'_a is a section of the feedback function f' which in no way is an inverse of f .

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

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

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

7.4 Proofs

Proof of Lemma 1 See Lemma 6. ■

Proof of Lemma 2 See Lemma 7. ■

Proof of Lemma 3 Suppose $\mu \in \Delta(\hat{\Sigma}_a(\sigma^*))$, i.e., $\text{supp } \mu \subseteq \hat{\Sigma}_a(\sigma^*)$. Let $\sigma \in \Sigma$. If $\mu(\sigma) = 0$, then $\mu(\sigma | a, m) = \mu(\sigma)$ for all $m \in \text{supp } \hat{f}_a(\sigma_\mu)$. If $\mu(\sigma) > 0$, then $\sigma \in \hat{\Sigma}_a(\sigma^*)$ and so $\hat{f}_a(\sigma)(m) = \hat{f}_a(\sigma^*)(m) = \hat{f}_a(\sigma_\mu)(m)$ for all $m \in M$. Therefore

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

Conversely, assume $\mu(\cdot | a, m) = \mu(\cdot)$ for all $m \in \text{supp } \hat{f}_a(\sigma_\mu)$. Let $\sigma \in \text{supp } \mu$. If $\hat{f}_a(\sigma_\mu)(m) > 0$, then

$$\mu(\sigma) \frac{\hat{f}_a(\sigma)(m)}{\hat{f}_a(\sigma_\mu)(m)} = \mu(\sigma)$$

and so $\hat{f}_a(\sigma)(m) = \hat{f}_a(\sigma_\mu)(m)$. Since $\text{supp } \mu \cap \hat{\Sigma}_a(\sigma^*) \neq \emptyset$, we have $\hat{f}_a(\sigma)(m) = \hat{f}_a(\sigma^*)(m)$. If $\hat{f}_a(\sigma_\mu)(m) = 0$, we have $\hat{f}_a(\sigma)(m) = 0$ as well as $\hat{f}_a(\sigma^*)(m) = 0$ because $\text{supp } \mu \cap \hat{\Sigma}_a(\sigma^*) \neq \emptyset$. Summing up, if $\sigma \in \text{supp } \mu$ then $\hat{f}_a(\sigma)(m) = \hat{f}_a(\sigma^*)(m)$ for all $m \in M$. We conclude that $\mu \in \Delta(\hat{\Sigma}_a(\sigma))$. ■

Proof of Lemma 4 Fix $a \in A$. Observability of consequences implies that $\rho_a(s) = g_a(f_a(s))$ for each $s \in S$, where $g_a : M \rightarrow C$ is $\mathcal{B}_M - \mathcal{B}_C$ -measurable; as $f_a : S \rightarrow M$ is $\mathcal{F}_a - \mathcal{B}_M$ -measurable, then $\rho_a : S \rightarrow C$ is $\mathcal{F}_a - \mathcal{B}_C$ -measurable. Moreover, $v : C \rightarrow \mathbb{R}$ is $\mathcal{B}_C - \mathcal{B}_{\mathbb{R}}$ -measurable and bounded above, and so $r_a = v \circ \rho_a : S \rightarrow \mathbb{R}$ is $\mathcal{F}_a - \mathcal{B}_{\mathbb{R}}$ -measurable and bounded above. Thus

$$R_a(\sigma) = \int_S r_a d\sigma = \int_S r_a d\sigma|_{\mathcal{F}_a} \quad \forall \sigma \in \Delta(S). \quad (50)$$

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

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

because, by Lemma 4, $R(a^*, \sigma) = R(a^*, \sigma^*)$ for all $\sigma \in \hat{\Sigma}_{a^*}(\sigma^*)$. \blacksquare

Proof of Proposition 4 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 (4) for a^* and Proposition 3 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 4, $R(b^*, \sigma) = R(b^*, \sigma^*)$ 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 3. \blacksquare

Proof of Proposition 5 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 $\mu^* \left(\hat{\Sigma}_{a^*}(\sigma^*) \right) = 1$.

\blacksquare

Proof of Proposition 6 The optimality condition (4) for a^* and Proposition 3 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 4, $R(b^*, \sigma) = R(b^*, \sigma^*)$ 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 3.

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 $\nu^* \left(\hat{\Sigma}_{b^*}(\sigma^*) \right) = 1$.

\blacksquare

Proof of Proposition 7 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 4, 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. \blacksquare

Proof of Corollary 2 Given a true model $\sigma^* \in \Sigma$, the result follows from Proposition 7 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 3 Given a true model $\sigma^* \in \Sigma$, the result follows from Proposition 7 since own-action independence of feedback implies $\hat{\Sigma}_a(\sigma^*) = \hat{\Sigma}_{a'}(\sigma^*)$ for every $a, a' \in A$. Hence, $\gamma(\sigma^*) = \arg \max_{a \in A} R(a, \sigma^*)$. ■

Proof of Proposition 9 First recall that a is fixed. Let F_q denote the bivariate cumulative distribution function of q , that is

$$F_q(w, \varepsilon) = q((-\infty, w] \times (-\infty, \varepsilon]) \quad \forall (w, \varepsilon) \in W \times E.$$

Some simple algebra shows that, for each $\theta \in \Theta$, the set $\hat{\Sigma}_a(\theta)$ is

$$\left\{ \theta' \in \Theta : F_q \left(\frac{u - \theta_0 - \theta_{1\mathbf{a}}a - \theta_{1\pi}\pi}{\theta_2}, \frac{\pi - a}{\theta_3} \right) = F_q \left(\frac{u - \theta'_0 - \theta'_{1\mathbf{a}}a - \theta'_{1\pi}\pi}{\theta'_2}, \frac{\pi - a}{\theta'_3} \right) \quad \forall (u, \pi) \in C \right\}$$

We first prove the inclusion \subseteq . Consider $\theta' \in \hat{\Sigma}_a(\theta)$. Recall that $\theta' \in \hat{\Sigma}_a(\theta)$ if and only if

$$\hat{\rho}_a(q \times \delta_{\theta'}) = \hat{\rho}_a(q \times \delta_{\theta}).$$

In particular, we have

$$\int_S h(\rho_a) d(q \times \delta_{\theta'}) = \int_S h(\rho_a) d(q \times \delta_{\theta}) \quad (51)$$

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

1. For $h(u, \pi) = \pi$ and $\theta'' \in \Theta$, we have that $\int_S \pi d(q \times \delta_{\theta''}) = a$.
2. For $h(u, \pi) = \pi^2$ and $\theta'' \in \Theta$, we have that $\int_S \pi^2 d(q \times \delta_{\theta''}) = a^2 + (\theta''_3)^2$.
3. For $h(u, \pi) = u$ and $\theta'' \in \Theta$, we have that $\int_S u d(q \times \delta_{\theta''}) = \theta''_0 + (\theta''_{1\pi} + \theta''_{1\mathbf{a}})a$.
4. For $h(u, \pi) = u^2$ and $\theta'' \in \Theta$, we have that $\int_S u^2 d(q \times \delta_{\theta''}) = (\theta''_0)^2 + (\theta''_{1\pi} + \theta''_{1\mathbf{a}})^2 a^2 + (\theta''_{1\pi})^2 (\theta''_3)^2 + (\theta''_2)^2 + 2\theta''_0 (\theta''_{1\pi} + \theta''_{1\mathbf{a}})a$.
5. For $h(u, \pi) = u\pi$ and $\theta'' \in \Theta$, we have that $\int_S u\pi d(q \times \delta_{\theta''}) = a(\theta''_0 + (\theta''_{1\pi} + \theta''_{1\mathbf{a}})a) + \theta''_{1\pi} (\theta''_3)^2$.

Given (51), note that point 2 gives $\theta'_3 = \theta_3$, then points 3 and 5 give $\theta'_{1\pi} = \theta_{1\pi}$, then point 4 gives $\theta'_2 = \theta_2$, finally point 3 again yields $\theta'_0 + \theta'_{1\mathbf{a}}a = \theta_0 + \theta_{1\mathbf{a}}a$. This concludes the proof of the first set inclusion and formalizes the moments heuristics described in the main text.

For the opposite inclusion, consider $\theta' \in \Theta$ such that $\theta'_0 + \theta'_{1\mathbf{a}}a = \theta_0 + \theta_{1\mathbf{a}}a$, $\theta'_{1\pi} = \theta_{1\pi}$, $\theta'_2 = \theta_2$, $\theta'_3 = \theta_3$. We have that

$$\left(\frac{u - \theta_0 - \theta_{1\mathbf{a}}a - \theta_{1\pi}\pi}{\theta_2}, \frac{\pi - a}{\theta_3} \right) = \left(\frac{u - \theta'_0 - \theta'_{1\mathbf{a}}a - \theta'_{1\pi}\pi}{\theta'_2}, \frac{\pi - a}{\theta'_3} \right) \quad \forall (u, \pi) \in C,$$

which implies

$$F_q \left(\frac{u - \theta_0 - \theta_{1\mathbf{a}}a - \theta_{1\boldsymbol{\pi}}\pi}{\theta_2}, \frac{\pi - a}{\theta_3} \right) = F_q \left(\frac{u - \theta'_0 - \theta'_{1\mathbf{a}}a - \theta'_{1\boldsymbol{\pi}}\pi}{\theta'_2}, \frac{\pi - a}{\theta'_3} \right) \quad \forall (u, \pi) \in C.$$

Hence, $\theta' \in \hat{\Sigma}_a(\theta)$. This proves the statement. \blacksquare

Proof of Lemma 5 Some simple algebra shows that

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

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

Proof of Proposition 10 It holds

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

and so $V(a, \mu^*)$ is, up to a constant, equal to

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

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

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

Putting $a = a^*$ we get

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

and so

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

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

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

as desired. ■

Proof of Proposition 11 It holds

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

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

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

and so

$$a^* \geq a^o \iff a^* \frac{\theta_{1\mathbf{a}}^* - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})}{\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})} \leq 0 \tag{52}$$

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 (29) $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) \neq -\hat{\beta}^*$ and so $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) < -\hat{\beta}^*$. By (52), $(\theta_{1\mathbf{a}}^* - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) / (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) < 0$, which in turn implies $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) < \theta_{1\mathbf{a}}^*$. Conversely, suppose $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) < \theta_{1\mathbf{a}}^*$. Since $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) \leq -\hat{\beta}^*$, by (29) it follows $a^* > 0$. Moreover, being $(\theta_{1\mathbf{a}}^* - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) / (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) < 0$, by (52) it holds $a^* > a^o$. (iii) Suppose $0 < a^* < a^o$. By (52), $(\theta_{1\mathbf{a}}^* - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) / (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) > 0$, that is, $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) \in (\theta_{1\mathbf{a}}^*, -\hat{\beta}^*)$. Conversely, suppose $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) \in (\theta_{1\mathbf{a}}^*, -\hat{\beta}^*)$. By (29), $a^* > 0$. Moreover, being $(\theta_{1\mathbf{a}}^* - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) / (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) > 0$, by (52) it holds $a^* < a^o$. \blacksquare

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

$$\ell(\mu^*, \theta^*) = \frac{\theta_0^{*2} (\theta_{1\mathbf{a}}^* - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}))^2}{\left(1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*)^2\right) \left(1 + (\hat{\beta}^* + \theta_{1\mathbf{a}}^*) (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}))\right)^2} \quad (53)$$

There is a zero welfare loss if and only if $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) = \theta_{1\mathbf{a}}^*$, that is, if and only if the monetary authority's expected value of the coefficient $\theta_{1\mathbf{a}}$ is correct. Otherwise, the loss is nonzero, as (53) shows.

Proof of Proposition 12 and eq. (53) First note that

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

and

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

Hence,

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

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

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

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

$$\ell(a^*, \theta^*) = \theta_0^{*2} (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}))^2 = \theta_0^{*2} (\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) - \theta_{1\mathbf{a}}^*)^2 \quad (54)$$

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

$$\ell(a^*, \theta^*) = 0 \quad (55)$$

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

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

In the next section we show that

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

Hence,

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

It is easy to check that, along with (54) and (55), this completes the proof. ■

Proof of eq. (42) It holds

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

as desired. ■

Proof of Proposition 13 Given any $\mu \in \Delta(\Sigma)$, the best reply is easily seen to take the form

$$B(\mu) = -\frac{\mathbb{E}_\mu(\theta_0 (\hat{\beta}^* + \theta_{1\mathbf{a}}))}{1 + \mathbb{E}_\mu(\hat{\beta}^* + \theta_{1\mathbf{a}})^2} = -\frac{\mathbb{E}_\mu(\theta_0) (\hat{\beta}^* + \mathbb{E}_\mu(\theta_{1\mathbf{a}})) + \text{Cov}_\mu(\theta_0, \theta_{1\mathbf{a}})}{1 + (\hat{\beta}^* + \mathbb{E}_\mu(\theta_{1\mathbf{a}}))^2 + \text{Var}_\mu(\theta_{1\mathbf{a}})}$$

If μ is dogmatic, then (46) holds. Suppose μ is non-dogmatic, so that both $\text{Var}_\mu(\theta_0)$ and $\text{Var}_\mu(\theta_{1\mathbf{a}})$ are not zero. Set $\kappa = \hat{\beta}^* + \mathbb{E}_\mu(\theta_{1\mathbf{a}})$ and $C_e(\mu) = -\mathbb{E}_\mu(\theta_0) \kappa / (1 + \kappa^2)$. The following holds:

$$\begin{aligned}
B(\mu) \leq C_e(\mu) &\iff \frac{\mathbb{E}_\mu(\theta_0) \kappa + \text{Cov}_\mu(\theta_0, \theta_{1\mathbf{a}})}{1 + \kappa^2 + \text{Var}_\mu(\theta_{1\mathbf{a}})} \geq \frac{\mathbb{E}_\mu(\theta_0) \kappa}{1 + \kappa^2} \\
&\iff \mathbb{E}_\mu(\theta_0) \kappa + \text{Cov}_\mu(\theta_0, \theta_{1\mathbf{a}}) + (\mathbb{E}_\mu(\theta_0) \kappa + \text{Cov}_\mu(\theta_0, \theta_{1\mathbf{a}})) \kappa^2 \\
&\geq \mathbb{E}_\mu(\theta_0) \kappa + \mathbb{E}_\mu(\theta_0) \kappa^3 + \mathbb{E}_\mu(\theta_0) \kappa \text{Var}_\mu(\theta_{1\mathbf{a}}) \\
&\iff \text{Cov}_\mu(\theta_0, \theta_{1\mathbf{a}}) + \mathbb{E}_\mu(\theta_0) \kappa^3 + \text{Cov}_\mu(\theta_0, \theta_{1\mathbf{a}}) \kappa^2 \geq \mathbb{E}_\mu(\theta_0) \kappa \kappa^2 + \mathbb{E}_\mu(\theta_0) \kappa \text{Var}_\mu(\theta_{1\mathbf{a}}) \\
&\iff \text{Cov}_\mu(\theta_0, \theta_{1\mathbf{a}}) (1 + \kappa^2) \geq \mathbb{E}_\mu(\theta_0) \kappa \text{Var}_\mu(\theta_{1\mathbf{a}}) \iff \frac{\mathbb{E}_\mu(\theta_0) \kappa}{1 + \kappa^2} \leq \frac{\text{Cov}_\mu(\theta_0, \theta_{1\mathbf{a}})}{\text{Var}_\mu(\theta_{1\mathbf{a}})} \\
&\iff C_e(\mu) \geq -\frac{\text{Cov}_\mu(\theta_0, \theta_{1\mathbf{a}})}{\text{Var}_\mu(\theta_{1\mathbf{a}})}
\end{aligned}$$

If $\mu \in M(\theta^*)$, then

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

We conclude that $h_{\theta^*}(\mu) = B(\mu) = C_e(\mu) = -\text{Cov}_\mu(\theta_0, \theta_{1\mathbf{a}}) / \text{Var}_\mu(\theta_{1\mathbf{a}})$ if $\mu \in M(\theta^*)$. Hence, $\text{Cov}_\mu(\theta_0, \theta_{1\mathbf{a}}) = 0$ implies $h_{\theta^*}(\mu) = 0$. If $\text{Cov}_\mu(\theta_0, \theta_{1\mathbf{a}}) \neq 0$, we can write

$$h_{\theta^*}(\mu) = -\frac{\text{Cov}_\mu(\theta_0, \theta_{1\mathbf{a}})}{\text{Var}_\mu(\theta_{1\mathbf{a}})} = -\text{Corr}_\mu(\theta_0, \theta_{1\mathbf{a}}) \sqrt{\frac{\text{Var}_\mu(\theta_0)}{\text{Var}_\mu(\theta_{1\mathbf{a}})}} = \sqrt{\frac{\text{Var}_\mu(\theta_0)}{\text{Var}_\mu(\theta_{1\mathbf{a}})}}$$

where the last equality holds because $\mu \in M(\theta^*)$ implies $\text{Corr}_\mu(\theta_0, \theta_{1\mathbf{a}}) = -1$ (recall Assumption 2). \blacksquare

7.5 Beyond Assumption 2

The only result of Section 5 that depends on Assumption 2 is Proposition 11. Without such assumption, the result takes the following more general form.

Proposition 14 *Consider a self-confirming equilibrium (a^*, μ^*) .*

- (i) $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) \notin [\theta_{1\mathbf{a}}^* \wedge (-\hat{\beta}^*), \theta_{1\mathbf{a}}^* \vee (-\hat{\beta}^*)]$ if and only if policy activism results in hyperactivism (i.e., $a^* > 0$ implies $a^* > a^o$);
- (ii) $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) \in (\theta_{1\mathbf{a}}^* \wedge (-\hat{\beta}^*), \theta_{1\mathbf{a}}^* \vee (-\hat{\beta}^*))$ if and only if policy activism results in hypoactivism (i.e., $a^* > 0$ implies $a^* < a^o$).

The proof relies on a lemma.

Lemma 8 *Consider a self-confirming equilibrium (a^*, μ^*) . The following conditions are equivalent:*

1. $a^* > 0$ implies $a^* > a^o$ (resp., $a^* < a^o$);
2. $(\theta_{1\mathbf{a}}^* - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) / (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) < 0$ (resp., > 0);
3. $a^* < 0$ implies $a^* < a^o$ (resp., $a^* > a^o$).

Proof (1.) implies (2.): Suppose that $a^* > 0$ implies $a^* > a^o$. By (52), $a^* > 0$ implies

$$a^* \frac{\theta_{1\mathbf{a}}^* - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})}{\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})} < 0$$

and so $(\theta_{1\mathbf{a}}^* - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) / (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) < 0$. (2.) implies (1.): Suppose $(\theta_{1\mathbf{a}}^* - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) / (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) < 0$. If $a^* > 0$, then $a^* (\theta_{1\mathbf{a}}^* - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) / (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) < 0$, and so, by (52), $a^* > a^o$. (2.) implies (3.): Suppose $(\theta_{1\mathbf{a}}^* - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) / (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) < 0$. If $a^* < 0$, then $a^* (\theta_{1\mathbf{a}}^* - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) / (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) > 0$, and so, by (52), $a^* < a^o$. (3.) implies (2.): Suppose that $a^* < 0$ implies $a^* < a^o$. By (52), $a^* < 0$ implies

$$a^* \frac{\theta_{1\mathbf{a}}^* - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})}{\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})} > 0$$

and so $(\theta_{1\mathbf{a}}^* - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) / (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) < 0$. ■

Proof of Proposition 14 By Lemma 8, $a^* > 0$ implies $a^* < a^o$ if and only if $(\theta_{1\mathbf{a}}^* - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) / (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) < 0$. In turn, this happens if and only if $\theta_{1\mathbf{a}}^* \vee (-\hat{\beta}^*) < \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})$ or $\mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) < \theta_{1\mathbf{a}}^* \wedge (-\hat{\beta}^*)$. By Lemma 8, $a^* > 0$ implies $a^* < a^o$ if and only if $(\theta_{1\mathbf{a}}^* - \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) / (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}})) > 0$. In turn, this happens if and only if $\theta_{1\mathbf{a}}^* \vee (-\hat{\beta}^*) > \mathbb{E}_{\mu^*}(\theta_{1\mathbf{a}}) > \theta_{1\mathbf{a}}^* \wedge (-\hat{\beta}^*)$. ■

References

- [1] Aliprantis, C. D. and K. C. Border, *Infinite dimensional analysis*, 3rd ed., Springer, Berlin, 2006.
- [2] Arrow, K. J. Alternative approaches to the theory of choice in risk-taking situations, *Econometrica*, 19, 404-437, 1951.
- [3] Battigalli P., *Comportamento razionale ed equilibrio nei giochi e nelle situazioni sociali*, unpublished thesis, Università Bocconi, 1987.
- [4] Battigalli, P. and D. Guaitoli, Conjectural equilibria and rationalizability in a macro-economic game with incomplete information, Quaderni di Ricerca, Università Bocconi, 1988 (published in *Decisions, Games and Markets*, Kluwer, Dordrecht, 97-124, 1997).
- [5] Battigalli, P., S. Cerreia-Vioglio, F. Maccheroni, and M. Marinacci, Self-confirming equilibrium and model uncertainty, *American Economic Review*, 105, 646-677, 2015.
- [6] Berberian, S. K., *Borel spaces*, mimeo, 1997.
- [7] Bernanke, B. S., Monetary policy under uncertainty, *Annual Economic Policy Conference*, Federal Reserve Bank of St. Louis, 2007.
- [8] Bertsekas, D. P. and S. E. Shreve, *Stochastic optimal control: the discrete-time case*, Academic Press, New York, 1978.

- [9] Billingsley, P., *Convergence of probability measures*, 2nd ed., Wiley, New York, 1999.
- [10] Billingsley, P., *Probability and measure*, anniversary ed., Wiley, New York, 2012.
- [11] Blinder, A. S., What central bankers could learn from academics - and vice versa, *Journal of Economic Perspectives*, 11, 3-20, 1997.
- [12] Blinder, A. S., *Central banking in theory and practice*, MIT Press, Cambridge, 1998.
- [13] Brainard, W., Uncertainty and the effectiveness of policy, *American Economic Review P&P*, 57, 411-425, 1967.
- [14] Cerreia-Vioglio S., F. Maccheroni, M. Marinacci, and L. Montrucchio, Classical subjective expected utility, *Proceedings of the National Academy of Sciences*, 110, 6754-6759, 2013.
- [15] Cho, I.-K. and K. Kasa, Learning and model validation, *Review of Economic Studies*, 82, 45-82, 2015.
- [16] Cho, I.-K., N. Williams, and T. Sargent, Escaping Nash inflation, *Review of Economic Studies*, 69, 1-40, 2002.
- [17] 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.
- [18] 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.
- [19] Denzau, A. T. and D. C. North, Shared mental models: ideologies and institutions, *Kyklos*, 47, 3-31, 1994.
- [20] Dudley, R. M., *Real analysis and probability*, Cambridge University Press, Cambridge, 2002.
- [21] Dynkin, E. B. and A. A. Yushkevich, *Controlled Markov processes*, Springer, Berlin, 1979.
- [22] Easley, D. and N. M. Kiefer, Controlling a stochastic process with unknown parameters, *Econometrica*, 5, 1045-1064, 1988.
- [23] Friedman, M., The effects of full employment policy on economic stability: a formal analysis, in *Essays in Positive Economics* (M. Friedman, ed.) 117-132, University of Chicago Press, Chicago, 1953.
- [24] Fudenberg, D. and D. Kreps, Learning in extensive games, I: self-confirming equilibria, *Games and Economic Behavior*, 8, 20-55, 1995.
- [25] Fudenberg, D. and D. K. Levine, Self-confirming equilibrium, *Econometrica*, 61, 523-545, 1993a.

- [26] Fudenberg, D. and D. K. Levine, Steady state learning and Nash equilibrium, *Econometrica*, 61, 547-573, 1993b.
- [27] Fudenberg, D., and D. K. Levine, Self-confirming equilibrium and the Lucas critique, *Journal of Economic Theory*, 144, 2354-2371, 2009.
- [28] Gaballo G. and R. Marimon, Breaking the spell with credit easing, mimeo, European University Institute, 2015.
- [29] Gaudard, M. and D. Hadwin, Sigma-algebras on spaces of probability measures, *Scandinavian Journal of Statistics*, 16, 169-175, 1989.
- [30] Jovanovic, B., Observable implications of models with multiple equilibria, *Econometrica*, 57, 1431-1437, 1989.
- [31] Kalai, E. and E. Lehrer, Subjective equilibrium in repeated games, *Econometrica*, 61, 1231-1240, 1993.
- [32] Kalai, E. and E. Lehrer, Subjective games and equilibria, *Games and Economic Behavior*, 8, 123-163, 1995.
- [33] Kurz, M., On rational belief equilibria, *Economic Theory*, 4, 859-876, 1994a.
- [34] Kurz, M., On the structure and diversity of Rational Beliefs, *Economic Theory*, 4, 877-900, 1994b.
- [35] Koopmans, T. C., Measurement without theory, *Review of Economic Statistics*, 29, 161-172, 1947.
- [36] Lucas, R. E. Jr., Expectations and the neutrality of money, *Journal of Economic Theory*, 4, 103-124, 1972.
- [37] Lucas, R. E. Jr. and E. C. Prescott, Investment under uncertainty, *Econometrica*, 39, 659-681, 1971.
- [38] Lucas, R. E. Jr., Expectations and the Neutrality of Money, *Journal of Economic Theory* 4(2), 103-124, 1972
- [39] Lucas, R. E. Jr., Methods and problems in business cycle theory, *Journal of Money, Credit, and Banking*, 12, 696-715, 1980.
- [40] Mackey, G. W., Borel structure in groups and their duals, *Transactions of the American Mathematical Society*, 134-165, 1957.
- [41] Marinacci, M., Model uncertainty, *Journal of the European Economic Association*, 13, 998-1076, 2015.
- [42] Marschak, J., Economic measurements for policy and prediction, in *Studies in Econometric Method* (W. Hood and T. J. Koopmans, eds.), 1-26, Wiley, New York, 1953.
- [43] Poole, W., Optimal choice of monetary policy instruments in a simple stochastic macro model, *Quarterly Journal of Economics*, 84, 197-216, 1970.

- [44] Puterman, M. L., *Markov decision processes: discrete stochastic dynamic programming*, Wiley, New York, 2014.
- [45] Rodrik, D., When ideas trump interests: preferences, worldviews, and policy innovations, *Journal of Economic Perspectives*, 28, 189-208, 2014.
- [46] Rothenberg, T. J., Identification in parametric models, *Econometrica*, 39, 577-591, 1971.
- [47] Samuelson, P. A. and R. M. Solow, Analytical aspects of anti-inflation policy, *American Economic Review P&P*, 50, 177-184, 1960.
- [48] Sargent, T. J., Rational expectations, the real rate of interest, and the natural rate of unemployment, *Brookings Papers on Economic Activity*, 429-472, 1973.
- [49] Sargent, T. J., *The conquest of American inflation*, Princeton University Press, Princeton, 1999.
- [50] Sargent, T. J., Evolution and intelligent design, *American Economic Review*, 98, 3-37, 2008.
- [51] Sargent T. J. and N. Williams, Impacts of priors on convergence and escapes from Nash inflation, *Review of Economic Dynamics*, 8, 360-391, 2005.
- [52] 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.
- [53] Savage, L. J., *The foundations of statistics*, Wiley, New York, 1954.
- [54] Soderstrom, U., Monetary policy with uncertain parameters, *Scandinavian Journal of Economics*, 104, 125-145, 2002.
- [55] Srivastava, S. M., *A course on Borel sets*, Springer, Berlin, 1998.
- [56] Theil, H., *Economic forecasts and policy*, North Holland, Amsterdam, 1961.
- [57] Tinbergen, J., *On the theory of economic policy*, North Holland, Amsterdam, 1952.