Monetary policies in self-confirming equilibria with uncertain models

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

March 14, 2014(preliminary, do not circulate)

1 Introduction

Rational expectations models remove pervasive inconsistencies in pre-rational expectations models between decision makers' beliefs and what observers can infer statistically from (long) histories of outcomes. But they do much more. Rational expectations models eradicate discrepancies between decision makers' beliefs and *what would be empirical distributions* under decisions that have not been and will not be taken.

The concept of a *self-confirming* equilibrium weakens rational expectations by requiring consistency only between decision makers' beliefs and what can be inferred statistically from (long) histories of outcomes *under decisions that they actually make.* By leaving room for decision makers to have "wrong" views about distributions of outcomes under decisions that they do not take, a self-confirming equilibrium can naturally confront decision makers with *model uncertainty.*

We study how model uncertainty affects sets of self-confirming beliefs and actions.

2 A recurrent decision problem

We consider decision makers who face a static decision problem under uncertainty, interpreted as a (stochastic) long run rest/steady state of a recurrent decision problem that they have been facing over a long period of time in a stationary environment.[1] In such (unmodelled) repeated problem, decision makers care only about current payoffs. Voluntary experimentation is thus not relevant, with purely informational linkages among problems.

At least two different interpretations are possible of this steady state setup. We can think of a single myopic decision maker that, repeatedly, is called upon to act in the recurrent problem. For instance, a policy maker that over a long period of time repeatedly holds an office and, myopically, cares only about the payoff that he can achieve within each term. Alternatively, we can consider a large population of decision makers, one of them being called upon to act in each occurrence of the decision problem. Actions are chosen based on the current payoff, because decision makers are either myopic or randomly selected (and so with negligible chances of a future selection), as well as on the information that the population gathered through its members that previously acted. For example, the population may be a political party, with each decision problem being an office term that a party member may be selected to hold.

Be that as it may, in the static problem that models the steady state the uncertain state s realization and the chosen action a jointly determine a payoff $c = \rho(a, s)$, as well as an information feedback $m = f_a(s)$ in terms of messages/observations.[2] Feedback is perfect if decision makers

[1] As Lucas (1980) p. 701 writes, at a stochastic rest state "external shocks remain fixed over a long period, so that households and firms would adjust to facing the same set of prices over and over again and attune their behavior accordingly."

[2] We use the terms message and observation interchangeably.

observe the realized state, that is, $m = s$. It is imperfect otherwise: for example they may only observe payoffs, that is, $m = c$.

States are generated by a chance mechanism described by a probability model p^*. The different states that, over time, model p^* may generate determine a distribution σ of messages, in which $\sigma(m)$ is interpreted as both the long run empirical frequency with which message m occurred and the probability with which it may obtain. This dual interpretation characterizes the steady state interpretation of our static decision problem. It relies, heuristically, on a "long run" that is long and stationary enough for asymptotic ergodic-type results to hold.

The distribution σ summarizes all that decision makers can have learned from the infinite history of messages that their actions generated via the feedback function. Under imperfect feedback, the decision problem might feature partial identification in the sense that different probability models are compatible with the same long run frequencies of messages σ presented to the decision makers under their action a and the true mechanism p^*. This defines the fundamental inference problem faced by our decision makers, and explains why uncertainty is pervasive also in steady states.

Denote by $\hat{\Sigma}_a(p^*)$ the collection of all such observationally equivalent models. Given the feedback f, the collection depends on both the selected action a and the true model p^*. Decision makers' beliefs about the relative likelihoods of probability models in $\hat{\Sigma}_a(p^*)$ are summarized by a subjective probability μ. The probability μ has to be consistent with what decision makers learned through action a. For this reason μ assigns positive probability only to models that, under action a, are observationally equivalent to the true model; that is, its support is included in $\hat{\Sigma}_a(p^*)$. This data confirmation requirement is the only condition we impose on μ. It is a much weaker requirement than assuming, as often done in the rational expectation literature, that decision makers know the true model. Rational expectations prevails in our setup only when there is perfect feedback, in which case full identification holds.

An action a is self-confirming if it is optimal with respect to belief μ. That is, a self-confirming action is optimal with respect to the information that it generates. Partial identification makes a nontrivial subjective probability μ play a key role despite the recurrent nature of the decision problem. But partial identification nevertheless restricts μ by requiring it to assign positive probability only to models that belong to $\hat{\Sigma}_a(p^*)$. The more partial is the identification, the weaker is the discipline.

We study steady states and have nothing to say about transient dynamics.

3 Preliminaries

3.1 Mathematics

Throughout we consider Polish spaces, that is, complete and separable metric spaces.[3] Let S be any such space and $\mathcal{S}$ its Borel σ-algebra. We denote by $\Delta(S)$ the collection of all (Borel) probability measures $\nu : \mathcal{S} \to [0, 1]$.[4] When S is finite, say with cardinality n, we assume that $\mathcal{S}$ is the power set 2^S and we identify $\Delta(S)$ with the simplex of $\mathbb{R}^n$.

We endow $\Delta(S)$ with the smallest σ-algebra that makes the real valued and bounded functions on $\Delta(S)$, of the form $\nu \mapsto \nu(E)$, measurable for all $E \in \mathcal{S}$. When S is finite this σ-algebra coincides with the relative Borel σ-algebra that $\Delta(S)$ inherits as the simplex of $\mathbb{R}^n$. Finally, we also endow any measurable subset Σ of $\Delta(S)$ with the relative σ-algebra inherited from $\Delta(S)$, and we denote by $\Delta(\Sigma)$ the collection of all probability measures defined on this σ-algebra.

Given any two pairs $(S, \mathcal{S})$ and $(Y, \mathcal{Y})$, we endow the Cartesian product $S \times Y$ with the product σ-algebra $\mathcal{S} \times \mathcal{Y}$. We denote by $\Delta(S) \otimes \Delta(Y)$ the collection of all *product probability measures.*

[3] Polish spaces, which are standard in probability theory, include finite spaces, compact metric spaces, and Euclidean spaces.

[4] Among them, δ_x denotes the Dirac measure concentrated on $x \in X$, that is, $\delta_x(A) = 1$ if $x \in A$ and $\delta_x(A) = 0$ if $x \notin A$.

Moreover, each measurable function $\varphi : S \to Y$ induces the pushforward map $\hat{\varphi} : \Delta(S) \to \Delta(Y)$ defined by

$$\hat{\varphi}(\nu) = \nu \circ \varphi^{-1} \qquad \forall \nu \in \Delta(S).$$

In other words, $\hat{\varphi}(\nu)(E) = \nu(\varphi^{-1}(E))$ for all $E \in \mathcal{Y}$. The following routine lemma describes a key feature of $\hat{\varphi}$.

Lemma 1 *$\hat{\varphi}$ is one-to-one if and only if φ is one-to-one.*

3.2 Classical subjective expected utility

Let S be a space of states of nature and C a consequence space. For simplicity, we assume that C is a convex subset of $\mathbb{R}^n$ (it becomes an interval of the real line when consequences are monetary). We consider a setting where A is a set of *actions* (or *controls*) available to the decision maker and $\rho : A \times S \to C$ is a *consequence function* that associates a consequence $\rho(a, s) \in C$ to each pair $(a, s) \in A \times S$ of action and state. The quartet (A, S, C, ρ) is the basic structure of the decision problem. As in Savage (1954), we assume that

$$\rho(a', s) = \rho(a, s) \quad \forall s \in S \Longrightarrow a = a' \tag{1}$$

that is, we identify actions that share the same consequences in all states.

The inherent randomness – often called aleatory or physical uncertainty – that characterizes states' realizations is described by probability models $p : \mathcal{S} \to [0, 1]$. They can be regarded as possible data generating mechanisms (see Neyman, 1957). Each action a is evaluated through its expected utility

$$\int_S (v \circ \rho)(a, s)\, dp(s)$$

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

$$R(a, p) = \int_S r(a, s)\, dp(s)$$

where $r : A \times S \to \mathbb{R}$ is the *payoff* (or *reward*) *function* $r = v \circ \rho$. In order to have all our integrals well defined, we assume throughout that r is bounded, i.e., $\sup_{(a,s)\in A\times S} |r(a, s)| < \infty$.

As in Cerreia-Vioglio et al (2011), we assume a la Neyman-Pearson-Wald (cf. Arrow, 1951, p. 418) that decision makers do not know the true probability model but that they know a collection $\Sigma \subseteq \Delta(S)$ of probability models that contains the true one. 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*.[5] They rank actions according to the *classical subjective expected utility* (SEU) criterion:

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

where $\mu \in \Delta(\Sigma)$ is a (*prior*) *subjective probability* of models in Σ that reflects some personal information on models that decision makers may have, in addition to the "physical" information behind Σ.

Representation (2) admits the reduced form

$$\int_\Sigma R(a, p)\, d\mu(p) = \int_S r(a, s)\, dp_\mu = R(a, p_\mu)$$

[5] Often model uncertainty is called epistemic uncertainty. However, any kind of uncertainty relevant for decision problems is, ultimately, epistemic (i.e., relative to decision makers' information). Physical uncertainty is thus epistemic as well.

where $p_\mu \in \Delta(S)$ is the subjective *predictive probability*

$$p_\mu(E) = \int_\Sigma p(E)\, d\mu(p) \qquad \forall E \in \mathcal{S}.$$

This reduced form is the original representation of Savage (1954), who elicited p_μ from betting behavior.

The decision problem can be summarized by the sextet $D = (A, S, C, \rho, \Sigma, v)$, which 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 $\{p\}$,[6] decision makers believe that p is the true model. The predictive probability trivially coincides with p and criterion (2) reduces to the Savage expected payoff criterion $R(a, p)$. Being a predictive probability, p here is a subjective probability, albeit a dogmatic one.

(ii) When Σ is a singleton $\{p\}$, decision makers have maximal structural information and, as result, know that p is the true model. There is only physical uncertainty, quantified by p, without any model uncertainty. Criterion (2) again reduces to the expected payoff criterion $R(a, p)$, 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 its balls, they would learn/know the urn composition, so Σ would be a singleton.

(iii) When $\Sigma = \{\delta_s : s \in S\}$, 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 (2) takes the form $R(a, \mu)$.

Finally, note that in the reduction operation that generates predictive probabilities some important probabilistic features may disappear. For instance, in a binary state space $S = \{s_1, s_2\}$ consider the two collections $\Sigma = \{(0,1), (1,0)\}$ and $\Sigma' = \{(\frac{1}{2} - \delta, \frac{1}{2} + \delta) : 0 \leq \delta \leq \varepsilon\}$, where $0 < \varepsilon < 1/2$ is an arbitrarily small quantity. If we take a uniform prior on each collection, we end up in both cases with the uniform predictive probability on S that assigns probability $1/2$ to each state. However, while the collection Σ consists of two very different models, the collection Σ' consists of many almost identical ones. Very different probabilistic scenarios thus reduce to the same predictive probability.

4 Partial identification

4.1 Partial identification correspondence

We model decision makers' information through a (*outcome*) *feedback function* $f : A \times S \to M$, where M is a message space. By selecting action $a \in A$ the decision maker receives a *message*

$$m = f_a(s)$$

when s occurs. The decision maker's information about the state is thus endogenous: the partition[7]

$$\left\{f_a^{-1}(m) : m \in M\right\} \subseteq 2^S$$

of the state space S that messages induce depends on the choice of action a.

When information does not depend on a, we say that there is *own-action independence* of feedback; formally, $f_a = f_{a'}$ for each $a, a' \in A$. The most important instance of own-action independence is

[6] We say that a prior μ has finite support if and only if it admits a finite carrier. In this case, we denote by $\operatorname{supp}\mu$ its minimal carrier. Note that $\mu \in \Delta(\Sigma)$ amounts to saying that $\operatorname{supp}\mu \subseteq \Sigma$.

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

perfect feedback, which occurs when the feedback function f is one-to-one. Without loss of generality, we can then write $f = \mathrm{Id}_S$, where Id_S is the identity function, i.e., $\mathrm{Id}_S(s) = s$ for each $s \in S$. In this case, messages reveal to the decision maker which state obtained, regardless of the chosen action. When this is not the case, information is imperfect (maximally imperfect when each f_a is constant).

An action a is *fully revealing* if $f_a = \mathrm{Id}_S$, that is, if it allows the decision maker to know 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 that consequences are observable, that is, for each $a \in A$,

$$f_a(s) = f_a(s') \Longrightarrow \rho_a(s) = \rho_a(s') \qquad \forall s, s' \in S \tag{3}$$

This condition ensures that feedback functions reflect the information on states ensured by outcome observability.

In our setting a message distribution $\sigma \in \Delta(M)$ can be interpreted as a long run frequency distribution of messages received by the decision maker, so that $\sigma(m)$ is the empirical frequency of message m. Given an action $a \in A$, consider the pushforward map $\hat{f}_a : \Sigma \to \Delta(M)$ given, for each $p \in \Sigma$, by

$$\hat{f}_a(p)(m) = p(s : f_a(s) = m) \qquad \forall m \in M$$

In other words, $\hat{f}_a(p)(m)$ is the empirical frequency with which the decision maker receives message m when he chooses action a and p is the true model. Its inverse correspondence $\hat{f}_a^{-1}$ partitions Σ into classes

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

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

The inverse correspondence $\hat{f}_a^{-1}(\sigma)$, which is compact valued, is a function if and only if $\hat{f}_a$ is one-to-one, that is, action a is fully revealing (Lemma 1). In this case the decision problem is identified under a since different models generate different message distributions, which thus uniquely pin down models. Accordingly, we have *partial identification* only when $\hat{f}_a$ is not one-to-one, that is, when the action is not fully revealing. 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.

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

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

For any fixed $p \in \Sigma$, the image

$$\hat{\Sigma}_a(p) = \left\{ p' \in \Sigma : \hat{f}_a(p') = \hat{f}_a(p) \right\} \tag{4}$$

is the collection of models that are observationally equivalent given the message distribution $\sigma = \hat{f}_a(p)$ that action a generates along with model p. In other words, $\hat{\Sigma}_a(p)$ is the estimated collection Σ given action and messages.

We can regard $\hat{\Sigma}_a$ as the *partial identification correspondence* determined by action a. It is compact valued, as well as nonempty since $p \in \hat{\Sigma}_a(p)$. It is convex if Σ is convex. Finally, it is a function if and only if $\hat{f}_a$ is one-to-one, that is, full identification characterizes fully revealing actions. In this case, $\hat{\Sigma}_a$ is the identity function, with $\hat{\Sigma}_a(p) = \{p\}$ for all $p \in \Sigma$, and so message distributions identify the true model. In contrast, when $\hat{\Sigma}_a(p)$ is a nonsingleton there is partial identification.

The partial identification correspondence has disjoint images:

Lemma 2 *Let $a \in A$. For any $p, p' \in \Sigma$, either $\hat{\Sigma}_a(p) = \hat{\Sigma}_a(p')$ or $\hat{\Sigma}_a(p) \cap \hat{\Sigma}_a(p') = \emptyset$.*

The collection $\{\hat{\Sigma}_a(p)\}_{p \in \Sigma}$ is thus a partition of Σ that we denote by Σ_a. Each element of this partition consists of probability models that are observationally equivalent under action a.

Example 3 Consider a decision maker who is asked to bet 1 euro on the color of the ball that will be drawn from an urn that contains 90 black, green, and yellow balls. After the draw, he is told whether or not he won 1 euro. 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, $f = \rho$. In particular, consider a decision maker who bets on black, that is, who chooses action b (similar properties hold for bets on either green or yellow). Then

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

that is, betting on b yields the partition $\{\{B\}, \{Y, G\}\}$ of S. As a result, the decision maker imperfectly observes the draws: he can only observe whether the color is black (message 1: he won) or not (message 0: he lost). His action prevents him from obtaining any evidence on the frequency of G and Y. In particular,

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

where $\Sigma \subseteq \Delta_q(\{B, G, Y\})$ is the set of possible urn compositions that he posits.[8] Hence, the evidence gathered through bet b only partially identifies the true model:

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

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

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

In other words, all probability models p' that assigns probability 1/3 to B are observationally equivalent. In particular, if we denote by p_β any probability model that assigns probability $\beta = n/90$ to B, then the partition $\Sigma_b = \{\hat{\Sigma}_b(p_\beta)\}_{\beta \in \{0,1,\ldots,90\}}$ has 91 elements, each consisting of the probability models that assigns probability β to B. All such models are observationally equivalent. ▲

Example 4 Consider the previous example again. We have perfect feedback when, after the draw, the decision maker is told the color of the ball. In this case $f = \mathrm{Id}_S$, that is, $f(b, s) = f(g, s) = f(y, s) = s$ for each $s \in \{B, G, Y\}$. Hence, $\hat{\Sigma}_b(p) = \hat{\Sigma}_g(p) = \hat{\Sigma}_y(p) = \{p\}$ for each $p \in \Sigma$. Regardless of the chosen action, the true model is identified. ▲

4.2 Comparative statics

Partial identification correspondences with larger images exhibit a higher degree of partial identification. We now show that the extent of partial identification ultimately depends on the underlying feedback functions. To this end, given any two feedback functions f and $\bar{f}$, say that f is *coarser* than $\bar{f}$ if, for each $a \in A$,

$$\bar{f}_a(s) = \bar{f}_a(s') \Longrightarrow f_a(s) = f_a(s') \qquad \forall s, s' \in S.$$

Decision makers with coarser feedback functions have worse information about states. Constant feedbacks are the coarsest, while perfect feedbacks are the least coarse.

[8] Here $\Delta_q(\{B, G, Y\})$ denotes the collection of all possible urn compositions, that is, the elements of $\Delta(\{B, G, Y\})$ that have rational values. Unlike $\Delta(\{B, G, Y\})$, the set $\Delta_q(\{B, G, Y\})$ is not convex.

Lemma 5 *If f is coarser than $\bar{f}$, then, for each $a \in A$,*

$$\hat{\bar{\Sigma}}_a(p) \subseteq \hat{\Sigma}_a(p) \qquad \forall p \in \Sigma.$$

Proof Since f is coarser that $\bar{f}$, there is a function $\varphi : \operatorname{Im} \bar{f} \to M$ such that $f = \varphi \circ \bar{f}$. Let $p^* \in \Sigma$. Given any $p \in \hat{\bar{\Sigma}}_a(p^*)$, we have

$$\begin{aligned} p \in \hat{\Sigma}_a(p) &\iff \hat{f}_a(p) = \hat{f}_a(p^*) \iff p(f_a = m) = p^*(f_a = m) \qquad \forall m \in M \\ &\iff p\left(\varphi \circ \bar{f}_a = m\right) = p^*\left(\varphi \circ \bar{f}_a = m\right) \qquad \forall m \in M \\ &\iff p\left(\bar{f}_a \in \varphi^{-1}(m)\right) = p^*\left(\bar{f}_a \in \varphi^{-1}(m)\right) \qquad \forall m \in M \\ &\iff \hat{\bar{f}}_a(\sigma)\left(\varphi^{-1}(m)\right) = \hat{\bar{f}}_a(p^*)\left(\varphi^{-1}(m)\right) \qquad \forall m \in M \end{aligned}$$

Since $\hat{\bar{f}}_a(p) = \hat{\bar{f}}_a(p^*)$, we conclude that $p \in \hat{\bar{\Sigma}}_a(p)$. ∎

Coarser feedback functions thus determine, for each action, coarser partial identification correspondences: worse information translates into a higher degree of partial identification. In particular, partition Σ_a is coarser than partition $\bar{\Sigma}_a$.

5 Self-confirming actions and beliefs

5.1 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 6 *Given a true model $p^* \in \Sigma$, a pair $(a^*, \mu^*) \in A \times \Delta(\Sigma)$ of actions and beliefs is a self-confirming equilibrium if*

$$V(a^*, \mu^*) \geq V(a, \mu^*) \qquad \forall a \in A \tag{5}$$

and

$$\mu^* \in \Delta\left(\hat{\Sigma}_{a^*}(p^*)\right) \tag{6}$$

The definition relies on two pillars: the optimality condition (5) that ensures that action a^* is optimal under belief μ^*; the data confirmation condition (6) that guarantees that belief μ^* is consistent with the data that action a^* reveals.

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

Example 7 In the urn setting of Example 3, suppose that $\Sigma = \{p^*, p_1, p_2\} \subseteq \Delta(\{B, G, Y\})$, where $p^* = (1/3, 0, 2/3)$ is the true model, while $p_1 = (1/3, 2/3, 0)$ and $p_2 = (1/3, 1/3, 1/3)$ Impose the normalization $u(0) = 0$ and $u(1) = 1$, so that

$$\begin{aligned} &R(b, p^*) = R(b, p_2) = R(b, p_1) = R(y, p_2) = R(g, p_2) = \frac{1}{3} \\ &R(y, p^*) = R(g, p_1) = \frac{2}{3} \quad ; \quad R(y, p_1) = R(g, p^*) = 0 \end{aligned} \tag{7}$$

(i) Consider a uniform probability $\mu^* \in \Delta(\Sigma)$ with support $\{p^*, p_1, p_2\}$ and $\mu^*(p^*) = \mu^*(p_1) = \mu^*(p_2) = 1/3$. The action and belief pair (b, μ^*) is self-confirming. Since $V(b, \mu^*) = V(g, \mu^*) = V(y, \mu^*) = 1/3$, the optimality condition (5) is satisfied. It is easy to check (cf. Example 3) that $\hat{\Sigma}_b(p^*) = \{p \in \Sigma : p(B) = 1/3\} = \Sigma$. Hence $\mu^* \in \Delta\left(\hat{\Sigma}_b(p^*)\right)$, and so the data confirmation condition (6) is also satisfied. The self-confirming equilibrium (b, μ^*) generates the message distribution $\sigma^* \in \Delta(\{0, 1\})$ with $\sigma^*(1) = 1/3$, that is, with a one third frequency of wins.

(ii) Consider a probability $\mu^* = \delta_{p^*}$ concentrated on the true model. The action and belief pair (y, μ^*) is self-confirming. The optimality condition (5) is easily seen to be satisfied, while the data confirmation condition (6) holds since $\hat{\Sigma}_y(p^*) = \{p \in \Sigma : p(Y) = 2/3\} = \{p^*\}$. The self-confirming equilibrium (y, μ^*) generates the message distribution $\sigma^* \in \Delta(\{0,1\})$ with $\sigma^*(1) = 2/3$, that is, with a two thirds frequency of wins.

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

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

Example 8 If in the previous example we suppose that only actions b and g are available, i.e., $A = \{b, g\}$, action g is still not part of any self-confirming equilibrium. But, if we further suppose that there is no structural information, so that $\Sigma = \Delta_q(\{B, G, Y\})$ consists of all possible urn compositions, this is no longer the case. For, the pair (g, δ_{δ_Y}) is self-confirming, where δ_{δ_Y} is the belief concentrated on the probability model $\delta_Y = (0, 0, 1)$. In fact, $V(b, \delta_{\delta_Y}) = V(g, \delta_{\delta_Y}) = 0$ and $\delta_Y \in \hat{\Sigma}_g(p^*) = \{p \in \Delta_q(\{B, G, Y\}) : p(G) = 0\}$. ▲

Remark We confine our analysis to statistically stationary states and so are silent about transition processes that might or converge to a steady state.

Under own-action independence of feedback (that is, actions do not affect information gathering), condition (6) becomes $\mu^* \in \Delta(\hat{\Sigma}(p^*))$, where $\hat{\Sigma}(p^*)$ is exogenously posited. We thus return to a traditional optimization notion (a la Nash in a game theoretic perspective) with a purely exogenous data confirmation condition (6). In particular, under perfect feedback (and so full identification), condition (5) becomes

$$R(a^*, p^*) \geq R(a, p^*) \qquad \forall a \in A \tag{8}$$

since condition (6) requires $\mu^* = \delta_{p^*}$. In this case, common in the rational expectations literature, decision makers know the true model. They confront only physical uncertainty.

Action a^* is *objectively optimal* if it satisfies condition (8). Objectively optimal actions are the ones that decision makers would select if they 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. In any case, the pair (a^*, δ_{p^*}), where action a^* is objectively optimal and belief δ_{p^*} is concentrated on the true model, is always self-confirming.

Example 9 In Example 7, bet y is the objectively optimal action. ▲

Condition (5) can be written in predictive form as $R(a^*, p_{\mu^*}) \geq R(a, p_{\mu^*})$ for each $a \in A$. Relatedly, the data confirmation condition (6) trivially implies that the predictive probability p_{μ^*} belongs to $\hat{\Sigma}_{a^*}(p^*)$. In turn this implies that

$$\hat{f}_{a^*}(p_{\mu^*}) = \hat{f}_{a^*}(p^*) \tag{9}$$

that is, the predictive probability and the true model both assign the same probability to messages. Condition (9) is an alternative data confirmation condition in terms of predictive distributions, weaker than (6). It is easily seen to be equivalent when we further require that $p_{\mu^*} \in \Sigma$, because this implies $p_{\mu^*} \in \hat{\Sigma}_{a^*}(p^*)$, and so $\mu^* \in \Delta(\hat{\Sigma}_{a^*}(p^*))$.

Finally, consider

$$p_{\mu^*} = p^* \tag{10}$$

a much stronger data confirmation condition than (6) and (9). If (a^*, μ^*) is a self-confirming pair in which μ^* satisfies condition (10), then a^* is easily seen to be objectively optimal. Any belief that satisfies condition (10) is thus equivalent, in terms of actions, to the belief $\mu^* = \delta_{p^*}$ 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.

5.2 Value

5.2.1 Self-confirming equilibria

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

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

associates to each possible true model p^* its collection $\Gamma(p^*)$ of self-confirming equilibria (a^*, μ^*).

Through the self-confirming correspondence we can characterize the value of self-confirming equilibria

Proposition 10 *$V(a^*, \mu^*) = R(a^*, p^*)$ for each self-confirming equilibrium $(a^*, \mu^*) \in \Gamma(p^*)$.*

The result is based on the following lemma of independent interest, based on Battigalli et al (2011), which shows that observationally equivalent models share the same expected utility. Formally, the function $R(a, \cdot) : \Sigma \to \mathbb{R}$ is Σ_a-measurable.

Lemma 11 *Given any a and p, $R(a, p') = R(a, p)$ for each $p' \in \hat{\Sigma}_a(p)$.*

Proof Since payoffs are observable, there exists a function $\bar{r} : \operatorname{Im} f \to \mathbb{R}$ such that $r = \bar{r} \circ f$. Let $p' \in \hat{\Sigma}_a(p)$. Since $\hat{f}_a(p') = \hat{f}_a(p)$, it is true that

$$\begin{aligned} R(a, p') &= \int_S \bar{r}(f(a, s))\, dp'(s) = \int_M \bar{r}(m)\, d\hat{f}_a(p')(m) \\ &= \int_M \bar{r}(m)\, d\hat{f}_a(p)(m) = \int_S \bar{r}(f(a, s))\, dp'(s) = R(a, p) \end{aligned}$$

as desired. ■

Proof of Proposition 10 By the previous lemma, $R(a, p) = k$ for each $p \in \hat{\Sigma}_a(p^*)$. Hence, $V(a^*, \mu^*) = \int_\Sigma R(a, p)\, d\mu^*(p) = k$. ■

5.2.2 Subjective model uncertainty

A basic subjective assessment of the decision maker is about which models in Σ he deems actually possible and, as a result, are assigned a strictly positive mass by his prior, and so belong to the prior support. Next we show that self-confirming equilibria with sharper basic subjective assessments have higher values.

Proposition 12 *If $\operatorname{supp} \mu^* \subseteq \operatorname{supp} \nu^*$ then $V(a^*, \mu^*) \geq V(b^*, \nu^*)$ for all self-confirming equilibria $(a^*, \mu^*), (b^*, \nu^*) \in \Gamma(p^*)$.*

Proof Since $\operatorname{supp} \mu^* \subseteq \hat{\Sigma}_{a^*}(p^*)$ and $\operatorname{supp} \nu^* \subseteq \hat{\Sigma}_{b^*}(p^*)$, by Lemma 11 it holds $R(a^*, p) = R(a^*, p^*)$ for each $p \in \operatorname{supp} \mu^* \supseteq \operatorname{supp} \nu^*$ and $R(b^*, p) = R(b^*, p^*)$ for each $p \in \operatorname{supp} \nu^*$. Hence,

$$V(a^*, \mu^*) = \int_{\operatorname{supp} \mu^*} R(a^*, p)\, d\mu^*(p) = R(a^*, p^*)$$

$$V(b^*, \nu^*) = \int_{\operatorname{supp} \nu^*} R(b^*, p)\, d\nu^*(p) = R(b^*, p^*)$$

Given the optimality of b^*, it holds $V(b^*, \nu^*) \geq \int_{\operatorname{supp} \nu^*} R(a^*, p)\, d\nu^*(p) = R(a^*, p^*) = V(a^*, \mu^*)$, proving the statement. ■

Priors μ^* and ν^* that share the same support are called equivalent, written $\mu^* \sim \nu^*$.[9] By the previous result, if $\mu^* \sim \nu^*$ then $V(a^*, \mu^*) = V(b^*, \nu^*)$ for all self-confirming equilibria $(a^*, \mu^*), (b^*, \nu^*) \in$

[9] When their supports are finite, priors are equivalent if and only if they assign zero probability to the same events.

$\Gamma(p^*)$. The value of self-confirming equilibria is thus pinned down by the decision maker basic subjective assessment. The specific shape of the prior is value irrelevant, only supports matter. But, more is actually true: actions can be exchanged across such self-confirming equilibria.

Proposition 13 *If $\mu^* \sim \nu^*$ then $(a^*, \nu^*), (b^*, \mu^*) \in \Gamma(p^*)$ for all self-confirming equilibria $(a^*, \mu^*), (b^*, \nu^*) \in \Gamma(p^*)$.*

Proof Using the notation of the previous proof, since $R(b^*, p) = R(b^*, p^*)$ for each $p \in \operatorname{supp} \nu^*$, it holds

$$\begin{aligned} V(b^*, \mu^*) &= \int_{\operatorname{supp}\mu^*} R(b^*, p)\, d\mu^*(p) = \int_{\operatorname{supp}\nu^*} R(b^*, p)\, d\mu^*(p) \\ &= R(b^*, p^*) = R(a^*, p^*) \geq R(a, p^*) \end{aligned}$$

for all $a \in A$. Hence, $(b^*, \mu^*) \in \Gamma(p^*)$. A similar argument shows that $(a^*, \nu^*) \in \Gamma(p^*)$. ■

5.2.3 Objective model uncertainty

The support of μ^* can be seen as the decision maker subjective model uncertainty, which in a self-confirming equilibrium must be consistent with the objective model uncertainty $\hat{\Sigma}_{a^*}(p^*)$ via the inclusion $\operatorname{supp}\mu^* \subseteq \hat{\Sigma}_{a^*}(p^*)$. Here we show that the value results that we just established for the subjective model uncertainty extend to objective model uncertainty. To this end, say that action a^* is *self-confirming* if there is a supporting belief μ^* such that $(a^*, \mu^*) \in \Gamma(p^*)$. By Proposition 10, the value of self-confirming actions a^* is $R(a^*, p^*)$, and so it is independent of the supporting belief. Define a (*equilibrium*) *action correspondence* $\gamma : \Sigma \to 2^A$ that associates to each possible true model p^* the associated self-confirming actions, that is, the actions a^* such that $(a^*, \mu^*) \in \Gamma(p^*)$ for some belief μ^*.

Example 14 Under full identification $\gamma(p^*) = \arg\max_{a \in A} R(a, p^*)$ for each $p^* \in \Sigma$. In the opposite case $\Sigma_a = \{\emptyset, \Sigma\}$, the set $\gamma(p^*)$ consists of all actions that are not strictly dominated. ▲

Next we show that self-confirming actions with better identification have higher values, regardless of which beliefs support them.

Proposition 15 *If $\hat{\Sigma}_{b^*}(p^*) \subseteq \hat{\Sigma}_{a^*}(p^*)$ then $R(b^*, \nu^*) \geq R(a^*, \mu^*)$ for all self-confirming actions $a^*, b^* \in \gamma(p^*)$.*

Proof It is enough to observe that in the integrals of the proof of Proposition 12 we can replace $\operatorname{supp}\mu^*$ and $\operatorname{supp}\nu^*$ with, respectively, $\hat{\Sigma}_{a^*}(p^*)$ and $\hat{\Sigma}_{b^*}(p^*)$. ■

Among other things, Proposition 15 implies that $V(a^*, \mu^*) = V(b^*, \nu^*)$ if $\hat{\Sigma}_{a^*}(p^*) = \hat{\Sigma}_{b^*}(p^*)$. 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.

We close with an example of self-confirming actions that have different identification properties.

Example 16 In Example 7 we saw that the actions b and y are self-confirming, with $\hat{\Sigma}_b(p^*) = \{p \in \Sigma : p(B) = 1/3\}$ and $\hat{\Sigma}_y(p^*) = \{p \in \Sigma : p(Y) = 2/3\}$. Hence, $\hat{\Sigma}_b(p^*) \neq \hat{\Sigma}_y(p^*)$, that is, the two self-confirming actions differ in their identification properties. In particular, $\hat{\Sigma}_b(p^*) \cap \hat{\Sigma}_y(p^*) = \{p^*\}$. They also have different values: $R(b, p^*) = 1/3 \neq 2/3 = R(y, p^*)$. ▲

5.3 Welfare loss

Let p^* be the true model. Since $V(a^*,\mu^*) = R(a^*,p^*) \leq \max_{a\in A} R(a,p^*)$, because of partial identification the decision makers experience a welfare loss

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

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

By Proposition 10, the loss $l(a^*,p^*)$ does not depend on either the feedback function f or the supporting belief μ^*. Only the true probability model p^* matters. Moreover, by Proposition 15

$$\hat{\Sigma}_{b^*}(p^*) \subseteq \hat{\Sigma}_{a^*}(p^*) \Longrightarrow l(b^*,p^*) \leq l(a^*,p^*) \qquad \forall a^*,b^* \in \gamma(p^*)$$

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 (and so optimal from a purely statistical viewpoint) to be self-confirming amounts to being objectively optimal. However, for an action that allows the decision maker to know the truth (or to get as close as possible to it), it has to be optimal in light of the decision maker's objective. Truth may be unpleasant.

Proposition 17 *Given a true model $p^* \in \Sigma$, suppose there is an action a such that $\hat{\Sigma}_a(p^*) \subseteq \hat{\Sigma}_{a'}(p^*)$ for each $a' \in A$. Then $a \in \gamma(p^*)$ if and only if it is objectively optimal.*

Proof The "if" is trivial since $a \in \arg\max_{a'\in A} R(a',p^*)$ implies $(a,\delta_{p^*}) \in \Gamma(p^*)$. As to the converse, suppose $(a,\mu^*) \in \Gamma(p^*)$. Since $\hat{\Sigma}_a(p^*) \subseteq \hat{\Sigma}_{a'}(p^*)$ for each $a' \in A$, by Lemma 11 for each such a' it is true that $R(a',p) = R(a',p^*)$ for all $p \in \hat{\Sigma}_a(p^*)$. Suppose to the contrary that there is $a' \in A$ such that $R(a',p^*) > R(a,p^*)$. Then, $V(a',\mu^*) = \int_{\hat{\Sigma}_a(p^*)} R(a',p)\,d\mu^*(p) > \int_{\hat{\Sigma}_a(p^*)} R(a,p)\,d\mu^*(p) = V(a,\mu^*)$, which contradicts the optimality of a. We conclude that $a \in \arg\max_{a'\in A} R(a',p^*)$. ■

Actions that are optimal from a statistical viewpoint may not be optimal from a decision theoretic one. The decision maker is not purely a statistician: he is not interested per se in discovering the true model unless the action that allows discovery is optimal in the decision problem. In this sense, there is no separation between estimation and decision in the present setup.

The next result further illustrates the point by showing that, regardless of the true model, fully revealing actions are self-confirming if and only if they are weakly dominant.[10] That is, unless they are weakly dominant, fully revealing actions may not be self-confirming.

Corollary 18 *Let a be a fully revealing action. Then $a \in \gamma(p)$ for each $p \in \Delta(S)$ if and only if a is weakly dominant.*

Proof We prove the "only if," the converse being trivial. Let a be a fully revealing action. By the previous result, $a \in \gamma(p)$ for each p if and only if $l(a,p) = 0$ for each p. Since S is finite, the latter condition implies $l(a,\delta_s) = 0$ for each $s \in S$, that is, $r(a,s) \geq r(a',s)$ for each $a' \in A$ and each $s \in S$. ■

The extent of the partial identification loss that arises when p^* is the true model is described by the set

$$L(p^*) = \{l(a^*,p^*) : a^* \in \gamma(p^*)\}$$

The set $L(p^*)$ is a singleton when $\gamma(p^*)$ is a singleton, that is, when there is a unique self-confirming action if p^* is the true model. Otherwise it is a set that accounts for the different losses that the different self-confirming actions in $\gamma(p^*)$ may cause.

[10] Action a is weakly dominant if, for each $a' \in A$, it is true that $r(a,s) \geq r(a',s)$ for each $s \in S$.

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 5, shows that the set of self-confirming pairs of actions and beliefs increases as ex post private information becomes coarser. The same behavior is featured by the partial identification loss.

Proposition 19 *Given decision problems (D, f) and $(D, \bar{f})$, if f is coarser than $\bar{f}$, then*

$$\bar{\Gamma}(p) \subseteq \Gamma(p) \text{ and } \bar{L}(p) \subseteq L(p) \qquad \forall p \in \Sigma$$

6 Phillips curve exploitation example

We now illustrate our machinery in the context of a 1970'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. 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.

6.1 Steady state model economies

We study a class $\theta \in \Theta$ of model economies at a (stochastic) steady state in which unemployment u and inflation π are affected by 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

$$u = \theta_0 + \theta_{1\pi}\pi + \theta_{1a}a + \theta_2 w \tag{11}$$

$$\pi = a + \theta_3\varepsilon \tag{12}$$

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 (11). Coefficients $\theta_{1\pi}$ and θ_{1a} are slope responses of unemployment to actual and planned inflation, 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 without policy interventions.[11]

Throughout the section we tacitly maintain the following assumption on 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 (non-trivial, 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

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

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 $(\boldsymbol{u}, \boldsymbol{\pi}) : A \times W \times E \times \Theta \to U \times \Pi$ defined by

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

[11] We identify nonintervention with the zero-target-inflation policy $a = 0$.

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 that, respectively, θ_{1a} and $\theta_{1\pi}$ quantify. 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 7.4).

Be that as it may be, in what follows $\Theta = \left\{\theta \in \mathbb{R}^5 : \theta_0 > 0,\ \theta_{1a} \leq -\theta_{1\pi},\ \theta_{1\pi} < 0,\ \theta_2 > 0 \text{ and } \theta_3 > 0\right\}$. Our monetary authority will pay special attention to the following two competing model economies.

6.1.1 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, unexpected part ε does. Up to unexpected shocks ε, the Phillips curve is vertical.

6.1.2 The Samuelson-Solow model

A second example of a 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 $\theta = (\theta_0, \theta_{1\pi}, 0, \theta_2, \theta_3)$, with $\theta_{1\pi} \neq 0$, and $\xi = (\theta_0, \theta_{1\pi}, \theta_{1\pi}\theta_3, \theta_2, \theta_3)$. In this Keynesian model the policy multiplier ξ_2 is strictly negative: monetary policies affect, at steady state, unemployment rates.

6.2 The policy problem: setup and identification

6.2.1 Setup

A 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.[12] The consequence space C consists of unemployment and inflation pairs (u, π), so we set $C = U \times \Pi \subseteq \mathbb{R}^2$. The outcome function $\rho : A \times S \to C$ is

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

[12] Section 7.1 further discusses the Cartesian structure of the state space.

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. Its matrix representation is

$$\rho(a, w, \varepsilon, \theta) = \begin{bmatrix} \theta_0 \\ 0 \end{bmatrix} + a \begin{bmatrix} \theta_{1\pi} + \theta_{1a} \\ 1 \end{bmatrix} + \begin{bmatrix} \theta_2 & \theta_{1\pi}\theta_3 \\ 0 & \theta_3 \end{bmatrix} \begin{bmatrix} w \\ \varepsilon \end{bmatrix} \tag{13}$$

Finally, given a reward function $r : A \times S \to \mathbb{R}$, the objective function of the monetary authority is

$$V(a, \mu) = \int_\Sigma R(a, p)\, d\mu(p) = \int_\Sigma \left(\int_{W \times E \times \Theta} r(a, w, \varepsilon, \theta)\, dp(w, \varepsilon, \theta) \right) d\mu(p) \tag{14}$$

6.2.2 Factorization

We assume that the monetary authority observes consequences, so that $f = \rho$. 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 $\rho_a^{-1} : C \to 2^{W \times E \times \Theta}$. In particular,

$$\begin{aligned} \hat{f}_a(p)(u, \pi) &= p\left(\rho_a^{-1}(u, \pi)\right) = p\left((w, \varepsilon, \theta) : \rho(a, w, \varepsilon, \theta) = (u, \pi)\right) \\ &= p\left((w, \varepsilon, \theta) : \boldsymbol{u}(a, w, \varepsilon, \theta) = u \text{ and } \boldsymbol{\pi}(a, w, \varepsilon, \theta) = \pi\right) \end{aligned} \tag{15}$$

and

$$\hat{\Sigma}_a(p) = \left\{ p' \in \Sigma : \hat{f}_a(p') = \hat{f}_a(p) \right\} \qquad \forall p \in \Sigma \tag{16}$$

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. This suggests factoring the probability models $p \in \Sigma \subseteq \Delta(W \times E \times \Theta)$ as

$$p(w, \varepsilon, \theta) = q(w, \varepsilon) \times \delta_{\bar{\theta}}(\theta) \tag{17}$$

where $q \in \Delta(W \times E)$ is known and $\delta_{\bar{\theta}} \in \Delta(\Theta)$ is a Dirac probability measure concentrated on a given economic model $\bar{\theta} \in \Theta$. We can indicate a probability model with $\bar{\theta}$ and write $p_{\bar{\theta}}$.

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). The resulting factorization (17) 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 p_{θ^*}); (ii) since q is known, it allows us to regard Σ as a subset of Θ, and so to define the prior μ on θ.[13]

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

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

In words, shocks are assumed to be uncorrelated and suitably normalized.

6.2.3 Identification

In this richer "factorized" setup, through the inverse correspondence $\rho_{a,\theta}^{-1} : C \to 2^{W \times E}$ defined by

$$\rho_{a,\theta}^{-1}(u, \pi) = \{(w, \varepsilon) : \rho(a, w, \varepsilon, \theta) = (u, w)\}$$

[13] Otherwise, μ should be defined on pairs (q, θ), i.e., $\mu \in \Delta(\Delta(W \times E) \times \Theta)$.

the monetary authority can infer combinations of shocks (w, ε) that model θ says might have occurred when policy a has been chosen and message (u, π) has been received. Since $\Sigma \subseteq \Theta$, the pushforward map (15) and the partial identification correspondence (16) are now given by

$$\hat{f}_a(\theta)(u, \pi) = q\left(\rho_a^{-1}(u, \pi)\right) = q\left((w, \varepsilon) : \rho(a, w, \varepsilon, \theta) = (u, \pi)\right)$$

and

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

Since $\theta_2 \neq 0$ (Assumption 1), the function $\rho_{a,\theta}$ is easily seen to be injective, so that the inverse correspondence $\rho_{a,\theta}^{-1}$ is actually a function: given model θ and policy a, messages (u, π) reveal which shocks (w, ε) occurred.

A sharp identification result (proved in the Appendix) holds.

Proposition 20 *The partial identification correspondence $\hat{\Sigma}_a : \Theta \to 2^{\Theta}$ is given by*

$$\hat{\Sigma}_a(\theta) = \left\{\theta' \in \Theta : \theta_0 + \theta_{1a} a = \theta'_0 + \theta'_{1a} a,\ \theta'_{1\pi} = \theta_{1\pi},\ \theta'_2 = \theta_2,\ \theta'_3 = \theta_3\right\} \tag{18}$$

Given a true model θ, the shocks 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. This important identification is made possible by some moment conditions, formally spelled out in the proof. We can, however, heuristically describe them via the bivariate random variable $(\boldsymbol{u}_a, \boldsymbol{\pi}_a) : W \times E \times \Theta \to C$ that, for a given policy a, represents the unemployment and inflation rates determined by the state (w, ε, θ).[14] The monetary authority observes in the long run the following moments:

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

Therefore,

$$\theta_{1\pi} = \frac{\mathrm{Cov}_\theta(\boldsymbol{u}_a, \boldsymbol{\pi}_a)}{\mathrm{Var}_\theta(\boldsymbol{\pi}_a)} \tag{19}$$

is the beta coefficient of the Phillips regression of unemployment over inflation,[15]

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

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

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

$$\theta_0 + \theta_{1a} a = \mathbb{E}_\theta(\boldsymbol{u}_a) - \frac{\mathrm{Cov}_\theta(\boldsymbol{u}_a, \boldsymbol{\pi}_a)}{\mathrm{Var}_\theta(\boldsymbol{\pi}_a)} \mathbb{E}_\theta(\boldsymbol{\pi}_a) \tag{20}$$

where the right hand side is the alpha coefficient of the Phillips regression. In the long run the alpha is observed by the monetary authority, but such observation depends on the policy a that the authority chose. In other words, the estimated value of the intercept of the Phillips regression depends on the authority behavior: what the authority learned depends on what it did.

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

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

6.2.4 Estimated model economy

The estimated 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 $\theta_{1\pi}$,[16] by $\hat{\sigma}_{\boldsymbol{u}|\boldsymbol{\pi}}$ the residual variance 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 $\theta_0 + \theta_{1a}a$ depends on policy a; for this reason, we denote it by $\hat{\alpha}(a)$.

Using this notation we can write

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

As a result, the long run estimated version of the model economy (11)-(12) that the monetary authority considers is

$$u = \hat{\alpha}(a) + \hat{\beta}\pi + \hat{\sigma}_{\boldsymbol{u}|\boldsymbol{\pi}}w \tag{21}$$

$$\pi = a + \hat{\sigma}_{\boldsymbol{\pi}}\varepsilon \tag{22}$$

$$\hat{\alpha}(a) = \theta_0 + \theta_{1a}a \tag{23}$$

In particular, (21) is the estimated aggregate supply equation and (22) is the estimated inflation equation. The intercept of the former equation depends on the policy a via the equality (23), which only partly identifies the two coefficients θ_0 and θ_{1a}. Next we address this key partial identification issue.

6.2.5 Partial identification line

The monetary authority is still 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, $\theta_0 = \mathbb{E}_\theta(\boldsymbol{u}_0)$, which would be observed if $a = 0$ were chosen; the latter is the "direct" impact of policy on unemployment.

Because of these two remaining unidentified coefficients, without loss of generality in what follows we can consider as parameter space $\tilde{\Theta} = \{(\theta_0, \theta_{1a})\} = \mathbb{R}^2$. The parameter space is now the plane. By (18), the partial identification correspondence $\hat{\Sigma}_a : \tilde{\Theta} \to 2^{\tilde{\Theta}}$ becomes

$$\hat{\Sigma}_a(\theta) = \left\{(\theta'_0, \theta'_{1a}) \in \mathbb{R}^2 : \theta'_0 = -\theta'_{1a}a + \theta_0 + \theta_{1a}a\right\} \tag{24}$$

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

Given a true model $\theta = (\theta_0, \theta_{1a})$, as the chosen policy a varies 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 $-a$. In each such line there is a unique Lucas-Sargent model, characterized by $\theta'_{1a} = -\hat{\beta}$, as well as a unique Samuelson-Solow model, characterized by $\theta'_{1a} = 0$. In other words,

[16] By Assumption 1, 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.

partial identification lines feature a unique specimen of each class of models.

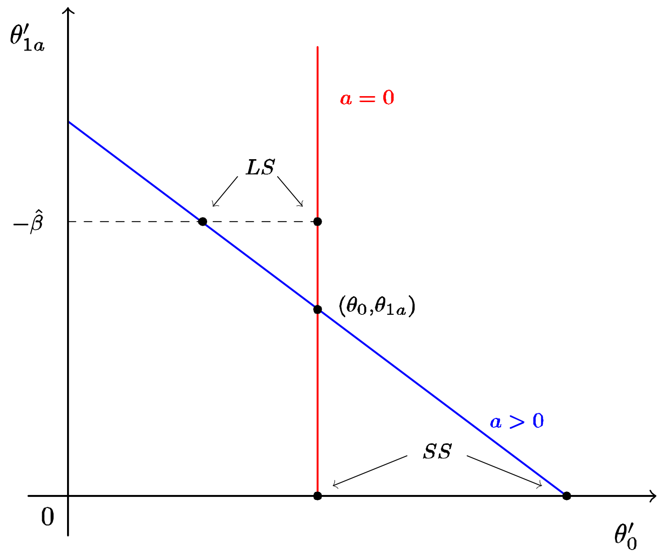

The figure illustrates the previous analysis. In particular, LS stands for Lucas-Sargent model and SS for Samuelson-Solow model, while the red (resp., blue) line is the partial identification line that correspond to policy $a = 0$ (resp., $a > 0$)

6.3 The policy problem: value and equilibria

6.3.1 Value

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

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

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

Lemma 21 *$R(a, \theta) = v(\mathbb{E}_\theta(\boldsymbol{u}_a), \mathbb{E}_\theta(\boldsymbol{\pi}_a)) + cost. \qquad \forall (\theta, a) \in \tilde{\Theta} \times A.$*

The linear quadratic framework thus allows us to express the expected reward as the utility of expectations. As a result, the objective function (14) becomes

$$V(a, \mu) = -\int_\Theta v(\mathbb{E}_\theta(\boldsymbol{u}_a)), \mathbb{E}_\theta(\boldsymbol{\pi}_a))\, d\mu(\theta) + cost. \tag{25}$$

Proof Some simple algebra shows that

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

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

6.3.2 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 $\tilde{\Theta}$. Under this notation, the partial identification line is

$$\hat{\Sigma}_a(\theta^*) = \left\{(\theta_0, \theta_{1a}) \in \mathbb{R}^2 : \theta_0 = \theta_0^* + (\theta_{1a}^* - \theta_{1a})\, a\right\}$$

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

$$\operatorname{supp} \mu^* \subseteq \hat{\Sigma}_{a^*}(\theta^*) = \left\{(\theta_0, \theta_{1a}) \in \mathbb{R}^2 : \theta_0 = \theta_0^* + (\theta_{1a}^* - \theta_{1a})\, a^*\right\}$$

Next we characterize self-confirming equilibria.

Proposition 22 *A policy and belief pair $(a^*, \mu^*) \in A \times \Delta(\tilde{\Theta})$ is self-confirming if and only if*[17]

$$a^* = -\frac{\theta_0^* \left(\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1a})\right)}{1 + \left(\hat{\beta}^* + \theta_{1a}^*\right)\left(\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1a})\right)} \tag{26}$$

and

$$\mu^*\left(\left\{(\theta_0, \theta_{1a}) \in \mathbb{R}^2 : \theta_0 = \theta_0^* - \frac{\theta_0^* \left(\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1a})\right)}{1 + \left(\hat{\beta}^* + \theta_{1a}^*\right)\left(\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1a})\right)}(\theta_{1a}^* - \theta_{1a})\right\}\right) = 1 \tag{27}$$

The proof of the result is in the Appendix. Here we can heuristically derive it in the special case of dogmatic beliefs. By (25), up to a constant the monetary authority value function is

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

It depends only on the expected values of inflation and unemployment according to the authority beliefs. Suppose, by way of example, that such beliefs are dogmatic: there is some model economy $\bar{\theta} \in \tilde{\Theta}$ such that $\mu = \delta_{\bar{\theta}}$.[18] For instance, a new-classical authority that believes that there are no systematically exploitable trade-offs between inflation and unemployment sets $\bar{\theta}_{1a} = -\hat{\beta}^*$ (and so the conjectured policy multiplier $\bar{\xi}_2$ is zero). In contrast, a Keynesian authority that believes in such trade-offs may set, for instance, $\bar{\theta}_{1a} = 0$ (the conjectured multiplier $\bar{\xi}_2 = \hat{\beta}^*$ is then strictly negative).

Based on the estimated model economy (21)-(23), a dogmatic authority conjectures that, according to the chosen policy a, the expected values of inflation and unemployment are constrained by the equation

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

This conjectured constraint is the version of the estimated aggregate supply equation (21) that the authority expects to face systematically given its dogmatic belief. If such belief happens to be correct, $\bar{\theta} = \theta^*$, the conjectured constraint is the true one.

The authority decision problem is

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

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

[18] Recall that $\delta_{\bar{\theta}}(\theta) = 1$ if and only if $\theta = \bar{\theta}$.

The Lagrangian is

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

and so the first order conditions are

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

By solving them we get

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

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

$$a^* = B\left(\bar{\theta}\right) \qquad \text{(subjective best reply)} \tag{28}$$

and

$$\bar{\theta}_0 = \theta^*_0 + \left(\theta^*_{1a} - \bar{\theta}_{1a}\right)a^* \qquad \text{(confirmed beliefs)} \tag{29}$$

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

$$a^* = -\frac{\theta^*_0\left(\hat{\beta}^* + \bar{\theta}_{1a}\right)}{1 + \left(\hat{\beta}^* + \theta^*_{1a}\right)\left(\hat{\beta}^* + \bar{\theta}_{1a}\right)} \tag{30}$$

and

$$\bar{\theta}_0 = \theta^*_0 - \frac{\theta^*_0\left(\hat{\beta}^* + \bar{\theta}_{1a}\right)}{1 + \left(\hat{\beta}^* + \theta^*_{1a}\right)\left(\hat{\beta}^* + \bar{\theta}_{1a}\right)}\left(\theta^*_{1a} - \bar{\theta}_{1a}\right) \tag{31}$$

which are the equilibrium relations (26) and (27) in the case of dogmatic beliefs.[19]

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 $\mathbb{E}_{\theta^*}(\boldsymbol{u}_a) = \theta^*_0$ 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 $(\mathbb{E}_{\bar{\theta}}(\boldsymbol{u}_a), a)$. This dogmatic belief is confirmed, and so condition (29) is satisfied, if $\mathbb{E}_{\bar{\theta}}(\boldsymbol{u}_a) = \theta^*_0$, 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 (28) 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

[19] Note 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 suggests a certainty equivalence principle for our linear quadratic setup, which will be spelled out in Section 7.3.

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

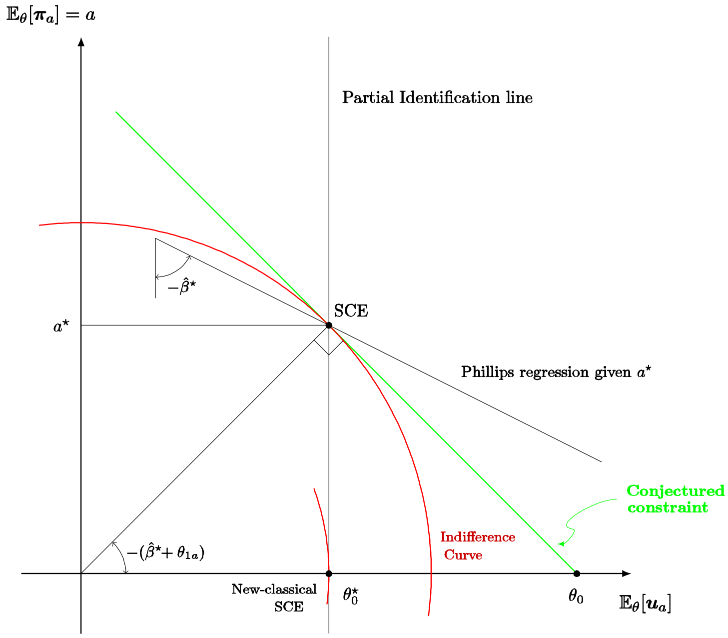

When the dogmatic belief is such that $\bar{\theta}_{1a} = 0$, and so $\bar{\xi}_2 = -\hat{\beta}^*$ becomes the conjectured multiplier, the monetary authority is Keynesian and the figure becomes:

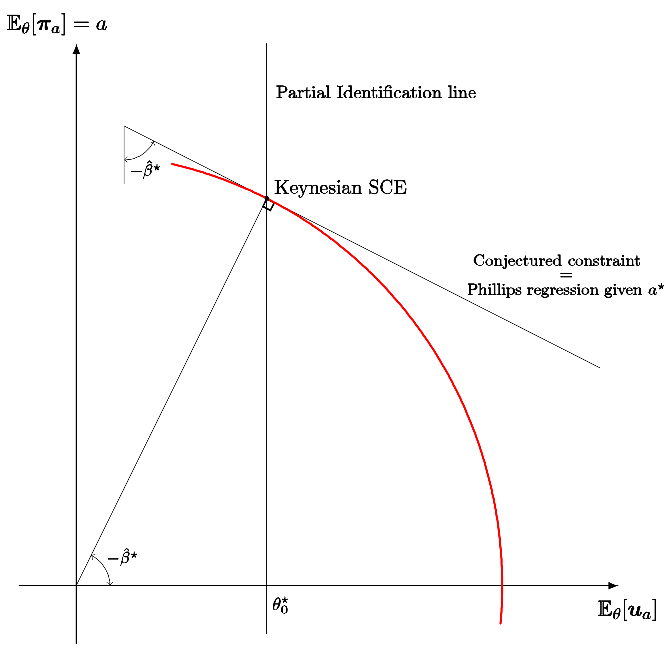

The conjectured constraint is $E_{\bar{\theta}}(\boldsymbol{u}_a) = \bar{\theta}_0 + \hat{\beta}^* E_{\bar{\theta}}(\boldsymbol{\pi}_a)$. Its slope is the beta coefficient of the Phillips regression, which represents the trade-off between inflation and unemployment that the Keynesian authority now believes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

6.3.3 Policy activism and welfare

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

To this end we need to consider the policy multiplier $\xi_2 = \theta_{1\pi} + \theta_{1a}$ that quantifies the impact of planned inflation on unemployment (Section 6.1). Given an equilibrium belief μ^*, its expected value is $\mathbb{E}_{\mu^*}(\xi_2) = \hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1a})$. The authority underestimates the multiplier when $\mathbb{E}_{\mu^*}(\xi_2) > \xi_2^*$ and overestimates it when $\mathbb{E}_{\mu^*}(\xi_2) < \xi_2^*$.[20] Since the authority learned $\hat{\beta}^*$, it holds $\mathbb{E}_{\mu^*}(\xi_2) \gtreqless \xi_2^*$ if and only if $\mathbb{E}_{\mu^*}(\theta_{1a}) \gtreqless \theta_{1a}^*$. For instance, when θ_{1a}^* and $\mathbb{E}_{\mu^*}(\theta_{1a})$ are positive this means that the multiplier is under/overestimated if and only if the direct impact of planned inflation on unemployment is over/underestimated.

The objectively optimal policy is

$$a^o = -\frac{\theta_0^* \left(\hat{\beta}^* + \theta_{1a}^*\right)}{1 + \left(\hat{\beta}^* + \theta_{1a}^*\right)^2} \tag{32}$$

It is immediate to see that $a^* = a^o$ if and only if $\mathbb{E}_{\mu^*}(\theta_{1a}) = \theta_{1a}^*$. The equilibrium action is objectively optimal when the monetary authority has a correct expected value of the direct impact of planned inflation on unemployment (and so of the policy multiplier). More generally, next we show that policy hyperactivism characterizes authorities that overestimate the policy multiplier, while hypoactivism characterizes authorities that underestimate such impact.[21]

Proposition 23 *Consider a self-confirming equilibrium (a^*, μ^*).*

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

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

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

(iv) $\mathbb{E}_{\mu^*}(\theta_{1a}) = -\hat{\beta}^*$ *if and only if policy a^* is zero-target-inflation, i.e., $a^* = 0$.*

For the monetary authority both kind of deviations from objective optimality, hyperactivism and hypoactivism, cause the same welfare loss. In fact, the welfare loss has, in general, the following form.

Proposition 24 *It holds $l(a^*, \theta^*) = (1 + (\hat{\beta}^* + \theta_{1a}^*)^2)(a^* - a^o)^2$.*

In the next section we will illustrate this result (proved in the Appendix) with a few examples.

6.4 Policy dogmatism

6.4.1 Equilibria

Assume that the monetary authority has dogmatic equilibrium beliefs $\mu^* = \delta_{\bar{\theta}}$. By Proposition 22, a pair $(a^*, \delta_{\bar{\theta}}) \in A \times \Delta(\tilde{\Theta})$ is self-confirming if and only if it satisfies relations (30) and (31), which we also heuristically derived. Two special cases are noteworthy.

[20] Both ξ_2^* and $\mathbb{E}_{\mu^*}(\xi_2)$ are negative (Assumption 2), and so $\mathbb{E}_{\mu^*}(\xi_2) \gtreqless \xi_2^*$ if and only $\left|\mathbb{E}_{\mu^*}(\xi_2)\right| \lesseqgtr |\xi_2^*|$.

[21] The result is proved in the Appendix. Being $\xi_2^* \leq 0$ (Assumption 2), the four cases it considers exhaust all possibilities. Section 7.4 will generalize the result to the case when Assumption 2 does not hold. Note that, being $\mathbb{E}_{\mu^*}(\theta_{1a}) \leq -\hat{\beta}^*$, in (iv) it holds $\mathbb{E}_{\mu^*}(\theta_{1a}) = -\hat{\beta}^*$ if and only if $\mu^*(\theta_{1a} = -\hat{\beta}^*) = 1$, i.e., $\mu^* = \delta_{(\theta_0, -\hat{\beta}^*)}$.

(i) When $\bar{\theta}_{1a} = -\hat{\beta}^*$, and so 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 23-(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 $\bar{\theta}_{1a} = 0$, and so the conjectured policy multiplier is strictly negative, we obtain the self-confirming equilibrium

$$a^* = -\frac{\theta_0^* \hat{\beta}^*}{1 + \hat{\beta}^* \left(\hat{\beta}^* + \theta_{1a}^*\right)} \quad \text{and} \quad \bar{\theta} = \left(\theta_0^* \left(\frac{1 + \hat{\beta}^{*2}}{1 + \hat{\beta}^* \left(\hat{\beta}^* + \theta_{1a}^*\right)}\right), 0\right) \tag{33}$$

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 23 such policy is hyperactive if $\theta_{1a}^* > 0$, hypoactive if $\theta_{1a}^* < 0$, objectively optimal if $\theta_{1a}^* = 0$.

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.

6.4.2 A new-classical world

So far we did not posit any true economic model. 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 $a^* = -\theta_0^*(\hat{\beta}^* + \bar{\theta}_{1a})$ and $\bar{\theta}_0 = \theta_0^*(1 - (\hat{\beta}^* + \bar{\theta}_{1a})^2)$. Hence, the policy and belief pair

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

is the dogmatic self-confirming equilibrium in a Lucas-Sargent model economy. By Proposition 23 policy a^* is hyperactive when $\bar{\theta}_{1a} < \theta_{1a}^*$ and objectively optimal when $\bar{\theta}_{1a} = \theta_{1a}^*$. The welfare loss is $l(a^*, \theta^*) = \theta_0^{*2}(\hat{\beta}^* + \bar{\theta}_{1a})^2$.

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

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

$$\left(0, \delta_{(\theta_0^*, -\hat{\beta}^*)}\right) \tag{34}$$

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 $\left(a^*, \delta_{(\bar{\theta}_0,0)}\right)$ is a self-confirming equilibrium if and only if $a^* = -\theta_0^* \hat{\beta}^*$ and $\bar{\theta}_0 = \theta_0^*(1 - \hat{\beta}^{*2})$. The policy and belief pair

$$\left(-\theta_0^* \hat{\beta}^*, \delta_{\left(\theta_0^*\left(1-\hat{\beta}^{*2}\right),0\right)}\right) \tag{35}$$

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

6.4.3 Welfare

What are the welfare implications of dogmatism? By way of example, we consider a new-classical authority in a Keynesian economy, as well as a Keynesian authority in a new-classical economy. The loss of a new-classical zero inflation policy in a Keynesian economy, with $\theta_{1a}^* = 0$, is $(\theta_0^* \hat{\beta}^*)^2/(1+\hat{\beta}^{*2})$, while the loss of a Keynesian nonzero inflation policy (33) in a new-classical economy, with $\theta_{1a}^* = -\hat{\beta}^*$, is $(\theta_0^* \hat{\beta}^*)^2$.

Hence, a mistaken new-classical authority has a lower welfare loss than a mistaken Keynesian one: doing nothing (new-classical nonintervention) when would be optimal to do something (Keynesian activism) is less harmful than the opposite. In particular, this welfare gap increases in $\hat{\beta}^*$, that is, in the apparent trade-off between inflation and unemployment that a Keynesian authority would believe in a new-classical economy.

6.5 Policy secularism

6.5.1 Equilibria

Suppose that the monetary authority is not dogmatic, but has 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 $(\theta_0^{ss}, 0)$. If $\mu_k^* \in [0,1]$ is the belief in the latter model, we can write belief μ^* as

$$\mu^* = (1 - \mu_k^*)\, \delta_{\left(\theta_0^{ls}, -\hat{\beta}^*\right)} + \mu_k^* \delta_{\left(\theta_0^{ss}, 0\right)} \tag{36}$$

Since $\mathbb{E}_{\mu^*}(\theta_{1a}) = -(1 - \mu_k^*)\hat{\beta}^*$, by Proposition 22 the pair (a^*, μ^*) is a self-confirming equilibrium if and only if

$$a^* = -\frac{\theta_0^* \hat{\beta}^* \mu_k^*}{1 + \hat{\beta}^* \mu_k^* \left(\hat{\beta}^* + \theta_{1a}^*\right)} \tag{37}$$

and

$$\theta_0^{ls} = \frac{\theta_0^*}{1 + \hat{\beta}^* \left(\hat{\beta}^* + \theta_{1a}^*\right) \mu_k^*} \quad ; \quad \theta_0^{ss} = \frac{\theta_0^* \left(1 + \hat{\beta}^{*2} \mu_k^*\right)}{1 + \hat{\beta}^* \left(\hat{\beta}^* + \theta_{1a}^*\right) \mu_k^*} \tag{38}$$

As a result, in this case the pair

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

is a self-confirming equilibrium for every $\mu_k^* \in [0,1]$. We thus have a continuum of equilibria parametrized by the belief μ_k^* in the Samuelson-Solow model. 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 6.4.1). More generally, the equilibrium policy is hyperactive when $\mu_k^* > -\left(\hat{\beta}^* + \theta_{1a}^*\right)/\hat{\beta}^*$, hypoactive when $0 < \mu_k^* < -\left(\hat{\beta}^* + \theta_{1a}^*\right)/\hat{\beta}^*$ (Proposition 23).

In equilibrium the coefficients (38) of the Lucas-Sargent and Samuelson-Solow models depend on the authority belief μ_k^*: different such beliefs correspond to different Lucas-Sargent and Samuelson-Solow equilibrium specifications. Though the support of the equilibrium belief (36) always contains a specimen of both classes of model economies, such specimen changes as belief μ_k^* changes.

Finally, the welfare loss is

$$l\left(a^*, \theta^*\right) = \frac{\theta_0^{*2}\left(\hat{\beta}^* \mu_k^* + \hat{\beta}^* + \theta_{1a}^*\right)^2}{\left(1 + \hat{\beta}^* \mu_k^* \left(\hat{\beta}^* + \theta_{1a}^*\right)\right)^2 \left(1 + \left(\hat{\beta}^* + \theta_{1a}^*\right)^2\right)} \tag{39}$$

6.5.2 A new-classical world

To further study two-models beliefs, let us posit a true model. As we did in the study of dogmatism, assume that a Lucas-Sargent model economy $\theta^* = (\theta_0^*, -\hat{\beta}^*)$ is the true model. If so, by (37) and (38) the pair (a^*, μ^*) is a self-confirming equilibrium if and only if $a^* = -\theta_0^* \hat{\beta}^* \mu_k^*$, $\theta_0^{ls} = \theta_0^*$ and $\theta_0^{ss} = \theta_0^*(1 + \hat{\beta}^{*2} \mu_k^*)$. Hence, in this case the pair

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

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

Again, 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 back to the dogmatic new-classical equilibrium (34), while if $\mu_k^* = 1$ we get back to the dogmatic Keynesian equilibrium (35). 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 $\theta_0 = \theta_0^*(1 - \hat{\beta}^{*2}\mu_k^*))$.

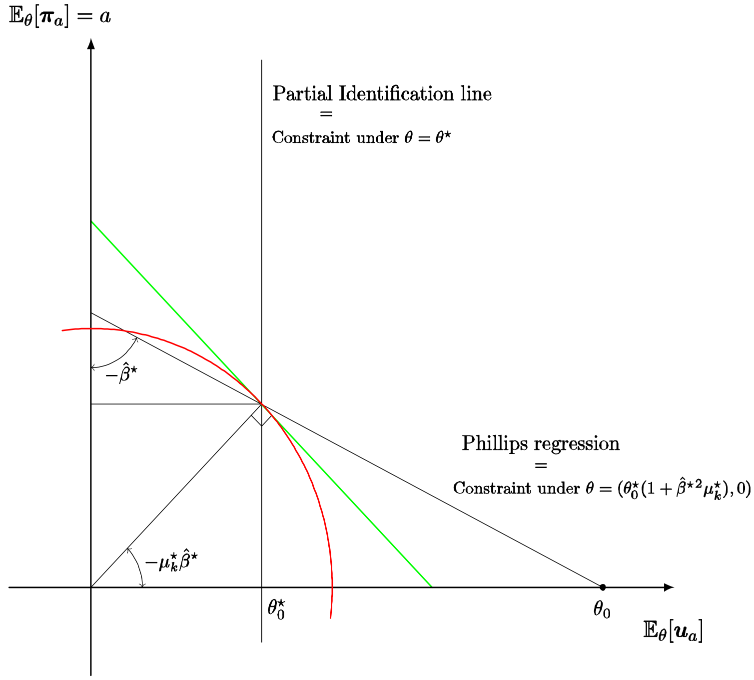

The second figure gives an alternative geometrical representation. Every policy a induces a pair of (objectively) expected rewards, the reward under model θ^*, $R(a, \theta^*)$, and the reward under model θ, $R(a, \theta)$. By changing a one obtains the locus of possible pair of rewards. If $R(a, \theta^*) \neq R(a, \theta)$, the monetary authority can infer which one of the two model 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 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 $(1-\mu_k^*)/\mu_k^*$.

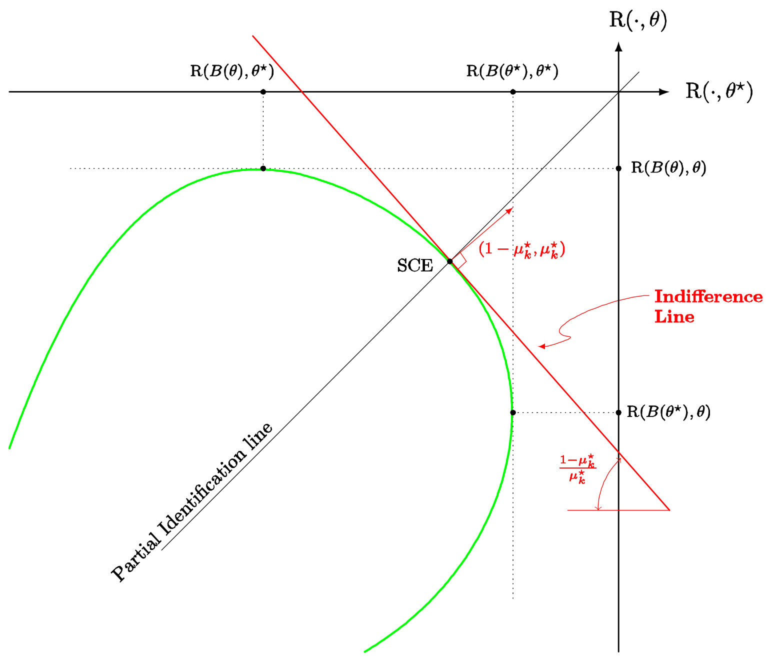

7 Discussion

7.1 Economic and stochastic model uncertainties

In the Phillips curve example we posited a state $s = (w, \varepsilon, \theta)$ in which we distinguish a shock pair (w, ε) and a model economy θ. The key factorization (17) builds on such distinction. A deeper distinction, however, holds. Policy problems feature two types of model uncertainties. First, there is uncertainty about the economics of the economic phenomenon studied, and so about the economic models that are actually able to explain it. Second, there is uncertainty about the statistical performance of such economic models, due to measurement errors and unexplained variation due to (possibly many) minor explanatory variables that the policy maker is "unable and unwilling to specify", as Marschak (1953) p. 12 remarks.

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.[22] This is why, by assuming q known, in the Phillips curve example we focused on economic model uncertainty.

In general, state spaces relevant for policy problems can be represented as Cartesian products of factors that represent the two types of uncertainty. If so, versions of the factorization (17) would apply.

7.2 Equilibrium in beliefs

A fixed point standpoint When best replies to beliefs are unique we can equivalently state self-confirming equilibria as equilibria in beliefs. Formally, say that a decision problem $D = (A, S, C, \Sigma, \rho, v)$ is *nice* if there is a best reply function $B : \Delta(\Sigma) \to A$ such that $B(\mu) = \arg\max_{a \in A} V(a, \mu)$ for each $\mu \in \Delta(\Sigma)$.

[22] A main instance of what Denzau and North (1994) and, more recently, Rodrik (2014) call "ideas".

In a nice decision problem, given a true model $p^* \in \Sigma$, 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

$$\mu^* \in \Delta\left(\hat{\Sigma}_{B(\mu^*)}(p^*)\right) \tag{40}$$

Hence, there is a unique equilibrium condition (40), which is in terms of beliefs (actions being pinned down by the best reply function). For this reason in nice problems we can view self-confirming equilibria as fixed points of beliefs: if we define the self-correspondence $\eta : \Delta(\Sigma) \to 2^{\Delta(\Sigma)}$ by $\eta(\mu) = \Delta\left(\hat{\Sigma}_{B(\mu)}(p^*)\right)$, we can say that belief $\mu^* \in \Delta(\Sigma)$ is self-confirming if and only if

$$\mu^* \in \eta(\mu^*) .$$

As a result, in nice problems the loss function can be defined in terms of beliefs by setting $l(\mu, p) = l(B(\mu), p)$.

Phillips curve Because of the quadratic objective function, the Phillips curve decision problem is nice. Hence, a belief $\mu^* \in \Delta(\Sigma)$ is self-confirming equilibrium if and only if $\mu^* \in \eta(\mu^*)$, where $\eta : \Delta(\Sigma) \to 2^{\Delta(\Sigma)}$ is given by $\eta(\mu) = \Delta(\hat{\Sigma}_{B(\mu)}(\theta^*))$. In particular, a dogmatic belief $\bar{\theta} \in \Sigma$ is self-confirming if and only if $\delta_{\bar{\theta}} \in \eta(\delta_{\bar{\theta}})$, that is, if and only if $\bar{\theta} \in \hat{\Sigma}_{B(\bar{\theta})}(\theta^*)$.

Finally, in the Appendix we show that

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

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

Conclusion This fixed point view of self-confirming equilibria can be natural in some applications, though it should be viewed within the general theoretical framework presented in Section 5.

7.3 A certainty equivalence principle

In our equilibrium result, Proposition 22, the expectation $\mathbb{E}_{\mu^*}(\theta_{1a})$ plays a key role. In particular:

(i) The equilibrium partial identification line $\hat{\Sigma}_{a^*}(\theta^*)$ depends on equilibrium beliefs only via the expectation $\mathbb{E}_{\mu^*}(\theta_{1a})$.

(ii) $(\mathbb{E}_{\mu^*}(\theta_0), \mathbb{E}_{\mu^*}(\theta_{1a})) \in \hat{\Sigma}_{a^*}(\theta^*)$, that is, the subjective expected values of the coefficients are observationally equivalent to their true values. Moreover, it holds $\mathbb{E}_{\mu^*}(\theta_0) = \theta_0^*$ if and only if $\mathbb{E}_{\mu^*}(\theta_{1a}) = \theta_{1a}^*$, and so both subjective expected values are correct if and only if either is.

By (i), given a self-confirming policy and belief pair (a^*, μ^*), another belief ν^* forms an action-equivalent self-confirming pair (a^*, ν^*) if and only if $\mathbb{E}_{\nu^*}(\theta_{1a}) = \mathbb{E}_{\mu^*}(\theta_{1a})$ and $\operatorname{supp} \nu^* \subseteq \hat{\Sigma}_{a^*}(\theta^*)$. The latter requirement automatically holds if $\nu^* \sim \mu^*$, that is, if the beliefs are equivalent. Such beliefs thus induce the same equilibrium policy as long as they share the same expected value for the coefficient θ_{1a}.

Jointly, (i) and (ii) imply the following lemma.

Lemma 25 *A pair (a^*, μ^*) is a self-confirming equilibrium if and only if the pair $(a^*, \delta_{\bar{\theta}})$ is, with $\bar{\theta} = (\mathbb{E}_{\mu^*}(\theta_0), \mathbb{E}_{\mu^*}(\theta_{1a}))$.*

In words, the dogmatic belief concentrated on the expected values of the coefficients induces the same equilibrium policy than the belief μ^* itself. Thus, every self-confirming equilibrium has a companion, action equivalent, equilibrium with dogmatic beliefs. It is a form of the certainty equivalence principle in the linear quadratic setup of Section 6.

7.4 Beyond Assumption 2

The only result of Section 6 that depends on Assumption 2 is Proposition 23. Without such assumption, the result takes the following more general form.

Proposition 26 *Consider a self-confirming equilibrium* (a^*, μ^*)*.*

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

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

8 Appendix

8.1 Lemma 1

For completeness we provide a proof of this known lemma. Since $\mathcal{S}$ and $\mathcal{Y}$ are Borel σ-algebras determined by Polish spaces, they contain all singletons and the function φ is bimeasurable.

Lemma 27 *For each* $s \in S$, $\hat{\varphi}(\delta_s) = \delta_{\varphi(s)}$.

Proof. Fix $s \in S$. For each $B \in \mathcal{Y}$,

$$\hat{\varphi}(\delta_s)(B) = \delta_s\left(\varphi^{-1}(B)\right) = \begin{cases} 1 & s \in \varphi^{-1}(B) \\ 0 & s \notin \varphi^{-1}(B) \end{cases} = \begin{cases} 1 & \varphi(s) \in B \\ 0 & \varphi(s) \notin B \end{cases} = \delta_{\varphi(s)}(B),$$

proving the statement. ■

Proof of Lemma 1 "If:" Let $p, q \in \Delta(S)$. If $p \neq q$ then there is $E \in \mathcal{S}$ such that $p(E) \neq q(E)$. Since φ is bimeasurable, it follows that $B = \varphi(E) \in \mathcal{Y}$. Since φ is injective, $\hat{\varphi}(p)(B) = p(E) \neq q(E) = \hat{\varphi}(q)(B)$. Hence, $\hat{\varphi}(p) \neq \hat{\varphi}(q)$, i.e., $\hat{\varphi}$ is injective. Next, consider $\mu \in \Delta(Y)$. Define $p : \mathcal{S} \to [0,1]$ by $p(E) = \mu(\varphi(E))$ for each $E \in \mathcal{S}$. Since φ is bimeasurable, μ is well defined and belongs to $\Delta(S)$. It follows that $\hat{\varphi}(p)(B) = p\left(\varphi^{-1}(B)\right) = \mu(B)$ for each $B \in \mathcal{Y}$, and so $\hat{\varphi}$ is surjective.

"Only if:" Consider $s, s' \in S$. By Lemma 27, $s \neq s'$ implies $\delta_{\varphi(s)} = \hat{\varphi}(\delta_s) \neq \hat{\varphi}(\delta_{s'}) = \delta_{\varphi(s')}$. Hence $\varphi(s) \neq \varphi(s')$, i.e., φ is injective. Finally, consider $y \in Y$. Since $\hat{\varphi}$ is surjective, there is $p \in \Delta(S)$ such that $\hat{\varphi}(p) = \delta_y$, i.e., $\delta_y(B) = \hat{\varphi}(p)(B) = p\left(\varphi^{-1}(B)\right)$ for each $B \in \mathcal{Y}$. Since φ is injective, $\varphi^{-1}(\{y\})$ is either empty or a singleton. Since all singletons are measurable, $1 = \delta_y(\{y\}) = p\left(\varphi^{-1}(\{y\})\right)$, yielding that $\varphi^{-1}(\{y\}) \neq \emptyset$, i.e., that φ is surjective. ■

8.2 Proposition 20

First recall that a is fixed. Some simple algebra shows that, for $\theta \in \Theta$, the set $\hat{\Sigma}_a(\theta)$ is

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

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

$$\int_S h(\rho_a)\, d(q \times \delta_{\theta'}) = \int_S h(\rho_a)\, d(q \times \delta_\theta) \quad \forall h : \mathbb{R} \times \mathbb{R} \to \mathbb{R} \tag{42}$$

Next observe that

1. For $h(u,\pi)=\pi$ and $\theta''\in\Theta$, we have that $\int_S \boldsymbol{\pi} d(q\times\delta_{\theta''})=a$.

2. For $h(u,\pi)=\pi^2$ and $\theta''\in\Theta$, we have that $\int_S \boldsymbol{\pi}^2 d(q\times\delta_{\theta''})=a^2+\left(\theta_3''\right)^2$.

3. For $h(u,\pi)=u$ and $\theta''\in\Theta$, we have that $\int_S \boldsymbol{u} d(q\times\delta_{\theta''})=\theta_0''+\left(\theta_{1\pi}''+\theta_{1a}''\right)a$.

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

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

Given (42), 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 $\theta_0'+\theta_{1a}'a=\theta_0+\theta_{1a}a$. This concludes the proof of the first set inclusion and formalizes the aforementioned heuristics.

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

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

yielding that

$$q\left(\frac{u-\theta_0-\theta_{1a}a-\theta_{1\pi}\pi}{\theta_2},\frac{\pi-a}{\theta_3}\right)=q\left(\frac{u-\theta_0'-\theta_{1a}'a-\theta_{1\pi}'\pi}{\theta_2'},\frac{\pi-a}{\theta_3'}\right)\quad\forall(u,\pi)\in C.$$

that is, $\theta'\in\hat{\Sigma}_a(\theta)$. This proves the statement. ■

8.3 Proposition 22

It holds

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

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

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

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

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

Putting $a=a^*$ we get

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

and so

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

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

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

as desired. ■

8.4 Proposition 23

Proof It holds

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

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

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

and so

$$a^* \gtrless a^o \Longleftrightarrow a^* \frac{\theta_{1a}^* - \mathbb{E}_{\mu^*}(\theta_{1a})}{\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1a})} \lesseqgtr 0 \tag{43}$$

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 (26) $\mathbb{E}_{\mu^*}(\theta_{1a}) \neq -\hat{\beta}^*$ and so $\mathbb{E}_{\mu^*}(\theta_{1a}) < -\hat{\beta}^*$. By (43), $(\theta_{1a}^* - \mathbb{E}_{\mu^*}(\theta_{1a})) / (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1a})) < 0$, which in turn implies $\mathbb{E}_{\mu^*}(\theta_{1a}) < \theta_{1a}^*$. Conversely, suppose $\mathbb{E}_{\mu^*}(\theta_{1a}) < \theta_{1a}^*$. Since $\mathbb{E}_{\mu^*}(\theta_{1a}) \leq -\hat{\beta}^*$, by (26) it follows $a^* > 0$. Moreover, being $(\theta_{1a}^* - \mathbb{E}_{\mu^*}(\theta_{1a})) / (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1a})) < 0$, by (43) it holds $a^* > a^o$. (iii) Suppose $0 < a^* < a^o$. By (43), $(\theta_{1a}^* - \mathbb{E}_{\mu^*}(\theta_{1a})) / (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1a})) > 0$, that is, $\mathbb{E}_{\mu^*}(\theta_{1a}) \in (\theta_{1a}^*, -\hat{\beta}^*)$. Conversely, suppose $\mathbb{E}_{\mu^*}(\theta_{1a}) \in (\theta_{1a}^*, -\hat{\beta}^*)$. By (26), $a^* > 0$. Moreover, being $(\theta_{1a}^* - \mathbb{E}_{\mu^*}(\theta_{1a})) / (\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1a})) > 0$, by (43) it holds $a^* < a^o$. ■

8.5 Proposition 24 and eq. 41

First note that

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

and

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

Hence,

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

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

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

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

$$l(a^*, \theta^*) = \theta_0^{*2} \left(\hat{\beta}^* + \mathbb{E}_{\mu^*}(\theta_{1a})\right)^2 = \theta_0^{*2} \left(\mathbb{E}_{\mu^*}(\theta_{1a}) - \theta_{1a}^*\right)^2 \tag{44}$$

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

$$l(a^*, \theta^*) = 0 \tag{45}$$

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

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

In the next section we show that

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

Hence,

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

It is easy to check that, along with (44) and (45), this completes the proof. ■

8.6 Eq. 39

It holds

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

as desired. ■

8.7 Proposition 26

Lemma 28 *Consider a self-confirming equilibrium (a^*,μ^*). The following conditions are equivalent:*

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

Proof (i) implies (ii) Suppose that $a^*>0$ implies $a^*>a^o$. By (43), $a^*>0$ implies

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

and so $\left(\theta_{1a}^*-\mathbb{E}_{\mu^*}(\theta_{1a})\right)/(\hat{\beta}^*+\mathbb{E}_{\mu^*}(\theta_{1a}))<0$. (ii) implies (i) Suppose $\left(\theta_{1a}^*-\mathbb{E}_{\mu^*}(\theta_{1a})\right)/(\hat{\beta}^*+\mathbb{E}_{\mu^*}(\theta_{1a}))<0$. If $a^*>0$, then $a^*\left(\theta_{1a}^*-\mathbb{E}_{\mu^*}(\theta_{1a})\right)/(\hat{\beta}^*+\mathbb{E}_{\mu^*}(\theta_{1a}))<0$, and so, by (43), $a^*>a^o$. (ii) implies (iii) Suppose $\left(\theta_{1a}^*-\mathbb{E}_{\mu^*}(\theta_{1a})\right)/(\hat{\beta}^*+\mathbb{E}_{\mu^*}(\theta_{1a}))<0$. If $a^*<0$, then $a^*\left(\theta_{1a}^*-\mathbb{E}_{\mu^*}(\theta_{1a})\right)/(\hat{\beta}^*+\mathbb{E}_{\mu^*}(\theta_{1a}))>0$, and so, by (43), $a^*<a^o$. (iii) implies (ii) Suppose that $a^*<0$ implies $a^*<a^o$. By (43), $a^*<0$ implies

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

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

Proof of Proposition 26 By Lemma 28, $a^*>0$ implies $a^*<a^o$ if and only if $\left(\theta_{1a}^*-\mathbb{E}_{\mu^*}(\theta_{1a})\right)/(\hat{\beta}^*+\mathbb{E}_{\mu^*}(\theta_{1a}))<0$. In turn, this happens if and only if $\theta_{1a}^*\vee\left(-\hat{\beta}^*\right)<\mathbb{E}_{\mu^*}(\theta_{1a})$ or $\mathbb{E}_{\mu^*}(\theta_{1a})<\theta_{1a}^*\wedge\left(-\hat{\beta}^*\right)$. By Lemma 28, $a^*>0$ implies $a^*<a^o$ if and only if $\left(\theta_{1a}^*-\mathbb{E}_{\mu^*}(\theta_{1a})\right)/(\hat{\beta}^*+\mathbb{E}_{\mu^*}(\theta_{1a}))>0$. In turn, this happens if and only if $\theta_{1a}^*\vee\left(-\hat{\beta}^*\right)>\mathbb{E}_{\mu^*}(\theta_{1a})>\theta_{1a}^*\wedge\left(-\hat{\beta}^*\right)$. ■

References

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

[2] Battigalli, P., S. Cerreia-Vioglio, F. Maccheroni, and M. Marinacci, Self-confirming equilibrium and model uncertainty, IGIER-Bocconi wp 428, 2011.

[3] 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.

[4] Denzau, A. T. and D. C. North, Shared mental models: ideologies and institutions, *Kyklos,* 47, 3-31, 1994.

[5] Gilboa, I. and D. Schmeidler, Maxmin expected utility with a non-unique prior, *Journal of Mathematical Economics*, 18, 141-153, 1989.

[6] Hansen, L. P. and T. J. Sargent, Robust control and model uncertainty, *American Economic Review*, 60-66, 2001.

[7] Hansen, L. P. and T. J. Sargent, *Robustness.* Princeton: Princeton University Press, 2008.

[8] Hewitt, E. and K. Stromberg, *Real and abstract analysis*, Berlin: Springer-Verlag, 1965.

[9] Klibanoff, P., M. Marinacci and S. Mukerji, A smooth model of decision making under ambiguity, *Econometrica*, 73, 1849-1892, 2005.

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

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

[12] Lucas, R. E. Jr., Methods and problems in business cycle theory, *Journal of Money, Credit, and Banking,* 12, 696-715, 1980.

[13] Marschak, J., Economic measurements for policy and prediction, in *Studies in Econometric Method* (W. Hood and T. J. Koopmans, eds.). New York: Wiley, 1953.

[14] Neyman, J., Inductive behavior as a basic concept of philosophy of science, *Review of the International Statistical Institute*, 25, 7-22, 1957.

[15] Phelps, R. R., *Lectures on Choquet's theorem*, Princeton: Van Nostrand, 1966.

[16] Rodrik, D., When ideas trump interests: preferences, worldviews, and policy innovations, *Journal of Economic Perspectives*, 28, 189-208, 2014.

[17] Rothenberg, T. J., Identification in parametric models, *Econometrica*, 39, 577-591, 1971.

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

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

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

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

[22] Savage, L. J., *The foundations of statistics.* New York: John Wiley and Sons, 1954.

[23] Sion, M., On general minimax theorems, *Pacific Journal of Mathematics*, 8, 171-176, 1958.

[24] Theil, H., *Economic forecasts and policy*, Amsterdam: North Holland, 1961.

[25] Tinbergen, J., *On the theory of economic policy*, Amsterdam: North Holland, 1952.