Journal of Industrial Organization Education

Manuscript 1049

A systematic procedure for finding Perfect Bayesian Equilibria in Incomplete Information Games

Felix Munoz-Garcia, *Washington State University*

DOI: 10.1515/1935-5041.1049

A systematic procedure for finding Perfect Bayesian Equilibria in Incomplete Information Games

Felix Munoz-Garcia

Abstract

This paper provides a non-technical introduction to a procedure to find Perfect Bayesian equilibria (PBEs) in incomplete information games. Despite the rapidly expanding literature on industrial organization that uses PBE as its main solution concept, most undergraduate and graduate textbooks still present a relatively theoretical introduction to PBEs. This paper offers a systematic five-step procedure that helps students find all pure-strategy PBEs in incomplete information games. Furthermore, it illustrates a step-by-step application of this procedure to a signaling game, using a worked-out example.

KEYWORDS: perfect bayesian equilibria, systematic procedure, incomplete information, signaling games

Author Notes: I would like to thank Ana Espinola-Arrendondo for her useful comments and suggestions.

Introduction

The literature on industrial organization and applied game theory has significantly contributed to our understanding of strategic interactions in sequential-move contexts where one agent has access to more accurate information than his/her opponents. Examples include Spence (1974), who uses signaling games for the study of labor market games; Milgrom and Roberts (1982 and 1986), Harrington (1986), and Bagwell and Ramey (1990, 1991), which analyze limit-pricing practices by one or multiple incumbents; Gal-Or (1989), who examines warranties; and, more recently, Ridley (2008) and Fong (2011), which consider entry-deterrence games and players' revelation of their altruism concerns, respectively.

Most studies in this literature use the Perfect Bayesian equilibrium (PBE) solution concept, since strategies in these information settings must be sequentially rational. Despite the wide use of PBE, most undergraduate and graduate textbooks on game theory still provide relatively theoretical definitions of PBE. Yet, they essentially lack a systematic exposition on how to find PBEs in incomplete information games using step-by-step examples. Furthermore, the PBE is often one of the most advanced solution concepts introduced in undergraduate game theory courses (as well as in certain Masters' programs), leading many students to especially struggle with this topic, ultimately deterring them from pursuing research in this rapidly expanding area of economics.

This paper introduces both undergraduate and graduate students to a systematic five-step procedure that allows for a search of all PBEs in pure strategies. Such a procedure is often used by many scholars in the field of industrial organization, but it is relegated to technical appendices, thereby limiting its dissemination among undergraduate and Master's students. Our paper first provides a non-technical introduction to the PBE solution concept, and then offers a step-by-step application of this procedure to a signaling game.[1] This paper includes graphical illustrations, in order to focus students' attention on the most relevant payoff comparisons at each of the PBE we examine. Furthermore, and for completeness, we emphasize the distinction between equilibrium and off-the-equilibrium beliefs, in order to familiarize non-technical readers with a topic several students find especially challenging.

The following section describes the PBE solution concept, separately discussing its two main ingredients: sequential rationality in incomplete information environments and consistency of beliefs. Section 3 then presents the five steps of the systematic procedure. Finally, section 4 applies this procedure to a signaling game between a monetary authority, who announces an inflation target, and a labor union, who responds demanding a wage increase.

[1] In order to facilitate the use of this procedure to other signaling games, such as those analyzing limit pricing or advertising, the game we consider is strategically similar to the labor-market signaling game introduced by Spence (1974).

Perfect Bayesian Equilibrium - Definition

A strategy profile for N players $(s_1, s_2, ..., s_N)$ and a system of beliefs over the nodes at all information sets are a PBE if:

a) Each player's strategies specify optimal actions, given the strategies of the other players, and given his beliefs.

b) The beliefs are consistent with Bayes' rule, whenever possible.

This definition hence emphasizes two elements we must find in a PBE. First, condition (a) resembles the condition for players' best responses in the standard definition of the Nash equilibrium solution concept, but applied to incomplete information settings, since players must find optimal actions given his beliefs about his opponents' types. Second, condition (b) stresses that beliefs must satisfy Bayes' rule. Furthermore, this property must hold both when players form beliefs along the equilibrium path (in this case, the application of Bayes' rule is straightforward, as we describe below), and off-the-equilibrium path (in this case, Bayes' rule cannot be applied as we illustrate below, and hence off-the-equilibrium beliefs must be arbitrarily specified). Let us separately examine each of the above conditions.

Sequential rationality in incomplete information contexts

In order to apply sequential rationality in an environment where players do not observe each others' types (e.g., production costs, abilities, etc.), we must extend the notion of sequential rationality applied in games of complete information (i.e., when using backward induction), to games of incomplete information. This implies, in particular, the need for every player to maximize his expected utility level, given his own beliefs about the other players' types. Specifically, at every information set at which a player is called on to move, he must choose the strategy that maximizes his expected utility, given that all other players will do the same, and given his own beliefs about the other players' types.

Example: Consider the following sequential-move game with incomplete information. A monetary authority (such as the Federal Reserve Bank, or the European Central Bank) privately observes its degree of commitment with maintaining low inflation levels. After observing its type (either Strong or Weak, with probabilities 0.6 and 0.4, respectively), the monetary authority decides to announce that the expectation for inflation during the next period will be either High or Low. Upon observing the message sent by the monetary authority, but without observing its true type, the labor union responds asking a high wage increase (denoted as H in the figure) or a low wage increase (represented with L). For compactness, μ denotes the labor union's belief about the monetary authority's type being Strong upon observing a High inflation announcement (in the vertical information set located on the left-hand side of the figure), i.e.,

$\mu \equiv \mu(Strong|HighInflation)$. Likewise, γ represents the labor union's beliefs after observing a Low inflation announcement (in the information set on the right-hand side of the figure), i.e., $\gamma \equiv \mu(Strong|LowInflation)$.

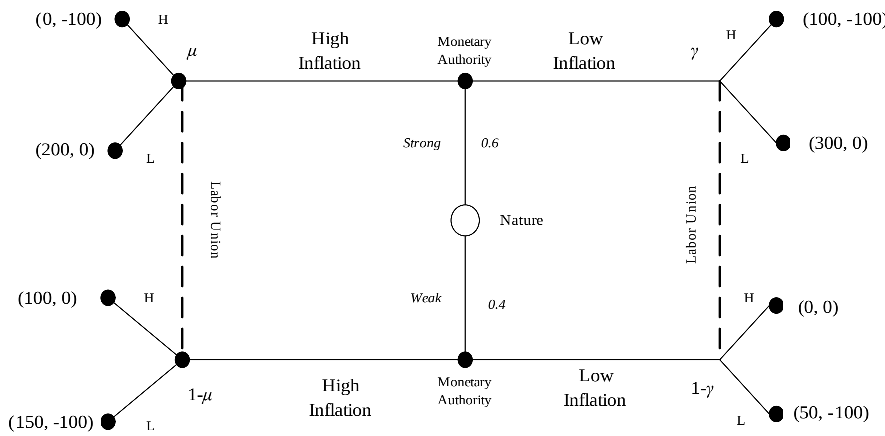

Fig 1. Monetary authority signaling game.

In this setting, after observing a low inflation announcement (in the right-hand side) the labor union responds with a high salary increase (H) if and only if $EU_{Labor}(H|LowInflation) > EU_{Labor}(L|LowInflation)$. That is, if

$$(-100)\gamma + 0(1-\gamma) > 0\gamma + (-100)(1-\gamma)$$

which holds for all $\gamma < \frac{1}{2}$. Similarly, if the monetary authority instead announces a high inflation target (in the left-hand side of the figure), the labor union responds with a high salary increase (H) if and only if $EU_{Labor}(H|HighInflation) > EU_{Labor}(L|HighInflation)$. That is, if

$$(-100)\mu + 0(1-\mu) > 0\mu + (-100)(1-\mu)$$

which holds for all $\mu < \frac{1}{2}$.

Conditional beliefs about types

Let us now examine a player's beliefs about his opponent's type. First note that player, by observing his opponent's action, might be able to infer something about the his opponent's type through such action. In this case, we say that a player (e.g., the labor union) *updates* his beliefs about his opponent's type (the monetary authority's type).

Such belief updating must, in addition, satisfy Bayes' rule. In order to understand the use of Bayes' rule in this context, let us apply it to the previous example. Let us hence denote by α^{Strong} the probability that the Strong type of monetary authority announces a high inflation, and by α^{Weak} the probability that the Weak type announces a high inflation. Then, after observing a high-inflation announcement (as illustrated in the figure below), the labor union's belief that such a message originates from a Strong monetary authority, μ, can be expressed as

$$\mu = \frac{p\alpha^{Strong}}{p\alpha^{Strong} + (1-p)\alpha^{Weak}} = \frac{0.6\alpha^{Strong}}{0.6\alpha^{Strong} + 0.4\alpha^{Weak}}$$

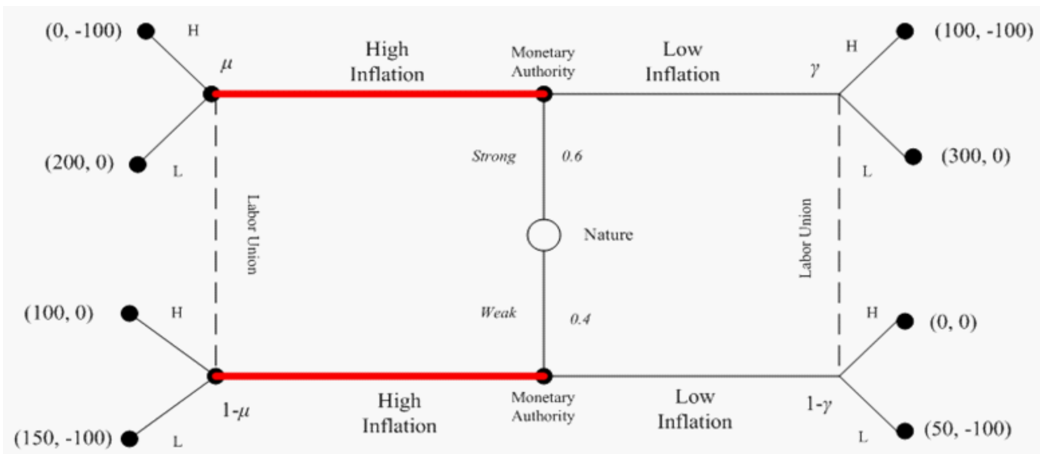

Fig. 2. Belief updating.

Intuitively, the labor union's beliefs are defined by a conditional probability: the probability that, conditional on observing a high-inflation announcement (which occurs with probability $0.6\alpha^{Strong} + 0.4\alpha^{Weak}$), such announcement originates from a Strong monetary authority. In other words, we divide the probability that the monetary authority is Strong *and* makes a high-inflation announcement, $0.6\alpha^{Strong}$, over the probability that any type of monetary authority (Strong or Weak) announces a high inflation level. Because Bayes' rule analyzes how a player updates his beliefs after observing his opponent's action, the literature often refers to a player's updated beliefs as its "posterior" beliefs, as opposed to his "prior" beliefs (which simply coincide with the initial probability distribution over types).

Example: Assume that $\alpha^{Strong} = \frac{1}{8}$ and $\alpha^{Weak} = \frac{1}{16.}$. Then, the labor union's posterior beliefs about the monetary authority being Strong after observing a high-inflation announcement, μ, are

$$\mu = \frac{0.6\alpha^{Strong}}{0.6\alpha^{Strong} + 0.4\alpha^{Weak}} = \frac{0.6\frac{1}{8}}{0.6\frac{1}{8} + 0.4\frac{1}{16}} = 0.75$$

Intuitively, while nature assigns a probability of 0.6 to the monetary authority being Strong, the labor union assigns a *larger* probability weight to the event that the observed high-inflation announcement originates from the Strong type. Indeed, $\alpha^{Strong} = \frac{1}{8}$ and $\alpha^{Weak} = \frac{1}{16.}$ implies that the Strong monetary authority announces high-inflation targets with twice the probability than its Weak counterpart. As a consequence, upon observing the high-inflation announcement the labor union updates its beliefs in favor of the Strong type of sender.

A remark about off-the-equilibrium beliefs: Consider a setting in which one of the responder's information sets is unreached. In the above example, this occurs when both types of monetary authorities choose the same announcement: for instance, when both announce a high inflation, i.e., $\alpha^{Strong} = \alpha^{Weak} = 1$, the vertical information set in the opposite side of the game tree (right-hand side) is unreached. Consider now that the labor union observes the event of a low-inflation announcement. According to the above strategy profile, such an announcement should never be observed. Despite being surprised by this off-the-equilibrium announcement, the labor union must still update its belief γ, as follows

$$\gamma = \frac{0.6\left(1-\alpha^{Strong}\right)}{0.6\left(1-\alpha^{Strong}\right)+0.4\left(1-\alpha^{Weak}\right)} = \frac{0}{0}$$

where $1-\alpha^{Strong}$ $\left(1-\alpha^{Weak}\right)$ denotes the probability that a Strong (Weak, respectively) monetary authority sends a high-inflation announcement. However, since $\alpha^{Strong} = \alpha^{Weak} = 1$, both the numerator and denominator become zero, yielding an indeterminate result for ratio γ. As a result, belief γ is thus indeterminate, and we are allowed to arbitrarily specify the value of γ, i.e., set any value between zero and one $\gamma \in [0, 1]$.[2]

Recall that condition (b) on the definition of a PBE stated that beliefs must be consistent with Bayes' rule "whenever possible." This qualification in condition (b) is related with our discussion of off-the-equilibrium beliefs. Indeed, when a player is called on to move at an information set that is reached along the equilibrium path, he can use Bayes' rule in order to update its posterior beliefs. However, when he is at an information set which should not be reached in equilibrium (off-the-equilibrium path), he cannot apply Bayes' rule in order to update his beliefs; and beliefs can be arbitrarily specified.

A natural question at this point is whether off-the-equilibrium beliefs are relevant, or a technicality that we can ignore in our search of equilibrium behavior in games of incomplete information. Off-the-equilibrium beliefs are important, since they determine the optimal response of a player after observing a particular message from his opponent. Depending on the optimal response we identify, the sender (e.g., monetary authority in our above example) can be induced to change his message (inflation announcement), thereby affecting our equilibrium results. As a consequence, we must be especially careful about off-the-equilibrium beliefs in our description of strategy profiles that can be sustained as a PBE.

[2] In this case, we refer to γ as "off-the-equilibrium" beliefs (also referred by some scholars as "out-of-equilibrium" beliefs), since it specifies beliefs about the probability of being in a node that belongs to an information set that is not reached in equilibrium.

Procedure to find PBEs

In this section we describe a systematic procedure to search for PBEs in incomplete information games where one player is privately informed about his type, while his opponent's type is common knowledge. In order to facilitate our analysis, note that we usually classify PBEs into two different classes: *separating* PBEs, where different types of the privately informed player behave differently, e.g., the Strong monetary authority announces a low inflation, while the Weak type of authority announces a high inflation. In contrast, in *pooling* PBEs all types of the privately informed player behave similarly, e.g., both the Strong and Weak type of monetary authority announce a low inflation. Let us next describe the procedure to check if a particular strategy profile (either separating or pooling) constitutes a PBE.

1. Specify a strategy profile for the privately informed player.

    - In our above example, there are only four possible strategy profiles for the privately informed monetary authority: two separating strategy profiles, $High^S Low^W$ and $Low^S High^W$, and two pooling strategy profiles, $High^S High^W$ and $Low^S Low^W$. (For future reference, one can shade the branches corresponding to the strategy profile we test.)

2. Update the uninformed player's beliefs using Bayes' rule at all information sets, whenever possible.

    - In our above example, we need to specify beliefs μ and γ, which arise after the labor union observes a high or a low inflation announcement, respectively.[3]

3. Given the uninformed player's updated beliefs, find his optimal response.

    - In our above example, we first determine the optimal response of the labor union (H or L) upon observing a high-inflation announcement (given its updated belief μ), and we then determine its optimal response (H or L) after observing a low-inflation announcement (given its updated belief γ). (Also for future reference, it might be helpful to shade the branches corresponding to the optimal responses we just found.)

4. Given the optimal response of the uninformed player, find the optimal action (message) for the informed player.

[3]Note that in a separating strategy profile both information sets are reached in equilibrium, and hence beliefs can be updated using Bayes' rule. In a pooling strategy profile, in contrast, one information set is unreached in equilibrium, thus implying that either μ or γ must be arbitrarily specified.

- In our previous example, we first check if the Strong monetary authority prefers to make a high or low inflation announcement (given the labor union's optimal response after receiving each possible message, as determined in step 3). We then operate similarly for the Weak type of monetary authority.

5. Then check if this strategy profile for the informed player coincides with the profile you suggested in step 1.

   (a) If it coincides, then this strategy profile, updated beliefs and optimal responses *can* be supported as a PBE of the incomplete information game.

   (b) Otherwise, we say that this strategy profile *cannot* be sustained as a PBE of the game.

The following section separately applies this procedure to test each of the four candidate strategy profiles: two separating strategy profiles, $Low^S High^W$, and $High^S Low^W$, and two pooling strategy profiles, $High^S High^W$ and $Low^S Low^W$.

Step-by-step example

Separating equilibrium with $Low^S High^W$

First step. We first specify the separating strategy profile $Low^S High^W$ for the informed player, i.e., the Strong monetary authority announces a low inflation while the Weak authority announces a high inflation level. For future reference, figure 3 shades branches Low^S and $High^W$.

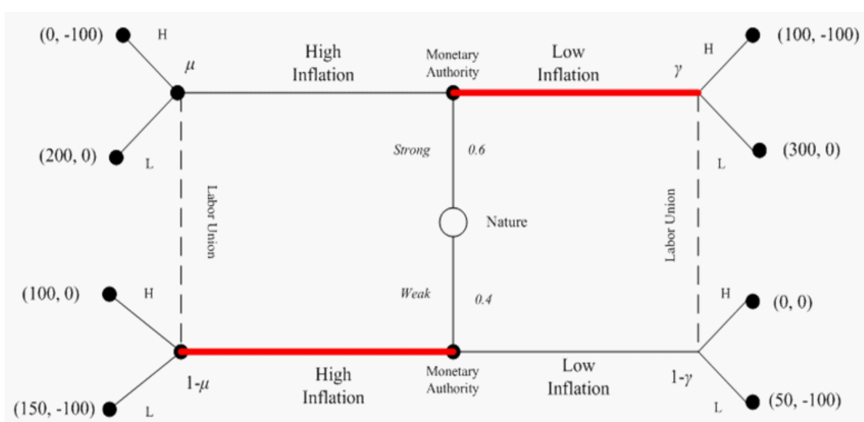

Fig 3. Separating strategy profile $Low^S High^W$.

Second step. We can now use Bayes' rule to update the uninformed player's (labor union) beliefs.

- Upon observing a high-inflation announcement (in the left-hand side of the figure), the labor union updates μ taking into account that in this separating strategy profile $\alpha^{Strong} = 0$ while $\alpha^{Weak} = 1$, i.e., only the weak type of authority announces a high-inflation target. More formally,

$$\mu = \frac{0.6\alpha^{Strong}}{0.6\alpha^{Strong} + 0.4\alpha^{Weak}} = \frac{0.6 \times 0}{0.6 \times 0 + 0.4 \times 1} = 0$$

  Intuitively indicating that, if the labor union observes a high-inflation announcement, it assigns zero probability to such announcement originating from a Strong type of monetary authority. Upon observing such announcement, in contrast, $1 - \mu = 1$, representing that the labor assigns full probability to this announcement being made by the Weak type of authority. Graphically, this belief updating entails that, along the vertical information set on the left-hand side of the figure, we focus on the lower node alone.

- Similarly, after observing a low-inflation announcement (on the right-hand side of the game tree), the labor union updates γ still considering that $\alpha^{Strong} = 0$ and $\alpha^{Weak} = 1$ in this strategy profile.

$$\gamma = \frac{0.6\left(1 - \alpha^{Strong}\right)}{0.6\left(1 - \alpha^{Strong}\right) + 0.4\left(1 - \alpha^{Weak}\right)} = \frac{0.6 \times 1}{0.6 \times 1 + 0.4 \times 0} = 1$$

  Intuitively implying that, if the labor union observes a low-inflation announcement, it believes that such a message must originate from a Weak type of authority, i.e., $\gamma = 1$, and never stem from a Strong authority, i.e., $1 - \gamma = 0$. Graphically, this belief updating entails that, along the vertical information set on the right-hand side of the tree, we focus on the upper node alone.

Third step. Optimal response of the uninformed player:

- Upon observing a *high-inflation announcement*, since $\mu = 0$, the labor union focuses on the lower node of this information set (see lower left-hand corner of the figure). Given this belief, the labor union responds with a high salary increase (H) since its associated payoff ($0) is larger than that from L (-$100). In order to keep track of this result, Figure 4 below shades the branch corresponding to the labor union's response of H after observing a high-inflation announcement. Importantly, H must be shaded for all nodes that belong to this information set, since the uninformed labor union cannot condition his response on the monetary authority's type.

- Upon observing a *low-inflation announcement*, since $\gamma = 1$, the labor union focuses on the upper node of this information set (see upper right-hand corner of the figure). Given this belief, the labor union responds with a low salary increase (L) since its associated payoff ($0) is larger than that from H (-$100). Similarly as above, Figure 4 shades the branch corresponding to L after observing a low-inflation announcement. Furthermore, L must be shaded in all nodes within this information set.

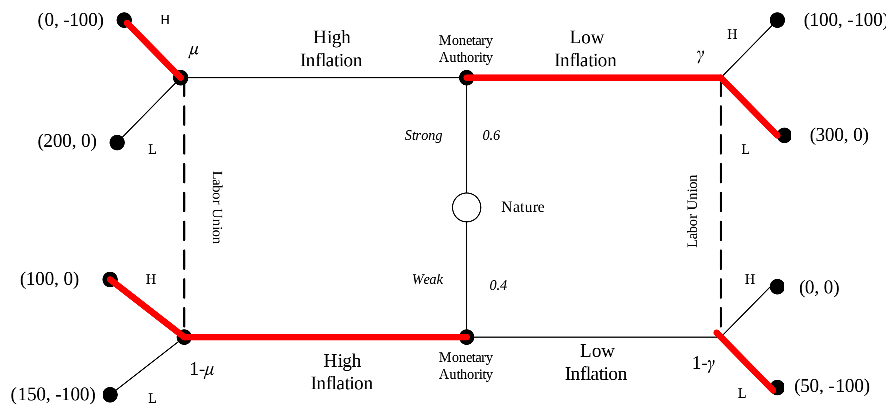

Fig 4. Separating strategy profile $Low^S High^W$.

Fourth step. Given the uninformed player's optimal responses, we can now determine the informed player's optimal messages.

- *Strong type*. If the Strong monetary authority behaves as prescribed by this strategy profile (announcing a low-inflation target), it can anticipate that such announcement will be responded with a low wage demand, L (we just need to follow the shaded branches in the upper part of the figure corresponding to the Strong type of authority), ultimately yielding a payoff of $300. If, instead, it deviates towards a high-inflation announcement, it can anticipate that such an announcement will be responded with a high-wage demand (H), implying a lower payoff of $0. Hence, the strong monetary authority does not have incentives to deviate from the separating strategy profile.

- *Weak type*. If the Weak monetary authority behaves as prescribed by this strategy profile (announcing a high-inflation target), it can anticipate that such announcement will be responded with a high wage demand, H, ultimately yielding a payoff of $100. If, instead, it deviates towards a low-inflation announcement, it can anticipate that such an announcement will be responded with a low-wage demand (L), implying a lower payoff of $50. Therefore, the weak monetary authority does not have incentives to deviate from the prescribed strategy profile either.

Fifth step. Therefore, no type of privately informed player has unilateral incentives to deviate from the prescribed separating strategy profile $Low^S High^W$, whereby the monetary authority announces a low inflation only when its type is Strong. As a consequence, the separating strategy profile $Low^S High^W$ *can* be sustained as a PBE of this incomplete information game.

Separating equilibrium with $High^S Low^W$

First step. We next specify the opposite separating strategy profile $High^S Low^W$, i.e., the Strong monetary authority announces a high inflation while the Weak authority announces a low inflation level. (You might suspect that this strategy profile is insensible, or literary crazy. Your suspicion was right since, as we next show, this strategy profile cannot be sustained as a PBE.) Following the same procedure as above, figure 5 shades branches $High^S$ and Low^W.

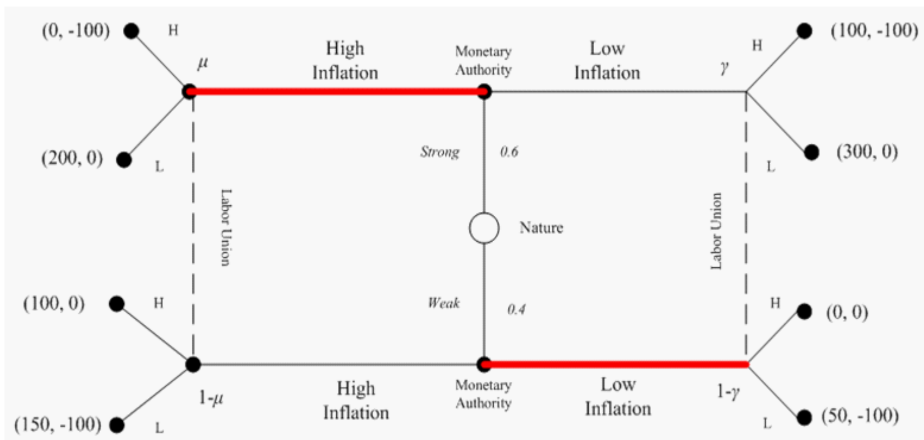

Fig 5. Separating strategy profile $High^S Low^W$.

Second step. We can now use Bayes' rule to update the uninformed player's (labor union) beliefs.

- Upon observing a high-inflation announcement (in the left-hand side of the figure), the labor union updates μ taking into account that in this separating strategy profile $\alpha^{Strong} = 1$ while $\alpha^{Weak} = 0$, i.e., only the strong type of authority announces a high-inflation target. More formally,

$$\mu = \frac{0.6\alpha^{Strong}}{0.6\alpha^{Strong} + 0.4\alpha^{Weak}} = \frac{0.6 \times 1}{0.6 \times 1 + 0.4 \times 0} = 1$$

  Intuitively indicating that, if the labor union observes a high-inflation announcement, it assigns full probability to such announcement originating from a Strong type of monetary authority.[4]

[4]Upon observing such an announcement, in contrast, $1 - \mu = 0$, representing that the labor union assigns zero probability to this announcement being made by the Weak type of authority. Graphically, this belief updating entails that, along the vertical information set on the left-hand side of figure 5, we now focus on the upper node.

- Similarly, after observing a low-inflation announcement (on the right-hand side of the game tree), the labor union updates γ as follows:

$$\gamma = \frac{0.6\left(1-\alpha^{Strong}\right)}{0.6\left(1-\alpha^{Strong}\right)+0.4\left(1-\alpha^{Weak}\right)} = \frac{0.6\times 0}{0.6\times 0+0.4\times 1} = 0$$

  Intuitively implying that, if the labor union observes a low-inflation announcement, it believes that such a message must originate from a Weak type of authority, i.e., $\gamma = 0$, and never originate from a Strong authority, i.e., $1-\gamma = 1$. Graphically, this belief updating entails that, along the vertical information set on the right-hand side of the tree, we focus on the lower node.

Third step. Optimal response of the uninformed player:

- Upon observing a *high-inflation announcement*, since $\mu = 1$, the labor union focuses on the upper node of this information set (see upper left-hand corner of the figure). Given this belief, the labor union responds with a low salary increase (L) since its associated payoff ($0) is larger than that from H (-$100). Figure 6 below shades the branch corresponding to the labor union's response of L after observing a high-inflation announcement.

- Upon observing a *low-inflation announcement*, since $\gamma = 0$, the labor union focuses on the lower node of this information set (see lower right-hand corner of the figure). Given this belief, the labor union responds with a high salary increase (H) since its associated payoff ($0) is larger than that from L (-$100). Similarly as above, figure 6 shades the branch corresponding to the labor union's response of H after observing a low-inflation announcement.

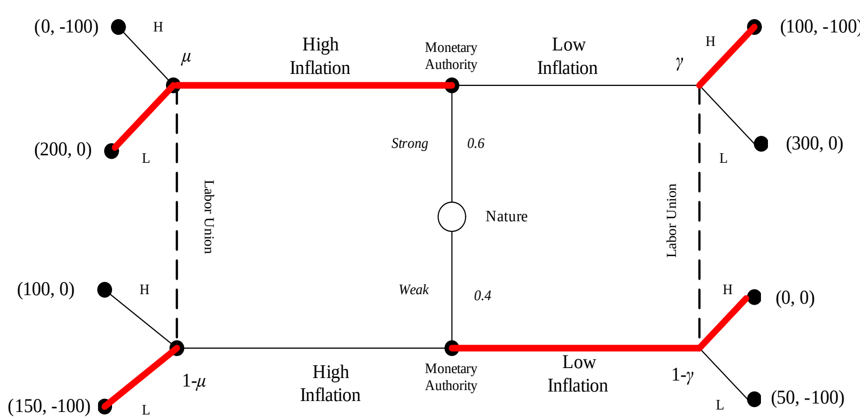

Fig 6. Separating strategy profile $High^S Low^W$.

Fourth step. Given the uninformed player's optimal responses, we can now determine the informed player's optimal messages.

- *Strong type*. If the Strong monetary authority behaves as prescribed by this strategy profile (announcing a high-inflation target), it can anticipate that such announcement will be responded with a low wage demand, L (we just need to follow the shaded branches in the upper part of figure 6 corresponding to the Strong type of authority), ultimately yielding a payoff of \$200. If, instead, it deviates towards a low-inflation announcement, it can anticipate that such an announcement will be responded with a high-wage demand (H), implying a lower payoff of \$100. Hence, the strong monetary authority does not have incentives to deviate from the prescribed strategy profile.

- *Weak type*. If the Weak monetary authority behaves as prescribed by this strategy profile (announcing a low-inflation target), it can anticipate that such announcement will be responded with a high wage demand, H (just follow the shaded branches in the lower part of figure 6 corresponding to the Weak type of authority), ultimately yielding a payoff of \$0. If, instead, it deviates towards a high-inflation announcement, it can anticipate that such an announcement will be responded with a low-wage demand (L), implying a higher payoff of \$150. Therefore, the weak monetary authority *has* incentives to deviate from the prescribed strategy profile.

Fifth step. Since we found that one type of privately informed player (the Weak type of monetary authority) has incentives to deviate from the prescribed strategy profile $Low^S High^W$, we can conclude that $Low^S High^W$ *cannot* be sustained as a PBE of this incomplete information game.

Pooling equilibrium with $High^S High^W$

First step. Let us next check if the pooling strategy profile $High^S High^W$ where both types of monetary authority announce a high inflation can be sustained as a PBE. Following the same approach as for the separating strategy profiles, figure 7 below shades branches $High^S$ and

$High^W$.

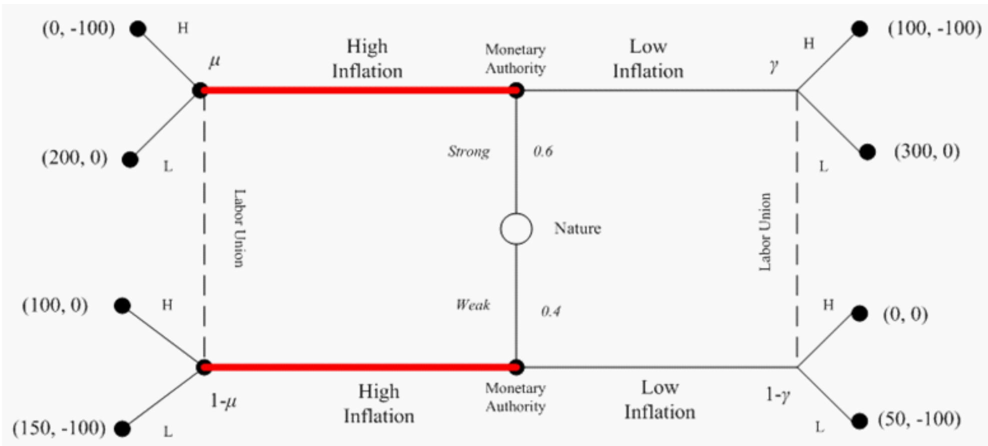

Fig 7. Pooling strategy profile $High^S High^W$.

Second step. We can now use Bayes' rule to update the uninformed player's (labor union) beliefs.

- Upon observing a high-inflation announcement, the labor union updates μ taking into account that in this pooling strategy profile $\alpha^{Strong} = 1$ and $\alpha^{Weak} = 1$, i.e., both the strong and weak types of authority announce a high-inflation target. More formally,

$$\mu = \frac{0.6\alpha^{Strong}}{0.6\alpha^{Strong} + 0.4\alpha^{Weak}} = \frac{0.6 \times 1}{0.6 \times 1 + 0.4 \times 1} = 0.6$$

  which coincide with the prior probability that the monetary authority is Strong. Intuitively, since both types of authorities select a high-inflation announcement in this strategy profile, the labor union's observation of a high-inflation announcement does not allow it to further restrict its posterior beliefs about the monetary authority's type, i.e., the announcement becomes uninformative. Hence, the posterior beliefs (updated using Bayes' rule) coincide with the prior probability distribution. This is a common result in pooling strategy profiles, whereby updated beliefs along the equilibrium path coincide with the prior probability distribution over types.[5]

- If the labor union observes a low-inflation announcement, the labor union must still update γ considering that $\alpha^{Strong} = 1$ and $\alpha^{Weak} = 1$. Note, however, that such an announcement only occurs off-the-equilibrium path according to this strategy profile. Indeed, if we use Bayes' rule to update the labor union's beliefs in this setting we obtain an

[5] Unlike in the case of separating strategy profiles, these beliefs entail that, along the vertical information set on the left-hand side of figure 7, we cannot focus on one of the nodes, since the probability of being in the upper node is still 0.6 and that of being in the lower node is 0.4 (both of them being different from zero).

indeterminate result,

$$\gamma = \frac{0.6\left(1-\alpha^{Strong}\right)}{0.6\left(1-\alpha^{Strong}\right)+0.4\left(1-\alpha^{Weak}\right)} = \frac{0.6\times 0}{0.6\times 0+0.4\times 0} = \frac{0}{0}$$

implying that this player's off-the-equilibrium beliefs can be arbitrarily specified, i.e., $\gamma \in [0,1]$.

Third step. Optimal response of the uninformed player:

- Upon observing a *high-inflation announcement*, since beliefs along the equilibrium path satisfy $\mu = 0.6$, the labor union cannot focus on a single node, and must select whether to respond with H or L by comparing the expected utility of each response, as follows.

$$\begin{aligned} EU_{Labor}\left(H|High\right) &= 0.6\times(-100)+0.4\times 0 = -60 \\ EU_{Labor}\left(L|High\right) &= 0.6\times 0+0.4\times(-100) = -40 \end{aligned}$$

  Hence, the labor union optimally responds with a low salary increase (L) since its associated expected payoff ($-40) is larger than that from H (-$60). Figure 8 below shades the branch corresponding to the labor union's response of L after observing a high-inflation announcement.

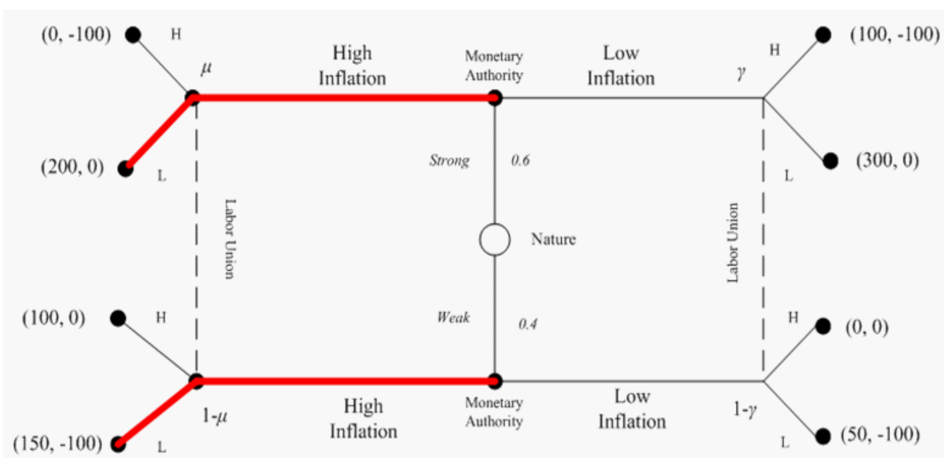

Fig 8. Pooling strategy profile $High^{S}High^{W}$.

- Upon observing a *low-inflation announcement*, since off-the-equilibrium beliefs satisfy $\gamma \in [0,1]$, the labor union cannot focus on a single node, and must select whether to respond with H or L by computing the expected utility of each response, as follows.

$$\begin{aligned} EU_{Labor}\left(H|Low\right) &= \gamma\times(-100)+(1-\gamma)\times 0 = -100\gamma \\ EU_{Labor}\left(L|Low\right) &= \gamma\times 0+(1-\gamma)\times(-100) = -100+100\gamma \end{aligned}$$

Given these expected payoffs, the labor union responds with a high salary increase (H) if and only if $-100\gamma > -100 + 100\gamma$, or $\frac{1}{2} > \gamma$. We will then need to divide our following step, where we analyze the optimal announcements of the monetary authority, into two cases: (1) $\gamma < \frac{1}{2}$, where the labor union responds with H after observing a low-inflation announcement; and (2) $\gamma \geq \frac{1}{2}$, where the labor union responds with L after observing a low-inflation announcement.

Fourth step. Given the uninformed player's optimal responses, we can now determine the informed player's optimal messages.

1. **CASE 1** $\gamma < \frac{1}{2}$: These off-the-equilibrium beliefs induce the labor union to respond with H after observing a low-inflation announcement (to facilitate comparison, figure 9 below shades the branches corresponding to H in the right-hand side of the figure). Let us next check if either type of monetary authority has incentives to deviate from the prescribed pooling strategy profile.

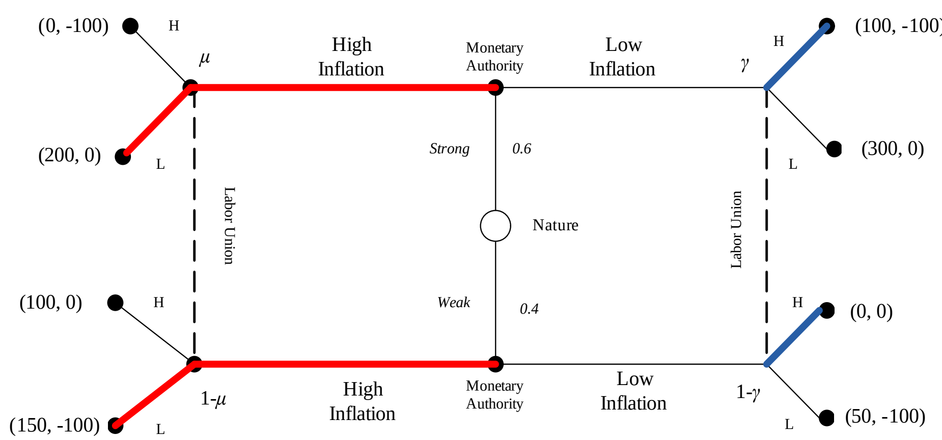

Fig 9. Pooling strategy profile $High^S High^W$ when $\gamma < \frac{1}{2}$.

- *Strong type.* If the Strong monetary authority behaves as prescribed by this strategy profile (announcing a high-inflation target), it can anticipate that such announcement will be responded with a low wage demand, L, ultimately yielding a payoff of \$200. If, instead, it deviates towards a low-inflation announcement, it can anticipate that such an announcement will be responded with a high-wage demand (H) as indicated in the shaded branches in the right-hand side of the figure given that we consider the case in which $\gamma < \frac{1}{2}$, implying a lower payoff of \$100. Hence, the strong monetary authority does not have incentives to deviate from the prescribed strategy profile.

- *Weak type*. If the Weak monetary authority behaves as prescribed by this strategy profile (announcing a high-inflation target), it can anticipate that such announcement will be responded with a low wage demand, L, ultimately yielding a payoff of \$150. If, instead, it deviates towards a low-inflation announcement, it can anticipate that such an announcement will be responded with a high-wage demand (H) as indicated in the shaded branches in the right-hand side of the figure given that we consider the case in which $\gamma < \frac{1}{2}$, implying a lower payoff of \$0. Therefore, the weak monetary authority does not have incentive to deviate from the prescribed pooling strategy profile either.
- Since no type of privately informed player (monetary authority) has incentives to deviate from the prescribed pooling strategy profile, we conclude that $High^S High^W$ *can* be supported as a PBE when off-the-equilibrium beliefs satisfy $\gamma < \frac{1}{2}$.[6]

2. **CASE 2** $\gamma \geq \frac{1}{2}$: These off-the-equilibrium beliefs induce the labor union to respond with L after observing a low-inflation announcement (to facilitate comparison, figure 10 below shades the branches corresponding to L in the right-hand side of the figure). Let us next check if either type of monetary authority has incentives to deviate from the prescribed pooling strategy profile.

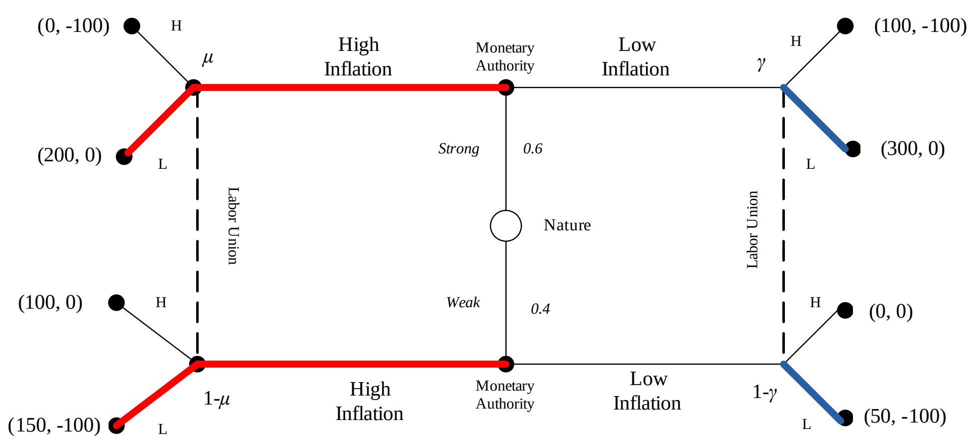

Fig 10. Pooling strategy profile $High^S High^W$ when $\gamma \geq \frac{1}{2}$.

[6]Nonetheless, this pooling PBE does not survive the Cho and Kreps' (1987) Intuitive Criterion. Indeed, if the labor union observes the off-the-equilibrium message of low inflation, it could infer that the only type of monetary authority for which such a message produces a higher payoff than in the pooling PBE is the Strong type of authority. Hence, the labor union's off-the-equilibrium beliefs would be restricted to $\gamma = 1$ (leading it to respond with L), thus providing the Strong monetary authority with incentives to deviate from the prescribed pooling strategy profile. For more details on the application of this refinement criterion to signaling games, see Espinola-Arredondo and Munoz-Garcia (2010).

- *Strong type.* If the Strong monetary authority behaves as prescribed by this strategy profile (announcing a high-inflation target), it can anticipate that such announcement will be responded with a low wage demand, L, ultimately yielding a payoff of \$200. If, instead, it deviates towards a low-inflation announcement, it can anticipate that such an announcement will be responded with a low-wage demand (L) as indicated in the shaded branches in the right-hand side of figure 10 given that we consider the case in which $\gamma \geq \frac{1}{2}$, implying a higher payoff of \$300. Hence, the strong monetary authority *has* incentives to deviate from the prescribed pooling strategy profile.[7]
- *Weak type*. If the Weak monetary authority behaves as prescribed by this strategy profile (pooling with the Strong type), it can anticipate that such an announcement will be responded with a low wage demand, L, ultimately yielding a payoff of \$150. If, instead, it deviates towards a low-inflation announcement, it can anticipate that such an announcement will be responded with a low-wage demand (L) as indicated in the shaded branches in the right-hand side of figure 10 given that we consider the case in which $\gamma \geq \frac{1}{2}$, implying a lower payoff of \$50. Therefore, the weak monetary authority does not have incentive to deviate from the prescribed pooling strategy profile.
- Since we found that one type of privately informed player (the Strong monetary authority) has incentives to deviate from the prescribed pooling strategy profile, we conclude that $High^S High^W$ *cannot* be supported as a PBE when off-the-equilibrium beliefs satisfy $\gamma \geq \frac{1}{2}$.

Fifth step. Therefore, the pooling strategy profile $High^S High^W$ can only be supported as a PBE when off-the-equilibrium beliefs satisfy $\gamma < \frac{1}{2}$.

Pooling equilibrium with $Low^S Low^W$

First step. Let us finally check if the pooling strategy profile $Low^S Low^W$ where both types of monetary authority announce a low inflation can be sustained as a PBE. Following the same

[7]Once we identify one type of privately informed player with incentives to deviate from the prescribed strategy profile, we can claim that such strategy profile cannot be sustained as a PBE. For completeness, we nonetheless include the analysis corresponding to the weak type of monetary authority below.

approach as above, figure 11 shades branches Low^S and Low^W.

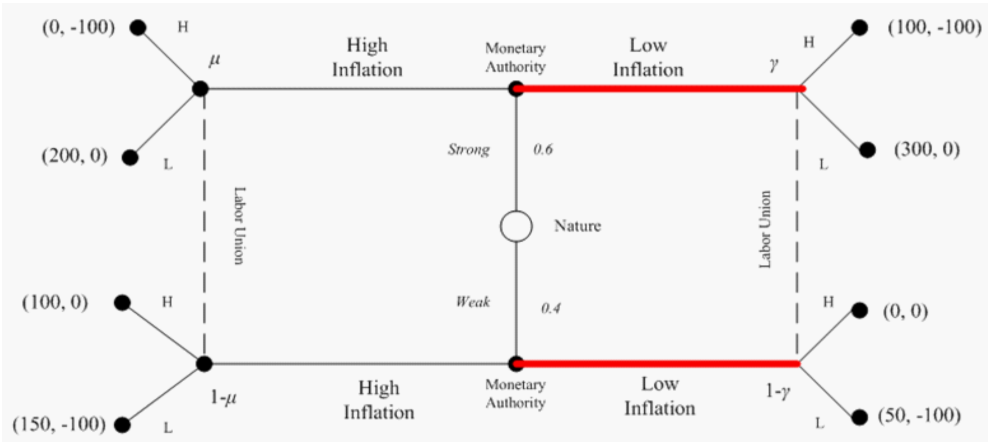

Fig 11. Pooling strategy profile $Low^S Low^W$.

Second step. We can now use Bayes' rule to update the uninformed player's (labor union) beliefs.

- Upon observing a high-inflation announcement (in the right-hand side of the figure), the labor union updates γ taking into account that in this strategy profile $\alpha^{Strong} = 0$ and $\alpha^{Weak} = 0$, i.e., neither the strong nor the weak type of authority announce a high-inflation target. More formally,

$$\gamma = \frac{0.6\left(1-\alpha^{Strong}\right)}{0.6\left(1-\alpha^{Strong}\right)+0.4\left(1-\alpha^{Weak}\right)} = \frac{0.6\times 1}{0.6\times 1+0.4\times 1} = 0.6$$

  which coincide with the prior probability that the monetary authority is Strong. Intuitively, since both types of authorities select a low-inflation announcement, the labor union's observation of a low-inflation announcement does not allow it to further restrict its posterior beliefs.

- If the labor union observes a high-inflation announcement (on the left-hand side of the game tree), the labor union must still update μ considering that $\alpha^{Strong} = 0$ and $\alpha^{Weak} = 0$. Such an announcement, however, only occurs off-the-equilibrium path according to this strategy profile. Indeed, if we use Bayes' rule to update the labor union's beliefs we obtain an indeterminate result,

$$\mu = \frac{0.6\alpha^{Strong}}{0.6\alpha^{Strong}+0.4\alpha^{Weak}} = \frac{0.6\times 0}{0.6\times 0+0.4\times 0} = \frac{0}{0}$$

  implying that this player's off-the-equilibrium beliefs can be arbitrarily specified, i.e., $\mu \in [0,1]$.

Third step. Optimal response of the uninformed player:

- Upon observing a *low-inflation announcement*, since beliefs along the equilibrium path satisfy $\gamma = 0.6$, the labor union must compare the expected utility of responding with H or L, as follows.

$$EU_{Labor}(H|Low) = 0.6 \times (-100) + 0.4 \times 0 = -60$$
$$EU_{Labor}(L|Low) = 0.6 \times 0 + 0.4 \times (-100) = -40$$

  Hence, the labor union optimally responds with a low salary increase (L) since its associated expected payoff is larger. Figure 12 below shades the branches corresponding to the labor union's response of L after observing a high-inflation announcement.

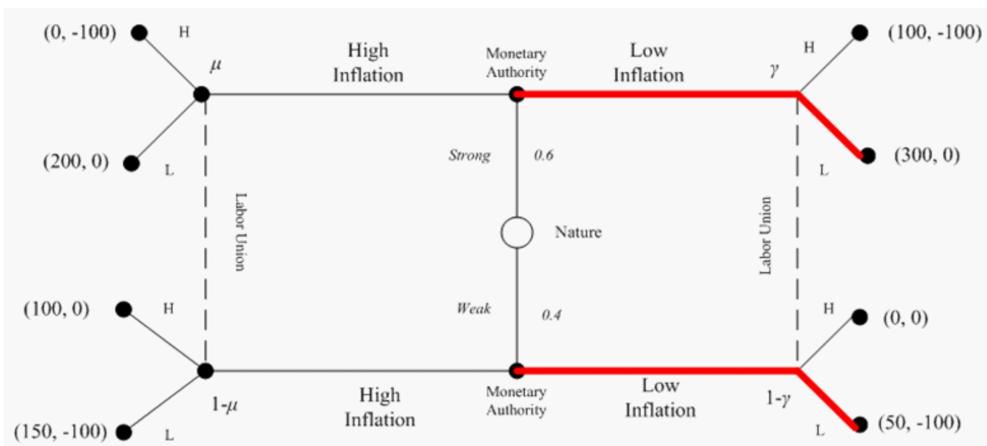

Fig 12. Pooling strategy profile $Low^S Low^W$.

- Upon observing a *high-inflation announcement*, since off-the-equilibrium beliefs satisfy $\mu \in [0, 1]$, the labor union must select whether to respond with H or L by computing the expected utility of each response, as follows.

$$EU_{Labor}(H|Low) = \mu \times (-100) + (1 - \mu) \times 0 = -100\mu$$
$$EU_{Labor}(L|Low) = \mu \times 0 + (1 - \mu) \times (-100) = -100 + 100\mu$$

  Given these expected payoffs, the labor union responds with a high salary increase (H) if and only if $-100\mu > -100 + 100\mu$, or $\frac{1}{2} > \mu$. We will then need to divide the next (fourth) step, where we analyze the optimal announcements of the monetary authority, into two cases: (1) $\mu < \frac{1}{2}$, where the labor union responds with H after observing a high-inflation announcement; and (2) $\mu \geq \frac{1}{2}$, where the labor union responds with L after observing a high-inflation announcement.

Fourth step. Given the uninformed player's optimal responses, we can now determine the informed player's optimal messages.

1. **CASE 1** $\mu < \frac{1}{2}$: These off-the-equilibrium beliefs induce the labor union to respond with H after observing a high-inflation announcement (to facilitate comparison, figure 13 shades the branches corresponding to H in the left-hand side of the figure). Let us next check if either type of monetary authority has incentives to deviate from the prescribed pooling strategy profile.

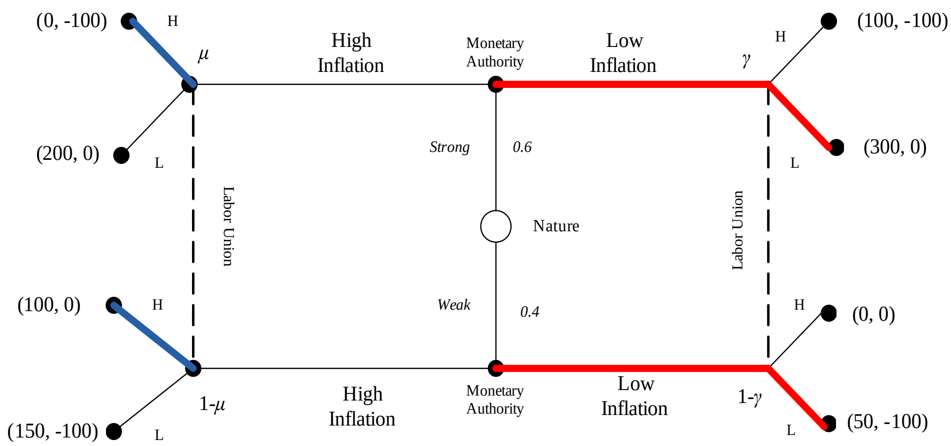

Fig 13. Pooling strategy profile $Low^S Low^W$ when $\mu < \frac{1}{2}$.

- *Strong type.* If the Strong monetary authority behaves as prescribed by this strategy profile (announcing a low-inflation target), it can anticipate that such announcement will be responded with a low wage demand, L, ultimately yielding a payoff of \$300. If, instead, it deviates towards a high-inflation announcement, it can anticipate that such an announcement will be responded with a high-wage demand (H) as indicated in the shaded branches in the left-hand side of figure 13 given that we consider the case in which $\mu < \frac{1}{2}$, implying a lower payoff of \$0. Hence, the strong monetary authority does not have incentives to deviate from the prescribed strategy profile.
- *Weak type*. If the Weak monetary authority behaves as prescribed by this strategy profile (pooling with the Strong type), it can anticipate that such announcement will be responded with a low wage demand, L, ultimately yielding a payoff of \$50. If, instead, it deviates towards a high-inflation announcement, it can anticipate that such an announcement will be responded with a high-wage demand (H) as indicated in the shaded branches in the left-hand side of the figure given that we consider the case in which $\mu < \frac{1}{2}$, implying a higher payoff of \$100. Therefore, the weak monetary authority *has* incentive to deviate from the prescribed strategy profile.

- Since one type of privately informed player (Weak monetary authority) has incentives to deviate from the prescribed pooling strategy profile, we conclude that $Low^S Low^W$ *cannot* be supported as a PBE when off-the-equilibrium beliefs satisfy $\mu < \frac{1}{2}$.

2. **CASE 2** $\mu \geq \frac{1}{2}$: These off-the-equilibrium beliefs induce the labor union to respond with L after observing a high-inflation announcement (to facilitate comparison, figure 14 shades the branches corresponding to L in the left-hand side). Let us next check if either type of monetary authority has incentives to deviate from the prescribed pooling strategy profile.

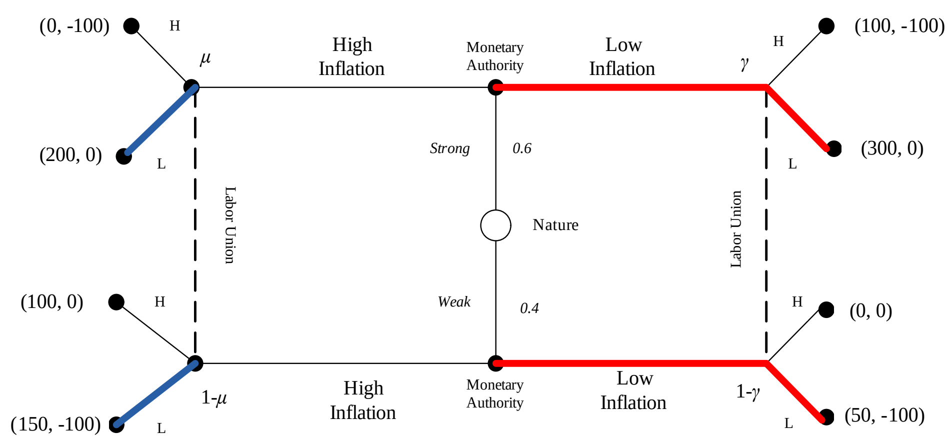

Fig 14. Pooling strategy profile $Low^S Low^W$ when $\mu \geq \frac{1}{2}$.

- *Strong type.* If the Strong monetary authority behaves as prescribed by this strategy profile (announcing a low-inflation target), it can anticipate that such announcement will be responded with a low wage demand, L, ultimately yielding a payoff of \$300. If, instead, it deviates towards a high-inflation announcement, it can anticipate that such an announcement will be responded with a low-wage demand (L) as indicated in the shaded branches in the left-hand side of figure 14 given that we consider the case in which $\mu \geq \frac{1}{2}$, implying a lower payoff of \$200. Hence, the strong monetary authority *does not* have incentives to deviate from the prescribed strategy profile.
- *Weak type*. If the Weak monetary authority behaves as prescribed by this strategy profile (pooling with the Strong type), it can anticipate that such an announcement will be responded with a low wage demand, L, ultimately yielding a payoff of \$50. If, instead, it deviates towards a low-inflation announcement, such an announcement will be responded with a low-wage demand (L) as indicated in the shaded branches in the left-hand side of figure 14 given that we consider the case in which $\mu \geq \frac{1}{2}$,

implying a higher payoff of \$150. Therefore, the weak monetary authority *has* incentives to deviate from the prescribed strategy profile.

- Since we found that one type of privately informed player (the Weak monetary authority) has incentives to deviate from the prescribed pooling strategy profile, we conclude that this strategy profile *cannot* be supported as a PBE when off-the-equilibrium beliefs satisfy $\mu \geq \frac{1}{2}$.

Fifth step. Therefore, the pooling strategy profile $Low^S Low^W$ cannot be supported as a PBE regardless of the off-the-equilibrium beliefs, i.e., cannot be sustained when $\mu < \frac{1}{2}$ or when $\mu \geq \frac{1}{2}$.

References and Notes

1. BAGWELL, K., AND G. RAMEY (1990). “Advertising and pricing to deter or accommodate entry when demand is unknown.” International Journal of Industrial Organization 8, pp. 93-113.

2. BAGWELL, K., AND G. RAMEY (1991). “Oligopoly limit pricing.” The RAND Journal of Economics 22, pp. 155-172.

3. CHO, I. AND D. KREPS (1987) “Signaling games and stable equilibrium,” Quarterly Journal of Economics 102, 179-222.

4. ESPINOLA-ARREDONDO, A. AND F. MUNOZ-GARCIA (2010) “The Intuitive and Divinity Criterion: Interpretation and Step-by-Step Examples,” The B.E. Journal of Industrial Organization Education, vol 5, article 7.

5. FONG, YUK-FAI (2009) “Private Information of Nonpaternalistic Altruism: Exaggeration and Reciprocation of Generosity,” The B.E. Journal of Theoretical Economics, vol. 9, article 1.

6. GAL-OR, ESTHER (1989) “Warranties as a signal of quality”, Canadian Journal of Economics, 22(1), pp. 30-51.

7. HARRINGTON, J.E. JR.. (1986) “Limit pricing when the potential entrant is uncertain of its cost function.” Econometrica 54, pp. 429-437.

8. MILGROM, P., AND J. ROBERTS (1982) “Predation, reputation, and entry deterrence.” Journal of Economic Theory 27, pp. 280-312.

9. Milgrom, P., and J. Roberts (1986) "Price and advertising signals of product quality." Journal of Political Economy 94, pp. 796-821.

10. Ridley, David B. (2008) "Herding versus Hotelling: Market entry with costly information," Journal of Economics and Management Strategy, 17(3), pp. 607-631.

11. Spence, M (1973) "Job market signaling," Quarterly Journal of Economics, 87, pp. 355-374.