Foresight in a Game of Leadership

Logan Perry¹ and Sergey Gavrilets^{1,2,*}

¹Department of Mathematics, Center for the Dynamics of Social Complexity, University of Tennessee, Knoxville, TN 37996 USA

²Department of Ecology and Evolutionary Biology, National Institute for Mathematical and Biological Synthesis, University of Tennessee, Knoxville, TN 37996 USA

*Corresponding author: gavril@utk.edu

ABSTRACT

Leadership can be effective in promoting cooperation within a group, but as the saying goes “heavy is the head that wears the crown.” A lot of debate still surrounds exactly what motivates individuals to expend the effort necessary to lead their groupmates. Evolutionary game theoretic models represent individual’s thought processes by strategy update protocols. The most common of these are random mutation, individual learning, selective imitation, and myopic optimization. Recently we introduced a new strategy update protocol - foresight - which takes into account future payoffs, and how groupmates respond to one’s own strategies. Here we apply our approach to a new 2x2 game, where one player, a leader, ensures via inspection and punishment that the other player, a subordinate, produces collective good. We compare the levels of inspection and production predicted by Nash Equilibrium, Quantal Response Equilibrium, level-k cognition, fictitious play, reinforcement learning, selective imitation, and foresight. We show that only foresight and selective imitation are effective at promoting contribution by the subordinate and inspection and punishment by the leader. The role of selective imitation in cultural and social evolution is well appreciated. In line with our prior findings, foresight is a viable alternative route to cooperation.

Introduction

The web of social and economic ties that weaves us together has never been thicker. Every day we make decisions based on our expectations of how others will act¹. While these waters are difficult to traverse, we are not without aid as there exists a plethora of mores to help us navigate this interpersonal maelstrom. The mores that coordinate political and economic relationships are referred to as institutions²⁻⁴. Whether it is a teacher corralling their students’ behavior or the United Nations placing sanctions on an entire nation, institutions serve as an effective tool for resolving conflict and shaping behavior. The role of institutions in modern society cannot be understated, with some claiming they are the determining factor of whether nations succeed or fail^{3,5}. However, their origins are much older than modern society as they existed even in hunter-gather groups⁶.

Here, we are concerned with the institution of leadership in small-scale societies^{7,8}. A leader is defined as an individual who has non-random differential influence over group behavior⁹. Leaders can take on several different roles within a group, e.g. role-models¹⁰, managers¹¹, punishers¹² or volunteers¹³. Leader-follower relationships are likely to emerge in groups of conspecifics that benefit from acting in a unified manner¹⁴. Examples of actions that necessitate high degrees of coordination include migration, hunting, deterring predation, resolving internal conflicts, and competing with neighboring groups¹⁴.

Joint actions often lead to the collective action problem¹⁵ when all individuals can benefit from an action but no one is willing to bear its cost. The collective action problem (CAP) is present in many animal and human groups¹⁵⁻¹⁹. Several solutions to the CAP exist. They include kin selection, direct and indirect reciprocity, punishment²⁰, selective incentives and institutional design^{15,21}, the presence of within-group heterogeneity^{17,22-24} and influential individuals²⁵.

Here we are focused on the solution by way of punishment and the institution of leadership. In particular, we are concerned with the question of why leaders would choose to enforce cooperation since abstaining from doing it could be less costly - the so called second-order free-rider effect. We investigate this by developing a modified inspection game²⁶.

Inspection games are typically concerned with modeling an *inspector* that seeks to verify the adherence to some pre-arranged contract and an *inspectee* that may be tempted to violate said contract. In the past, inspection games have been used to better understand relationships between law enforcement and criminals²⁷⁻³⁰, employers and their employees³¹, and countries dealing with nuclear armament³². Here we modify the standard inspection game²⁶ to account for interactions between a leader and a subordinate. We do this by altering basic assumptions so that the subordinate has no incentive to produce a good, while the leader has a vested interest in seeing the good produced. Our motivation for doing this is to mirror a standard CAP with punishment in a two person game. In particular we are interested in our model having properties that parallel the first- and second-order free-rider effects. The simplicity of our proposed model allows us the opportunity to obtain analytical results, which can in turn be used to provide a baseline understanding of more complicated models.

Parameter	Interpretation
b	Benefit produced by a subordinate
c	Cost of contributing to a subordinate
d	Punishment for shirking
h	Cost of inspecting to a leader
k	Cost of punishing to a leader
θ	Taxation rate
λ	Precision parameter
ω	Foresight parameter

Table 1. Parameters in the leadership game.

To accurately portray the dynamics of the leader-subordinate relationship we must consider their thought processes. While both individuals are interested in bettering their own positions, the leader has sway over the subordinate and must make decisions based on this influence. This means that the leader must anticipate how their subordinate will react and hence must take into consideration how their subordinate makes decisions. This consideration of other's cognitive processes is referred to as a theory of mind^{33,34}. A theory of mind is thought to be key to promoting cooperation within groups³⁴, has been linked to the size of individuals' social networks³⁵, their propensity for social cooperation³⁶, and the extent to which they are agreeable³⁷. Standard approaches of myopic optimization, adjustment through error and mutation, reinforced learning, and selective imitation fail to account for this theory. Here we continue to develop the strategy update protocol of foresight³⁸, which, we argue, incorporates some aspects of the theory of mind and in doing so solves the second-order free-rider effect.

Foresight works by altering agents' utility functions to consider not only how a strategy will do in this round, but how it will affect the next round as well³⁸. Fundamental to foresight is the fact that agents consider how their actions will shape the future behaviors of others. Thus it accounts both for human's theory of mind and the "shadow of the future"³⁹. The major focus of our earlier work was the impact of foresight on a population facing a collective action problem in the presence of peer punishment. We showed that in heterogeneous groups possessing foresight a division of labor will develop where the strong will specialize in enforcing the contributions of the weak. We now seek to further develop our theory of foresight and obtain analytical results.

Our approach here is broken into several phases. We begin by developing a two-player game based on the classical inspection game²⁶ to investigate the relationship between a leader and their subordinate, which we dub the leadership game. As in our original paper³⁸, we assume heterogeneity, but now take it to be systemic. This is done by pre-designating a leader who serves as an enforcer over a subordinate. We show that the leader lacks motivation to enforce the contribution of their subordinate despite their vested interests in seeing goods produced. Our investigation of the leadership game is structured as follows. First, we consider some typical approaches to analyzing the leadership game, e.g. Nash equilibria, Quantal Response Equilibria (QRE) models, level- k models, and fictitious play models. We follow this up by investigating the notion of foresight, which we previously introduced³⁸. Our results show that the ability of foresight successfully motivates leaders to punish subordinates, and in turn motivates subordinates to produce. Finally, we compare and contrast foresight with two different learning protocols: Cross's reinforcement learning model⁴⁰ and selective imitation⁴¹.

Results

Leadership Game

We consider a simple 2×2 game played between a leader and a subordinate, which is based on the inspection game²⁶ described in the Supplementary Information (SI). The subordinate is tasked with producing a good or benefit at a personal cost to themselves, while the leader has a vested interest in seeing the good is produced. Since we are interested in drawing parallels with collective action problems we make assumptions in such a way that the subordinate has no incentive to see the good be produced unless they are facing punishment.

The subordinate can either produce the good ($x = 1$) or shirk on the production of the good ($x = 0$). In the case the subordinate produces the good, they pay a cost of c to produce a good of value b . Any benefit produced by the subordinate is split with the leader in a $\theta : 1 - \theta$ ratio. Here $0 \leq \theta \leq 1$ can be thought of as a taxation rate. The strategies available to the leader are to enforce production via inspection ($y = 1$) or to not inspect ($y = 0$). Inspection costs the leader h , but in the event that a leader inspects a non-producing subordinate they inflict a punishment of d at a cost of k . We assume that all parameters are positive (see Table 1). Table 2 describes the corresponding payoff matrix.

Subordinate	Leader		
		Inspect	Don't Inspect
	Produce	$(1 - \theta)b - c, \theta b - h$	$(1 - \theta)b - c, \theta b$
	Shirk	$-d, -h - k$	$0, 0$

Table 2. Payoff matrix for the basic leadership game.

The payoff functions for the leader and the subordinate are then

$$\pi_S(x, y) = [(1 - \theta)b - c]x - dy(1 - x), \quad (1a)$$

$$\pi_L(x, y) = b\theta x - [h + k(1 - x)]y. \quad (1b)$$

We will make three assumptions. First, given the subordinate contributes (i.e., $x = 1$), the benefit to the leader exceeds its cost of inspection, i.e., $\theta b > h$. Second, given our aforementioned interest in CAP, we assume that without punishment (i.e., if $y = 0$) the subordinate is not motivated to contribute, i.e., $(1 - \theta)b - c < 0$. Third, facing the threat of punishment (i.e., if $y = 1$), the subordinate however is motivated to contribute, i.e., $(1 - \theta)b - c > -d$.

Pure strategies. Next we derive best response functions for the subordinate, $BR_S(y)$, given the leader's action y and for the leader, $BR_L(x)$, given the subordinate's action x . If the leader inspects ($y = 1$), the subordinate prefers to produce ($x = 1$). If, however, the leader doesn't inspect, then the subordinate's best option is to do nothing. Therefore, $BR_S(y) = y$. This implies that the subordinate can be motivated to produce the good. On the other hand, the leader's best response is not to inspect no matter what the subordinate does: $BR_L(x) = 0$. Therefore, the only Nash equilibrium is $(x^*, y^*) = (0, 0)$.

Mixed strategies. Suppose that the subordinate chooses to produce the good with probability p while the leader opts to inspect with probability q . Then the expected payoffs are

$$E_S(p, q) = p[(1 - \theta)b - c] - (1 - p)qd, \quad (2a)$$

$$E_L(p, q) = p\theta b - q[h + (1 - p)k]. \quad (2b)$$

From the above we can see that the subordinates best response depends upon the leaders strategy. In particular, there is a critical inspection rate of the leader,

$$q_c = \frac{c - (1 - \theta)b}{d}, \quad (3)$$

($0 \leq q_c \leq 1$), such that if $q < q_c$, the subordinate is best off always doing nothing: $BR_S(q) = 0$. If $q > q_c$, the subordinate is best off always producing the good: $BR_S(q) = 1$. If $q = q_c$, the subordinate will receive the same payoff no matter what they do. The case of the leader is much simpler, as the leader's best response is always to do nothing: $BR_L(p) \equiv 0$. Hence, the only Nash equilibrium is at $(p^*, q^*) = (0, 0)$ when the subordinate does not contribute and the leader does not inspect.

Implications for evolutionary dynamics. Assume that individuals are bounded rational and attempt to increase their payoffs by evaluating some "candidate" strategies, which they generate mentally, and choosing one with probabilities proportional to estimated payoffs. (In the terminology of Ref.⁴², this is a direct strategy revision protocol.) The results above imply that the game will converge to the Nash equilibrium $(0, 0)$ of nothing being done.

Quantal Response Equilibrium (QRE)

Next, we generalize our results for the case when agents make errors in evaluating payoffs. We do this by relaxing the assumption that individuals are best responders and replace it with the assumption that individuals are better responders. In this new paradigm, all strategies are played with non-zero probabilities, but the rate at which they are played is proportional to their payoff. This approach leads to investigate what is known as the Quantal Response Equilibrium⁴³ of our model.

Let p be the probability the subordinate contributes, and q be the probability the leader inspects. Let $E_{S,0} = \pi_S(1, q)$ and $E_{S,1} = \pi_S(0, q)$ be the expected payoff of not contributing and contributing to the subordinate, respectively. Define $E_{L,0} = \pi_L(p, 0)$ and $E_{L,1} = \pi_L(p, 1)$ for the leader similarly. In the QRE approach, p and q are specified as

$$p = \frac{1}{1 + \exp[\lambda(E_{S,0} - E_{S,1})]}, \quad q = \frac{1}{1 + \exp[\lambda(E_{L,0} - E_{L,1})]}, \quad (4)$$

where λ is the precision parameter (e.g., with $\lambda = 0$, the players' decision are random: $p = q = 0.5$, while as $\lambda \rightarrow \infty$, both player use myopic best response). Note that we assume both players have the same precision. The QRE solutions for p and q

109 satisfy the equalities⁴³

$$\frac{\ln(p(1-p))}{\lambda} = E_{S,0} - E_{S,1}, \quad \frac{\ln(q(1-q))}{\lambda} = E_{L,0} - E_{L,1}. \quad (5)$$

110 We solve the above equations numerically. Figure S1 in the SI shows the impact that the precision parameter λ has on the QRE
 111 values (p^*, q^*) . For $\lambda = 0$ play is perfectly random so that $p^* = q^* = \frac{1}{2}$, as we would expect. As we increase λ , play converges
 112 to the single Nash equilibrium at $(0, 0)$, again as we would expect. So unless error is very high (i.e., λ is small), there will be
 113 not much inspection or contribution.

114 Before introducing foresight, we show next that just incorporating a theory of mind in our model is not enough to overcome
 115 the free-rider-like effects. We consider two different models attempting to capture some aspects of this theory.

116 Level- k Approach

117 A common way for capturing a theory of mind in models of decision-making is by supposing the agents utilize level- k rationality,
 118 which is a hierarchical way of thinking based on iterative logic⁴⁴. In the most basic model of cognition is level-0 rationality,
 119 which makes agents to just play the strategies available to them at random. It is important to note, that players are never
 120 assumed to be level-0, but rather it is the simplest model of others a player may have. A level-1 player will assume that all
 121 other individuals are level-0 and select a strategy which best responds to their predicted actions. Likewise, a level-2 player will
 122 assume all others are level-1 and select a strategy accordingly. In this way we can iteratively define a level- k player who will
 123 assume all others are level- $k - 1$ and choose a strategy which best responds to their predicted action.

124 We begin by supposing both agents are level-1, which means they assume the other is level-0. Generalizing slightly, we
 125 suppose that a level-0 subordinate is assumed to contribute with probability p_0 , while a level-0 leader is assumed to contribute
 126 with probability q_0 (e.g. $p_0 = q_0 = 0.50$). Now the expected payoffs are given by $E_S(p, q_0)$ and $E_L(p_0, q)$ (See Equations (2a)
 127 and (2b)). This implies the best response for a subordinate is to not contribute provided the expected cost of being punished is
 128 less than the net cost of producing, i.e. $q_0 d < c - (1 - \theta)b$. Likewise, it is best to contribute provided the expected cost of being
 129 punished is greater than the net cost of producing. Meanwhile the leader's expected payoff $E_L(p_0, q)$ is always maximized by
 130 setting $q = 0$ implying it is always best for the leader to not inspect.

131 Now a level-2 subordinate will assume that the leader is level-1, and thus anticipates that the leader will never inspect. This
 132 in turn means that a level-2 subordinate will never contribute. A level-2 leader on the other hand will expect a subordinate to
 133 contribute sometimes, but will always be better off doing nothing. So a level-2 leader will always opt to not inspect. Finally, for
 134 levels 3 onward we have by similar logic that neither player will do anything. Hence, level- k modeling is unable to overcome
 135 the free-rider-like problem in our model.

136 Fictitious Play

137 In the model of fictitious play⁴⁵, every player assumes their opponents are playing strategies drawn from a certain stationary
 138 distribution which the player attempts to estimate via observation. Each player then chooses their action (i.e., a value of x or y)
 139 in an attempt to maximize their payoff given a prediction or assessment of their opponent's strategy.

140 The leader assumes that the subordinate uses a mixed strategy contributing with a certain probability. Let $\tilde{p}(t)$ be the
 141 leader's estimation of this probability at time t . The subordinate assumes that the leader uses a mixed strategy inspecting with
 142 a certain probability. Let $\tilde{q}(t)$ be the subordinate's estimation of this probability at time t . We take $\tilde{p}(0)$, and $\tilde{q}(0)$ to be the
 143 initial beliefs. Let $x(t)$ and $y(t)$ be the action taken, i.e. 0 or 1, by the subordinate and leader, respectively, in round t . Now we
 144 define a system of recurrence equations describing how the leader and subordinate adjust their beliefs based on observations
 145 of previous actions

$$\tilde{p}(t+1) = (1 - \ell)\tilde{p}(t) + \ell x(t), \quad (6a)$$

$$\tilde{q}(t+1) = (1 - \ell)\tilde{q}(t) + \ell y(t). \quad (6b)$$

146 Here, ℓ is a parameter which scales the impact of the most recent action on the agent's estimation. In general, ℓ can depend on t .
 147 For example, $\ell = \frac{1}{1+t}$ corresponds to the original approach⁴⁵. The case of $\ell = 1$ corresponds to best response.

148 Fictitious play itself is then defined as any rule the agent uses to choose a response from the set of best responses to his or her
 149 estimation of the opponent's strategy. For our case, the natural choice of the rule is given by the best response functions $BR_S(q)$
 150 and $BR_L(p)$ established above. Since our game is dominance solvable via iteration (see the SI for details), we know from Ref.⁴⁶
 151 that it will converge to an equilibrium asymptotically. As it has only one Nash equilibrium at $(0, 0)$, it will converge to it. This
 152 makes sense as a rational leader would never choose to inspect as not inspecting offers a higher payoff in all circumstances, and
 153 a rational subordinate would quickly learn this is the case and thus choose to not contribute. Hence, fictitious play modeling is
 154 unable to overcome the free-rider-like problem in our model.

Foresight

We have shown that under myopic optimization, level- k modeling, or fictitious play leaders will fail to enforce and subordinates will fail to contribute. One method to overcoming this is to introduce foresight³⁸. If we assume that the leader is willing to suffer a cost this round in order to make a gain in future rounds, then they could be motivated to inspect the subordinate. More specifically, we introduce the foresight parameter $\omega \in (0, 1)$, which measures the weight placed on next round's *forecasted* payoffs versus this round's anticipated payoff. This averaging of payoff now with payoff later can be compared with the typical practice of discounting future payoffs. Where foresight is particularly novel is how we account for the leader forecasting future payoffs. These forecasted payoffs depend upon the leader's *model* of their subordinate. This consideration of how their subordinate reasons is where our leader's theory of mind is on display. We assume that the leader's model of their subordinate's behavior is based on a best response and focus on the effect foresight has on a leader's strategy selection.

We will assume that only the leaders use foresight. (The SI shows that allowing for the subordinate to use foresight does not change our conclusions.) Consider a weighted sum of the leader's payoffs for this and the next rounds:

$$(1 - \omega) \times (\theta b x - [h + k(1 - x)]y) + \omega \times (\theta b \tilde{x} - [h + k(1 - \tilde{x})]\tilde{y}),$$

where $\tilde{x}$ and $\tilde{y}$ are the subordinate and leader's efforts in the next round. The leader expects that their action this round y will affect the subordinate's action $\tilde{x}$ in the next round. If the subordinate uses best response, as we will assume, $\tilde{x} = y$. At the same time, y has no effect on the benefit to be produced by the subordinate this round, $\theta b x$, or the cost of the inspection in the next round, $[h + k(1 - \tilde{x})]y'$. Therefore we can define the leader's utility function as a weighted sum of the costs of inspection and punishment this round and the benefit next round

$$u_L(x, y) = (1 - \omega) \times (-[h + k(1 - x)]y) + \omega \times (\theta b y). \quad (7)$$

For the subordinate who uses best response, the utility function is equal to the expected payoff given by equation (1a). Table 3 defines the utilities of different actions in this model.

Pure strategies. We can see that the state $(0, 0)$ is still a Nash equilibrium but only if $(1 - \omega)(h + k) > \omega \theta b$, that is, if the foresight parameter is small enough: $\omega < \omega^* \equiv \frac{h+k}{\theta b + h + k}$. However there is a new possibility: if $\omega \theta b > (1 - \omega)h$, or, equivalently, if the foresight parameter is large enough: $\omega > \omega^{**} \equiv \frac{h}{\theta b + h}$, there is another Nash equilibrium $(1, 1)$. For intermediate values of ω , i.e. if $\omega^{**} < \omega < \omega^*$, these two equilibria coexist. The leader does the best at the $(1, 1)$ equilibrium where the payoffs are $u_S = (1 - \theta)b - c$ and $u_L = \theta b - h$. The subordinate does the best at the $(0, 0)$ equilibrium where the payoffs are 0 and 0.

Mixed strategies. Assume that the two players make efforts with probabilities p and q , respectively, and consider formally the corresponding expected utilities:

$$U_S(p, q) = [(1 - \theta)b - c]p - dq(1 - p), \quad (8a)$$

$$U_L(p, q) = (1 - \omega)(-(h + k(1 - p))q) + \omega \theta b q. \quad (8b)$$

Computing the derivative $\frac{\partial U_S(p, q)}{\partial p}$, one finds that the subordinate's utility u_S increases with p for $q > q_c$ and decrease otherwise, where q_c is defined by equation (3). Similarly, computing the derivative $\frac{\partial U_L(p, q)}{\partial q}$, the leader's utility increases with q if $p > p_c$ and decreases otherwise. The critical value

$$p_c = \frac{1}{k} \left(h + k - \frac{\omega \theta b}{1 - \omega} \right). \quad (9)$$

Note that $p_c < 0$ and thus the condition $p > p_c$ is always satisfied if $\omega \theta b > (1 - \omega)(h + k)$, i.e., if the expected future benefit is larger than the current cost of inspection and punishment (or, equivalently, if $\omega > \omega^{**}$). There is a mixed Nash equilibrium ($p^* = p_c, q^* = q_c$), but this equilibrium is unstable: if one player deviates from it, the other player will be motivated to change their strategy as well.

Implications for evolutionary dynamics. These results imply that in corresponding evolutionary models utilizing direct strategy revision protocols⁴², depending on parameters and initial conditions, the system can go to either $(0, 0)$ or $(1, 1)$ state.

To illustrate these possibilities, assume that the subordinate always plays the best response to the leader's previous action, i.e. $\tilde{x} = y$. Let the leader use a mixed strategy q . There are four possibilities for a combination of x and y entering the leader's utility equation (7): $(0, 0)$, $(0, 1)$, $(1, 0)$ and $(1, 1)$ with probabilities $(1 - q)^2$, $q(1 - q)$, $q(1 - q)$ and q^2 , respectively. [Note that x is equal to y in the previous time step.] Therefore the expected utility to strategy q

$$U_L(q) = (1 - q)^2 u_L(0, 0) + q(1 - q) u_L(0, 1) + q(1 - q) u_L(1, 0) + q^2 u_L(1, 1) = k(1 - \omega)q(q - q^*).$$

where $q^* = p_c$. That is, function $U_L(q)$ is quadratic with a maximum at $q = 1$.

Subordinate	Leader		
		Inspect	Don't Inspect
	Produce	$(1 - \theta)b - c, -(1 - \omega)h + \omega\theta b$	$(1 - \theta)b - c, 0$
	Shirk	$-d, -(1 - \omega)(h + k) + \omega\theta b$	$0, 0$

Table 3. Utility matrix for the leadership game with foresight.

If candidate strategies evaluated by the leader deviate only slightly from their current strategy, the dynamics will proceed in the direction of the gradient of $U_L(q)$. That is, if $q^* < 0$, or equivalently $\theta b > (1 - \omega)(k + h)$, q will evolve to 1 for any initial condition. If $\theta b < (1 - \omega)(k + h)$, q evolves to 1 if it exceeds $q^*/2$ initially. These conclusions are not affected qualitatively if the leader make errors in predicting the subordinate's behavior (see the SI). If candidate strategies can deviate from the current strategy substantially, reaching the state $q = 1$ can happen quickly and for any initial condition. The evolution towards the state where both players always make maximum efforts is the new dynamical feature made possible by the leader's foresight.

Learning

Finally we compare foresight with two models of learning: reinforcement learning⁴⁰ and payoff-biased selective imitation⁴¹.

Reinforcement Learning

In reinforcement learning, agents form opinions of strategies based on the payoffs received when those strategies are implemented. Following Ref.⁴⁰, let $u_X(x, y)$ be the utility to player $X \in L, C$ when the subordinate plays x and the leader plays y . Let $p(t)$ be the probability the subordinate contributes and $q(t)$ be the probability the leader inspects at time step t . Then in Cross' learning process⁴⁰ version of our model the probabilities p and q change according to stochastic equations

$$p(t+1) = xu_S(x, y) + (1 - u_S(x, y))p(t), \quad (10a)$$

$$q(t+1) = yu_L(x, y) + (1 - u_L(x, y))q(t). \quad (10b)$$

What this means is that, after players observe how their current action (i.e. x or y) did, they update their state (i.e., p or q) by taking a weighted average between their old state and the state that puts all the weight on the current action (either 0 or 1), where utility $U_X(x, y)$ serves as the weight. This approach requires⁴⁰ that all utilities are scaled to be between 0 and 1. We can achieve this, e.g. by defining them as

$$u_S(x, y) = \frac{\exp\{\lambda \pi_S(x, y)\}}{\sum_x \exp(\lambda \pi_S(x, y))}, \quad u_L(x, y) = \frac{\exp\{\lambda \pi_L(x, y)\}}{\sum_y \exp(\lambda \pi_L(x, y))},$$

where λ is a parameter. By this construction, utility increases with the payoff, and all utilities fall between 0 and 1.

In the continuous time limit, stochastic system (10) can be approximated^{40, 47} by deterministic differential equations:

$$\frac{dp(t)}{dt} = p(t)(U_S(x=1) - \bar{U}_S), \quad (11a)$$

$$\frac{dq(t)}{dt} = q(t)(U_L(y=1) - \bar{U}_L). \quad (11b)$$

where

$$U_S(x=1) = (1 - q)u_S(1, 0) + qu_S(1, 1),$$

$$U_L(y=1) = (1 - p)u_L(0, 1) + pu_L(1, 1)$$

are the expected utilities of strategies $x = 1$ and $y = 1$, and

$$\bar{U}_S = pq u_S(1, 1) + p(1 - q) u_S(1, 0) + (1 - p)q u_S(0, 1) + (1 - p)(1 - q) u_S(0, 0),$$

$$\bar{U}_L = pq u_L(1, 1) + p(1 - q) u_L(1, 0) + (1 - p)q u_L(0, 1) + (1 - p)(1 - q) u_L(0, 0)$$

are the expected utilities of subordinates and leaders, respectively.

We analyzed both the stochastic and deterministic versions of this model. Stochastic numerical simulations of equations (10) show that the system always converges to equilibrium (0, 0) (see the SI). This conclusion is supported by linear stability analysis of equilibria of equations (11): the only stable equilibrium is (0, 0) (see the SI). We conclude that reinforcement learning is unable to overcome the free-rider-like problem in our model.

Selective Imitation

Here we assume that individuals compare their payoff with that of a peer and choose to either copy the selected individual (if their payoff is higher than the focal individual's) or keep their own strategy⁴¹.

Consider a population of pairs each consisting of a subordinate and a leader. Let the i -th pair's actions at time t be denoted by $(x_i(t), y_i(t))$ where all variables retain their usual interpretations. For each individual i , randomly select a peer j (i.e. leader for leader and subordinate for subordinate) and have the focal individual observe that peer's payoff and action. Then the probability that at time t subordinate i continues playing their current action given they observed subordinate j is

$$p_t(i \rightarrow i|j) = \frac{\exp\{\lambda \pi_S(x_i(t), y_i(t))\}}{\exp\{\lambda \pi_S(x_i(t), y_i(t))\} + \exp\{\lambda \pi_S(x_j(t), y_j(t))\}},$$

which of course means the probability i switches to mimic j is $p_t(i \rightarrow j|j) = 1 - p_t(i \rightarrow i|j)$. A similar equation describes changes in the leader's probability q_i given they observe leader j' . As above, λ measures precision in payoff comparisons.

Assume that leaders and subordinates update their strategies at the same rate. In numerical simulations (not shown) the leaders and subordinates opt to do nothing. However if subordinates always play the best response to their leader's previous action (i.e., if $x = y_{\text{prev}}$), the system can evolve to a state with nonzero efforts. (See the SI for numerical illustrations.)

Pure strategies. Consider the case of pure strategies: inspect and not inspect. Given our assumptions about parameters and best response in the subordinates, the former strategy always has a higher payoff than the latter. Therefore the frequency of leaders who inspect will always increase (subject to stochastic errors). The larger precision parameter λ , the faster it happens.

Mixed strategies. If leaders use mixed strategies, then using an approach similar to the one we applied to analysing foresight, there are four possible combinations of x and y in the equation for the leader's payoff $\pi_L(x, y)$: $(0, 0)$, $(0, 1)$, $(1, 0)$ and $(1, 1)$ with probabilities $(1 - q)^2$, $q(1 - q)$, $q(1 - q)$ and q^2 , respectively. Therefore the expected payoff to the leader's strategy q is

$$\Pi_L(q) = (1 - q)^2 \pi_L(0, 0) + q(1 - q) \pi_L(0, 1) + q(1 - q) \pi_L(1, 0) + q^2 \pi_L(1, 1) = kq(q - q^{**}).$$

where $q^{**} = 1 - \frac{\theta b - h}{k}$. This is a quadratic maximized at $q = 1$. If variation in q in the population is small, q will evolve in the direction of the gradient of $\Pi_L(q)$. That is, if the cost of punishment is small (i.e., $k < \theta b - h$), q^{**} is negative, and $\Pi_L(q)$ is always increasing with $q \in [0, 1]$. Thus, q is expected to evolve by selective imitation to $q = 1$ for any initial value. If the cost of punishment is large (i.e., $k > \theta b - h$), then $0 \leq q^{**} \leq 1$. So q will increase to one for initial $q > q^{**}/2$, but will decrease to zero for initial $q < q^{**}/2$. If variation in q in the population is large, the population will always evolve towards increasing q .

That is, with the best response in subordinates and selective imitation in leaders the dynamics are similar to those under foresight. In both cases, the system can evolve to state $(1, 1)$.

Discussion

Here we have studied the impact of foresight on leader-subordinate dynamics in some simple models. Our aim in doing this was to shed light on what can motivate individuals to enforce contribution to production of a collective good. Typically, such enforcement comes with an inherent cost that discourages group members from being coercive as they seek to avoid the cost. This is known as the second-order free-rider problem. Earlier work highlighted several mechanisms such as meta-punishment¹², conformism⁴⁸, signaling⁴⁹, and group-selection^{50,51} as potential routes to overcoming the second-order free-rider problem. We have shown here as well as in Ref.³⁸ that, in addition to these mechanisms, foresight is an effective way of motivating enforcement of cooperation. "Foresight" refers to a novel strategy update protocol, which stresses two key components³⁸. First, that individuals care about their future payoffs. Second, that individuals consider how others will respond to their present actions in future interactions. Both of these are fairly intuitive assumptions that make few requirements of agent's cognitive abilities. By developing foresight we sought to incorporate the deterrence theory⁵² into our model, which is the notion that punishment is used to modify the future behavior of the target.

We approached this problem by altering the payoffs and assumptions of the inspector game²⁶. In particular, we were concerned with modeling the interaction between a single leader interested in enforcing production and a single commoner tasked with producing some good. In this way we were able to incorporate characteristics of the general collective action problem (namely the first- and second-order free-rider effects) into a simple 2x2 game. While our earlier work³⁸ has relied exclusively on numerical simulations, the simplicity of our models here has allowed us to get some analytical results.

Our models can be interpreted as describing a simple case of institutionalized punishment. There are both similarities but also differences with earlier evolutionary studies of social institutions. In our paper, the evolving part of the institution of leadership was the level of monitoring which translated into punishment levels in a way similar to that in Ref.^{53,54}. In Ref.^{11,55} it was the tax imposed by the leaders while in Ref.⁵⁶ it was the proportion of public goods invested into the group's growth rate. In Ref.^{11,55,56} players inherited their strategies from parents (subject to rare random mutation). In Ref.^{53,54} players used payoff-biased imitation. In contrast, we have considered and compared a number of different strategy revision protocols.

We started by considering several different ways of simulating human behaviour, namely Nash equilibria, Quantal Response Equilibria, level- k cognition and fictitious play. Our results show that each of these methods were vulnerable to the second-order free-rider problem. That is, in these basic models while the subordinate could be motivated to produced the good, the leader was not inclined to enforce production and as a result nothing got done. We proceeded by analyzing the effect foresight had on the best response functions and the Nash equilibrium. Upon introducing foresight, we saw that the leader now viewed punishment as an utility increasing action and thus (provided sufficient emphasis on future payoffs) willingly enforced production of their subordinate. Foresight in the subordinate only served to lessen the magnitude of the their payoff. This difference in impact is due the fact that the subordinate's action do not directly influence the leader on the same scale as the leader's action directly influence the subordinate. Our main results are that the introduction of foresight produced new Nash equilibria at which leaders led and subordinates followed. These new equilibria were found to be dependent upon the emphasis placed on future payoffs measured by parameter ω . Additionally, we found that in the repeated leadership game that foresight could effectively overcome the second-order free-rider effect. Even when error was introduced into the leader's predictions, they were motivated to inspect provided certain conditions were met. Our final task was to compare foresight with two other strategy update protocols: reinforcement learning and selective imitation. Our results show that reinforcement learning is not able to overcome the second-order free-rider effect. In contrast, selective imitation is able to accomplish that.

Earlier Ref.^{53,54} studied similar models of institutionalized punishment with multiple subordinates per leader. They showed that selective imitation can lead to the evolution of punishment if leaders update their strategy at a much slower rate than subordinates. This happens because a low update rate prevents the leaders from abandoning a costly punishment strategy before subordinates have learned to contribute to avoid punishment. In contrast, in our model of selective imitation subordinates do not have to learn from others via incremental improvements to adapt but rather they use the best response to the current strategy of the leader. This introduces a new Nash equilibrium which can be then discovered by some leaders via random innovation and then spread across the whole system by imitation. In a similar way, foresight in leaders results in the appearance of a new Nash equilibrium discoverable by leaders via, say, a process of mental scenario building by considering several candidate strategies and comparing their expected utilities (i.e., by using a direct strategy revision protocol *sensu* Ref.⁴²).

Our approach is related to models of level- k cognition^{44,57}. Specifically, best response utilized by subordinates can be viewed as a level-1 strategy to level-0 players who do not change their strategies while foresight in leaders is related to level-2 reasoning. Typically level- k model assume that level-0 players choose their strategies uniformly randomly. Were we to make this assumption, neither inspection nor production would happen in our model. Thus, our work shows that the exact assumptions placed on the level-0 players strongly impacts the overall dynamics of the game.

Overall, our work highlights the importance of strategy revision protocols in evolutionary dynamics⁴². While the free-rider problems exists regardless of the strategy revision protocol employed, the assumptions made on how people think can impact how effective groups are at overcoming these problems. Our protocol of foresight is a new way to consider how people think, which can be used in conjuncture with existing strategy revision protocols.

There are several different questions of interest that must be answered by future work. First and foremost is the question of whether foresight would evolve in a population where it is initial absent. In our current and prior paper, we have taken for granted that foresight is present and sought only to show how it could be an effective route to overcoming the second-order free-rider problem. Having proved its efficacy we should now turn our attention to whether or not its emergence is a reasonable assumption. Secondly, here we considered the leadership game for only two players. A reasonable extension would be to assume multiple agents acting in the role of subordinates (and potential in the role of leaders as well). Thirdly, our results indicate that foresight can affect the basic dynamics of a game (in that it alters the Nash equilibria). It would be a worthwhile endeavour to investigate the impact foresight has on a wider range of classical games.

References

1. Schick, F. Democracy and interdependent preferences. *Theory and Decision* (1972).
2. Schelling, T. *The Strategy of Conflict* (Harvard University Press, 1960).
3. North, D. *Institutions, Institutional Change and Economic Performance (Political Economy of Institutions and Decisions)* (Cambridge University Press, 1990).
4. Hurwicz, L. Institutions as families of game forms. *The Jpn. Econ. Rev.* **47**, 113–132 (1996).
5. Acemoglu, D. & Robinson. *Why Nations Fail: The Origins of Power, Prosperity and Poverty* (Crown Publishers, New York, 2011).
6. Kaplan, H., Gurven, M., Hill, K. & Hurtado, A. The natural history of human food sharing and cooperation: A review and a new multi-individual approach to the negotiation of norms. In Gintis, H., Bowles, S., Boyd, R. & Fehr, E. (eds.) *Moral sentiments and Material Interests: The Foundations of Cooperation in Economic Life*, 75–113 (MIT Press, 2005).

- 321 7. Glowacki, L. & von Rueden, C. Leadership solves collective action problems in small-scale societies. *Philos. Transactions*
322 *Royal Soc. Lond. B* **370**, 20150010 (2015).
- 323 8. Garfield, Z. H., Hubbard, R. L. & Hagen, E. H. Evolutionary models of leadership: Tests and synthesis. *Hum. Nat.* **30**,
324 23–58 (2019).
- 325 9. Smith, J. E. *et al.* Leadership in mammalian societies: emergence, distribution, power, and payoff. *Trends Ecol. & Evol.*
326 **31**, 54–66 (2016).
- 327 10. Zimmerman, M. & Eguiluz, V. Cooperation, social networks, and the emergence of leadership in a prisoner's dilemma
328 with adaptive local interactions. *Phys. Rev. Lett.* **72**, 056118 (2005).
- 329 11. Hooper, P. L., Kaplan, H. S. & Boone, J. L. A theory of leadership in human cooperative groups. *J. Theor. Biol.* **265**,
330 633–646 (2010).
- 331 12. Boyd, R. & Richerson, P. J. Punishment allows the evolution of cooperation (or anything else) in sizable groups. *Ethol.*
332 *Sociobiol.* **13**, 171–195 (1992).
- 333 13. Archetti, M. & Scheuring, I. Coexistence of cooperation and defection in public goods games. *Evolution* **65**, 1140–1148
334 (2011).
- 335 14. van Vugt, M. Evolutionary origins of leaders and followership. *Pers. Soc. Psychol. Rev.* **10**, 354–371 (2006).
- 336 15. Olson, M. *The logic of collective action: Public Goods And The Theory of Groups* (Harvard University Press, Cambridge,
337 MA, 1965).
- 338 16. Bowles, S. & Gintis, H. *A cooperative species: Human reciprocity and its evolution* (Princeton University Press, Princeton,
339 NJ, 2011).
- 340 17. Gavrillets, S. & Fortunato, L. A solution to the collective action problem in between-group conflict with within-group
341 inequality. *Nat. Commun.* **5**, article 3526 (2014).
- 342 18. Henrich, J. P. *Foundations of human sociality: Economic experiments and ethnographic evidence from fifteen small-scale*
343 *societies* (Oxford University Press on Demand, Oxford, United Kingdom, 2004).
- 344 19. Willems, E. P. & van Schaik, C. P. Collective action and the intensity of between-group competition in nonhuman primates.
345 *Behav. Ecol.* **26**, 625–631 (2015).
- 346 20. McElreath, R. & Boyd, R. *Mathematical models of social evolution. A guide for the perplexed* (Chicago University Press,
347 Chicago, 2007).
- 348 21. Ostrom, E. Collective action and the evolution of social norms. *J. Econ. Perspectives* **14**, 137–158 (2000).
- 349 22. Andreoni, J. Privately provided public goods in a large economy: the limits of altruism. *J. Public Econ.* **35**, 57–73 (1988).
- 350 23. Gavrillets, S. Collective action problem in heterogeneous groups. *Proc. Royal Soc. Lond. B* **370**, 20150016 (2015b).
- 351 24. McGinty, M. & Milam, G. Public goods provision by asymmetric agents: experimental evidence. *Soc. Choice Welf.* **40**,
352 1159–1177 (2013).
- 353 25. Gilby, I. C. *et al.* “Impact hunters” catalyse cooperative hunting in two wild chimpanzee communities. *Philos. Transactions*
354 *Royal Soc. B* **370**, article 20150005 (2015).
- 355 26. Fudenberg, D. & Tirole, J. *Game Theory* (The MIT Press, Cambridge, MS, 1992).
- 356 27. Tsebelis, G. The abuse of probability in political analysis: The Robinson Crusoe fallacy. *The Am.* **83**, 77–91 (1989).
- 357 28. Tsebelis, G. Are sanctions effective? *J. Confl. Resolut.* **34**, 3–28 (1990).
- 358 29. Tsebelis, G. Penalty has no impact on crime. *Ration. Soc.* **2**, 255–286 (1990).
- 359 30. Holler, M. J. Fighting pollution when decisions are strategic. *Public Choice* **76**, 347–356 (1993).
- 360 31. Borch, K. Insuring and auditing the auditor. In Deistler, M., Furst, E. & Schwodiauer, G. (eds.) *Games, Economic*
361 *Dynamics, and Time Series Analysis* (Physica, Heidelberg, 1982).
- 362 32. Avenhaus, R., Caty, M., Kilgour, D. M., von Stengel, B. & Zamir, S. Inspection games in arms control. *Eur. J. Oper. Res.*
363 **90**, 383–394 (1996).
- 364 33. Premack, D. & Woodruff, G. Does the chimpanzee have a theory of mind? *Behav. Brain Sci.* **1**, 515–526 (1978).
- 365 34. Tomasello, M., Carpenter, M., Call, J., Behne, T. & Moll, H. Understanding and sharing intentions: The origins of cultural
366 cognition. *Behav. Brain Sci.* **28**, 675–691 (2005).

35. Stiller, J. & Dunbar, R. Perspective-taking and memory capacity predict social network size. *Soc. Networks* **29**, 93–104 (2007).
36. Paal, T. & Bereczkei, T. Adult theory of mind, cooperation, Machiavellianism: The effect of mindreading on social realtions. *Pers. Individ. Differ.* **43**, 541–551 (2007).
37. Nettle, D. & Liddle, B. Agreeableness is related to social-cognitive, but not social-perceptual. *Eur. J. Pers.* **22**, 323–335 (2008).
38. Perry, L., Shrestha, M. D., Vose, M. D. & Gavrillets, S. Collective action problem in heterogeneous groups with punishment and foresight. *Journal of Statistical Physics* **172**, 293–312 (2018).
39. Axelrod, R. *The Evolution of Cooperation* (Basic Books, New York, NY, 1984).
40. Borgers, T. & Sarin, R. Learning through reinforcement and replicator dynamics. *Journal of Economic Theory* (1997).
41. Hofbauer, J. & Sigmund, K. *Evolutionary games and population dynamics* (Cambridge University Press, Cambridge, 1998).
42. Sandholm, W. H. *Population games and evolutionary dynamics* (MIT press, Cambridge, MA, 2010).
43. Goeree, J. K., Holt, C. A. & Smith, A. M. An experimental examination of the volunteer's dilemma. *Games Econ. Behav.* **102**, 303–315 (2017).
44. Stahl, D. O. & Wilson, P. W. On players models of other players: theory and experimental evidence. *Games Econ. Behav.* **10**, 218–254 (1995).
45. Fudenberg, D. & Levin, D. *The Theory of Learning in Games* (MIT Press, 1998).
46. Nachbar, J. “Evolutionary” selection dynamics in games: Convergence and limit properties. *Int. J. Game Theory* **19**, 59–89 (1990).
47. Taylor, M. *Anarchy and cooperation* (Wiley, New York, NY, 1976).
48. Henrich, J. & Boyd, R. Why people punish defectors: weak conformist transmission can stabilize costly enforcement of norms in cooperative dilemmas. *J. Theor. Biol.* **208**, 79–89 (2001).
49. Gintis, H., Smith, E. A. & Bowles, S. Costly signaling and cooperation. *J. Theor. Biol.* **213**, 103–119 (2001).
50. Boyd, R., Gintis, H., Bowles, S. & Richerson, P. J. The evolution of altruistic punishment. *Proc. Natl. Acad. Sci. United States Am.* **100**, 3531–3535 (2003).
51. Sober, E. & Wilson, D. S. *Unto Others: The Evolution and Psychology of Unselfish Behavior* (Harvard University Press, Cambridge, MC, 1999).
52. Axelrod, R. An evolutionary approach to norms. *American Political Science Review* **80**, 1095–1111 (1986).
53. Isakov, A. & Rand, D. The evolution of coercive institutional punishment. *Dyn. Games Appl.* **2**, 97–109 (2012).
54. Roithmayr, D., Isakov, A. & Rand, D. Should law keep pace with society? Relative update rates determine the co-evolution of institutional punishment and citizen contributions to public goods. *Games* **6**, 124–149 (2015).
55. Powers, S. & Lehmann, L. An evolutionary model explaining the neolithic transition from egalitarianism to leadership and despotism. *Proc. Royal Soc. Lond.* **281**, 20141349 (2014).
56. Powers, S. & Lehmann, L. The co-evolution of social institutions, demography, and large-scale human cooperation. *Ecol. letters* **16**, 1356–1364 (2013).
57. Nagel, R. Unraveling in guessing games: an experimental study. *The Am. Econ. Rev.* **85**, 1313–1326 (1995).
58. Hastings, A. *et al.* Transient phenomena in ecology. *Science* **361**, eaat6412 (2018).

Acknowledgements

We thank A. Bentley, M. Berry, C. Strickland, and Y. Chen for comments and suggestions. Supported by the U. S. Army Research Office grants W911NF-14-1-0637 and W911NF-18-1-0138 and the Office of Naval Research grant W911NF-17-1-0150.

Author contributions statement

L.P and S.G. conceived the original idea, performed the analysis and simulations, and wrote the paper.

Competing Interest Statement

The author's have no competing interests.

Supplementary Information for “Foresight in a Game of Leadership” Logan Perry and Sergey Gavrillets

Here we provide additional information on the following topics discussed in the main paper:

- Classic Inspection Game
- Quantal Response Equilibrium (QRE) solution
- Foresight for subordinate
- Phase plane analysis of the foresight model
- Errors in foresight
- Reinforcement learning dynamics
- Selective imitation dynamics

Classic Inspection Game

In the classic Inspection Game, there is a worker who works for a principal²⁶. The worker can choose to shirk or work, where working costs the worker g and produces output of value v for the principal. The principal can either inspect or not inspect. The act of inspecting costs the principal h but it allows the principal to determine if the worker shirked or not. Now the principal is set to pay the worker a wage of w unless they have evidence that the worker shirked. Additional assumptions are that $w > g > h > 0$. This yields the game as described in Table S1. This game has no pure-strategy equilibrium, but does have a mixed Nash equilibrium at $(x, y) = (\frac{w-h}{w}, \frac{g}{w})$, where x is the probability that the worker shirks, and y is the probability that the principal inspects.

	Principal	
	Inspect	Don't Inspect
	Works	Don't Inspect
Worker	Works	$w - g, v - w - h$
	Shirks	$0, -h$
		$w, -w$

Table S1. Payoff matrix for the classic Inspection Game²⁶.

Quantum Response equilibrium (QRE) solution

Figure S1 shows the impact of the precision parameter λ on the QRE values (p^*, q^*) as found numerically by solving equation () of the main text. For $\lambda = 0$, play is perfectly random so that $p^* = q^* = \frac{1}{2}$, as we would expect. As we increase λ , play converges to the single Nash equilibrium at $(0, 0)$, again as we would expect.

Foresight for subordinate

Assume that the subordinates cares about the future payoff and predicts that the leader will use the best response. Because the leader's best response is always 0, the subordinate predicts that there will be no monitoring and punishment and thus will choose not to contribute.

Phase plane analysis in the foresight model

Let p and q be the probabilities of making an effort (i.e., the mixed strategies) for the subordinate and leader. Their expected utilities are

$$U_S = ((1 - \theta)b - c)p - dq(1 - p) \quad (\text{S1a})$$

$$U_L = (1 - \omega)(-h + k(1 - p))q + \omega\theta bq \quad (\text{S1b})$$

This corresponding derivatives are given by

$$\partial_p U_S = ((1 - \theta)b - c) + dq \quad (\text{S2a})$$

$$\partial_q U_L = (1 - \omega)(-(h + k(1 - p))) + \omega\theta b. \quad (\text{S2b})$$

Figure S1. Effect of the precision parameter λ on the Quantal Response Equilibrium values p^* and q^* . Parameters: $b = 1, c = 2, d = 2, h = 0.1, k = 1$, and $\theta = 0.5$.

As discussed in the main text, $\partial_q U_L < 0$ whenever $p < q^{**} = 1 + \frac{1}{k} \left(h - \frac{\omega \theta b}{1 - \omega} \right)$, and $\partial_q U_L > 0$ whenever $p > q^{**}$, and $\partial_p U_S < 0$ whenever $q < q_c = \frac{c - (1 - \theta)b}{d}$ and $\partial_p U_S > 0$ whenever $q > q_c$. This yields the expected flow on the phase plane (p, q) as outlined in Figure S2.

Figure S2. Flow associated with equations (S1a). Parameters: $b = 1, c = 2, d = 3, k = 1, h = .25, \theta = 0.5$, and $\omega = 0.6$. Red dots mark Nash equilibria two of which are stable $(0, 0)$ and $(1, 1)$ and one unstable at (q^{**}, q_v) .

Errors in foresight

Here we incorporate errors into the leader's predictions of the actions of the subordinate. We can interpret this error in one of two ways. First, it can be that the leader makes mistakes when attempting to decide what the subordinate will do. Second, it can be that the leader anticipates that the subordinate will make mistakes and takes this into account. Either way, we define the probability $p(y)$ the leader will predict the subordinate contributes given the leader inspects with probability q as

$$p(y) = \frac{1}{1 + \exp[\lambda(\pi_S(0, y) - \pi_S(1, y))]} = \frac{1}{1 + \exp[\lambda(-dy - b(1 - \theta) + c)]}.$$

449 That is, we allow for logit errors with precision parameter λ as above.

Using the same approach as in the main text, we write the leader's utility as

$$U_L(y_{prev}, y) = (1 - \omega) \times (-[h + k(1 - p(y_{prev}))]y) + \omega \times (\theta b p(y))$$

where y_{prev} and y are two subsequent values of the leader's y . Then, for a leader using strategy q , the expected utility is

$$U(q) = (1 - q)^2 U_L(0, 0) + q(1 - q) U_L(0, 1) + q(1 - q) U_L(1, 0) + q^2 U_L(1, 1).$$

450 From the above equation we can find the critical value q^* such that for initial values of q above q^* , q is expected to evolve to
 451 1. Figure S3 illustrate the dependence of q^* on λ and ω . This Figure shows that increases ω increases decreases q^* and thus
 452 increases the range of conditions leading to production and inspection. As $\lambda \rightarrow \infty$, the predicted value of q^* converges to that
 453 one found in the main text.

Figure S3. The critical value q^* in the model with errors in the leader's foresight for different values of precision parameter λ . The three curves correspond to $\omega = 0.75, 0.5$ and 0.25 (from top to bottom). Other parameters: $b = 1, c = 2, d = 2, h = 0.25, k = 1$, and $\theta = 0.5$.

454 Reinforcement learning dynamics

455 The system of differential equations (11) in the main text has four equilibria: $(0, 0)$, $(1, 0)$, $(0, 1)$ and $(1, 1)$. Standard stability
 456 analysis shows that the only locally stable equilibrium is $(0, 0)$. For example, the eigenvalues of the Jacobian matrix at this state
 457 are $u_S(1, 0) - u_S(0, 0)$ and $u_L(0, 1) - u_L(0, 0)$, which are always negative by our assumptions. The corresponding eigenvalues
 458 of state $(1, 1)$ are $u_S(0, 1) - u_S(1, 1)$ and $u_L(1, 0) - u_L(1, 1)$, which are always positive by our assumptions.

459 Numerical simulations show that the system may spend very long time in a neighborhood of each of the other three
 460 equilibria, that is, the dynamics are characterized by long transients⁵⁸. For example, Figure S4 shows the stochastic trajectories
 461 in the reinforcement learning model for a range of random initial conditions. With small precision λ , fluctuations in p and q
 462 persist for a long time. Even with large λ some subordinate-leader pairs can stay close to state $(1, 1)$ for a long time.

463 Selective Imitation

464 Consider first the case where both leaders and subordinates update their actions x and y via selective imitation. Figure S5 shows
 465 the trajectories associated with a range of random initial conditions drawn from the uniform distribution over $[0.1, 0.9]$. We see
 466 that regardless of the initial condition, the system eventually settles into the equilibrium of $(0, 0)$.

467 Next consider the case where leaders update via selective imitation, but subordinates update via myopic optimization.
 468 Figure S6 shows the trajectories associated with a range of random initial conditions drawn from the uniform distribution over
 469 $[0.1, 0.9]$. With small precision λ the system may vocationally end up at $(0, 0)$ state. [This is analogous to fixing a deleterious
 470 allele by random genetic drift in population genetics.] With increasing precision λ , the system always evolves to $(1, 1)$ state.

Figure S4. Time-series in the case of reinforcement learning for different values of λ . The graphs show 100 independent runs with initial conditions drawn from the uniform distribution over $[.1, .9]$. The thin lines represent the individual runs while the thick lines represent the mean over all runs with the red corresponding the the leader and the blue to the subordinate. Parameters: $b = 1, c = 2, d = 2, k = .5, h = .25$, and $\theta = 0.5$.

Figure S5. Time-series in the case of selective imitation in both players. The thin lines represent the individual runs while the thick lines represent the mean over all runs with the red corresponding the the leader and the blue to the subordinate. Note that all trajectories tend toward the equilibrium $(0,0)$. Parameters: $b = 1, c = 2, d = 2, k = .5, h = .25, \theta = 0.5$, and $\lambda = 1$.

Figure S6. Time-series in the case of myopic optimization in the subordinate and selective imitation in the leader for different values of λ . The results summarize 100 independent runs with initial conditions drawn from the uniform distribution over $[0.1, 0.9]$. The thin lines represent the average over all 1000 groups while the thick lines represent the mean over all runs with the red corresponding the the leader and the blue to the subordinate. Parameters: $b = 1, c = 2, d = 2, k = .5, h = .25$, and $\theta = 0.5$.