

## DESCRIPTION AND THE PROBLEM OF PRIORS

ABSTRACT. Belief-revision models of knowledge describe how to update one's degrees of belief associated with hypotheses as one considers new evidence, but they typically do not say how probabilities become associated with meaningful hypotheses in the first place. Here we consider a variety of Skyrms-Lewis signaling game [Lewis (1969)] [Skyrms (2010)] where simple descriptive language and predictive practice and associated basic expectations coevolve. Rather than assigning prior probabilities to hypotheses in a fixed language then conditioning on new evidence, the agents begin with no meaningful language or expectations then evolve to have expectations conditional on their descriptions as they evolve to have meaningful descriptions for the purpose of successful prediction. The model, then, provides a simple but concrete example of how the process of evolving a descriptive language suitable for inquiry might also provide agents with conditional expectations that reflect the type and degree of predictive success in fact afforded by their evolved predictive practice. This illustrates one way in which the traditional problem priors may simply fail to apply to one's model of evolving inquiry.

### DESCRIPTION, PREDICTION, AND EXPECTATION

Belief-revision models of knowledge describe how to update one's degrees of belief as one considers new evidence. On a Bayesian model, for example, one fixes a descriptive language, sets coherent prior probabilities over a set of hypotheses expressed in the language, then updates one's degrees of belief as one conditions on new evidence. While such an account of reflective inquiry has many virtues, it has nothing to say concerning how to assign prior probabilities to meaningful hypotheses. This is the problem of priors.

On reflection, one might, however, find the problem of priors itself puzzling. While it is indeed unclear what procedure one should adopt in assigning prior expectations to hypotheses expressed in a fixed descriptive language, it is similarly unclear how one might ever come to use such a language without already having a rich set of expectations. The symmetry of these reflections suggests a strategy.

We will consider how it might be possible for basic expectations to coevolve with a simple descriptive language. More specifically, we will consider how simple

descriptive and predictive practice might coevolve with basic expectations in the context of a sender-predictor game, a variety of Skyrms-Lewis signaling game.<sup>1</sup>

Rather than assigning prior probabilities to hypotheses then conditioning on new evidence, the modeled agents begin with no meaningful language or expectations, then come to form basic expectations conditional on their simple descriptions as they coevolve descriptions that may be used as the basis for successful prediction. These basic conditional expectations reflect the degree of success in fact exhibited in their descriptive and predictive practice and, as such, may evolve to inform other decisions.

Reflective Bayesian inquiry is only possible for agents who have the capacity to represent alternative possible states of the world and associate expectations with these representations. The story here concerns how agents might evolve such capacities in the first place. With regard to the problem of priors, it provides a simple but concrete example of how the evolution of a descriptive language suitable for coordinated prediction might also provide a core set of well-tuned posterior expectations that might then be available to constrain subjective degrees of belief should the agent turn to reflective inquiry.

We will consider the story in three parts.

#### PART I: THE COEVOLUTION OF DESCRIPTION AND PREDICTION

A sender-predictor game is a variety of Skyrms-Lewis signaling game.<sup>2</sup> In a sender-predictor game, however, the agents coevolve both descriptive and predictive dispositions. We will refer to these as the agents' first-order dispositions. Their second-order dispositions determine how they update their first-order dispositions as they learn from experience.<sup>3</sup>

In a two-state/two-signal/two-act sender-predictor game, the sender observes a prior state of nature, then sends a signal. The probability of particular signal being sent is determined by the state of nature and the sender's first-order dispositions.

---

<sup>1</sup>See [Lewis (1969)] for Lewis' characterization of signaling games and [Skyrms (2006)], [Skyrms (2010)], [Barrett (2009)], and [Barrett (2007)] for examples of such games in an evolutionary context.

<sup>2</sup>David Lewis [Lewis (1969)] introduced signaling games in the context of classical game theory as a way to study the possible nature of linguistic convention. Brian Skyrms [Skyrms (2006)] [Skyrms (2010)] later considered Lewis' signaling games in the context of evolutionary game theory. Such evolutionary games have a number of virtues over Lewis' conventional games as one need not assume common knowledge nor any special rational faculties. Rather, the agents have only simple conditional dispositions. The sender-predictor games described in the present paper were coded and run as C++ simulations.

<sup>3</sup>The distinction between first- and second-order dispositions is less clear for agents who might learn how to learn. The agents here are not so subtle. Also note that the evolutionary story here involves the evolution of the dispositions of particular agents rather than the evolution of types of agent in a population. It is often possible to translate results between such models. See [Skyrms (2010)] for a discussion of this point.

On observing the signal, the receiver performs a predictive action that is either successful or unsuccessful depending on the posterior state of nature at a later time. The probability of a particular predictive action being performed is determined by the signal and the predictor's first-order dispositions. If the predictive action is successful, which will depend on both the regularities exhibited by nature and what counts as success for the agents as determined by their second-order dispositions, then the sender and predictor reinforce the first-order dispositions that led to the action they just took conditional on the state of nature and the signal type, respectively. If the predictive action is unsuccessful, they might weaken the first-order dispositions that led to the actions they just took. Precisely how they reinforce or weaken their first-order dispositions is determined by their second-order dispositions.<sup>4</sup>

Consider a sender who checks the water temperature off Newport Pier each morning. If the water is cooler than normal, she draws a ball from one urn; and if it is warmer than normal, she draws a ball from another urn. Suppose that the cool urn and the warm urn each initially contain one blue and one green ball. The sender waves either a blue or green flag depending on the type of ball drawn. Her predicting friend, standing on the beach on Catalina Island, sees the color of the flag through his spyglass and draws a ball from an urn that corresponds to the flag color. Suppose that each of the predictor's urns begins with one ball each of two predictive action types corresponding to going night fishing or to staying home to repair fishing nets. He then performs the action corresponding to the type of ball he drew. Night fishing is successful if there are clear skies that evening; staying home to repair nets is successful if there is fog.<sup>5</sup> If the predictive action is successful, then each agent returns his ball to the urn from which it was drawn and adds a new ball of the same type; otherwise, each agent simply returns his ball to the urn from which it was drawn.<sup>6</sup>

---

<sup>4</sup>See [Barrett (2012)] for more details regarding the set-up and interpretation of such a games.

<sup>5</sup> In terms of the endogenous norms of the model, all it means for a predictive action to be successful is that, given the regularities between the prior and posterior states of nature and the agent's second-order dispositions to update their first-order dispositions, it produces an event that in fact leads to a reinforcement of the agents' first-order dispositions that led to the act. Otherwise, the predictive action is unsuccessful.

<sup>6</sup>This sort of simple reinforcement learning was introduced by Herrnstein in his discussion of the law of effect [Herrnstein (1970)]. More sophisticated learning dynamics also allow for punishment and forgetting. They typically do much better than simple Herrnstein reinforcement in games like those discussed in this paper [Barrett and Zollman (2009)], and, saliently, they often much better model the actual behavior of learners [Roth and Erev (1995)] [Bereby-Meyer, Yoella and Erev (1998)]. The methodological thought is that Herrnstein reinforcement learning requires only relatively weak dispositional resources and if it allows for successful coordinated action in a particular context, then one can expect a broad class of more sophisticated reinforcement dynamics to allow for similar success.




FIGURE 1. The water-fog sender-predictor game

While the sender’s signals are initially meaningless and the receiver’s predictive actions are correspondingly random, as the sender’s dispositions to signal conditional on the state of nature and the receiver’s dispositions to act conditional on the signal evolve, the sender’s signals become meaningful precisely insofar as they may serve as the basis for successful coordinated predictive action. If there is in fact a natural correlation between morning water temperature and evening weather, the agents will be able to exploit this natural regularity for the purpose of successful action if they are able to coevolve appropriately interrelated descriptive and predictive dispositions.

Suppose that nature is such that the morning water temperature simply determines whether or not there will be fog in the evening. More specifically, suppose that cool water guarantees clear weather and warm water guarantees fog. In this case, if lower and higher than normal water temperatures are unbiased, independent, and randomly distributed in nature, then, in this particular game, the agents are guaranteed to coevolve successful descriptive and predictive dispositions in the limit. When they do, one type of signal will almost always be sent when the water is cool and will almost always lead to night fishing, and the other will almost always be sent when the water is warm and will almost always lead to staying home to repair nets.<sup>7</sup>

The game and the agents’ evolved behavior is more subtle when the relationship between morning water temperature and evening fog is stochastic rather than deterministic. This is also a natural context to describe how well-tuned posterior expectations might coevolve with the agents’ successful descriptive and predictive practice.

<sup>7</sup>The deterministic case is formally equivalent to a two-state/two-signal/two action signaling game. See [Argiento, Pemantle, Skyrms, and Volkov (2009)] for a proof of convergence to a signaling system in this case. If there are more than two states, signals, and actions or if nature is biased, then the agents’ dispositions sometimes fail to converge to a signaling system. See [Barrett (2006)] and [Huttegger (2007)] for discussions of such suboptimal pooling equilibria.

## PART II: TRACKING PREDICTIVE SUCCESS

Consider the following implementation of the stochastic water-fog game where the evening weather is only randomly determined by the morning water temperature. Here the sender has one urn for each possible prior water temperature state of nature labeled 0 and 1. Each of her urns begins with one ball of each type of signal she might send labeled *blue* and *green*. The receiver has one urn corresponding to each possible signal type (color). Each of his urns begins with one ball of each type of predictive action he might take labeled  $0'$  and  $1'$ , each of which is optimal in the corresponding posterior weather state.

Suppose that the prior states of nature are equally likely and randomly distributed. Suppose further that nature is such that when prior water temperature state 0 obtains, the most likely posterior evening weather state is  $0'$ , occurring with probability  $T_{00'}$ ; and when prior water temperature state 1 obtains, the most likely posterior evening weather state is  $1'$ , occurring with probability  $T_{11'}$ . The probability of a transition in nature from prior state 0 to posterior state  $1'$  then is  $T_{01'} = 1 - T_{00'}$ , and the probability of a transition in nature from prior state 1 to posterior state  $0'$  is  $T_{10'} = 1 - T_{11'}$ . When the predictor gets a signal of type  $\mathcal{S}$  and performs an action of type  $\mathcal{A}'$ , the action counts as a success for signal type  $\mathcal{S}$  if and only if the posterior state of nature is of type  $\mathcal{A}$ .




FIGURE 2. Crossover probabilities and optimal actions

In this version of the game, we will suppose as before that the agents learn by simple reinforcement. We will suppose further, however, that the predictor has expectation weights associated with each possible signal type and that those weights evolve by bounded reinforcement with forgetting. More specifically, in addition to his action urns, the predictor has an expectation urn for blue signals and an expectation urn for green signals. Each expectation urn begins the game with a random assortment of a thousand white and black balls. If an act is successful, in addition to reinforcing the act he took, the predictor adds a white ball to the expectation urn corresponding to the signal type he just used; if the act was unsuccessful, he adds a black ball to that expectation urn. Then the predictor draws a ball at random out

of the expectation urn he just put a ball in, and discards it. This gradual random forgetting maintains a constant number of balls in each expectation urn.<sup>8</sup>

For now, we will simply suppose that the weight of white balls in each expectation urn represents the degree to which the predictor expects that his predictive action will be successful conditional on his receiving the corresponding signal. By stipulation, then, if the sender’s signals evolve successful operational meanings for the purpose of prediction, then the expectation weights will evolve to indicate the predictor’s conditional expectations on meaningful descriptions of prior states.




FIGURE 3. Water-fog game with expectations

On simulation, the sender and receiver begin by randomly signaling and predicting, and any initial success is the result of blind luck. Over time, however, the agents very often learn to exploit the statistical regularities exhibited by nature. And when they do, the predictor’s conditional expectations on the signal types typically coevolves to agree with the degree of predictive success the agents’ evolved descriptive and predictive practice in fact affords.

In the symmetric case, where each prior state of nature is equally likely and the failure of the most likely posterior state occurring is independent of the prior state (i.e.  $T_{01'} = T_{10'} = T$ ), the agents nearly always evolve on simulation to signal and to predict as well as theoretically possible.<sup>9</sup> More specifically, the sender evolves to associate a different signal type with each of the prior states and nearly always sends that type of signal when the corresponding prior state obtains; and, when the predictor receives the signal, he nearly always performs the predictive action

<sup>8</sup>Note that the learning rule for signaling behavior (unbounded without forgetting) is different from the learning rule for expectations (bounded with forgetting). The thought here is to use the simplest learning rule that allows the agents at least a modest level of success. Since the evolution of successful signaling behavior is possible without bounds or forgetting, the simpler learning rule is used. The agents do much better yet in evolving successful signaling behavior on bounded learning with forgetting [Barrett and Zollman (2009)]. Unbounded learning without forgetting does not typically do well at closely tracking signaling success rates as they evolve.

<sup>9</sup>Suppose  $T_{01'} = T_{10'} = 0.25$ . The theoretically best possible success rate is 0.75. On simulation, the agents evolve to do better than a success rate of 0.70 in 0.881 of the runs with  $1 \times 10^6$  plays/run, 0.924 of the runs on  $1 \times 10^7$  plays/run, and 0.959 of the runs on  $1 \times 10^8$  plays/run with 1000 runs in each case. And the predictor’s expectation conditional on each signal type evolves to agree well with their actual relative frequency of successes on each signal type.

that has been most likely to occur when the prior state corresponding to the signal has obtained.<sup>10</sup>

If clear weather has typically followed cool water and fog has typically followed warm water, then the sender evolves to nearly always send one type of signal whenever she observes cool water and the other whenever she observes warm water; and the predictor evolves to perform the predictive action that is most likely to be successful on a clear evening when he receives the signal associated with cool water and to perform the predictive action that is most likely to be successful on a foggy evening when he receives the signal associated with warm water. The agents then evolve a perfectly sharp descriptive language for prior states, their success rate in predicting posterior states approaches the theoretically optimal rate of  $1 - T$ , and the predictor comes to expect that the probability of an action being successful, conditional on each signal, is  $1 - T$ , which reflects the agents' actual success rate.<sup>11</sup> In this case, then, the agents evolve to make the best possible predictions and their posterior conditional expectations in fact tell them just how reliable their predictions are. Finally, significant to the problem of priors, the evolved conditional expectations here are independent of the initial contents of the predictor's expectation urns. By the time the agents have evolved sharp meanings and clear predictive dispositions, the initial configuration of the expectation urns has been washed out under bounded reinforcement learning with random forgetting.<sup>12</sup>

In the asymmetric case, where the failure of the most likely posterior state occurring depends on the prior state (i.e.  $T_{01'} \neq T_{10'}$ ), the simple-minded reinforcement learning implemented by the simulated agents here may lead them to a suboptimal equilibrium where they do not evolve to signal and predict as well as theoretically possible. But if the statistical asymmetry is modest, even agents who learn in this primitive way typically do evolve to signal and to predict as well as theoretically possible. And when they do, they also coevolve conditional expectations that match the asymmetric success rates for each of their evolved terms.<sup>13</sup> In particular, the

<sup>10</sup>The conjecture here is that this is true for the symmetric case as an extension of the results of [Argiento, Pemantle, Skyrms, and Volkov (2009)].

<sup>11</sup>More precisely, the conditional expectations quickly evolve to approximate the agents' actual success rate, then fluctuate under the random mechanism for forgetting but spend most of the time close to the actual evolving success rate. More precisely, once they begin to track the actual success rate closely, their conditional expectations on the simulations described here nearly always within about 5% of their actual success rate. One would expect yet closer tracking with more balls in the expectation urns.

<sup>12</sup>This is because bounded reinforcement learning with forgetting is much faster here at tracking the evolving success rate than simple reinforcement learning is at evolving successful descriptive and predictive dispositions. See [Barrett and Zollman (2009)] for a description of the relative virtues of learning dynamics that incorporate a relatively strong variety of forgetting.

<sup>13</sup>Even when the agents do not evolve to make predictions in a theoretically optimal way, they still do quite well since they typically evolve a language that successfully communicates the most common state(s) of nature reliably [Huttegger (2007)]. Significantly for the issue of well-tuned

receiver evolves to nearly always do predictive action  $0'$  on the signal that is sent on prior state 0 and predictive action  $1'$  on the signal that is sent on prior state 1; and he also evolves to expect that his action will be successful with probability  $T_{00'} = 1 - T_{01'}$  in the first case and  $T_{11'} = 1 - T_{10'}$  in the second.<sup>14</sup>

In both the symmetric and asymmetric cases, when successful, the agents co-evolve intertwined linguistic and predictive dispositions. The signals are initially meaningless and the predictions made on the basis of these signals are consequently wildly unreliable. Successful prediction coevolves with the meanings of the signals, and the predictor's conditional expectation weights for the success of each type of predictive action conditional on each signal type reflect the type and degree of predictive success. Neither the evolved meaning of the agents' descriptive language nor their predictive dispositions nor the associated expectation weights are somehow prior; rather, the agents' descriptive and predictive practice and their expectations are interdependent and forged by precisely the same empirical evidence in the context of the same evolutionary process.

Since the predictor's expectation weights typically do in fact track the type and degree of predictive success afforded by the agents evolved language and predictive practice, whatever that success may be, if their past relative frequencies of successful coordinated prediction are representative of their future success, then their expectation weights can be expected to reliably indicate what their expectations should be.

### PART III: HOW EXPECTATION WEIGHTS MAY EVOLVE TO REPRESENT EXPECTATIONS

While we have simply supposed that the predictor's expectation weights represent the degree to which he expects that his predictive action will be successful conditional on his receiving the corresponding signal, one might want to say how expectation weights might come to play the role of expectations even as they evolve

---

expectations, their coevolved conditional expectations track their actual degree of predictive success on each signal type for both optimal and suboptimal evolutions. So while the simulated agents may not know that their evolved descriptive and predictive practices are suboptimal, their expectations will reflect the reliability of the predictions based on those practices whatever this may be.

<sup>14</sup>Suppose  $T_{01'} = 0.25$  and  $T_{10'} = 0$ . The theoretically best possible success rate is 0.875. On simulation, the agents evolve to do better than a success rate of 0.825 in 0.751 of the runs with  $1 \times 10^6$  plays/run, 0.777 of the runs on  $1 \times 10^7$  plays/run, and 0.773 of the runs on  $1 \times 10^8$  plays/run with 1000 runs in each case. A little less than a quarter of the time, then, the agents fail to evolve those distinctions that would allow them to make optimal predictions. But when they do evolve to capture the relevant distinctions, they also learn that the term associated with prior state 0 leads to a successful predictive action about 0.75 of the time and the term associated with prior state 1 nearly always leads to a successful predictive action. the modest success of the agents in asymmetric cases is the result of their simple-minded learning dynamics. If they use bounded reinforcement learning with forgetting to evolve their signaling behavior, they typically involve nearly optimal dispositions on simulation.

to better track the coevolving success of the agents' evolving descriptive and predictive practice.

A natural thought is to appeal to the fact that the predictor will do well when he uses his expectation weights as expectations since they will in fact evolve to represent his actual degree of predictive success. But to say more precisely why predictor might evolve to use expectation weights as expectations as opposed to doing something else, one must provide a specific potential use for his evolving expectation weights and at least one alternative to his not using the expectations weights in the standard way. A simple concrete example is perhaps instructive.

Suppose that the predictor is not only concerned with whether he should go fishing or mend nets but also in making side wagers concerning the success of his predictive actions.<sup>15</sup> Suppose further that his second-order dispositions are such that he begins by typically flipping a fair coin to determine whether to accept a particular wager but that he also starts with a small positive probability of using his evolving expectation weights as if they were expectations to determine whether to accept the wager. Suppose that the probability that he will flip a coin or use expectation weights to determine whether to accept a wager in future plays increases or decreases slightly, proportional to his winnings, when each method is used to regiment his behavior. And, finally, suppose that the probability of regimenting his wagers in each way always remains positive.

Such dispositions might be represented in a learning game as follows. The predictor has a decision urn that determines whether he flips a coin or uses expectation weights to regiment his betting behavior on a particular play. The decision urn starts with 999,999 balls that indicate that the predictor flip a coin and one ball that indicates that he use expectation weights as expectations. When presented with a wager, the predictor draws a ball from the decision urn. If the ball indicates that he flip a coin, he does so and accepts the wager on heads and rejects in on tails, whatever it is. If the ball indicates that he use expectation weights as expectations, he accepts a wager of  $a$  of his dollars against  $b$  of his friend's if and only if the ratio of the number of white balls to the total number of balls in the expectation urn corresponding to the signal he just received is greater than  $a/(a + b)$ . Suppose that the stakes of the wager are always \$10 and that each bet is whole valued and randomly determined with equal likelihood.

---

<sup>15</sup>One of the anonymous reviewers suggested a slightly more complicated, but much better, story. On this telling, the side wagers might be represented in the predictor's wish to anticipate future catches in order to reserve adequate, but not unnecessary, transportation of his catch to market. What matters to the model is just that one allow for the possibility that the predictor's evolved expectations in fact be relevant to his actions and that the way in which they are relevant might itself evolve subject to the success and failure of those actions.

In order to capture the predictor's second-order dispositions as characterized above, we will suppose that the contents of the decision urn evolve by proportional reinforcement with forgetting. In particular, if the predictor flips a coin and wins the wager, he adds  $b$  balls of the type that indicate that he flip a coin to the decision urn, then he discards  $b$  balls from the urn at random. And if he flips a coin and loses, he adds  $a$  balls of the type that indicate that he use expectation weights as expectations to the decision urn, then discards  $a$  balls from the urn at random. Similarly, if the predictor uses expectation weights as expectations and wins, he adds  $b$  balls of the type that indicate that he use expectation weights as expectations to the decision urn, then he discards  $b$  balls from the urn at random. And if he uses expectation weights as expectations and loses, he adds  $a$  balls of the type that indicate that he flip a coin to the decision urn, then discards  $a$  balls from the urn at random. Finally, suppose that the agent never discards from the decision urn the last ball of any type.

If there are in fact regular statistical correlations between prior and posterior states, then the predictor typically evolves to make the most accurate statistical predictions possible and the expectation weights coevolve to track how accurate his predictions are. Since expectation weights evolve to indicate the agents' actual type and degree of success, one should expect the predictor to do better than even in the long run when he uses expectation weights as expectations to determine whether to accept a wager. Since the wagers offered the predictor are unbiased and since flipping a fair coin leads to an unbiased acceptance of wagers, one should expect the predictor to be even in the long run when he flips a coin to determine whether to accept a wager. Hence the proportion of balls in the decision urn that indicate that the predictor use expectation weights as expectations tends to increase over time. The result is that the predictor typically ends up using expectation weights as expectations.<sup>16</sup>

The upshot is that on this simple model the agents end up with relatively stable, well-tuned expectation weights that typically get used as expectations without there ever being anything like an assignment of prior expectations over possible descriptions. Note, in particular, that the expectation weights are not initially

---

<sup>16</sup>Note, however, that since the predictor might use expectation weights as expectations to accept a wager or flip a coin to accept a wager and win, and since neither type of ball is ever allowed to go to extinction on the model, the predictor's behavior is never fully stable. Also, note that the argument here depends upon the agents evolving predictive and descriptive practices that are at least to some degree successful. The more successful they are, the faster the predictor will evolve to use expectation weights as expectations. If the predictor never does better than even in his predictions, he will never evolve to use his expectation weights systematically as expectations in the side wagers. It is in this sense that the evolution of expectations here is contingent on the co-evolution of successful descriptive and predictive practice.

associated with meaningful descriptions insofar as there are no meaningful descriptions to begin. Indeed, the expectation weights do not initially represent expectations at all, but rather, only gradually come to play such a role. And finally, since the expectation weights closely track the actual type and degree of success of the agents' evolving descriptive and predictive practice, any initial bias in the expectation weights washes out by the time the agents evolve a successful descriptive and predictive practice.

#### DISCUSSION

The proper morals to draw from the present model are modest. It shows that a simple dispositional model sufficiently sophisticated to coevolve successful descriptive and predictive practice might also coevolve conditional expectations that reflect the type and degree of the evolved success. Such expectations might then serve as the basis for more sophisticated predictive action.

The simple model described here does not seek to represent how actual inquirers in fact evolved the preconditions for reflective inquiry as practiced. Rather, it illustrates how basic descriptive language, predictive practice, and conditional expectations might coevolve in such a way that the problem of priors, as a problem regarding how one ought to assign initial expectations to a set of meaningful descriptive hypotheses, simply fails to apply. Here one only has meaningful descriptive hypotheses if one also has expectations over such hypotheses that track the type and degree of predictive success they have afforded.

The present model then does not address the problem of priors by providing a rule for setting rational priors nor does it provide any special justification for one setting of priors rather than another. The evolved inquirers here are just never faced with the task of assigning new expectations to a set of meaningful descriptive hypotheses. Rather, the very existence of a set of meaningful hypotheses presupposes a rich evolutionary history of success and failure in predictive action.

The expectations the agents use to inform their actions are the expectations they have in fact evolved. This is just part of what it means here for an expectation weight to represent an expectation. Since the evolved expectations are forged in the context of the agents actual descriptive and predictive practice, one might expect such use to be as successful as the practice. But there is no special justification for the reliability of their evolved expectations. The modeled expectations have tracked and will continue to track the type and degree of success in fact exhibited in the agents coevolving descriptive and predictive practice, whatever this may be. But there is no special assurance that present conditional expectations will track the future success of one's evolving practice.

The present model shows how it is possible that the process of evolving a descriptive language suitable for predictive inquiry might also provide agents with conditional expectations that are well-tuned to the type and degree of predictive success their practice has so far afforded. Such conditional expectations might be taken to represent the evolved commitments of the inquirers. As with the evolved application to side wagers, such commitments might be extended to inform increasing rich types of future action including reflective inquiry itself.

#### CONCLUSION

On the present model neither the inquirers' descriptive language and predictive practice nor their associated expectations are somehow prior. Since having meaningful descriptions for the purpose of prediction means also having conditional expectations, the model provides an account of how inquiry might evolve such that the problem of priors simply fails to apply.

The conditional expectations that evolve on the model are associated with those descriptions that were in fact used in making successful predictions. Consequently, they can be expected to reflect whatever type and degree of predictive success the agents' evolved descriptive and predictive practice has so far afforded.<sup>17</sup>

---

<sup>17</sup>I would like to thank Brian Skyrms, Jim Weatherall, and Seamus Bradley in particular for helpful discussions on this topic. I would also like to thank the two anonymous reviewers for their very helpful comments and suggestions.

## REFERENCES

- [Argiento, Pemantle, Skyrms, and Volkov (2009)] Argiento, Raffaele, Robin Pemantle, Brian Skyrms and Stas Volkov (2009) Learning to Signal: Analysis of a Micro-Level Reinforcement Model. *Stochastic Processes and their Applications* 119(2):373–390
- [Barrett (2012)] Barrett, Jeffrey A. (2012a) On the Coevolution of Theory and Language and the Nature of Successful Inquiry. Forthcoming in *Erkenntnis* DOI 10.1007/s10670-013-9466-z
- [Barrett (2009)] Barrett, Jeffrey A. (2009) Faithful Description and the Incommensurability of Evolved Languages. *Philosophical Studies* 147(1):123–137
- [Barrett (2007)] Barrett, Jeffrey A. (2007) Dynamic Partitioning and the Conventionality of Kinds. *Philosophy of Science* 74:527–546
- [Barrett (2006)] Barrett, Jeffrey A (2006) Numerical Simulations of the Lewis Signaling Game: Learning Strategies, Pooling Equilibria, and the Evolution of Grammar. UC Irvine Institute for Mathematical Behavioral Sciences Preprint MBS 06-09.
- [Barrett and Zollman (2009)] Barrett, J. A. and K. Zollman (2009) The Role of Forgetting in the Evolution and Learning of Language. *Journal of Experimental and Theoretical Artificial Intelligence* 21(4): 293-309.
- [Bereby-Meyer, Yoella and Erev (1998)] Bereby-Meyer, Yoella and Ido Erev (1998) On Learning to become a Successful Loser: A Comparison of Alternative Abstractions of Learning Processes in the Loss Domain. *Journal of Mathematical Psychology* 42:226-286
- [Lewis (1969)] Lewis, David (1969) *Convention*. Harvard University Press, Cambridge, MA
- [Herrnstein (1970)] Herrnstein, R. J. (1970) On the Law of Effect. *Journal of the Experimental Analysis of Behavior* 13: 243–266.
- [Huttegger (2007)] Huttegger, Simon (2007) Evolution and the Explanation of Meaning. *Philosophy of Science* 74:1–27
- [Roth and Erev (1995)] Roth, Al and Ido Erev (1995) Learning in Extensive Form Games: Experimental Data and Simple Dynamical Models in the Intermediate Term. *Games and Economic Behavior* 8:164-212
- [Skyrms (2010)] Skyrms, Brian (2010) *Signals Evolution, Learning, & Information*. Oxford University Press, New York
- [Skyrms (2006)] Skyrms, Brian (2006) Signals. *Philosophy of Science* 75(5):489–500