A Bayesian model of quasi-magical thinking can explain observed cooperation in the public good game

Joanna Masel[1,2]

[1]Department of Biological Sciences, Stanford, CA 94305, USA
[2]Department of Ecology & Evolutionary Biology, University of Arizona, Tucson, AZ 85721, USA
Email: masel@u.arizona.edu

Abstract

People often contribute to the public good, contradicting the assumption of perfect self-interest in traditional economic theory. Models of learning, reciprocity and altruism have been proposed to explain this cooperation, but none can explain all quantitative aspects of observed contributions in the public good game. Here a new model is described in which players recognize that “other players may be like me”, and so assume a correlation between their own contribution and the likely contributions of other players. The extent of the correlation is calculated by treating a player’s own conjectured contribution in the same way as an observed contribution of another player within a Bayesian model of learning. Although players recognize that this correlation is not causal, they nevertheless choose to maximize expected utility conditional on their own action rather than standard expected utility. Ignoring causality in this way is consistent with the documented psychological phenomenon of “quasi-magical thinking”. This mathematical model of quasi-magical thinking generates results in agreement with data that preference functions have not been able to explain.

Classification numbers: C71, H40
Keywords: conditional expected utility, rationality

Acknowledgements

I thank Aviv Bergman, Sam Bowles, Marc Feldman, Lilach Hadany, Paul Higgins, Ben Kerr, Adi Livnat, John Miller, Henrique Pereira, Leo Simon and Russell Thomson for helpful discussions and comments on the manuscript. Research was supported by NIH grant GM28016 and the Center for Computational Genetics and Biological Modeling at Stanford.

Biographical note

As of August 2, 2004, I am an assistant professor in Ecology and Evolutionary Biology at the University of Arizona. I am a mathematical biologist with expertise in evolution, and have recently become interested in human behavior.

1 Introduction

A number of games, such as the Prisoner's dilemma, have been designed to study the choice between cooperation and defection. Here we will primarily consider the public good game, in which intermediate choices between complete cooperation and complete defection are both possible and common. The prevalence of quantitative intermediate choices means that complex data on human decision-making are available. No model is yet able to explain all aspects of these data.

In the public good game, each of n players is given 20 tokens, which can be redeemed for cash. Player j chooses a discrete number of tokens $0 \leq s_j \leq 20$ to contribute to a common pool. Total contributions to the pool are then multiplied by some factor $mn>1$, and distributed equally back to all players. Whatever the other players do, a player is always better off contributing nothing. The principle of dominance states that under such circumstances "rational" players should contribute nothing. In experimental situations, however, players consistently contribute around half their available funds to the common pool (Ledyard, 1995).

A variety of experiments have shown that none of the current explanations for this cooperation is satisfactory, as we describe in Section 2. In Section 3, we introduce the concepts behind a novel model that can explain the data described in Section 2, and briefly review supporting experiments from other contexts. This model is based on maximizing conditional expected utility within a Bayesian learning model. In Section 4 we develop the model mathematically for a single round of the public good game. In Section 5 we analyze the effects of the parameters in this model, and where possible, compare these predicted effects to experimental data. In Section 6, we extend the mathematical model to the case when multiple rounds of the game are played. In Section 7, we summarize how the model explains data on cooperation that cannot be explained by preference functions, and suggest new experiments to further test the model.

2 Models of learning, reciprocity and altruism

One possible explanation for observed cooperation is that players do not immediately understand the incentives of an artificial game, and need repetition in order to learn. In agreement with the learning hypothesis, contributions fall when the game is repeated and players observe the choices made by others in previous rounds (Isaac et al., 1985; Kim and Walker, 1984; Ledyard, 1995). Nevertheless, by the final round of a repeated game, some contributions are still made. In addition, when players complete ten rounds and are only then told that they will start a second set of ten rounds, contributions return to relatively high levels not much below those seen in the first round, rather than continuing where they left off (Andreoni, 1988; Burlando and Hey, 1997; Croson, 1996; Isaac and Walker, 1988). A similar effect is seen when experienced rather than novice players are recruited to start a set of rounds (Isaac et al., 1984). Further evidence against the learning hypothesis comes from comparative experiments with 10, 40 or 60 rounds, in which contributions start from similar levels and decline to similar levels (Isaac et al., 1994). Since it is unlikely that subjects playing more rounds learn more slowly, some other explanation must be behind cooperation.

These data could be explained by strategic play involving some form of reciprocity. According to this hypothesis, it is not common knowledge that all players are playing “rationally”. In this case, free-riding may no longer be the best option, as it may encourage other players to free-ride in subsequent rounds (Kreps et al., 1982). To test this hypothesis, the game can be played repeatedly in an anonymous computer laboratory such that two players rarely or never meet more than once, and have no information about the choices a particular player has made in previous rounds. This set-up is known as a “strangers” treatment, while a game repeated between the same group of players is known as a “partners” treatment.

The strategies hypothesis predicts that no cooperation should occur in the strangers treatment, in contrast to the partners treatment.

Andreoni (1988) found that, contrary to the prediction, strangers contribute slightly more than partners do, while Brandts and Schram (2001) and Weimann (1994) found that there was no difference. Croson (1996) and Keser and van Winden (2000) performed more rigorous statistical analyses, in which different rounds from the same pool of players were not treated as independent measurements, and showed that partners contributed slightly more. In no case, however, were contributions by strangers close to zero. Contributions retained many quantitative traits, such as decline over time and the restart effect. The strategies hypothesis can explain some, but by no means all of the observed cooperation.

Another explanation is that some players are genuinely altruistic rather than selfish. Altruism (Becker, 1974) or “warm glow” effects (Andreoni, 1989) exist when a player’s true utility function is not simply equal to the player’s personal monetary reward. Instead, the utility includes a term for the monetary reward received by the group as a whole (Becker, 1974) or for the act of contributing (Andreoni, 1989). Typically, a player has a “cut-off decision rule” to contribute if the marginal benefit from contributing, given by m in our formulation, is greater than some cut-off m' (Anderson et al., 1998; Palfrey and Rosenthal, 1988; Palfrey and Prisbey, 1996). A perfectly self-interested subject has $m' = 1$. Deviations from the decision rule are interpreted as being due to subject confusion and error (Anderson et al., 1998; Andreoni, 1995; Palfrey and Prisbrey, 1997; Palfrey and Prisbey, 1996). For example, decline of contributions over time is interpreted as a decline in errors in an experiment where the only possible error is to contribute too much.

Altruism cannot explain all the quantitative patterns seen in the data, however. For example, players frequently split their tokens so as to give some to the public good and to keep some for private consumption. Brandts and Schram (2001) showed that splitting is a robust phenomenon which cannot easily be explained by the combination of simple cut-off decision rules and random noise.

Other explanations invoke more complex utility functions. For example, Fischbacher et al. (2001) showed that some players are conditionally cooperative, exhibiting a form of non-strategic reciprocity. This behavior could be explained by including a term for inequity aversion in players' utility functions. If such considerations of "fairness" favor contributions close to expectations of others' contributions, this can explain why tokens are split.

Other experiments on cooperation remain difficult to explain within a model of complex preference functions, however. In the one-shot Prisoner's dilemma, 3% of subjects cooperate when told their partner has defected, 16% cooperate when told their partner has cooperated, and 37% cooperate when they are not told their partner's choice (Shafir and Tversky, 1992). According to a general model of preference functions, cooperation of those who are not told of their partner's choice should be intermediate between 3% and 16%. The high observed cooperation of 37% and other puzzling data can be explained by the model presented below.

3 Outline of conditional expected utility model

We hypothesize that a major force behind contributions to the public good is the reasoning "What if everyone else thought like me?". Even when players understand that there is no causal link between their own behavior and that of others, they may still recognize a probabilistic link, and believe that this is sufficient reason to contribute. Similar reasoning can be applied to practical situations such as the decision to vote in an election (Quattrone and Tversky, 1984). The probability that a single vote will affect the result is incredibly low. Nevertheless the reasoning "If I don't go out and do it, who will?" has some power in motivating people, despite the lack of a causal relationship between one person voting and the number of other like-minded people who also vote.

The terms “magical thinking” and the “illusion of control” refer to erroneous beliefs that people can influence outcomes through their actions. The term “quasi-magical thinking” describes “cases in which people act as if they erroneously believe that their action influences the outcome, even though they do not really hold that belief” (Shafir and Tversky, 1992). Quasi-magical thinking can explain the experiment mentioned above in which 3% of subjects cooperate in the Prisoner’s dilemma when told their partner has defected, 16% cooperate when told their partner has cooperated, and 37% cooperate when they are not told their partner’s choice (Shafir and Tversky, 1992). In this case, a player acts as though she believes that her cooperation makes her partner more likely to cooperate, even though she knows that this cannot be true. Quasi-magical thinking is also supported by a range of other data (Quattrone and Tversky, 1984).

In conventional economic theory, rational players maximize their expected utility function *EU* as a function of their strategy choice *s*. Expected utility may also depend on a set of perceived probabilities $p(\theta)$ that the mean level of contribution in the population is equal to θ. In other words, they maximize

$$EU(s) = \int p(\theta)U(s,\theta)d\theta\,.$$

An alternative idea is to maximize conditional expected utility (CEU) (Jeffrey, 1983). In this case, perceived probabilities $p(\theta \mid s)$ vary according to the strategy chosen, whether or not there exists a causal link. Players maximize

$$CEU(s) = \int p(\theta \mid s)U(s,\theta)d\theta. \qquad (1)$$

The axiomatic foundations of CEU differ from the Savage axioms, and have been described by Bolker (1967) and Jeffrey (1983). In this paper, we do not consider whether it is a good idea to follow this principle, but instead discuss whether the principle can explain the observed behavior of experimental subjects in public good games. Players who maximize CEU exhibit quasi-magical thinking, since they behave as if all conditional probabilities had a causal basis, even though they sometimes know that none exists.

If a player assumes that other players may be like her, then she will assume a positive correlation between her own contribution s and the mean contribution θ of the other players. A player maximizing CEU will contribute to the public good if she believes this correlation is sufficiently large. In support of such a perceived correlation, players who cooperate in the Prisoner's dilemma are significantly more optimistic about the probability that other players will cooperate than players who defect, although it is not clear which is the cause and which the effect (Dawes et al., 1977). Further evidence shows that the variation in the predictions of subjects who play a round of the Prisoner's dilemma is greater than the variation in the predictions of subjects who merely observe a round, demonstrating that the act of choice can modify a player's expectations (Dawes et al., 1977).

A previous model invoked a non-causal perceived correlation in order to explain cooperation, but the extent of the perceived correlation was arbitrary (Orbell and Dawes, 1991). Here we extend this model by describing a natural way to calculate conditional probabilities in the public good game. Players enter the game with a Bayesian prior distribution corresponding to their expectations of human behavior. When behavior is observed, a player's prior distribution is updated. A weighting function makes recent data more significant than data from earlier rounds, to reflect the fact that behavior may change over time. A player's own hypothetical choice of strategy s is treated as an observed data point to calculate the conditional probability $p(\theta \mid s)$. This updating follows Bayesian principles, and we refer to it as "conjectural updating".

A player's belief that other players may be like her is often taken to be a flawed mode of reasoning and is referred to as the "false consensus effect" (Ross et al., 1977). It has been argued that this reasoning may in fact be reasonable in many circumstances, since a player is a member of the same population as all other players. According to this argument, a "truly false consensus effect" occurs only when a player treats her own contribution differently from an observation of any other player's contribution (Dawes, 1989, 1990; Dawes and Mulford, 1996). In our model of conjectural updating, a player treats her own contribution in exactly the same way as she would treat an observation of another player's contribution, as a piece of data to be analyzed like any other.

4 Mathematical Model of Single Round Game

Here we consider a single round of the public good game. Since the model is based on a learning mechanism, the interpretation of a single round game is not entirely intuitive. In Section 6, the model will be extended to describe a multiple round game by including the possibility of a trend over time and by factoring in payoffs for future rounds. This section is best regarded as a slightly simplified model of the final round of a set of repeated strangers public good games. A summary of the main parameters of the model is given in Table I. The monetary payoff for each player is

$$\pi(s_j) = 20 - s_j + m\sum_{k=1}^{n} s_k \text{ where } 1/n < m < 1. \qquad (2)$$

Each player has a Bayesian prior distribution corresponding to her estimate $p(\theta)$ of the probability that the mean contribution of other players is θ. The prior distribution has mean $\hat{\theta}$ and variance σ^2. According to (1) and (2), the expected payoff is therefore

$$CEU(s) = 20 + (m-1)s + m(n-1)(\hat{\theta} \mid s) \tag{3}$$

A simple estimate of θ is the mean of all contributions that have ever been observed. This leads to a linear relationship between s and $\hat{\theta} \mid s$, and so it cannot explain why players split their tokens. This estimate is also slow to adapt to a changing environment since it gives too much weight to events in the distant past. Players are unlikely to expect or to be exposed to a static environment, and so a better estimate would discount the past. The size of the optimal discount factor α depends on how fast the environment is changing.

Let players assume that the mean contribution θ_i at round i changes between rounds according to $\theta_{i+1} = \theta_i + W$, where W is drawn from a random distribution with mean 0 and standard deviation of ε. The mean $\bar{x}$ of set of n observed contributions at time i is assumed to be drawn from a distribution with mean θ_i and variance v/n. For conjectural updating, n=1. Equation (4) calculates the optimal weight α to place on the old estimate $\hat{\theta}$, whose uncertainty is $\sigma_i^2 + \varepsilon^2$, relative to the newly observed data point $\bar{x}$, whose uncertainty is v/n.

$$\begin{gathered} \hat{\theta} \mid \bar{x} = \alpha\hat{\theta} + (1-\alpha)\bar{x} \text{ where } \alpha = \frac{v/n}{v/n + \sigma_i^2 + \varepsilon^2} \\ \frac{1}{\sigma_{i+1}^2} = \frac{1}{\sigma_i^2 + \varepsilon^2} + \frac{n}{v} \end{gathered} \tag{4}$$

Equation (4) is a special case of the Kalman filter. A detailed derivation and discussion of (4) in the context of animal memory and learning can be found elsewhere (McNamara and Houston, 1987). McNamara and Houston (1987) developed Equation (4) to calculate the optimal use of information in a changing environment in the special case where v and ε are constant and known. In this case σ approaches a steady state, after which time α is constant and the data are weighted exponentially according to how long ago they were observed. Note that if a changing environment is not assumed i.e. ε=0, then each player's prior becomes progressively narrower through the rounds i.e. $\sigma \to 0$ and $p(\theta \mid s)$ loses its dependence on s. This has been acknowledged as a weakness of earlier related models (Orbell and Dawes, 1991). A changing environment is a reasonable assumption in the public good game, where change over time can be clearly observed by the experimentalist.

Based on the steady state solution of (4), a simple model would be for each player j to update her prior for θ according to a constant discount factor α_j. Exponential discounting is used in many models of learning (e.g., Camerer and Ho, 1999). A constant discount factor leads to a linear relationship between a player's own choice s and her estimate $\hat{\theta} \mid s$. This model has two major faults: firstly, the extent of the correlation between s and $\hat{\theta} \mid s$ remains highly arbitrary, and secondly, this model cannot explain why players split their tokens so as to give only some to the public good while also keeping some for private consumption.

Both of these faults can be addressed by a model in which players estimate v and ε from the data, and then use $\hat{v}$ and $\hat{\varepsilon}$ to estimate θ by following (4). The simplest way to estimate v and ε is for players to assume that they are constant over time. Players' priors for v and ε can be considered equivalent to the observation of w data points, where w is a positive real number. In other words, w is the relative weight or certainty that a player assigns to her estimates of v and ε, and serves an equivalent role in the model to the uncertainty factor σ^2 that she assigns to her estimate $\hat{\theta}$.

When only one data point *s* is considered during conjectural updating, a player cannot independently update both $\hat{v}$ and $\hat{\varepsilon}$. Instead she holds the ratio $\hat{v}/\hat{\varepsilon}^2$ constant, and updates $\hat{v}+\hat{\varepsilon}^2$ according to

$$\hat{v} \mid s + \hat{\varepsilon}^2 \mid s = \frac{w\left(\hat{v}+\hat{\varepsilon}^2\right)+\left(\hat{\theta}-s\right)^2}{w+1} \tag{5}$$

For each hypothetical strategy choice *s*, a player calculates $\hat{v} \mid s$ and $\hat{\varepsilon}^2 \mid s$ using (5), then uses (4) with n=1 to calculate $\hat{\theta} \mid s$ and then (3) to calculate the CEU of strategy *s*. This is done for each possible strategy, and the strategy giving the highest CEU is played.

Note that a player does not assume that other players follow the same decision rule as herself. Nor does she assume a particular set of decision rules or distribution of player types. A player has no information on the history of individual players, which could potentially be used to formulate complex expectations about the behavior of others. Instead, the player merely keeps track of simple estimates of population parameters, namely an estimate $\hat{\theta}$ of the mean level of contribution θ in the population, an estimate $\hat{v}$ of the variance v between contributions within the population, and an estimate $\hat{\varepsilon}$ of the drift ε in the mean level of contribution, and assumes that her own choices are indicative of these parameters.

Note that the initial value of $\hat{\varepsilon}$ is based on the expected difference between mean contributions in the first and "zeroth" rounds, and should be regarded as a mathematical extrapolation. Note too that the initial value of σ^2 refers to the variance of the prior for θ in the "zeroth" round. The variance of the prior for θ in the first round is given by $\sigma_0{}^2 + \hat{\varepsilon}_0{}^2$.

5 Effect of parameters in single round game

Here we examine how the parameters m, $\hat{\theta}$, σ, n, $\hat{\varepsilon}$, $\hat{v}$ and w affect the level of cooperation chosen by an individual, and compare the predictions of the model to experimental observations.

The parameter m represents both the group and the individual return on cooperation, and predicted cooperation increases substantially and monotonically with m, as shown in Fig. 1a. This is in agreement with a range of data and has been replicated many times (Ledyard, 1995).

In general, contributions are affected by the weight players place on their previous estimate $\hat{\theta}$ relative to the weight they place on their conjectured strategy i.e. by the parameter α in (4). Anything that increases α in this model should decrease cooperation. Most parameters affect contributions via their effect on α.

For example, cooperation increases monotonically with the uncertainty parameter σ, as shown in Fig. 1a. As a player is more certain in her estimate of other players' behavior, she places a lower relative weighting on the additional data point s corresponding to her own conjectured choice. The correlation she then expects between her own behavior and that of others is therefore lower, leading her to cooperate less. Note that σ^2 should not be less than $\left(2\hat{v}/n\right)\Big/\left(1+\sqrt{1+\left(4\hat{v}/n\right)/\hat{\varepsilon}^2}\right)$, which is the steady state value of σ^2 after many observations at constant $v = \hat{v}$ and $\varepsilon = \hat{\varepsilon}$ (McNamara and Houston, 1987). Nor should σ^2 be too high: for reference as an appropriate upper bound, a flat probability distribution between 0 and 20 has $\sigma^2 = 33$.

Contributions also increase with the ratio $\hat{\varepsilon}^2/\hat{v}$ between drift and noise in the population, as shown in Fig. 1b. High drift $\hat{\varepsilon}$ leads a player to discount past observations and put more weight on her own conjectured strategy, leading to more cooperation, while high noise $\hat{v}$ has the opposite effect. Full cooperation is found in the unlikely event that $\hat{\varepsilon}^2 > \hat{v}$.

The influence of some other parameters is more complicated. From (3), (4) and (5) we see that

$$\begin{gathered} CEU(s) = 20 + (mn-1)\hat{\theta} + (mn - 1 - \alpha m(n-1))(s-\hat{\theta}), \\ \text{where } \alpha = \frac{\hat{v}}{\hat{v}+\hat{\varepsilon}^2}\left(1 - \frac{\sigma^2(w+1)}{A + (s-\hat{\theta})^2}\right) \text{ and } A = w(\hat{v}+\hat{\varepsilon}^2) + \sigma^2(w+1)\text{'} \end{gathered} \qquad (6)$$

combining the direct contribution of the term $s-\hat{\theta}$ on CEU(*s*) with the indirect contribution of the term $(s-\hat{\theta})^2$ through α. This leads to some interesting predictions of the model. If a player initially estimates $\hat{\theta} = 10$, then by contributing 11 rather than 10 tokens, the increased term $(s-\hat{\theta})^2$ leads to a small increase in α, favoring the choice of 10 tokens, but this effect may be outweighed by the increase in $s-\hat{\theta}$, favoring 11 tokens. In contrast, contributing 20 rather than 19 tokens has an equivalent effect on the linear term $s-\hat{\theta}$, but a proportionally greater effect on $(s-\hat{\theta})^2$ and hence on α, leading the player to discount her conjectured choice to a greater extent, and reducing the perceived marginal reward for cooperation. The balance between the two effects may lie at some intermediate number of tokens. Splitting the tokens is in contrast with many other models, and in agreement with the data (Brandts and Schram, 2001).

Note that when a player contributes fewer than $\hat{\theta}$ tokens, she decreases $s-\hat{\theta}$ while increasing $(s-\hat{\theta})^2$ and α. The fewer tokens she contributes, the more she will discount her choice, and the smaller the perceived penalty for defection. Once a player decides to contribute fewer than $\hat{\theta}$ tokens, the best strategy is then to contribute none at all.

In Fig. 1c we see that cooperation rises with the expected mean level of contribution $\hat{\theta}$ in certain parameter ranges. This agrees with the observation from (6) that the effect of $\hat{\theta}$ on relative utility comes through the term $s - \hat{\theta}$. Keser and van Winden (2000) found that players in a strangers treatment who contributed more than the group average decreased their next contribution, while players who contributed less increased it. This in agreement with our model in which a player's contribution is related in a linear fashion to her estimate $\hat{\theta}$, and so a rise in $\hat{\theta}$ following an observation of higher cooperation will generally lead to increased cooperation in the next round.

Once $\hat{\theta}$ is higher than a certain cut-off level, however, complete free-riding is observed, as shown in Fig. 1c. This corresponds to the point at which s=0 leads to a high enough value of $(s - \hat{\theta})^2$ for a player to discount her own choice to such an extent that U(s=0) > U($s \geq \hat{\theta}$). It appears that with a high enough expectation $\hat{\theta}$ a player becomes sufficiently confident that "the others will contribute anyway, even if I don't" that she no longer perceives that the correlation between her own contribution and that of others is large enough to justify contributing. In agreement with this interpretation, the cut-off level is closely related to σ. It is also affected to a lesser extent by the values for other parameters, as shown in Fig. 1d. The existence of this cut-off can explain the observation that some players make a contribution of zero for a while, then one or two high contributions, then revert to zero.

Note from term A in (6) that the effect of large $(s - \hat{\theta})^2$ on α and hence on reducing a player's contributions is smaller when w is large, and so increasing w leads to increased cooperation, as shown in Fig. 1e. This is because the discount factor α for conjectural updating depends on the ratio $(\sigma^2 + \hat{\varepsilon}^2 \mid s)/\hat{v} \mid s$, as can be seen in (4). Since a player's hypothetical choice s changes $\hat{\varepsilon}^2 \mid s$ and $\hat{v} \mid s$ by an equal factor, according to (5), the effect of s on $\hat{v} \mid s$ will be proportionally greater than the effect on

$\sigma^2 + \hat{\varepsilon}^2 \mid s$. For large *w*, the extent to which large $\left(s-\hat{\theta}\right)^2$ increases $\hat{v} \mid s$, and hence decreases contributions, is smaller.

The hypothesis that behavior affects expectations is sometimes put as an alternative to the hypothesis that expectations affect behavior (e.g., Dawes et al., 1977). In fact, the correlation between behavior and expectations may be due to a combination of both effects. Here we have developed a model based on the assumption that behavior affects expectations, and have used this model to derive an effect of prior expectations on behavior.

The influence of group size is complicated, as there is more than one legitimate way to hold other parameters constant while increasing group size. One way is simply to change the parameter *n*: this holds the marginal per capita return constant (Isaac et al., 1984). The marginal per capita return is defined as the ratio of benefits to costs to an individual contributing an additional token to the public good based on (2), and is therefore equal to *m*. The public good benefit provided per player by the contribution of a single token is not diluted by the addition of additional players, and so the total amount of public good provided increases with the number of players. The maximum possible reward corresponding to full cooperation of all players increases with group size in this case, and the level of cooperation predicted by our model also increases, as shown in Fig 1f by the difference between the solid and dot-dashed lines. This is consistent with experimental data using a partners treatment (Isaac et al., 1994), in contrast to some other models (Andreoni, 1989; Cornes and Sandler, 1984).

An alternative is to vary *m* and *n* together so that the total amount of public good provided stays constant as *m* varies. This means that when a player contributes a dollar to the public good, the amount that each player receives back from that dollar changes, including the amount received back by the contributing player herself. This makes altruism “more expensive” for large numbers of players. In this case the contributions decrease slightly with increasing population size, as shown in Fig. 1f by the difference between the solid and dashed lines. This is also consistent with experimental data where more cooperation was seen with n=4, m=0.75 than with n=10, m=0.3 using a partners treatment (Isaac et al., 1994). In contrast, a simple model of altruism (Isaac et al., 1994) and an evolutionary model of learning (Miller and Andreoni, 1991) predict no difference.

It has not yet been determined experimentally what factors need to be held constant as n varies for the level of contribution to remain constant. One hypothesis is to hold constant the cut-off $(1-m)/(m(n-1))$ at which contributing i tokens leads to a payoff greater than that at the complete free-riding equilibrium (Isaac et al., 1994). A second hypothesis is that players have a utility function that is a linear combination of individual reward and group reward (Ledyard, 1995). In this case contributions should be related to the term $(1-m)/(mn-1)$ (Ledyard, 1995). A third hypothesis, based on fairness, suggests that contributions should be related to the term $mn-1$(Ledyard, 1995). In our system for a single round game, we can see from (3) that the cost-benefit ratio of contributing is consistent with the first hypothesis, as can be seen by the superimposed graphs in Fig. 1f.

6 Mathematical model of repeated strangers game

When the public good game is repeated under the strangers treatment, contributions decline markedly over time. Players are likely to notice this, and to factor it into their calculations. The model can be extended accordingly by letting players assume that $\theta_{i+1} = \theta_i + W$, where W has mean δ rather than zero. In practice, the decline of contributions over time is fairly linear, so the players' assumption is reasonable. Players have a prior for δ, which we assume for simplicity has the same weight or confidence w that they assign to their priors for v and ε. In the repeated game, a player takes into account the expected payoffs for all future rounds. The payoff for the next round is given by (3). Since players have not yet decided their contribution for subsequent rounds, they assume that they will achieve the average group payoff i.e. the payoff achieved when all the players contribute θ. The total estimated payoff for rounds i+1 through to k is therefore given by

$$CEU(s) = 20 + (m-1)s + m(n-1)\hat{\theta}_{i+1} + \sum_{j=i+2}^{k}\left(20 + (mn-1)\hat{\theta}_j\right) \qquad (3b)$$
$$\text{where } \hat{\theta}_j = \max\left\{0, \min\left\{20, \hat{\theta}_i \mid s + (j-i)\left(\hat{\delta}_i \mid s\right)\right\}\right\}$$

Equation (3b) implies that players are "forward looking". This means that the more rounds left to play, the larger the reward a player perceives by increasing $\hat{\theta}$. On the other hand, the cost of increasing $\hat{\theta} \mid s$ through increasing the contribution s is based only on the current round of the game, and so it is constant. This can explain why contributions decline over time during the repeated strangers game. In later rounds, there are fewer rounds remaining, and so a player has less incentive to cooperate when other factors are held constant. This emphasis on the number of rounds remaining might also help explain why the rate of decline in contributions depends on the total number of rounds in the game.

For data on the second and subsequent rounds of a repeated strangers game, a player begins by updating her estimate $\hat{\delta}$. $\hat{\delta}$ cannot be updated from information on a single round. The initial estimate $\hat{\delta}$ has weight $\max\{0, w-n\}$, and so updating in round i+1 based on j data points follows

$$\hat{\delta}_{i+1} \text{or } \hat{\delta}_i \mid s = \frac{\left((i-1)n + \max\{0, w-n\}\right)\hat{\delta}_i + j\left(\bar{x}_{i+1} - \hat{\theta}_i\right)}{(i-1)n + \max\{0, w-n\} + j}$$

Next, (5b) (for conjectural updating) or (5c) (after a full round, first discarding the results of conjectural updating) is used instead of (5) to update $\hat{v}$ and $\hat{\varepsilon}$.

$$\hat{v}_i \mid s + \hat{\varepsilon}_i^{\,2} \mid s = \frac{n\left(i+\frac{w}{n}\right)\left(\hat{v}_i + \hat{\varepsilon}_i^{\,2}\right) + \left(\hat{\theta}_i + \hat{\delta}_i \mid s - s\right)^2}{n\left(i+\frac{w}{n}\right)+1} \tag{5b}$$

$$\begin{aligned} \hat{v}_{i+1} &= \frac{(n-1)\left(i+\frac{w}{n}\right)\hat{v}_i + \sum_{j=1}^{n}\left(\bar{x}_{i+1} - x_{i+1,j}\right)^2}{(n-1)\left(i+\frac{w}{n}+1\right)} \\ \hat{\varepsilon}_{i+1}^{\,2} &= \frac{n\left(i+\frac{w}{n}\right)\hat{\varepsilon}_i^{\,2} + n\left(\hat{\theta}_i + \hat{\delta}_{i+1} - \bar{x}_{i+1}\right)^2 - \frac{1}{n-1}\sum_{j=1}^{n}\left(\bar{x}_{i+1} - x_{i+1,j}\right)^2}{n\left(i+\frac{w}{n}+1\right)} \end{aligned} \tag{5c}$$

The prior for θ_i is now updated according to

$$\hat{\theta}_{i+1} = \alpha\left(\hat{\theta}_i + \hat{\delta}_{i+1}\right) + (1-\alpha)\bar{x}_{i+1} \tag{4b}$$

Equation (4b) is an approximation only, since the variance in the prior for δ has been neglected. Note that the expected value of θ in the first round is not the initial value $\hat{\theta}_0$ but $\hat{\theta}_0 + \hat{\delta}_0$.

To determine the effect of group size, note from (3b) that the cost-benefit ratio of contributing is now proportional to

$$\frac{1-m}{m(n-1)\hat{\theta}_{i+1} + \sum_{j=i+2}^{k}(mn-1)\hat{\theta}_j}.$$

This approximates the second hypothesis on group size described in the Section 5 above when δ=0 and the number of rounds remaining is large.

Note that further theoretical work is needed on the multiple round game. For example, players following the model presented here assume that δ, v and ε are constant and independent of $\hat{\theta}_i$. This assumption is perhaps unreasonable as $\hat{\theta}_i$ approaches the boundaries of 0 and 20. The mathematical framework put forward here is imperfect and needs further development, but the principle of the system is clear.

7 Discussion

Models explaining cooperation in the one-off or strangers public good game can be divided into two broad categories. In the first, it is assumed that a player's true utility is not equal to the player's monetary payoff, but instead includes terms representing emotions such as altruism, warm glow, fairness, spite and shame (Bowles and Gintis, 2003). In the second, it is assumed that a player's beliefs and/or decision-making process deviate from "true" rationality. An example of this second class of model is to relax the assumption of common knowledge of rationality, allowing strategic reciprocity in the partners treatment. Many models in the second class assume that players are "confused" and the deviation of their choices from the maximization of their utility is modeled as a random error term (Anderson et al., 1998; Andreoni, 1995; Palfrey and Prisbrey, 1997; Palfrey and Prisbey, 1996).

A single change to the conventional model of rationality is made in this paper: quasi-magical thinking is introduced so that players treat all correlations as though they were causal, even when they know full well that causality is absent. The formal simplicity of this change means that abandoning conventional rationality does not imply abandoning all systematic analysis. The effects of the change are highly systematic and cannot be modeled by a random error term.

Players assume that their own conjectured level of cooperation is one more statistic like any other in a population from which they can acquire data, and not a special case that should be ignored. In the absence of other information, a player will place considerable weight on her own decision when estimating the behavior of the population. When she has acquired other data making her more confident in her estimate of the population, then she will place less weight on her own conjectural decision, in the same way as she would place less weight on any other piece of additional information. We have used Bayesian methods to propose a consistent method of assuming that other members of the population might be similar to oneself, and have avoided arbitrary correlations.

Philosophers have explored at length the principle of maximizing CEU with reference to a dilemma known as Newcomb's problem (Nozick, 1969). This in itself suggests that maximizing CEU is not so clearly irrational as not to cross people's minds. Genuine responses to Newcomb's problem, which invokes an omniscient being, are not easily accessible experimentally, although other experimental evidence strongly suggests that many people exhibit quasi-magical thinking and maximize CEU (Shafir and Tversky, 1992). Here we have shown that data on the public good game should show certain quantitative trends if players follow the principle of maximizing CEU. The evidence suggests that they do.

Could a sufficiently complicated preference function be sufficient to explain observed cooperation, or do we need to invoke the principle of maximizing CEU? Here we list four pieces of evidence that a complicated preference function cannot explain. Firstly, there is direct evidence that players exhibit quasi-magical thinking in the Prisoner's dilemma (Shafir and Tversky, 1992). Secondly, the act of choice is known to modify expectations (Dawes et al., 1977). Thirdly, it is not clear how a preference function can explain the restart effect or the dependence of the rate of contribution decline on the number of rounds remaining in the strangers treatment. Finally, players who are asked to estimate the action of other players before making their own choice contribute less than those who are not asked (Croson, 2000). When a player has already stated a single set of beliefs about other players' actions, she is less likely to modify those beliefs in the light of her own conjectural strategy. The act of stating beliefs in advance is likely to weaken the extent to which players follow the principle of maximizing CEU, leading to lower contributions in agreement with the data. Similarly, players are less likely to maximize CEU in the Prisoner's Dilemma when told that their partner's choice was made in the past rather than made simultaneously or in the future, a condition also likely to weaken quasi-magical thinking (Morris and Sim, 1998).

Further experimental evidence would still be desirable. For example, one player could be allowed to repeatedly play novices rather than other continuing players. This player is exposed to a constant rather than changing environment and so her expectations of the population parameters would become both accurate and certain over time. In addition, the loss of symmetry may reduce the perceived correlation between the experienced player and the novices. The model of CEU predicts that the experienced player's contribution will rapidly drop to zero in this experiment, far more rapidly than in a symmetric game. There is no obvious reason why a complicated preference function should predict this behavior.

More complicated experiments are more difficult to interpret within the mathematical framework presented here. For the repeated strangers game, more parameters are needed. In addition to allowing $\delta \neq 0$, heterogeneous initial parameters must be assumed for each of the interacting players. Note that player heterogeneity is due entirely to heterogeneous priors in this system. Other sources of heterogeneity may also be present, for example the presence of genuine free-riders, those with altruistic or "warm-glow" preferences, those with preferences for a more equal distribution of payoffs, those who play strategically, and of course an element of complete randomness. Yamagishi et al., (1999) found that 11 out of 39 players agreed that "your own group members would allocate you more if you allocated more to a member of your own group" and therefore appeared to exhibit magical thinking. Shafir and Tversky

(1992) found that 25% of subjects exhibit behavior clearly indicative of magical or quasi-magical thinking in the Prisoner's dilemma Clearly, the model presented here describes some but not all players in a given experiment.

Still more complicated experimental designs include the possibility of strategic reciprocity within a partners condition, a threshold effect, heterogeneous payoffs and the opportunity for punishment. Many of these more complicated designs, such as the opportunity for punishment, increase levels of cooperation. Nevertheless, a substantial level of cooperation persists when these effects are excluded by the experimental design. Clearly, the veracity of the model described here is best explored when as many of these effects as possible are excluded. The explanation for cooperation given here need not exclude other explanations, however, even for a single individual. Strategic reciprocity and fear of punishment are clearly significant in other experimental designs. Ultimately, other utility functions can be combined with the CEU approach described here, since the two models are not mutually exclusive.

Parameters relating to players' prior expectations could be experimentally isolated by focusing on observers of the repeated strangers game, who would receive a monetary payoff depending on how accurately they predict behavior. Data from this set-up could be used to verify or improve the updating rule proposed here and to estimate the distribution of initial values of $\hat{\theta}$, σ, $\hat{\varepsilon}$, $\hat{v}$ and w in a population. These parameter estimates could then be used to analyze data from the full repeated strangers game.

The framework described here can also provide alternative interpretations for old data. For example, allowing players to communicate with one another before they decide their contributions increases contributions, even though the discussion is not binding and many players lie about their intentions (Dawes et al., 1977). Communication might increase $\hat{\theta}$ and/or decrease $\hat{v}$, in addition to or instead of acting as a signal of an implied social contract, although this effect may be partly offset by a decrease in

σ. In support of this mechanism of increasing cooperation, players who discuss their choices in the Prisoner's dilemma estimate that a higher proportion of players will cooperate than those who do not discuss their choices (Dawes et al., 1977). This appears to be a sensible adjustment, since players' estimates of the proportion of players cooperating become more accurate after discussion (Dawes et al., 1977).

Other puzzling data come from experiments on group identity (Karp et al., 1993; Rabbie et al., 1989). Players are divided into two groups based on highly arbitrary criteria such as the alleged tendency to underestimate or overestimate dots on a screen or aesthetic preference between the artists Klee and Kandinsky. A player chooses how to allocate money between a recipient player of the same group and a recipient player of a different group. When the allocating player is potentially a recipient from a member of the same group, then she tends to favor members of the same group. When the allocating player is potentially a recipient from a member of the other group, then she tends to favor members of the other group. When the allocating player is not also a recipient, then no bias is shown. These data are difficult to explain in terms of favoritism towards one's own group, but make sense when considered in terms of the model of quasi-magical thinking presented here.

Clearly, the Bayesian learning model as presented here is not an accurate reflection of players' conscious deliberations. Even if a player did reason according to the assumptions of the model given here, she may well fail to choose precisely the optimal strategy within that framework due to the difficult computations involved. Nevertheless, the Bayesian model may offer an approximation of human decision-making, just as conventional rationality is normally only assumed to be an approximation. It is possible that players make systematic as well as random errors in Bayesian updating, such as placing too much or too little weight on their own choice (Dawes, 1989), and this could be a focus for further research.

In summary, we have presented a tractable model of quasi-magical thinking that avoids arbitrary correlations. We find that it can explain at least four pieces of data on human cooperation that cannot be explained by preference functions, and we propose further experiments to test the model.

References

Anderson, S. P., Goeree, J. K. and Holt, C. A., 1998. A theoretical analysis of altruism and decision error in public goods games. Journal of Public Economics, 70(2), 297-323.

Andreoni, James, 1988. Why Free Ride? Strategies and Learning in Public-Goods Experiments. Journal of Public Economics, 37(3), 291-304.

Andreoni, James, 1989. Giving With Impure Altruism: Applications to Charity and Ricardian Equivalence. Journal of Political Economy, 97(6), 1447-1458.

Andreoni, James, 1995. Cooperation in Public-Goods Experiments: Kindness or Confusion? American Economic Review, 85(4), 891-904.

Becker, G.S., 1974. A Theory of Social Interactions. Journal of Political Economy, 82(6), 1063-1093.

Bolker, E.D., 1967. A simultaneous axiomatization of utility and subjective probability. Philosophy of Science, 34, 333-340.

Bowles, Samuel and Gintis, Herbert, 2003, Prosocial Emotions, SFI Working Paper.

Brandts, Jordi and Schram, A., 2001. Cooperation and noise in public goods experiments: applying the contribution function approach. Journal of Public Economics, 79(2), 399-427.

Burlando, R. and Hey, J. D., 1997. Do Anglo-Saxons free-ride more? Journal of Public Economics, 64(1), 41-60.

Camerer, Colin and Ho, T. H., 1999. Experience-weighted attraction learning in normal form games. Econometrica, 67(4), 827-874.

Cornes, R. and Sandler, T., 1984. The theory of public goods: Non-Nash behaviour. Journal of Public Economics, 23, 367-379.

Croson, Rachel T. A., 1996. Partners and strangers revisited. Economics Letters, 53(1), 25-32.

Croson, Rachel T. A., 2000. Thinking like a game theorist: factors affecting the frequency of equilibrium play. Journal of Economic Behavior & Organization, 41(3), 299-314.

Dawes, Robin M., 1989. Statistical Criteria for Establishing a Truly False Consensus Effect. Journal of Experimental Social Psychology, 25, 1-17.

Dawes, Robin M., 1990, The Potential Nonfalsity of the False Consensus Effect. In: R. M. Hogarth (Ed.), Insights in decision making: A tribute to Hillel J. Einhorn. University of Chicago Press, Chicago and London, pp. 179-199.

Dawes, Robin M., McTavish, J. and Shaklee, H., 1977. Behavior, Communication, and Assumptions About Other Peoples Behavior in a Commons Dilemma Situation. Journal of Personality and Social Psychology, 35(1), 1-11.

Dawes, Robin M. and Mulford, M., 1996. The false consensus effect and overconfidence: Flaws in judgment or flaws in how we study judgment? Organizational Behavior and Human Decision Processes, 65(3), 201-211.

Fischbacher, Urs, Gächter, Simon and Fehr, Ernst, 2001. Are people conditionally cooperative? Evidence from a public goods experiment. Economics Letters, 71(3), 397-404.

Isaac, R. M., McCue, K. F. and Plott, C. R., 1985. Public-Goods Provision in an Experimental Environment. Journal of Public Economics, 26(1), 51-74.

Isaac, R. M. and Walker, J. M., 1988. Communication and Free-Riding Behavior - the Voluntary Contribution Mechanism. Economic Inquiry, 26(4), 585-608.

Isaac, R. M., Walker, J. M. and Thomas, S. H., 1984. Divergent Evidence On Free Riding - an Experimental Examination of Possible Explanations. Public Choice, 43(2), 113-149.

Isaac, R. M., Walker, J. M. and Williams, A. W., 1994. Group-Size and the Voluntary Provision of Public-Goods - Experimental-Evidence Utilizing Large Groups. Journal of Public Economics, 54(1), 1-36.

Jeffrey, R. C., 1983, The logic of decision. University of Chicago Press, Chicago.

Karp, D., Jin, N., Yamagishi, T. and Shinotsuka, H., 1993. Raising the Minimum in the Minimal Group Paradigm. The Japanese Journal of Experimental Social Psychology, 32(3), 231-240.

Keser, C. and van Winden, F., 2000. Conditional cooperation and voluntary contributions to public goods. Scandinavian Journal of Economics, 102(1), 23-39.

Kim, O. and Walker, M., 1984. The Free Rider Problem: Experimental Evidence. Public Choice, 43(1), 3-24.

Kreps, D.M., Milgrom, P., Roberts, J. and Wilson, R., 1982. Rational Cooperation in the Finitely Repeated Prisoners' Dilemma. Journal of Economic Theory, 27, 245-252.

Ledyard, John O., 1995, Public goods: a survey of experimental research. In: J.H. Kagel and A.E. Roth (Eds.), The handbook of experimental economics. Princeton University Press, Princeton, pp. 111-194.

McNamara, J. M. and Houston, A. I., 1987. Memory and the efficient use of information. Journal of Theoretical Biology, 125(4), 385-395.

Miller, John H. and Andreoni, James, 1991. Can Evolutionary Dynamics Explain Free Riding in Experiments? Economics Letters, 36(1), 9-15.

Morris, M. W. and Sim, D. L. H., 1998. Distinguishing sources of cooperation in the one-round prisoner's dilemma: Evidence for cooperative decisions based on the illusion of control. Journal of Experimental Social Psychology, 34(5), 494-512.

Nozick, R., 1969, Newcomb's problem and two principles of choice. In: N. Rescher (Ed.), Essays in Honor of Carl G. Hempel. Reidel, Dordrecht, pp. 114-146.

Orbell, J. and Dawes, Robin M., 1991. A Cognitive Miser Theory of Cooperators Advantage. American Political Science Review, 85(2), 515-528.

Palfrey, T. R. and Prisbrey, J. E., 1997. Anomalous behavior in public goods experiments: How much and why? American Economic Review, 87(5), 829-846.

Palfrey, T. R. and Rosenthal, H., 1988. Private Incentives in Social Dilemmas - the Effects of Incomplete Information and Altruism. Journal of Public Economics, 35(3), 309-332.

Palfrey, T.R. and Prisbey, J.E., 1996. Altruism, Reputation and Noise in Linear Public Goods Experiments. Journal of Public Economics, 61(3), 409-427.

Quattrone, G. A. and Tversky, A., 1984. Causal Versus Diagnostic Contingencies - On Self-Deception and On the Voters Illusion. Journal of Personality and Social Psychology, 46(2), 237-248.

Rabbie, J. M., Schot, J. C. and Visser, L., 1989. Social Identity Theory - a Conceptual and Empirical Critique from the Perspective of a Behavioral Interaction-Model. European Journal of Social Psychology, 19(3), 171-202.

Ross, L., Greene, D. and House, P., 1977. False Consensus Effect - Egocentric Bias in Social-Perception and Attribution Processes. Journal of Experimental Social Psychology, 13(3), 279-301.

Shafir, E. and Tversky, A., 1992. Thinking Through Uncertainty: Nonconsequential Reasoning and Choice. Cognitive Psychology, 24(4), 449-474.

Weimann, J., 1994. Individual Behavior in a Free Riding Experiment. Journal of Public Economics, 54(2), 185-200.

Yamagishi, Toshio, Jin, Nobuhito and Kiyonari, Toko, 1999. Bounded Generalized Reciprocity: Ingroup Boasting and Ingroup Favoritism. Advances in Group Processes, 16, 161-197.

Figure 1.

Predicted levels of cooperation in the single round public good game given different parameters for the rules of the game and for the prior estimates of the player. Parameters not shown as variables in the graphs are given by m=0.5, n=4, w=4, $\hat{\theta}$=8, σ^2=8, $\hat{v}$=4, $\hat{\varepsilon}$=1. Parameter choices for m and n are typical of experiments reported in the literature, and choice of $\hat{\theta}$ represents an accurate estimate of the true value. Behavior was explored for a wide range of values for the less accessible parameters w, σ^2, $\hat{v}$ and $\hat{\varepsilon}$. **a.** Cooperation increases monotonically with the return of cooperation m and a player's initial uncertainty in others' behavior σ^2. **b.** High estimated drift $\hat{\varepsilon}^2$ relative to estimated noise $\hat{v}$ leads a player to discount the past and therefore cooperate more. **c.** Chosen levels of cooperation increase with estimated contribution level $\hat{\theta}$ for values of $\hat{\theta}$ below a cut-off. Above this cut-off, complete free-riding occurs. **d.** This cut-off is highly dependent on σ, and also varies slightly with other parameters. **e.** When confidence w in $\hat{\varepsilon}$ and $\hat{v}$ is low, a player discounts hypothetical strategies which diverge from $\hat{\theta}$ as being due to noise, and so contributes less. **f.** Effect of group size n. When n is varied while m is held constant (solid and dot-dashed lines) cooperation increases with increasing n. When n is varied while mn is held constant (solid and dashed lines) cooperation decreases slightly with increasing n.

Table I: Definition of main parameters

θ	mean level of contribution in the population
$\hat{\theta}_i$	a player's estimate of θ at the end of round i
$\hat{\theta}_i \mid s$	a player's estimate of θ at the end of round i, revised by conjectural updating based on a hypothetical strategy choice s in round i+1
σ^2	variance of a player's prior for θ i.e. a player's uncertainty in her estimate $\hat{\theta}$
v	variance in contributions between players in the population
δ	bias in the drift of the mean contribution level between rounds
ε^2	variance in the drift of the mean contribution level between rounds
$\hat{v}, \hat{\delta}, \hat{\varepsilon}^2$	a player's estimates of v, δ and ε^2
w	a player's confidence in her estimates $\hat{v}$, $\hat{\delta}$ and $\hat{\varepsilon}$, equivalent to prior observation of w data points
n	number of players in a group
m	an individual's monetary return on contributing relative to not contributing for a single round

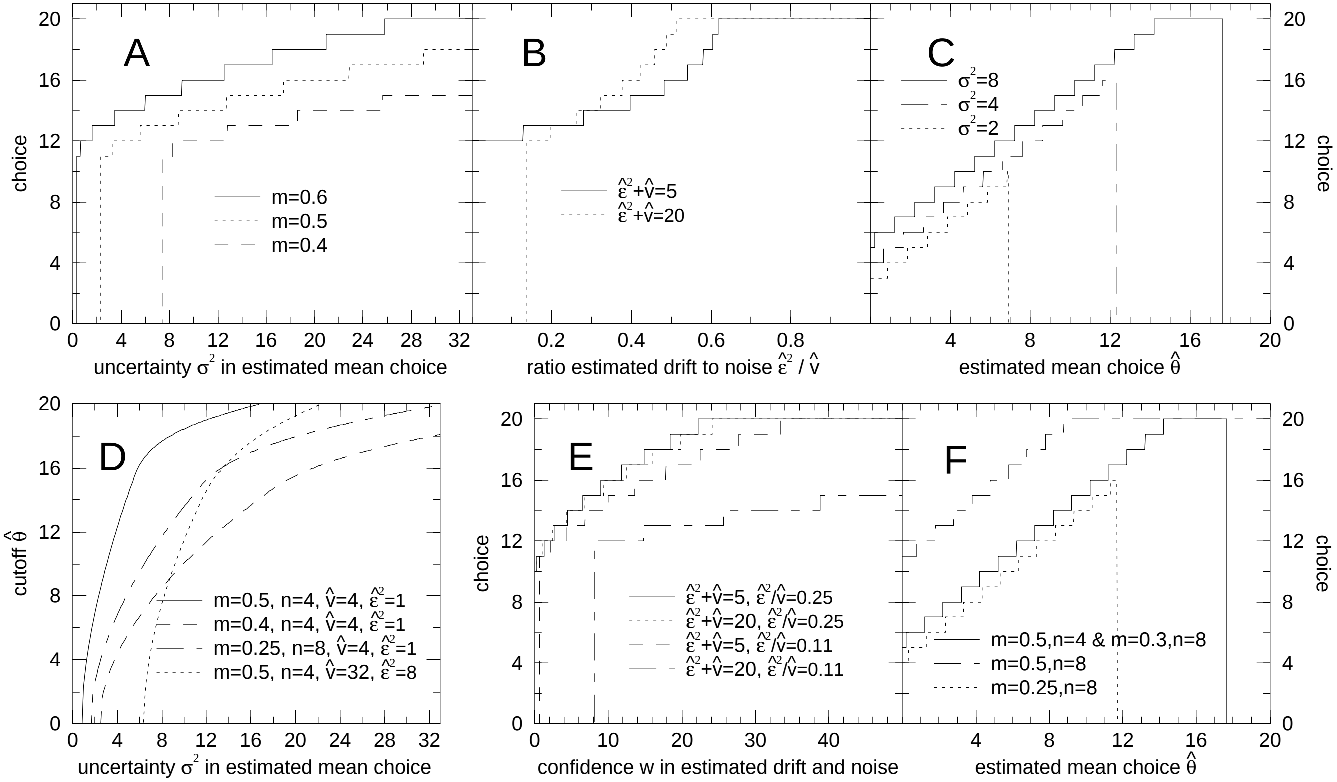
A
m=0.6
m=0.5
m=0.4
choice
0
4
8
12
16
20
uncertainty σ^2 in estimated mean choice
0
4
8
12
16
20
24
28
32
B
$\hat{\varepsilon}^2+\hat{v}=5$
$\hat{\varepsilon}^2+\hat{v}=20$
ratio estimated drift to noise $\hat{\varepsilon}^2 / \hat{v}$
0.2
0.4
0.6
0.8
C
$\sigma^2=8$
$\sigma^2=4$
$\sigma^2=2$
estimated mean choice $\hat{\theta}$
4
8
12
16
20
choice
D
cutoff $\hat{\theta}$
m=0.5, n=4, $\hat{v}=4$, $\hat{\varepsilon}^2=1$
m=0.4, n=4, $\hat{v}=4$, $\hat{\varepsilon}^2=1$
m=0.25, n=8, $\hat{v}=4$, $\hat{\varepsilon}^2=1$
m=0.5, n=4, $\hat{v}=32$, $\hat{\varepsilon}^2=8$
uncertainty σ^2 in estimated mean choice
E
choice
$\hat{\varepsilon}^2+\hat{v}=5$, $\hat{\varepsilon}^2/\hat{v}=0.25$
$\hat{\varepsilon}^2+\hat{v}=20$, $\hat{\varepsilon}^2/\hat{v}=0.25$
$\hat{\varepsilon}^2+\hat{v}=5$, $\hat{\varepsilon}^2/\hat{v}=0.11$
$\hat{\varepsilon}^2+\hat{v}=20$, $\hat{\varepsilon}^2/\hat{v}=0.11$
confidence w in estimated drift and noise
0
10
20
30
40
F
m=0.5,n=4 & m=0.3,n=8
m=0.5,n=8
m=0.25,n=8
estimated mean choice $\hat{\theta}$
choice