



AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search

<http://ageconsearch.umn.edu>

aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

Isolation, Assurance and Rules:
Can Rational Folly Supplant Foolish Rationality?

Peter J. Hammond

No 842

WARWICK ECONOMIC RESEARCH PAPERS

DEPARTMENT OF ECONOMICS

THE UNIVERSITY OF
WARWICK

Isolation, Assurance and Rules: Can Rational Folly Supplant Foolish Rationality?*

Peter J. Hammond
Department of Economics, University of Warwick,
Coventry CV4 7AL, U.K.
p.j.hammond@warwick.ac.uk

February 2008

Abstract

Consider an “isolation paradox” game with many identical players. By definition, conforming to a rule which maximizes average utility is individually a strictly dominated strategy. Suppose, however, that some players think “quasi-magically” in accordance with evidential (but not causal) decision theory. That is, they act as if others’ disposition to conform, or not, is affected by their own behavior, even though they do not actually believe there is a causal link. Standard game theory excludes this. Yet such “rational folly” can sustain “rule utilitarian” cooperative behavior. Comparisons are made with Newcomb’s problem, and with related attempts to resolve prisoner’s dilemma.

*My intellectual debts to Amartya Sen are accompanied by a friendship deep enough to presume it will survive not only well into his fourth quarter century, but also beyond my taking serious liberties with the titles, if not the ideas, of Sen (1967, 1977).

1 Introduction

1.1 Isolation versus Assurance

Sen (1967) introduced two simple games to discuss what was an appropriate rate of discount, and what policy measures might help produce an optimal rate of saving or investment in an intertemporal economy. The first game, which he called the “isolation paradox”, is a multi-person version of prisoner’s dilemma in which each player has two pure strategies, one of which (say D) strictly dominates the other (say C). Moreover, if all players chose their dominated strategy C , the outcome would be unanimously preferred to the outcome when all choose their dominant strategy D . The second game, which he called the “assurance problem”, is a significant variant of the first in which C would become a best response just in case all other players choose C ; otherwise, D is always better.

1.2 Foolish Rationality?

Of course, both games have a pure strategy Nash equilibrium where all players choose D . In the isolation paradox, moreover, this is the unique equilibrium. As in prisoner’s dilemma, it is the only strategy profile in which all players avoid the strictly dominated strategy C . Indeed, the prisoner’s dilemma is a special case of an isolation paradox involving only two players. As in prisoner’s dilemma, there is sharp contrast between individual rationality and achieving any sensible collective objective that satisfies the usual Pareto rule. Finally, note that any one player who deviates from equilibrium to play C instead of D is indeed likely to be “isolated”.

In the assurance problem, however, there is one other Nash equilibrium in pure strategies, where all play C . This is the best possible outcome that can be reached by a symmetric pure strategy profile. The “problem”, of course, is one of coordination. All players would like to be assured that all others will choose C , to which C is the unique best response. But without sufficient confidence that this is what other players will choose, the unique best response reverts to D .

Rather similar games to the isolation paradox also arise in Laffont’s (1975) discussion of “Kantian economics”, as well as in Harsanyi’s discussions of rule utilitarianism — see especially Harsanyi (1977). That is, although there may be a richer strategy space, the pure strategy rule that is optimal if all follow it is to choose the strategy C , since this is Pareto superior to the alternative pure strategy rule where each chooses D .

A utilitarian ethicist should want to transform any isolation paradox into an assurance problem, in the hope of achieving a unanimously preferred outcome. The rest of this paper discusses various attempts to do this. It will include ways of trying to achieve the optimal symmetric mixed strategy, in case this is better than any pure strategy. Also, the computations turn out to be simpler in the limiting case when the games have infinitely many players, regarded as a random sample of points from a non-atomic probability measure space.

1.3 Outline

After laying out the game model in Section 2, the next Section 3 briefly considers the kind of “warm glow” effect that Andreoni (1990) introduced to explain some individuals’ willingness to donate towards a public good. It also considers additional “rule consequences”.

The main idea of the paper, however, is to consider whether the concept of “quasi-magical” thinking due to Shafir and Tversky (1992) may represent a useful form of rational folly in this context. This is the subject of Section 4, which begins with some remarks on Newcomb’s problem, and on attempts to motivate cooperative behavior in prisoner’s dilemma. The section concludes with a very brief discussion of evidential and causal decision theory, and of the difference between the two.

Next, Section 5 revisits the isolation game, and considers what is needed to reach the optimal proportion of cooperation that was considered in Section 2.2. The final Section 6 offers a few concluding remarks and suggestions for future research.

2 Two Games with Many Players

2.1 Players, Strategies, and Payoffs

The isolation paradox and assurance problem are two games with many features in common. In the limiting versions of both games presented here, there is a large set of many players represented in the usual way by a non-atomic probability space $(N, \mathcal{N}, \lambda)$ of ex ante identical agents. This space may as well be the Lebesgue unit interval $[0, 1]$ of the real line \mathbb{R} . We will consider only symmetric strategies and correlation devices, so there is no need to bother with agents’ labels.

In both games, each agent has two alternative actions denoted by $a \in A := \{C, D\}$, where C signifies cooperation and D signifies defection. Al-

ternatively, C can be regarded as conforming to a rule, and D as deviating from it.

Let $\gamma \in [0, 1]$ denote the proportion of agents who cooperate by following the rule in question. Each agent has preferences over lotteries with outcomes in $A \times [0, 1]$ that can be represented by the expected value of a von Neumann–Morgenstern utility function $u(a, \gamma)$. We assume that $u(D, \gamma) > u(C, \gamma)$ for each fixed $\gamma \in [0, 1)$, so strategy D is better than C unless $\gamma = 1$.

For each fixed $a \in A$, assume that the mapping $\gamma \mapsto u(a, \gamma)$ is twice continuously differentiable on $[0, 1)$, with partial derivatives w.r.t. γ that satisfy $u'(a, \gamma) > 0$ and $u''(a, \gamma) < 0$. Assume too that $u'(D, \gamma) \geq u'(C, \gamma)$, signifying that the (positive) private benefit $u(D, \gamma) - u(C, \gamma)$ of deviating does not decrease as γ increases within the semi-closed interval $[0, 1)$. At $\gamma = 1$, however, $u(C, \gamma)$ may not even be continuous. Using $\lim_{\gamma \uparrow 1}$ to denote a limit as γ tends to 1 from below, we assume that $u(C, 1) \geq \lim_{\gamma \uparrow 1} u(C, \gamma)$. If there is a discontinuity at $\gamma = 1$, therefore, it is due to an upward jump.

2.2 An Optimal Symmetric Strategy Rule

Since all agents have identical preferences, an optimal symmetric mixed strategy rule is to have each agent play C with probability γ^* , chosen to maximize

$$U(\gamma) := \gamma u(C, \gamma) + (1 - \gamma)u(D, \gamma)$$

w.r.t. γ . Because

$$U'(\gamma) = u(C, \gamma) - u(D, \gamma) + \gamma u'(C, \gamma) + (1 - \gamma)u'(D, \gamma),$$

it is evident that

$$U''(\gamma) = 2[u'(C, \gamma) - u'(D, \gamma)] + \gamma u''(C, \gamma) + (1 - \gamma)u''(D, \gamma) < 0$$

for all $\gamma \in [0, 1)$, given our earlier assumptions.

In principle, there can be a corner maximum at $\gamma^* = 0$. This holds if

$$0 \geq U'(0) = u(C, 0) - u(D, 0) + u'(D, 0)$$

or $u'(D, 0) \leq u(D, 0) - u(C, 0)$. In the latter case, however, the isolation paradox completely disappears because it is best for everybody to deviate. So we assume that $u'(D, 0) > u(D, 0) - u(C, 0)$ from now on.

This leaves two other possible kinds of optimum. First, there is a corner maximum at $\gamma^* = 1$ provided that

$$0 \leq \lim_{\gamma \uparrow 1} U'(\gamma) = \lim_{\gamma \uparrow 1} [u(C, \gamma) - u(D, \gamma) + u'(C, \gamma)]$$

or

$$\lim_{\gamma \uparrow 1} u'(C, 1) \geq \lim_{\gamma \uparrow 1} [u(D, \gamma) - u(C, \gamma)] > 0.$$

This is the case when the symmetric optimum is for everybody to follow the pure strategy C . It occurs when, for each $\gamma < 1$, the marginal social gain per head of increasing γ exceeds the marginal private loss per head of having more people cooperate.

An alternative kind of optimum is to have only some people obeying the rule, with others not bothering to do so. For example, it is rarely optimal to have every person vote; only enough so that the (diminishing) marginal benefit of having more vote no longer exceeds the marginal cost, which could be considerable. In the general isolation paradox we are considering, there would then be an interior maximum at the unique $\gamma^* \in (0, 1)$ where $U'(\gamma^*) = 0$. Such an interior optimum γ^* exists if and only if both $u'(D, 0) > u(D, 0) - u(C, 0)$ at $\gamma = 0$, and $u'(C, \gamma) < u(D, \gamma) - u(C, \gamma)$ for all γ close enough to 1. Then the symmetric optimum is for everybody to follow the mixed strategy of choosing C with probability γ^* .

2.3 Distinguishing between Isolation and Assurance

The only difference between the isolation paradox and the assurance problem comes in the relative values of $u(C, 1)$ and $u(D, 1)$. The isolation paradox is a generalized prisoners' dilemma where each player has D as a dominant strategy. In particular, $u(D, \gamma) > u(C, \gamma)$ even when everybody else chooses C and so $\gamma = 1$.

By contrast, in the assurance problem one has $u(D, 1) < u(C, 1)$. Thus, provided that almost every other player chooses C (but only in this case), each player prefers C to D . At $\gamma = 1$, therefore, at least one of the functions $\gamma \mapsto u(C, \gamma)$ and $\gamma \mapsto u(D, \gamma)$ must be discontinuous. For simplicity, we assume that $u(D, \gamma)$ is continuous at $\gamma = 1$, but that $\lim_{\gamma \uparrow 1} u(C, \gamma) < u(C, 1)$.

3 Avoiding Isolation: Modified Preferences

3.1 Warm Glow Effects

To transform an isolation paradox into an assurance problem, one can try to modify the players' common preference ordering. One way of doing this is to add some sort of "warm glow" effect, to use the term that Andreoni (1990) introduced to explain some individuals' willingness to donate towards a public good.

For example, when C is chosen, each player's payoff function could be modified from $u(C, \gamma)$ to $u(C, (1 - \epsilon)\gamma + \epsilon)$, as if the player were placing extra weight on his own cooperation compared with all the other players'. Then, if $\epsilon \in (0, 1)$ is sufficiently large and $\gamma \in [0, 1)$ is sufficiently small, one will have

$$u(C, (1 - \epsilon)\gamma + \epsilon) > u(D, \gamma).$$

So if each player experiences this warm glow effect strongly enough, then D will no longer dominate C . But the point of this paper is not to delve into what might work to produce such a warm glow.

3.2 Rule Consequences

A second way of transforming an isolation paradox into an assurance problem is to include a "rule consequence" as an extra argument in each player's utility function. If more players adhere to the rule by choosing C , this somehow reinforces the rule and makes conformity more likely. Also, if there is a succession of opportunities to conform or deviate in a repeated game, players may observe how many have conformed to the rule in past periods, and assign the rule some valuation or reputation accordingly.

Indeed, an obvious way of arriving at an assurance problem is to increase the payoff $u(C, 1)$ from cooperating in the case when the whole population plays C so that it exceeds $u(D, 1)$. Yet in a large population, where an obvious measure of conformity each period is the observed proportion $\gamma \in [0, 1]$ of players who chose C , each individual has a negligible effect on γ , so D may well remain a dominant strategy unless $u(C, \gamma)$ increases discontinuously as $\gamma \uparrow 1$. That is precisely why such a discontinuity at $\gamma = 1$ was allowed.

3.3 Limitations

In any case, all that any such modification achieves is a different payoff function $u(a, \gamma)$ defined on $A \times [0, 1]$. There are two difficulties with this kind of approach. First, though conceivably it might work when $\gamma^* = 1$, it would take a rather sophisticated form of rule consequentialism to achieve an interior optimum $\gamma^* \in (0, 1)$. Indeed, a rigid rule utilitarian may still choose C even when $\gamma^* < 1$. Yet if there are many other players who all choose D , independent of what this rigid player does, the objective of maximizing the average utility of the whole population may not be achieved; the rigid player's gain from bending to play D may well outweigh any (small) negative effect on the other players of having one more person choose D .

Second, consider a Harsanyi (1986) world of “very half-hearted altruists”, many of whom may choose C with a probability below γ^* . Then even sophisticated altruists have to estimate the proportions π and $1 - \pi$ of the population who will and will not follow a modified rule. They also need to estimate what proportion $q < \gamma^*$ will play C among the fraction $1 - \pi$ who will not follow a modified rule. Knowing these statistics would enable them to choose their optimal mixed strategy, which consists of playing C with a compensating higher probability $p > \gamma^*$ chosen to satisfy $\pi p + (1 - \pi)q = \gamma^*$, if possible. In fact, the optimum is when

$$p = \min\left\{1, q + \frac{1}{\pi}(\gamma^* - q)\right\}.$$

4 Avoiding Isolation: Two Forms of Folly

4.1 Newcomb’s Problem

Another way to try converting an isolation paradox into an assurance problem is to alter players’ probabilistic beliefs. This possibility is prominently illustrated by Newcomb’s problem, to which philosophers in particular have devoted much attention. Indeed, the first published account is due to Nozick (1969) — see also Nozick (1993).

The problem postulates an infallible predictor who brings out two boxes. One is transparent, and can be seen to contain \$1,000. The second box is opaque; its content when opened will depend on what is predicted about the subject’s choice between the following two alternatives:

either take both boxes, in which case the opaque box will be empty;

or refuse the transparent box (with its \$1,000), in which case the opaque box will contain \$1 million.

Note that the predictor has to fill the opaque box, or not, before the subject’s choice. So obviously taking both boxes is a dominant strategy. Yet when the opaque box is found empty, might the subject not regret having tried passing up the \$1,000 in the transparent box?

Indeed, consider an alternative version of the paradox. There is no penalty for taking the \$1,000 from the transparent box. Before opening the opaque box, however, the subject must decide whether or not to surrender the whole \$1,000 as a bet on what the opaque box contains. The subject will get the contents of the opaque box whether the bet is made or not. If the bet is made, however, the predictor will have foreseen this and

put \$1 million in the opaque box; but if no bet is made, the opaque box will be found empty. Now placing such a bet for \$1 million at odds of 1000 to 1 against seems quite tempting, except to a subject who has an urgent immediate need for the \$1,000. Yet the alternative version produces exactly the same monetary outcomes as the original version — either \$0 or \$1,000 if the opaque box is empty, and either \$1 million or \$1,001,000 if it is full.

4.2 Twin Prisoner’s Dilemma

Philosophers such as Gibbard and Harper (1978) and Lewis (1979) were quick to realize that the prisoner’s dilemma which is so familiar to game theorists is closely related to Newcomb’s problem. Indeed, suppose two prisoners each have a dominant strategy D of confessing or a cooperative strategy C of staying silent. Suppose however that, before being arrested, both were told that the other could perfectly predict their choice, and would choose C if and only if they predict that the other player will play C . Then each faces the Newcomb problem, in effect, with tension between the dominant strategy D and, if the prediction claim is believed, the realization that the choice is between (C, C) and (D, D) .

Twin’s dilemma is a particularly stark version of this, where the two players are identical twins, who both expect the other twin to choose exactly the same as they do. Indeed, Howard (1988) considers an interesting variant of prisoner’s dilemma where the players are two computer programs that are allowed to read each other. This allows each program to check whether the other program is an exact copy of itself. Then, of course, it would be rational to cooperate if the two programs really are the same.

4.3 Magical and Quasi-Magical Thinking

A brief discussion of quasi-magical thinking by an economist can be found in Shiller (1999). He first explains the standard psychological phenomenon of “magical” thinking by referring to the work of the behavioral psychologist Skinner (1948), who discovered that pigeons that were first starved and then fed a small amount at regular intervals of 15 seconds were prone to develop strange routines, such as turning completely around several times while waiting for food. Skinner explained this as a form of superstition. Having noticed that they were fed every time soon after they followed the routine, the pigeons confused correlation with causation and acted as if they believed that their routine magically caused the food to arrive.

As for “quasi-magical” thinking, and its contrast with magical thinking, Shafir and Tversky (1992, p. 463) wrote as follows:

Magical thinking refers to the erroneous belief that one can influence an outcome (e.g., the roll of a die) by some symbolic or other indirect act (e.g., imagining a particular number) even though the act has no causal link to the outcome. We introduce the term quasi-magical thinking to describe 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.

Recently the biologist Masel (2007) has revived this idea. She did so in order to explain observed cooperation in public good game experiments, where offering nothing is a dominant strategy.

4.4 Evidential versus Causal Decision Theory

Shafir and Tversky (1992, p. 463) also quote Gibbard and Harper (1978):

a person . . . wants to bring about an indication of a desired state of the world, even if it is known that the act that brings about the indication in no way brings about the desired state itself.

That is, in quasi-magical thinking choosing the act is seen as *evidence* that a desired state of the world will occur. Whereas in magical thinking choosing the act is seen as the *cause* of that desired state. A corresponding distinction between “evidential” and “causal” decision theory has received considerable attention in the philosophical literature (for example, Hurley 1991, 2005; Joyce & Gibbard 1998; Joyce 1999, 2000, 2007). The point of view taken here, however, is due to Shin (1991), who argues that one can represent Jeffrey’s (1983) key notion of ratifiability using counterfactual beliefs, in a way that is more familiar to game theorists — especially following Aumann (1987).

5 Rational Folly in the Isolation Game

5.1 Universal Cooperation

We now revert to the isolation game considered in Section 2, where $u(C, \gamma)$ is twice continuously differentiable throughout the closed interval $[0, 1]$. Suppose that, although strategy D dominates C , nevertheless every agent has quasi-magical beliefs implying that the proportion of other agents who will

choose C is γ_C if that agent chooses C himself, but only $\gamma_D < \gamma_C$ if he chooses D . Suppose too that $u(C, 1) > u(D, 0)$, so that it is better to have everybody conform than everybody defect. Then, provided that γ_C is close enough to 1 and γ_D is close enough to 0, continuity evidently implies that $u(C, \gamma_C) > u(D, \gamma_D)$. Of course, this is enough to ensure that every agent prefers C to D , so universal cooperation can be sustained.

5.2 Limited Cooperation

Achieving universal cooperation clearly requires that all agents find a reason not to play their dominant strategy D . In our setting, this requires all agents to have quasi-magical beliefs. Suppose, however, that only a proportion $\rho \in (0, 1)$ of randomly selected agents have such beliefs. Also, as customary in Bayesian game theory, suppose this proportion ρ is commonly knowledge. Then the remaining fraction $1 - \rho$ of the population will choose their dominant strategy D . Accordingly, even an agent with quasi-magical beliefs will presumably have $\gamma_C \leq \rho$.

Of course, in this case cooperation by at most a proportion ρ of agents can be sustained. Even for this to be possible, however, requires that $u(C, \gamma_C) > u(D, \gamma_D)$, which is impossible unless $u(C, \rho) > u(D, 0)$. Because dominance implies that $u(C, 0) < u(D, 0)$ even when $u(C, 1) > u(D, 0)$, this is possible only if ρ is sufficiently close to 1.

5.3 Optimal Cooperation

In section 2.2 it was shown that the optimal proportion γ^* of cooperators might well be in the interior of the interval $[0, 1]$. Only in the very special case when γ^* coincides with the proportion ρ of agents with quasi-magical beliefs will this optimum be achievable with the kind of pure strategies considered above.

An obvious next step is to consider mixed strategies. But these raise conceptual difficulties when some agents have quasi-magical beliefs. This is because it remains unclear whether the expected proportion γ of cooperators depends upon the pure strategy that results ex post from randomization, or upon the mixed strategy that was selected ex ante. To avoid this difficulty, we postulate that each agent $i \in N$ receives a personal *disposition signal* $s_i \in A$, suggesting whether they should be disposed to cooperate or not.

In case this signal suggests playing D , we assume that each agent simply chooses what is anyway a dominant strategy. Also, even when the signal suggests playing C , each of the proportion $1 - \rho$ of agents who lack quasi-

magical beliefs will still play D . The proportion ρ of agents with quasi-magical beliefs, however, each expect the proportion of cooperators to be $\gamma_C \leq \rho$ if they play C , but only γ_D if they play D . So the agents with quasi-magical beliefs who receive the disposition signal C will indeed choose C provided that $u(C, \gamma_C) > u(D, \gamma_D)$. As before, this is impossible unless $u(C, \rho) > u(D, 0)$, which requires ρ to be sufficiently close to 1.

To achieve the optimal proportion γ^* of cooperators, the probability σ that any agent receives the disposition signal C must be selected to satisfy $\rho\sigma = \gamma^*$, where $\sigma \in [0, 1]$. So for γ^* to be achievable, we need $\rho \geq \gamma^*$. In other words, there must be enough quasi-magical thinkers. Or enough agents who, for whatever reason, are expected to choose the dominated strategy C instead of the dominant strategy D .

6 Concluding Remarks

Sen (1967) distinguished between the isolation paradox and assurance games. In the former, each player has a dominant strategy of departing from a rule which, if observed by enough people, would benefit all. In the latter, cooperating is better provided everybody else is cooperating; otherwise, defecting is always better.

This paper considered the quasi-magical beliefs that Shafir and Tversky (1992) postulated to explain why experimental subjects might be willing to choose a dominated strategy. It was shown that such “rational folly” may indeed be able to supplant the kind of “foolish rationality” that excludes dominated strategies in the isolated paradox.

Nevertheless, quasi-magical beliefs involve a form of self-deception that should leave game theorists and economists feeling considerable discomfort. To make such beliefs appear more palatable, it is perhaps worth investigating the evolutionary stability of this kind of behavior. Especially in the light of work such as Conley et al. (2006) on the role of “memetics”.

Meanwhile, however, it is perhaps enough to remind ourselves that many of the issues raised in Sen (1967) and in his later work remain at best incompletely resolved. They still represent challenges that may only be met if philosophers, decision and game theorists, psychologists, even economists, combine all their talents.

References

- Andreoni, J., 1990. Impure altruism and donations to public goods: a theory of warm-glow giving. *Economic Journal*, 100, p.464–477.

- Aumann, R.J., 1987. Correlated equilibrium as an expression of Bayesian rationality. *Econometrica*, 55, p.1–18.
- Conley, J.P., Toossi, A. & Wooders, M., 2006. Memetics and voting: how nature may make us public spirited. *International Journal of Game Theory*, 35, p.71–90.
- Gibbard, A. & Harper, W.L., 1978. Counterfactuals and two kinds of expected utility. In C.A. Hooker, J.J. Leach & E.F. McClennen, eds. *Foundations and applications of decision theory, vol. 1* Reidel: Dordrecht. p.125–162.
- Harsanyi, J.C., 1977. Rule utilitarianism and decision theory. *Erkenntnis*, 11, p.25–53.
- Harsanyi, J.C., 1986. Utilitarian morality in a world of very half-hearted altruists. In W.P. Heller, R.M. Starr & D. Starrett, eds. *Essays in honor of Kenneth J. Arrow, vol. I: Social choice and public decision making*. New York: Cambridge University Press. Ch. 3.
- Howard J.V., 1988. Cooperation in the prisoner’s dilemma. *Theory and Decision*, 24, p.203–21.
- Hurley, S.L., 1991. Newcomb’s problem, prisoner’s dilemma, and collective action. *Synthese*, 86, p.173–196.
- Hurley, S.L., 2005. Social heuristics that make us smarter. *Philosophical Psychology*, 18, p.585–612.
- Jeffrey, R.C., 1983. *The logic of decision*. 2nd. ed. Chicago: University of Chicago Press.
- Joyce, J.M., 1999. *The foundations of causal decision theory*. Cambridge & New York: Cambridge University Press.
- Joyce, J.M., 2000. Why we still need the logic of decision. *Philosophy of Science*, 67, p.S1–S13.
- Joyce, J.M., 2007. Are Newcomb problems really decisions? *Synthese*, 156, p.537–562.
- Joyce, J.M. and Gibbard, A., 1998. Causal decision theory. In S. Barberà, P.J. Hammond & C. Seidl, eds. *Handbook of utility theory, vol. 1: Principles*. Dordrecht: Kluwer Academic. Ch. 13.

- Laffont J.-J., 1975. Macroeconomic constraints, economic efficiency and ethics: an introduction to Kantian economics. *Economica*, 42, p.430–437.
- Lewis, D., 1979. Prisoner’s dilemma is a Newcomb problem. *Philosophy and Public Affairs*, 8, p.235–240.
- Masel, J., 2007. A Bayesian model of quasi-magical thinking can explain observed cooperation in the public good game. *Journal of Economic Behavior and Organization*, 64, p.216–231.
- Nozick, R., 1969. Newcomb’s problem and two principles of choice. In N. Rescher, ed. *Essays in honor of Carl G. Hempel*. Dordrecht: D. Reidel, p.107–133.
- Nozick, R., 1993. *The nature of rationality*. Princeton: Princeton University Press.
- Sen, A.K., 1967. Isolation, assurance and the social rate of discount. *Quarterly Journal of Economics*, 81, p.112–124.
- Sen, A.K., 1977. Rational fools: a critique of the behavioral foundations of economic theory *Philosophy and Public Affairs*, 6, p.317–344.
- Shafir, E. and Tversky, A., 1992. Thinking through uncertainty: nonconsequential reasoning and choice. *Cognitive Psychology*, 24, p.449–474.
- Shiller, R.F., 1999. Human behavior and the efficiency of the financial system. In J.B. Taylor & M. Woodford, eds. *Handbook of macroeconomics: volume 1, part C*. Amsterdam: North-Holland. Ch. 20, p.1305-1340.
- Shin, H.S., 1991. A reconstruction of Jeffrey’s notion of ratifiability in terms of counterfactual beliefs. *Theory and Decision*, 31, p.21–47.
- Skinner, B.F., 1948. Superstition in the pigeon. *Journal of Experimental Psychology*, 38, p.168–172.