LOTTERY ARGUMENTS AGAINST RELIABILISM*

1.) INTRODUCTION

Over the last approximately forty-three years numerous philosophers have directed versions of the lottery paradox argument against reliabilist theories of knowledge.[1] Philosophers like Armstrong and Dretske, aim versions of the argument at imperfect reliabilism (versions of reliabilism according to which indicators or processes with less than perfect reliability can produce knowledge). Other philosophers direct lottery arguments at all forms of reliabilism. This paper surveys these arguments, and concludes that lottery paradox arguments fail to refute reliabilism--even imperfect reliabilism. In the course of the defense of reliabilism, one also, I believe, comes to realize that epistemologists have been remarkably naive in their thinking about knowledge. Thus, as an added benefit one can learn from these oversights, though I do not pretend to make substantive theoretical contributions regarding these issues.

Proponents and opponents of imperfect reliabilism assume that lottery cases constitute a challenge to reliabilism independent of Gettier and Irrationality examples. As with Gettier and Irrationality examples, those adapting the lottery paradox wish to prove that reliability is insufficient for knowledge by showing that cases not counting as knowledge meet reliabilist conditions for knowledge. Unlike Gettier cases, one does not form the troublesome lottery ticket belief on the basis of false or misleading evidence. Unlike Irrationality cases, one does not employ a seemingly irrational belief forming process to generate the lottery ticket belief. But most importantly, unlike both Gettier and Irrationality cases, intuitions regarding the epistemic status of the resulting lottery ticket belief do not appear to provide what philosophers presume to be clear-cut intuitive evidence of a failure to know.[1] Indeed, all versions of the lottery paradox argument offered against reliabilism have a common set of premises, which includes what I call "the no-knowledge premise." Since the no-knowledge premise is contentious, it requires a supporting argument. I consider the supporting arguments offered for the no-knowledge premise, arguing that the imperfect reliabilist need accept none of the supporting arguments as sound. I go on to suggest that both common practice and considerations concerning the nature of knowledge in cognitive task performance make the no-knowledge premise seem dubious at best. Finally, I note that traditional internalist theories of justification are no immune to the paradox than reliabilist theories.

2.) IMPERFECT RELIABILISM AND THE BASIC LOTTERY PARADOX ARGUMENT

Versions of reliabilism number almost as many as its adherents. I do not attempt to represent every version of reliabilism. I do not even present my own somewhat complicated theory. Rather, I stick to a basic, naive reliabilism that represents almost everyone almost well enough. The simple version of reliabilism I have in mind defines knowledge as true belief generated by a reliable but imperfect process. For the purposes of this paper, and because lottery paradox cases are independent of other challenges to reliabilist theories, I assume that the objections raised by Gettier, Irrationality, and Evil Demon and any others may be addressed within this general framework. The thesis of this paper is that one does not need complications to address the lottery paradox. I avoid mention of justification when possible for two reasons. First, I have no stake in defending a reliabilist theory of justification. Second, talk of justification tends to lead to talk of internalism, which goes beyond the scope of this paper.

All of the various versions of the basic lottery paradox argument directed against reliabilism share the same set of assumptions. The basic lottery paradox argument works as a reductio:

(BL1) Suppose one has a fair lottery of some very large number of tickets (say, one in one hundred ten million [lotto odds]).

(BL2) The odds of any ticket winning will clearly be vanishingly small.

(BL3) Reliabilism implies that every time one truly believes based solely upon the probabilities that some ticket will lose, one knows that the ticket will lose.

(BL4) But, no one knows, based solely on the probabilities, that any particular ticket will win or that it will lose.

(BL5) Reductio Ad Absurdum: Reliabilism is refuted thus.

I must stop here to make two points, as both points emerge repeatedly throughout the paper. First, to adapt the lottery paradox to reliabilism one must assume that high probability is sufficient for reliability. This last assumption does not necessarily constitute a problem. However, one has to carefully select one's notions of probability, process, and the relevant population. If one takes subjective probability, then frequencies will not necessarily determine reliability. If one adopts a more objective notion of probability, then the probability associated with the process described as “the process generating the belief that the winning ticket won” equals 1, and that “the losing tickets won”, falls to 0. In similar fashion, one can alter the probability associated with a process by one's choice of the relevant population. For instance, take the relevant population to consist of a single instance and the probability associated with the process will be either 1 or 0.[2]

Second, Dretske (1981) has articulated a version of reliabilism which I term "perfect reliabilism." Since perfect reliabilism requires a truth-ratio of 1, perfect reliabilists can dismiss the lottery argument out of hand. Inferring the belief that a particular ticket will lose from the probability of a winning ticket within the population is an imperfect process, as one ticket, one would hope, has won. Since perfect reliabilists require the belief generating process to have perfect reliability, they can reject premise (BL3).

Perfect reliabilism has its attractions, but it also has its drawback. One must become "pragmatic" about what processes/indicators one counts as having perfect reliability or price oneself out of the knowledge market.[5] That is, after defining knowledge in terms of perfect reliability, one must go back and explain how it is that the ordinary ostensively imperfect processes/indicators one employs in everyday cognition actually count as perfectly reliable processes/indicators. Not even the brilliance and insight of Dretske can transform the notion of "pragmatic certainty" to something more coherent than notion of "almost infinite." Thus, I find perfect reliabilism unattractive. Still, I acknowledge that perfect reliabilists need go no farther for a solution to the lottery paradox. Imperfect creatures must press on.

For the imperfect reliabilist, the crucial premise in the basic lottery paradox argument--and the point where various versions of the argument diverge--is premise (BL4), the no-knowledge premise. Of course, no imperfect reliabilist will readily assent to the no-knowledge premise. However, the reductio does not go through without the no-knowledge premise. In order, therefore, to avoid begging the question, by unfairly assuming a premise that is inconsistent with imperfect reliabilism, critics of imperfect reliabilism must forward an argument in support of the no-knowledge premise.

One can separate the various supporting arguments for the no-knowledge premise according to the strategy employed by their creators. Some philosophers assume principles about the nature of knowledge, arguing that the lottery ticket belief(s) violate such principles. Some philosophers assume principles about the interrelationship of knowledge and rationality, arguing that the lottery ticket belief(s) violate those principles. I start by reviewing this last batch of supporting arguments.

3.) KNOWLEDGE, RATIONALITY, AND THE NO KNOWLEDGE PREMISE

The following supporting argument for the no-knowledge premise was once suggested in a response to an earlier version of this paper. This argument, I think, best represents the oral objections I have received in presenting this paper at conferences and in discussions with colleagues. Thus, I'll call this argument the ***dichotomy argument***, and I'll discuss it a length. I paraphrase the ***dichotomy argument***, as to the best of my knowledge it is not published:

(G1) If one knows that a lottery ticket will not win, one would act irrationally in buying the ticket.

(G2) One acts rationally in buying a ticket if the payoff is greater than the odds against winning.

(G3) The payoff is greater than the odds against winning after a certain number of tickets are sold (assuming pay-offs increase with the number to tickets sold).

(G4) Given (G2) and (G3), one sometimes acts rationally in buying a ticket.

(G5) Hence, by (G1) one does not know that a lottery ticket will not win.

All of the supporting arguments relating knowledge and rationality turn on two key premises. I call one premise ***the rational action premise***, as it asserts the rationality of some course of action. I call the other premise ***the irrational action premise***, as it asserts the irrationality of some course of action. According to the argument's irrational action premise (G1), one acts irrationally whenever one acts out of accordance with what one knows to be the case. According to the argument's rational action premise (G2), one acts rationally whenever the potential benefits of one's action outweigh the risks. Unfortunately for the argument, the two premises are rooted in two different conceptions of rationality, conceptions of rationality having the potential for conflict no matter what one's theory of knowledge.[4] As a result, anyone accepting the truth of these premises has a tool for refuting every theory of knowledge to date!

The irrational action premise is based upon an strong epistemic conception of rationality. Its adherents assert that rational action must be consistent with knowledge. The rational action premise is a utility-based conception of rational action. It instructs agents to pursue the course of action that maximizes their expected benefits at the least risk. The argument's two principles come into conflict whenever one rigs the risks and benefits so that by acting out of accordance with what they know, agents can derive great benefit while incurring little risk.

The strategy goes as follows: Choose a target theory that you wish to refute. Pick a case of knowledge sanctioned by the target theory. Hold a gun to the knower's head (pacifists can offer large sums of money). Demand that the knower declare the falsehood of the relevant belief. If one is careful not to select a belief too dear to the knower's heart, the benefit associated with the action of decrying the belief (continued existence, great material wealth) is profoundly greater than the risk of refusal (inconsistency, looking silly).

In such extreme cases as the one above, human nature as well as the utility-based rational action premise instruct the knower to act in ways inconsistent with what they know to be the case (decry the belief). The knower's action is rational according to the utility-based rational action premise. In decrying the belief, therefore, the knower refutes the target theory of knowledge: The epistemic-based irrational action premise asserts that if the belief really were knowledge, the knower would be irrational to decry the belief. The target theory must be false, since its classifying the belief as knowledge would result in one's acting both rationally and irrationally at the same time. The appropriate moral in this case is that expected utility theories require a more sophisticated notion of rationality than that expressed in the combination of (G1) and (G2). The moral is **not** that reliabilism is flawed.

I have encountered two types of objections to my treatment of the dichotomy argument.[6] The first of the two objections runs as follows: Your claim (that all theories are refuted by gun to the head argument) is false. Taken in isolation, the action of decrying the belief is inconsistent with one's knowledge, making one's decrying of the belief irrational. However, once one adds knowledge of the death threat to the knowledge set, decrying the belief becomes consistent with what one knows. Hence, one can decry the belief without acting irrationally.

I have two responses to the first objection. First, and most obviously, I do not endorse the gun to the head argument as a refutation of any theory of knowledge. Rather, I view the argument as indicating that something is wrong with the dichotomy argument's combination of rational action and irrational action premises. To the extent that readers reject the gun to the head argument, they seem to agree.

Second, I have no qualms with the claim that decrying the belief proves consistent with a larger knowledge set containing information about the potential harmful results of not decrying the belief. I should to note, however, that in the context of the objection the above claim is nothing more than an unsupported (though plausible) assumption. One can buttress the objection to the gun to the head argument only by giving a theory of how/when an action is consistent with a knowledge set. More importantly, I insist that should one accept the idea that decrying the belief is consistent with the larger knowledge set, then *prima facie* it looks as if one must also accept that buying a lottery ticket is consistent with a larger reliabilist knowledge set. One might plausibly take such a knowledge set to include the belief that the ticket will lose, the belief that it is nevertheless remotely possible that the ticket will win, and the belief that the probable loss of a dollar is insignificant when weighed against even the remotest possibility of winning millions of dollars. So, *prima facie*, the consistency move made in the first objection defuses the gun to head argument only if it also defuses the lottery argument against imperfect reliabilism. For the objection to work, the objector must offer a concrete and plausible notion of consistency between actions and knowledge sets which allows the gun to the head case, but which excludes the lottery ticket purchase. No such notion of consistency has been offered.

The second objection to my treatment of the above argument goes as follows: You falsely equate willingness to gamble with making statements under duress. The former case outlines conditions under which one would willingly acknowledge that the truth-value of the belief remains in question. The latter case outlines conditions under which one would say something, regardless of one's opinion of its truth-value, to avoid duress. Only the former case does the person’s action undermine the epistemic status of the belief. A torture-induced confession of a Gulf War P.O.W. would not undermine the epistemic status of his or her beliefs about the Unitied States. A willingness on the P.O.W.'s part to acknowledge that they could be wrong about the United States would undermine the epistemic status of their beliefs. If one knows that a ticket will not win, then no matter how large the jackpot, one can never rationally choose to buy a ticket.

I have three responses. First response; suppose for now that the two cases are different for the above-outlined reason (i.e., decrying the belief is rational, but no ticket purchase is rational). I doubt seriously that one can hold the above-outlined view without requiring certainty for knowledge.[8] According to the view, knowing a ticket will not win precludes the possibility of a rational ticket purchase. One also claims (according to the view) that the rationality of a ticket purchase is a function of the risk/benefit ratio. Yet, in order to insure that the benefit *never* exceeds the risk in cases where one knows the ticket will not win, one must insist that one be certain that the ticket will lose if one knows that the ticket will lose. Otherwise, one can always increase the benefit to counteract the risk. Of course, requiring certainty begs the question against the imperfect reliabilist by insisting that imperfect reliability is not sufficient for knowledge. I also argue in the last section of this paper that a certainty requirement looks implausible on independent grounds.

Second response; The cases are not importantly different. I acknowledge that one can make a statement denying the truth of a belief under duress without thereby doubting the belief in one's mind. However, I claim that the statement and betting cases are analogous for the purposes of the argument. Each case illustrates my claim that, unless one begs the question by requiring certainty for knowledge (and perhaps not even then), one can always find a point (the epistemic-utility fulcrum) at which instrumental concerns, not the truth/falsity or the warrant/lack of warrant of a belief, play the crucial role in determining action. In both cases, that point is reached when risk/benefit concerns override concerns about epistemic warrant--when being right becomes less useful instrumentally than living or having even a ridiculously slim chance at millions of dollars. In the gun to head case, the override point is not directly tied to the subject's perception of the epistemic status of the belief, as that status plays no direct role in the accruement of benefits. In the lottery case the subject's perception of the epistemic status of the belief plays a role in determining the override point since the perception of potential accruement of benefits is tied directly to epistemic status. Nevertheless, the important point for the argument is that in both cases epistemic concerns are eventually overridden by instrumental concerns in determining action.

Third response; there are examples where the perception of potential accruement of benefits is tied directly to epistemic status. For example, when a young atheistic post-doctoral fellow at the University of Rochester I distrusted the sincerity of many professed atheists. Thus, I once asked a colleague if he knew that God did not exist. When the colleague responded yes, I asked if him if he would at that moment ask Satan to claim his soul. This colleague told me, "No, I know God doesn't exist, but it's not a zero probability." According to the irrational action premise my colleague was irrational in refusing to acquiesce to my ideological challenge. He would not choose to, by his description, accept an incredibly minuscule risk of an eternity of torture for no gain. Was my colleague irrational? I have never thought so, nor has the profession judged him so as he is currently one of the most respected figures in his field.

D.M. Armstrong's version of the dichotomy argument also assumes a relationship between rationality and knowledge. Armstrong, however, only wants to refute imperfect reliabilism (1974, pp.184-90):

(A1) If one knows some set of propositions are each individually true, then it is rational to believe that the corresponding conjunction is true.

(A2) Suppose that one knows of each of some large number of propositions, n, that they are true.

(A3) Suppose further that one knows of the truth of these propositions because the process that generated the beliefs that the propositions are true has some level of reliability, m, where $m < 1$.

(A4) The probability of the conjunction of beliefs being true given the reliability of the process is $\boldsymbol{m^n}$ (assuming independence of each application of the process).

(A5) Given (A3), (A4), and a large enough set of beliefs, the probability of the truth of the conjunction of beliefs will be very low.

(A6) One is irrational to believe in the truth of the conjunction of such a large set of beliefs, even if all the beliefs are true, since, by (A5), the probability of the conjunction is very low.

(A7) Hence, (A3) must be false.

As with the first version of the dichotomy argument, the crucial premises of Armstrong's supporting argument are the rational action and irrational action premises. However, in Armstrong’s case the action is the forming or sustaining of a belief or beliefs. Therefore, I will call (A1) is the ***rational belief-forming premise*** and (A6) the ***irrational belief-forming premise***. Armstrong has erred, if by the irrational belief-forming premise (A6) he means that one's belief in a conjunction with low ***objective*** probability makes one irrational. Standard Gettier cases provide easy illustrations: Based their stellar reputations and his having watched them in philosophical forums many times, Wallis believes that his friends, Neiver Aspersion and Nomenda City, are brutally honest and deeply insightful philosophers. Wallis, therefore, accepts their assurances his counterexamples are clever and telling. But, unbeknownst to Wallis, to neither Neiver nor Nomenda can bring themselves to inform Wallis of his sad and complete bemusement. Were one to interpret the above case, in accordance with and objective probability reading of Armstrong’s irrational belief-forming premise (A6), Wallis’ believing the conjunction expressing Neiver’s and Momenda’s approvals would be irrational–despite his excellent but misleading reasons for his unreliably formed beliefs. Few people would assert that having false or unreliably formed beliefs for subjectively excellent reasons is irrational. Rather, Armstrong must mean that one is irrational to believe in a conjunction that has a low ***subjective*** probability.[7]

But, now there are two interpretations of the antecedent in the rational belief-forming premise (A1). On the first interpretation, the reliabilist is unaffected by the argument. On the second interpretation, Armstrong's rational belief-forming premise looks false.

Interpreters of the first stripe take the phrase ‘knows some set of propositions are each individually true’ to imply that one's subjective probability for the events associated with each belief equals 1. The subjective probability for the conjunction, then, ought to equal 1. (A1), therefore, undermines premises (A3) through (A6). So, on the first interpretation of the rational belief-forming premise, it is rational to believe in a conjunction when one knows that the conjuncts are true, ***even given imperfect reliabilism,*** since $\boldsymbol{p}^{n} \rightarrow 0$ as $\boldsymbol{n} \rightarrow \infty$ iff $\boldsymbol{p} < 1$.

Interpreters of the second stripe assert that (A1)'s antecedent phrase, ‘knows some set of propositions are each individually true’, means that the knower believes the propositions ***and*** the propositions happen to be true, but the subjective probability for the state of affairs associated with the belief does not equal 1. On the second interpretation, one can have knowledge without having psychological certainty, i.e., a subjective probability of 1. Imperfect reliabilists can accept such a requirement. But, if one does not require psychological certainty for knowledge, and thus for the associated beliefs, the rational belief-forming premise, (A1), looks false. I have a great number of beliefs of which I am subjectively less than certain. It is possible that unbeknownst to me all my beliefs are true and count as knowledge. But, this does not mean that I believe or should believe in the truth of the conjunction of all of these beliefs. Given my subjective probabilities and the probability calculus, I would be a fool to believe that all of my beliefs are true. In fact, I would violate Armstrong's irrationality premise (A6).

<u>4.) THE NATURE OF KNOWLEDGE AND THE NO KNOWLEDGE PREMISE</u>

Philosophers in the last section assume principles relating rationality to knowledge in their supporting arguments of the no-knowledge premise. The philosophers in this section assume principles concerning the nature of knowledge in their supporting arguments for the no-knowledge premise. As with the philosophers in the last section, some of the figures in this section wish to argue against reliabilism in all forms, while others only want to argue against imperfect reliabilism. For instance, Herbert Heidelberger is interested in refuting probabilism as a theory of knowledge. I extend the argument to imperfect reliabilism. Whereas the arguments in the first section attempt to forge a link between rationality and knowledge, Heidelberger wants to break such links and argue that ascriptions of rationality have a different logic than ascriptions of knowledge. Heidelberger's version of the supporting argument for the no-knowledge premise goes as follows (1963, pp.244-5):

(H1) If two beliefs count as knowledge, then a belief in their conjunction must also count as knowledge.

(H2) But, suppose that someone comes to know two propositions on the basis of a process having a some minimum reliability, $\boldsymbol{m} < 1$.

(H3) The probability of the conjunction of beliefs given their respective reliability is $\boldsymbol{m}^2$ (assuming independence).

(H4) Given (H2) and (H3), a belief in the conjunction of the antecedently known beliefs will not count as knowledge as the probability of the conjunctive belief will be less than $\boldsymbol{m}$.

(H5) Hence, (H2) must be false.

Heidelberger's argument turns on the plausibility of premise (H1). I call (H1) the ***closure under conjunction premise***. I think Alvin Goldman (1986, pp.83-4) has shown, the closure under conjunction premise fails to hold up under scrutiny: Suppose, for example, that Sheila knows that A and she knows that B. Suppose further that Sheila comes to believe the conjunction of A and B, not on the basis of her belief in A or in B, but on the basis of her beliefs that C and (A & B) → C. Since Sheila has committed the fallacy of affirming the consequent, it does not follow that she knows that (A & B)--even if (A & B) is merely the conjunction of two beliefs of which she independently has knowledge.

Examples such as the one above seem to show that having knowledge of conjuncts does not imply that one has knowledge that the conjunction holds. Instead, they suggest that knowledge of the conjunction depends upon the manner in which one generates the conjunctive belief. If Goldman and I are right about conjunctive beliefs, then whether or not a conjunctive belief counts as knowledge depends upon the process of belief generation. If the process involved in forming the conjunctive belief does not count as knowledge producing (i.e., have a sufficiently high reliability), then the conjunctive belief does not count as knowledge (even when I know each of the conjuncts).

As with the previous arguments, I have encountered objections to my treatment of the Heidelberger-type argument.[9] The first objection to my treatment of the Heidelberger argument goes as follows: Your argument against (H1) is somewhat of a red herring since one can modify (H1) to accommodate the Sheila case. Suppose one modifies (H1) in either of the ways below:

> (H1A) If two beliefs count as knowledge, then a belief in their conjunction formed on the appropriate basis of those beliefs must also count as knowledge.
>
> (H1B) If two beliefs count as knowledge, then a belief in their conjunction formed on the basis of those beliefs and a belief in the conjunction principle must also count as knowledge.

For the imperfect reliabilist, it remains indeterminate whether or not Sheila's conjunctive belief counts as knowledge--even if formed using the conjunction principle. To assert that the use of the conjunction principle always yields knowledge of a conjunction, therefore, begs the question against the imperfect reliabilist. Of course, if one could offer independent reasons to suppose that the conjunction principle held for all cases of knowledge, the charge of begging the question would ring hollow. In fact, the literature reflects the opposite verdict. It is widely acknowledged that holding an unrestricted closure under conjunction principle precludes knowledge in most cases of inclusive evidence or if one does not require perfect reliability for warrant.[9]

Peter Klein (1981 pp.190-201) has offered one of the few defenses of the rejection of lottery beliefs based upon the closure under conjunction principle. Klein considers the closure under conjunction principle as a principle of justification. He argues that one need not reject the principle. Rather, one needs to recognize that in cases where one has only "intrinsically probabilistic" evidence--evidence that explicitly assigns a numerical probability to the relationship between the state of affairs corresponding to one's belief and one's evidence--such evidence contains

> ...*unabsorbed* counterevidence for p_1 [citation omitted]. The evidence, e_1, assigns a high probability to p_1; but it also contains some evidence against p_1--namely, that in a specifiable number of cases, ticket 1 will not lose. The high probability, n-1/n, that ticket 1 will lose does not absorb the counterevidence that there is a 1/n chance that ticket 1 will win. (Klein 1981, p.196)

Klein suggests that in cases of inconclusive evidence of an exclusive and explicitly probabilistic nature, the belief is sufficiently undermined by the very evidence which supports it. However, Klein realizes that he too allows knowledge on the basis of inconclusive evidence, for instance, on the basis of testimonials. His rejection of the lottery paradox beliefs, as a result, seemingly undermines his own positive position. His response to such objections

> ...is that, although 'e_t&Q' [testimonial evidence, e_t, and an estimate of the likelihood of p_1 given e_t, Q] contains reference to the probability of the truth of p_1 on e_t, it is *not* intrinsically probabilistic evidence. ... [meaning] The proposition 'e_t&Q' contains probabilistic evidence for p_1, but it *also* contains evidence, namely e_t, which is nonintrinsically probabilistic, and Q is completely absorbed by e_t. (Ibid., p.198)

For Klein, evidence for a proposition which includes evidence that is not explicitly probabilistic is not intrinsically probabilistic evidence and, as a result, avoids lottery paradox problems. There are two serious problems with Klein's response. First, lottery cases will contain non-intrinsically probabilistic evidence, so that they do not differ from Klein's cases in that they involve only intrinsically probabilistic evidence. For example, one will likely hold the belief that there is a lottery ticket. **Second, the dynamics of conjunctive belief formation are in no way mediated by avoiding the explicit assignment of numeric probabilities to evidence**. When one forms a belief upon inconclusive evidence, there is always a chance that the evidence is misleading. A conjunctive belief inferred from two such beliefs inherits the chance of error from each belief. Nothing about these facts is altered by avoiding the explicit assignment of numeric probabilities to evidence relations. In short, no theory of knowledge or justification avoids summation of potential error in conjunction unless it either requires certainty or it separates justification from likelihood of truth so that justification does not track truth.

Consider Klein's own example. One forms a belief that $\boldsymbol{p_1}$ upon the basis of testimonial evidence, $\boldsymbol{e_t}$. There is a chance that the testimonial evidence is misleading with regard to $\boldsymbol{p_1}$, and one might represent one's awareness of this chance by expressing the evidence as a conjunction of $\boldsymbol{e_t \& Q}$. Nothing about the fact that Jones testified that $\boldsymbol{p_1}$ (i.e., $\boldsymbol{e_t}$) alters the fact that Jones' testimony may be misleading. If the mere chance that the evidence is misleading is sufficient to undermine the resultant belief's claim to knowledge, as Klein suggests above, then one's belief in

p_1 is undermined along with one's lottery beliefs. I argue in the last section of this paper that the chance of error need not undermine either the testimonial or lottery beliefs. So, it appears that Heidelberger's assumption of an unrestricted transitivity of conjunction principle is question-begging in that there is little reason to suppose that any possible theory will prove consistent with the principle as stated.

I turn now to a second objection against my response to Heidelberger-type arguments. The real issue, according to adherents of the second objection, concerns the dynamics of conjunctive belief formation on reliabilist accounts: If the reliability of a process producing conjunctive beliefs based upon beliefs in the individual conjuncts *always* comes out lower than that of the process which produces the inputs, then imperfect reliabilism results in the unlikely consequence that one can *never* have knowledge of a conjunction based upon knowledge of the conjuncts.

There are two problems with the second objection. First, lower probability need not always result in a failure to know. If the probability of the conjunction given the conjuncts is high enough, then the process will yield knowledge. Moreover, when the performance of real systems depends upon the truth of a conjunction, one generally finds that the system hedges its bets by employing tactics like large, redundant sample sizes and multiple samplings which have the effect of increasing the probability of the conjunction. In the eye, for instance, there are 100,000,000 rods, each of which pools its outputs, and can sample its environment every 1.2 seconds.

Second, not all cases of conjunctive belief formation involve processes having lower reliabilities than the reliabilities of the process associated with the formation of beliefs in the individual conjuncts. Suppose that I learn that P, Jones has a liver, and Q, Jones has a heart, each by independent processes having a reliability of .9. The fact that all living creatures with hearts have livers and vice versa means that my conjunctive belief that Jones has both a heart and liver results from a process that is more reliable than the ones which give rise to my beliefs in the original conjuncts (The probability of both process giving false answers equals $(.1)^2$, or .01. For both of the above-stated reasons, the reliabilist is better understood as accepting a restricted closure under conjunction principle rather than rejecting the possibility of conjunctive knowledge. In this case, the reliabilist is consistent with the practice of most epistemologists who place restrictions on such deductive closure principles.

Pollock's (1984) version of the supporting argument has the same general form as both Lehrer's (1974) and Kyburg's (1961 and 1983) versions.[11] Like Heidelberger, Pollock wants to employ a transitivity principle in his supporting argument. Pollock's principle is the transitivity of justification. As I said earlier, I am not particularly interested in justification. I, nevertheless, consider Pollock's argument, as his argument represents an important version of the lottery paradox (1984, pp.105-6):

> (P1) If "Louie the Loser" believes for each ticket that it will lose, then the reliabilist will tell us that Louie's beliefs are justified.
>
> (P2) "If each of a set of beliefs is justified for a person, and he correctly reasons from them to a conclusion, then either the conclusion becomes justified for him as well or else some of the initial beliefs become unjustified." (Ibid., p.106)
>
> (P3) But, if for each and every ticket Louie justifiably believes that that ticket will lose, then Louie could justifiably believe that all the tickets will lose.
>
> (P4) But, Louie cannot be justified in believing that no ticket will win.
>
> (P5) By (P2), then, one or more of Louie's beliefs concerning the individual tickets must not be justified after such a universal generalization.
>
> (P6) But, Louie's having made the universal generalization does not affect the

probabilities of any of the tickets winning. So, the reliabilist still must hold that the beliefs are justified.

(P7) Hence, reliabilism must be false.

Premise (P2), what I call *Pollock's principle*, turns out to be the crucial premise in Pollock's supporting argument. As with (H1A), what counts as "correctly reasoning" in (P2) constitutes the subject of the debate. The above formulation of (P2) proves useless against the imperfect reliabilist without a nonquestion-begging specification of "correctly reasoning." Unfortunately, Pollock's explicit commitment in (P1) to a reasoning strategy that treats each inference from the lottery odds to the belief that the individual ticket will loose as a series of reliable inferences is a remarkably inept mistake for someone with the intelligence and formal sophistication of Dr. Pollock. Just as the odds of picking a king from a card deck without replacement changes with each draw (i.e., 4/52 first draw, 4/51on the second draw....), the odds of, for example, the last ticket being a loser given that one believes all the previous tickets are losers is a mere 1/110,000,000–hardly reliable.

In the context of Pollock's example, moreover, (P2) seems to beg the question against the imperfect reliabilist once again: In the example Pollock's principle seems to assert that justification is closed under inference rules like those of predicate logic. Yet, as discussed above, the imperfect reliabilist is committed to denying such a closure principle.

Imperfect reliabilists who take an interest in justification assert that justification is a function of the process that generated the belief. If the process that generates the belief is not reliable, then regardless of the justificational status of the beliefs serving as inputs to the process, the generated belief fails to count as justified. Universal generalization is reliable only in those cases where the process reliably generates a true premise set.[12] The process generating Louie's premise set *never* generates a true premise set. The reliabilist need not, therefore, accept Louie's universal generalization as justifying. Justification need not be transitive on the reliabilist account.

But, does the imperfect reliabilist have to admit that Louie's isolated individual ticket belief must lose its justificational status in light of the justificational status of Louie's generalization? No. As Pollock notes, the reliability of the belief forming process for an isolated individual ticket remains unaffected by Louie's ill-fated generalization.

Actually, begging the question is not the only criticism one can offer against Pollock's principle. The principle also looks *prima facie* implausible on its own grounds. Pollock seems to be criticizing the reliabilist using a principle that he likely does not accept. The principle has two problems. First, if, as an internalist, one holds that the justificational status of a belief is something about which the person has conscious access, and as Pollock's principle suggests, at least one of one's initial beliefs becomes unjustified upon one's having made a correct inference to an unjustified conclusion, then one should be able to locate such a belief without further investigation. One should simply review the premise beliefs and note the unjustified belief. Of course, it is never that simple. A belief does not simply lose its justification. Its justification must be taken away as the result of further investigation.

However, suppose one interprets (P2) as asserting that if one of the premise beliefs from which one infers an unjustified belief becomes unjustified upon making the inference to the unjustified belief, then there should be some fact of the matter as to which belief is unjustified. One might then claim that (P2) requires only that the unjustified belief need be discoverable (as unjustified) without new evidence. Even on this interpretation (P2) seems false. Suppose, for instance, that a scientific theory predicts an outcome that scientists justifiably believe to be false.

The scientists do not immediately return to a now unjustified hypothesis and reject it. Nor do they immediately question all of the theory's hypotheses as well as the hypotheses that act as bridge principles in the operationalization of the predicted parameter. Locating the suspect premise(s) is, in fact, a difficult problem in the philosophy of science. The difficulty lies precisely in determining how to go about distributing blame given that all of the hypotheses are to some degree or another justified.[13] Almost always hypotheses about what belief(s) is unjustified are formed and tested by gathering further evidence specific to the hypotheses before any belief is deemed unjustified.

Second, if Pollock holds (P2), he should consider the following case: I infer the period of a pendulum from my beliefs about the length of the arm, the acceleration due to gravity, and the pendulum law. However, the periodicity turns out to be false when measured. According to Pollock, I must deny justification to at least one of my beliefs. However, I might well look successfully to frictional resistance in the pivot arm, etc.--never questioning my justification for believing in the law and the values for the initial conditions. This case seems to suggest not only that beliefs need not immediately lose their justificational status when used to infer an unjustified belief, but also that they may never lose that status.

I conclude that in the context of Pollock's example, (P1) is false, (P2) begs the question against the imperfect reliabilist in that it assumes an unacceptable closure principle for justification. Finally, (P2) also looks implausible on its own grounds.

Fred Dretske's argument is the last of the arguments I consider. Though Dretske formulates his argument within the framework of his theory, I find it illuminating because it gets at what I think philosophers really find disturbing about lottery paradox scenarios. Dretske wants to define *de re* non-inferential perceptual knowledge as beliefs caused by perfectly reliable signals. Dretske notices that lowering the minimum probability (of the state of affairs holding given the signal) required for knowledge below 1 seems to countenance as knowledge beliefs that offend our epistemic intuitions. Dretske's reasoning goes as follows (1981, pp.99-102):

(D1) Suppose one were to set the threshold probability (of the state of affairs given the signal) for information transmission (and thus, knowledge) at .9999999875 (or 1 in 80,000,000).

(D2) Suppose one were to come to believe, as a result of seeing a ticket in a twenty-three million ticket lottery, that the ticket did not win.

(D3) The belief that the ticket did not win counts as knowledge as the probability that the ticket did not win is .9999999875.

(D4) In the lottery case, one could come to know for each an every ticket (taken in isolation) that that ticket did not win.

(D5) In the case of the winning ticket, therefore, one could come to *know* that the winning ticket lost.

(D6) But, no one can *know* a falsehood.

(D7) Hence, the probability of the state of affairs given the signal must equal 1.

Of course, Dretske's version of the supporting argument ((D1)-(D5)) proves sound only if one does not have a truth clause in their definition of knowledge.[14] Dretske acknowledges this in a footnote. To Dretske's mind lottery cases would constitute situations where only the truth or falsity of a belief, and **no other factors**, determine whether or not the belief counts as knowledge. Despite the widespread acceptance of truth as a distinct component of a definition of knowledge, I think that many philosophers share Dretske's intuition as well as the sentiment he uses to describe the possibility:

> Although I say Malcolm "comes close" to endorsing this strange view, it is fairly clear that, in this passage at least, he does not actually accept it. He says that knowledge that P *can* differ from belief that P only with respect to the truth of "P." He does not say that it *always* does or that it does in situations such as that described in the lottery examples. (1981, p.250)

The idea that a cognizer could follow a process in two different cases, getting a true belief that counts as knowledge in one case, and a false belief not counting as knowledge in the other, curls many a philosopher's hair. I disagree. In the last section, therefore, I want to present a case for disregarding such intuitions. Specifically, I want to consider the question of whether information about populations and samples can lend epistemic warrant to a belief forming process. It is neither appropriate nor possible to give a full explication/defense of the positive view suggested in the final section. However, the sketch that I offer provides one with, I hope, a compelling reason to resist Dretske's intuition. I do not offer the reasons for rejecting the lottery as a blanket argument that beliefs formed on the basis of information about populations and samples universally have the status of knowledge. For instance, the reliabilist must still address Gettier, Irrationality, and Evil Demon examples.

4.) FACTS ABOUT POPULATIONS AND SAMPLES

The Dretske's intuition is just that, an intuition. One can strengthen or diminish its weight against imperfect reliabilism by evaluating it as a correct or a misleading intuition. Is the Dretske's intuition like the intuition that two plus two equals four, or is it like the intuition that space must be Euclidean? One's answer to such questions determines the attitude (accepting or dismissive) one should adopt towards the intuition.

In evaluating an intuition, one ought to determine if the intuition is systematic and widespread (uniform with regard to similar cases), or if the intuition is isolable from real practice and from theory. If one finds the latter, one ought to consider the possibility that the intuition is an artifact of one's thought experiments. Likewise, if one can explain the force of the intuition without attributing veracity to its content, one has yet another reason to consider it an artifact of one's thought experiments.

In evaluating the lottery intuition, therefore, I first wish to present a few examples that make it seem very plausible indeed that we do in fact make epistemically praiseworthy judgments in practice on the basis of information about populations and samples, and that we often accept such cases of uncertain inference as *prima facie* legitimate. Suppose, for instance, I want to know if I have disease **x**. I go to the doctor to take the test. I do not have disease **x** and my test comes back negative. Do I know that I do not have disease **x**? That would be the standard conclusion evinced in our day to day activities.

Yet, the test for disease **x** has a one in eighty million false negative rate. So, based solely upon the test results, the belief that I do not have disease **x** is only very highly probable (assuming that the probability is not dramatically impacted by the base rate). Moreover, taking multiple tests does not alter the dynamics of the situation. Multiple false negatives are just less probable. In other words, the only difference between myself and diseased patients with false negatives is the truth of my belief and the falsity of theirs. So, the only difference between the disease **x** and the lottery cases, seems to be that the same probability engenders a different epistemic judgement.

I also think that practice would allow that I knew that at the time of my first writing this paper, less than fifty percent of the citizens of Minnesota thought that their governor was doing a good job. I have not come to know of the governor's dismal approval rating because I asked every Minnesota resident. I have made an inference about the general population based upon a random sampling of that population. In other words, I believe that the governor had a low approval rating because, given the sampling, it is very improbable that more than fifty percent of the population thought that the governor was doing a good job.

Similarly, when I park my car and go inside for the evening, practice would seem to dictate that I know where my car is parked when asked five minutes later. Moreover, it would seem that the warrant of my belief is not undermined by the generic possibility of my car's having been stolen.

One might suppose that there is a difference between cases like the car case and the lottery case. In his paper, "Knowledge, Assertion, and Lotteries," Keith DeRose argues that we can account for the difference between our intuitions about all such cases of *prima facie* legitimate uncertain inferences and our lottery intuitions with the following solution:

> (a) although you believe you are a loser [of a lottery], we realize that you would believe this even if it were false (even if you were the winner), and (b) we tend to judge that S doesn't know that P when we think that S would believe that P even if P were false. (p.569)

However, DeRose's solution fails to capture many intuitions with regard to *prima facie* legitimate uncertain inferences, and hence his solution should be rejected. Specifically, DeRose's solution classifies beliefs that *prima facie* seem to count as knowledge, which are in fact treated as knowledge in real-world uncertain inference, as cases that do not count as knowledge. The sorts of intuitions about uncertain inference that DeRose's solution fails to capture are cases where we form a belief using processes that gives us good reason to believe something, but either that process, or another process that we would use in the relevant counterfactual situation, would not aid us in forming the contrary belief. In short, DeRose's solution fails to shed light on our intuitions with regard to the lottery paradox because it presupposes that lottery beliefs are judged not to be knowledge for reasons that apply to beliefs that we do count as *prima facie* knowledge. Consider a slight reformulation of the above medical test case:

Jones believes that he is HIV free on the basis of the fact that he has used condoms and selected his partners carefully. However, in order to satisfy his fiancee Jones goes to the doctor and takes a test. The test is a very cheap and effective screening test, having a very low false negative. That is, if the test says that Jones is negative, then it is almost certain that Jones is not infected. Jones' test comes back negative. Jones continues to believe that he is HIV negative based upon his test and his own knowledge of his sexual practices. However, the test has a shockingly high false positive (say 99%), and the doctor informs Jones of the false positive rate. ***In other words, though all the positives test positive so do large numbers of people who are not infected with HIV.*** If Jones had been positive the test would have come back positive. However, Jones would still have believed that he was HIV negative on the basis of his knowledge of his own sexual practices and the doctor's information that a positive test provides no evidence that he is positive more predictive than the base-rate in the population (i.e., it provides no good evidence). Hence, the process by which Jones forms his belief is such that he would have believed he was negative on the basis of the process even if he were positive (because the process provides him with no evidence he is positive, though it provides excellent evidence he is negative).

It seems clear to me that we would judge Jones' belief that he is HIV negative as a case of *prima facie* knowledge. Yet, the case violates (b), since Jones would not (for good reason) believe that he is positive on the basis of a positive test even if he were positive. Lest one suppose that this is a contrived counterexample, I should note that such a false positive scenario occurred when some friends of mine took a test to see if their unborn baby had Down's Syndrome. They continued to believe that their baby was not a Down's Syndrome baby in spite of the positive result on the basis of the fact that the test was no more reliable when positive than the base-rate (i.e., very unreliable). I would judge that I knew that my baby was not a Down's

Syndrome baby on the basis of the same test, though I would have maintained that belief in the face of a positive (even a true positive). In fact, such asymmetrical tests are commonly used as screening test in medicine precisely because they are very cheap relative to symmetrical tests (or tests to detect the condition) and the incidence of the disease is low in the population. In fact, this sort of asymmetrical reasoning can occur even in the most mundane case of uncertain inference, like my car example:

I park my car in front of my house and go inside. I believe that my car is parked out front on the basis of my having parked it there and my knowledge that in my virtually crime-free suburb the odds that my car has been stolen in the few minutes since I parked it are astronomically low. My belief would seem to count as *prima facie* knowledge. However, I would have believed my car was parked out front even if it had been stolen in the few minutes that had elapsed.

DeRose might claim that the above examples violate the spirit of the principle because in each of the cases the truth would come out in the end (i.e., other evidence would come to light, or the evidence would be different, etc.). Hence, Jones would *not* believe that X if X were false. Unfortunately, this claim seems to clash with the intuition that DeRose claims to capture in the original lottery case. Suppose that Jones believes that the ticket he just bought is a loser. In the end, the lottery board would disabuse him with regard to his lottery belief. So, such a move would undermine DeRose's claim to have captured common judgements with regard to lottery cases since he could not capture the original lottery intuition.

I conclude that DeRose's solution fails to capture the difference between cases occurring frequently in actual practice that people intuitively accept as *prima facie* cases of legitimate uncertain inference and cases that people do not accept as *prima facie* legitimate. For this reason his solution is not a solution.

So, these cases do seem to be cases of *prima facie* legitimate uncertain inference, and as such suggest that our intuitions with regard to lottery cases are unsystematic. Dretske's intuition also runs counter to philosophical theorizing about knowledge. For nearly two thousand years philosophers have thought of knowledge as something akin to justified true belief. Even post-Gettier maneuvering has lead very few philosophers to drop the truth clause of their definitions. But if one has a truth clause in one's definition of knowledge, then one presumably holds that one's definition allows for cases where only the truth of your belief and the falsity of another's separates your knowledge from their lack of knowledge. In short, any definition of knowledge with a truth clause either contains a redundant clause or sanctions cases of knowledge that violate the lottery intuition. Dretske recognizes the consequences of a truth clause, and does not have a truth clause in his definition of *de re* noninferential perceptual knowledge. Most philosophers, however, have kept truth clauses in their definitions. As a result, they must reject either their definitions or their intuitions.

So, practice and theory seem to run contrary to the lottery intuition. Can one explain the force of the intuition without attributing veracity to the intuition? I think I can. I note that one's subjective attitudes towards events, one's confidence in one's belief, one's feelings of safety, and so on, are not explicable merely by appealing to one's beliefs about the probability of the event. Clinicians who are given more information when diagnosing a patient, for instance, have a much higher level of subjective certainty in their diagnosis, even though their accuracy given the new information either remains constant or actually decreases slightly.[15] Likewise, physicians are more comfortable recommending a therapy when described as having a seven percent survival rate, then when described as having a ninety-three percent mortality rate.[16] I am never nervous when I drive to the airport, but I am always nervous on the flight. I have this reaction even though I know that the probability of a fatal crash per unit time when driving to the airport is greater than when flying out of the airport.

I think that the above phenomena drive Dretske's intuition. Philosophers worry about the error possible when following a reliable process. Reliable processes are suspect because mere reliability provides no guarantee against error, and when thinking about knowledge in philosophical contexts, one's subjective attitude towards a case is exclusively, or almost exclusively, shaped by worries about error. But, when focusing so intently upon error, philosophers fail to consider both (1) that there is a probability of error even when following explicit or lawlike rules of reasoning (once again, Hume sees this the most clearly), and (2) that the seriousness of error is at least as important as the possibility of error.

I devote the rest of this section to outlining my own reasons for supposing that information about populations and samples can produce knowledge. In other words, I have presented reasons to suppose that Dretske's intuition is both unsystematic with regard to real epistemic practice as well as theory, and explicable/debunkable by reference to one's misleading subjective attitudes towards events. I now want to offer a means of understanding how knowledge about populations and samples *could* lend epistemic warrant to beliefs.

I suggest that to understand knowledge, one must understand the role of knowledge in the performance of cognitive tasks. If one's task is to determine whether or not a patient has a disease one needs to generate veridical beliefs about the patient's health. Processes are strategies for generating an appropriate level of veridical beliefs from input. The doctor's input is the test result. Her strategy is to infer the patient's health from the test result and her knowledge about the probabilities associated with the test result.

Different tasks require more or less perfect strategies--higher or lower levels of reliability. The more false beliefs that are allowed the less perfect the strategy needed. Consider the designer of the now familiar lottery: He wants to make almost 107,000,000 dollars. He charges one dollar per ticket. One ticket will win three million. So, he can successfully achieve his goal merely by following the process of believing of each ticket in isolation that that ticket will lose. He will not always be right. However, he will achieve the prerequisite number of veridical beliefs (He will be right 106,999,999 times and wrong only once). He will successfully balance epistemic risk against epistemic need.

On the other hand, consider the lottery player: The lottery player wants to win the three million dollars. She has only finite resources, say, five dollars. The lottery designer has deliberately given the player only general win/lose information about the population of tickets. To achieve her goal, therefore, the player must employ a strategy that will generate a true belief about which ticket will win before she has spent her five dollars. The lottery player cannot avail herself of the believe-each-ticket-a-loser process due to its obvious inadequacy. She is clearly at a distinct epistemic disadvantage.

The same is true for Pollock's Louie. Louie wants to make a universal generalization about the earning power of a ticket, based upon his beliefs about individual tickets. In order to reliably make such a generalization Louie needs a process that will reliably generate a true *set* of beliefs about the tickets. Louie has no such process available to him, and given the situation and the information given Louie, no such process exists. Louie gains no knowledge about the potential earning power of an individual ticket by such an inference.

Dretske sees imperfect reliabilism as taking indefensible epistemic risk. That is, it utilizes a process that provides no guarantee against error. I hold that imperfect processes can generate knowledge--even though they do not guarantee truth--so long as they guarantee truth enough of the time. In other words, my suggestion is that the evaluation of the risk-taking involved in belief-forming processes must be made relative to the need for veridical beliefs associated with the successful performance of a cognitive task. While I too hold the epistemic goal of truth, I also recognize that it is often more important epistemically to get the *right* truths, even if one does not get them all the time, and that this may well require sacrificing some reliability for speed or power.[17] An imperfect process may or may not generate knowledge; the determining factor will be the assessment of the epistemic risk relative to the epistemic need.

BIBLIOGRAPHY

Armstrong, D. (1974). *Belief, Truth, and Knowledge*. Cambridge: Cambridge University Press.

Conee, E. and Feldman, R. (1998). "The Generality Problem for Reliabilism," in *Philosophical Studies* **89:** 1–29.

Derksen, A. (1978). "The Alleged Lottery Paradox Resolved," in *American Philosophical Quarterly*. 15: 67-74.

DeRose, K. (1996). "Knowledge, Assertion, and Knowledge," in *Australasian Journal of Philosophy*. 74: 568-579.

Dretske, F. (1981). *Knowledge and the Flow of Information*. Cambridge: MIT Press.

Dretske, F. (1981a). "The Pragmatic Dimension of Knowledge," in *Philosophical Studies*. 40:363-78.

Faust, D. (1984). *The Limits of Scientific Reasoning*. Minneapolis, MN: University of Minnesota Press.

Feldman, R. (1985). "Reliability and Justification," in The Monist. 68: 159-74.

Goldman, A. (1986). *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.

Heidelberger, H. (1963). "Knowledge, Certainty, and Probability," in *Inquiry* 6: 242-252.

Klein, P. (1981). *Certainty: A Refutation of Scepticism.* Minneapolis, MN: University of Minnesota Press.

Kyburg, H. (1961). *Probability and the Logic of Rational Belief.* Middletown, CN: Wesleyan University Press.

Kyburg, H. (1983). *Epistemology and Inference.* Minneapolis, MN: University of Minnesota Press.

Lehrer, K. (1974). *Knowledge.* Oxford: Clarendon Press.

Lehrer, K. (1990). *Theory of Knowledge*. Boulder, CO: Westview Press.

Nichols,S., Stich, S., and Weinberg, J. (2003). "Meta-Skepticism: Meditations in Ethno-Epistemology," in Luper, Steven (ed.) *The Skeptics: Contemporary Essays*. Ashgate-Publishing : Aldershot, England.

Piattelli-Palmarini, M. (1991). "Probability: Neither Rational nor Capricious," in *Bostonia* March/April: 28-35.

Pollock, J. (1983). "Epistemology and Probability," in *Synthese*. 55: 231-252.

Pollock, J. (1984). "Reliability and Justified Belief," in *Canadian Journal of Philosophy*. 14: 103-114.

Sanford, D. (1981). "Knowledge and Relevant Alternatives: Comments of Dretske," in *Philosophical Studies*. 40: 379-388.

Swain, M. (1981). *Reasons and Knowledge*. Ithica, NY: Cornell University Press.

Wallis, C. (1994). "Truth-Ratios, Process, Task, and Knowledge," in *Synthese* 98:243-69.

Yourgrau, P. (1983). "Knowledge and Relevant Alternatives," *Synthese*. 55: 175-190.

NOTES

*This paper was developed in part during my tenure as a Mellon Post-Doctoral Fellow at the University of Rochester. I have presented the paper at both the Northwest Philosophical conference and the Central Division American Philosophical Association. I owe special thanks to Prasanta Bandyopadhyay, Paul Bartha, Ronald Giere, Alvin Goldman, Henry Kyburg, and George Pappas.

1.) Alan Goldman (1990), John Pollock (1983 and 1984), Fred Dretske (1981), Keith Lehrer (1974), D.M. Armstrong (1973), Herbert Heidelberger (1963). I discuss all of these figures in this paper. Henry Kyburg is, of course, the philosopher who first thought about the lottery paradox in relation to epistemic issues. His *Probability and the Logic of Rational Belief* (Middletown, CN: Wesleyan University Press, 1961) is universally cited. However, Kyburg's interest in the lottery paradox focuses upon rational belief formation, not upon reliabilism or knowledge.

2.) Nichols, Stich, and Weinberg have shown that intuitions on Getter cases are niether univocal nor universal.

3.) Specifying the relevant population is one of the three central issues for imperfect reliabilism. I call this the relevance class problem in both my dissertation and a *Synthese* paper, where I propose a solution. Feldman (1985, 1998) discusses these issues under the “single case,” “no distinction”, and “generality” problems.

4.) See Dretske (1981 and 1981a) for one such attempt at pragmatism. See David Sanford (1981), Palle Yourgrau (1983) for critiques.

5.) James Maffie deserves thanks for suggesting to me that my point was best made by talking about conceptions of rationality.

6.) The various objections given here have been suggested and/or inspired by an anonymous commentator or during discussion of a version of this paper at the Northwest Philosophical Conference and Central Division APA.

7.) Interestingly, it seems to me that even if one requires subjective certainty for knowledge one can generate a bet that seems rational, at the very least, it does not seem irrational. Suppose that Descartes’ evil demon plans to torture you with excruciating pain for all eternity. You have one chance to escape this fate, winning a bet. The bet, however, is that something of which you are subjectively certain is actually false. If you lose, then you lose nothing. In this case, it seems false to suppose that accepting the bet is irrational. I would even suggest that it is rational, since you have nothing to lose and everything to gain.

8.) I must stop here for two digressions: First, as it turns out people are not as apt as one might like in applying the conjunction rule. Second, it has been suggested to me that the relevant sense of probability for the argument is epistemic probability. I do not consider epistemic probability in the body of the paper since: (1) I doubt that Armstrong had epistemic probability in mind. (2) Epistemic notions of probability are usually developed so as to defuse the lottery paradox. For instance, Henry Kyburg's notion of epistemic probability involves the rejection of the conjunction rule. Kyburg, therefore, would reject (A4). It is possible that Armstrong's argument might work given the right theory of epistemic probability. I will reject the argument until given that theory.

9.) Again, these objections have been suggested or inspired by commentators.

10.) Many philosophers have come to reject the transitivity of justification as the result of the lottery paradox. Kyburg's position in *Epistemology and Inference* (Minneapolis, MN: University of Minnesota Press, 1983), pp.36-7, clearly entails the rejection of transitivity for rational belief formation. Derksen (1978) also rejects transitivity. Lehrer (1974) adopts a more complicated position.

11.) This is not to say that Pollock uses the argument towards the same end as Kyburg and Lehrer. Lehrer is interested arguing against equating high probability with complete justification in relation to foundational theories of knowledge. So, it is plausible to count Lehrer among the anti-reliabilists. Kyburg concerns himself with developing a rational corpus of beliefs, and wants to reject the lottery paradox. I think it plausible to count among him the reliabilist sympathizers.

12.) Pollock claims that we cannot attack his argument at this point because we cannot specify in general a subset of the conditions under which a method is reliable, i.e., pick out the relevant population. According to Pollock such specifications inevitably lead us down a slippery slope to the position that "...a non-discursive belief is justified only if it is true." (Pollock 1984, p.109) Pollock's argument is unconvincing since many reliabilist, including myself, have offered theories concerning the specification of relevant populations which Pollock does not address. The objection is also beyond the scope of this article.

13.) Glymour's *Theory and Evidence* contains a good discussion of exactly this point in chapters 3 and 6.

14.) Occasionally someone objects that I have misrepresented Dretske here. Fortunately, Dretske himself agrees with my characterization of his argument. I should also note that Dretske's requirement of a probability of 1 may not insure that he has no false beliefs under some circumstances. For instance, the probability that a true random number generator (ours are biased towards the lower numbers) will not generate a prime number below ten is 1. But, it could still happen. One *might* be able to avoid this problem on a propensity interpretation.

15.) This example is documented in Faust (1984) pp.70-2.

16.) This example is documented in Piattelli-Palmarini (1991) p.30.

17.) Here again, I agree with Goldman 1986 (pp.122-142)