

Imprudence and Immorality: A Kantian Approach to the Ethics of Financial Risk

Tobey K. Scharding

Bloomsburg University of Pennsylvania

ABSTRACT: This paper takes up recent challenges to consequentialist forms of ethically evaluating risks and explores how a non-consequentialist form of deliberation, Kantian ethics, can address questions about risk. I examine two cases concerning ethically questionable financial risks: investing in abstruse financial instruments and investing while relying on a bailout. After challenging consequentialist evaluations of these cases, I use Kant's distinction between morality and prudence to evaluate when the investments are immoral and when they are merely imprudent. I argue that the investment practices are imprudent when they do not take adequate precautions to secure the firm's long-term flourishing. They are immoral in a Kantian sense when they risk the destruction of the financial system upon which the firms depend. The upshot of my analysis is that moral actions require more risk aversion than prudent actions and prudent actions require more risk aversion than expected-value-maximizing actions.

KEY WORDS: Ethics of Risk, Financial Risk, Kantian Ethics, Formula of Universal Law, Investment Ethics, Too Big To Fail

0. INTRODUCTION

IN MAKING RECOMMENDATIONS ABOUT WHAT TO DO IN RISKY CIRCUMSTANCES, which present the possibility of material loss, theorists generally calculate the probability that various outcomes—adverse, neutral, or positive—will come to pass. They establish each action's "expected value" by multiplying the numerical utility of each possible outcome by the probability that the outcome will occur and summing the probable-utility calculations. With these numbers, decision makers can begin debating what to do. Typically, the action with the highest overall expected value is preferred. As fruitful as this decision-making procedure is as regards questions concerning risk, it represents only a "consequentialist" viewpoint, focused on the outcomes of action and excluding other features of action. In this paper, I take up recent challenges to consequentialist forms of ethically evaluating risks and explore how a non-consequentialist form of deliberation, Kant's moral theory—which offers a strategy for ethically evaluating actors' *plans* for acting, rather than the *consequences* of their actions—can be used to address ethical questions about risk.¹

I begin by explaining the prevailing consequentialist strategy for evaluating risk. My focus is financial risk;² the risk takers that concern me are investment firms, especially as regards their brokerage activities of buying and selling financial securities

that present a risk of financial loss.³ I am particularly interested in the role of such risks in precipitating the 2008 financial crisis and I use examples from the run up to this crisis to illustrate and develop my theoretical concerns. These are, primarily, two ethical questions about financial risk. First, may investment firms risk money in abstruse financial instruments that are highly complicated and interconnected in intricate, hard-to-parse ways? Second, may investment firms rely on being bailed out, or rescued by an outside agency, when they put money at risk? I begin by exploring several challenges that consequentialist decision making faces in answering these questions. The worries involve (1) two ways in which consequentialist answers to these questions rely on probabilities that are controversial and (2) two ways in which consequentialist answers to these questions fail to address relevant values, such as how the financial risk is distributed.

Next, I explain how Kant's distinction between acting for reasons of prudence and acting for reasons of morality can avoid many of the problems associated with consequentialism.⁴ Drawing on Kant, I build an analytic framework for addressing questions of morals and prudence as they relate to business activities. I then use this framework to reconsider the two ethical questions about financial risk, evaluating when (if ever) these financial activities are immoral and when (if ever) they are imprudent.

My thesis is that these investment practices are sometimes imprudent, when they do not take adequate precautions to secure the firm's long-term flourishing. The practices are sometimes immoral in a Kantian sense—violating a rule of right conduct that the firm simultaneously recognizes as being generally binding, i.e., to which it makes itself an exception—when they risk the integrity of the financial system upon which the firm depends. An upshot of my analysis is that Kantian moral reasoning appears to require more risk aversion than prudential reasoning and prudential reasoning appears to require more risk aversion than expected-value reasoning. Relatedly, acting morally assures that key standards of prudence are met and acting prudently captures important aspects of expected value maximization.⁵

1. CONSEQUENTIALISM AND RISK

In this section, I set forth (1.1) the predominant, consequentialist form of reasoning about risk, expected value theory. I perform some rudimentary expected value calculations on two examples: (1.1.1) an abstruse investment and (1.1.2) a bailed-out investment. I then consider (1.2) objections to this way of reasoning about risk.

1.1 Expected Value Theory and Risk

“Expected value” decision makers rely on probabilistic, consequentialist forms of reasoning to deliberate the advisability of risky actions. Using this form of reasoning to decide questions about risk seems eminently appropriate. The problem that risky actions present concerns the probability of good or bad consequences; in deliberating whether to undertake risky actions, decision makers focus on achieving the former and avoiding the latter. Expected value relies on three analyses, related to the advisability of a particular action: (1) what are the possible consequences of the

action, (2) what is the probability that each of these consequences will occur, and (3) what is the utility of each of the possible consequences.⁶ Decision makers multiply the probability of each consequence by its utility and add together these values for all of the action's possible consequences to determine the "expected value" of the action. After computing the expected values of each of the possibilities for action that they face, decision makers choose the action with the highest expected value.

1.1.1 The Abstruse Investment

To demonstrate this kind of deliberation, consider the case of the *abstruse investment*.⁷ In this scenario, a firm questions whether it should risk money in an investment in which highly complicated financial instruments are interconnected in intricate, hard-to-parse ways that make the implications and risk profile of the investment difficult to discern (hereafter, in an investment that is "abstruse"). To use expected value theory to answer this question, decision makers perform calculations to determine which action, risking money in the abstruse investment or not risking money in the abstruse investment, has the higher expected value.

As an example, take a firm's decision about whether to invest in one of the very complicated financial securities that were heavily traded prior to the 2008 financial crisis, the synthetic collateralized debt obligation (or "synthetic CDO"). Banks create synthetic CDOs in reference to other financial securities, such as mortgage-backed securities (MBS). Whereas MBS are based in real assets—the collateral-backed mortgages that homeowners owe on their properties—synthetic CDOs only *refer* to MBS: they are bets about whether borrowers will default on their loans and are not based in real assets. As wagers, CDOs require both "long" investors, who bet that the borrowers will pay, and "short" investors, who bet that the borrowers will default. "Long" investors include both "funded" long investors—who buy the securities referenced by the CDO and receive interest as the loans are repaid, losing their investment if the borrowers default—and "unfunded" long investors, who wager on the loan's repayment but have no assets invested in the loan itself. Unfunded long investors receive premium payments from the short investors as long as the loan is being repaid but must pay premiums if the borrowers default (FCIC, 142).

In 2004, the U.S. bank, Goldman Sachs, created a synthetic CDO, Abacus 2004-1, referencing residential MBS, existing CDOs, and commercial MBS. The German bank, IKB, the asset management firm, TCW Group, and the U.S. bank, Wachovia, invested \$195 million in the funded long position. Goldman itself was the largest funded long investor, holding \$1.8 billion in that position. These investors stood to make millions if the loans were repaid. TCW and GSC Partners, another asset management firm, became unfunded long investors (FCIC, 142-43). These unfunded long investors did not pay initial investments but would owe short investors if the borrowers defaulted. As such, they could potentially gain, or lose, millions. The short investor was also Goldman Sachs. In its role as an investment firm engaging in brokerage activities, Goldman bet \$2 billion that the borrowers referenced by Abacus 2004-1 would default. This short investment limited Goldman's loss potential as a long funded investor.

It is obvious from the above description that the investors faced a complicated decision about whether, and how, to invest in Abacus 2004-1. I consider, in particular, GSC's decision to take an unfunded long position. The financial engineering behind this investment—including the combination of so many different kinds of MBSs and the investment's independence from cash assets—made it hard to understand and price. Using expected value theory, GSC would calculate the probabilities that various loans would default and use these calculations to determine whether the probability of profits outweighs the probability of losses. As an unfunded long investor, GSC's potential for profit was huge—\$2 billion at stake in the loans—but its potential for loss was equally high. In a simplified expected value calculation about Abacus 2004-1—deciding between actions (A1) invest long, (A2) invest short, and (A3) do not invest—it is obvious how expected value theory could recommend that GSC take the unfunded long position. It would recommend that position whenever investing unfunded long in Abacus 2004-1 had the highest expected value of GSC's possible actions, such as when borrowers were more likely to pay off their loans than to default. Because the instrument was abstruse, though, GSC's reasons for believing that the unfunded long position had the highest expected value would have been hard to discern—perhaps even to itself.

1.1.2 The Bailed-out Investment

In the case of a *bailed-out investment*, a firm questions whether it should rely on being “bailed out,” or rescued by an outside agency, when risking money. Expected value theory would recommend relying on the bailout—and risking money that would not otherwise be risked—whenever doing so has a higher expected value than not relying on the bailout (and not risking the money).

As an example, consider decisions made by the Federal National Mortgage Association (“Fannie Mae”) shortly before the 2008 crisis. During this time, the firm was actively purchasing mortgages from commercial banks. The U.S. government created Fannie Mae to encourage home ownership by establishing a “secondary market” for mortgages, such that commercial banks could sell their existing mortgages, receiving funds with which to make more mortgages. It became a publicly traded corporation in 1968 but maintained close ties with the U.S. Treasury as a “government-sponsored enterprise” (GSE). These close ties provided Fannie Mae—and a related GSE, the Federal Home Loan Mortgage Corporation (“Freddie Mac”)—with exemptions from state and local taxes and \$2.25 billion credit lines from the Treasury. The Federal Reserve System (“the Fed”) also allowed Fannie Mae and Freddie Mac to borrow money at very low rates, almost as low as the Treasury itself. Because of such privileges, many investors believed that the government “implicitly” guaranteed the GSEs (FCIC, 38).

In the summer of 2007, commercial banks became increasingly unwilling to grant new mortgages. The U.S. home mortgage market had weakened: home prices had declined and more borrowers were defaulting on their loans. At this time, the GSEs were already exposed to significant risk: they were responsible for \$5.3 trillion in mortgages, very close to the limit of how many loans they could legally fund. Even in the weakening housing market, though, there were opportunities for profit. Because fewer investors wished to purchase MBSs, the available MBSs would cost the GSEs less than they would in a booming market. On August 1, 2007, Fannie Mae CEO

Daniel Mudd requested that the U.S. government lift its limits and allow Fannie to buy \$150 billion of additional loans (FCIC, 309-12).

In this situation, we can see how considerations about bailouts might figure into an investment decision. Mudd wished to purchase more MBSs in order to maximize Fannie's profits. Using expected value theory—fueled by information about profits on other mortgages, by information about the housing market, and by estimations of how the housing market was likely to change—Mudd could calculate the profit potential of purchasing additional MBS. As part of these calculations, he would also consider possible gains or losses associated with the MBSs that Fannie already owned or guaranteed. If Mudd can rely on being bailed out, the possible losses need not advise against making the further investments. Should Mudd rely on being bailed out when deciding whether to purchase \$150 billion more in loans?

In order for the decision to buy more mortgages to make sense for Fannie Mae from an expected value standpoint, Mudd must calculate that making more loans is likely to maximize profit. The consequences that concern him are: (C1) the loans are repaid, producing large profits, (C2) the loans are not repaid and the government provides a bailout, producing neither profits nor losses, and (C3) the loans are not repaid and the government does not provide a bailout, producing large losses. Mudd must estimate both exactly how large the profits or losses are likely to be as well as how probable it is that the government will provide a bailout if the loans are not repaid.

Assume, for the purposes of demonstrating how expected value theory would address this decision, that there is a 50% chance that the loans will be repaid and a 50% chance they will not be repaid. Assume that profits and losses have the same magnitude: 10% of the original investment, or \$15 billion. Under these conditions, if the possible actions are (A1) make \$150 billion in additional loans (relying on a bailout) or (A2) don't make the loans (assuming that the government will not provide a bailout), expected value theory is indifferent only when there is a 0% chance that the government will provide a bailout. Whenever there is greater than 0% probability of a bailout, expected value theory advises Fannie to make the loans, given the assumptions that there is a 50/50 chance the loans will be repaid and possible profits and losses have the same magnitude.

The latter assumption, though, could be false when Fannie is already responsible for \$5.3 trillion in mortgages, some of which could produce losses when the \$150 billion in new mortgages produce losses. In a second expected value calculation, then, continue to assume that the possible profits are 10% but assume now that the possible losses are larger: 20% of initial investment. For the two possible actions, (A1) add \$150 billion in loans to the \$5.3 trillions in loans that have already been made and (A2) do not add \$150 billion in loans, the expected value calculations are as follows, given the assumption of 10% profits and 20% losses and using the three consequences defined above:

$$\begin{aligned} U(C1|A1) &= \$545 \text{ billion} \\ U(C2|A1) &= 0 \\ U(C3|A1) &= -\$1,090 \text{ billion} \end{aligned}$$

$$\begin{aligned} U(C1|A2) &= \$530 \text{ billion} \\ U(C2|A2) &= 0 \\ U(C3|A2) &= -\$1,060 \text{ billion} \end{aligned}$$

Given these assumptions, expected value theory is indifferent between the two actions when the probability of a bailout is 50%. Expected value theory prefers A1, making the further \$150 billion in loans, whenever there is a greater-than-50% probability of a bailout. With the admonition that many informed people thought (prior to the 2008 crisis) that it was extremely likely that Fannie would be bailed out, even assumptions of far larger losses could still produce this expected-value-theory recommendation: make the further \$150 billion in loans in the declining housing market.

1.2 Problems With Expected Value Calculations As Regards Risk

It is clear that there are many areas of potential weakness in the consequentialist reasoning performed above. In this section, I discuss two main problems, each of which has two subparts. First, it is difficult for expected value decision makers to make their probabilities precise because of problems concerning (1.2.1) judgment: either in assigning the reference group by which the probabilities are estimated or in estimating probabilities when there is no reference group. Second, expected value neglects values that seem ethically relevant to decisions about risk, including (1.2.2) the distribution of the risk and (1.2.3) the “free-floating fear” created by the risk.

1.2.1 Problems In Judgment

In order to determine the probability that abstruse investments will produce profits, decision makers require data from many comparable situations. The data must be analyzed and the abstruse investments assigned to a group of similar investments. Decision makers then examine how many of the similar investments were profitable and how many lost money in order to establish the probability that their abstruse investments will succeed.

The problem in assigning the abstruse investments to a suitable reference group (hereafter the problem of *assigning reference*) is that many different kinds of investments are “comparable” to the abstruse investments and the “suitable” reference group can be construed in many different ways. Some investments will have been made in similar domains: commercial real estate, or residential. Some will have been made in similar geographical areas: the Midwest, San Francisco, Georgia. Some will have been similar sizes, numerically. Whichever reference group the decision maker assigns, the assignment will be to some extent a matter of the decision maker’s own judgment. As Perry notes, each probability-affected entity “belongs to an indefinitely large number of reference classes” (2007, 335). Thus, the decision maker’s calculation of probabilities could be controversial. Indeed, this problem is associated with virtually all assignments of probability. Selecting a (uncontroversially) “best” reference class is often practically infeasible.

A related problem is that each investment is to some extent unique and governed by forces not found in the same combination in any previous investment. For abstruse

investments, this difficulty is acute. As an example of the problem, consider the strategies that firms used to manage risk prior to the 2008 crisis. The predominant model was “Value at Risk” (VaR) developed by J.P. Morgan, the American bank. VaR claimed to predict, with 95% certainty, how much a firm could lose if market prices changed. The model used historical data to make its predictions; with respect even to MBSs and especially to the more abstruse financial instruments to which MBSs lead, such as synthetic CDOs, the models had almost no data from which to draw (FCIC, 44). This left firms poorly prepared to assess their risk exposure prior to the crisis. Even if they had been inclined to limit exposure, then, they were ill equipped to do so. Call this the problem of *no reference*.

The challenge facing bailed-out investments is also serious. Very few banks have been considered “too big to fail” (TBTF)—or eligible to be bailed out when they become overextended—in the history of Wall Street. Prior to the 2008 bailouts, only one bank had received federal assistance because of its systemic importance to the overall economy.⁸ In 1984, the U.S. government loaned the Continental Illinois Bank \$4 billion in order to keep it from collapsing. This bailout inaugurated the idea of TBTF but no doctrine was ever clearly set forth. In particular, the government never specified criteria for what makes a bank TBTF. The practice was employed for a second time in 1998 when the Fed arranged for fourteen banks to provide \$3.65 billion to a hedge fund, Long-Term Capital Management (LTCM) when LTCM’s investments produced huge losses at the end of the 1990s (Williams 2010, 87-88). This extremely weak track record challenges expected value theory’s ability to determine the probability that a particular bank will be bailed out.

In response to such worries, an expected value theorist might note that the theory does have resources to account for such concerns. In particular, expected value calculations can include not only (a) estimated probabilities of consequences but also (b) estimated probabilities of how confident the decision makers are about (a), their estimated probabilities of outcomes. In this sense, expected value theorists can specify when their calculations are more uncertain, such as in those concerning abstruse and bailed-out investments.

Adding additional probability calculations can clearly mitigate the problems of assigning reference and no reference. But such additional probability figures will be difficult to calculate—and dependent on judgment—for the reasons discussed above. As such, they will not remove judgment from expected value calculations.

1.2.2 Problems As Regards The Distribution of Risk

Secondly, consequentialist reasoning fails to account for other values that seem relevant to decisions about investment risk. In particular, expected value reasoning does not consider the distribution of risks of harm (Hayenhjelm and Wolff 2012, 14) (hereafter, the problem of *distribution*). The calculations performed above consider risk to the investment firm, only. But the firm’s actions may affect many other parties. In particular, the calculations neglect possible consequence (C4): the investments produce catastrophic losses for the economy as a whole, undermining the economic system itself.

Consider the 2008 crisis, in which ordinary homeowners were permitted to suffer foreclosures, depleting their life savings, even as large financial institutions—including ones that were partially responsible for the crisis—received billions in bailout funds. Long-lasting catastrophe was avoided for the country as a whole by the government bailout: but it was taxpayers who bore the burden, while many wealthy investors were spared.

1.2.3 Problems About “Free-Floating Fear”

The consequentialist framework also does not account for the “free-floating fear” created by risky environments, as discussed by Nozick (1974, 65-69). People—including those who do not suffer material harms—can be traumatized by the risk of harm. A risk-filled environment might directly undermine welfare or might disincline people from taking welfare-producing actions. Needless to say, it is difficult to establish a clear line of causality from such an environment to the harms that people affected by the environment experience. Such difficulties challenge consequentialism to the extent that consequentialist calculations must be able reliably to predict the consequences of actions in order to take them into account.

Consider, for example, the various kinds of fearful behavior that Americans exhibited following the financial crisis of 2008. As one commentator noted, Americans bought more guns, took more money out of the bank, bought more gold, and went to church more (Noonan, 2009, A9). This behavior probably influenced some firms’ profitability but it would be strange for those firms to take social effects, whose occurrence and effect on productivity is highly uncertain, into account in their expected value calculations.

In response, an expected value theorist might argue that expected value calculations can simply assign a disutility value to free-floating fear. Again, although this strategy would help to mitigate the problem presented by free-floating fear, it would be very difficult to calculate precisely the disutility value (due to the problems in establishing chains of causality, as discussed above). In this sense, the problem can be mitigated but not fully addressed.

2. PRUDENCE AND MORALS IN KANT

Although expected value theory is clearly very helpful in examining questions about risk, the discussion offered above demonstrates that it is subject to certain weaknesses. These weaknesses motivate interest in other strategies for addressing ethical questions about risk, which might be free from the problems considered above. In this section, I investigate how one of the dominant forms of non-consequentialist reasoning, Kantian ethics, can evaluate risk. I am especially interested in Kant’s distinction between standards of prudence and standards of morality as regards the permissibility of risky actions. Moral actions, for Kant, are motivated by respect for the moral law and seek to conform to the moral law, without regard for the actions’ consequences. “Merely” prudent actions, by contrast, focus on consequences, seeking to achieve long-term prosperity and avoid catastrophe (rather than maximize expected value). In this section, I explain the textual basis for the distinction and describe

how I will use prudential decision-making and Kantian moral decision-making to recommend decisions about financial risk.

2.1 *Textual Basis for Prudential and Moral Decision-Making*

Kant illustrates⁹ the difference between moral actions and actions that serve merely prudent ends via the case of the *honest shopkeeper*. The honest shopkeeper returns the proper change to his young customers even though the customers are unable to count that change.¹⁰ Kant points out that an observer cannot determine whether the shopkeeper's action in returning the correct change is *moral*, or undertaken because the businessman recognizes that it is his moral duty to return the correct change, i.e., because his store depends upon a moral norm in which shopkeepers never intentionally fail to return the correct change to their customers. An observer cannot confirm that the action is moral because this action is also *prudent* for the shopkeeper. Returning the correct change to customers who are incapable of discerning fraud serves the shopkeeper's long-term material self-interest by preserving his good reputation: though the children won't complain about being shirked on change, their parents would.

2.2 *Prudential- and Moral Decision-Making Procedures in Kant*

2.2.1 Prudential Decision-Making Procedures

Based on this example, I understand prudential reasoning as recommending actions that serve the decision maker's long-term material prosperity. In particular, prudent actions must not undermine the decision makers' abilities to satisfy their future interests in the same manner by which they satisfy their present interests, e.g., by way of their present good reputations or their present businesses. Considering whether an action is prudent appears to be a form of consequentialist thinking by Kant's lights. The shopkeeper, in acting prudently, is concerned to assure that the consequences of his action—returning the correct change to a customer who cannot count the change—are positive.

I add two provisos to the understanding of prudence offered above, not discussed by Kant but consistent with his account of the honest shopkeeper. First, prudent actions do not guarantee good outcomes; rather, they establish appropriate safeguards against bad outcomes. The shopkeeper could still lose his business. Second, assessments of prudence are relative; whether a decision maker's precautionary activities count as prudent varies from decision maker to decision maker based on, among other things, the decision maker's individual objectives and preferences. What would be prudent for one shopkeeper might not be prudent for another.

To make a decision about what is the right thing to do according to this conception of prudential reasoning, decision makers must consider what are their long-term interests and must act so as to assure that these interests remain viable. Interestingly, prudent decision makers may avail themselves of expected value calculations; it might even be prudent, in the sense of promoting their long-term material prosperity, for them to do so. The answer to their question, "What is the right investment for

me at this time?" will likely differ from the one offered by expected value theory, however, as discussed further in sections 3.1 and 3.2, below.

2.2.2 Moral Decision-Making Procedures

To understand whether an action is moral in a Kantian sense, decision makers must first formulate the "rule" governing the action, or describe the action to be performed in the form of a general counsel that anyone in sufficiently similar circumstances could follow. There is rough agreement among Kantians that the rule should include the action taken and the circumstances in which the action is taken.¹¹ The rule should be formulated neutrally: in a way that both decision makers and outside observers can endorse. Then, decision makers test whether everyone in similar circumstances can follow the rule.¹² If the action becomes self-defeating when the rule is universalized—such that the actor would be unable to achieve his or her aims in acting if everyone acted in this way—then the action is an exception to a universally binding moral rule, i.e., is morally prohibited. If the rule can be universally followed, however, such that the rule holds the actor to the standard the actor expects from others, then the action is morally permissible.

In evaluating the action to be performed from a non-consequentialist, Kantian perspective, decision makers may consider the idea that the outcome of their action is risky: it could turn out well or badly. Kantian decision makers may even avail themselves of expected value calculations to determine exactly how risky is the proposed action. The difference between using expected value calculations to reason about risk from a non-consequentialist, as opposed to a consequentialist, perspective is that in the case of the former, the outcomes of the expected value calculations are not what determine the permissibility of the action. From the non-consequentialist, Kantian perspective, what determines the permissibility of the action is whether the action can be carried out universally.

3. IMPRUDENCE IN A CONTEXT OF RISK

In this section, I begin using my Kantian framework for prudential decision-making to evaluate questions about investment risk. I begin by revisiting the cases introduced above: (3.1) the abstruse investment and (3.2) the bailed-out investment. Then, I consider some (3.3) problems with prudential reasoning about risk.

3.1 Is the abstruse investment prudent?

Whereas expected value calculations, such as those discussed in section one, generally seek to maximize financial profit, considerations of prudence need not aim for a materially optimal outcome. Rather, in resolving to act prudently, decision makers seek primarily to preserve their ability to pursue similar courses of action in the future. This naturally motivates them to avoid disastrous outcomes¹³ but does not require them to seek optimal outcomes. The honest shopkeeper, for example, acts to avoid ruining his business by sullyng his reputation. He does not aim to maximize profit but, rather, to stay in business. Indeed, shirking some of his customers on change would be more likely to maximize the shopkeeper's profit, at least in the short term.

In this sense, an investment firm, reasoning prudently, would take an interest in the probabilities and utilities discussed above. Those calculations can help the firm to promulgate a prudent investment plan. The expected value analysis recommended, for example, that the firm invest in abstruse financial instruments when the value of profitable investments multiplied by the likelihood of achieving those profits was greater than the value of unsuccessful investments multiplied by the probability that the investments would lead to losses.

From a prudential perspective, a wider range of numbers could be interesting. A risk-averse investment firm, for example, might prefer not to invest in abstruse financial instruments when there was only a slight likelihood that the investments would produce profits. The risk-averse firm might choose to invest in such financial instruments only when they were very likely to produce profits. Similarly, a profit-hungry, risk-loving investment firm might choose to make the risky, complicated investments even when there was some likelihood that the investments would not produce profits. Such an investment firm would pursue the investments because of their inherent risk—assuming, of course, that the firm had a theory about why the investments were likely to succeed despite that risk—and, as such, would not be deterred by (someone else's calculation of) a high probability of loss.

Furthermore, evaluating the investments from the standpoint of prudence rather than expected value encourages the firm to take a wider perspective on the risky investments. As noted above in the discussion of expected value calculations, it is somewhat arbitrary which possible consequences the expected value theorist selects to evaluate. Because the prudent investment firm makes decisions that aim to secure its long-term prosperity—and stave off disaster—it is more motivated to consider consequences that have a lower probability but a larger downside.

The prudent investment firm would not necessarily include possible consequence C4—that the investments produce catastrophic losses for the economy as a whole—in its decision-making process. Because the prudent firm is (like the honest shopkeeper) self-interested, the fate of the overall economy enters its considerations only insofar as this possibility would produce catastrophic losses for the firm itself. However, the prudent investment firm is more motivated to include the possibility of financial disaster in its considerations: for the reason that financial disasters deplete the firm's money along with everyone else's.

Moreover, the prudent investment firm would worry more about the challenges to probabilistic calculation than the decision maker who calculates expected values without the further aim of making prudent decisions. As the prudent investment firm wishes to avoid catastrophe, it would be more likely to include the additional probabilistic calculations, discussed in section 1.2.1, above, that help to mitigate the weaknesses in expected value theory discussed in that section.

Perhaps most importantly, the prudent investment firm would modify the expected value consequentialist calculations discussed above by demarcating its risk tolerance—how much money it is willing to put at risk and at how much risk—and holding itself to this limit. This is not to say that the prudent investment firm can avoid all of the problems with consequentialist, expected value thinking, as discussed above. Some of these problems are endemic to consequentialist thinking and cannot

be avoided, as discussed further in section 3.3, below. However, I have discussed several ways in which a prudent frame of mind would help investment firms to use consequentialist, expected value thinking in a less problematic way.

For an example of prudent reasoning, consider Goldman's investments in the synthetic CDO discussed above, Abacus 2004-1. As the bank underwriting \$1.8 billion of the loans associated with Abacus, Goldman was an unfunded long investor. Perhaps due to the abstruse nature of these investments, including difficulties in making reliable predictions about the likelihood that borrowers would repay those loans, Goldman also took a short position on the synthetic CDO, betting \$2 billion that the loans would fail. These actions helped Goldman to assure that its long-term outcome would be satisfactory—it would stay in business and avoid catastrophe—to a greater extent than other investors in the abstruse Abacus 2004-1.

3.2 Is the bailed-out investment prudent?

In section one, above, I performed expected value calculations to evaluate whether an investment firm should rely on an external agency to bail out its investments if they become unprofitable. This analysis determined (for a certain set of assumptions) that investment firms should rely on bailouts whenever there is almost any probability that an external agency will bail out the firm if its investments became unprofitable.

Examining this problem from a prudential perspective—thinking consequentially about the long-term viability of an investment plan, with an eye to assuring sustainability and avoiding disaster rather than seeking to maximize present value, merely—brings up several concerns. As in the abstruse investment firm's prudential calculations, the bailed-out firm would worry that the lack of past evidence of bailouts makes it difficult to determine whether it will be considered too important to fail and be rescued from insolvency. The established practice and track record of TBTF, as noted above, suggests that it would not have been beyond reason for a large investment firm to presume that it would be rescued. But determining a probability with any degree of precision based on such scant evidence is very difficult. As such, most investment firms (reasoning prudently) would not count on external agencies to rescue them if their investments prove unsuccessful.¹⁴

Fannie Mae CEO Mudd's decision, discussed above, to take on additional, very risky loans—while owning or guaranteeing \$5.3 trillion in mortgages with only \$100 billion in real assets—might seem somewhat different from this estimation, however. At the time of that decision, it seemed close to certain that the U.S. government would bail out Fannie Mae if its investments failed. In this sense, Mudd's decision to buy \$150 billion of bad loans, or loans that were very likely bad, might seem prudent. Although this point is somewhat subtle, it seems to me that Mudd's decision was not prudent, in the sense of aiming to preserve long-term sustainability. As discussed above, the decision was probably favored by expected value theory and, as a matter of historical fact, did not cause Fannie Mae to fail (because the GSE was subsequently bailed out). The decision to rely on a bailout while pursuing ruinous practices seems imprudent, though, for the reason that the ruinous practices do in fact threaten—even if they do not destroy—the long-term profitability of the firm.

3.3 Problems With Prudential Reasoning About Risk

In this section, I consider how well prudential decision-making can address the weaknesses in expected value theory that I discussed above, concerning: (3.3.1) judgment, (3.3.2) paralysis (free-floating fear), and (3.3.3) distribution. I argue that prudential reasoning is adept at addressing some of these worries but that certain problems persist.

3.3.1 Judgment

Whereas expected value calculations rely on judgment to determine which possible consequences will be evaluated and which sources of probability estimates are the most reliable, the prudent decision maker employs judgment to question those probability estimates in depth. Using judgment may, in this sense, help the prudent investment firm to reach a better decision about whether to make an abstruse investment or rely on a bailout. But the reliance on judgment also renders the decision more personal, less objective, and more controversial.

In response, a prudent decision maker might note that judgment is what gives prudent decision making its decisive advantage over expected value calculations. Because decision makers can refer to their own experiences and expertise to determine what course of action is most likely to secure their long-term interests, they are more likely to achieve their personal goals. In this sense, the worry helps to establish the nature of prudent reasoning: it *is* more subjective and personal than expected value maximization. Although this challenges prudential reasoning as an objective decision-making strategy—one that will offer recommendations to everyone about what it is best for them to do—it strengthens prudential reasoning as a decision-making strategy for personal decision making that seems well suited to help decision makers reach their distinctive goals.

3.3.2 Distribution

In the discussion of problems with expected value theory as regards risk, above, I noted that expected value reasoning does not consider the distribution of risks of harm in its calculation. This means that it could recommend actions that make a small number of people extremely wealthy while imposing losses and risks on a much larger group. Prudential reasoning seems to have a similar problem.

Moreover, prudential reasoning, like expected value theory, is not directly concerned with consequence C4, in which the economic system is undermined. Prudential reasoning does do more than expected value theory, though, to avoid C4, in the sense that all prudent decision makers explicitly attempt to secure their long-term prosperity and C4 undermines everyone's long-term prosperity, as discussed above.

3.3.3 Paralysis (Free-floating Fear)

A third worry is that the prudent investment firm—which is deeply concerned with the problems in consequentialist thinking—will be paralyzed by these worries and will be unable to continue investing. This worry is related to, though distinct from, the concern about free-floating fear discussed above. As the prudent investment firm

fully recognizes how the problems in consequentialist thinking challenge human abilities to determine the probabilities that various outcomes will occur or even what outcomes are possible, the prudent firm recognizes that any investment plan is deeply fallible.

The prudent investment firm might resolve simply to err on the side of caution when performing prudential analyses—given that so much is uncertain—without yet halting action. In this sense, prudential reasoning offers a strategy for addressing the worries about free-floating fear discussed above: it is imprudent to be immobilized by fear about risk. A prudent decision maker would strive to move beyond this fear. This point is again subtle, though, in the sense that prudential reasoning offers no firm guidelines about how to address fear or paralysis. As such, even to the extent that prudential thinking can mitigate some of the problems associated with consequentialist decision making, as discussed in section one, it does not eliminate them.

4. IMMORALITY IN A CONTEXT OF RISK

In this section, I consider whether the (4.1) abstruse and (4.2) bailed-out investments are immoral independent both of their prudence or imprudence and of their expected value. I use Kantian moral reasoning to evaluate whether risking money in these investments is moral, i.e., whether in so investing firms make themselves exceptions to moral norms upon which the investments simultaneously rely. I aim to ascertain whether adding moral evaluations to decisions about risk can help to avoid the problems with expected value- and prudential reasoning discussed above. A further aim is to explore the extent to which Kantian moral reasoning—with its non-consequentialist orientation on the plans for acting rather than the actions' consequences—can address questions about risk, which unavoidably involve the consequences of actions. I conclude by considering several (4.3) objections to Kantianism.

4.1 Is the abstruse investment immoral?

As noted above, Kantian moral decision-making analyzes decisions in terms of the rules underlying the proposed actions, seeking to determine whether those rules can be universally followed. For the abstruse investment decision, I propose to test the following rule, R1: “When I believe that it will maximize profit, I invest in abstruse financial instruments.” This could have been the rule underlying GSC’s decision to take a long unfunded position in the Abacus 2004-1 synthetic CDO. If R1 were universalized then in every opportunity in which an investor could bet on an abstruse financial instrument, the investor would invest if it thought—e.g., because of expected value calculations—that the investment would maximize profits.

To evaluate this rule from a non-consequentialist, Kantian perspective, the decision maker considers whether the proposed action requires that other actors take different actions. Could all firms invest in abstruse financial instruments? Or do the abstruse firm’s risky investments depend on a social norm in which most firms take a different action—namely, refraining from abstruse investments—in order for it to be possible for the abstruse firm to make its risky investments in the first place?

If all firms invest in financial instruments whose risk levels and implications are difficult to ascertain due to their abstruse nature, it seems plausible that there would be many more failed investment strategies and many more failed firms. Those bad *consequences*, though, do not establish that the rule is immoral from a Kantian perspective. To do that, decision makers would need to show that the action of making abstruse investments requires other firms to act less rashly and take more precautions in their investing activities: perhaps by making sure that they grasp the risk and implications of every investment they make, perhaps by hedging against abstruse bets, as Goldman did with respect to Abacus 2004-1, discussed above.¹⁵

Indeed, a firm can only recklessly bet on financial instruments whose potentially catastrophic implications for the economy as a whole¹⁶ are opaque to the firm if other firms follow a more conservative investment strategy, on the following rationale. If all firms take abstruse risks, some investments (and some firms) will fail, while others will succeed. Because the investments are interconnected, however—firm A depends on premiums paid by firm B in order to pay the premiums it owes to firm C—even firms that win their abstruse-investment bets will lose money when R1 is a universal law. The significance of this bad consequence from a non-consequentialist, Kantian perspective is that it causes the (deeply interconnected) financial system, in which investment bets are placed, to break down. Winning firms have counted on losing firms for loan repayments. Because losing firms are now insolvent, however, they will not honor their obligations. Thus, firms that have money (i.e., winning firms) will be unwilling to loan, as those who need money (i.e., losing firms) are the same as those who have just failed to meet repayment obligations. Trust will evaporate. No one will be able to borrow money and the financial system, which depends on a ready supply of cash to pursue investment opportunities, will be at a standstill.

Abstruse investing requires, in this sense, some firms to adopt a more informed, risk-averse approach to investing, so that those firms will be there to provide trust—and credit—when abstruse investments fail. Such risk-averse firms refrain from relying on abstruse-investing firms and, as such, are not stung when (some of) the abstruse investments fail. The risk-averse firms will still be willing to make loans. In order to seek high profits, then, abstruse investors make themselves exceptions to the moral standard to which they hold the firms with which they interact. If everyone invested with the willful abandon of the abstruse investor, no one could achieve their aim of maximizing profit through investing because trust would break down and investment activities would cease. Abstruse investing becomes self-defeating when universalized.

4.2 *Is the bailed-out investment immoral?*

With respect to the second case, the bailed-out investment—deciding whether to rely on a bailout from an external agency when risking money—I propose the following rule (R2): “When I believe that I have a financial rescuer, I make riskier investments than I would if I did not believe that I had a rescuer.”

If this rule were universalized, then whenever an investment firm believes it has a guarantor to underwrite its financial responsibilities, it will take risks that it otherwise would not take. As in the case of the abstruse investment, we consider

the outcomes of these risky actions without allowing the outcomes to determine the permissibility of the risky action. In some cases, the risks will allow the firm to earn greater profits than it would have absent the risks. In other cases, the additional risks will not prove profitable but will not prove catastrophic: the investment firm will be able to meet its financial obligations. In the final kind of case, the additional risks will prove catastrophic and the investment firm will be forced to depend upon its presumed rescuer. As noted in section one, above, there are two possibilities: the rescuer will rescue the firm or the rescuer will be unable, or unwilling, to rescue the firm. If the rescuer is unable or unwilling to rescue the investment firm then the firm will fail to meet its financial obligations.

In this scenario, the integrity of the economic system would be straightforwardly undermined if the firm's insolvency cascades—the firms to which it owes money become unable to pay the firms to which they owe money and, ultimately, all firms become unable to invest—like that of the abstruse investment firm described in section 4.1, above. R2, like R1, would refute itself when universalized. Because I have not assumed in the case of the bailed-out investment that financial obligations are deeply interconnected, however, the self-refutation cannot be obtained this way. Two alternative analyses, consistent with my assumptions, each secure the self-refutation.

First, in an economic system, whenever someone loses money in an investment—say, Fannie Mae or GSC partners—someone else makes money: such as the commercial banks that sold Fannie mortgages or Goldman Sachs. Even following transactions that cause some parties to become insolvent, then, the economic system may appear intact: some parties have lost but others have won. The problem here is that in order for those who have lost to be rescued, the rescuer—typically, the federal government—must obtain rescue money. Governments obtain money by taxing their citizens. But the only citizens who have money in this economy are the ones who have earned profits. Thus, the “winners” must bail out the “losers.” If this were the normal course of investing, though, investing would be impossible: if all parties ended up with the same resources, no one would risk money in the first place. R2 is self-refuting when universalized.

Second, in actual economic system like the U.S. economy, the U.S. government is the guarantor of last resort. But the U.S. government would not be unable, or even unwilling, to bail out a firm when that firm's insolvency would undermine the U.S. economy as a whole. Widespread bailouts would, however, change the nature of the economy: in particular, the economic system would cease being based in private property, directed by individual self-interest, and conducive to investments. The government would own key economic players; the economy would be centralized.¹⁷ As people in centralized economies lack private property to invest, making investments (that riskily rely on rescuers, or otherwise) in the context of such economies is impossible. Again, R2 refutes itself when universalized.

4.3 Problems With Kantian Moral Reasoning About Risk

In this section, I consider Kantian moral reasoning's resources to address the weaknesses in expected value theory and prudential reasoning that I discussed above, including: (4.3.1) distribution, (4.3.2) paralysis (free-floating fear), and

(4.3.3) judgment. Like prudent reasoning, Kantian moral reasoning is adept at addressing some of these worries but certain problems remain irresolvable.

4.3.1 Distribution

Like expected value theory and prudential reasoning, Kantian moral theory says nothing explicit about the distributional consequences of its recommendations. Uniquely among these three decision-making strategies, however, Kantianism explicitly considers the decision-making strategies available to *each person* if the rule under examination were followed universally. Moreover, unlike expected value theory and prudential reasoning, Kantianism considers consequence C4, even though Kantianism is the theory, among the three, that is least concerned with the consequences of action.

4.3.2 Paralysis (Free-floating Fear)

Kantianism is also the theory, of the three, most likely to take an interest in paralysis, or free-floating fear. Consider R3, “When I believe that it will maximize profits, I make risky investments that create free-floating fear,” and R4, “When my prudential reasoning is unable to vanquish all uncertainty, I am paralyzed by uncertainty and refrain from investing.” To the extent that the risky investments would produce paralysis and free-floating fear, and to the extent that paralysis and free-floating fear would undermine the economic system, Kantian morality would prohibit making investments that create free-floating fear or undertaking prudential reasoning that produces paralysis.

4.3.3 Judgment

Expected value calculations relied on judgment to determine which possible consequences will be evaluated and which sources of probability estimates are the most reliable. Prudent decision makers used judgment to question the expected value probability estimates in greater depth. Kantian moral reasoning relies on judgment to express the rule to be evaluated and to consider the rule’s universalization. (It also relies on probabilistic judgments to determine that an action under consideration is risky, as discussed in section 1.2.1, above.) Guidelines for writing and universalizing the rule were offered above. Although such guidelines limit the extent to which personal biases can distort the moral evaluation, it seems indisputable that decision makers must employ (subjective, personal) judgment in writing and evaluating the rules, at least to some extent.

A further concern about judgment arises in determining how risky is a particular investment. Morally speaking, investors may not risk money in abstruse or bailed-out investments any time the investments risk C4. It is immoral for firms to make complicated investments when those investments risk the integrity of the economic system itself for the reason that, in those circumstances, the investment firms except themselves from a moral rule that they regard as being generally binding. If the investments do not involve systemic risk, then the above analysis does not explicitly prohibit them. (And, as broached above, some risks are highly beneficial.)

So it is especially important to ascertain when investments risk economic ruin. Its risk profile is one of the things that is abstruse about abstruse investments, however, so it might be difficult to establish that a particular abstruse investment is too risky and is, thus, immoral by Kantian lights. Given the difficulties in determining whether an abstruse investment risks consequence C4—and thus whether the rule underlying that investment decision is morally permissible—risk-involving rules seem to require enhanced scrupulousness from the standpoint of Kantian morality.

There is a further, related worry about judgment in Kantian moral theorizing about risk that is important to address, though I will not be able to do so fully here. Just as it could be morally permissible to risk money in abstruse and bailed-out investments when the overall economy is not also risked, it could be morally permissible to risk undermining the economy so long as the risk is trivial.¹⁸ In this sense, the extent of the risk is not all that matters to our calculation, but also how likely it is that the harm threatened by the risk will actually occur. If the probability of a catastrophe is truly trivial, investments that risk this outcome could be morally permissible.

In response to this worry, I develop the conclusions broached above and propose that investment strategy rules that (a) risk the economic system itself (b) above probability P are prohibited by the Kantian morality; investment strategies that risk recoverable loss—or irrecoverable loss below probability P—are not (on that ground) morally impermissible. Although different levels of risk tolerance are morally permissible—and might be prudent—for various firms, we need to know whether the worst-case scenario is threatened with some probability P. When the worst-case scenario is threatened with probability P, the risk is morally prohibited.

5. CONCLUSION

A key insight of Kant's moral law is that a firm may never act in a way that would destroy, if universalized, what the firm needs in order to act. Imprudence focuses only on the firm's own success or failure; morality concerns conditions for the possibilities of success or failure more broadly construed. In this conclusion to the paper, I sketch some guidelines as regards that limit on permissible action and discuss the relationship between immorality and imprudence that I have developed in this paper.

The abstruse investment puts money into complicated financial instruments that are interconnected in intricate and hard to parse ways, i.e. that are very risky. The intricate nature of the instruments—and the fact that their particular risks cannot be formulated in detail—prevents the abstruse investment firm from having an adequate conception of the risk it undertakes in buying and selling these instruments. Expected value maximization has a difficult time grappling with this problem, due to the difficulties attendant upon formulating the probabilities correctly. It can make general recommendations about when the investments are to be preferred but given the difficulties of making precise calculations, it offers little certainty as to the best course of action.

In order to address this problem, I turned to Kant's distinction between prudence and morality. According to Kant, prudence and morality sometimes intersect but we can conceptually distinguish choosing the plan for action that best furthers long-term

self-interest from choosing a moral plan for action. I used the idea of prudence to interpret the expected value calculations in a way that allowed them to gain more certainty in the context of risk. My prudential evaluation relied on the consequentialist expected value reasoning but was not identical to that reasoning. One addition that prudential thinking brings to the consequentialist calculation is the idea of a risk tolerance. Although the prudent investment firm will not be able to determine with certainty to how much risk the abstruse investments exposes it, the firm can decide upon a risk tolerance as part of its prudential calculations to help decide when it is prudent, according to the firm's lights, to make the investments and when it is not.

Next, I examined the investments from the point of view of Kantian morality, seeking to understand if the abstruse investment firm, in trading the complicated financial instruments, makes itself an exception to a rule to which it holds other firms, e.g. by relying on them to make less risky trades. My examination was notable for the way in which it included the consequences of an action while evaluating the action in a non-consequentialist manner. I determined that it is immoral for the abstruse investment firm to trade complicated financial instruments when such investments risk the destruction of the entire economy.

Regarding the second case, the bailed-out investment, I questioned when it would maximize expected value to rely on a bailout if one's investments fail, when it would be imprudent to do so, and when it would be immoral to do so. Expected value calculations recommend that investment firms rely on a bailout when (assuming a 50% chance of investment success and that the size of probable profits and probable losses is identical) there is a greater than 0% probability that the external agency will bail out the investment firm if its investments fail.

Turning to prudential calculations to try to achieve more helpful guidance (especially in light of uncertainties about probabilities and other aspects of consequentialist reasoning), I determined that firms, thinking prudently, should assume that it is highly improbable that an external agency will bail them out when their investments prove unsuccessful. Prudent firms should rely on outside rescuers only when it is possible to minimize the possible adverse consequences of doing so.

When I examined moral issues concerning the bailed-out firms' decisions, I resolved that a moral firm might rely on an external rescuer only when doing so would not risk an economic catastrophe. As in the prudent firm's evaluations, the circumstances in which a moral firm may permissibly rely on a bailout will be highly circumscribed.

The primary aim of this paper has been to motivate interest in non-consequentialist accounts of the ethics of risk and describe one such account, based in Kantian morality. One interesting upshot of this work concerns the different levels of risk aversion associated with expected value-, prudential-, and moral reasoning. Although I have not fully established the relationships among levels of risk aversion in these various forms of reasoning, my paper suggests that the requirements of (2) prudence vis-a-vis risk aversion seem more demanding than the requirements of (1) expected value maximization and the requirements of (3) morality seem more demanding than (2) prudence as regards the regulation of risk. Expected value permits any investment that is likely to maximize profit. Of these, prudence prohibits investments that threaten the investor's long-term interests. Of prudent investments,

the moral law prohibits those that risk of the integrity of the economic system (on the grounds that, if universalized, such risks would undermine the economy, making any further investments impossible). In this sense, the Kantian moral standard for investments would impose more risk aversion on the financial system than what currently exists, helping both to address the theoretical problems with consequential reasoning addressed in this paper and, more perhaps more crucially, the crises to which they lead.

ACKNOWLEDGEMENTS

I am grateful to audiences at the 2014 Annual Meeting of the Society for Business Ethics, High Point University, and Bloomsburg University, for useful feedback on previous versions of this paper. Matthew Brophy and Jeffrey Moriarty provided invaluable written comments. Conversations about earlier drafts with Waheed Hussain, Michael Santoro, Eric Orts, Amy Lynn MacArthur, Joe Blosser, Robert Moses, Scott Lowe, Steven Hales, Kurt Smith, Gary Hardcastle, Wendy Lynne Lee, Richard Brook, and Michael Simonetti were especially helpful. I also wish to acknowledge Alan Strudler and two anonymous reviewers at *Business Ethics Quarterly* for extremely valuable comments on earlier drafts of this paper.

NOTES

1. No one has yet defended an approach to questions about risk based in Kant's formula of universal law, as I do in this paper. The paper does build on a number of criticisms of consequentialist decision-making frameworks as regards ethical question about risk. In the interest of motivating interest in and introducing the Kantian view, I must unfortunately forgo detailed responses to much of this excellent work. Oberdiek has objected that consequentialist frameworks objectionably trade "lives for convenience" and societies should seek a less objectionable way of making decisions that affect large numbers of people (2004, 201). Hansson raises two crucial objections: (1) consequentialist decision-making frameworks omit issues that are relevant to ethical evaluation, such as the risk-taker's intention in imposing a risk on someone else (2010, 587), and (2) consequentialist decision-making frameworks unrealistically assume that actors are certain of the consequences of their various possible actions, along with the likelihood that those actions will come about (2010, 587-88). Hayenhjelm and Wolff also object that consequentialist decision-making frameworks omit morally relevant aspects of the decision: in particular, the distribution of the risk of harm across the population that is put at risk (2012, 11-20). Oberdiek, Hansson, Hayenhjelm and Wolff are interested in contractualist solutions to ethical questions about risk in the Scanlonian tradition: risks are ethically permissible only when everyone whom they expose to a risk of harm would consider the risk worth running.

2. As discussed in Hansson, the appropriate form of risk regulation may vary among different kinds of risk. Risks arise in different social sectors and regulation should consider the particular aspects of society—workplace organization, community planning, the provision of basic goods—that most directly relate to the risk (2005, 12). Financial risk, as elaborated below, addresses concerns about the organization and regulation of the economic system.

3. Throughout this paper, I assume that the investment firm is a rational actor—which forms intentions and acts on them—as discussed in French (1979). The actual agents are usually corporate executives, such as the CEO and the firm's risk advisory board, who act on behalf of the firm's shareholders. Where relevant, I discuss how individual employees contribute to the firm's risky investing activities. Various challenges can be raised against French's view. The most influential is Jackall (1988, 2009), who argues that corporate organization is decentralized and bureaucratic; thus, the firm (qua firm) cannot act rationally. Another important worry is Altman (2007), who argues that businesses are incapable of the kind of agency Kant requires for ethical evaluation. I believe that such worries can be adequately addressed, though I do not attempt to do so here.

4. This is, of course, not the use to which Kant puts his distinction. My paper appropriates for its own purposes, rather than interpreting in the context of Kantian ethics, the distinction between acting for reasons of prudence and acting for reasons of morality. For a discussion of how Kant intends the distinction in his own exegesis, see Wood (1999, 27–33).

5. This aspect of my argument—if standards of Kantian morality (as regards risk aversion) are met then standards of prudence (as regards risk aversion) are met—may be formulated logically as: $M(r) \rightarrow P(r)$. Note that this counterfactual has the standard logical entailments. In particular, I do not seek to challenge the (sadly) well-established fact that many immoral actions can be prudent.

6. I address a rudimentary form of expected value theory here. More sophisticated forms may include additional analyses. For example, in my discussion of an objection to expected value theory (section 1.2.1, below), I include an estimate of decision makers' probability of correctness regarding the probability estimates they employ.

7. In this context, abstruse can be seen as a relative term. For some very well informed people, such as very knowledgeable research scientists, certain highly complicated investments are much less abstruse. My example focuses on investments that are truly abstruse: either in the sense that they are abstruse even to the very well informed or that they are abstruse because decision makers are not very well informed.

8. The U.S. government has rescued other financial institutions—including hundreds of savings and loan (S&L) associations in the S&L crisis of the 1980s–90s—and other businesses, including Lockheed Aircraft and the Chrysler Corporation. These institutions were not bailed out because of concerns about systemic risk, however.

9. As broached above in note 4, most Kant scholars hold that Kant's aim in discussing this example is to isolate a key point in his moral philosophy: the phenomenon of acting from duty rather than inclination. My discussion of this example does not challenge that interpretation of the passage. Rather, I seek to use Kant's distinction to demonstrate his view's resources to evaluate questions about financial risk.

10. Although no one has examined this case in the context of the ethics of financial risk, the example has been extensively discussed. Wood raises the problem of determining in which proportion the 23 two motives—morality and prudence—influenced the shopkeeper's action. Kant believes that the answer to this question is unknowable (1989, 472–73). Thomas examines how the example relies upon certain knowledge about human life (1993, 16). Herman (1981, 366), Smith (1991, 288), and Hills (2009, 114–15) examine what would need to be the case (counterfactually) to make this example a matter of morality rather than prudence. McCarthy locates the wrong in the shopkeeper's action in the fact that in a possible world where the shopkeeper will no longer interact with the youthful customer or his affiliates, the shopkeeper's rule commits him to giving the incorrect change and thus acting immorally (2002, 637). Of these, the interpretation offered in Herman (1981) is most crucial to the present work.

11. E.g., Rawls (2009, 168), Korsgaard (1996, 57–58), Scharding (forthcoming).

12. As is well known, Kant gives three formulations of the moral law as he understands it: the formula of universal law (FUL), the formula of humanity (FH), and the formula of the realm of ends (FRE). My analysis focuses on FUL, only. For purposes of convenience, I refer to FUL as Kant's decision making strategy for morality, setting aside worries that FUL expresses the moral law only incompletely. For a discussion of those worries, see (Wood 1999).

13. This idea about prudence resembles Rawls's maximin. Whereas maximin nudges decision makers towards risk aversion by requiring them to choose whatever course of action has the *least-bad* worst outcome—in a distributional sense, as discussed below—my account of prudence avoids bad outcomes more generally and in a more self-interested sense. Maximin would reject an outcome in which the bottom quintile of the population was very badly affected while decision makers—such as investment firms engaging in brokerage activities—flourished. My simpler account of prudence could endorse that decision, prohibiting very bad outcomes for the bottom quintile of the population only when such bad outcomes make it impossible for the decision makers to continue their brokerage activities in the future. For an insightful account of maximin's role in risk aversion for Rawls, see (Freeman 2014).

14. This is not to say that relying on bailouts is not often beneficial. For example, bankruptcy laws allow people to take risks and then start again if the risks have bad consequences. Social welfare, like health and pension guarantees, allows people to take more physical risks. Even the federal funding of research is a kind of bailout: providing a safety net for researchers that encourages innovation and creativity. My point is that it is (typically) imprudent to rely on these benefits.

15. Most firms—not just Goldman—typically hedge by purchasing short and long positions in a pattern that is likely to allow them to meet their financial obligations even if some investments are unsuccessful. Goldman's strategy with respect to Abacus 2004–1 was more risk averse than these typical hedges, though, in the sense that Goldman bet both for and against the same synthetic CDO.

16. Only abstruse investments that risk consequence C4 seem self-refuting. Consider, for example, an abstruse investment that does not risk C4, such as when someone makes a \$10 bet on a horse race without understanding horse racing. This investment is “abstruse” to the investor in the relative sense discussed in note 7, above. This investor is likely to lose \$10. While it would be rejected by expected value theory and prudence, the investment seems permissible under Kantian moral reasoning. It would be possible for everyone to invest in this manner; the bets do not threaten to undermine the institution of horse racing.

17. This is not to say that the economy becomes state run whenever there is a bailout. Certain, more limited forms of bailouts—such as bankruptcy protection or limited liability for people who invest in publicly-traded firms—could be available universally without causing the economy to become state-run.

18. Thomson (1986, 185–86) sketches this point, though she does not fully establish it. Thomson is not working in a Kantian framework when she discusses the possible moral permissibility of trivial risks of serious harm.

REFERENCES

- Altman, Matthew. 2007. “The Decomposition of the Corporate Body: What Kant Cannot Contribute to Business Ethics.” *Journal of Business Ethics* 74: 253–266.
- Freeman, Samuel. “Original Position.” In *The Stanford Encyclopedia of Philosophy*, edited by E. N. Zalta. <http://plato.stanford.edu/archives/fall2014/entries/original-position/>
- The Financial Crisis Inquiry Commission (FCIC). *The Financial Crisis Inquiry Report: Final Report of the National Commission on the Causes of the Financial and Economic Crisis in the United States*. Pursuant to Public Law 111–21. Date issued: February 25, 2011. <http://www.gpo.gov/fdsys/pkg/GPO-FCIC/pdf/GPO-FCIC.pdf>
- French, Peter A. 1979. “The Corporation as a Moral Person.” *American Philosophical Quarterly* 16: 207–215.
- Hansson, Sven Ove. 2003. “Ethical Criteria of Risk Acceptance.” *Erkenntnis* 59: 291–309.
- . 2005. “Seven Myths of Risk.” *Risk Management* 7: 7–17.
- . 2010. “The Harmful Influence of Decision Theory on Ethics.” *Ethical Theory and Moral Practice* 13: 585–593.
- Hayenhjelm, Madeleine and Wolff, Jonathan. 2012. “The Moral Problem of Risk Impositions.” *European Journal of Philosophy* 20: E1–E142.
- Herman, Barbara. 1981. “On the Value of Acting from the Motive of Duty.” *The Philosophical Review* 90: 359–382.
- Hills, Alison. 2009. “Moral Testimony and Moral Epistemology.” *Ethics* 120: 94–127.
- Jackall, Robert. 2009. *Moral Mazes: The World of Corporate Managers*. New York: Oxford University Press, 1988.
- Kant, Immanuel. 2002. *Groundwork for the Metaphysics of Morals*. Translated by Allen W. Wood. New Haven: Yale University Press.
- Korsgaard, Christine. 1996. “Kant’s Formula of Universal Law.” *Pacific Philosophical Quarterly* 66: 24–47.
- McCarthy, David. 2002. “Intending Harm, Forseeing Harm, and Failures of the Will.” *Noûs* 36: 622–642.
- Oberdiek, John. 2004. “The Ethics in Risk Regulation: Towards a Contractualist Re-Oriented.” *Rutgers Law Journal* 36: 199–204.
- Noonan, Peggy. 2009. “There’s No Pill for This Kind of Depression: Six Months After the Collapse, A ‘Pandemic of Fear’.” *The Wall Street Journal*, March 13, A9.
- Nozick, Robert. 1974. *Anarchy, State, and Utopia*. New York: Basic Books.

- O'Neill, Onora. 2003. "Autonomy: The Emperor's New Clothes." *Proceedings of the Aristotelian Society* 77: 1–21.
- Perry, Stephen R. 2007. "Risk, Harm, Interests, and Rights" in Lewens, Ted, ed., *Risk: Philosophical Perspectives*. New York: Routledge.
- Rawls, John. 2009. *Lectures on the History of Moral Philosophy*. Cambridge: Harvard University Press.
- Scharding, Tobey K. "Crafting Maxims." *Teaching Ethics* (forthcoming).
- Smith, Holly M. 1991. "Varieties of Moral Worth and Moral Credit." *Ethics* 101: 279–303.
- Thomas, Lawrence. 1993. "The Reality of the Moral Self." *The Monist* 76: 3–21.
- Thomson, Judith J. 1986. "Imposing Risks" in Parent, William, ed. *Rights, Restitution, and Risk: Essays in Moral Theory*. Cambridge: Harvard University Press.
- Williams, Mark T. 2010. *Uncontrolled Risk: The Lessons of Lehman Brothers and How Systemic Risk Can Still Bring Down the World Financial System*. New York: McGraw-Hill.
- Wood, Allen W. 1989. "The Emptiness of the Moral Will." *The Monist* 72: 454–483.
- . 1999. *Kant's Ethical Thought*. New York: Cambridge University Press.