Corporate Social Responsibility as Reputation Insurance: Theory and Evidence*

Dylan B. Minor
Haas School of Business
UC Berkeley

July 2010

Abstract

We posit Corporate Social Responsibility (CSR) is being used by firms as a reputation insurance mechanism, helping them better withstand the tumult of negative business shocks. Guided by our theory, we empirically test CSR as insurance in the setting of product markets. We find higher type firms experience events significantly less often than low type firms. In addition, higher type firms do indeed experience a lesser change in firm value: we find higher CSR types enjoy $600 million of saved firm value after a product recall compared with low types.

*The author would like to thank Daniel Diermeier, Bryan Hong, John Morgan, Orie Shelef, Steve Tadelis, and Terry Taylor for their helpful comments. Also thanked are the seminar participants at UC Berkeley, University of Michigan, and Northwestern University.

"CSR is best seen as the management of risk, as the avoidance of damages to the company's reputation." *Financial Times*, July 7, 2004.

There have been a plethora of past studies examining the relationship of a firm's financial performance with its level of corporate social responsibility (CSR). In short, the studies show there is little relation between the two (see Elfenbein (2007) for an extensive survey). Meanwhile, CSR seems to be increasingly important to firms. Indeed, a recent survey by the *Economist* magazine[1] reports some 56% of managers consider CSR as a "high" or "very high" priority. This compares with roughly 34% three years ago and an expected 69% three years hence. Further, they report 87% of firms now have a CSR firm program. Echoing firm sentiment, many MBA program ranking schemes now include a standalone category for CSR. Why are firms so concerned with CSR?

Recent work has offered various reasons from managers seeking "warm" glow from their CSR activities (e.g., Fisman et al. 2006) to firms simply appeasing the demands of NGOs to prevent boycott, or even forestalling looming governmental regulation (e.g., Baron & Diermeier 2007). Additionally, it has been suggested firms use CSR to signal a whole swath of different messages, most all revolving around the trustworthiness of the firm and its effect on *current* revenue (e.g., Goyal 2006), or even firms use CSR as a form of penance, offsetting the firm's past irresponsible behavior (e.g., Kotchen & Moon 2007).

However, *when managers are actually asked why they engage in CSR*, they claim it is to secure a better brand and reputation. However, only some 6.5% of managers report CSR increases revenue. What is the value of increasing (brand) reputation if in the end it does not increase revenue? Similarly, some 63% of managers say their adoption of sustainability practices either does not change or even decreases profit. How does this make business sense? *We posit a primary value of CSR is that of an insurance mechanism for the firm's value.* That is, investing in CSR can help build social reputation that softens the blow of future business shocks on a firm's value. The primary benefit comes *after* an event. Thus, CSR is engaged in not to increase a

[1] January 7th, 2008

firm's value *but to protect it.* This is much in the spirit of Hermalin (2008) that shows higher levels of corporate governance are a result of firms wanting to protect their profits as opposed to better governance yielding higher profits. However, whereas he is concerned with corporate governance protecting current profits, we are concerned with protecting the value of the firm through a *contingent* future benefit when facing shocks.

The mechanism by which CSR investment is expected to "payoff" during an event is at least twofold. First, there are NGOs that will make demands of firms in terms of CSR commitments. To the extent such firms meet these requests, it is expected they will receive limited wrath from these NGOs *after* a "bad" event occurs (see Vogel (2005))[2]. A second reason could be that when firms invest in CSR, they indicate their level of CSR related issue effort (e.g., environmental effort), helping improve investors and regulators' posteriors after a negative issue related event.

For example, if a firm overtly incurs extra expense to have superior environmental management systems, when it faces an environmental accident, such an *ex-ante* commitment could help tip the scale from the view of the investor and regulator that such an event was really due to bad luck instead of negligence. Hence, assuming negligence is more costly than noise, the firm's market value would be less punished and it would be less likely to be pursued by a regulator than had the firm not shown such commitment. For this paper, we examine the mechanism of this latter form.

For a very recent concrete example of this mechanism consider British Petroleum (hereafter, BP) and Johnson and Johnson (hereafter, J&J). Leading up to its recent devastating oil spill on April 20th, 2010, BP had begun developing a reputation of carelessness through outsourcing safety and cost cutting, providing few actions that showed a genuine care for environmental safety. Indeed, as publicly noted by the US Chemical safety board, BP was "cutting the costs for safety and maintenance to increase profits[3]." Consequently, although on April 20th, 2010 analysts best estimates of *total* spill costs were around \$3 to \$12 billion, the stock market wiped

[2]See also Baron and Diermeier (2007) for a further exploration of this possibility.

[3]See Wall Street Journal story "Drilling Down: A Troubled Legacy in Oil" May 1,2010.

out $32 billion of firm value[4]. Investors weighed the scale of negligence versus noise and it fell heavily to negligence.

In contrast, J&J faced a product recall of various over the counter children's medicines based on serious manufacturing problems. However, J&J had already begun hiring outside experts to help improve its quality control on the problem plant, as it detected a need to do so. Thus, although J&J will certainly have to pay for faulty manufacturing, J&J was showing it was taking substantive steps to improve safety and quality, reducing total expected long run firm cost. Hence, when a recall was triggered investors instead tipped the scaled more in the direction the event was noise over gross negligence[5], providing nominal change in firm value.

We see the contribution of this paper as twofold. To our knowledge, this is the first paper to formally[6] examine CSR as an insurance mechanism. Second, this is the first paper to empirically test the notion of CSR as reputation insurance. In short, we aim to open the black box of CSR decision making by examining CSR as reputation insurance, examining the efficacy of the primary reason given by managers for engaging in CSR.

Our paper is organized as follows. In the first section we present our model of CSR investment, exploring the mechanism of CSR benefit and identifying when it is expected to payoff. The next section provides our empirical analysis, which suggests CSR does generally provide a substantial insurance benefit. Our final section provides a concluding discussion.

1 CSR as Insurance

The central idea of CSR as an insurance mechanism is again that a firm first makes a CSR investment to obtain a higher CSR reputation. By CSR reputation

[4]See http://money.cnn.com story "BP loses $32 billion in value on spill."

[5]See Wall Street Journal story "J&J Lapses Are Cited in Drugs for Kids" on 5/27/2010.

[6]For an informal discussion of CSR as an insurance mechanism for firm value see Minor and Morgan (2010).

we mean the firm's reputation of how conscientiously they go about the production and selling of their goods, which in turn then creates a sense of how trustworthy the firm is.

For example, consider the realm of product markets. Product recalls can provide shocks to firm value and reputation. Now CSR in product markets can be thought of as a firm's superior reputation earned from the conscientious creation, marketing, and distribution of its products. A firm with high CSR would, for example, embrace superior quality assurance procedures in the development and production of its products, conduct ethical marketing campaigns, provide products with extra social value, provide products to disadvantaged demographic groups, and generally face product recalls voluntarily. Contrarily a bad CSR firm could often be involved in regulation fights, suffer safety violation fees, accept lower product safety standards, and conduct limited due diligence on their supply chain. One key distinguishing characteristic between high and low CSR activities is the latter will tend to be less costly in the short term, and particularly less costly if the bad event never occurs, which, for the current example, is a product recall. Since some of the low CSR activities are unobservable by the public, if the event never occurs, it is difficult to know the extent to which such activities occur. Nevertheless, to the extent these unobservable activities (e.g., limited supply chain due diligence) are correlated with other activities that are observable (e.g., regulatory fines), it is still possible to develop reasonable ex-ante reputations.

Now if a bad business shock should occur, the firms then with higher CSR reputations will not be punished as badly. This net result is similar to if we examined a group of workers with typical incomes of $50,000 and homeowner insurance premiums of $500. The insured and uninsured would only have a 1% difference in annual income, which would likely be difficult to distinguish statistically. However, upon an event (e.g., a house fire), there would be a very significant difference between an uninsured and insured in terms of cash flow via the insurance benefit. Hence, to find if the homeowners insurance acted like insurance, we would again have to make sure we capture negative events and not simply time series of ordinary cash flows.

Therefore, here the expected CSR insurance benefit, just as with homeowners insurance, is both a function of the probability of the event occurring and the expected net benefit received, conditional on the event occurring. Also, CSR is not expected to offer 100% insurance. Typically, with insurance there is an deductible and/ or coinsurance. Thus, a firm will still likely lose value during an adverse event, but they will less value than had they not been insured.

Therefore, in our formulation, it is not expected that CSR will ever payoff in a current revenue or profit sense, just as with traditional insurance. In fact, we conjecture a slight negative contemporaneous effect since CSR investment is costly, though there may be some offsetting of current value (e.g., high CSR might attract or retain higher quality employees). Moreover, if a higher CSR reputation lessens the chance of an adverse event (and thus its cost), this too does not show in a time series. This explains how the past literature could have had such difficulty establishing a statistical relationship between CSR and accounting financial performance: most firms for a typical time series experience no CSR benefit. To find the value of CSR, we must instead examine periods of unlikely events.

Now the precise benefit of CSR reputation during an event is the differential of lost firm value between a higher and lower CSR type firm. This differential is driven by how investors ex-post asses the probability the firm willfully caused the event versus it simply being due to bad luck, which will be determined by the firm's ex-ante CSR reputation. This probability assessment then affects how likely the event will be found to come from negligence. The firm's punishment is then increasing in the likelihood it will be found to be negligent, which is determined by an enforcer. We now turn to formally modelling this idea.

1.1 The Model

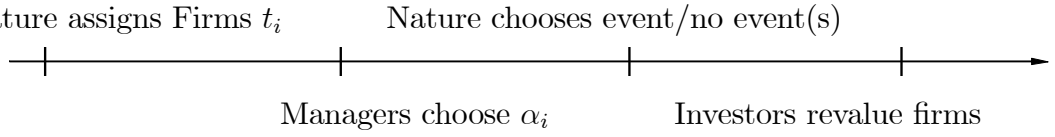

We being with a stylized model that will inform our empirical analysis. For our model, we begin by nature assigning firms a publicly *observable* type[7] $t_i \in \{H, L\}$, where H (L) type has high (low) CSR. Types determine a firm's probability of facing a negative event for a given level of negligence. Again, we can imagine when a firm experiences a bad event, it can simply come from an exogenous shock (i.e., an act of nature), or also from an endogenous act (i.e., a willful act of the firm). This is the quandary of a regulator and investors ex-post an event: how to balance the scale of an event's cause- did the event come about more as a function of willfulness or simply bad luck?

For our analysis, we endogenize this concept by having firms choose their level of negligence through an *unobservable* variable $\alpha \in [0, 1]$, which determines just how likely an event is caused by a willful act versus exogeneity. We can then think of the level of negligence α as a firm's propensity to commit negligence, as risks being taken that are more or less likely to result in negligence. In particular, we consider $\alpha \equiv \Pr(N = 1)$, the probability negligence has been committed. We then write the probability of an event E for firm i and type t as:

$$\begin{aligned} \Pr(E) &= \Pr(N = 1)\Pr(E|N = 1) + \Pr(N = 0)\Pr(E|N = 0) \\ \Pr(E) &\equiv P_i = \alpha_i \times \gamma_t + (1 - \alpha_i) \times \gamma_0 \end{aligned}$$

The variable $\gamma_t \equiv \Pr(E|N = 1)$ then determines how likely negligence causes an event, literally defined as the probability an event occurs given negligence has occurred. Similarly, $\gamma_0 \equiv \Pr(E|N = 0)$ is the likelihood an event is caused even when there is no negligence. Thus, at one extreme, if the firm is always negligent, the probability of an event occurring is simply γ_t. At the other extreme, if the firm is never negligent, its probability of an event is γ_0. Clearly, it should be the case $\gamma_t > \gamma_0$, which we assume throughout.

[7] See the appendix for some assumptions that yield heterogeneous firms spawning from homogeneous ones.

To capture the notion of CSR types we assume $\gamma_L > \gamma_H$. This means if both firm types choose the same level of negligence α_i, it is always more likely the low type firm will trigger an event. Practically, we can think of this as a high type firm's management process allowing a greater threshold before negligence triggers a public event (e.g., from more diligent monitoring of supply chains). We can also think of γ_t as being related to *observable* firm activities that affect the incidence of events; contrarily, α_i can be thought of as *unobservable* firm activities that lead to an event. In summary, we then have $\gamma_L > \gamma_H > \gamma_0$, which just says endogenous acts are more likely to result in an event versus an exogenous act, and a low type's endogenous act is more likely than the high type's to result in an event.

Per the model's time line above, a firm is first assigned its type. We abstract away from how firms arrive at being different types, but simply assume the existence of some of both types. It may be an amalgam of firm culture, competitive advantages, and simply path dependency. Whatever the case, we do provide some sufficient conditions to yield an equilibrium of different firm types beginning with homogeneous firms in the appendix. The manager of the firm next chooses the level of negligence α it will take on to maximize expected net profit. We also abstract away from any agency issues between firm and manager to focus on reputation concerns.

After managers choose α_i, nature then draws at the probability of P_i if firm i faces an event. If an event does occur, investors and regulators will be assessing how likely the event was brought about by negligence versus noise. The former is more costly to the the firm than the latter. For this model, we will have the cost of an event be simply a function of the chance negligence was committed given the firm is now facing an event. For a more complicated model, where we endogenous the firm target choice of an enforcer *ex-post* an event, as well as generalize to an arbitrary number of firms and CSR types, see Minor (2010). However, this fuller model adds little to our current analysis and so we omit it.

The probability of negligence *conditional* on an event occurring is then denoted η and is calculated via Baye's rule as:

$$\begin{aligned}\Pr(N &= 1|E) = \frac{\Pr(N=1)\Pr(E|N=1)}{\Pr(N=1)\Pr(E|N=1)+\Pr(N=0)\Pr(E|N=0)} \\ \Pr(N &= 1|E) \equiv \eta_i(t) = \frac{\alpha_i \times \gamma_t}{\alpha_i \times \gamma_t + (1-\alpha_i)\times\gamma_0}\end{aligned}$$

Of course, after an event, either a firm has or has not committed negligence, and knows it so. However, investors do not know which case it is, and thus must form a belief via Baye's rule of how to revalue a firm, as we assume the ultimate cost of the event (on average) will depend on whether it was due to negligence or not.

1.2 Firm Objective Functions

The firms have the following objective function:

$$\begin{aligned}\pi_i &= -(1-\alpha_i)^2 - I\cdot\alpha_i\cdot a - \Pr(E_i)\cdot\eta_t\cdot C \\ &= -(1-\alpha_i)^2 - I\cdot\alpha_i\cdot a - \alpha_i\cdot\gamma_t\cdot C\end{aligned}$$

That is, firm type $i \in \{L, H\}$ maximizes profit by choosing a level of negligence α_i. We see then profit is increasing in α_i in a concave fashion. This can be thought of as a firm increases its level of negligence, it also enjoys greater return via profit, but at a decreasing rate, and assuming it does not experience an event. That is, a little negligence (e.g., a bit of "cutting corners") can improve profits considerably- if an event never occurs. However, as negligence becomes too great it can start offsetting such benefit through, for example, excessively deteriorating production quality.

However, we can also interpret α as the firm's distance from its ideal point of creating an externality (i.e., $\alpha = 1$). If a firm was never to face an event, it would never have to internalize the cost of its externality, and thus it would then always impose the externality. However, this model has a *contingently* imposed cost of the externality.

The middle term captures the (net) cost of CSR, which must be paid by the high type. That is, a high type is more conscientious and thus less likely to encounter a public event given a level of negligence, though to be so is costly. I is an indicator function that has value one if the firm is a high type. As would be expected, we see the greater the negligence risk being taken, the greater the cost of insurance.

Our final term, the event cost, is then simply $\Pr(E)$, the probability of an event happening for firm i, times η_i, the probability of an event given negligence, times C, an event cost (for some $C > 0$) given negligence.

We make the assumption of *no arbitrage*, defined thus:

No Arbitrage: $(\gamma_L - \gamma_H) \cdot L - a \leq 0$

This condition simply says if a low type firm is always negligent (i.e., $\alpha = 1$), CSR insurance is no better than actuarially fair for that firm. If this assumption is not met, our model goes against the notion of insurance, which in practice is usually *less* than actuarially fair. This is also an assumption that justifies having different types in equilibrium. Indeed, if it were not met, it would say it is strictly profitable for a low type to always be negligent and then purchase CSR insurance, which would then provide additional surplus over the cost of insurance, despite being negligent for certain.

1.3 Determining Equilibria

To find our equilibria we begin with our first order conditions for best response functions for each type:

$\pi_H = -(1-\alpha_H)^2 - a \cdot \alpha_H - \alpha_H \cdot \gamma_H \cdot C$

$\Rightarrow FOC : 2(1-\alpha_H) - a - \gamma_H \cdot C \equiv 0$

$\Rightarrow$

$$\alpha_H^* = 1 - \frac{a + \gamma_H \cdot C}{2}$$

Similarly, we have $\pi_L = -(1-\alpha_L)^2 - \alpha_L \cdot \gamma_L \cdot C$

$\Rightarrow$

$$\alpha_L^* = 1 - \frac{\gamma_L \cdot C}{2}$$

Per our probability Lemmas in the appendix, if in equilibrium we have $\alpha_H^* \leq \alpha_L^*$, high type firms will have events less often and also suffer less of loss in firm value upon an event. The latter follows because the conditional probability of negligence is lower given an event for the high than for the low type with $\alpha_H^* \leq \alpha_L^*$, and thus there is less change in firm value to reflect a lower expected event cost given an event has happened. It so happens, the high type firm is indeed always weakly less negligent than the low type, which gives us the following proposition.

Proposition 1 *Assuming no arbitrage, high CSR types are weakly less negligent than low CSR types (i.e.,* $\alpha_H^* \leq \alpha_L^*$*). Consequently, high type firms experience events less often (i.e.,* $\Pr(E|low) > \Pr(E|high)$ *) and suffer less loss in firm value (i.e.,* $\eta_L > \eta_H$*) when facing an event.*

Proof:

As discussed above, from our probability Lemmas in the appendix, we only need to show we have $\alpha_H^* \leq \alpha_L^*$.

$\alpha_H^* \leq \alpha_L^*$
$\iff 1 - \frac{a+\gamma_H \cdot C}{2} \leq 1 - \frac{\gamma_L \cdot C}{2}$
$\iff \frac{\gamma_L \cdot C}{2} \leq \frac{a+\gamma_H \cdot C}{2}$
$\iff \gamma_L \cdot C \leq a + \gamma_H \cdot C$
$\iff (\gamma_L - \gamma_H) \cdot C - a \leq 0$, which proves our result. The second line comes by definition and the final line by our no arbitrage assumption. In the appendix we provide 2 conditions on our primitives to assure $\alpha_i^* \in [0, 1]$, since α_i^* is a probability measure for $i \in \{L, H\}$. $\square$

This proposition simply says, in equilibrium, the high type is weakly less negligent. But this then means based on Baye's rule, conditional on event occurring, the high type firm will experience less event cost. And by the definition of event

probabilities, the high type will face events less often. Thus, this provides two crisp predictions to test empirically, which we turn to now.

2 Empirical Examination

2.1 General Strategy

Our primary empirical aim is to test the notion CSR works as reputation insurance. That is, high CSR type firms enjoy a buffering of their firm value vis-a-vis a low-type firm during an event. From Proposition 1, we have the high type firms loose less value upon an event compared to a low type firm. It should also be clear from the previous arguments, we can readily expand our two types to three (or more) types to get the same predictions in monotonic ordering (see Minor (2010) for a formal argument). That is, the higher the type, the lower the event rate and the lesser the change in firm value. We will here empirically study three different types, as we describe in detail below.

Our empirical setting is product markets where the event is then a product recall. These events are often seen by the investment committee as a potential shock to a firm's value and reputation due to their signalling nature (Davidson and Worrell (1992) provide a review of past product recall literature. See also Hartman (1987) for a hedonic model treatment of recalls).

CSR reputation in product markets can be thought of as a firm's superior reputation earned from the conscientious creation, marketing, and distribution of its products. "Good" CSR will typically mean a firm will embrace superior quality assurance procedures in the development and production of its products, conduct ethical marketing campaigns, provide products with extra social value, provide products to disadvantaged demographic groups, and generally face product recalls voluntarily. Contrarily "Bad" CSR means firms are usually involved in regulation fights, suffer safety violation fees, accept lower product safety standards, and conduct limited due diligence on their supply chain.

Thus, we categorize our three firm types as follows. The lowest type, which we will call "Irresponsible" types are involved in "Bad" things ex-ante an event. The next type, "Responsible" types are not involved with "Bad" things, but neither are they involved in "Good" things—they are simple responsible corporate citizens. Finally, there are some exceptional firms that not only avoid being involved in "Bad" things, but are also participating in some extra "Good" things. These firms we dub "Stellar" types. This typology aligns with the notion it takes effort or cost to move from one to the other: as a firm becomes more conscientious in its activities it moves from Irresponsible to Responsible, and then with even further conscientious it becomes a Stellar type.

Our empirical strategy is to first calculate the abnormal change in firm values (i.e., after controlling for firm heterogeneity) during an event and to then regress these varying percent changes in firm value on the level of ex-ante CSR reputation, as well as various time, financial, and industry controls. We begin by reviewing our data characteristics and then turn to our event study methodology and regression model.

2.2 Data

Our Data consist of three components. The first part is the abnormal returns of various firms during our product recalls, which we describe in detail in the next section. The event returns are then merged with Compustat, our second set of data. For firm control data of the S&P 500 firms we have: annual sales ("Sales (net)"), asset value ("Assets-total"), market value ("common shares outstanding"$\times$"price-calender year-closing"), and percent of profits per share ("EPS (Basic) - Exclude Extra. Items "$\div$"price-calender year-closing"). Actual product recall events were obtained from manual collection of product recall events of S&P 500 firms as indexed by the Wall Street journal from 1991 through 2006. Although this categorization of product recalls is certainly not perfect, it is the primary source used by past product recall literature. Further, we wanted to have an ex-ante fixed criteria of selecting

recalls to prevent subjective inclusion or exclusion on the part of the researcher. We do note recalls included in the Wall Street journal press announcements are biased towards larger event recalls. However, our theory predicts it is these large scale recalls where we will see any effects, if any exist, from ex-ante CSR reputation.

Occasionally some firms had more than one event announcement in a year, most often a later press announcement related to the same event. For our data collection, we simply summed the abnormal returns together, following the methodology as shown below, by summing abnormal returns over event window days. Having more than one event in a year for a given firm occurred for 25 of the firm/ event years for an average of 1.5 additional events for each occurrence. This excludes autos, and so is out of a total of 147 firm event years. In other words, roughly 17% of the firm years had multiple events, each averaging an additional 1.5 events.

An important exception was automobile firms (GM, Ford, and Chrysler pre-1999). These firms are very different in that they have a product recall every year, and typically multiple recalls in a given year. Hence, we dummy for these three firms since the probability of recall is 100% every year versus less than a 3% chance of recall for all firms. Our results are also robust to simply dropping automobiles from the data.

Our final component of data is CSR ratings from KLD analytics. KLD is considered the "gold standard" of CSR ratings by social investment firms. It is also most commonly used in past related academic studies (see Chatterji et al. (2007) for a review). KLD conducts proprietary research to assign annual CSR ratings to publicly held firms across various dimensions .

For KLD's CSR ratings on the product dimension, analyst's score a firm on four areas of positive (i.e., "product strengths") and negative (i.e., "product concerns") CSR. The four areas of CSR strengths include "Product Quality," "R&D," "Benefits economically disadvantaged people," and "Other. " The concerns areas include "Product Safety," "Marketing Controversy," "Antitrust Concerns," and "Other." One can think of their rating scheme as a latent variable model: every firm is rated by analysts on various factors unobserved by the econometrician. Once a firm has a value above some threshold, they receive an outcome of one for each of 8 categories

(i.e., four strengths and four concerns), and zero otherwise. Finally, KLD then provides a Product Strengths and Product Concerns rating that are each simply coded 0,1,2,3 or 4, measuring the number of ones earned in each of the respective categories. Now we consider the probability of event for an "Irresponsible," "Responsible," and "Stellar" firms. "Irresponsible" is a firm that has at least total product concerns of 1 or greater before an event occurs. A "Responsible" firm has avoided bad marks (i.e., no product concerns marks) but neither does it have any exceptional marks (i.e., product strengths marks). Finally, "Stellar" firms have avoided bad marks while additionally obtaining exceptional marks. The empirically likelihood[8] (i.e., probability) of a Irrepsonsible, Responsible, and Stellar type firm of having an event over the entire 15 year period is 3.5%, 2.2%, 2.1%, respectively. Both of the higher types (i.e., Responsible and Stellar types) are statistically different from the Irresponsible type. However, the higher types are not statistically different from each other. Nonetheless, the monotonic ordering theory suggests in event rates is preserved. We now turn to the estimation of abnormal event returns for the core of our study.

2.3 Event Study

The particular event study methodology we use is a financial events study[9]. The idea behind a (financial) event study is to measure the effect of an event on firm value. This approach relies on finance theory's notion of market efficiency: firms are priced based on all currently available public information. Thus, once new public information is released it is almost immediately absorbed into the value of the firm via its stock price.

The event study methodology procedurally has the first step of estimating how a particular company's stock price changes in relation to various market factors *before* the event occurs. The particular factor model we use is the most commonly

[8] We exclude the Auto industry, as previously mentioned it has exceptionally high recall rates compared with all other firms.

[9] For a thorough review see MacKainly (1997).

used Fama/ French model. Expected return in this setting is estimated as an OLS specified by:

$$R_{i,t} = \alpha_i + \beta_i R_{M,t} + S_i SMB_t + H_i HML_t + \varepsilon_{i,t}$$

That is, the return of the stock equals a firm fixed effect , plus a sensitivity to the general market return R_M, sensitivity to small stocks versus large stocks (SMB), and finally a sensitivity to high versus low book to market type stocks. Coefficients are estimated from a time series just before but disjoint to the particular event of interest; here, following common practice, the estimation period begins 8 months prior and ends 30 days prior to the event. These coefficient estimates are then used to predict the return during the event period. That is, our predicted return around the event period becomes:

$$\widehat{R}_{i,t} = \widehat{\alpha}_i + \widehat{\beta}_i R_{M,t} + \widehat{S}_i SMB_t + \widehat{H_i} HML_t$$

The next step is to then use this estimated returns model from the first step to predict what the expected returns are during the event of interest and then calculate the "abnormal return," defined as the difference in *actual* return from the *predicted* return: $AR_{i,t} = \widehat{R}_{i,t} - R_{i,t}$. The cumulative abnormal return is then simply the sum of these returns. For our study, we used the day before and the day of the event announcement as our "event window." This is the most stringent of windows; however, we wanted to minimize the effect of any other previous or subsequent news confounds. We begin the window the day before, as is practice, to capture any "news leakage" the day before the event announcement. Thus cumulative abnormal return is then simply: $CAR_i = \sum_{t=-1}^{0} AR_{i,t}$, where 0 is the event day. Below is a graph of the event study estimation method.

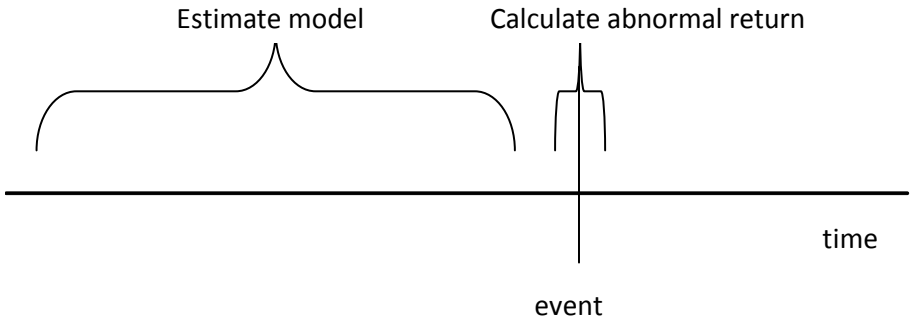

Financial Event Study Methodology

Our next figure reports the distribution of our CARs in relation to an estimated normal distribution of the data. As expected, the actual data has a fatter negative tail than a normal distribution. That is, we expect there to be more severe negative events and fewer severe positive events than a normal distribution would predict, since recalls are generally costly. There is also a much greater grouping of events at the center of just below zero, which seems to be close to the direct cost of a typical event. Hence, it seems there are routine recalls that are presumably of an exogenous type, thus only changing firm value by close to the direct recall cost effect. However, there a good number of events that have huge negative costs, much beyond routine recall expenses. It are these events against which we hypothesize superior CSR reputation will help protect firm value. We now turn to examining the relationship of event returns to ex-ante CSR reputation.

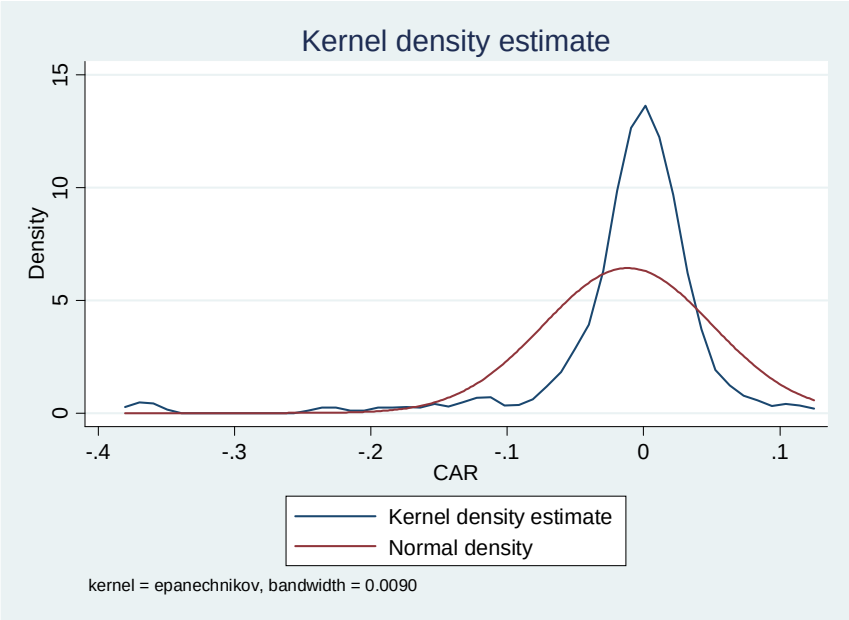

Actual Distribution of CAR versus Normal Distribution

2.4 Regression Model and Results

Once we calculate our abnormal returns (CAR), our final step is to examine any relationship between the ex-ante product CSR level and the respective CAR via a cross sectional regression. In particular, we specify the following:

$$CAR_i = \alpha + \beta_1 Auto_i + \beta_2 L.responsible_i + \beta_3 L.stellar_i + \beta_4 \overrightarrow{YEAR}_i + \beta_5 L.\overrightarrow{FIRM}_i + \varepsilon_i$$

CAR_i is again the cumulative abnormal return for firm i as calculated in the previous section. $Auto_i$ is simply a dummy for if the firm is an automobile manufacture. $L.responsible_i$ is a dummy for an average product CSR type, as defined in our former section, but its value is for the year *prior* to the event. Similarly, $L.stellar_i$ denotes the firm was a high type firm the year *before* the event. Thus, our low types, Irresponsible firms, are our baseline. That is, the former two dummies will tell us how much better compared with a low type the higher types fare under an adverse event.

We added year fixed effects for some specifications, denoted by the vector $\overrightarrow{YEAR_i}$. Finally, our firm financial and industry controls captured in the vector $L.\overrightarrow{FIRM_i}$,again valued based on year prior to an event.

Industry controls are made by dummying the naics industry code to the 2 digit level. We do not dummy at the 3 digit level because most dummies could not be estimated since we only have 184 firm event years over 15 years. Even with 2 digit level industry codes about half of our dummies cannot be estimated due to a paucity of observations. Below we report our regression results. When we control for lagged financials we lose some observations due to their not being in the index previously or recent mergers.

Dependent Variable: Cumulative Abnormal Return, baseline irresponsible type					
Variable	**None**	**Time**	**Financial**	**Industry**	**All**
Auto	.0315 (.0110)***	.0306 (.0103)***	.0314 (.0183)*	.0389 (.0196)**	.0157(.0768)
Responsible Type	.0279 (.0091)***	.0287 (.0095)***	.0305 (.0108)***	.0240 (.0098)**	.283 (.0134)**
Stellar Type	.0274 (.0097)***	.0251 (.0097)***	.0314 (.0104)***	.0190 (.0105)*	.0163 (.0158)
Year Control	NO	YES	NO	NO	YES
Firm Control	NO	NO	YES	NO	YES
Industry Control	NO	NO	NO	YES	YES
N	184	184	156	184	156
R Squared	0.0719	0.1413	0.0855	0.1191	0.2266

It Pays for Firms to be Responsible Ex Ante

Each column of results differs based on controls included. As can be seen, if a firm is able to carefully avoid being involved in bad activities (i.e., a Responsible type), it will save close to 3% of abnormal firm value should it face an event compared with the losses of a low type firm. This effect is also economically significant, as it amounts to an average saved firm value of over $600 million for the median firm (market) value of $23 billion. Meanwhile, if a firm is one of those exceptional firms that not only is careful to avoid bad activities, but also is involved in good ones—a

Stellar type, the savings in firm value is similar to the Responsible firm type. Even for the final column, using all controls, the difference in coefficients is not statistically different.

What this says is, it really pays to carefully avoid harmful activities in building reputation. However, it does not seem to pay, at least in an insurance sense, to build reputation of going the extra mile and providing some additional "good" activities.

Of course, there could be contemporaneous benefits in provisioning some extra "good" activities. However, as past literature has found, there has not been any financial benefit identified in doing so. In short, at least financially speaking, it really pays to be responsible by avoiding bad, but it doesn't pay to also be exceptionally good.

It might seem curious that adding time, financial, and industry controls does not change our CSR coefficient estimates. However, recall in the event study, we already controlled for time and financial characteristics as a well as a fixed firm effect for each firm. The main effect of adding all the controls is to introduce more noise while not affecting coefficient estimates. Indeed, when estimating all our coefficients (i.e., the final column) we are estimating a parameter per fewer than every 5 observations. Thus not surprisingly, our final column of estimates have our coefficients less significant.

One concern in estimating these abnormal returns is it could be the case abnormal return is simply the expected direct cost (e.g., the cost of replacing faulty automobile tires) of the product recall. If ex-ante CSR level is related with actual event cost, this in itself would be interesting, as it says CSR predicts the level of an event. Nonetheless, it would not support our CSR reputation story where there is uncertainty over the degree of negligence and its punishment that the market must price in immediately after an event. Unfortunately, the expected direct cost of a product recall is seldom made public (nor is it commonly disclosed ex-post). However, for our sample, roughly 10% of the announcements were accompanied by estimates of the direct event costs. For this subsample, the direct costs explain roughly 16% of the variation in CAR. Further, when a loss is sustained by a firm (i.e., a negative CAR), the direct costs represent 38% of the total loss on average. In absolute value

terms (i.e., because sometimes a firm has a positive CAR during an event), direct costs represent 26% of the value of CAR. Thus, though this sub-sample is only a small portion of the events, it suggests it is not the expected direct cost of an event driving differences in CAR. Further, the direct costs have a very narrow band of cost difference, whereas the change in abnormal firm value varies widely, suggesting there is much more than just product recall direct cost embedded in the CAR. Our theory suggests the CAR should be a combination of direct recall loss and (expected) financial loss if negligent.

In the end, the evidence is quite suggestive that ex-ante CSR reputation affects ex-post value of a firm. But it is not so much being involved with exceptionally good things as it is in carefully avoiding bad things. Thus, effective CSR in insurance terms seems to be more about making sure a firm is being careful to avoid harmful behavior, simply a responsible corporate citizen rather than an exceptional one.

One can even estimate what a firm should pay for such carefulness. As far as benefit, firms moving from Irresponsible to Responsible (or Stellar), saves some $600 million of firm value. With an incidence rate of roughly 2.5%, this means a risk neutral firm should be willing to pay over $15 million per annum to be a higher type.

3 Concluding Discussion

We have offered a (partial) solution to the puzzle of why firms invest in CSR when it apparently has no effect on current profits or firm value and yet is a costly activity. In particular, we proposed firms use CSR as a reputation insurance mechanism. We developed a model that showed we generally expect a firm with high ex-ante CSR to better weather a shock to firm value, as well as experience such events less often.

Empirically, we also found support for CSR acting as insurance, though the results are nuanced. In particular, it really pays to be responsible as a firm to avoid bad behavior—this tends to save over $600 million of abnormal should a firm face an adverse event. However, then becoming an *exceptional* corporate citizen by

engaging in additional stellar behavior, does not seem to pay additional dividends in an insurance sense or financial sense.

Hence, with many business schools and firms focusing on "doing well by doing good," a better mantra would be doing well by carefully avoiding bad. Or perhaps a more succinct way to put it: *primum non nocere*—first do no harm. In fact, if a manager becomes too focused on doing good, she very well could miss the seemingly more important task of avoiding harm.

This is the very advice BP missed as it spent significant sums of money to rebrand itself with its sunburst logo and the tagline "Beyond Petroleum." It was involved with some "good" environmental projects. However, it completely missed the call to avoid harm, to be careful to avoid bad events. Consequently, as its number came up, and nature drew a bad event for BP, investors and regulators swung the scale of responsibility firmly to negligent. And so it is, completely uninsured, BP will have to pay substantial sums of money.

"Most of the rhetoric on CSR may be about doing the right thing and trumping competitors, but much of the reality is plain risk management. It involves limiting the damage to the brand and the bottom line that can be inflicted by a bad press and consumer boycotts, as well as dealing with the threat of legal action."

Economist , January 7,2008

4 Appendix

4.1 Probability Lemmas

We now prove two helpful lemmas. The first says if both firms of different types decide on the same level of negligence, given an event does occur, for a low type firm, it is more likely the event was caused by negligence over noise versus the high type firm. Further, low type firms will have events more often than high type firms, given they both choose the same level of negligence. Since for these Lemmas we assume each firm is a different type, we will denote $i \in \{L, H\}$ which then corresponds to such firms respectively choosing $t \in \{L, H\}$.

Lemma 1 *If $\alpha_L = \alpha_H$, then $P_L > P_H$ and $\eta_L > \eta_H$.*

Proof: The first claim follows immediately by construction.

For the second implication, we want to show:$\eta_L \equiv \frac{\alpha_L \times \gamma_L}{\alpha_L \times \gamma_L + (1-\alpha_L) \times \gamma_0} > \frac{\alpha_L \times \gamma_H}{\alpha_L \times \gamma_H + (1-\alpha_L) \times \gamma_0} \equiv \eta_H$, invoking our assumption of $\alpha_L = \alpha_H$. However, this is always true since $\eta_i(t)$ is increasing in γ_t and $\gamma_L > \gamma_H$. Hence, conditional on an event, and assuming both firms had chosen the same α_i, it is always more likely the low type firm actually committed negligence that caused the event. ■

Our second lemma says if *conditional on an event* it is equally likely either a high or low type firm committed negligence, then ex-ante the high type firm selected a higher level of negligence. However, the low type firms *will still have a higher incidence of events,* even though they are choosing a lower level of negligence.

Lemma 2 *If $\eta_L = \eta_H$, then $\alpha_L < \alpha_H$ and $P_L > P_H$.*

The first implication follows from the analysis of the previous lemma: starting with $\alpha_L = \alpha_H \Rightarrow \eta_L > \eta_H$, which means to get to $\eta_L = \eta_H$ we require $\alpha_L < \alpha_H$, since $\frac{\partial}{\partial \alpha_H}(\eta_H) > 0$.

For the second implication, we first fix $\alpha_L < \alpha_H$ and γ_L. Note if $\alpha_L \times \gamma_L = \alpha_L \times \gamma_H$, our equality does not hold since $(1 - \alpha_H) \times \gamma_0 < (1 - \alpha_L) \times \gamma_0$. And in

particular with $\alpha_L \times \gamma_L = \alpha_H \times \gamma_H$, we get $\frac{\alpha_L \times \gamma_L}{\alpha_L \times \gamma_L + (1-\alpha_L) \times \gamma_0} < \frac{\alpha_H \times \gamma_H}{\alpha_H \times \gamma_H + (1-\alpha_H) \times \gamma_0}$. Per Lemma 1, $\frac{\alpha_H \times \gamma_H}{\alpha_H \times \gamma_H + (1-\alpha_H) \times \gamma_0}$ is strictly increasing in γ_H, which then means it must be the case that $\alpha_L \times \gamma_L > \alpha_H \times \gamma_H$ to maintain our needed equality. This then give us: $\alpha_L \times \gamma_L > \alpha_H \times \gamma_H \Rightarrow \alpha_L \times (\gamma_L - \gamma_0) > \alpha_H \times (\gamma_H - \gamma_0) \iff \alpha_L \times (\gamma_L - \gamma_0) + \gamma_0 > \alpha_H \times (\gamma_H - \gamma_0) + \gamma_0 \iff P_L > P_H$. That is, the probability of an event is greater for a low type, even though the high type chooses a greater level negligence. ■

4.2 Regularity Conditions for a_i^*

Here we derive conditions our primitives to assume in equilibrium a_i well defined. That is, since it is a probability measure we have: $a_i^* \in [0, 1]$ for $i \in \{H, L\}$. Since we know $a_H^* \leq a_L^*$, we need to solve for just two conditions.

First, we need $0 \leq a_H^*$. This happens when:

$\alpha_H^* = 1 - \frac{a + \gamma_H \cdot C}{2} \geq 0$

$\iff 2 \geq a + \gamma_H \cdot C$

$\iff \frac{2-a}{\gamma_H} \geq C$

That is, we need the event cost to not be too costly to make sure firms will commit a positive level of negligence.

Second, we need $\alpha_L^* \leq 1$.

$\alpha_L^* = 1 - \frac{\gamma_L \cdot C}{2} \leq 1$

$\iff \frac{\gamma_L \cdot C}{2} \geq 0$

But this last condition is always met. Hence, we are left with a single condition of $\frac{2-a}{\gamma_H} \geq C$ to eliminate the uninteresting case of both firms never committing any negligence.

5 The Assumption of Heterogeneous Firms

If we do not have heterogeneous firms, our problem is not interesting, as we have no firms of differential CSR levels to compare. Thus, we now consider how we could

begin with homogeneous firms but in equilibrium have heterogeneous firms. Below we provide a sketch of several independent sufficient conditions that would yield this outcome.

1) "Standards" increase in the number of high type firms.

This is where there is an increasing cost of the race to the top. That is, as more firms become high types, the standard to then become a high type becomes even greater. Hence, we have some maximal standard level where some marginal firm is just indifferent now between bearing the cost of being a high type or remaining a low type. If all firms are homogeneous, any mix of high and low types is an equilibrium. If however, firms have differential cost structures in providing CSR, then all the higher cost types above the marginal type will remain low types while the rest become high types.

2) There is a fixed pool of consumer demand for CSR firm provided products.

Here we have the same argument as above but now it is the potential benefit rather than cost separating firms. As more firms choose to become high types, profit is driven to zero until no other firms can profitably serve CSR demanding customers. We again get a bifurcation of firm types.

3) CSR inputs have market power.

If the CSR input markets—e.g., specialized labor or safety devices— have market power, then firms taking on high CSR, will bid up the cost of inputs until profit is driven to zero. Consequently, some population of firms will be high types and the rest low types.

4) Managers are heterogenous in that some get a "warm" glow from provisioning CSR.

This argument says managers differentially value managing a "responsible" firm and thus we get variation in firm levels of responsibility simply from managers heterogeneous preferences, though firms are the same. See Fisman et al. (2007) for an expansion of this argument.

References

[1] Baron, David P and Daniel Diermeier. 2007. "Strategic Activism and Nonmarket Strategy." *Journal of Economics & Management Strategy* 16 (3): 599–634

[2] Chatterji, Aaron, David I. Levine, and Michael W. Toffel. 2007. "How Well Do Social Ratings Actually Measure Corporate Social Responsibility?" Working Paper 07-051. Harvard Business School.

[3] Diamond, D. 1989. "Reputation Acquisition in Debt Markets." *Journal of Political Economy* 97(4):828-62.

[4] Davidson, Wallace and Dan L. Worrell. 1992. "The Effect of Product Recall Announcements on Shareholder Wealth." *Strategic Management Journal,* 13 (6): 467-473.

[5] Elfenbein, Hillary. 2007. "Does it Pay to Be Good?" Working Paper. UC Berkeley.

[6] Fisman, Ray, Geoffrey Heal, and Vinay B. Nair. 2007. "A Model of Corporate Philanthropy." Working Paper. Columbia University.

[7] Goyal, Ashima. 2006. "Corporate Social Responsibility as a Signalling Device for Foreign Direct Investment." *International Journal of the Economics of Business* 13(1): 145-163

[8] Hartman,R .S. 1987. "Product Quality and Market Efficiency The Effect of Product Recalls on Resale Prices and Firm Valuation." *Review of Economics and Statistics*, 6 (9):367-72

[9] Hermalin, Benjamin E. 2008. "Firm Value and Corporate Governance: Does the Former Determine the Latter?." Working Paper. UC Berkeley.

[10] Kotchen, Matthew, and Jon J. Moon. 2007. "Corporate Social Responsibility for Irresponsibility." Working Paper. UC Santa Barbara.

[11] MacKinlay, A. Craig. 1997. "Event Studies in Economics and Finance." *Journal of Economic Literature* 35: 13-39.

[12] Minor, Dylan 2009. "CSR as Reputation Insurance: Adding Endogenous Enforcement." Working Paper. UC Berkeley.

[13] Minor, Dylan and John Morgan 2010. "CSR as Reputation Insurance: Primum Non Nocere." Working Paper. UC Berkeley.

[14] Tadelis, Steve. 2002. "The Market for Reputations as an Incentive Mechanism." *Journal of Political Economy* 110(4):854-82.

[15] Vogel, David. 2005. *The Market for Virtue*. Washington D.C.: Brookings.