Do Managers Do Good With Other People's Money? *

Ing-Haw Cheng[†] Harrison Hong[‡] Kelly Shue[§]

First Draft: November 15, 2011
This Draft: August 15, 2013

Abstract

We find support for two key predictions of an agency theory of unproductive corporate social responsibility. First, increasing managerial ownership decreases measures of firm goodness. We use the 2003 Dividend Tax Cut to increase after-tax insider ownership. Firms with moderate levels of insider ownership cut goodness by more than firms with low levels (where the tax cut has no effect) and high levels (where agency is less of an issue). Second, increasing monitoring reduces corporate goodness. A regression discontinuity design of close votes around the 50% cut-off finds that passage of shareholder governance proposals leads to slower growth in goodness.

*Hong acknowledges support from the National Science Foundation through grant SES-0850404. We thank Jeremy Stein, Vicente Cuñat, Uday Rajan, E. Han Kim, Amy Dittmar, Nadja Guenster and seminar participants at the Swedish Institute for Financial Research, St. Gallen, the Berkeley-ECGI Conference on CSR, U of Texas at Dallas, 2013 Meetings of the American Finance Association, 2013 UC Davis-CalPERS Sustainability Symposium and Gamma Foundation Conference for helpful comments. Online appendix available at http://www.dartmouth.edu/∼icheng/.

[†]Dartmouth College, Tuck School of Business, email: ing-haw.cheng@tuck.dartmouth.edu.

[‡]Princeton University Department of Economics and Bendheim Center for Finance, and NBER, email: hhong@princeton.edu.

[§]University of Chicago, Booth School of Business, email: kelly.shue@chicagobooth.edu.

1. Introduction

Milton Friedman (1970), in a biting op-ed for the New York Times, declared that the only social responsibility of corporations is to make money. His view regarding spending on corporate social responsibility (CSR) is one of agency problems, managerial perks and pet projects. In other words, managers are doing good with other people's money. A contrasting popular view of the business case for corporate goodness is that firms do well by doing good. Firm goodness spending generates a halo effect that increases firm profits and insulates firms from risks associated with litigation or regulation.[1]

Despite the many papers that have been written on this topic, it is difficult to draw causal conclusions regarding the motives for corporate goodness spending because of a lack of a clear identification strategy. Notably, the large literature on "doing well by doing good" explores the correlation between firm performance and CSR and concludes in favor of the profit thesis. Evidence from experiments and field studies also suggest that there are potential halo effects from being charitable or good (see, e.g., Elfenbein, Fisman, and McManus, 2012; List, 2006). However, altruistic managers may invest in CSR to protect stakeholders such as employees or the broader community, in ways that do not directly contribute to shareholder wealth. As Tirole (2001) points out, a stakeholder instead of shareholder maximization paradigm can very quickly lead to mission creep and agency perils.

Indeed, recent research on the thesis of doing well by doing good finds that market reactions to increased goodness announcements are mixed, consistent with the idea that some forms of goodness may be value destroying (Kruger, 2013). Even more troubling, using financial shocks such as the bursting of the Internet bubble and the credit crisis of 2007, Hong, Kubik, and Scheinkman (2011) show that the correlation between performance and goodness is related to omitted variables bias in the form of heterogeneity in firm financial constraints.

[1]The thesis of "doing well by doing good" has long been advanced in the management literature and is reviewed in Benabou and Tirole (2010) and Heal (2005). A number of theories implicitly rely on the idea that firms are well-positioned to deliver warm-glow feelings (Becker, 1974; Andreoni, 1989) to consumers. For instance, Besley and Ghatak (2005) models strategic complementarities involving goodness in the production function. Baron (2001) models strategic deterrence of regulation through the use of corporate goodness. Goodness may also pay by improving employee efficiency, reducing conflicts among stakeholders, mitigating litigation risk, deterring potential regulation, signaling product quality, and improving investor and consumer relations by preventing product or capital market boycotts by socially responsible consumers or investors.

Yet, the stakes in this debate have never been higher. Anecdotal evidence suggests that some firms, especially large corporations, invest hundreds of millions of dollars annually on energy conservation projects, employee and community development programs, or other altruistic endeavors.[2] Many institutional investors are increasingly adopting socially responsible factors, also known as "environmental, social and governance" (ESG) factors, in their portfolio choices. For example, the United Nations Principles of Responsible Investing (UNPRI) has institutions that manage around 1.5 trillion dollars globally as signatories. The ESG movement tends to make the argument for socially responsible investing on the basis of a combination of the business case for corporate goodness and moral principles.

In this paper, we provide a first attempt to identify whether there is an agency motive for CSR. Namely, we test the two key predictions of an agency theory in which managers engage in unproductive CSR as a way to enjoy private benefits. These two key predictions typically emerge from most variants of agency models in which low managerial ownership stakes and imperfect monitoring lead to managers who maximize a combination of private benefits and firm profits. Our novel perspective here is that agency problems need not simply manifest themselves as managerial selfish perks, as has traditionally been framed since Jensen and Meckling (1976), but also as managerial altruism or social preferences (see, e.g., Fehr and Schimdt, 1999; Charness and Rabin, 2002).

The first and most basic prediction is that increasing the ownership stakes of a manager leads to a reduction in firm goodness. Since managerial ownership stakes regulate the degree of agency conflicts between the manager and shareholders, these models predict that larger stakes lead to less of an agency conflict and lower unproductive CSR spending. The second prediction is that an increase in monitoring leads to less firm goodness. Monitoring is costly and is a substitute for larger managerial ownership stakes, so increasing monitoring through improved governance yields the same comparative statics as increasing managerial stakes.

We measure corporate social responsibility using the most comprehensive scores in the literature, the Kinder, Lydenberg and Domini (KLD) scores of CSR. KLD scores are developed

[2]Hong, Kubik, and Scheinkman (2011) provide some figures, including $100 million for global education programs and energy conservation by Intel, $160 million to community and employee philanthropic programs and earmarked billions more for the development of eco-friendly products by General Electric, and Google's 1% program to invest 1% of its profits in philanthropic and non-profit interests. In addition, firms may switch to greener and more costly production processes or voluntarily choose to pay living wages to employees.

by a for-profit company, akin to a credit rating agency. The scores measure firm-level social responsibility along the lines of community relations, product characteristics, environmental impact, employee relations, and diversity.[3] KLD scans public databases, such as those on employee strikes and Environmental Protection Agency (EPA) violations, and uses a team of analysts to measure these and other social responsibility dimensions of firm production. We explain in Section 4 why these scores are a reasonable, albeit imperfect, proxy for socially responsible expenditures. We also show that KLD scores are correlated with firm donation dollars but offer a more informative measure of corporate goodness. Admittedly, we are not trying to break new ground in this paper in terms of measuring CSR. Our analysis uses KLD scores as the left-hand side dependent variable as opposed to a right-hand side control variable, as has predominantly been the case in the literature. Therefore, noisiness in our left-hand side variable will raise standard errors but should not bias our estimates.

Where we try to break new ground in the CSR literature is to more cleanly identify a motive for CSR. To test our first prediction, we expand on Chetty and Saez (2005) and Chetty and Saez (2010), who argue that the Jobs and Growth Tax Relief Act of 2003 (commonly known as the 2003 Dividend Tax Cut) raised the effective ownership stakes of managers and insiders by cutting the highest statutory dividend tax rate from 35% to 15%. The 2003 Dividend Tax Cut was largely unanticipated and led to a surge in dividend payouts, which they argue was a result of insiders substituting away from perk projects towards dividends after their effective ownership stakes increased. Consistent with their agency perspective, evidence from Poterba (2004), Auerbach and Hassett (2006) and Auerbach and Hassett (2007) point to higher equity prices as a result of the tax cut.

Our work also helps fill a gap in the existing 2003 Dividend Tax Cut literature: lack of direct evidence of cuts in perk spending along with the rise in dividend payouts. Given the debate over whether the higher dividend payouts were due to firms simply substituting from share repurchases to dividends, evidence on cuts in perk spending during the post-2003 years would be extremely informative. By examining how the tax cut affected firm goodness, we shed light on a motive behind CSR and also shed light on whether agency was a force in the rise in

[3]Beginning in 2010, KLD scores were re-tooled into what are now known as MSCI ESG scores, after MSCI acquired RiskMetrics in 2009.

dividends post-2003. It is important to highlight that our analysis does not test whether all or average CSR is unproductive. Instead, we test whether the the extra marginal dollar spent on CSR is a form of perk spending by exploring how KLD scores change following the tax cut.

To identify the effect of the tax cut on CSR, we test a canonical agency model that predicts that the sensitivity of firm perk spending to the dividend tax cut should be a non-monotonic function of pre-tax insider ownership levels. In general, the dividend tax cut increases the effective ownership stake of management because each share's claim to after-tax profits increases. If the marginal dollar spent on goodness is unproductive and the result of agency conflicts, then the tax cut should lead to an average decline in any unproductive goodness. However, the dividend tax cut has a very small effect on effective managerial ownership if ex-ante managerial share ownership is very low. When managers own close to zero percent of the firm, even a large dividend tax cut will not change their incentives. At the same time, managers with very high levels of ownership are likely to have incentives that are already closely aligned with those of shareholders. These high ownership managers spend very little on inefficient goodness to begin with, so there is little scope to cut agency-motivated goodness following a dividend tax cut. Hence, the model predicts that managers with intermediate firm ownership stakes should respond more to a dividend tax cut than managers with very low or very high ownership stakes.

Consistent with this prediction, we find that firms' goodness scores fell on average after the 2003 Dividend Tax Cut, and that the scores of medium insider ownership firms fell by relatively more. We employ a differences-in-differences methodology that flexibly allows for the effect of the tax cut to vary through time in order to capture its long-run dynamic effects. We split firms into ownership terciles based on 2001 and 2002 executive ownership data. From 2003 through 2004, average goodness among the medium ownership firms fell more compared to both high and low ownership firms. Relative to low ownership firms, medium ownership firms' goodness scores fell more by 0.14-standard deviations of their 2002 distribution. Relative to high ownership firms, medium ownership firms' goodness scores fell more by 0.22-standard deviations. These effects are statistically significant at the 10% and 1% levels, respectively. By the end of 2006, the difference-in-difference with high ownership firms widens to 0.29-standard deviations, although the difference-in-difference with low ownership firms attenuates to 0.10-standard deviations.

The non-monotonicity with respect to ex-ante managerial ownership is central to the identification strategy of the first prediction. It rules out that our findings are driven by (1) changes in firm goodness around the tax cut that are common across all firms, and (2) changes in corporate goodness around the tax cut that are monotonic with respect to insider ownership or other correlated variables. The only alternative time pattern that we cannot rule out would have to separately affect changes in goodness scores of medium as opposed to low and high ownership firms specifically at the time of the tax cut. However, one may be concerned that the stronger response among medium ownership firms is spuriously generated through the functional form of a linear regression using a lucky grouping of tercile breakpoints. To address this concern, we consider a number of additional analyses, including fitting non-parametric local polynomials to better capture the shape of the response. Overall, we find robust evidence in favor of this non-monotonicity.

To test the second prediction that improved monitoring leads to a decline in unproductive CSR, we exploit a regression discontinuity (RD) experiment using close proxy contests regarding shareholder-initiated governance proposals. The identifying assumption is that close votes around the 50 percent cut-off are random in terms of whether a governance proposal is passed and represent plausibly exogenous shocks to the monitoring of managers. We build on earlier work, which finds that, while shareholder proposals are non-binding, close votes around the 50% cut-off lead to discontinuously substantial changes in the implementation of governance proposals (Ertimur, Ferri, and Stubben, 2010; Thomas and Cotter, 2007). Our analysis builds most closely on Cuñat, Gine, and Guadalupe (2010), which uses a vote share RD approach to show that passage of shareholder proposals increase firm value. Cuñat, Gine, and Guadalupe (2010) also shows that the passage of shareholder proposals can improve governance through channels other than through changes in governance provisions, such as by empowering shareholders.

We find that firms in which shareholder proposals narrowly pass experience significantly slower growth in goodness scores than firms in which the proposals narrowly fail. Indeed, the economic magnitudes are sizeable. We find that firms that just failed to pass proposals experienced greater annual growth of KLD scores in the year of the vote than those that just-passed: the difference in the change in KLD scores is around 0.6. This difference in growth rates for firms around the cut-off is one-third of a standard deviation of the dependent variable,

an effect that is statistically significant at the 5 percent level.

Overall, improvements in managerial incentives and governance lead to a reduction in firm goodness, implying that the marginal dollar spent on goodness is a result of agency problems. We emphasize that our results apply to marginal rather than average goodness spending. Some goodness may be productivity-enhancing, but on the margin, managers over-invest in goodness due to agency problems, and, as a result, reduce goodness in response to shocks that make them care more about shareholder wealth.

2. Model

We consider the following stylized two-period agency model following Chetty and Saez (2010). The manager can decide how to spend the firm's cash Γ at $t = 0$: invest K in a productive project that yields $f(K)$ net profits for shareholders at $t = 1$, invest G in a corporate goodness project that gives the manager private benefits of $g(G)$ at $t = 1$, or pay out dividend D at $t = 0$. We assume that f and g are positive, strictly increasing and concave functions, with $f(0) = g(0) = 0$.[4] The function g is a reduced form that captures any managerial private benefits generated by G.

We are interested in whether any observed changes in goodness in response to the Dividend Tax Cut and shareholder voting are consistent with the model's predicted response of G, and for brevity refer to G as goodness spending going forward. It is worth noting at the outset that we allow for productive goodness spending to be included in K, whereas G captures any unproductive goodness spending. If managers adjust measures of total goodness in response to shocks such as the Dividend Tax Cut or the passage of shareholder proposals in a manner consistent with the model's comparative statics about G, then managers had invested in goodness to the point where marginal goodness spending was unproductive.

The fraction of insider ownership is denoted by α and is treated as an exogenous parameter. The manager's ownership stake represents his claim to a fraction of the payouts from the firm,

[4]We follow the convention of Chetty and Saez (2010) in writing the model in terms of net profits and in assuming that the manager returns the capital G used for investment in private benefits during period 1. The results thus do not depend on the capital G being burned up, although one may also assume that it is without changing any predictions. Details of the model are discussed in the Online Appendix.

which is reduced by the dividend tax rate τ.[5] We also assume that shareholders monitor the firm by increasing the manager's weight on firm payouts relative to private benefits by an amount η. The parameter r is the discount rate (there is no uncertainty in this model). The manager maximizes a linear combination of after-tax payouts from production and his private benefits from goodness spending, subject to the constraint that his spending on capital and goodness equal his cash on hand: $\Gamma = K + G + D$. More specifically, the manager's problem is given by:

$$\max_{K,D\geq 0} \alpha\left(1-\tau\right)\left(1+\eta\right)\left[D+\frac{f\left(K\right)+\Gamma-D}{1+r}\right]+\frac{g\left(\Gamma-K-D\right)}{1+r}.$$

Let $\omega = \alpha(1-\tau)(1+\eta)$ denote the effective ownership stake of the manager in the firm. It increases with the insider ownership stake (α), decreases with the dividend tax rate (τ) and increases with the amount of monitoring (η). Our identification strategies focus on these three parameters. Since increasing η is the same as decreasing τ, we assume for simplicity in this exposition that η is zero. We further assume $g(G) = BG$, which simplifies the intuition, and that $B < (1-\tau)r$, which ensures that there are both dividend and non-dividend-paying firms in equilibrium. In the Online Appendix, we derive the model under more general conditions.

A decrease in the dividend tax rate τ increases the effective ownership stake ω of the manager. A basic implication of the model is then that the dividend tax cut raises the marginal cost of goodness spending and leads to less goodness spending: $\partial G/\partial\tau \geq 0$. As we will discuss shortly, the dividend tax cut leads to strictly less goodness spending among a subset of firms where management owns a moderate amount of the shares. The average goodness spending across the ownership distribution will thus fall after the tax cut.

Overall, an increase in ω means that the manager makes better decisions for shareholders and should invest less in his pet projects. This implication is verified by comparing corporate goodness measures before and after the 2003 Dividend Tax Cut. However, this implication consists of a single difference and is difficult to cleanly test because concurrent changes to firm conditions (e.g., time trends or changes in other government policies or macro conditions) may spuriously drive the relationship between the Dividend Tax Cut and the subsequent decline in

[5]The model assumes that production and capital stock are paid out to shareholders as dividends in period 1. Our results remain similar if the firm instead pays out a fixed portion using dividends and the remainder through share repurchases. We also assume that the corporate tax rate is zero as it will not affect our analysis.

corporate goodness. We state it here simply to confirm that the aggregate trend supports the agency perspective.

Therefore, we exploit a more nuanced prediction of the agency model. The first prediction centers on how the extent to which managers cut goodness spending following the tax cut should vary according to the ex ante insider ownership level α. The model predicts that managers who have an intermediate level of ownership will cut more goodness relative to those with higher and lower levels of ownership.

To see why, consider the following three cases. A firm with zero insider ownership, $\alpha = 0$, invests nothing at all in productive capital ($K = 0$), pays no dividends in period $t = 0$, and invests everything in goodness ($G = \Gamma$). The sensitivity of ω, the effective ownership stake, to the dividend tax rate τ is proportional to α. When α is zero, managers will be very insensitive to the tax cut and will not decrease their goodness spending.

Medium insider ownership firms, or those with $\alpha \in (0, \bar{\alpha})$ for $\bar{\alpha} = \frac{B}{(1-\tau)r}$, do not pay dividends in period $t = 0$, and have an interior solution for capital and goodness spending given by $\alpha(1-\tau) f'(K) = B$ and $G = \Gamma - K > 0$. They cut goodness in response to the tax cut as it makes the after-tax marginal return on capital more attractive, and substitute this towards productive investment K. In the extreme case where a discrete tax cut is large enough to induce the manager to start paying dividends in period $t = 0$, goodness spending is reduced to zero.

High insider ownership firms, or those with $\alpha \in (\bar{\alpha}, 1]$, are invested the first-best level of capital K^*, defined as the solution to $f'(K^*) = r$. All remaining cash is paid out as dividends in period $t = 0$, since the after-tax risk-free rate is higher than the return on goodness, B. Firms with extremely high ownership stakes do not have large agency problems, so there is little scope for reducing G relative to firms with intermediate levels of ownership, since $G = 0$.

Therefore, medium ownership firms will cut goodness more than high ownership and zero ownership managers. Under mild regularity conditions on f, we show that medium ownership managers also cut goodness in response to a tax cut more than managers with positive yet lower ownership levels, so that the effect is not limited to a comparison to zero ownership managers who are at a corner solution. We also show that the prediction holds when considering a broader class of increasing and concave functions f and g. The intuition is the same: the marginal effect

of the tax cut on effective ownership is proportional to share ownership, and thus is lower for managers with lower ownership. High ownership managers have little goodness spending to cut; therefore, medium ownership managers are the most sensitive to the tax cut. We provide all details in the Online Appendix for the interested reader. We summarize with:

Prediction 1. *In response to a dividend tax cut, medium ownership managers cut goodness G more than low ownership managers and high ownership managers.*

We empirically test this prediction using a difference-in-differences approach comparing the change in goodness following the Dividend Tax Cut across firms with different ex-ante managerial ownership stakes. We exploit the non-monotonicity in the relationship between insider ownership and how aggressively firms should cut goodness in response to the dividend tax cut.

When testing Prediction 1, we assume that ownership stakes α are exogenous in some short horizon. That is, when the Dividend Tax Cut passed, the firm did not have sufficient time to re-contract and adjust ownership levels. An additional caveat is that we ignore the possibility that firms can repurchase shares in lieu of paying dividends. We follow Chetty and Saez (2005) in assuming that frictions exist which lead firms to pay out at least part of their cash flows through dividends. In practice, over three-quarters of the firms in our sample do pay dividends and other firms may be affected by the tax cut because they anticipate paying dividends in the future. Investigation of why firms pay dividends when an alternative payout option exists is outside the scope of this paper.

Our approach also provides a complementary test of the real effects of the Dividend Tax Cut by examining whether it induced firms to alter the composition of their investments. Examining whether total investment (K plus G) responds to the Dividend Tax Cut suffers from the fact that total investment mixes together both productive and unproductive investments, and that the tax cut affects these two types of investments differently. Hence, our analysis can be thought of as an out-of-sample test for whether the rise in dividend payouts following the 2003 Dividend Tax Cut was in fact due to a reduction in agency problems.

A substitute for giving the manager a larger effective ownership stake is to engage in costly monitoring, which can be modeled with a parameter η. As monitoring η increases, the manager

puts less weight on his pet projects and more weight on profits. As such, a second prediction of our model is that improvements in monitoring or governance of firms should reduce goodness spending.

Prediction 2. *An increase in the governance parameter η leads firms to decrease their corporate goodness G.*

We empirically test this prediction using a regression discontinuity approach focusing on close governance proxy votes. The RD design allows us to estimate the effect of quasi-random changes in governance on corporate goodness.

With these two predictions in hand, we turn to the empirical analysis. Note that this model is highly stylized and cannot speak to the magnitude of reactions with respect to shifts in the incentives or governance parameters. We use the model to generate qualitative directional guidance.

3. Data

Our study uses data from four sources. Both quasi-experiments rely on ratings of corporate social responsibility from the Kinder, Lydenberg, Domini, & Co. (KLD) database. Both experiments also use stock prices and shares outstanding from the Center for Research in Security Prices (CRSP) and accounting variables from CompuStat. KLD's coverage of begins in 1991; our analysis uses KLD information from 1991 to 2012. Following Hong, Kubik, and Scheinkman (2011), we focus on five dimensions of firm goodness documented by KLD: community activities, diversity, employee relations, environmental policies, and the social benefits of their products. KLD reports for each firm its number of strengths and concerns across these five dimensions. As pointed out by Hong, Kubik, and Scheinkman (2011), principal components analysis places roughly equal weights across these five dimensions. We therefore construct a firm's goodness score by aggregating the total number of CSR strengths and subtracting the total number of CSR concerns across these five dimensions. By aggregating these scores, we show below that this does a better job of picking up firm goodness than any of the individual measures alone, as each individual measure can be noisy.

KLD also documents issues associated with human rights, corporate governance, and whether the business itself is controversial. Most human rights scores are only available during a select few years in the 1990's, so we do not include them; in general, we only include scores that have been consistently rated by KLD every year from 1991 through 2009.[6] We do not include corporate governance as this may be directly related to insider ownership, which we use to form our portfolios. Whether the business or industry is controversial is a firm fixed effect and thus not applicable to our analysis which focuses on changes in goodness over time.

For the quasi-experiment using the 2003 Dividend Tax Cut, we use data on top-five total executive share ownership from S&P's ExecuComp database, with dividend data from the CRSP Monthly File (aggregated to an annual frequency) and other firm variables from CRSP and CompuStat. To be included in our sample, a firm must be in the S&P 500 in the year 2001 or 2002 and have inside share ownership data from ExecuComp. We focus on S&P 500 firms as KLD began a staggered expansion of their sample universe in 2001, which had previously included the S&P 500 and Domini 400 Social index constituents (which overlapped significantly), and we do not wish to introduce this selection issue into our analysis. All together, we have 503 firms, which is slightly more than 500 because of turnover in the S&P 500 index.

For the quasi-experiment involving close governance votes, our data on proxy contests for governance proposals comes from Riskmetrics, which covers all shareholder-initiated proposals at S&P 1500 firms from 1997 to 2011. The data includes detailed voting records for 5,262 governance-related proposals in areas such as anti-takeovers measures, compensation oversight, board structure, and auditing. Riskmetrics provides data on the type of proposal, the shareholder proponent, the recommendation of the board of directors, and the percentage of votes in favor of the proposal.[7]

[6]KLD changed their ratings methodology significantly beginning in 2010. We deal with this in Section 6.

[7]The Riskmetrics Shareholder Proposals dataset contains information on the number of votes in favor for each proposal as a percentage of all votes for and against (excluding abstentions and broker-nonvotes), rounded to the nearest percentage point. Because we are interested in votes very close to the 50 percent cutoff and would benefit from more precise vote share data, we supplement our main data with more detailed information about the exact number of votes for, against, and abstained using two additional data sources also provided by Riskmetrics: Voting Analytics and Voting Results.

4. Measuring Corporate Goodness

We begin our empirical analysis by stating the advantages and disadvantages of our left-hand side variable, the KLD score, as a measure of firm goodness. The disadvantage is that, like most ratings produced by commercial firms, there is a black-box aspect to the KLD score. Ideally, one would like to have data on dollar amounts spent on corporate social responsibility. Direct donations are one obvious measure, but these amounts can be small in comparison to expenditures on cleaner production technologies or employee benefits and living wages. For example, Google announced a plan to spend 1% of its profits on the development of socially responsible products. This 1% figure for Google dwarfs the direct donations numbers for Google and other firms.

There is mounting evidence in the literature that the equal-weighted KLD scores are indeed informative of corporate social responsibility. First, Chatterji, Levine, and Toffel (2009) find that KLD scores capture the past environmental performance of firms and also forecast the future environmental performance of these firms reasonably well. Environmental performance is the most straightforward of the categories to cross-validate, as the authors can reference EPA fines and whether a firm is mentioned on hazardous waste lists. It is more difficult for researchers to independently validate other dimensions of social performance, and this is what KLD was designed to do.

Second, we provide some anecdotal evidence for the timeliness of KLD scores in picking up the social responsibility of firms. In Figure 1, we plot the KLD scores for Apple and Google through time. We also plot the equal-weighted KLD scores for firms in the S&P 500. First, consider Apple. Apple is rated as one of the most socially responsible firms from 1991-1994. It scores range from a high of 6 to a low of 4 during this period. The average S&P 500 score during this period is between zero and 1. However, in 1996, Apple's KLD score falls dramatically to 1 and remains substantially lower than its previous average for most of the remaining sample. The drop coincides with Steve Jobs returning as CEO and taking control of Apple in 1996. The company was also nearing bankruptcy during this period. Steve Jobs famously decided to cut all of Apple's corporate social responsibility programs. He subscribed to the view that the only responsibility of a corporation is to make and sell great products for society, and presumably

earn its shareholders a lot of money in the process (Greenfield, 2011).

Google went public in 2004, and its first score is 1. But there is a jump in scores in subsequent years, as Google announced its 1% profit-for-social responsibility program soon after going public (Strom and Helft, 2011). The mantra of Google's founders is "don't be evil," and Google's KLD scores seem to be correlated with both the public image and programs that Google promote. The important take-away from both of these cases is that the KLD scores reflect not only the differences between these two companies, but also changes in firm goodness over time.

Third, in analysis which we omit for brevity, we gather donation data from the Chronicles of Philanthropy for approximately 100 large firms each year, chosen from Fortune magazine's list of top revenue-producing firms in the US. The sample covers 2000-2011. Donations data covers cash donations as well as product donations (e.g., donating drugs if the company is in the pharmaceutical industry). We find, for this limited sample, that equal-weighted KLD scores predict donations well in annual levels. The median firm in our sample donates around 30-40 million dollars a year. A regression of the logarithm of donations on KLD scores yields a coefficient of 0.0728 with a statistical significance level of around 5%. Firm-level annual changes in KLD scores are also positively correlated with annual changes in giving, but this relationship is weaker and insignificantly different from zero. The examples of Apple and Google above illustrate that KLD scores can change quite dramatically when there are changes in management or investment strategy and suggest that KLD scores capture more timely variation in overall firm CSR behavior than do donations dollars alone.[8]

Fourth, KLD scores are widely used by socially responsible investment (SRI) funds to screen out irresponsible companies from their indexes. SRI funds typically own stocks with the highest KLD scores within an industry. Additionally, Hong and Kostovetsky (2012) find that money managers of non-SRI funds who have, on net, contributed towards Democratic candidates in elections, and whose political values are thus likely to favor social responsibility, tilt their portfolios toward firms with the highest KLD scores within industries. In follow-up work, DiGiuli and Kostovetsky (2011) find that firms with Democratic CEOs are also more likely

[8]Indeed, Card, Hallock, and Moretti (2010) point to donations by company employees as being more important than donations by the corporations themselves.

to have higher KLD scores. In sum, KLD scores, while imperfect, are correlated with and informative about corporate social responsibility.

5. Prediction 1: The 2003 Dividend Tax Cut and CSR

To test Prediction 1, we compare changes in corporate goodness scores before and after the 2003 Dividend Tax Cut across portfolios of firms with different tiers of managerial ownership. We focus on firms in the S&P 500 in 2001 and 2002, just before the tax cut was enacted in 2003. In both 2001 and 2002, we compute the fractional insider share ownership for the top five executives and take an average over both years in order to smooth temporal irregularities.

Following Chetty and Saez (2005), we focus on share holdings rather than option holdings for a few reasons. First, most options are not dividend protected, so that the effect of the Dividend Tax Cut is ambiguous (Lambert, Lanen, and Larcker, 1989). Second, firms grant options rather than shares for a host of reasons other than to improve managerial incentives, including tax treatment, strategic management of financial constraints and even accounting manipulation. The manipulation motive is especially pronounced during the early 2000's when we measure insider ownership in light of the number of options-related scandals during this period. A number of papers show that, during the period of 1996-2002, firms which granted large amounts of options were more likely to engage in back-dating by retroactively granting options when the firm's stock price was at a temporary low and before a sharp reversal (e.g., Lie, 2005). Hence, option grants may not predict better alignment between managers and shareholders in the same way as insider ownership stakes of shares.

Table 1 provides summary statistics for our firms. In Panel A, for the ranking period of 2001-2002, we report summary statistics for firms by low, medium and high ownership portfolios based on tercile breakpoints. The median fractional ownership in each portfolio is 0.001, 0.003 and 0.021, with tercile breakpoints of 0.002 and 0.006 dividing a distribution that reaches a maximum of 0.297. Unconditionally, the median fractional ownership stake was 0.3%, which translates to $27.8M in dollar ownership.

Panel B reports statistics on goodness scores for the pre-period of 1999-2002, while Panel C reports the analogous statistics for the post-period of 2003-2006. Prior to the Dividend Tax

Cut, there were dividend increases or initiations in 11.7% of firm-years from 1999 through 2002. After the Dividend Tax Cut in 2003 through 2006, this fraction was 21.1%. In other words, firms significantly increased dividends after the 2003 Dividend Tax Cut, consistent with Chetty and Saez (2005).[9]

Turning to firm goodness scores, we find strong empirical support for the basic implication of our agency model: a dividend tax cut should lead to a decline in goodness. The goodness score, defined as total strengths minus concerns, significantly drops on average after the 2003 Dividend Tax Cut. The mean score is 0.657 in the pre-period and is 0.333 in the post-period, which represents a drop of 14% of one standard deviation in pre-period goodness scores. Firm goodness decreased in 24% of firm-years in the pre-period compared to 31% in the post period. In other words, the decline in goodness was spread across many firms.

To address the possible endogeneity of concurrent changes or time trends driving this decrease in goodness, we use a difference-in-differences approach to test Prediction 1, that the magnitude of the change in goodness in response to the Dividend Tax Cut should be largest for firms with intermediate levels of managerial ownership and smaller for firms with low or high levels of managerial ownership.

Figure 2 plots the time series of goodness scores from 1999 through 2006 for the three portfolios sorted by managerial ownership. The figure reveals that from 1999 to 2002, medium-ownership firms had goodness scores similar to those of low and high-ownership firms, with slightly higher scores in the 2000-2002 period. From 2003 onwards, that gap narrows substantially, which suggests there may be an effect from the tax cut. However, the figure itself is inconclusive for two reasons. First, unobserved heterogeneity correlated with insider ownership such as size may be driving the changes rather than insider ownership itself. Second, there are time trends in the data, with goodness scores declining for all groups throughout this period.

We use regression analysis to tease out any effects more formally. In order to allow our estimates to capture long-run changes in goodness, we form an annual panel of firm goodness scores from four years prior to the tax cut to four years after the tax cut, including the year of the tax cut itself: 1999-2006. We allow for the effect of the tax cut to vary each year and

[9]Our dividend increase and initiation definitions follow that of Chetty and Saez (2005). For four firm-years observations with KLD scores, there was insufficient data to assess whether the firms increased dividends.

examine whether the cumulative effect of the tax cut for low and high ownership firms differs systematically from that of medium ownership firms. We estimate the following specification via OLS:

$$\begin{aligned} goodness_{it} &= \alpha_i + \beta_t + \sum_{s \neq 2002} \gamma_s \left[LowOwn_i \times 1_{\{t=s\}} \right] + \sum_{s \neq 2002} \delta_s \left[HighOwn_i \times 1_{\{t=s\}} \right] \\ &+ \sum_{s \neq 2002} \xi_s \left[LowMktCap_i \times 1_{\{t=s\}} \right] + \sum_{s \neq 2002} \lambda_s \left[HighMktcap_i \times 1_{\{t=s\}} \right] + u_{it}, \end{aligned} \tag{1}$$

where α_i are a series of firm fixed effects, β_t are a series of year effects, $LowOwn_i$ ($HighOwn_i$) is an indicator for whether firm i is in the low (high) ownership tercile j during 2001-2002, $LowMktCap_i$ ($HighMktCap_i$) is an indicator for whether firm i is in the low (high) market capitalization tercile j during 2001-2002, and $1_{\{t=s\}}$ represent indicators for whether year t equals s. We allow for the effect of the tax cut to vary heterogeneously through time by the market capitalization tercile of the firm since it is well-known that fractional insider ownership is negatively correlated with firm size. Flexible controls for firm size also help to soak up any potential variation associated with firms of different sizes attracting more or less scrutiny from investors in response to the tax cut. We take tercile 2 (medium) as the omitted reference tercile for both ownership and market capitalization. We take 2002, the year before the tax cut, as the omitted reference year (i.e., there is no β_{2002} term). Note that the level effects for the ownership and market capitalization are subsumed into the firm fixed effects. Firm fixed effects also control for any time-invariant industry or firm characteristics across firms.[10]

The coefficients $\{\beta_t\}$ represent the difference in goodness measured at the end of year t relative to goodness measured at the end of 2002, the year just prior to the tax cut, for medium ownership firms (who are also in the medium market capitalization tercile). The coefficients $\{\gamma_t\}$ represent the differences-in-differences of goodness at the end of year t minus goodness at the end of 2002 for low ownership firms minus medium ownership firms. Finally, the coefficients $\{\delta_t\}$ represent the differences-in-differences of goodness at the end of year t minus goodness at the end of 2002 for high ownership firms minus medium ownership firms.

The model predicts that $\gamma_t > 0$ and $\delta_t > 0$ for $t \geq 2003$. We can examine the assumption of

[10]Our specification is similar to that used in the literature evaluating the effect of job displacement or job training programs on earnings, for example in Jacobson, LaLonde, and Sullivan (1994).

no pre-trend differences by testing whether $\gamma_t = \delta_t = 0$ for $t < 2002$. We cluster standard errors by firm, which allows for heteroskedasticity across firms and within-firm serial correlation in u_{it}.

The first set of columns in Table 2 presents the results of this regression. We find strong empirical support for Prediction 1 of our agency model. Although the immediate change in goodness in 2003 is muted, the differences become larger through time. Medium ownership firms see scores fall throughout the tax cut period, as evidenced by the negative β_t coefficients after the tax cut. More importantly, this decrease in goodness is $\gamma_{2004} = 0.32$ more than the decrease in goodness for low ownership firms, a difference-in-difference that is statistically significant at the 10% level and represents a 0.14-standard deviation decrease in medium ownership firms' 2002 KLD scores.[11] High ownership firms cut $\delta_{2004} = 0.49$ goodness points less (0.22-standard deviations) over this period, an effect statistically significant at the 1% level.

These effects are largely persistent. The difference-in-difference for low ownership firms minus medium ownership firms from the end of 2002 through the end of 2006 remains positive, with $\gamma_{2006} = 0.23$, or 0.10-standard deviations. Over this same period, the difference-in-difference for high ownership firms minus medium ownership firms is $\delta_{2006} = 0.64$, or 0.29-standard deviations, which is statistically significant at the 1% level.

Table 2 also shows that there were few significantdifferences between the firms prior to the tax cut. The γ_t and δ_t coefficients are insignificant before 2002. These results are highlighted in Figure 3, which plots the sequence of γ_t and δ_t coefficients, scaled by the standard deviation of KLD scores among medium ownership firms in 2002, through time along with their 95% confidence bands. In general, there are few differences across the ownership terciles before the tax cut. To summarize, corporate goodness declines on average following the Dividend Tax Cut, but goodness declines the most among medium ownership firms.

5.1. Robustness

An alternative explanation for this pattern would need to explain both the divergence of medium ownership firms' goodness scores from high and low ownership firms, as well as the fact that

[11] We compute economic significance throughout this section by scaling coefficients by the cross-sectional standard deviation of KLD scores among medium ownership firms in 2002, which is 2.22.

this divergence appears during the time of the tax cut. The non-monotonicity of the predicted effect is particularly helpful in identifying the effects of the model, as a monotonic pattern in which higher ownership firms always cut more would be easier to rationalize with alternative time trend stories or with an omitted variable correlated with ownership.

However, non-linearities may appear spuriously due to the functional form of equation (1). Furthermore, the unobservability of the functional form of the manager's utility function makes it difficult to pin down exactly where non-linearities with respect to ex-ante managerial ownership should occur. As such, we consider a number of additional analyses to verify that this non-linear response: 1) is not specific to our use of tercile groupings, 2) holds non-parametrically across the distribution of ownership, 3) survives placebo tests, and that 4) our standard errors are appropriately sized to measure the probability of a Type I error.

First, the second set of columns in Table 2 examines whether our results hold across an expanded set of quintile ownership portfolios rather than terciles. If our model's prediction holds in the data, we should see even larger effects when comparing medium ownership firms to high and low ownership firms at further extremes. To test this, we first form portfolios based on quintiles of ownership, and then modify equation (1) so that the quintile 3 represents the medium ownership portfolio, quintile 1 (the lowest) represents the low ownership portfolio, and quintile 5 (the highest) represents the high ownership portfolio. We then add a full set of interactions between year effects and indicators for portfolio 2 and portfolio 4, although we omit reporting these effects for brevity. The differences-in-differences between low and medium ownership firms from 2002 through 2004 and 2005 were $\gamma_{2004} = 0.43$ and $\gamma_{2005} = 0.49$, or 0.19 and 0.22 standard deviations, respectively, and are significant at the 10% level. The effects between high and medium ownership were even larger, cumulating to a total difference-in-difference of $\delta_{2006} = 0.77$, or 0.35-standard deviations, from 2002 through 2006.

Second, we check that our results hold non-parametrically by fitting the relationship between changes in KLD and the percentile ranking of a firm's ex-ante managerial ownership using local polynomial regression techniques (see Fan and Gijbels, 1996 for background and an extensive review). Figure 4 plots, for each year t, the fitted relationship between the change in KLD from 2002 to t against the percentile ranking of ownership for each firm in our sample. Local polynomial smoothers involve many choices by the econometrician in terms of the kernel used,

its bandwidth, and the order of the polynomial, that trade off bias in the estimated conditional expectation with its variance. We plot results for four sets of choices, starting with the baseline of a local linear regression with an Epanechnikov kernel set at the data-driven "rule-of-thumb" (ROT) optimal constant bandwidth outlined in Fan and Gijbels (1996, Section 4.2).[12]

Overall, the plots show that, from 2003 onwards, a substantial U-shape develops in the previously-flat distribution of the change in KLD scores over the baseline year of 2002. An analysis that first residualizes changes in KLD scores and ownership percentiles against size deciles in a semi-parametric version of equation (1) also reveals a U-shape.

Third, we scramble the tax cut year but maintain the same ownership portfolio assignment methodology, with the idea that we should only observe results in 2003. We estimate equation (1) assuming counterfactual "placebo tax cuts" in every year before 2003. In general, because our goodness scores go back to 1991 and we examine the evolution of goodness across four-year pre- and post-periods, we are able estimate equation (1) assuming different counterfactual tax cut years beginning with 1995. For each one of these counterfactual experiments, we form ownership portfolios by averaging fractional ownership in the two years before the counterfactual tax cut year and estimate equation (1). Table 3 reports these results. We find no effects for any placebo tax cut year before 2003.

Fourth, we scramble our ownership portfolio assignment (within true tax cut year of 2003) to assess the likelihood that random portfolio assignment among the firms in our sample would have generated our results. If our standard errors are properly constructed, random portfolio assignment should only lead to rejections of the null hypothesis at the 5% confidence level with probability 0.05. However, Bertrand, Duflo, and Mullainathan (2004) find evidence that serial correlation leads to excessive Type I errors relative to the prescribed size of the test in difference-in-differences estimates when using traditional OLS standard errors. One solution recommended to handle this size distortion is the use of clustered standard errors. We follow their paper to check for whether there is any evidence of a size distortion in our tests of $\gamma_t \neq 0$ and $\delta_t \neq 0$ by estimating equation (1) with 10,000 bootstrapped samples where portfolios

[12]The optimal bandwidth is determined for each year t. We also check our results using bandwidths of 75% and 125% of the optimal bandwidth. Fan, Gasser, Gijbels, Brockmann, and Engel (1997) show the optimality of the Epanechnikov kernel, although we have also checked our results using Gaussian and biweight kernels. Finally, we check our results using a local cubic regression with constant ROT bandwidth. Fan and Gijbels (1995) and Ruppert and Wand (1994) show that it is optimal to use odd-ordered polynomials.

of ownership are randomly assigned. We form 10,000 random samples by drawing random vectors of firm goodness scores, with replacement, from our observed sample, where one vector represents observed goodness scores for one firm through time. For each random sample, 503 such vectors are drawn, which corresponds with the sample size of our main analysis. We then randomly assign ownership portfolios among these vectors, estimate equation (1), and tabulate the fraction of rejections at the 5% significance level for each coefficient for our 10000 random samples. Table 4, Panel A reports the results. We find that our actual rejection rates are very close to 5% and conclude that random portfolio assignment would have rejected the null hypothesis at the 5% confidence level only 5% of the time.[13]

Panel B performs a second robustness check suggested by Bertrand, Duflo, and Mullainathan (2004) by averaging goodness scores in the two years pre-tax-cut (2001-2002) and two years post-tax-cut (2003-2004) for each firm. Then we compare the average change across terciles of ownership. This averaging is a very conservative way to address the serial correlation problem. We find that, although statistical significant is more limited, medium ownership firms reduce goodness more relative to the low and high ownership terciles.

6. Prediction 2: Close Governance Votes and CSR

We now test Prediction 2 of our agency model: corporate goodness should decline following improvements in firm governance. In general, correlations between governance and goodness may reflect the influence of a third unobserved factor instead of a causal relationship. We isolate the causal effect of changes in governance on goodness using a regression discontinuity approach focusing on close governance proxy votes. The assumption is that firms close to either side of the 50% vote share cutoff for passage are identical, except that firms are sharply more likely to pass a governance measure when the vote share in favor of a proposal slightly exceeds 50%. We use these close votes to generate random assignment of governance among firms around the cut-off and test whether firms in which these proposals just passed experienced slower growth in goodness scores than firms in which the proposals just failed.

[13] As in Bertrand, Duflo, and Mullainathan (2004), we also perform this analysis where we do not sample with replacement and thus are simply assigning portfolios randomly among the 503 firms in our sample for 10000 replications. Results are nearly identical.

Implicit in our RD design is the assumption that the 50% cut-off generates meaningful exogenous variation in governance. Shareholder proposals are non-binding, so firms are not legally obligated to implement proposals that receive with majority support. However, boards face strong pressure to honor the voting cutoff. For example, Institutional Shareholder Services, the leading provider of corporate governance ratings and advice for institutional investors, recommends that investors vote to remove all existing board members (in a voting process that can be binding) if the board fails to act on a shareholder proposal that received the support of a majority of shares outstanding in the previous year.

The impact of the 50% cutoff on governance has also been established empirically in previous work by Cuñat, Gine, and Guadalupe (2010), who use close votes to estimate the effect of governance on firm value. They find a small but significant effect on firm value. They further estimate that passing a proposal around the discontinuity leads to a discrete 31% increase in the probability of implementation (measured as the change in the number of anti-takeover provisions the firm has in place). We replicate and confirm their findings in later tables. Moreover, narrowly passing a proposal also leads to an increase in the probability of proposing and passing future shareholder-initiated governance proposals. Hence, passing the proposal can affect governance through more than just passing a new governance provision. For example, shareholders seem to become more empowered and are more likely to be activist in the future.

In related work, Ertimur, Ferri, and Stubben (2010) find, using data from an earlier period, that passing a proposal increases the probability of implementation by 27% and that the effect of passing proposals on implementation has steadily increased over time, although the statistical significance of the upward trend is weak. Thomas and Cotter (2007) also look at implementation rates and show that over the 2002-2004 period, implementation rates for proposals with greater than 50% of vote share sharply increased over time.

Using these earlier studies as our backdrop, we apply this RD design to examine the relative growth rates of KLD scores around the 50% vote share cut-off. The agency perspective predicts lower growth rates to the right of the cut-off (firms with improved governance) compared to firms just to the left of the cut-off. We restrict our analysis to shareholder initiated, governance-related proposals, which have the dual advantages of being directly targeted at managerial agency problems and of being heavily contested, resulting in a higher density of vote shares

close to the 50 percent vote share cutoff. We also exclude proposals that directly target CSR. In theory, one could use close votes on proposals that directly seek to increase CSR to measure the effect of CSR on firm value. However, proposals that directly target CSR overwhelmingly receive weak shareholder support, with insufficient observations near the cutoff for passage for estimation. Therefore, we instead study how exogenous improvements in governance affect CSR investment.[14] We also distinguish between different subcategories of governance proxy proposals, following the categorization developed by Cuñat, Gine, and Guadalupe (2010).

First, we empirically support the validity of our design. The RD assumes that vote shares near the cutoff are not manipulated. Previous research by Listokin (2008) and Smith (2013) has shown that vote shares of *management* sponsored proposals may indeed be manipulated. However, the same research suggests that vote shares of *shareholder* sponsored proposals (which are the only types of proposals studied in this paper) are not subject to the same manipulation. This is because management can influence the voting behavior of shares held by brokers in the case of management sponsored proposals but not in the case shareholder sponsored proposals. We support these findings by plotting the distribution of vote shares in Figure 5. The x-axis represents vote share in excess of the passage cutoff of 50%. Therefore, $x = 0$ is the location of the discontinuity. The distribution of vote shares around the cut-off is smooth, suggesting that outcomes to either side of the cut-off are likely to be random as opposed to being manipulated.

In all, we have 5,262 governance-related shareholder proposals spanning the 1997-2011 time period. In our baseline analysis, we exclude proposals targeted at limiting managerial compensation because our model predictions with respect to these proposals are ambiguous. While improvements in governance are expected to reduce perk spending, limits on compensation may lead managers to substitute toward perk spending and private benefits. After removing proposals targeted at compensation, we are left with 3853 proposals. 157 of these are within 1 percentage point of the 50% cutoff, 545 are within 5 percentage points, and 1044 are within 10 percentage points. For our RD analysis, we use an optimal bandwidth around the discontinuity point to appropriately downweight votes that are farther away from the cutoff.

In Table 5, we further support the validity of the RD design by showing that firms on either

[14] In our sample, a regression discontinuity using close votes for proposals directly focused on CSR yields noisy and insignificant estimates.

side of the regression discontinuity do not differ significantly in terms of ex-ante characteristics.[15] We use a local linear regression discontinuity estimation in the style of Imbens and Lemieux (2007) by fitting local linear trends in vote share on either side of the regression discontinuity using the optimal bandwidth for each ex-ante outcome. We determine the optimal bandwidth around the discontinuity point using kernel methods developed in Imbens and Kalyanaraman (2009). KLD is the KLD score of firms in the year before the shareholder proposal is voted on. $Change\,in\,KLD$ is KLD in $t-1$ minus KLD in $t-2$. The others ex-ante outcomes are firm characteristics measured in $t-1$. In general, there are no significant jumps and firms appear comparable on either side of the boundary. This supports the RD assumption that firms near the vote share discontinuity point are similar on all dimensions except that firms just to the right of the discontinuity point are discretely more likely to experience improvements in governance.

Next, we estimate the effect of governance changes on measures of firm goodness using the same RD estimation method. Table 6 estimates our dependent variable of interest – the change in KLD scores $(t+1)-(t)$ around the regression discontinuity. We again follow Imbens and Lemieux (2007) and estimate the jump by fitting local linear trends in vote share on either side of the regression discontinuity using the optimal bandwidth. The first column in Panel A represents our baseline estimate. We find that improvements in governance lead firms to reduce their annual growth in corporate goodness by 0.6. Given that the standard deviation of the growth rate in corporate goodness is 1.6, this represents an economically significant decline of one-third of a standard deviation. The other columns in Panel A show that the estimates are robust to alternative bandwidth selections.

This discontinuity can be seen in Figure 6, where we plot the smoothed values for these changes around the cut-off. We look at votes between 40% (bounded by -10 on the x-axis) and 60% (bounded by the +10 on the x-axis). Indeed, there is an apparent jump in the KLD change to the left of the cut-off compared to the right. The magnitude is comparable to the one estimated in Table 6 and the magnitude of the jump at the vote share cutoff is significantly different from zero with a p-value of 0.035.

[15]To increase power, we include all proposals (including compensation-related proposals) in our sample. The results do not change if we limit the sample to non-compensation proposals.

In Panel B of Table 6, we look at how the effect of narrowly passing a proposal varies depending on the type of proposal. As previously discussed, compensation-related proposals seek to limit managerial pay. Board-related proposals target the membership and committee roles of the board of directors, for example, mandating an increase in the number of independent directors. Riskmetrics governance provisions measure a set of formal governance provisions measured by Riskmetrics, most of which relate to anti-takeover defenses. Voting-related proposals seek to increase the voting rights of shareholders. Of the 1062 proposal votes which are within 10% of the cutoff, 26% are compensation-related, 6% were board-related, 49% were Riskmetrics governance provision-related proposals, 14% were voting-related, and the remainder were other types of votes.

In general, analyses using sub-categories of proposals lead to noisy estimates that are not significant. The lower statistical significance is not uncommon for regression discontinuity designs which heavily weight the smaller set of observations in which the vote share in favor of each sub-category of proposal is very close to the cut-off. Nevertheless, the results suggest that our baseline estimates are not driven by only one particular kind of governance proposal. We find that narrowly passing board, Riskmetrics governance provisions, and voting related proposals leads to substantial, albeit noisily estimated, declines in the growth rate of goodness (all estimates exceed one-third of a standard deviation the growth rate of the goodness score). Meanwhile, narrowly passing compensation-related proposals lead to increases in goodness, consistent with the idea that managers may substitute toward goodness as a form of perk spending when their compensation is cut.

In Table 7, we support the robustness of our baseline results. The KLD ratings methodology was modified significantly in 2010. This should not bias our results because the ratings change should not differentially affect firms to the left and right of the 50% vote share cutoff. Nevertheless, we present results using alternative measures of goodness, which offer great continuity over time. The first column in the top panel measures KLD as a standardized z-score, calculated using all firms rated in each year. The second column measures KLD using subcategories of goodness that have been consistently rated over time. The third column measures KLD demeaned among all firms rated in each year. The last column uses the baseline KLD score, but with controls for the logarithms of firm assets, market cap, return, and total compensation in

the year of the shareholder proposal vote. Controlling for additional covariates is non-standard in RD estimation because the cutoff itself should generate random variation. Nevertheless, we find that the inclusion of additional covariates does not affect the results substantively.

In the second row, for columns (1) and (2), we consider two placebo tests for discontinuities in changes in goodness scores: at a placebo cutoff equal to the median of vote shares conditional on the vote share being less than the true cutoff and greater than the true cutoff, respectively. Reassuringly, both of these placebo estimates are insignificant and closer to zero in comparison to the estimated 0.64 point decline in goodness estimated using the true vote share cutoff.

Finally, in columns (3) and (4) of the second panel, we empirically support the assumption that narrow passage of shareholder proposals leads to improvements in governance. We focus on shareholder proposals that seek to alter formal governance provisions that are measured by Riskmetrics in their governance ratings. Most of these proposals seek to remove anti-takeover defenses. The threat of takeovers may provide a disciplining role on managers (Bertrand and Mullainathan, 2003), so the removal of anti-takeover defenses likely represents an improvement in governance. We measure the quality of governance provisions using the entrenchment index, developed in Bebchuk, Cohen, and Ferrell (2009).[16] Consistent with previous findings in Cuñat, Gine, and Guadalupe (2010), Ertimur, Ferri, and Stubben (2010), and Thomas and Cotter (2007), we find that narrow passage of shareholder proposals significantly reduces entrenchment and therefore correspond to improvements in governance.

An important consideration when interpreting the results of the regression discontinuity is that not all shareholders necessarily wish to improve governance,[17] so narrow passage of a subset of shareholder proposals may not push executives toward firm value creation. In this analysis, we implicitly assume that the average governance-related shareholder proposal around the cutoff

[16]There is debate in the literature regarding best practices in terms of measuring governance provisions. Bebchuk, Cohen, and Ferrell (2009) focus on the six governance provisions that are most predictive of firm value: staggered boards, limits to shareholder bylaw amendments, poison pills, golden parachutes, and supermajority requirements for mergers and charter amendments. These are the same six provisions that are most frequently referenced in our sample of shareholder proposals. Gompers, Ishii, and Metrick (2003) develop an earlier measure of governance provisions (GIM index) that represents the sum of 24 possible governance provisions. We also test whether narrow passage of shareholder proposals affect the GIM index and find similar results, but which are insignificantly different from zero.

[17]For instance, Barber (2007) and Heinkel, Kraus, and Zechner (2001) point out that institutions such as CalPERS pursue social activism by pushing firms to adopt CSR, and that this activism is not necessarily consistent with shareholder value maximization.

is intended to reduce agency problems. We believe this assumption is valid because of previous work, e.g., Cuñat, Gine, and Guadalupe (2010) showing that, on average, proposals near the vote share cutoff increase firm value. Moreover, while some shareholders may seek objectives other than value maximization, it is less likely that a significant fraction of pivotal shareholders would vote to reduce value-enhancing corporate goodness (which would be necessary to explain our results). Nevertheless, we caution readers that we are measuring a local average treatment effect which may mask heterogeneity in the goals of specific shareholder proposals.

We also caution that the magnitude of our regression discontinuity measure may be specific to our sample period, even if the qualitative direction of our results hold more generally. Most of the close votes in our sample occur between 2008 and 2011. Following the financial crisis, there were many shareholder-initiated proposals. During this contentious period, management may have been particularly responsive to improvements in governance as represented by shareholder proposals. As a result, one may not wish to extrapolate the magnitude of our measure too aggressively into other time periods.

Finally, it is important to note that our model takes a simple, reduced form, view of how agency problems can motivate unproductive CSR: goodness increases the manager's private benefits, so managers will spend less on goodness when ownership stakes increase or governance improves. However, along the lines suggested by Zwiebel (1996) and Kuhnen and Zwiebel (2008), goodness spending may also be a way for management to entrench themselves by currying favor with other stakeholders such as employees who might help managers resist shareholder discipline. Such an entrenchment motive for goodness is also agency driven but outside our simple model.[18]

7. Conclusion

We find that improvements in managerial incentives and governance lead to a reduction in firm goodness, implying that the marginal dollar spent on goodness is a result of agency problems.

[18]It is possible that managers at firms to the left of the vote share cutoff (where proposals just fail) increase goodness in order to entrench themselves and prevent future shareholder activism while managers to the right of the cutoff (where proposals just pass) reduce unproductive CSR because of improved monitoring. This would lead to a decline in goodness across the vote share cutoff that is consistent with an broad agency-related interpretation CSR, but in a more nuanced way than specified in our model.

It is important to highlight that some forms of goodness investment, not on the margin, may increase firm value. However, managers *overspend* on goodness because they wish to do good with other people's money. As we alluded to in the introduction, our findings link traditional agency problems with the burgeoning literature on social preferences and altruism. Much more work remains to be done on whether this is being driven by social pressure or a purer altruistic motive, as considered in DellaVigna, List, and Malmendier (2012). The economic analysis of corporate social responsibility is likely to become even more important in the future, as society's concerns about global warming or income inequality are likely to confront the traditional doctrine of shareholder value maximization.

References

Andreoni, J., 1989, "Giving with Impure Altruism: Applications to Charity and Ricardian Equivalence," *Journal of Political Economy*, 97(1), 1447–1458.

Auerbach, A. J., and K. Hassett, 2006, "Dividend Taxes and Firm Valuation: New Evidence," *American Economic Review*, 96(2), 119–123.

——— , 2007, "The 2003 Dividend Tax Cuts and the Value of the Firm: An Event Study," in *Taxing Corporate Income in the 21st Century*, ed. by A. J. Auerbach, J. Hines, and J. Slemrod. pp. 93–126.

Barber, B., 2007, "Monitoring the Monitor: Evaluating Calpers' Activism," *Journal of Investing*, pp. 66–80.

Baron, D. P., 2001, "Private Politics, Corporate Social Responsibility, and Integrated Strategy," *Journal of Economics & Management Strategy*, 10(1), 7–45.

Bebchuk, L., A. Cohen, and A. Ferrell, 2009, "What Matters in Corporate Governance?," *Review of Financial Studies*, 22(2), 783–827.

Becker, G., 1974, "A Theory of Social Interaction," *Journal of Political Economy*, 82, 1063–1093.

Benabou, R., and J. Tirole, 2010, "Individual and Corporate Social Responsibility," *Economica*, 77, 1–19.

Bertrand, M., E. Duflo, and S. Mullainathan, 2004, "How Much Should We Trust Differences-in-Differences Estimates?," *Quarterly Journal of Economics*, 119(1), 249–275.

Bertrand, M., and S. Mullainathan, 2003, "Enjoying the Quiet Life? Corporate Governance and Managerial Preferences," *Journal of Political Economy*, 111(5), 1043–1075.

Besley, T., and M. Ghatak, 2005, "Competition and Incentives with Motivated Agents," *American Economic Review*, 95(3), 616–636.

Card, D., K. F. Hallock, and E. Moretti, 2010, "The geography of giving: The effect of corporate headquarters on local charities," *Journal of Public Economics*, 94(3-4), 222–234.

Charness, G., and M. Rabin, 2002, "Understanding Social Preferences With Simple Tests," *The Quarterly Journal of Economics*, 117(3), 817–869.

Chatterji, A. K., D. I. Levine, and M. W. Toffel, 2009, "How Well Do Social Ratings Actually Measure Corporate Social Responsibility?," *Journal of Economics & Management Strategy*, 18(1), 125–169.

Chetty, R., and E. Saez, 2005, "Dividend Taxes and Corporate Behavior: Evidence from the 2003 Dividend Tax Cut," *The Quarterly Journal of Economics*, 120(3), 791–833.

——— , 2010, "Dividend and Corporate Taxation in an Agency Model of the Firm," *American Economic Journal: Economic Policy*, 2(3), 1–31.

Cuñat, V., M. Gine, and M. Guadalupe, 2010, "The Vote is Cast: The Effect of Corporate Governance on Shareholder Value," NBER Working Papers 16574, National Bureau of Economic Research, Inc.

DellaVigna, S., J. A. List, and U. Malmendier, 2012, "Testing for Altruism and Social Pressure in Charitable Giving," *The Quarterly Journal of Economics*, 127(1), 1–56.

DiGiuli, A., and L. Kostovetsky, 2011, "Are Red or Blue Companies More Likely to Go Green? Politics and Corporate Social Responsibility," *SSRN Working Paper no. 2084900*.

Elfenbein, D. W., R. Fisman, and B. McManus, 2012, "Charity as a Substitute for Reputation: Evidence from an Online Marketplace," *Review of Economic Studies*, 79(4), 1441–1468.

Ertimur, Y., F. Ferri, and S. R. Stubben, 2010, "Board of directors' responsiveness to shareholders: Evidence from shareholder proposals," *Journal of Corporate Finance*, 16(1), 53–72.

Fan, J., T. Gasser, I. Gijbels, M. Brockmann, and J. Engel, 1997, "Local Polynomial Regression: Optimal Kernels and Asymptotic Minimax Efficiency," *Annals of the Institute of Statistical Mathematics*, 49(1), 79–99.

Fan, J., and I. Gijbels, 1995, "Adaptive Order Polynomial Fitting: Bandwidth Robustification and Bias Reduction," *Journal of Computational and Graphical Statistics*, 4(3), 213–227.

——— , 1996, *Local Polynomial Modelling and Its Applications*. Chapman & Hall, London.

Fehr, E., and K. Schimdt, 1999, "A Theory of Fairness, Competition and Cooperation," *The Quarterly Journal of Economics*, 114(3), 817–868.

Friedman, M., 1970, "The Social Responsibility of Business Is To Increase Its Profits," *New York Times Magazine*, pp. 122–126.

Gompers, P., J. Ishii, and A. Metrick, 2003, "Corporate Governance And Equity Prices," *The Quarterly Journal of Economics*, 118(1), 107–155.

Greenfield, R., 2011, "Yes, You Can Think Less of Steve Jobs for Not Being a Philanthropist," *The Atlantic Wire*, Available: http://www.theatlanticwire.com/technology/2011/08/yes-you-can-think-less-steve-jobs-not-being-philanthropist/41885/ [Last accessed: July 2013].

Heal, G., 2005, "Corporate Social Responsibility—Economic and Financial Perspectives," *Geneva Papers*, 30, 387–409.

Heinkel, R., A. Kraus, and J. Zechner, 2001, "The Effect of Green Investment on Corporate Behavior," *Journal of Financial and Quantitative Analysis*, 36(04), 431–449.

Hong, H., and L. Kostovetsky, 2012, "Red and Blue Investing: Values and Finance," *Journal of Financial Economics*, 103(1), 1–19.

Hong, H., J. D. Kubik, and J. A. Scheinkman, 2011, "Financial Constraints on Corporate Goodness," Working papers, Princeton University.

Imbens, G., and T. Lemieux, 2007, "Regression Discontinuity Designs: A Guide to Practice," NBER Technical Working Papers 0337, National Bureau of Economic Research, Inc.

Imbens, G. W., and K. Kalyanaraman, 2009, "Optimal Bandwidth Choice for the Regression Discontinuity Estimator," IZA Discussion Papers 3995, Institute for the Study of Labor (IZA).

Jacobson, L. S., R. J. LaLonde, and D. G. Sullivan, 1994, "Earnings Losses of Displaced Workers," *American Economic Review*, 83(4), 685–789.

Jensen, M. C., and W. H. Meckling, 1976, "Theory of the firm: Managerial behavior, agency costs and ownership structure," *Journal of Financial Economics*, 3(4), 305–360.

Kruger, P., 2013, "Corporate Goodness and Shareholder Wealth," Working papers, University of Geneva.

Kuhnen, C. M., and J. Zwiebel, 2008, "Executive Pay, Hidden Compensation and Managerial Entrenchment," *SSRN Working Paper no. 972622*.

Lambert, R. A., W. N. Lanen, and D. F. Larcker, 1989, "Executive Stock Option Plans and Corporate Dividend Policy," *Journal of Financial and Quantitative Analysis*, 24(04), 409–425.

Lie, E., 2005, "On the Timing of CEO Stock Option Awards," *Management Science*, 51(5), 802–812.

List, J. A., 2006, "The Behavioralist Meets the Market: Measuring Social Preferences and Reputation Effects in Actual Transactions," *Journal of Political Economy*, 114(1), 1–37.

Listokin, Y., 2008, "Management Always Wins the Close Ones," *American Law and Economics Review*, 10(2), 159–184.

Poterba, J., 2004, "Taxation and Corporate Payout Policy," *American Economic Review*, 94(2), 171–175.

Ruppert, D., and M. P. Wand, 1994, "Multivariate Locally Weighted Least Squares Regression," *The Annals of Statistics*, 22(3), 1346–1370.

Smith, E. E., 2013, "Do Shareholders Want Less Governance?," Working papers, New York University.

Strom, S., and M. Helft, 2011, "Google Finds It Hard to Reinvent Philanthropy," *The New York Times*, Available: http://www.nytimes.com/2011/01/30/business/30charity.html [Last accessed: July 2013].

Thomas, R. S., and J. F. Cotter, 2007, "Shareholder proposals in the new millennium: Shareholder support, board response, and market reaction," *Journal of Corporate Finance*, 13(2-3), 368–391.

Tirole, J., 2001, "Corporate Governance," *Econometrica*, 69(1), 1–35.

Zwiebel, J., 1996, "Dynamic Capital Structure under Managerial Entrenchment," *American Economic Review*, 86(5), 1197–1215.

Table 1: Summary Statistics

Panel A provides summary statistics for the cross-section of S&P 500 firms in 2001-2002 used in our analysis. For each variable, we first average over the two years 2001-2002 for each firm when possible and then report summary statistics by terciles of average insider ownership in the cross-section. Panels B and C provide summary statistics in an annual panel of these firms from 1999 through 2006.

Variable	Mean	SD	Min.	Max.	Median	N
Panel A: Ranking Period Cross-Sectional Statistics (Firms, 2001-2002 Average)						
Low Ownership						
Insider Ownership	0.001	0.000	0.000	0.002	0.001	168
Market Capitalization ($B)	31.6	51.4	1.2	366.7	11.0	168
Medium Ownership	0.000	0.000	0.000	0.000	0.000	0
Insider Ownership	0.003	0.001	0.002	0.006	0.003	168
Market Capitalization ($B)	12.9	23.6	0.5	225.5	7.3	168
High Ownership	0.000	0.000	0.000	0.000	0.000	0
Insider Ownership	0.044	0.056	0.006	0.297	0.021	167
Market Capitalization ($B)	13.3	31.0	0.6	311.9	5.7	167
Panel B: Pre-Period (Firm-Years, 1999-2002)						
Dividend Increase or Initiation Indicator	0.117	0.321	0	1	0	1826
Goodness Score (Total Strengths - Concerns)	0.657	2.317	-7	9	1	1829
Change in Goodness Score < 0 Indicator	0.239	0.427	0	1	0	1829
Panel C: Post-Period (Firm-Years, 2003-2006)						
Dividend Increase or Initiation Indicator	0.211	0.408	0	1	0	1842
Goodness Score (Total Strengths - Concerns)	0.333	2.437	-7	10	0	1842
Change in Goodness Score < 0 Indicator	0.309	0.462	0	1	0	1842

Table 2: Dividend Tax Cut Difference-in-Differences

This table reports the results from a difference-in-differences estimation in the annual panel where the dependent variable is the level of the goodness score. The right-hand side variables are firm-specific fixed effects, year-specific effects, and interaction effects between year effects and indicators for all ownership portfolios. The first set of columns report results where we group ownership by terciles while the second set of columns report results where we group ownership by quintiles. Ownership portfolios are formed based on average ownership in the two-years prior to the shock year. The omitted category is always the medium ownership portfolio (tercile 2 or quintile 3). We also control for a set of market capitalization portfolio indicators (tercile or quintile) fully interacted with all year effects. The omitted year among market capitalization portfolios is the middle portfolio. The omitted year is 2002, and the coefficients are cumulative changes from the end of 2002 through the end of year *t*. In the quintile results, we do not report the interaction effects for ownership portfolios 2 and 4 for brevity. Standard errors clustered by firm are presented in brackets. */**/*** indicates significant at the 10%, 5% and 1% levels, respectively.

	Terciles			Quintiles		
	Year	Interaction Effects		Year	Interaction Effects	
	Effects	Low Own.	High Own.	Effects	Low Own.	High Own.
Year *t:*	β(t)	γ(t)	δ(t)	β(t)	γ(t)	δ(t)
1999	0.074	0.129	0.336	0.113	0.021	0.389
	[0.47]	[0.68]	[1.90]*	[0.50]	[0.08]	[1.78]*
2000	0.312	0.063	0.112	0.261	-0.044	0.114
	[2.67]***	[0.41]	[0.79]	[1.78]*	[-0.20]	[0.70]
2001	-0.030	0.112	0.123	-0.067	0.045	0.103
	[-0.34]	[0.99]	[1.33]	[-0.53]	[0.27]	[0.89]
2003	-0.183	0.095	0.099	-0.158	0.194	0.202
	[-1.88]*	[0.73]	[0.85]	[-1.07]	[1.07]	[1.31]
2004	-0.405	0.320	0.489	-0.246	0.427	0.548
	[-2.88]***	[1.90]*	[3.17]***	[-1.28]	[1.85]*	[2.68]***
2005	-0.736	0.243	0.590	-0.812	0.485	0.698
	[-4.35]***	[1.13]	[3.07]***	[-3.24]***	[1.72]*	[2.83]***
2006	-0.849	0.227	0.635	-0.959	0.368	0.774
	[-4.32]***	[0.93]	[2.86]***	[-3.52]***	[1.14]	[2.74]***
N		3671			3671	
R-Squared		0.054			0.058	
Firms		503			503	

Table 3: Placebo Difference-in-Differences

This table reports the results from a difference-in-differences estimation in the annual panel where the dependent variable is the level of the goodness score, but where we analyze different placebo shock years. The right-hand side variables are firm-specific fixed effects, year-specific effects, and interaction effects between year effects and indicators for ownership tercile portfolios. Ownership terciles are formed based on average ownership in the two years prior to the placebo shock year. The omitted category is always the medium ownership portfolio, tercile 2, and we include a set of market capitalization tercile portfolio indicators fully interacted with all year effects. The omitted year is the year before the placebo shock year, and the coefficients are cumulative changes from the end of the omitted year through the end of year *t*. Standard errors clustered by firm are presented in brackets. */**/*** indicates significant at the 10%, 5% and 1% levels, respectively.

Panel A: Placebo Tax Cut Years 1995-1998

	Base (t=0): 1994 Shock (t=+1): 1995		Base (t=0): 1995 Shock (t=+1): 1996		Base (t=0): 1996 Shock (t=+1): 1997		Base (t=0): 1997 Shock (t=+1): 1998	
	Interaction Effects		Interaction Effects		Interaction Effects		Interaction Effects	
	Low Own.	High Own.	Low Own.	High Own.	Low Own.	High Own.	Low Own.	High Own.
Year *t:*	γ(t)	δ(t)	γ(t)	δ(t)	γ(t)	δ(t)	γ(t)	δ(t)
-3	0.112	-0.130	0.193	0.006	0.016	0.072	-0.072	-0.015
	[0.52]	[-0.71]	[0.92]	[0.03]	[0.07]	[0.38]	[-0.35]	[-0.08]
-2	0.260	0.077	0.247	-0.024	-0.037	0.183	-0.132	-0.256
	[1.34]	[0.52]	[1.27]	[-0.15]	[-0.19]	[1.07]	[-0.74]	[-1.60]
-1	0.129	-0.007	0.072	0.016	-0.232	-0.083	0.027	-0.194
	[0.82]	[-0.05]	[0.49]	[0.12]	[-1.49]	[-0.68]	[0.22]	[-1.70]*
+1	0.050	0.003	0.021	-0.033	-0.038	0.014	-0.175	-0.047
	[0.33]	[0.02]	[0.14]	[-0.27]	[-0.29]	[0.12]	[-1.27]	[-0.39]
+2	0.063	-0.053	0.050	0.083	-0.158	0.036	-0.087	-0.019
	[0.32]	[-0.31]	[0.26]	[0.51]	[-0.91]	[0.25]	[-0.49]	[-0.13]
+3	0.066	0.080	-0.080	0.123	-0.032	0.133	0.286	-0.055
	[0.29]	[0.39]	[-0.36]	[0.68]	[-0.16]	[0.78]	[1.29]	[-0.29]
+4	-0.128	0.005	-0.156	0.067	0.139	-0.118	0.234	-0.143
	[-0.52]	[0.02]	[-0.61]	[0.33]	[0.60]	[-0.58]	[0.94]	[-0.66]
N	3557		3544		3468		3457	
R-Squared	0.035		0.027		0.018		0.022	
Firms	478		486		484		492	

Table 3, Continued

Panel B: Placebo Tax Cut Years 1999-2002

	Base (t=0): 1998 Shock (t=+1): 1999		Base (t=0): 1999 Shock (t=+1): 2000		Base (t=0): 2000 Shock (t=+1): 2001		Base (t=0): 2001 Shock (t=+1): 2002	
	Interaction Effects		Interaction Effects		Interaction Effects		Interaction Effects	
	Low Own.	High Own.	Low Own.	High Own.	Low Own.	High Own.	Low Own.	High Own.
Year *t:*	γ(t)	δ(t)	γ(t)	δ(t)	γ(t)	δ(t)	γ(t)	δ(t)
-3	-0.156	-0.300	0.226	-0.098	0.203	0.109	-0.116	0.125
	[-0.75]	[-1.62]	[1.16]	[-0.56]	[0.92]	[0.62]	[-0.59]	[0.69]
-2	0.156	-0.243	0.256	0.029	-0.042	0.022	-0.114	0.226
	[0.97]	[-1.65]*	[1.42]	[0.19]	[-0.23]	[0.15]	[-0.65]	[1.45]
-1	0.245	0.113	-0.062	-0.066	-0.043	0.066	0.039	0.159
	[1.79]*	[0.92]	[-0.51]	[-0.62]	[-0.29]	[0.61]	[0.31]	[1.40]
+1	0.081	-0.011	0.229	0.097	-0.056	-0.152	-0.063	-0.085
	[0.65]	[-0.11]	[1.47]	[0.84]	[-0.44]	[-1.27]	[-0.55]	[-0.88]
+2	0.307	0.048	0.276	0.002	-0.059	-0.151	0.099	0.062
	[1.62]	[0.30]	[1.47]	[0.01]	[-0.37]	[-1.10]	[0.60]	[0.41]
+3	0.316	-0.122	0.195	0.015	-0.065	-0.156	0.118	0.158
	[1.49]	[-0.64]	[0.95]	[0.08]	[-0.33]	[-0.89]	[0.60]	[0.87]
+4	0.209	-0.117	0.095	0.035	-0.051	-0.063	-0.040	0.197
	[0.91]	[-0.57]	[0.40]	[0.17]	[-0.22]	[-0.31]	[-0.17]	[0.89]
N	3454		3418		3531		3599	
R-Squared	0.031		0.040		0.041		0.045	
Firms	499		493		509		502	

Table 4: Robustness in Dividend Tax Cut

Panel A presents rejection rates where we repeat our experiment with 2003 as the shock year but bootstrap placebo treatment portfolios as follows. For each replication, we form a random sample by first drawing random vectors of firm goodness scores from our primary sample with replacement, where one vector represents observed goodness scores for one firm through time. For each replication, 503 such vectors are drawn. We then assign ownership portfolios randomly among these vectors, estimate equation (1), and tabulate the fraction of rejections at the 5% significance level for each coefficient for 10000 such replications. Panel B presents the change in goodness scores for firms in ownership Terciles 1 and 3 relative to Tercile 2. For each firm, we average scores in the 2003-2004 period and 2001-2002 period before computing the change in this average. We then project these changes in the cross-section onto indicators for low and high ownership portfolios. The coefficients represent the average difference in these changes relative to Tercile 2, the omitted category. White-heteroskedasticity robust standard errors are presented in brackets. */**/*** indicates significant at the 10%, 5% and 1% levels, respectively.

Panel A: Rejection Rates for Random Placebo Treatment Groups

	Terciles		Quintiles	
	Interaction Effects		Interaction Effects	
Year *t*:	Low Tercile	High Tercile	Lowest Quintile	Highest Quintile
-3	0.0507	0.0508	0.0515	0.0528
-2	0.0502	0.0516	0.0473	0.0509
-1	0.0488	0.0506	0.0453	0.0508
+1	0.0469	0.0466	0.0518	0.05
+2	0.0518	0.0471	0.0504	0.0475
+3	0.0478	0.0496	0.0492	0.0529
+4	0.0508	0.0507	0.0513	0.0496
Replications	10000		10000	

Panel B: Average Change in Average Scores, 2001-2002 through 2003-2004

Ownership:	
Low	0.161
	[1.14]
High	0.212
	[1.67]*
Constant	-0.209
	[-2.22]**
N	482
R-Squared	0.006

Table 5: Regression Discontinuity Validity Analysis

This table supports the validity of the RD design by showing that firms on either side of the regression discontinuity do not differ significantly in terms of ex-ante characteristics. We use local linear methods with an optimal bandwidth to estimate the discontinuous jump in the ex-ante outcome. KLD is the goodness score of firms in the year before the shareholder proposal is voted on. Change in KLD is the goodness score in t-1 minus the goodness score in t-2. The others variables are firm characteristics measured in t-1. Standard errors are presented in brackets. */**/*** indicates significant at the 10%, 5% and 1% levels, respectively.

	KLD	Change in KLD	Log Assets	Log Market Cap	Firm Return	Log Total Compensation
Pass	0.004	-0.181	-0.322	-0.207	0.0310	0.024
	[0.407]	[0.407]	[0.335]	[0.313]	[0.060]	[0.145]
Bandwidth	4.946	4.946	4.005	4.481	2.306	3.731
Obs	4864	4864	5182	5259	5253	4867

Table 6: Baseline Regression Discontinuity Analysis

This table estimates the change in KLD scores (t+1) - (t) around the regression discontinuity. We again follow Imbens and Lemieux (2010) and estimate the jump by fitting local linear trends in vote share on either side of the regression discontinuity using the optimal bandwidth as well as multiples of the optimal bandwidth. In Panel B, we present estimates by type of proposals following the categorization in Cunat, Gine, and Guadalupe (2011). Standard errors are presented in brackets. */**/*** indicates significant at the 10%, 5% and 1% levels, respectively.

Panel A: All Non-Compensation Governance Proposals

	Optimal Bandwidth	75% Optimal Bandwidth	125% Optimal Bandwidth	50% Optimal Bandwidth	150% Optimal Bandwidth
Pass	-0.635	-0.700	-0.543	-0.771	-0.440
	[0.302]**	[0.349]**	[0.349]**	[0.444]*	[0.251]*
Bandwidth	4.436	3.327	5.545	2.218	6.654
Obs	3412	3412	3412	3412	3412

Panel B: Variation by Proposal Type

	Compensation	Board	Riskmetrics Governance Provisions	Voting	All Other
Pass	1.052	-1.475	-0.473	-0.432	-1.811
	[0.581]	[0.927]	[0.371]	[0.592]	[1.394]
Bandwidth	4.847	6.51	4.889	5.731	7.001
Obs	1316	652	1780	412	568

Table 7: Robustness

This table supports the robustness of our baseline results. All outcomes are annual changes between periods (t+1) and period (t). The KLD ratings methodology was modified in 2009. While this change should not differentially affect firms to the left and right of the 50% vote share cutoff, we present results using alternative measures of KLD, which offer great continuity over time. The first column in the top panel measures KLD as a standardized z-score, calculated using all firms rated in each year. The second column measures KLD using only subcategories of goodness that have been consistently rated over time. The third column measures KLD demeaned among all firms rated in each year. The fourth column uses the baseline KLD score, but controls for firm log assets, log market cap, return, and log total compensation in the year of the shareholder proposal vote. The two placebo cutoffs test for jumps in changes in KLD scores (t+1) - (t), at a placebo cutoff equal to the median of vote shares conditional on vote share being less than the true cutoff and greater than the true cutoff. The last two columns in the lower panel explore whether narrow passage of shareholder proposals directly affects governance, as measured using the entrenchment index described in Bebchuk, Cohen, and Ferrell (2009). Because the entrenchment index rating methodology changed in 2007, we present results using the index prior to 2007 as well as an annual standardized index rating covering the full sample period. Standard errors are presented in brackets. */**/*** indicates significant at the 10%, 5% and 1% levels, respectively.

	KLD standardized annually	KLD adjusted measure	KLD demeaned by year	RD with covariate controls
Pass	-0.424	-0.595	-0.667	-0.585
	[0.211]**	[0.307]*	[0.315]**	[0.309]*
Bandwidth	3.613	3.49	4.098	4.436
Obs	3412	3412	3412	3412

	Placebo: Cutoff at Lower Median	Placebo: Cutoff at Upper Median	Entrenchment Index 1997- 2006	Entrenchment Index 1997-2011
Pass	-0.251	0.157	-0.481	-0.329
	[0.324]	[0.301]	[0.199]**	[0.165]**
Bandwidth	3.813	4.36	2.481	2.477
Obs	3412	3412	1315	1652

Figure 1: KLD Scores for Apple, Google and the S&P 500

This figure plots the KLD score each year for Apple, Google, and the averages score for the S&P 500.

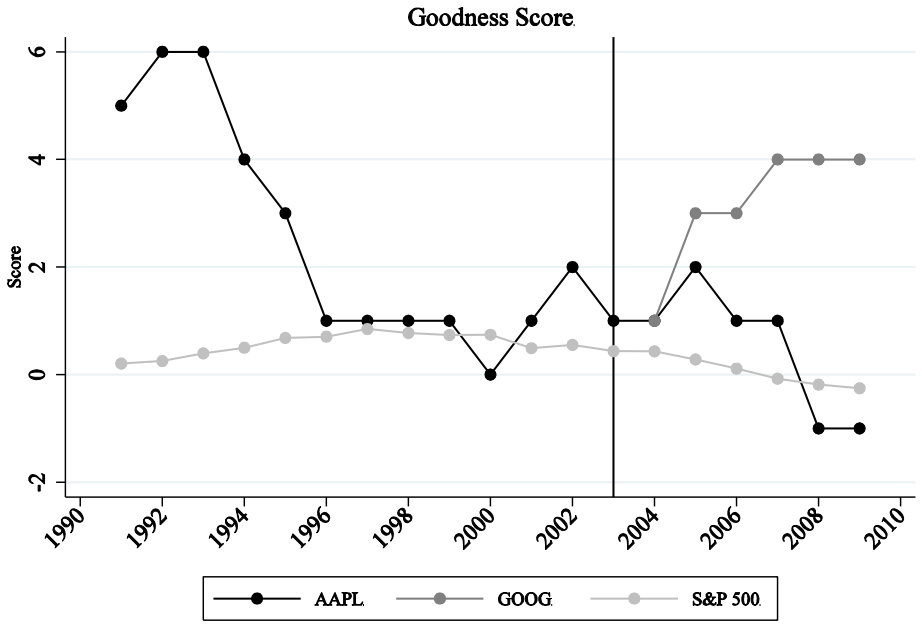

Figure 2: KLD Scores

This figure plots the average KLD score in each ownership tercile portfolio for each year 1999-2006 for our main sample of firms in the S&P 500 during 2001-2002.

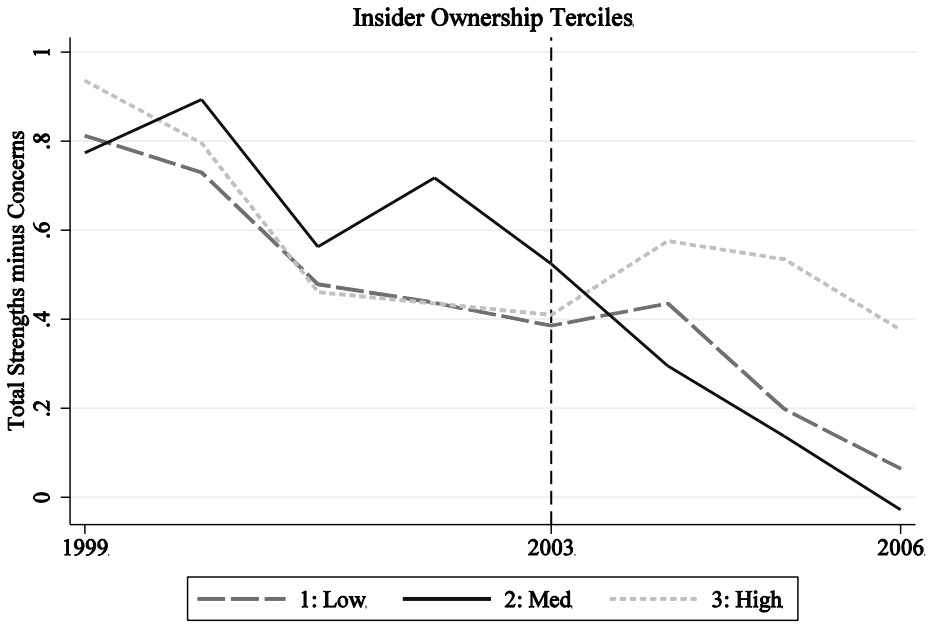

Figure 3: Dividend Tax Cut Differences-in-Differences

This figure plots the cumulative difference-in-differences of goodness scores for the low ownership portfolio minus the medium ownership portfolio in light grey, and high ownership minus medium ownership in dark grey, from Table 2. The coefficients are scaled by the standard deviation of medium ownership KLD scores in 2002. 95% confidence bands based on clustered standard errors are plotted as dashed lines. 2002 is taken as the base reference year (t=0) and 2003 is taken as the shock year (t=+1).

Panel A: Tercile Ownership Portfolios

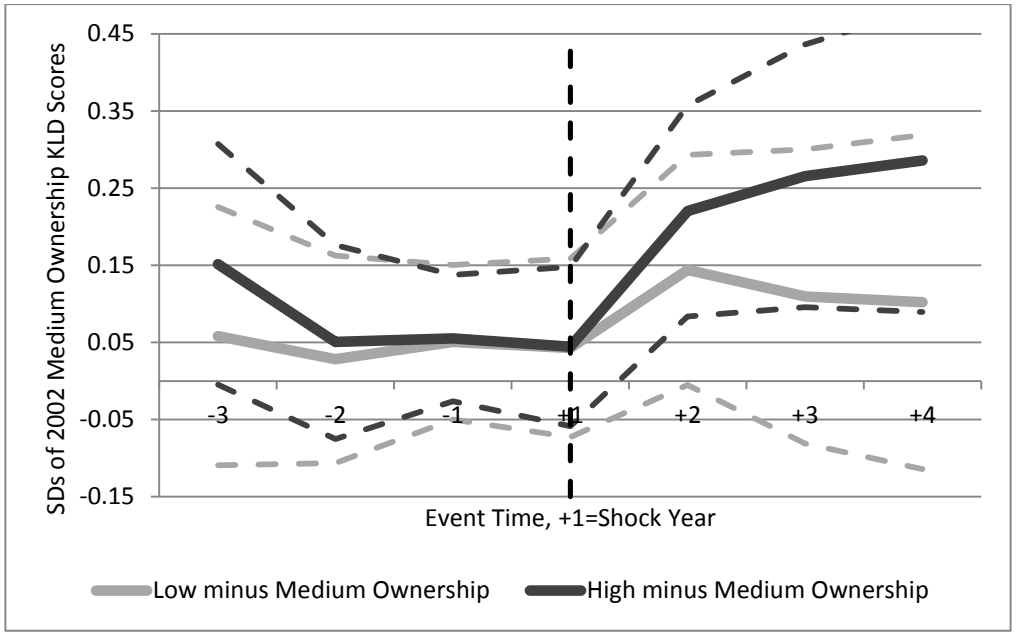

Panel B: Quintile Ownership Portfolios

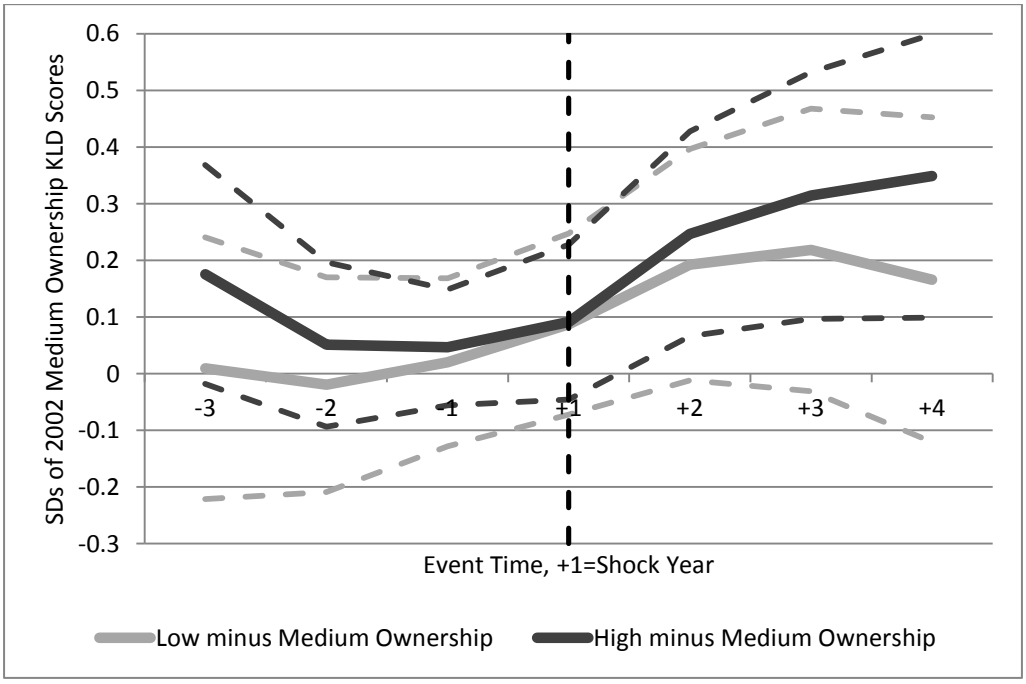

Figure 4: Non-Parametric Analysis of Dividend Tax Cut

This figure plots local polynomial fits of changes in KLD between 2002 and time *t* relative to ownership percentiles. The black line indicates a local linear fit using an Epanechnikov kernel with the "rule of thumb" optimal bandwidth of Fan and Gijbels (1996, Section 4.2). The dark grey and moderate grey lines indicate a fit using a bandwidth of 125% and 75% times this optimal bandwidth, respectively. The light grey line indicates a local cubic polynomial fit.

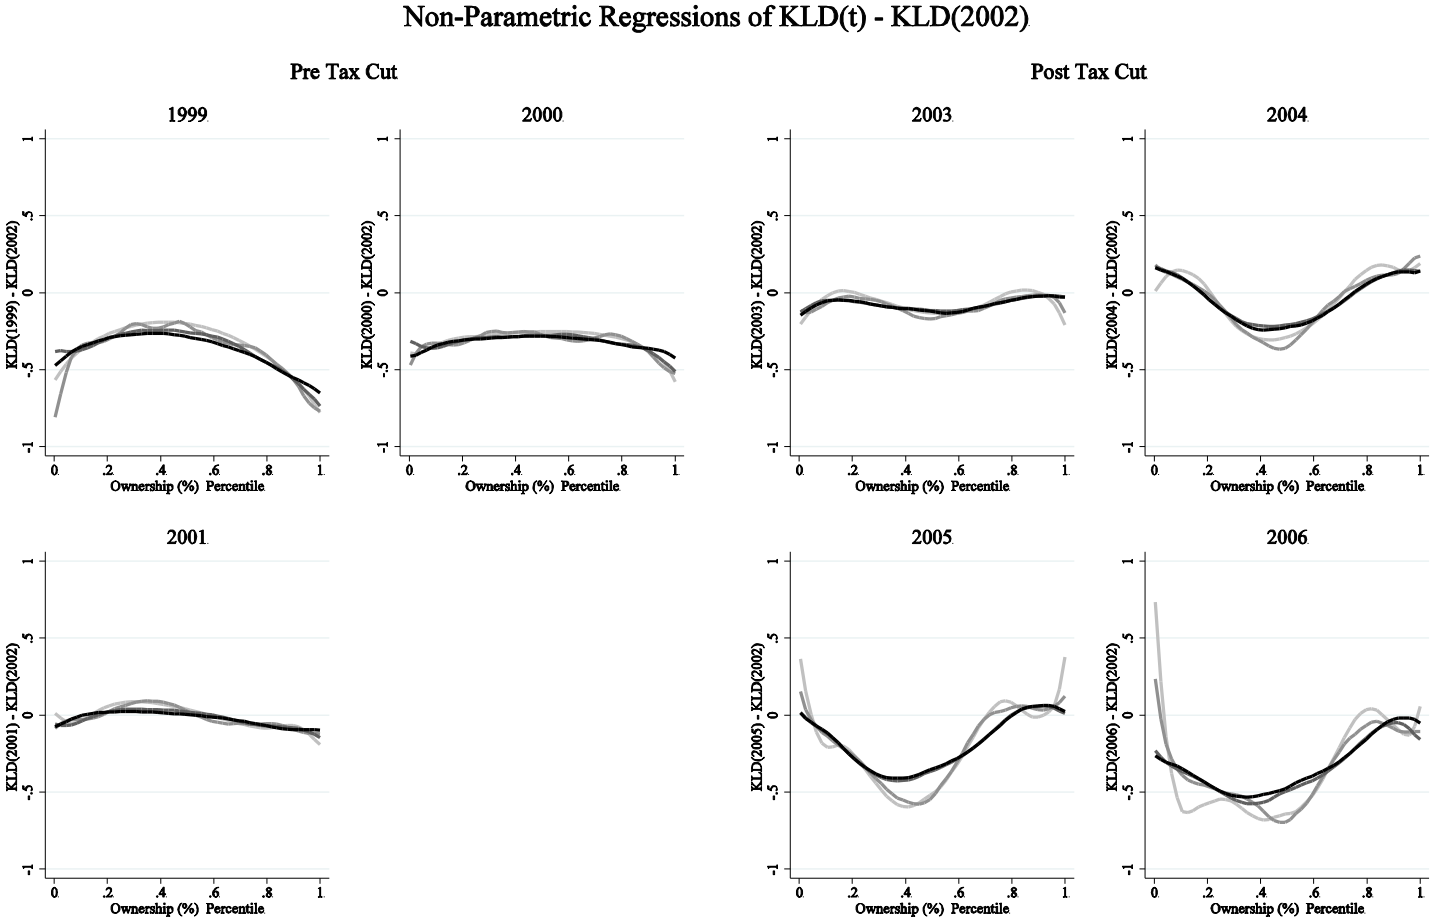

Figure 5: Distribution of Vote Shares

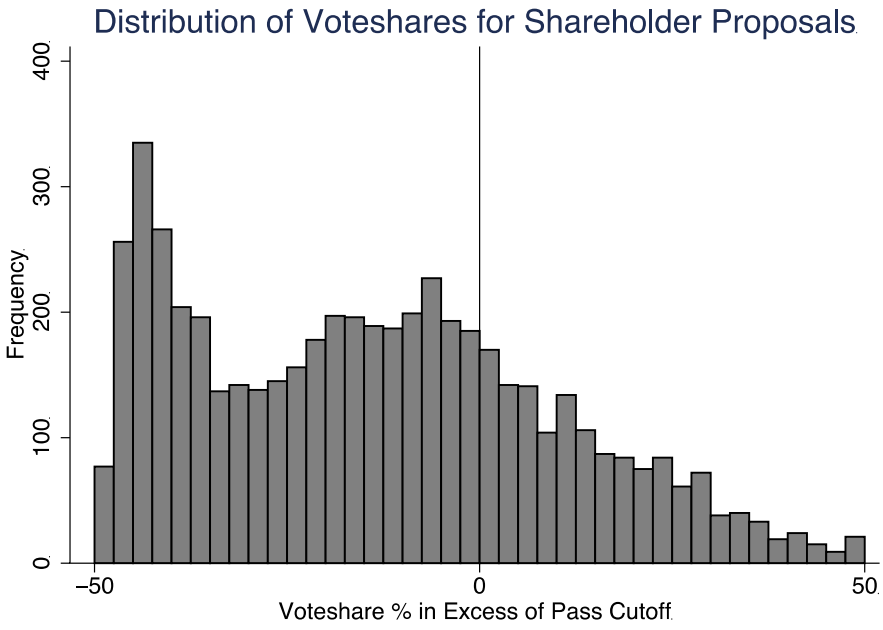

Figure 6: Change in KLD Scores at 50% Vote Share Discontinuity

This figure plots the baseline regression discontinuity result reported in Table 6 Panel A for the optimal bandwidth. Gray bands reports 90 percent pointwise confidence intervals. Although the confidence intervals overlap at the cutoff, the difference at the discontinuity is significantly different from zero with a p-value of 0.035.