

OVERSIGHT LIABILITY FOR RISK- MANAGEMENT FAILURES AT FINANCIAL FIRMS

ROBERT T. MILLER*

ABSTRACT

Many people believe that excessive risk taking at large financial firms was an important cause of the financial crisis of 2007–2008 and that preventing another crisis requires improving risk-management systems at such institutions. One way to do this would be to use board oversight liability to hold directors personally liable for failing to properly monitor the risks that their firms are running. The purpose of this Article is to determine what role director oversight liability can efficiently play in improving risk-management practices at large financial firms.

A key contention of this Article is that previous treatments of this problem have largely failed to appreciate what risk managers at large financial firms actually do, and so the Article begins by explaining some of the financial models that risk managers typically use to measure the market risk and credit risk on portfolios of assets. A realistic appreciation of these models shows that the measurements of risk that they yield must necessarily incorporate paradigmatic business judgments, most importantly because these models aim to predict future results on the basis of historical data. In

* Associate Professor of Law, Villanova University School of Law. For helpful comments and discussion, I thank Richard A. Booth, William Bratton, Thomas C. Crimmins, John Y. Gotanda, Joseph Hoffman, John F. Murphy, Jennifer O'Hare, Richard E. Redding, Chaim Saiman, and Mark A. Sargent, as well as the participants in the conference "Navigating the Storm: The Regulatory Response and Reform Agenda in the Wake of a Financial Crisis," held at the University of Pennsylvania Law School on February 15, 2010, at which I presented some of the ideas contained in this Article. I thank Courtney Magnarella and Ariana Gillies for their excellent work as my research assistants, and I thank Amy Spare and Matthew McGovern of the Villanova Law Library for invaluable assistance in obtaining some of the sources I cite. Most of all, I thank Jennifer L. Miller, whose assistance and comments on all aspects of this Article were absolutely indispensable to its completion.

other words, the predictive ability of the models is founded on the business judgment that the future will resemble the past in relevant respects. Risk-management decisions are therefore always business decisions.

With this conclusion firmly established, the Article reviews the principles of director oversight liability, the most relevant in this context being that oversight liability requires a showing that the directors were consciously disregarding their duties. This scienter-based standard practically guarantees that oversight claims based on alleged failures to detect and prevent weaknesses in the firm's risk-management systems will fail, as in fact happened in the Citigroup case, the most important oversight case predicated on alleged risk-management failures thus far litigated. The result in Citigroup has been subjected to much academic criticism. This Article considers these criticisms and argues that, in light of the actual nature of risk management and the financial models used therein, these criticisms are generally misguided.

The result in Citigroup has also prompted proposals for expanding oversight liability in ways that would allow courts to review substantive risk-management decisions by corporate boards. The Article argues that such proposals are meritless for two reasons. First, because risk-management decisions are always business decisions, and because any business decision leading to losses for the company can be characterized as a risk-management failure, allowing courts to review risk-management decisions in oversight liability cases would, in effect, repeal the business judgment rule. Second, although such proposed expansions of oversight liability are aimed at limiting the excessive risk taking that supposedly contributed to the financial crisis, risk taking can be excessive in several distinct ways, and the sense in which excessive risk taking represents a genuine failure of risk-management systems (namely, risk taking in excess of the risk tolerance of the firm as set by the board) is not the sense in which excessive risk taking may have been a cause of the financial crisis (namely, socially inefficient risk taking or excessive systemic risk). Hence, even if expanding oversight liability for risk-management failures did not otherwise involve the tremendous inefficiency of effectively repealing the business judgment rule, such an expansion would still not be well calculated to address the perceived problem of excessive risk taking as a cause of the financial crisis.

This Article concludes that director oversight liability has little or no role to play in improving risk-management practices at major financial firms.

TABLE OF CONTENTS

I. INTRODUCTION.....	50
II. RETURN, RISK, AND RISK-MEASUREMENT MODELS.....	57
A. RETURN AND RISK ON A SINGLE INVESTMENT.....	57
B. MEASURING RISK IN THE FIRM'S PORTFOLIO.....	60
1. Market Risk Models.....	61
a. VAR Models: Back Simulations.....	62
b. VAR Models: Monte Carlo Simulations.....	66
c. Stress Tests.....	67
2. Credit Risk.....	73
a. RORAC and RAROC.....	74
b. Option Models of Default Risk.....	77
c. The KMV Portfolio Manager Model.....	79
III. LIABILITY FOR RISK-MANAGEMENT FAILURES UNDER <i>CAREMARK AND STONE</i>	81
A. OVERSIGHT LIABILITY UNDER <i>CAREMARK AND STONE</i>	82
B. THE <i>CITIGROUP</i> CASE.....	87
1. The Court's Opinion in <i>Citigroup</i>	87
2. Evaluation of the <i>Citigroup</i> Case.....	91
a. Plaintiffs' Ignorance of Risk-Management Systems.....	91
b. Plaintiffs' Misunderstanding of Risk-Management Systems.....	93
3. Academic Criticisms of <i>Citigroup</i>	95
a. Cognizability of Oversight Claims for Risk- Management Failures.....	96
b. Judicial Setting of Risk-Management Standards.....	97
c. Adapting <i>Caremark</i> to the Risk-Management Context.....	100
IV. THE INEFFICIENCY OF EXPANDING OVERSIGHT LIABILITY FOR RISK-MANAGEMENT FAILURES.....	102
A. EXPANDING <i>CAREMARK</i> WOULD REPEAL THE BUSINESS JUDGMENT RULE.....	103
B. THE VARIOUS MEANINGS OF "EXCESSIVE RISK".....	109
1. Risks in Excess of Return-Adjusted Systematic Risk.....	110
2. Risks in Excess of the Risk Tolerance of the Firm as Set by the Board.....	111
3. Risks in Excess of the Risk Tolerances of Shareholders....	114
4. Socially Inefficient Risk Taking and Excessive Systemic Risk.....	114
V. CONCLUDING OBSERVATIONS.....	120

I. INTRODUCTION

Many people think that excessive risk taking by major financial institutions was an important cause of the financial crisis of 2007–2008, and, consequently, that there were significant failures of risk-management systems at such firms. President Barack Obama recently told the Business Roundtable that “we have a financial system with the same vulnerabilities that it had before this crisis began,” and “if there aren’t rules in place to guard against the recklessness of a few . . . it starts a race to the bottom that results in all of us losing.”¹ Similarly, former Federal Reserve Chairman Alan Greenspan has written that “the most sophisticated private sector risk management was unable to neutralize the burst” of the asset bubble precipitating the crisis, and the “incontrovertible evidence of underpricing of risk did not prod private sector risk management to tighten the reins.”² AIG’s former chairman and chief executive officer, Maurice “Hank” Greenberg, thinks that one important reform needed is to change irresponsible risk-management systems that failed to prevent the crisis.³

Academic commentary is in broad agreement. For example, Stephen Bainbridge thinks that the financial crisis “revealed serious risk management failures on an almost systematic basis throughout the business community.”⁴ Judge Richard Posner explains such failures by saying that, in the years leading up to the financial crisis, financial firms tended to give more weight to the views of risk-creating traders than risk-reducing risk managers because the contributions of the former to the profits of the firm are easier to measure than those of the latter.⁵ Advancing a somewhat different explanation, Lucian Bebchuk and Holger Spamann argue that the structure of bankers’ compensation packages created incentives for

1. President Barack Obama, Remarks by the President to the Business Roundtable (Feb. 24, 2010), available at <http://www.whitehouse.gov/the-press-office/remarks-president-business-roundtable>. See also RICHARD A. POSNER, A FAILURE OF CAPITALISM: THE CRISIS OF '08 AND THE DESCENT INTO DEPRESSION 322–23 (2009) (“[N]o single bank, in the highly competitive financial-intermediation industry, could justify to its shareholders reducing its risk taking . . . and therefore their return on equity, merely because the risks that it and its competitors were taking might precipitate a financial crisis There would be only one effect of the bank’s altruism . . . the bank would lose out in competition with its daring competitors.”).

2. ALAN GREENSPAN, THE AGE OF TURBULENCE 523 (2008).

3. Maurice R. Greenberg, *Rescuing American Capitalism*, NAT’L INT., Jan.–Feb. 2010, at 4; Maurice R. Greenberg, *Six Steps Toward Financial Reform*, WALL ST. J., Mar. 5, 2010, at A21.

4. Stephen M. Bainbridge, *Caremark and Enterprise Risk Management*, 34 J. CORP. L. 967, 968 (2009). See also Melanie Roldier, *Reasonable Risk*, WALL ST. & TECH., Dec. 31, 2008, at 15 (stating that the financial crisis has been blamed on “[p]oor risk management practices”).

5. POSNER, *supra* note 1, at 80.

excessive risk taking.⁶ More generally, Mark Zandi thinks that the financial crisis resulted not simply from problems in the United States' residential real estate market, but because "[g]lobal investors had taken on too much risk, not simply in their subprime mortgage security holdings, but arguably in all their investments," with riskiness in all kinds of investments being systematically underpriced.⁷ There is thus a broad consensus among academics that financial firms should improve risk-management systems to limit future risk taking.

Risk management, however, is difficult in practice.⁸ Designing and applying the financial models needed to compute risk levels requires significant mathematical sophistication, and determining the numbers to input into these models entails making difficult business judgments under conditions of great uncertainty. In fact, as Greenspan has written, "Financial regulators . . . know far less than private-sector risk managers," and "the open secret about regulation in the free-market world is that regulators take their cues from private-sector practitioners."⁹ Thus, the "Federal Reserve and other supervisory institutions continually seek the advice of the best and brightest risk-management professionals," and so new regulation, "if it comes, will reflect the private sector's already revised market practices."¹⁰ Thus, subject to certain limitations, regulators

6. Lucian A. Bebchuk & Holger Spamann, *Regulating Bankers' Pay*, 98 GEO. L.J. 247, 247 (2009) (arguing that "banks' compensation structures have produced incentives for excessive risk taking"). See also Paul Krugman, Editorial, *Reform or Bust*, N.Y. TIMES, Sept. 21, 2009, at 23 (arguing that the compensation packages of bank executives encourage excessive risk taking); Lucian A. Bebchuk & Holger Spamann, *Reducing Incentives for Risk Taking*, N.Y. TIMES DEALBOOK (Oct. 12, 2009, 9:00 AM), <http://dealbook.blogs.nytimes.com/2009/10/12/reducing-incentives-for-risk-taking/> ("[C]ompensation structures in financial firms should be devised to avoid excessive incentives for risk-taking and . . . doing so requires tying executive compensation to long-term results and preventing cashing out of large amounts of compensation on the basis of short-term results.").

7. MARK ZANDI, FINANCIAL SHOCK: A 360° LOOK AT THE SUBPRIME MORTGAGE IMPLOSION, AND HOW TO AVOID THE NEXT FINANCIAL CRISIS 177 (2009).

8. For academic treatment, see generally MICHEL CROUHY, DAN GALAI & ROBERT MARK, THE ESSENTIALS OF RISK MANAGEMENT (2005) (describing the entire financial risk-management field, including policies, methodologies, data, and technological infrastructure), and JAMES LAM, ENTERPRISE RISK MANAGEMENT (2003). For the leading study among industry professionals seeking to establish best practices, see Comm. of Sponsoring Orgs. of the Treadway Comm'n, *Enterprise Risk Management—Integrated Framework* (Sept. 2004), http://www.coso.org/Publications/ERM/COSO_ERM_ExecutiveSummary.pdf. Even a cursory perusal of either of the books mentioned above, not to mention the multivolume Committee of Sponsoring Organizations study, conveys the immense complexity of the subject. For general treatment in legal literature, see Susan Schmidt Bies, *Keynote Address to the International Symposium on Risk Management and Derivatives*, 8 FORDHAM J. CORP. & FIN. L. 81 (2003), and Betty Simkins & Steven A. Ramirez, *Enterprise-Wide Risk Management and Corporate Governance*, 39 LOY. U. CHI. L.J. 571 (2008).

9. GREENSPAN, *supra* note 2, at 524.

10. *Id.*

generally allow large financial institutions to develop and use their own proprietary models for calculating risk.¹¹

But if those outside a major financial firm are unlikely to be able to improve on the decisions of the firm's own risk managers, the effort to improve risk management at such firms naturally focuses on how to make those *inside* the firm more attentive to risk management. For this reason, many of the proposals to reduce perceived excessive risk taking have centered on giving the board of directors increased incentives to monitor and control the risks that the business is running. As Martin Lipton has said, "the risk oversight function of the board of directors . . . has taken center stage . . . and expectations for board engagement with risk are at all-time highs."¹² Indeed, the Securities and Exchange Commission has already promulgated new rules requiring all public companies to provide greater disclosure about their risk oversight practices, including information about the board's role in managing risk.¹³ Senator Charles Schumer has introduced a bill that would require all public companies to

11. See, e.g., Capital Adequacy Guidelines for Bank Holding Companies: Internal-Ratings Based and Advanced Measurement Approaches, 12 C.F.R. pt. 225 app. G, pt. 1 § 1(a) (2010) (establishing "[m]inimum qualifying criteria for bank holding compan[ies] using bank-holding company-specific internal risk measurement and management processes for calculating risk-based capital requirements"); ANTHONY SAUNDERS & MARCIA MILLON CORNETT, FINANCIAL INSTITUTIONS MANAGEMENT: A RISK MANAGEMENT APPROACH 598–99 (6th ed. 2008) (discussing how the Federal Reserve allows bank holding companies meeting certain criteria to use their own internal risk-measurement models and risk-management processes to comply with regulatory requirements, "subject to strict methodological and disclosure standards"). Although in the absence of implementing regulations from the Federal Reserve the matter is partly unclear, the Dodd-Frank Wall Street Reform and Consumer Protection Act probably modifies this situation in some respects. See Dodd-Frank Wall Street Reform and Consumer Protection Act of 2010, Pub. L. No. 111-203, 124 Stat. 1376 (to be codified in scattered sections of 12 and 15 U.S.C.). In particular, the Dodd-Frank Act requires the Federal Reserve to establish enhanced risk-based capital, leverage, and liquidity requirements, and overall risk-management requirements (among other requirements) for certain systemically important financial firms, including bank holding companies with \$50 billion or more in assets. How such required enhanced standards will interact with the Federal Reserve's criteria under Appendix G for approving internal, company-specific models, and standards is currently unclear. See DAVIS POLK & WARDWELL LLP, SUMMARY OF THE DODD-FRANK WALL STREET REFORM AND CONSUMER PROTECTION ACT 7–10 (2010), http://www.davispolk.com/files/Publication/7084f9fe-6580-413b-b870-b7c025ed2ecf/Presentation/PublicationAttachment/1d4495c7-0be0-4e9a-ba77-f786fb90464a/070910_Financial_Reform_Summary.pdf.

12. Martin Lipton, *Risk Management and the Board of Directors*, HARV. L. SCH. FORUM ON CORP. GOVERNANCE & FIN. REG. (Dec. 17, 2009, 9:33 AM), <http://blogs.law.harvard.edu/corpgov/2009/12/17/risk-management-and-the-board-of-directors-2/>.

13. Proxy Disclosure Enhancements, Exchange Act Release Nos. 33-9089, 34-61175, IC-29092, 74 Fed. Reg. 68,334 (Dec. 29, 2009) (amending 17 C.F.R. §§ 229, 239, 240, 249, 274). See also N.Y. STOCK EXCH., NYSE LISTED COMPANY MANUAL § 303A.07(b)(i)(D) (2009) (mandating that listed companies have an audit committee charged with, among other things, "discuss[ing] policies with respect to risk assessment and risk management").

establish board committees to supervise enterprise-wide risk-management practices,¹⁴ a proposal that would enshrine in law requirements similar to those already imposed by stock exchange rules.¹⁵ In part duplicating and in part expanding existing legal requirements, the Dodd-Frank Wall Street Reform and Consumer Protection Act, enacted on July 21, 2010, requires that certain systematically important publicly traded financial companies, as well publicly traded bank-holding companies with consolidated assets of \$10 billion or more, establish risk committees that include independent directors and at least one risk-management expert having experience in risk management at large, complex companies.¹⁶

Going further, Warren Buffett has said that “a board of directors of a huge financial institution is *derelect* if it does not insist that its CEO bear full responsibility for risk control,” and “if [the CEO] fails at it . . . the financial consequences for him and his board should be severe.”¹⁷ Buffett’s reference to adverse financial consequences introduces the issue of board oversight liability—personal liability for directors who fail adequately to monitor and manage the risks that the corporation is bearing. The economic idea here is simple: improve risk-management practices by giving a strong financial incentive to those best placed to design, implement, and monitor those practices. Legally, such a strategy seems feasible because, under Delaware law and the analogous corporate laws of other states, the board’s authority to manage the business and affairs of the corporation¹⁸ implies a fiduciary duty to monitor the activities of the corporation.¹⁹ If the board breaches this duty to monitor, it exposes itself to so-called oversight

14. Shareholder Bill of Rights Act of 2009, S. 1074, 111th Cong. § 5 (2009). *See also U.S. Sen. Schumer Unveils Shareholder Bill of Rights*, REUTERS (May 19, 2009, 12:37 PM), <http://www.reuters.com/article/idUSTRE54I4PF20090519>. For criticism of the bill, see Martin Lipton, Jay W. Lorsch & Theodore Mirvis, *Schumer’s Shareholder Bill Misses the Mark*, WALL ST. J., May 12, 2009, at A15.

15. *See, e.g.*, N.Y. STOCK EXCH., *supra* note 13, § 303.A07(b)(iii)(D) (providing that a listed company’s audit committee shall have among its duties and responsibilities “discuss[ing] policies with respect to risk assessment and risk management”). Perhaps not surprisingly in light of the conclusions of this Article, the Act does virtually nothing else to affect the role of the board of directors supervising the risk-management practices of a financial firm. For a helpful summary of the immensely long Dodd-Frank Act, see DAVIS POLK & WARDWELL LLP, *supra* note 11.

16. Dodd-Frank Wall Street Reform and Consumer Protection Act of 2010 § 165(h), 124 Stat. at 1429–30 (to be codified at 12 U.S.C. § 5365).

17. Letter from Warren Buffett, chairman and CEO of Berkshire Hathaway Inc., to the shareholders of Berkshire Hathaway Inc. 16 (Feb. 26, 2010), *available at* <http://www.berkshirehathaway.com/letters/2009ltr.pdf>.

18. DEL. CODE ANN. tit. 8, § 141 (2001).

19. *In re Caremark Int’l Inc. Derivative Litig.*, 698 A.2d 959, 961, 970 (Del. Ch. 1996). More recently, the Delaware Supreme Court explicitly adopted and elaborated on *Caremark* in *Stone v. Ritter*, 911 A.2d 362, 369 (Del. 2006).

liability.

Oversight liability suits, often referred to as *Caremark* claims after the leading case,²⁰ have historically concerned failures by the board to detect and prevent criminal wrongdoing and fraud by subordinate employees that subject the corporation to liability. Nevertheless, it may seem reasonable that the board's duty to monitor under *Caremark* could also include a duty to monitor how much risk the business is bearing, thus allowing shareholders to hold the directors liable if they fail to detect and prevent the corporation's assumption of excessive risks. Such a system, it seems, would not only place the incentive where it would do the most good, but also would comport with basic notions of fairness. As Buffett puts it, "It has not been the shareholders who have botched the operations of some of our country's largest financial institutions," and so if these institutions "are harmed by the[] recklessness [of their CEOs and directors], [the CEOs and directors] should pay a heavy price."²¹

Such was the plaintiffs' theory in *In re Citigroup Inc. Shareholders Derivative Litigation*,²² a case decided by the Delaware Court of Chancery (Chancellor William B. Chandler III) in 2009. In *Citigroup*, the plaintiff shareholders alleged that some of Citigroup's current and former directors had "breached their fiduciary duty by failing to properly monitor and manage the risks the Company faced from problems in the subprime lending market,"²³ with the result that Citigroup was overexposed to such risks,²⁴ and consequently suffered severe losses when the subprime market collapsed.²⁵ In deciding the *Citigroup* case, Chancellor Chandler applied settled principles of Delaware law,²⁶ holding that directors can be liable in oversight only if they consciously disregarded their duties—that is, only if the directors actually knew (not merely should have known) that they were

20. See, e.g., *Stone*, 911 A.2d at 364 (noting that the lower court "characterized the allegations as a 'classic *Caremark* claim,' a claim that derives its name from *In re Caremark Int'l Deriv. Litig.*").

21. Letter from Warren Buffett, *supra* note 17, at 16.

22. *In re Citigroup Inc. S'holder Derivative Litig.*, 964 A.2d 106, 111 (Del. Ch. 2009). I have discussed *Citigroup* and some of the academic criticisms made of the case in Robert T. Miller, *The Board's Duty to Monitor Risk After Citigroup*, 12 U. PA. J. BUS. L. 1153, 1158–61 (2010).

23. *Citigroup*, 964 A.2d at 111.

24. *Id.* at 121.

25. *Id.* at 113 (detailing billions of dollars in losses Citigroup suffered from its exposure to subprime debt).

26. The leading case is *Stone v. Ritter*, 911 A.2d 362, 370 (Del. 2006) (holding that "imposition of [oversight] liability requires a showing that the directors knew that they were not discharging their fiduciary obligations"). I have discussed the rationale underlying this rule in Robert T. Miller, *Wrongful Omissions by Corporate Directors: Stone v. Ritter and Adapting the Process Model of the Delaware Business Judgment Rule*, 10 U. PA. J. BUS. & EMP. L. 911, 940–41 (2008).

not properly monitoring the business.²⁷ Applying this standard, Chancellor Chandler ruled for the defendant directors on their motion to dismiss the complaint.²⁸ The Chancellor went even further, however, for some language in the opinion suggests that oversight claims based on failures to monitor risks that the business is running should fail as a matter of law.²⁹

The result in *Citigroup* has led to significant academic criticism. Some of this criticism has centered on the result in the particular case, and thus Kelli Alces says that Citigroup's "board could easily be considered asleep at the switch when corporate catastrophe occur[red]."³⁰ Other criticism has centered on both the principles of Delaware law that entailed the result in the case and the possibility that *Caremark* claims based on failures to monitor risk-management systems may fail as a matter of law. For example, while not necessarily disagreeing with the result in *Citigroup*, Bainbridge says, "There is no doctrinal reason that *Caremark* claims should not lie in cases in which the corporation suffered losses . . . due to lax risk management."³¹ While acknowledging important differences between *Caremark* claims based on failures of risk management and such claims based on failures to ensure compliance with law, Bainbridge concludes that "courts need to develop a modified regime for deciding *Caremark* claims" based on alleged failures of risk management.³²

Going further, Eric Pan thinks that "[i]t seems fantastic that the duty to monitor . . . incentivizes boards to take no responsibility for the business results of the company—a complete disregard for the principle that the corporation shall be managed by or under the direction of the board."³³ He concludes that Delaware judges should "begin speaking out about the importance of a board's duty to monitor and to back up their exhortations by expanding the scope and application of the duty to monitor in future cases."³⁴ Adopting a more extreme view, J. Robert Brown notes that there is "extraordinary frustration with the excessive degree of risk taken by companies currently in trouble," and concludes that the Court of Chancery

27. *Citigroup*, 964 A.2d at 125, 128.

28. *Id.* at 139–40.

29. See, e.g., *id.* at 131 (stating that "[w]hile it may be tempting to say that directors have the same duties to monitor and oversee business risk" as they do to monitor the business to prevent fraud and illegality, "[o]versight duties under Delaware law are not designed to subject directors . . . to personal liability for failure to predict the future and to properly evaluate business risk").

30. See Kelli A. Alces, *Debunking the Corporate Fiduciary Myth*, 35 J. CORP. L. 239, 252 (2009).

31. Bainbridge, *supra* note 4, at 968.

32. *Id.*

33. Eric J. Pan, *A Board's Duty to Monitor*, 54 N.Y.L. SCH. L. REV. 717, 739 (2010).

34. *Id.* at 740.

squandered “an opportunity to define meaningful responsibilities by the board,” thus demonstrating that “Delaware cannot be trusted to define fiduciary obligations for directors.”³⁵

The primary purpose of this Article is to determine what role director oversight liability can play in improving risk-management practices at large financial institutions. Since I believe that previous treatments of this problem have largely failed to appreciate what risk managers at large financial firms actually do, I begin Part II by explaining some of the financial models such managers typically use to measure market risk and credit risk. My goal in Part II is not only to convey a sense of the extreme complexity and sophistication of contemporary risk-assessment models, but also to emphasize that, at key points, all such models inevitably rely on business judgments about future market conditions—judgments that are typically made by extrapolating from historical conditions. For some, the phrase “risk management” seems to denote an intellectual discipline applying scientific methods to obtain highly reliable results, and my intention in Part II is to make it clear that risk management, for all its mathematical complexity, is essentially and inevitably a matter of business judgments about future market conditions, not a scientific mechanism to prevent losses.

Having set out in Part II this realistic understanding of risk-measurement models used in contemporary risk management, I turn in Part III to the law of oversight liability, first reviewing the current doctrine under *Caremark* and then considering the *Citigroup* case. I conclude that, because under current law oversight liability requires a showing of conscious disregard of duty by corporate directors, the plaintiffs in *Citigroup*, who alleged virtually nothing about the states of mind of the Citigroup directors, never had any serious chance of prevailing. More importantly, however, I conclude in Part III that the arguments advanced by the plaintiffs in *Citigroup*, as well as most of those made by their academic supporters, are untenable in light of the description in Part II of the typical risk-management systems that major financial firms use. That is, both the plaintiffs and their academic supporters seem to have little understanding of the nature, purposes, and inherent limitations of risk-management systems.

35. J. Robert Brown, *Delaware Courts and Exonerating the Board from Supervising Risk: In re Citigroup Derivative Litigation (Introduction)*, [THERACETOTHEBOTTOM.ORG](http://www.theracetothetbottom.org/preemption-of-delaware-law/delaware-courts-and-exonerating-the-board-from-supervising-r-4.html) (March 12, 2009, 9:00 AM), <http://www.theracetothetbottom.org/preemption-of-delaware-law/delaware-courts-and-exonerating-the-board-from-supervising-r-4.html>.

Next, since *Citigroup* was decided under the existing *Caremark* doctrine, and since many commentators have argued for an expansion of that doctrine, in Part IV I consider whether an expansion of oversight liability would likely improve risk-management practices at major financial firms. I reach two main conclusions. First, I argue that any attempt to expand oversight liability beyond the limits of *Caremark* would effectively repeal the business judgment rule, which I assume is undesirable. Second, I argue that expanding oversight liability to control excessive risk taking is misguided in principle because there are several different senses in which risk taking can be called excessive, and the way in which excessive risk taking might conceivably be amenable to control by oversight liability (risk taking in excess of the board's risk tolerance) is not the way in which excessive risk taking may have contributed to the financial crisis (risk taking in activities that, through a tragedy of the commons effect, have positive expected value for the firm but negative expected value for society as a whole). In Part V, I make some concluding remarks.

II. RETURN, RISK, AND RISK-MEASUREMENT MODELS

In subsequent parts of this Article, I argue that popular and legal-academic treatments of risk management generally reflect a failure to understand the actual practices of risk managers at financial firms, resulting in commentary that is significantly disconnected from the realities it purports to address. My purpose in this part, therefore, is to supply some much-needed context by describing the problems risk managers face and some of the techniques they use in addressing them. To do that, I begin with some elementary financial theory presented with somewhat more mathematical precision than corporate lawyers and judges generally employ; this precision is necessary in order to present the underlying concepts of the risk-assessment models I later describe. That said, I emphasize that, despite the introduction of a few statistical concepts in this Article, the Article's description of risk-assessment models is a drastically simplified account of models and theories in a complex, mathematically challenging and constantly evolving specialty in modern financial theory.

A. RETURN AND RISK ON A SINGLE INVESTMENT

Consider an investor thinking about making a particular investment, such as investing \$100 in a certain stock. Having determined a holding period for the investment, the investor starts by computing the investment's *expected value*, its weighted-average outcome. To do this, the investor identifies the possible values of the stock at the end of the selected holding

period (assume for simplicity that the stock pays no dividends during this period), and estimates the probability that each value will result. Multiplying each possible outcome by the probability that it will occur and adding the results, the investor obtains the expected value of the investment at the targeted date. For example, if the investor concludes that, on the selected future date, there are three equally likely outcomes in which the stock is worth \$180, \$120, and \$90, then the expected value of the investment on that date is $(\frac{1}{3})(\$180) + (\frac{1}{3})(\$120) + (\frac{1}{3})(\$90) = \130 . Expressing the expected gain on the investment ($\$130 - \$100 = \$30$) as a percentage of the amount invested (\$100), the investor computes the *expected rate of return*—in this case, 30 percent.

Of course, the mathematical computation here is the easy part. The hard part is correctly predicting the future—knowing what the possible outcomes are and the probabilities with which they may result. When sophisticated investors try to do this, they usually begin by looking at the historical returns of similar investments,³⁶ a procedure that clearly depends on the assumption that the future will resemble the past.³⁷ For example, when investing in particular stocks, investors may begin by determining the distribution of historical returns over the past ten years and then calculating the average annual return on the stock for this period.³⁸ They may then adjust the average annual return over this period to reflect whatever factors they think will make the future different from the past. Determining the expected value of the investment is thus nothing more than educated guesswork. As Aswath Damodaran writes, “When valuing an asset at any point in time, we make forecasts for the future. Since none of us possess crystal balls, we have to make our best estimates given the information that we have at the time.”³⁹ Thus, “there will always be uncertainty associated

36. See, e.g., ASWATH DAMODARAN, INVESTMENT VALUATION 64 (2d ed. 2002) (“[E]xpected returns . . . that we run into in practice are almost always estimated using past returns rather than future returns.”).

37. *Id.* (“The assumption made when using historical [values] is that past distributions are good indicators of future return distributions. When this assumption is violated, as is the case when the asset’s characteristics have changed significantly over time, the historical estimates may not be good measures of risk.”).

38. This procedure can introduce various biases. For example, on the one hand, it may seem that the investor would get a better estimate by including as many years as possible in the historical average, but, on the other hand, the years in the distant past may represent conditions that are very different from the present and so would be less likely to be representative of future results. So there is a tension between including *more* years in the average and including *more relevant* years. Also, the year that is chosen as the first year to include in the average may be a year in which the market as a whole was at historically high or historically low levels; if the former, returns on subsequent years may tend to understate future results, and if the latter, such returns may tend to overstate them.

39. ASWATH DAMODARAN, DAMODARAN ON VALUATION 5 (2d ed. 2006).

with valuations, and even the best valuations come with a substantial margin for error.”⁴⁰

Now, if the investor were sure to earn the expected return, the investment would be risk free. In fact, however, the actual return may—indeed, likely will—differ from the expected return.⁴¹ This fact makes the investment risky. Intuitively, risk is the spread of the possible returns around the expected return. In legal discussions, the matter is often left at this level, if it is stated correctly at all.⁴² For our purposes, however, more mathematical precision will be essential. In particular, the intuition that risk is the spread of possible returns around the expected return is made precise by defining the risk as the *variance* or *standard deviation* of the possible returns.⁴³

To calculate the variance, we start by taking, for each possible return, the difference between such possible return and the expected return. This difference is called the deviation of the possible return from the expected return. For instance, in the example above, the possible returns are +80% (when the value is \$180), +20% (when the value is \$120), and –10% (when the value is \$90), and so the deviations from the expected return of +30% are +50%, –10%, and –40%, respectively. We would like to take an average of these, but since some are positive and some are negative, adding them would result in some offsetting and thus understate how much the possible returns deviate from the expected return. To overcome this problem, we square these deviations,⁴⁴ which makes all of the values positive and gives us, in this case, the values 25%, 1%, and 16%, respectively. Averaging these, we get the variance, which is 14%. Taking the square root (intuitively, this undoes the squaring we needed to do to make all the numbers positive), we get the standard deviation, which in this case is approximately 37.4%.

40. *Id.* at 4.

41. DAMODARAN, *supra* note 36, at 61 (“[A]ctual returns . . . may be very different from the expected returns, and it is this difference between actual and expected returns that is a source of risk.”).

42. For example, in *Citigroup*, the court incorrectly stated that “[r]isk has been defined as the chance that a return on an investment will be different than [sic] expected.” *In re Citigroup Inc. S’holder Derivative Litig.*, 964 A.2d 106, 126 (Del. Ch. 2009). Risk is not the *probability* that the actual return will differ from the expected return, but rather a measure of how much potential returns differ from expected returns—a measure that incorporates not only the probability of the possible return differing from the expected return, but also how much these will differ.

43. DAMODARAN, *supra* note 36, at 62 (“The spread of the actual returns around the expected return is measured by the variance or standard deviation of the distribution; the greater the deviation of the actual returns from the expected return, the greater the variance.”).

44. STEPHEN A. ROSS, RANDOLPH W. WESTERFIELD & JEFFREY F. JAFFE, *CORPORATE FINANCE* 272–73 (3d ed. 1993).

Now, the crucial point is that calculating the risk involves the same highly uncertain numerical inputs involved in calculating the expected return—namely, all the possible outcomes of the investment and the probability that each such outcome will occur. In computing the return, we started with these numbers and took their weighted average. In calculating the risk, we again started with these numbers and did more complex calculations to get their variance and the standard deviation. In both cases, therefore, we had to begin with educated guesses about the future. Of course, we never really know the future. Given our limited cognitive powers, the best that human beings can do is consider similar cases in the past and extrapolate from these by assuming that the future will resemble the past.⁴⁵ Damodaran's point about the inherent uncertainty of valuation applies to risk as well⁴⁶: since exactly the same forecasts for the future underlie the computation of the risk in an investment as underlie its expected value, even the best estimates of risk come with a substantial margin of error.

B. MEASURING RISK IN THE FIRM'S PORTFOLIO

So far I have discussed only the estimating of risk and return on a single investment. Financial firms, of course, have large numbers of investments, including both long and short positions, and they are not as concerned with the risk and return on individual trades as with the risk and return on the firm's portfolio as a whole. Although how best to measure and control the risk of the firm's portfolio depends on the composition of that portfolio, large financial firms typically distinguish between *market risk* and *credit risk* in their risk-management systems.⁴⁷ Of course, such

45. Of course, our hypothetical investor evaluating the investment described in the text would also have to discount the expected value of the investment at the selected time horizon back to present value, and to do that he would have to determine the appropriate discount rate. That rate will depend on the investor's weighted-average cost of capital ("WACC") for the investment, which will depend on the riskiness of the investment. DAMODARAN, *supra* note 36, at 712–13 (referring to "the fundamental investment principle that it is not who raises the money that determines the cost of equity as much as what the money is raised for," and stating that "[t]he same firm will face a higher cost of equity for funds raised to finance riskier projects and a lower cost of equity to finance safer projects"). Hence, once again, the investor would have to rely on the empirical assumptions that underlie the computation of the expected return and the risk as explained in the text. In fact, since the WACC will depend on the systematic component of the riskiness of the investment and not the unsystematic component, and since separating these depends on making assumptions about market returns and market risk, the investor will have to make analogous assumptions about the distribution of future market prices, which introduces even more uncertainty. *See, e.g., id.* at 69 (stating that the capital asset pricing model is the "risk and return model that has been in use the longest and is still the standard in most real-world analyses").

46. *See id.* at 64.

47. The difference between market risk and credit risk is discussed in SAUNDERS & CORNETT,

firms are subject to many other kinds of risks as well,⁴⁸ but these two are probably the most important. Further, they are undoubtedly the most studied and well-developed areas of risk management. Accordingly, the purpose of this section is to explain some of the more important market-risk and credit-risk models, and to indicate the kinds of business judgments needed to apply them. As I mentioned above, the presentation here is drastically simplified in comparison to the models that professional risk managers at large financial firms actually use.

1. Market Risk Models

One of the ways major financial firms typically seek to profit is by holding and trading on organized financial markets certain assets, liabilities, and derivatives, including long and short positions in equity securities, debt securities, commodities, and various options and swaps.⁴⁹ The set of assets, liabilities, and derivatives that the firm holds for such purposes is its trading portfolio,⁵⁰ and since the values of the individual positions in its portfolio depend on market conditions, so does the value of the whole portfolio. Such a firm thus faces market risk, the risk to the firm's earnings caused by changes in market conditions affecting the values of the assets, liabilities, and derivatives in its portfolio,⁵¹ including changes

supra note 11, at 168–75 (defining “market risk” as “[t]he risk incurred from assets and liabilities in [a financial institution’s] trading book due to changes in interest rates, exchange rates, and other prices,” and “credit risk” as “[t]he risk that promised cash flows from loans and securities held by [financial institutions] may not be paid in full”). Such distinctions are so common and well established that they are reflected in law, including the Federal Reserve’s regulations regarding capital adequacy for bank holding companies. See Capital Adequacy Guidelines for Bank Holding Companies: Risk-Based Measure, 12 C.F.R. pt. 225 app. A (2010); Capital Adequacy Guidelines for Bank Holding Companies: Internal-Ratings-Based and Advanced Measurement Approaches, 12 C.F.R. pt. 225 app. G (2010).

48. For example, most large financial firms are also subject to the following risks: (1) off-balance sheet risks, which result from contingent assets and liabilities that, precisely because of their contingent nature, are not recorded on the firm’s balance sheet; (2) foreign exchange risks, which result from changes in exchange rates affecting the value of the firm’s assets and liabilities denominated in nondomestic currencies; (3) sovereign risks, which arise from the actions of sovereign states interrupting or restricting payments to the firm; (4) operational risks, which result from failures of the firm’s information, auditing, reporting, monitoring, or support systems; and (5) liquidity risks, which arise from sudden withdrawals from the firm necessitating the sale of assets in short periods of time at less than their fair market value. See SAUNDERS & CORNETT, *supra* note 11, at 168–84 (defining and describing various kinds of risks faced by financial firms).

49. *Id.* at 266 (noting that a financial firm’s “trading portfolio can be differentiated from its investment portfolio on the basis of time horizon and liquidity” because the former “contains assets, liabilities, and derivative contracts that can be quickly bought or sold on organized financial markets (such as long and short positions in bonds, commodities, foreign exchange, equity securities, interest rate swaps, and options)”).

50. *Id.*

51. *Id.* at 171–72 (“Market risk arises when [financial firms] actively trade assets and liabilities

in trading prices, interest rates, market volatility, and market liquidity. Major financial firms have developed several kinds of financial models to measure market risk,⁵² the most important of which are value at risk (“VAR”) models. Below, I describe the simplest kind of VAR models (back simulations) and then explain, in a very general way, how these simple models can be improved by certain statistical techniques (Monte Carlo methods). Along the way, I emphasize how the construction and application of these models requires, at key junctures, the application of business judgment. I then turn to stress tests, which are in some respects similar to VAR models, but involve even more business judgment in their design and application.

a. VAR Models: Back Simulations

Most financial firms concerned about market risk use at least some models based on historic or back simulation approaches.⁵³ Conceptually, back simulations are relatively straightforward. The firm takes its current portfolio of positions, which has a certain net value based on current market prices, and revalues it on the basis of actual market prices as they existed yesterday, the day before yesterday, the day before that, and so on,⁵⁴ going back for a predetermined number of days, commonly five hundred days.⁵⁵ Since a large firm typically holds thousands or tens of thousands of positions, the computational task involved here may be

(and derivatives) rather than hold them for longer-term investment, funding, or hedging purposes. . . . Market risk is the incremental risk incurred by [a financial firm] when interest rate, foreign exchange, and equity return risks are combined with an active trading strategy”). Anthony Saunders and Marcia Millon Cornett define “market risk” as “[t]he risk incurred from assets and liabilities in [a financial firm’s] trading book due to changes in interest rates, exchange rates, and other prices,” *id.* at 168, or as “the risk related to the uncertainty of [a financial firm’s] earnings on its trading portfolio caused by changes, and particularly extreme changes, in market conditions such as the price of an asset, interest rates, market volatility, and market liquidity,” *id.* at 266. Since changes in interest rates can affect the values of positions in the firm’s trading portfolio, market risk includes interest rate risk with respect to the assets affected by market risk.

52. For example, see the RiskMetrics models first developed at J.P. Morgan in 1994. J.P. MORGAN/REUTERS, RISKMETRICS—TECHNICAL DOCUMENT (4th ed. 1996), available at http://pascal.iseg.utl.pt/~aafonso/eif/rm/TD4ePt_4.pdf (“[Providing] a detailed description of RiskMetrics, a set of technique and data to measure market risks in portfolios of fixed income instruments, equities, foreign exchange, commodities, and their derivatives issued in over 30 countries.”). See also JAMES T. GLEASON, RISK 186 (2000) (describing RiskMetrics as “the most commonly used parametric VaR model”); SAUNDERS & CORNETT, *supra* note 11, at 268–77.

53. SAUNDERS & CORNETT, *supra* note 11, at 279 (“[T]he large majority of [financial firms] that have developed market risk models have employed a historic or back simulation approach.”).

54. GLEASON, *supra* note 52, at 187 (noting that, in performing a historic simulation, “In a nutshell, you run your daily [profit and loss] process over and over again, using the history of past daily market price moves”).

55. SAUNDERS & CORNETT, *supra* note 11, at 279.

immense.⁵⁶ But because the computations are based on publicly available market data, there are no significant business judgments involved at this stage. All the factors affecting the values of the positions in the portfolio are accounted for because the values being used are actual historical values.

Once the firm has determined the value that its portfolio would have had on each of the last five hundred days, for example, it ranks the days from worst to best, the worst being the day the portfolio had the lowest value and the best being the day the portfolio had the highest value. It then determines a worst case. For example, the 5 percent worst case is the value the portfolio had on the twenty-fifth worst day among the last five hundred days. With respect to that day, the firm's portfolio was worth more on 95 percent of the days in the sample and less on 5 percent of the days in the sample. The firm then computes the VAR of its portfolio as the difference between the current market value and the value on the 5 percent worst case day. For instance, if this difference is \$50 million, the firm will say that the VAR on its portfolio is \$50 million.

Thus far, the whole exercise has been a purely mathematical one based on historical market data. The next step is critical: the VAR is interpreted as a measure of the firm's current market risk. That is, if, as in the example above, the 5 percent worst case yields a VAR of \$50 million, the model is interpreted as meaning that the probability that the firm's portfolio will lose more than \$50 million in value tomorrow is less than 5 percent. Equivalently, risk managers will say that the VAR on the portfolio tomorrow is \$50 million with 95 percent confidence, meaning that the firm can be 95 percent certain that it will not lose more than \$50 million in the next day's trading because of changes in market conditions.

The underlying assumption in this interpretation is that the past days used in the back simulation are collectively an accurate predictor of what tomorrow will be like—equivalently, that the distribution of past market conditions is an accurate predictor of the probability distribution of future market conditions. Here is where business judgment and uncertainty return, for market conditions change over time, sometimes very quickly, and so the distribution of past conditions is not a perfect predictor of future

56. The computations are still simpler than in many other financial models because there is no need to compute variances, standard deviations, covariances, correlations, and so forth. *See id.* at 279 (noting that the "advantages of [a back simulation] are that (1) it is simple, (2) it does not require that asset returns be normally distributed, and (3) it does not require that the correlations or standard deviations of asset returns be calculated"). It is worth noting, however, that such measures are not really excluded from the analysis. Rather, the market prices used in running the back simulation already impound the market's estimates of these things.

conditions. For example, if over the period of the historical sample equity prices were high by historical standards, or if interest rates were unusually low by these standards, then these unusual conditions will be reflected in the model as if they were normal. The model will thus tend to underpredict changes in these conditions. Similarly, very unusual but nonetheless not unprecedented events (for example, a stock market crash) may not be reflected in the historical sample because these events have not occurred during the sample period.⁵⁷ The possibility of such events occurring in the future will thus not be reflected in the predictions of the model. Notice too the similarity between using standard valuation techniques to price an asset and using back simulations to measure risk: just as in the former we assume that the distribution of historical returns on the asset predicts future returns, so too in the latter we assume that the hypothetical distribution of values of a portfolio based on actual historical conditions predicts the portfolio's future value. In both cases, we are assuming that the future will resemble the past, which, we know from experience, it often does not.

Since there is often not a more reliable way of predicting the future than to assume it will resemble the past, even knowledgeable people sometimes speak as if the predictions of back simulations are as certain as predictions based on the laws of physics. For example, people sometimes speak of the firm's VAR as being a worst-case scenario loss that it might sustain.⁵⁸ In fact, the VAR as usually calculated is not even the worst case but only the 5 percent (or 1 percent, and so forth) worst case, and even that assumes that nothing worse will happen tomorrow than happened during the period covered by the sample used in the back simulation.

Furthermore, even when these obvious mistakes are avoided, using back simulations to predict the future also assumes that the days included in the sample are a genuinely random sample, which they are not. On the contrary, the days in the sample reflect conditions in the recent past, and so whatever atypical conditions may have prevailed during that period—such as a raging bull market, or a recession, or a bubble in technology stocks or the residential real estate market—will be reflected in the sample.⁵⁹ This

57. *E.g.*, PHILIPPE JORION, *VALUE AT RISK: THE NEW BENCHMARK FOR MANAGING FINANCIAL RISK* 357 (3d ed. 2007) (“In practice, VAR measures based on recent historical data can fail to identify extreme unusual situations that could cause severe losses.”).

58. *See* Christine Harper, *Death of VaR Evoked as Risk-Taking Vim Meets Taleb's Black Swan*, BLOOMBERG (Jan. 27, 2008, 7:12 PM), <http://www.bloomberg.com/apps/news?pid=20670001&sid=axo1oswvqx4s>.

59. *E.g.*, GLEASON, *supra* note 52, at 188 (“Whatever market trends are in the history you use are built into the VaR you calculate. Some risk professionals see this as a problem.”). James T. Gleason also notes that, when the purpose is to calculate the VAR over a longer interval (for example, the VAR

means that predictions based on the sample will tend to overestimate the probability of the continuation of current conditions and underestimate the probability that future conditions will be different from recent ones.

In addition to these problems, the five hundred or so days typically included in a back simulation are a relatively small sample for statistical purposes.⁶⁰ The obvious solution would be to increase the size of the sample by adding more days and stretching further into the past, say, by expanding the sample to ten thousand days. The problem with this solution, however, is that days further in the past are less relevant for predicting the future, because conditions in the distant past tend to be more unlike conditions today than are conditions in the more recent past.⁶¹ For instance, extending the sample to ten thousand days would include dates more than twenty-five years in the past. That long ago, many of the equity securities in the firm's portfolio did not exist because their issuing companies had not yet sold securities to the public or perhaps had not even been founded, some relevant exchange rates are unobtainable because the euro had yet been created, and many of the derivative securities in the firm's portfolio had not yet even been invented, much less been traded on publicly observable markets. The relevance of such long-ago days for predicting tomorrow's results, therefore, is very low.

Risk managers have responded to this problem in two main ways. One involves taking a larger sample of days but weighting them for relevance,⁶² giving more weight to recent days and less weight to more remote days.⁶³ Another, probably the best regarded among professional risk managers nowadays⁶⁴ and the most commonly used at large financial firms,⁶⁵ uses

over the coming ten days, not the VAR over just the next trading day), because overlapping periods cannot be used, the sample required will stretch far into the past (for example, to get sixty observations for a ten-day VAR, we need to go back three years). This results in a sample that includes observations of highly dubious relevance. *Id.*

60. See, e.g., DAVID FREEDMAN ET AL., STATISTICS 335–38 (2d ed. 1991) (discussing the appropriate size of statistical samples).

61. SAUNDERS & CORNETT, *supra* note 11, at 282.

62. See GLEASON, *supra* note 52, at 187.

63. See, e.g., LINDA ALLEN, JACOB BOUDOUKH & ANTHONY SAUNDERS, UNDERSTANDING MARKET, CREDIT AND OPERATIONAL RISK: THE VALUE AT RISK APPROACH 36 (2004); JORION, *supra* note 57, at 230.

64. GLEASON, *supra* note 52, at 185 (“Conventional wisdom on which [VAR] forecasting technique is best has shifted . . . to Monte Carlo.”); JORION, *supra* note 57, at 307 (“Because of its flexibility, the [Monte Carlo] simulation method is by far the most powerful approach to VAR.”).

65. GLEASON, *supra* note 52, at 189 (stating that as of the year 2000, U.S. banks and securities firms were trending toward using Monte Carlo methods in calculating VAR); JORION, *supra* note 57, at 307 (“Securities firms commonly use simulation techniques, known as *Monte Carlo methods*, . . . to measure risk.”).

Monte Carlo simulations.

b. VAR Models: Monte Carlo Simulations⁶⁶

In Monte Carlo simulations, risk managers obtain a larger sample of days by creating artificial but realistic days.⁶⁷ The mathematics soon become extremely complex, but the concept is reasonably straightforward.⁶⁸ We begin by identifying the risk factors that can affect the value of the firm's portfolio, such as changes in interest rates, exchange rates, securities prices, and so forth. In general, there will be important relationships among these factors. For example, when the federal funds rate increases, foreign exchange rates tend to change in favor of the dollar, and when inflation increases, bond prices fall, and so forth. Based on the real days included in some recent historical sample, we can compute the covariances of the various risk factors—a statistically precise measure of how the variables vary against one another. We do this for each pair of the identified risk factors, arranging them in a matrix. We then assume that a set of market conditions is realistic if the differences between it and current market conditions respect the covariances in the matrix—intuitively, if the changes between current conditions and the hypothetical conditions are in accordance with the pattern of changes recorded in the matrix.⁶⁹ We then generate a very large number of random⁷⁰ but realistic cases—ten thousand days is a common measure—and calculate the VAR in the portfolio on the basis of the values that the firm's portfolio would have under such conditions. That is, we look at the five-hundredth worst day to find the 5 percent worst case, the one-hundredth worst day to find the 1 percent worst case, and so forth, and calculate the VAR accordingly.

66. Despite the colorful name, Monte Carlo methods are a serious branch of mathematics used for a variety of purposes beyond financial theory. Invented by Stanislaw Ulam and developed by John von Neumann, they were first used by scientists involved in the Manhattan Project to solve certain mathematical problems related to nuclear chain reactions. JORION, *supra* note 57, at 308.

67. For examples of Monte Carlo simulations in calculating VAR, see CHRISTOPHER L. CULP, *THE RISK MANAGEMENT PROCESS* 353–57 (2001); GLEASON, *supra* note 52, at 195–203; CHRIS MARRISON, *THE FUNDAMENTALS OF RISK MEASUREMENT* 118–31 (2002); and SAUNDERS & CORNETT, *supra* note 11, at 282–83.

68. “The basic concept behind the Monte Carlo approach is to simulate repeatedly a random process for the financial variable of interest covering a wide range of possible situations. These variables are drawn from prespecified probability distributions that are assumed to be known . . .” JORION, *supra* note 57, at 308.

69. The text here elides several significant mathematical complexities. See MARRISON, *supra* note 67, at 119–26 (discussing Cholesky decompositions and Eigenvalue decompositions of the covariance matrix).

70. Generating truly random numbers, it turns out, is not a simple thing to do, and professional risk managers rely on sophisticated algorithms to produce draws of numbers that pass certain statistical tests for randomness. See JORION, *supra* note 57, at 313.

Monte Carlo methods thus solve the problem of the small size of the relevant sample available in back simulations. That said, however, Monte Carlo simulations come with limitations of their own. For one, they are very computationally intensive, requiring significant computing power and coding.⁷¹ For another, the usual implementations of Monte Carlo methods assume that the distribution of risk factors have certain convenient statistical properties that, at least sometimes, the actual data do not have.⁷²

Additionally, there are more fundamental limitations to Monte Carlo simulations. First, in creating the matrix needed to generate the artificial trading days, we have to identify the relevant risk factors; factors not included in the selection will not be reflected in the results. Hence, there is an important business judgment here—a judgment about which factors can affect the value of the firm’s portfolio. In practice, however, this problem is probably relatively small because major financial firms use hundreds of risk factors in their Monte Carlo simulations.⁷³ Second, the accuracy of the predictions from the Monte Carlo simulations depends on, among other things, the future covariances of the identified risk factors being consistent with their historical covariances. We are assuming, in other words, that future changes in market conditions will be like past changes in relevant respects. This is now the third time we have seen this: just as in valuing assets we assume that the distribution of historical returns predicts future returns, and just as in simple back simulations we assume that hypothetical values of a portfolio based on past conditions predict its future value, so too in Monte Carlo simulations we assume that past covariances of risk factors predict their future covariances. The mathematics become increasingly complex, but the business judgment is fundamentally the same: we assume that the past is a guide to the future, realizing that it is, at best, a very imperfect guide.

c. Stress Tests

Back simulations, as noted above, fail to take account of events not reflected in the historical sample on which such simulations are based, and even when supplemented by Monte Carlo methods, VAR models generally

71. *Id.* 307 (stating that Monte Carlo methods “involve[] costly investments in intellectual and systems development” as well as “substantially more computing power than simpler methods”); MARRISON, *supra* note 67, at 126 (discussing computational intensity of Monte Carlo simulations and various solutions).

72. *See* MARRISON, *supra* note 67, at 118 (“The usual implementation of Monte Carlo assumes a stable, Joint-Normal distribution for the risk factors.”).

73. The large number of factors can create significant mathematical issues, not merely increased demand for computing power to perform the necessary calculations. *See id.* 122–23 (discussing difficulties in decomposing the covariance matrix that arise as more risk factors are included).

fail to take account of very rare but plausible events that could cause the firm to suffer significant losses.⁷⁴ One widely used method of addressing these limitations of VAR models is *stress testing*,⁷⁵ a procedure that came to the attention of the general public when, after the height of financial crisis had passed, the Federal Reserve required certain major banks in the United States to perform certain stress tests it has designed.⁷⁶

A stress test is essentially an exercise in what-if analysis. That is, in performing a stress test, a financial firm first assumes that certain exceptional but plausible events occur that affect the values and volatilities of securities and other assets and liabilities—such as a stock market crash, a sudden and large change in interest rates, or a default on its sovereign

74. The Committee on the Global Financial System (“CGFS”), a discussion forum on global financial markets established by the central banks of the Group of Ten nations, has conducted two studies on stress testing in risk management at major financial institutions. *See generally* COMM. ON THE GLOBAL FIN. SYS., *STRESS TESTING AT MAJOR FINANCIAL INSTITUTIONS: SURVEY RESULTS AND PRACTICE* (2005) [hereinafter 2005 CGFS SURVEY] (describing results of a study of stress tests at sixty-four large banks and securities firms from sixteen countries as of May 31, 2004); COMM. ON THE GLOBAL FIN. SYS., *A SURVEY OF STRESS TESTS AND CURRENT PRACTICE AT MAJOR FINANCIAL INSTITUTIONS* (2001) [hereinafter 2001 CGFS SURVEY] (describing results of a census of stress test scenarios in use at forty-three banks from ten countries as of May 31, 2000). In its 2001 survey, the CGFS notes the following:

Firms recognise the limited ability of statistical models such as VaR to accurately capture what happens in exceptional circumstances. In part, this is due to modeling assumptions that make it easier to compute VaR. However, there is a more fundamental problem with using statistical models like VaR for assessing risks in exceptional circumstances. By definition, exceptional circumstances occur rarely, and statistical inference is imprecise without a sufficient number of observations. Stress tests partially fill this gap, and thus complement VaR, by offering a quantitative picture of the exposure associated with a possible extreme event. In the absence of a reliable statistical measure of the probability of such an event, stress testing calls on the informed judgement of risk managers and senior executives to assess whether, and to what degree, the firm should move to limit or modify such exposure.

2001 CGFS SURVEY, *supra*, at 13 (quoting COMM. ON THE GLOBAL FIN. SYS., *STRESS TESTING BY LARGE FINANCIAL INSTITUTIONS: CURRENT PRACTICE AND AGGREGATION ISSUES 7–8* (2000)). *See also* JORION, *supra* note 57, at 357 (“VAR measures . . . can fail to identify extreme unusual situations that could cause severe losses.”).

75. Certain kinds of stress testing are required under Basel II and applicable regulations of the Federal Reserve. *See* Capital Adequacy Guidelines for Bank Holding Companies: Internal-Ratings Based and Advanced Measurement Approaches, 12 C.F.R. pt. 225 app. G (2010); JORION, *supra* note 57, at 357 (“Stress testing is indeed required by the Basel Committee as one of seven conditions to be satisfied [for a bank] to use internal models [to measure risk].”); *infra* note 284 and accompanying text. On stress testing generally, see Martin Čihák, *Introduction to Applied Stress Testing* (Int’l Monetary Fund, Working Paper No. 07/59, 2007), available at <http://www.imf.org/external/pubs/ft/wp/2007/wp0759.pdf>.

76. *See* BD. OF GOVERNORS OF THE FED. RESERVE SYS., *THE SUPERVISORY CAPITAL ASSESSMENT PROGRAM: DESIGN AND IMPLEMENTATION 1–3* (2009) [hereinafter *DESIGN AND IMPLEMENTATION*]. *See generally* BD. OF GOVERNORS OF THE FED. RESERVE SYS., *THE SUPERVISORY CAPITAL ASSESSMENT PROGRAM: OVERVIEW OF RESULTS* (2009) [hereinafter *OVERVIEW OF RESULTS*] (discussing the results of the stress tests).

debt by a given nation.⁷⁷ The firm then determines the value of its portfolio under the financial conditions that would result from the stressing event, computing the amount of its losses. It can then determine the VAR if the stressing event were to occur. By considering the effects of such rare but reasonably possible events, stress testing seeks to determine the consequences for the firm of events that would not occur under standard VAR models.⁷⁸ Stress tests are thus an important complement to VAR methods,⁷⁹ and long before the financial crisis, they were standard tools in risk management at large financial firms.⁸⁰

In some stress tests, the source of the shock or stressing event is a well-defined scenario,⁸¹ a set of specifiable events in the real world that would affect financial conditions, such as a war on the Korean peninsula, a default by Greece on its sovereign debt, or the termination of the Economic and Monetary Union of the European Union (the agreement that underlies the euro). Such stress tests are generally called *scenario stress tests*.⁸² Scenario stress tests can be historical, when based on actual events in the past, or hypothetical, when based on imaginable events that have never occurred but that might occur in the future.⁸³ Thus, such tests typically involve simultaneous changes in a number of financial variables, such as equity prices, foreign exchange rates, and interest rates.⁸⁴ Other stress tests, generally called *sensitivity stress tests*, do not posit any cause in the real

77. *E.g.*, 2005 CGFS SURVEY, *supra* note 74, at 3–4 (“Stress testing is a risk management tool used to evaluate the potential impact on a firm of a specific event and/or movement in a set of financial variables,” and is used to “[analyze] the impact on a portfolio of exceptional but plausible large loss events.”); 2001 CGFS SURVEY, *supra* note 74, at 7 (defining a “stress test” as any of “various techniques used by financial firms to gauge their potential vulnerability to exceptional but plausible events”).

78. *See* JORION, *supra* note 57, at 360 (“The goal of stress testing is to identify unusual scenarios that would not occur under standard VAR models.”).

79. 2005 CGFS SURVEY, *supra* note 74, at 2 (“Stress testing works as a complement, rather than as a supplement, to major risk management tools such as value-at-risk. It is therefore becoming an integral part of the risk management frameworks of banks and securities firms.”); 2001 CGFS SURVEY, *supra* note 74, at 6 (noting that “stress testing has become an integral part of banks’ risk management”).

80. *See* 2005 CGFS SURVEY, *supra* note 74; 2001 CGFS SURVEY, *supra* note 74; JORION, *supra* note 57, at 357–75.

81. 2005 CGFS SURVEY, *supra* note 74, at 3–4 (“In scenarios, the source of the shock, or stress event, is well defined, as are the financial risk parameters which are affected by the shock.”).

82. 2001 CGFS SURVEY, *supra* note 74, at 7 (noting that a “stress test scenario contains simultaneous moves in a number of risk factors . . . reflecting an event that the firm’s risk managers believe may occur in the foreseeable future”).

83. *Id.* at 7 (“[A] stress test scenario can be based on a significant market event experienced in the past (a historical scenario) or on a plausible market event that has not yet happened (a hypothetical scenario).”).

84. *Id.*

world for the stressing event,⁸⁵ but merely inquire what would happen to the firm's portfolio in the event of a sudden change in one or more specified financial variables,⁸⁶ such as an upward parallel shift of one hundred basis points of the yield curve for U.S. Treasury securities or a 20 percent drop in the dollar-to-euro exchange rate.⁸⁷

It should be obvious that stress testing involves even more business judgment than VAR methods, both as to the design and implementation of stress tests and to the interpretation of their results. As to design and implementation, in VAR methods, business judgments are essential because risk managers have to assume that certain historical data (for instance, about market prices in back simulations or about correlations of changes in risk factors in Monte Carlo simulations) will correctly predict future market conditions. Thus, risk managers must exercise business judgment in extrapolating from historical conditions to future conditions. In stress testing, however, risk managers must quite literally try to predict the future by identifying unlikely but possible future events that could adversely affect the value of the firm's portfolio.

In using historical scenarios, risk managers do this by looking to exceptional past events on the theory that extreme events that occurred in the past might occur again in the future. For example, in the survey conducted by the Committee on the Global Finance System ("CGFS") in 2005 of stress testing at major financial institutions, some of the most common historical episodes used in the surveyed firms' historical scenario stress testing were the 1994 bond market crash, the 1997 Asian financial crisis, the 1998 Long-Term Capital Management and Russian bond crisis, the 2001 terrorist attacks in the United States, the 1987 Black Monday stock market crash, and the 2000 burst of the technology stock bubble.⁸⁸ In such historical scenarios, the analogy with VAR models is apparent: risk managers are assuming that certain future events will resemble identified

85. 2005 CGFS SURVEY, *supra* note 74, at 4 ("[W]hile sensitivity tests specify financial risk parameters, the source of the shock is not identified.").

86. A key difference between sensitivity tests and scenario stress tests is that, while many sensitivity tests consider changes in just one financial variable, scenario stress tests consider the correlated changes in multiple financial variables that would be expected to result from the underlying stressing event posited in the scenario test. See JORION, *supra* note 57, at 365–71 (discussing multidimensional scenario analysis).

87. 2001 CGFS SURVEY, *supra* note 74, at 7 ("A sensitivity stress test isolates the impact on a portfolio's value of one or more predefined moves in a particular market risk factor or a small number of closely linked market risk factors. . . . The most common sensitivity stress test reported on the census was a parallel yield curve shift."). See also JORION, *supra* note 57, at 362–65 (discussing principles governing sensitivity tests).

88. 2005 CGFS SURVEY, *supra* note 74, at 18.

past events, but even more business judgment is required than in VAR methods because the risk managers must identify the relevant events and make judgments about the possibility of their recurrence.

In hypothetical scenarios, yet more business judgment is required because risk managers are trying to predict the future more directly— inventing possible scenarios subject only to the limits of human imagination. In practice, the hypothetical scenarios stress tested by major financial firms tend to be fairly obvious ones. In the 2005 CGFS Survey, for instance, some of the most commonly tested hypothetical scenarios were drastic downturns in the economies of the United States⁸⁹ or Europe, changes in monetary policy by the Federal Reserve, downturns in the global economy, increased inflation expectations, collapses of certain pegged currencies, terrorist attacks, geopolitical unrest in the Middle East, adverse events in the Chinese economy and related changes in the foreign exchange rates, and natural disasters.⁹⁰ Clearly, most of these events are more or less based on historical events.⁹¹ The reason for this is apparent: human beings are simply not very good at predicting unusual—let alone unprecedented—events before they happen.⁹² Not surprisingly, once such an unusual event occurs, it tends to enter the repertoire of stress tests employed by major financial firms. For example, after the terrorist attacks of September 11, 2001, terrorist attacks became a common scenario in stress testing.⁹³ As the CGFS has noted, even for hypothetical scenarios, the “design of stress tests draws heavily on historical events.”⁹⁴

Thus, without the ability to foresee the future, there can be no science of determining which events a financial firm should stress test, and no mechanical procedure to identify stressing events to test. In reality, sometimes risk managers formulate scenarios at the request of senior

89. Such was the stressing event in the stress tests required by the Federal Reserve of major American banks after the financial crisis. In particular, the key scenario in the exercise was “a recession that [was] longer and more severe than the [then-current] consensus expectation” regarding the recession following the financial crisis. DESIGN AND IMPLEMENTATION, *supra* note 76, at 3. For the results, see generally OVERVIEW OF RESULTS, *supra* note 76.

90. See 2005 CGFS SURVEY, *supra* note 74, at 18.

91. See 2001 CGFS SURVEY, *supra* note 74, at 20 (“The design of stress tests draws heavily on historical events, even for hypothetical scenarios.”).

92. JORION, *supra* note 57, at 374 (“The history of some firms has shown that people can be very bad at predicting extreme situations.”).

93. Indeed, although terrorist attacks were completely absent from the scenarios banks were using in 2001, they were commonly stress tested by 2005. Compare 2001 CGFS SURVEY, *supra* note 74, at 25–27, with 2005 CGFS SURVEY, *supra* note 74, at 18–19.

94. 2001 CGFS SURVEY, *supra* note 74, at 20.

managers.⁹⁵ Other times, recent events—such as significant moves in market prices—can suggest events to stress test.⁹⁶ It is entirely a matter of business judgment to divine, before they occur, very unusual events that could adversely affect a firm’s portfolio. According to the 2005 CGFS Survey, “key risk managers in a firm initially discuss and identify vulnerabilities in the portfolio held by the firm,” and then “[h]aving determined these vulnerabilities, risk managers work backwards and formulate plausible scenarios under which these vulnerabilities are stressed.”⁹⁷ Obviously, there is no algorithmic way to do this.

Even once a scenario is identified, however, running a stress test requires additional business judgment. In scenario stress tests, stressing events are likely to affect many financial variables simultaneously, but how they do so is often highly uncertain. For example, if the nations of the European Union decided to abolish the euro and revert to their national currencies, determining the effect on the firm’s portfolio would require making certain assumptions about the various initial exchange rates between the revived national currencies and the dollar. Of course, there is no way to know with certainty what these rates would be, but running the stress test would require that precise exchange rates be assigned; hence, the firm’s risk managers will have to use their business judgment and assign projected exchange rates. If the stressing event is a stock market crash in the United States (for instance, a 20 percent drop in the prices of equity securities), what would the effect be on bond prices, interest rates, and the dollar-euro exchange rate? Running a stress test based on a stock market crash requires making these assumptions. Hence, one of the striking results of the 2001 CGFS Survey was “that the magnitude of shocks [assumed by financial firms in running stress tests] varies substantially among scenarios that, on the surface, look quite similar.” As Philippe Jorion observes, stress testing is highly subjective,⁹⁸ so this is “true even for historical scenarios . . . whose shocks are common knowledge.”⁹⁹

Finally, even after a scenario has been identified and its effects on the firm’s portfolio calculated, interpreting the results of the exercise requires additional business judgment. Most importantly, since it is generally

95. 2005 CGFS SURVEY, *supra* note 74, at 6.

96. *Id.*

97. *Id.* (noting that this tactic is a version of the portfolio-driven approach).

98. JORION, *supra* note 57, at 374.

99. 2001 CGFS SURVEY, *supra* note 74, at 15. The CGFS explains that one source of the difference is that different firms may be using different time horizons to measure the shocks; for example, one firm may use a one-day shock, another a two-week shock, a third a peak-to-trough shock, and so forth. *Id.*

impossible to assign any definite probability to the scenario's occurrence, stress test results are generally presented without an attached probability.¹⁰⁰ Managers know that the event is within the realm of reasonable possibility, but they have no quantitative assessment of its probability.¹⁰¹ Thus, what a firm should do with the results of a stress test—for example, a determination that the stressing event would result in specified large losses for the firm—is far from clear. The CGFS has aptly stated that “stress testing calls on the informed judgement of risk managers and senior executives to assess whether, and to what degree, the firm should move to limit or modify [its] exposure” to the stressing event.¹⁰² Indeed, if a stress test shows that the stressing event would cause losses greater than the firm is willing to tolerate, then “there is a range of responses that a bank could take.”¹⁰³ These responses might range from “a discussion between the risk manager and the business unit about why the stress test limit was breached” to “mandatory unwinding of positions,” with many options in between, such as a decision to hedge some or all of the identified risk.¹⁰⁴ Firms will generally deal with the results of stress tests on a case-by-case basis, “depend[ing] on the bank’s general attitude toward stress test limits, the level within the firm of the limit breach, and market conditions.”¹⁰⁵ In other words, given the results of the stress tests, the firm’s managers have to make certain business decisions.

2. Credit Risk

Besides market risk, financial firms also face credit risk—the risk that cash flows to which they are legally entitled will not be paid when due. For

100. JORION, *supra* note 57, at 374 (“Generally, stress-test results are presented without an attached probability, which makes them difficult to interpret.”). This is true when the stress test is being used to assess market risk on a portfolio-wide basis. When the stress test is being used to, for example, allocate economic capital, it is necessary to assign a probability. Doing so, of course, would involve a paradigm business judgment. *See* 2005 CGFS SURVEY, *supra* note 74, at 7.

101. In the usual economic jargon, the term “uncertainty” is used instead of “risk.” Following a distinction first drawn by Frank Knight, a situation is said to involve “risk” if the randomness presents itself in the form of scientifically calculable objective probabilities, as with gambles on a (fair) roulette wheel or pair of dice. A situation is said to involve “uncertainty” if the randomness presents itself in the form of alternative possible events to which such objective probabilities cannot be assigned, such as the results of a horse race or, for that matter, an investment. *See* FRANK H. KNIGHT, RISK, UNCERTAINTY AND PROFIT 197–232 (1921); RICHARD A. POSNER, THE CRISIS OF CAPITALIST DEMOCRACY 288–91 (2010) (describing uncertainty in Knight’s sense); Mark J. Machina & Michael Rothschild, *Risk*, in 7 THE NEW PALGRAVE DICTIONARY OF ECONOMICS 190–91 (Steven N. Durlauf & Lawrence E. Blume eds., 2d ed. 2008).

102. 2001 CGFS SURVEY, *supra* note 74, at 13.

103. *Id.* at 17.

104. *Id.*

105. *Id.*

example, a financial firm holding a portfolio of loans knows that in all probability, some of them will go into default—that is, the borrower will be unable to pay the interest or principal on the loan. When that happens, the firm will have to recover whatever it can in bankruptcy, often much less than one hundred cents on the dollar.¹⁰⁶ Credit risk thus affects assets involving promised, fixed cash flows, including not only loans, but also debt securities and certain derivatives.¹⁰⁷ Indeed, in performing their function of transforming financial assets, financial firms create value in large part because of their willingness to bear—and ability to estimate and manage—credit risk.¹⁰⁸ Measuring that risk turns out to be very difficult.

a. RORAC and RAROC

Start with the risk on individual loans. Because different borrowers present different credit risks, financial firms in the business of lending money would like to be able to compare loans controlling for the credit risk of the borrowers. There are many ways this can be done. For example, the firm could revalue the loans, applying higher discount rates to more risky loans and lower discount rates to less risky loans, thus lowering the expected returns for riskier loans and obtaining risk-adjusted rates of return for the loans. Another way of accomplishing the same result is to adjust the capital the firm must deploy to make the investment. That is, the firm can calculate the return on the investment using not the capital the firm actually deploys in making the investment (for example, the amount of cash it

106. SAUNDERS & CORNETT, *supra* note 11, at 173–75 (discussing the probability distribution of returns on risky debt).

107. *Id.* at 168 (defining “credit risk” as the “risk that promised cash flows from loans and securities . . . may not be paid in full”); CULP, *supra* note 67, at 19 (defining “credit risk” as “the risk of the actual or possible nonperformance by a firm” of its financial obligations); MARRISON, *supra* note 67, at 5–6 (stating that credit risk “arises from defaults, when an individual, company or government fails to honor a promise to make a payment”); Frank J. Jones & Benjamin Wolkowitz, *The Determinants of Interest Rates on Fixed-Income Securities*, in THE FINANCIAL ANALYSTS HANDBOOK 273, 302 (Sumner N. Levine ed., 2d. ed. 1988) (“[T]he credit risk of a security is a measure of the likelihood that the issuer of the security, the borrower, will be unable to pay the interest or principal on the security when due. Credit risk is thus a measure of the creditworthiness of the issuer of the security.”). The boundary between credit risk and market risk is not sharp. For example, when the spread between corporate bonds and treasury bonds varies, the change in the value of corporate bonds (1) can reasonably be viewed as a market risk for a holder of those bonds since it results in changes in the value of the holder’s portfolio, but (2) can also reasonably be viewed as a credit risk for the holder because the change results from changing market perceptions as to the likelihood that the issuer will default. See MARRISON, *supra* note 67, at 5. See generally GLEASON, *supra* note 52, at 113–22 (discussing the recent history of credit risk management and the shift from traditional qualitative methods to quantitative methods and modeling).

108. SAUNDERS & CORNETT, *supra* note 11, at 295 (stating that the ability of a financial firm “to evaluate information and to control and monitor borrowers allows it to transform [financial] claims [such as consumer deposits into personal and commercial loans] at the lowest possible cost”).

actually lends to the borrower), but a greater amount—an amount adjusted upward to reflect the credit risk of the loan. This gives us the return on risk-adjusted capital (“RORAC”). The RORAC of a loan is thus a fraction—the numerator of which is the net income from the loan (usually for one year) including interest and fees, but less the cost of funds, and the denominator of which is the amount lent adjusted to reflect the credit risk of the loan:

$$RORAC = \frac{\text{Net Income for One Year}}{\text{Risk Adjusted Capital}}$$

Computing the denominator is difficult. But once that problem is solved, RORAC provides a metric to compare loans of different credit risk. The result is that a firm can decline to lend if the loan’s RORAC is below a rate determined by senior managers.

Another method of evaluating credit risk commonly used by major banks involves comparing the expected return on the loan to the capital the firm expects to have at risk on the loan, thus determining the risk-adjusted return on capital (“RAROC”) of the loan.¹⁰⁹ The RAROC of a loan is again a fraction—the numerator of which, as in RORAC, is the net income from the loan (usually for one year), including interest and fees, but less the cost of funds.¹¹⁰ In RAROC, however, the denominator is the amount in dollars that the firm estimates to be at risk on the loan in the event of default—a quantity generally called the capital at risk. Thus, we have the following:

$$RAROC = \frac{\text{Net Income for One Year}}{\text{Capital at Risk}}$$

By comparing the expected return to the capital at risk, RAROC gives risk managers a measure of the risk-adjusted return on capital used in making the loan. Thus, if the firm is considering making two loans with identical expected returns but different RAROCs, the loan with the higher RAROC is more attractive on a risk-adjusted basis because it puts less of the firm’s

109. *Id.* at 328 (noting that the RAROC methodology “was pioneered by Bankers Trust (acquired by Deutsche Bank in 1998) and has now been adopted by virtually all the large banks in the United States and Europe, although with some significant proprietary differences between them”).

110. CULP, *supra* note 67, at 465 (stating that the net income in the numerator of RAROC “is the period revenue on an activity less the economic costs of that activity, where the period is based on the frequency of capital allocations and rebalancing,” and costs “may include funding costs, operating costs, bonuses, salaries, and other costs of doing business” and “may be adjusted to reflect the opportunity cost of regulatory capital tied up” in the transaction). *See also* SAUNDERS & CORNETT, *supra* note 11, at 328–29 (defining “net income” in the numerator of RAROC to be “net income (accounting for the cost of funding the loan) on the loan”). In financial jargon, the numerator here is the all-in spread.

capital at risk in order to earn the same expected return.¹¹¹ Hence, most major financial firms have a policy of making loans only if the loan's RAROC is above a certain specified level—if the return is high enough to justify the risk.¹¹²

Computing a loan's RAROC, however, involves significant business judgments. As in RORAC, computing the numerator is the easy part. Computing the denominator—the capital at risk—is the hard part. There are several ways of doing this,¹¹³ but one common way starts from the intuition that, as the creditworthiness of the borrower changes, the value of the loan changes as well. That is, if the credit of the borrower deteriorates, the risk premium on the loan (the spread between the interest rate the market would charge the borrower over the risk-free rate) will increase, thus decreasing the market value of the loan. If the creditworthiness of the borrower improves, the risk premium on the loan will decrease, thus increasing the market value of the loan. Therefore, holding everything else constant, the capital at risk on the loan, and thus the RAROC, will be a function of the change in the spread of the borrower's borrowing rate over the risk-free rate.¹¹⁴

But how does one estimate the maximum possible change in this spread—the largest possible interest rate shock due to changes in the creditworthiness of the borrower? The most common way is to look at historical interest rate shocks and use them to predict possible future interest rate shocks. Ideally, the lending firm would like to consider historical data about changes in risk premiums on loans to companies similar to the borrower, but since the interest rates that large banks charge their customers are not generally publicly disclosed, obtaining a large sample of such data is difficult.¹¹⁵ Instead, financial firms commonly look at changes in risk premiums on publicly traded corporate bonds. The lender

111. See CULP, *supra* note 67, at 464 (describing RAROC as “appealing [as a measure of risk-adjusted return] because it can be consistently applied to and compared across business units, regardless of the nature of the businesses”).

112. See *id.* at 470 (“The comparison of RAROC to a hurdle rate enables a firm to determine whether a business is viable and thus entitled to new capital (or to keep its existing risk capital). One such hurdle rate is the requirement that risk-adjusted returns on a business be at least as high as the risk-adjusted returns required as compensation for the systematic risk of the business.”).

113. See *id.* at 465–69 (discussing various ways of defining capital at risk for use in RAROC); SAUNDERS & CORNETT, *supra* note 11, at 331 (explaining how some financial firms measure capital at risk on the basis of loan default rates).

114. Here, the text considerably simplifies the ideas and methods of calculation for RAROC models. See SAUNDERS & CORNETT, *supra* note 11, at 328–31 (discussing more fully the ideas and methods of calculation behind RAROC models).

115. *Id.* at 330 (stating that “publicly available data on loan risk premiums are scarce”).

would first estimate the credit rating that the loan would be assigned by credit rating agencies based on the identity of the borrower and the point of attachment of the loan on the borrower's balance sheet (that is, the relative seniority of the loan). Next, the lender would look at all changes in risk premiums for publicly traded bonds of the same rating over some recent period, say the last year, thus obtaining a frequency distribution of changes in risk premiums. The lender would then define the maximum possible interest rate shock as being a shock greater than a certain percentage of the shocks in the sample—commonly 99 percent. Thus, the maximum shock would be a shock greater than 99 percent and less than 1 percent of the changes in risk premiums in the sample, meaning that the capital at risk in the loan's RAROC is the capital lost by the firm in the 1 percent worst-case scenario. Or it would be, if historical data about interest rate shocks were an infallible guide to future interest rate shocks, which, of course, they are not.

b. Option Models of Default Risk

Another way of measuring credit risk is based on the theory of option pricing.¹¹⁶ The concept here is that the debt and equity positions in a firm can be viewed as options on the underlying assets of the firm. On the one hand, the equity holders have a call option on these assets, because they can, in effect, obtain a 100 percent ownership interest in the assets by paying the debt holders the face amount of the firm's outstanding debt.¹¹⁷ The strike price of this option is thus the face value of the debt. If the value of the firm's assets is greater than the face amount of the debt, the equity holders will, in effect, exercise the option by repaying the debt. If the value of the firm's assets falls below the face amount of the debt, the equity holders will not exercise the option—that is, they will not repay the debt—and will allow the debt holders to take possession of the firm's assets through bankruptcy.

From the point of view of the debt holders, the value of their position will be the lesser of the face amount of the debt and the value of the firm's assets. For if the value of the firm's assets exceeds the face amount of the debt, the equity holders will repay the loan at face value, but if the value of

116. The seminal theoretical paper is Robert C. Merton, *On the Pricing of Corporate Debt: The Risk Structure of Interest Rates*, 29 J. FIN. 449 (1974), building on the classic paper by Fisher Black & Myron Scholes, *The Pricing of Options and Corporate Liabilities*, 81 J. POL. ECON. 637 (1973).

117. The original idea appears in Black & Scholes, *supra* note 116, and was greatly expanded in Merton, *supra* note 116. See also GLEASON, *supra* note 52, at 159 ("Equity holders have an implicit call on firm value, which has been sold to them by the debtors—the debtors are short a put. The strike price on both these options is the firm's debt.")

the firm's assets falls below the face amount of the debt, the equity holders will not repay the debt and the debt holders will take possession of the assets. For debt holders, the economic effect is the same as having written a put option on the assets of the firm with the face amount of the debt as the exercise price. If the value of the assets falls below the exercise price, the equity holders exercise the option and put the assets to the debt holders,¹¹⁸ with the debt holders buying them by discharging the debt via bankruptcy. If the value of the assets exceeds the exercise price, the equity holders keep the assets and repay the debt, allowing the put to expire unexercised, and the repayment of the debt being the equivalent of the debt holder not having to purchase the assets.¹¹⁹

This means that, theoretically, the value of a loan—the debt position of the lender—can be calculated using option-pricing models such as Black-Scholes.¹²⁰ In the 1990s, KMV Corporation (now a unit of Moody's),¹²¹ developed a practical way of doing this for public company borrowers.¹²² In KMV's model, the market value of the common stock of the firm is interpreted to be the value of an option on its underlying assets,¹²³ with the face amount of its debt being the strike price on the option. The volatility of the firm's common stock over a recent historical period is interpreted as the leverage-adjusted volatility of the value of the firm's assets over the same period. Using these values, which are observable in the market, it is possible to use the equations of the Black-Scholes option-pricing model to calculate both the value of the firm's

118. See GLEASON, *supra* note 52, at 159 (describing how “if the firm's market value falls below its obligations, the equity holders can put the firm's value to the lenders/debt holders in lieu of repayment”).

119. See generally CULP, *supra* note 67, at 385–89 (explaining the various models within the traditional credit risk-measurement framework); GLEASON, *supra* note 52, at 157–62 (describing the models for credit risk); SAUNDERS & CORNETT, *supra* note 11, at 332–35 (describing in greater detail the option models of default risk).

120. See CULP, *supra* note 67, at 385–89; SAUNDERS & CORNETT, *supra* note 11, at 334–35; Merton, *supra* note 116, at 5–6.

121. SAUNDERS & CORNETT, *supra* note 11, at 332 (noting that KMV was acquired by Moody's in 2002).

122. See GLEASON, *supra* note 52, at 158 (describing the history of KMV and its development of credit risk models based on the work of Robert C. Merton, Fischer Black, and Myron Scholes and noting that the KMV model “is notable in its simplicity” but the “complexity and challenge occurs in converting the concept into a practical application”). Note that KMV also has a related model for private company borrowers. See *id.* at 164–65 (discussing KMV's public company and private company models and noting that since “[m]ore information is available for public firms, and it typically is available sooner, . . . the private model uses more proxy information to estimate values and volatilities”).

123. *Id.* at 162–65 (discussing use of market data in KMV's model).

assets and volatility of the firm's assets over the same period.¹²⁴ Assuming that historical volatilities are accurate predictors of future volatilities, we can use them to compute the probability that, during a given future period, say the coming year,¹²⁵ the value of the firm's assets will dip below the face value of the firm's liabilities coming due during the same period. That is, we can compute the expected default frequency ("EDF").¹²⁶ Of course, the value of this prediction depends on the assumption that historical volatilities of the firm's common stock are accurate predictors of future volatilities. In fact, KMV's model predicts defaults more accurately than many of the alternative models.¹²⁷

c. The KMV Portfolio Manager Model

Risk managers care not only about the credit risk of individual positions that the firm takes, but also about the credit risk of the firm's whole portfolio of positions. Measuring and controlling this risk is more complex than merely aggregating the risks on individual positions because some credit risk is diversifiable (for instance, the bankruptcy risks of aerospace companies and those of retail restaurant chains are not perfectly correlated);¹²⁸ thus, when a firm holds a sufficiently large pool of assets exposed to credit risk, it can eliminate the nonsystematic component of credit risk.¹²⁹ Although financial firms use many models to measure portfolio-wide credit risk,¹³⁰ probably the most commonly used model is

124. See SAUNDERS & CORNETT, *supra* note 11, at 335 n.50 (explaining the mathematical details of the KMV model).

125. GLEASON, *supra* note 52, at 163 (noting that KMV's expected default frequency is usually calculated for a one-year period).

126. See SAUNDERS & CORNETT, *supra* note 11, at 335–36 (stating that, using implied values of the firm's assets and the volatility of such assets, "the likely distribution of possible asset values of the firm relative to its current debt obligations can be calculated over the next year," and "the expected default frequency . . . that is calculated reflects the probability that the market value of the firm's assets . . . will fall below the promised repayments on its short-term debt liabilities . . . in one year").

127. GLEASON, *supra* note 52, at 165–66 (discussing empirical studies measuring the accuracy of various default predictors); SAUNDERS & CORNETT, *supra* note 11, at 336 ("Simulations by KMV have shown that EDF models outperform both Z score-type models and [Standard & Poor's] ratings changes as predictors of corporate failure and distress.").

128. Put another way, what are the odds that both Google and Boeing go bankrupt in the same year? Such an eventuality is very unlikely, more unlikely than only one of them going bankrupt, unless there is an extremely severe recession. See SAUNDERS & CORNETT, *supra* note 11, at 355 (noting that "default correlations tend to be low and lie between .002 and .15," and asking, in a somewhat unfortunate example, "[W]hat is the probability that both IBM and General Motors will go bankrupt at the same time?").

129. See *id.* at 348 ("[T]he risk-return characteristics of each loan in the portfolio are a concern for the [financial firm], but the risk-return of the overall loan portfolio, with some of the risk of the individual loans diversified, affects [a financial firm's] overall credit risk exposure.").

130. For example, RiskMetrics has a model called CreditMetrics, which it describes as a "value-

KMV's Portfolio Manager,¹³¹ which builds on KMV's option model for measuring default risk as explained above.¹³²

The purpose of the model is to estimate the *efficient frontier* of the portfolio—that is, those combinations of loans that produce the highest return for a given level of risk.¹³³ Given this model, senior managers at a financial firm can determine the maximum level of credit risk that the firm will tolerate in its portfolio, and then risk managers can use the model to identify the combination of loans (varying the amounts of loans to various borrowers) to achieve the highest possible return for the given level of risk.¹³⁴ Whether with the KMV model or otherwise, this involves estimating the expected return to the lender on each loan in the portfolio, the default risk on each loan in the portfolio, and the correlation of default risks among the loans in the portfolio.¹³⁵ Computing the first involves calculating not only the net income from the loan, but also the expected loss on the loan—that is, the EDF as determined by the option model explained above, multiplied by the amount that the lender expects to lose if the loan goes into default (usually called the “loss given default” or “LGD”). The hardest part is computing the covariances of the default risks. KMV does this on the basis of historical data currently extending back

at-risk (VaR) framework applicable to all institutions worldwide that carry credit risk in the course of their business.” *The CreditMetrics Framework for Quantifying Credit Risk*, RISKMETRICS GROUP (April 2, 1997), <http://www.riskmetrics.com/publications/techdocs/cmtdovv.html>. See GLEASON, *supra* note 52, at 167–69 (discussing RiskMetric's CreditMetrics, Credit Suisse First Boston's CreditRisk+ model, and McKinsey and Company's Portfolio View model); ANTHONY SAUNDERS & LINDA ALLEN, *CREDIT RISK MANAGEMENT IN AND OUT OF THE FINANCIAL CRISIS* (3d ed. 2010) (discussing various models for measuring portfolio credit risk).

131. See *Moody's KMV Portfolio Manager*, MOODY'S KMV, http://www.moodykmv.com/products/pc_portfolioManager.html (last visited Nov. 6, 2010) [hereinafter *Moody's Portfolio Manager*] (describing general characteristics of the Moody's KMV Portfolio Manager); GLEASON, *supra* note 52, at 167 (discussing KMV's portfolio model and stating that “[m]any banks have created a credit portfolio management process based on the portfolio measures of risk and return that KMV's model provides”); SAUNDERS & CORNETT, *supra* note 11, at 353–56.

132. SAUNDERS & CORNETT, *supra* note 11, at 354 (“KMV's Portfolio Manager Model . . . uses the default probability on each loan in a portfolio to identify the overall risk of the portfolio.”).

133. See EDWIN J. ELTON ET AL., *MODERN PORTFOLIO THEORY AND INVESTMENT ANALYSIS* 68 (8th ed. 2010).

134. On efficient frontiers and portfolio theory generally, see *id.* Elton and his coauthors define the efficient frontier as “that subset of portfolios that will be preferred by all investors who exhibit risk avoidance and who prefer more return to less.” *Id.* The subject is considerably more complex than the text implies—the efficient frontier depends on, among other things, whether the portfolio holder can borrow or lend with or without risk, sell assets short, and so forth. See *id.*

135. SAUNDERS & CORNETT, *supra* note 11, at 354 (“Any model that seeks to estimate an efficient frontier for loans . . . needs to determine and measure . . . the expected return on a loan . . . , the risk of a loan . . . , and the correlation of default risks between loans made to borrowers . . .”).

through two business cycles,¹³⁶ the assumption being that future covariances will be similar to past ones. Much as with back simulations used to measure market risk, KMV's model to measure default covariances can be improved by Monte Carlo methods.¹³⁷ Thus, by using the KMV Portfolio Manager, it is possible to measure the overall credit risk in the firm's portfolio of loans and then adjust the composition of the portfolio in order to obtain the overall level of risk desired by senior managers.¹³⁸

III. LIABILITY FOR RISK-MANAGEMENT FAILURES UNDER *CAREMARK* AND *STONE*

The purpose of Part II was to convey a sense, albeit drastically simplified, of what risk-management practices at major financial firms are actually like. In short, such practices are highly sophisticated exercises in applied financial theory. Although the details vary from one type of risk to another and from model to model, in almost all cases, measurements of risk or VAR depend on the assumption that patterns in certain historical data will continue into the future. The questions of which data to select, how to adjust them, and how much confidence should be placed in the assumption that data used will accurately predict the future are necessarily business judgments. In stress testing, the role of business judgment is even more obvious. With these conclusions clearly in mind, Part III addresses the question of whether, or under what circumstances, directors of financial firms should be liable in oversight for failing to properly monitor the firm's risk-management practices. Part III begins to answer this question by (1) reviewing the principles of director oversight liability as articulated in *Caremark* and confirmed and elaborated on in *Stone v. Ritter*, and then (2) considering how those principles were applied in *Citigroup* to alleged failures by the company's board to properly oversee the firm's risk-management practices leading up to the financial crisis of 2007–2008. In particular, this part argues that the plaintiffs in *Citigroup*, as well as many of the academic commentators discussing the case, clearly do not engage

136. See GLEASON, *supra* note 52, at 167 (noting that computing default correlations is “the key that KMV uses to move from default predictions for each individual credit to predictions of default across the whole portfolio”); *Moody's Portfolio Manager*, *supra* note 131 (referring to “validated data on correlations through two business cycles”).

137. See *Moody's Portfolio Manager*, *supra* note 131 (referring to the incorporation of Monte Carlo simulations in the KMV Portfolio Manager Model).

138. For the sake of brevity and because of the increasing complexity of the mathematics, I do not discuss stress testing for credit risks in the text, but it is possible to use this method, and many large financial firms routinely do. See generally Antonella Foglia, *Stress Testing Credit Risk: A Survey of Authorities' Approaches*, 5 INT'L J. CENT. BANKING 9 (2009) (discussing stress tests for credit risks in greater detail).

with the realities of risk management as described in Part II.

A. OVERSIGHT LIABILITY UNDER *CAREMARK* AND *STONE*

Suits alleging that a board failed to properly monitor the business of the corporation are governed by the principles announced by the Delaware Court of Chancery in *Caremark*¹³⁹ and recently approved by the Delaware Supreme Court in *Stone*.¹⁴⁰ Under *Stone*, to prevail on an oversight claim, the plaintiff shareholder must prove either (1) “the directors utterly failed to implement any reporting or information system or controls” to monitor the business, or else (2) “having implemented such a system or controls, [the directors] consciously failed to monitor or oversee its operations thus disabling themselves from being informed of risks or problems requiring their attention.”¹⁴¹ There is thus a scienter element—a requirement of conscious wrongdoing—in either kind of oversight claim.¹⁴² That is, the “imposition of liability requires a showing that the directors *knew* that they were not discharging their fiduciary obligations”; that is, that they were “demonstrating a *conscious* disregard for their responsibilities.”¹⁴³

Oversight claims have fallen into two main classes. In the first group of cases, plaintiffs have argued that the board should be liable for failing to detect and prevent actions by junior employees that have caused the corporation to violate the law, thus exposing the company to criminal or civil liability.¹⁴⁴ Thus, in the classic case of *Graham v. Allis-Chalmers Manufacturing Co.*, junior employees had caused the company to engage in a price-fixing conspiracy in violation of the federal antitrust laws.¹⁴⁵ In *Caremark*, actions by the company employees had violated provisions of the Medicare and Medicaid statutes,¹⁴⁶ and in *Stone* the violations were of the Bank Secrecy Act and related regulations.¹⁴⁷ Somewhat similarly, in *Guttman v. Huang*, the corporation incurred costs when it had to restate some of its financial statements, allegedly because the directors failed to

139. *In re Caremark Int'l Inc. Derivative Litig.*, 698 A.2d 959, 970 (Del. Ch. 1996).

140. *Stone v. Ritter*, 911 A.2d 362 (Del. 2006).

141. *Id.* at 370.

142. See generally Jennifer Arlen, *The Story of Allis-Chalmers, Caremark and Stone: Directors' Evolving Duty to Monitor*, in CORPORATE LAW STORIES 323 (J. Mark Ramseyer ed., 2009) (describing the scienter requirements established for oversight cases); Hillary A. Sale, *Delaware's Good Faith*, 89 CORNELL L. REV. 456 (2004) (providing an overview of directors' fiduciary duties and discussing the duty of directors to act in good faith).

143. *Stone*, 911 A.2d at 370 (emphases added).

144. Miller, *supra* note 22, at 1157.

145. *Graham v. Allis-Chalmers Mfg. Co.*, 188 A.2d 125 (Del. 1963).

146. *In re Caremark Int'l Inc. Derivative Litig.*, 698 A.2d 959, 961–65 (Del. Ch. 1996).

147. *Stone*, 911 A.2d at 365.

prevent the officers from engaging in certain accounting irregularities.¹⁴⁸ In the second group of cases, plaintiffs have alleged that the board failed to detect and prevent fraud or similar forms of wrongdoing perpetrated by junior employees against the corporation. Thus, in *David B. Shaev Profit Sharing Account v. Armstrong*, the company (Citigroup, again, as it happens) paid out substantial sums in a settlement related to allegedly fraudulent transactions by Enron and Worldcom that its employees allegedly caused the company to facilitate.¹⁴⁹ In *Desimone v. Barrows*, the plaintiffs alleged that the board failed to detect and prevent a scheme by executives to backdate stock options.¹⁵⁰ More dramatically, in *ATR-Kim Eng Financial Corp. v. Araneta*, much as in the classic oversight case of *Francis v. United Jersey Bank*,¹⁵¹ the plaintiffs alleged that the directors failed to prevent officers or controlling shareholders from looting the corporation.¹⁵²

In any event, the crucial feature in oversight cases after *Stone* is scienter—the necessity of proving that directors knew that they were not discharging their fiduciary obligations. This is not simply a matter of proving that the board knew it had a fiduciary duty to oversee the business of the corporation and in fact failed to do so adequately; rather, it requires showing that the board knew that it was failing to adequately oversee the business of the corporation.¹⁵³ It is not knowledge of the duty that matters,

148. *Guttman v. Huang*, 823 A.2d 492, 493 (Del. Ch. 2003).

149. *David B. Shaev Profit Sharing Account v. Armstrong*, No. 1449-N, 2006 Del. Ch. LEXIS 33, at *1–2 (Del. Ch. Feb. 13, 2006), *aff'd*, 911 A.2d 802 (Del. 2006).

150. *Desimone v. Barrows*, 924 A.2d 908, 921–22 (Del. Ch. 2007).

151. *Francis v. United Jersey Bank*, 432 A.2d 814 (N.J. 1981).

152. *ATR-Kim Eng Fin. Corp. v. Araneta*, No. 489-N, 2006 Del. Ch. LEXIS 215, at *1 (Del. Ch. Dec. 21, 2006).

153. In 2008, this issue arose in the *Revlon* context as well: in *Ryan v. Lyondell Chemical Co.*, No. 3176-VCN, 2008 Del. Ch. LEXIS 105, at *3–6 (Del. Ch. July 29, 2008), *rev'd*, 970 A.2d 235 (Del. 2009), plaintiff shareholders alleged that defendant directors had breached their duty of loyalty by knowingly failing in their duty under *Revlon, Inc. v. MacAndrews & Forbes Holdings, Inc.*, 506 A.2d 173, 179 (Del. 1986), to take reasonable steps to get the best price reasonably available for shareholders when negotiating a change-of-control transaction. Although the matter is somewhat unclear, Vice Chancellor Jon W. Noble implies that, since the complaint alleged that (1) the board knew it had a *Revlon* duty to attempt to get the best price reasonably available, and (2) the board did not do enough to get this price, the complaint stated a sufficient duty of loyalty claim that the board had consciously disregarded its duties. For instance, he writes, “It suffices that, on this limited record, there exists apparent and unexplained director inaction despite [the directors] knowing that the Company was ‘in play’ and their knowing that *Revlon* and its progeny mandated certain conduct” on their part. *Ryan*, 2008 Del. Ch. LEXIS 125, at *25 (Noble, V.C.). The Delaware Supreme Court reversed the Chancery Court, stating that “the trial court equated an arguably imperfect attempt to carry out *Revlon* duties with a knowing disregard of one’s duties that constitutes bad faith.” *Lyondell Chem. Co. v. Ryan*, 970 A.2d 235, 241 (Del. 2009). I discuss the case at length in Robert T. Miller, *Lyondell Chemical Co. v. Ryan: Good Faith Comes to Revlon-Land*, 11 ENGAGE 14 (2010).

but knowledge of the failure to perform the duty in the particular circumstances. As Vice Chancellor Leo E. Strine stated in *Desimone*,

in order to state a viable *Caremark* claim . . . a plaintiff must plead the existence of facts suggesting that the board knew that internal controls were inadequate, that the inadequacies could leave room for illegal or materially harmful behavior, and that the board chose to do nothing about the control deficiencies that it knew existed.¹⁵⁴

Clearly, malfeasance of this level is rare, and proof of such malfeasance is even rarer. Hence, as Chancellor William T. Allen famously stated in *Caremark*, oversight liability “is possibly the most difficult theory in corporation law upon which a plaintiff might hope to win a judgment.”¹⁵⁵

The Delaware courts have thus set a high standard in oversight liability cases, much higher than in ordinary business judgment cases. In business judgment cases, when the plaintiff is challenging a particular decision that the board has made, the plaintiff can prevail without showing scienter by showing that the board made the business decision without first informing itself of all the material facts reasonably available.¹⁵⁶ The plaintiff need not show that the board knew it was making a decision without knowing the material facts reasonably available to it. Gross negligence, in a wholly objective sense, in gathering information before making the decision will suffice.¹⁵⁷ Why then do courts impose the scienter standard in oversight contexts?

The key reason, presumably, as Chancellor Allen remarked in *Caremark* itself, is that, in reviewing duty-of-care claims, Delaware courts do not consider the substance, content, or merits of an active business decision by the board, but only the process leading up to the decision.¹⁵⁸ That is, courts will not review whether the decision was substantively right or wrong, rational or irrational, prudent or foolish. Courts will only evaluate whether the directors informed themselves of all the material

154. *Desimone*, 924 A.2d at 940.

155. *In re Caremark Int'l Inc. Derivative Litig.*, 698 A.2d 959, 967 (Del. Ch. 1996).

156. *Smith v. Van Gorkom*, 488 A.2d 858, 872 (Del. 1985).

157. *Id.* at 873 (holding that “gross negligence is . . . the proper standard for determining whether a business judgment reached by a board of directors was an informed one”); *Aronson v. Lewis*, 473 A.2d 805, 812 (Del. 1984) (stating that the standard for director liability under the business judgment rule “is predicated upon concepts of gross negligence”).

158. *Caremark*, 698 A.2d. at 967–68. *See also* *Brehm v. Eisner*, 746 A.2d 244, 264 (Del. 2000) (stating that “‘substantive due care’ . . . [is] a concept . . . foreign to the business judgment rule,” for courts “do not measure, weigh or quantify directors’ judgments” and “do not even decide if [such decisions] are reasonable,” and so in the business judgment context, “[d]ue care in the decisionmaking context is *process* due care only”).

information reasonably available before deciding¹⁵⁹ and then made an honest judgment as to what was in the best interest of the company.¹⁶⁰ As the Delaware Supreme Court later stated, “Due care in the decisionmaking context is *process* due care only.”¹⁶¹ The reasons for not reviewing the substance of the board’s business decisions on the merits include all the traditional rationales underlying the business judgment rule—the incompetence of courts to judge business issues, the dangers of hindsight bias in reviewing failed business decisions, and the insulation of directors against the consequences of honest risk taking that has a poor result for the business.¹⁶² The theory of Delaware law is that such rationales are not applicable—or have substantially less weight—when courts review the process of decisionmaking, including the process of gathering information before making a decision.¹⁶³

159. *Van Gorkom*, 488 A.2d at 872.

160. *See Aronson*, 473 A.2d at 812 (stating that the business judgment rule is “a presumption that in making a business decision the directors of a corporation acted on an informed basis, in good faith and in the honest belief that the action taken was in the best interests of the company”).

161. *Brehm*, 746 A.2d at 264.

162. *See* William T. Allen, Jack B. Jacobs, & Leo E. Strine Jr., *Realigning the Standard of Review of Director Due Care with Delaware Public Policy: A Critique of Van Gorkom and Its Progeny as a Standard of Review Problem*, 96 NW. U. L. REV. 449, 451–52 (2002) (discussing rationales for the business judgment rule, including dangers of hindsight bias and encouraging directors to take business risks); Stephen M. Bainbridge, *The Business Judgment Rule as Abstention Doctrine*, 57 VAND. L. REV. 83, 114–15 (2004) (explaining how hindsight bias can lead a reviewing court to mistake rational decisions that result in bad outcomes for irrational decisions); Daniel R. Fischel, *The Business Judgment Rule and the Trans Union Case*, 40 BUS. LAW. 1437, 1440–42 (1985) (discussing how, in business judgment cases, it is difficult to distinguish between managers who made good decisions that had bad outcomes and managers who made bad decisions).

163. This theory is, however, open to serious question. Judge Frank H. Easterbrook and Daniel Fischel have argued that the distinction between a court reviewing business decisions on the merits and a court reviewing the process of decisionmaking leading up to the decision is illusory, because the decision as to how much information to gather before deciding is itself a business decision. In this context, value-maximizing directors will gather information only to the point that the marginal cost of gathering information equals the marginal benefit of doing so in the form of better decisionmaking. *See* FRANK H. EASTERBROOK & DANIEL R. FISCHEL, *THE ECONOMIC STRUCTURE OF CORPORATE LAW* 107–08 (1991) (“Judicial inquiry into the amount of information managers should acquire before deciding creates the precise difficulties that the business judgment rule is designed to avoid. Information is necessary for corporate managers to maximize the value of the firm. But there is a limit to how much managers should know before making a decision. . . . Information is costly, and investors want managers to spend on knowledge only to the point where an additional dollar generates that much in better decisions.”). Chancellor Allen made the same point in *In re RJR Nabisco, Inc. Shareholders Litigation*, No. 10389, 1989 Del. Ch. LEXIS 9, at *57–58 (Del. Ch. Jan. 31, 1989) (noting that “information has costs” and “the amount of information that it is prudent to have before a decision is made is itself a business judgment of the very type that courts are institutionally poorly equipped to make”). I think this is correct. The distinction that Delaware law makes may perhaps be preserved, however, by arguing that (1) decisions about which information to gather before deciding, although undoubtedly made on a cost-benefit basis, are nevertheless different from business decisions generally because they are subject to less uncertainty; and, for this reason, (2) courts are competent to review the

Now, in an oversight claim, the alleged wrongful conduct is not a decision by the board. Rather, it is mere inaction, “an unconsidered failure of the board to act.”¹⁶⁴ Since there is no decision at issue in an oversight claim, there is no procedure leading up to that decision for the court to review. As a result, the business judgment standard of procedural due care is necessarily inapplicable, and a wholly different standard must therefore apply.¹⁶⁵

In the oversight context, we are considering whether the board was at fault for its mere failure to act. If we approached this question applying an objective standard, however, the question would inevitably become whether the board in fact knew what it should have known—that is, whether the information and reporting systems the board had designed and implemented were objectively adequate or reasonable.¹⁶⁶ Regardless of whether this question is understood in a negligence or even gross negligence sense, applying the standard would require a court to review the content of the board’s decision as to which kinds of information and reporting systems to put in place.¹⁶⁷ If the point of such systems is to maximize value for the shareholders, then the issue in designing such systems is whether the benefits of the systems—in the form of improved decisionmaking by the board—exceed the costs of designing and

former, at least under a gross negligence standard, even though they are not competent to review the latter under the same standard. Nevertheless, following *Van Gorkom*, the Delaware General Assembly enacted section 102(b)(7) of the Delaware General Corporation Law, which authorizes Delaware corporations to include in their certificates of incorporation provisions that eliminate personal liability for directors for breaches of the duty of care, including the duty to be informed of all the reasonably available, material facts relevant to a business decision. See DEL. CODE ANN. tit. 8, § 102(b)(7) (2001). See generally E. Norman Veasey, Jesse A. Finkelstein & C. Stephen Bigler, *Delaware Supports Directors with a Three-Legged Stool of Limited Liability, Indemnification, and Insurance*, 42 BUS. LAW. 399 (1987) (discussing section 102(b)(7) exculpation). Since the vast majority of public companies have section 102(b)(7) provisions, the net effect, similar to the oversight context under *Caremark* and *Stone*, is that plaintiffs cannot prevail in a suit for damages if they allege nothing more than a violation of the duty of care. See *Malpiede v. Townson*, 780 A.2d 1075, 1094 (Del. 2001) (holding that a complaint that unambiguously and solely asserts only a due care claim is dismissible once the corporation’s section 102(b)(7) provision is invoked); Miller, *supra* note 22, at 1161.

164. *In re Caremark Int’l Inc. Derivative Litig.*, 698 A.2d 959, 967 (Del. Ch. 1996) (emphasis omitted). See Miller, *supra* note 22, at 1165. See generally Miller, *supra* note 26 (discussing this important difference between directorial liability for decisions actively made and unconsidered failures to act). The actively-made decision to do nothing is itself a business decision, reviewable under the business judgment rule, not principles of oversight liability. *Id.* at 944.

165. See Miller, *supra* note 22, at 1165; Miller, *supra* note 26, at 912–13.

166. See generally Miller, *supra* note 26 (discussing at length why the business judgment standard of procedural due care is inapplicable to oversight claims where there is no decision at issue).

167. See EASTERBROOK & FISCHER, *supra* note 163, at 107; Miller, *supra* note 22, at 1166. Of course, if Easterbrook and Fischel are correct, this is exactly what courts do when they review the adequacy of a board’s information gathering prior to making a business decision.

implementing them.¹⁶⁸ As Chancellor Allen put it, “the level of detail that is appropriate for such an information system is a question of business judgment.”¹⁶⁹ Thus, all the reasons that underlie the business judgment rule’s effective prohibition on courts’ reviewing the content of business decisions would apply as much to this decision as to any other. In other words, if courts were to review oversight claims under an objective standard, they would necessarily become involved in reviewing certain business decisions on the merits—namely, decisions as to what kind of information and reporting system the company should have. Since it is fundamental in Delaware law that courts will not review business decisions on the merits, applying an objective standard in oversight cases is impossible. With this objective standard unworkable, Delaware courts naturally turned to the only other kind of standard apparently available—a subjective standard. Hence, in *Caremark* and *Stone*, the rule became that directors are subject to oversight liability only if they knowingly breached their fiduciary duties, that is, only if they actually knew that they were not doing what they believed they ought to do to monitor the company.

A result of this is that any oversight liability suit is daunting for plaintiffs because they must prove scienter. Suits based on alleged failures by the board to detect and prevent excessive risk taking are novel applications of this very high standard and therefore were never very likely to succeed. As we shall see below, in the most important case brought thus far, the directors won easily.

B. THE *CITIGROUP* CASE

Part III.B discusses the *Citigroup* case, including (1) Chancellor Chandler’s opinion in the case, (2) the criticisms of that opinion, which come mostly from academic commentators, and (3) an evaluation of the case in light of the discussion of risk-measurement practices described in Part II above.

1. The Court’s Opinion in *Citigroup*

In *Citigroup*, plaintiff shareholders sued current and former directors

168. See Miller, *supra* note 26, at 939–40 (arguing that “whether an information and reporting system is reasonable is itself a substantive business decision [because,] assuming that a *reasonable* system is one that maximizes shareholder value in the long run, a system will be reasonable if the benefits of the system, in the form of improved decision-making by the board, exceed the costs of designing and implementing it [and] [m]easuring these costs and benefits and balancing the one against the other is exactly the kind of highly uncertain judgment at issue in business decisions”).

169. *Caremark*, 698 A.2d at 970.

and officers of Citigroup, alleging that they had breached their fiduciary duties by “failing to adequately oversee and manage Citigroup’s exposure to the problems in the subprime mortgage market,”¹⁷⁰ which ultimately resulted in Citigroup suffering billions of dollars in losses.¹⁷¹

These losses arose from two main sources. First, Citigroup had been in the business of purchasing asset-backed securities (“ABSs”), including subprime residential mortgage-backed securities (“RMBSs”), and repackaging them into collateralized debt obligations (“CDOs”). They then sold many of the tranches of these CDOs to others.¹⁷² This exposed Citigroup to several different kinds of risk. For one thing, creating CDOs requires large quantities of ABSs, and amassing such quantities of securities takes time. Hence, banks creating CDOs generally have to warehouse—that is, hold in their own names—significant quantities of ABSs that they intend to package into CDOs.¹⁷³ While warehousing such securities, Citigroup bore the risk that the securities would lose value, a risk usually called warehousing risk.¹⁷⁴ Because this risk arises from banks actively buying and selling (albeit in this case after repackaging) securities, warehousing risk is a form of market risk. In addition, although Citigroup sold many tranches of its CDOs to third parties, it often retained the most senior tranches (the so-called super-senior tranches), directly exposing Citigroup to market risk arising from changes in the value of such securities.¹⁷⁵ Finally, many of the tranches of the CDOs that Citigroup sold to others included a right (called a “liquidity put”) to sell the securities in the tranche back to Citigroup.¹⁷⁶ Thus, Citigroup was, in effect, guaranteeing the value of the securities. This guarantee exposed Citigroup to the risk that the value of the CDO tranches would fall and Citigroup

170. *In re Citigroup Inc. S’holder Derivative Litig.*, 964 A.2d 106, 114 (Del. Ch. 2009).

171. *Id.* at 113–14 (detailing Citigroup’s losses on collateralized debt obligations (“CDOs”) and losses arising from bringing its structured investment vehicles (“SIVs”) back onto its balance sheet).

172. *Id.* at 113 (detailing Citigroup’s involvement with CDOs). For general information on CDOs in the financial crisis, see GARY B. GORTON, *SLAPPED BY THE INVISIBLE HAND* 97–103 (2010).

173. GORTON, *supra* note 172, at 138–41. See also Robert T. Miller, *Morals in a Market Bubble*, 35 U. DAYTON L. REV. 113, 129 (2009); Gary B. Gorton, *The Subprime Panic* 28 (Yale Int’l Ctr. for Fin., Working Paper No. 08-25, 2008), available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1276047.

174. See GORTON, *supra* note 172, at 138–39; ANNUAL REPORT ON FORM 10-K OF CITIGROUP, INC., FOR 2007, at 38 (2008) [hereinafter 2007 CITIGROUP REPORT], available at <http://www.citigroup.com/citi/fin/data/k07c.pdf> (noting warehousing risk); *id.* at 49 (“As a result of downgrades and market developments during the fourth quarter of 2007, the fair value of the CDO warehouse inventory . . . declined significantly . . .”).

175. 2007 CITIGROUP REPORT, *supra* note 174, at 48 (detailing Citigroup’s retention of super-senior tranches of its CDOs).

176. *Citigroup*, 964 A.2d at 113.

would have to repurchase them at prices above their market value.

Second, Citigroup sponsored several structured investment vehicles (“SIVs”).¹⁷⁷ SIVs are special-purpose, off-balance entities controlled by a bank that acts something like traditional savings and loans. The SIV borrows short-term, usually by selling commercial paper, and invests long-term, often in RMBSs, profiting from the usual spread between long-term and short-term interest rates.¹⁷⁸ Although SIVs are most obviously exposed to interest rate risk, they are also exposed to market risk on the value of the assets they hold. When the subprime crisis began, the value of RMBSs backed by subprime mortgages declined sharply. Other market participants thus worried that SIVs holding RMBSs of declining value may have suffered losses sufficient to reduce the value of their assets below that of their liabilities. When the SIVs needed to refinance their short-term debt, they found that there was no market for their commercial paper.¹⁷⁹ When this happened to Citigroup’s SIVs, the bank provided billions of dollars of funding to the SIVs and eventually consolidated them into its own balance sheet, resulting in additional substantial losses.¹⁸⁰ Therefore, albeit somewhat indirectly, Citigroup was exposed to market risk on the value of the assets held by its SIVs.

The plaintiffs’ suit against Citigroup’s directors was derivative, the plaintiffs alleging that the directors were liable to the corporation.¹⁸¹ Therefore, when the defendant directors moved to dismiss the suit, the question before Chancellor Chandler under the applicable doctrines concerning derivative actions was whether the plaintiffs had pled particularized facts showing that the board’s conduct was “so egregious on its face that . . . a substantial likelihood of director liability therefore exists.”¹⁸² Chancellor Chandler thus began by recounting the doctrine of

177. *Id.* at 114.

178. Citigroup’s SIVs were of this kind. *See id.* On SIVs in the financial crisis generally, see GORTON, *supra* note 172, at 105–07; POSNER, *supra* note 101, at 57–58.

179. This is what happened to Citigroup’s SIVs. *Citigroup*, 964 A.2d at 114 (“The SIVs held subprime mortgages that had decreased in value, and the normally liquid commercial paper market became illiquid.”).

180. *Id.*

181. *Id.*

182. *Id.* at 121 (quoting *Aronson v. Lewis*, 473 A.2d 805, 815 (Del. 1984)). The text here elides a multitude of details irrelevant to the analysis. In derivative actions shareholder-plaintiffs “must either (1) make a pre-suit demand by presenting the allegations to the corporation’s directors, requesting that they bring suit, and showing that they wrongfully refused to do so, or (2) plead facts showing that demand upon the board would have been futile.” *Id.* at 120. *See also* *Stone v. Ritter*, 911 A.2d 362, 366–67 (Del. 2006). In *Citigroup*, the plaintiffs had not made demand on the board and so were pleading that demand would have been futile. *Citigroup*, 964 A.2d at 121. When the defendant directors moved to dismiss the suit, Chancellor Chandler was considering the motion under Court of Chancery

Caremark as approved in *Stone*, concluding that “to establish oversight liability a plaintiff must show that the directors knew they were not discharging their fiduciary obligations or that the directors demonstrated a conscious disregard for their responsibilities such as by failing to act in the face of a known duty to act.”¹⁸³

Applying this standard, Chancellor Chandler dismissed the complaint.¹⁸⁴ Chancellor Chandler first observed that the plaintiffs did “not contest that Citigroup had procedures and controls in place that were designed to monitor risk,” including an Audit and Risk Management Committee charged with “assist[ing] the board in fulfilling its oversight responsibility relating to . . . risk assessment and risk management,” which met approximately monthly.¹⁸⁵ Hence, the plaintiffs had not alleged that the directors knowingly “utterly failed to implement any”¹⁸⁶ risk-management system. Nor did the plaintiffs adequately allege that, having implemented a risk-management system, the directors “consciously failed to monitor or oversee its operations thus disabling themselves from being informed of risks or problems requiring their attention.”¹⁸⁷ Specifically, Chancellor Chandler said that the “plaintiffs’ allegations do not even specify how the board’s oversight mechanisms were inadequate or how the director defendants knew of these inadequacies and consciously ignored them.”¹⁸⁸ Instead, the plaintiffs alleged facts, mostly drawn from publicly available sources, that tended to show that it was widely known during the relevant

Rule 23.1 relating to demand futility, under which, to survive the motion to dismiss, “the complaint must plead with particularity facts showing that a demand would have been futile.” *Id.* at 120 (citing *Stone*, 911 A.2d at 367 n.9). When the underlying claim is one of oversight liability, the plaintiff must allege particularized facts that create a reasonable doubt whether the board of directors “could have properly exercised its independent and disinterested business judgment in responding to a demand.” *Id.* (quoting *Rales v. Blasband*, 634 A.2d 927, 934 (Del. 1993)). This means that plaintiffs must properly plead either that a majority of the board was interested or lacked independence (which the plaintiffs in *Citigroup* did not allege), or else plead particularized facts showing that the board’s conduct was “so egregious on its face that board approval cannot meet the test of business judgment, and a substantial likelihood of director liability therefore exists.” *Id.* at 121 (quoting *Aronson*, 473 A.2d at 815). Hence, the plaintiffs in *Citigroup* had to plead with particularity facts that, if true, would show that the Citigroup board faced a substantial threat of oversight liability. See Miller, *supra* note 22, at 1160.

183. *Citigroup*, 964 A.2d at 123 (emphasis omitted).

184. *Id.* at 123–31.

185. *Id.* at 127. In fact, like all major financial firms, Citigroup had elaborate risk-management systems in place under the direction of a chief risk officer. See 2007 CITIGROUP REPORT, *supra* note 174, at 39–65. This is hardly surprising since various elaborate kinds of risk assessment and risk management procedures are legally required of bank holding companies such as Citigroup under applicable regulations of the Federal Reserve. See 12 C.F.R. pt. 225 apps. A, G (2010) (providing guidelines for bank holding companies).

186. *Citigroup*, 964 A.2d at 123 (quoting *Stone*, 911 A.2d at 370).

187. *Stone*, 911 A.2d at 370.

188. *Citigroup*, 964 A.2d at 128.

period that the subprime market was deteriorating.¹⁸⁹ Chancellor Chandler, however, held that such facts did not show that the directors knew that the risk-management systems they had created were inadequate and were consciously disregarding a duty to prevent Citigroup from suffering losses.¹⁹⁰ At most, the plaintiffs had alleged “that the directors had made bad business decisions.”¹⁹¹ Although Chancellor Chandler never explicitly states it, “the plaintiffs simply did not plead facts tending to show scienter, and so their complaint did not state a claim under *Caremark* and *Stone*.”¹⁹²

2. Evaluation of the *Citigroup* Case

It would be difficult to argue that the *Citigroup* plaintiffs had adequately pleaded scienter and thus also difficult to argue with the outcome of the case. Not surprisingly, therefore, most criticism of this decision, as discussed below in Part III.B.3, focuses on possible changes to the *Caremark* standard itself, not Chancellor Chandler’s application of that standard. Nevertheless, before turning to those criticisms, it is important to examine carefully what the plaintiffs did not plead in *Citigroup*, beyond their failure to plead scienter. In this case, what the plaintiffs did *not* say is actually more illuminating than what they did say.

a. Plaintiffs’ Ignorance of Risk-Management Systems

Although plaintiffs’ failure to plead scienter was legally dispositive in *Citigroup*, from the point of view of someone familiar with contemporary risk-management practices at major financial firms, the most striking aspect of the complaint in the case is its utter and complete failure to criticize in any meaningful way any of Citigroup’s actual risk-measurement models or risk-management practices. Citigroup’s annual reports describe these practices in considerable detail, distinguishing various kinds of risks that the company faces and explaining some of the techniques that Citigroup uses to measure and manage such risks.¹⁹³ There are separate sections on market risk¹⁹⁴ and credit risk (which is broken down into consumer credit and corporate credit),¹⁹⁵ as well as sections on real estate

189. *Id.* at 127–28.

190. *Id.* at 128.

191. *Id.*

192. Miller, *supra* note 22, at 1160.

193. 2007 CITIGROUP REPORT, *supra* note 174, at 39–65; ANNUAL REPORT ON FORM 10-K OF CITIGROUP, INC. FOR 2006, at 59–77 (2007) [hereinafter 2006 CITIGROUP REPORT], available at <http://citigroup.com/citi/fin/data/k06c.pdf>.

194. 2007 CITIGROUP REPORT, *supra* note 174, at 61–64; 2006 CITIGROUP REPORT, *supra* note 193, at 71–74.

195. 2007 CITIGROUP REPORT, *supra* note 174, at 40–60; 2006 CITIGROUP REPORT, *supra* note

risks,¹⁹⁶ risks from derivatives,¹⁹⁷ operational risks,¹⁹⁸ and country and cross-border risks.¹⁹⁹ Citigroup also provided extensive quantitative disclosure, for example, explaining the distribution of its loan portfolio by types of borrower (consumer versus corporate), geography, lines of business for corporate borrowers, and so forth.²⁰⁰ From the quantitative disclosure that Citigroup provided, it would certainly have been possible to compute, at least in some approximate way, Citigroup's exposure to certain kinds of risks and then to argue that these were too high (for example, perhaps Citigroup's exposure was greater than that of many or most of its peer banks).²⁰¹ The plaintiffs never attempted to do anything like this.

Furthermore, although Citigroup did not disclose the details of its risk-measurement models, it did disclose enough information about them to give the plaintiffs (or anyone else reading its securities laws filings) a general idea of the kinds of risk analyses it was performing. For example, the section on market risk, which was the most relevant section to the allegations in the *Citigroup* case, discloses that each business unit in Citigroup "is required to establish, with approval from independent market risk management, a market risk limit framework, including risk measures and controls, that clearly defines approved risk profiles and is within the parameters of Citigroup's overall risk appetite."²⁰² Across business units, however, "Market risks are measured in accordance with established standards to ensure consistency . . . and the ability to aggregate risk."²⁰³ This means that Citigroup was using firm-wide market-risk measurement models like those discussed in Part II.B above. Indeed, Citigroup discloses that its primary market risk model was a VAR model that identifies risk factors, computes their volatilities and correlations (presumably on the basis of historical data), and then calculates (presumably using a Monte

193, at 60–70.

196. 2007 CITIGROUP REPORT, *supra* note 174, at 48–53.

197. *Id.* at 57–59; 2006 CITIGROUP REPORT, *supra* note 193, at 68–70.

198. 2007 CITIGROUP REPORT, *supra* note 174, at 64; 2006 CITIGROUP REPORT, *supra* note 193, at 75.

199. 2007 CITIGROUP REPORT, *supra* note 174, at 65; 2006 CITIGROUP REPORT, *supra* note 193, at 76–77.

200. 2007 CITIGROUP REPORT, *supra* note 174, at 40–48.

201. Other major banks provide disclosure about their risk-management practices comparable to Citigroup's. *See, e.g.*, ANNUAL REPORT ON FORM 10-K OF BANK OF AMERICA CORPORATION FOR 2006, at 58–81 (2007), *available at* http://media.corporate-ir.net/media_files/irol/71/71595/reports/2006AR.pdf; ANNUAL REPORT ON FORM 10-K OF JPMORGAN CHASE AND CO. FOR 2006, at 61–82 (2007), *available at* http://files.shareholder.com/downloads/ONE/986281963x0x86632/a2d29e0f-1aa6-40d7-bc00-d762cbca5c6b/JPMC_AR06_complete.pdf.

202. 2006 CITIGROUP REPORT, *supra* note 193, at 71.

203. *Id.*

Carlo or similar simulation) a one-day VAR for Citigroup's portfolio "at a 99% confidence level."²⁰⁴ Citigroup's risk managers also performed back simulations to check the accuracy of the VAR computations produced in its primary model.²⁰⁵

If the plaintiffs in *Citigroup* had consulted the risk-management literature or experts familiar with that literature, they could have included in their complaint recitations of the well-known limitations of the kinds of models Citigroup was using.²⁰⁶ Moreover, if the plaintiffs had availed themselves of their right as shareholders to inspect the books and records of the company,²⁰⁷ they could have learned some of the details of Citigroup's VAR model. Then, presumably with the help of professional risk managers as expert witnesses, they could have attempted to make detailed criticisms of the model. Plaintiffs also could have criticized some of the obvious business judgments involved in applying the models. For example, in applying the VAR market risk model, Citigroup calculated the 1 percent worst case. Plaintiffs could have said that Citigroup should have been calculating the .01 percent worst case, so that its confidence level in the VAR would have been 99.99 percent rather than just 99 percent. Plaintiffs did none of this. Thus, they never even attempted to formulate a meaningful criticism of Citigroup's risk-management systems. The conclusion is that the plaintiffs in *Citigroup* simply did not know the first thing about risk management at financial firms.

b. Plaintiffs' Misunderstanding of Risk-Management Systems

Rather than making a serious effort to understand risk-management systems, the plaintiffs in *Citigroup*, as Chancellor Chandler very aptly put it, "seem[ed] to hope the Court [would] accept the conclusion that since the Company suffered large losses, and since a properly functioning risk-management system would have avoided such losses, the directors must have breached their fiduciary duties in allowing such losses."²⁰⁸ Now, Chancellor Chandler did not accept that argument, mostly because the

204. *Id.* at 73.

205. *Id.* at 74.

206. This is not to say that Citigroup's models were deficient, merely that all models have known limitations, and the plaintiffs could have identified these and included explanations of them in their complaint.

207. See DEL. CODE ANN. tit. 8, § 220 (2001) (allowing inspection by stockholders of books and records of the corporation). The Delaware courts have repeatedly suggested that plaintiffs, rather than filing complaints lacking sufficient factual allegations to withstand a motion to dismiss, should use section 220 to obtain relevant information on which to base a complaint. See, e.g., *Saito v. McKesson HBOC, Inc.*, 806 A.2d 113, 117–18 (Del. 2002).

208. *In re Citigroup Inc. S'holder Derivative Litig.*, 964 A.2d 106, 128 (Del. Ch. 2009).

relevant legal standard under *Caremark* required a showing of scienter. What mattered legally was not whether Citigroup had a properly functioning risk-management system, but whether the directors knew that the system was not functioning properly. But—and this is the more important point—the argument presupposes a profound misunderstanding of the nature and purpose of risk-management systems, for the second premise in the argument—that “a properly functioning risk-management system would have avoided such losses”—is simply, indeed flagrantly, false.

On the contrary, properly functioning risk-management systems do not eliminate risk and do not eliminate the possibility of losses, even catastrophic losses.²⁰⁹ At best, they measure risk and assist the firm in taking on the level of risk it desires. Imagine a financial firm with capital of about \$100 billion using a VAR model of the kind discussed in Part II.B.1 above. Its board of directors decides that a tolerable daily risk level is a \$100 million VAR with 99 percent confidence. Even if all the assumptions on which such models are based are correct—in particular that the historical pattern of covariances of risk factors will continue into the future—this means that the firm is taking a 1 percent chance every day of losing more than \$100 million. Put another way, every three to four months, the firm can expect to have a trading day in which its losses exceed \$100 million. Large financial firms also use VAR models to calculate their annual risk of bankruptcy. For example, such a firm might accept a 0.5 percent annual chance of a loss sufficient to bankrupt the company.²¹⁰ Thus, the bank is accepting the risk that it should expect to go bankrupt once in two hundred years.²¹¹ If the bank happens to suffer such a loss this

209. See Bainbridge, *supra* note 4, at 982 (“At best, an effective risk management program can prevent risks from materializing and perhaps limit the impact of those that do materialize.”). Generally speaking, risk-management systems at financial firms never *prevent* risks from materializing, for most of the risks with which they are concerned are well beyond the ability of the firm to control (for example, changes in interest rates, changes in market conditions, bankruptcies or other credit impairments of counterparties, and so forth). By correctly identifying and quantifying risks, risk-management systems can assist managers in adjusting the firm’s holdings to reduce exposure to risks or in hedging against risks retained.

210. The numbers in this example are intended to be realistic and were derived from off-the-record conversations with risk managers at large commercial banks.

211. Although large financial firms routinely use VAR models to measure the risk of bankruptcy, and although the boards of such firms routinely set risk tolerances for VARs sufficient to bankrupt the firm, these facts are apparently not widely appreciated. For example, even such a knowledgeable and sophisticated observer as Judge Richard A. Posner says, “Though I am not certain of this, I believe that the senior executives of the major banks were aware that there might be at least a small probability of bankruptcy as a result of the risks they were taking.” POSNER, *supra* note 1, at 321. In reality, such executives set the firm’s tolerance of the risk of bankruptcy and monitor the VAR in relation to that risk

year, that is no proof that its risk-management systems were at fault.

Consequently, the materialization of risks and consequent losses prove nothing about supposed deficiencies in risk management.²¹² In reality, proving the existence of such deficiencies is much more difficult. Such proof would necessarily involve criticizing the decisions of risk managers in relation to the standards generally upheld in the risk-management profession. For example, on a technical level, we might be able to show that a company's risk-management models incorporate statistical or other mathematical errors, though this is extraordinarily unlikely for a major financial firm. More plausibly, we might be able to show that the models incorporate statistical assumptions that can be shown to be false with respect to the data to which they are applied. For instance, some versions of the RiskMetrics market-risk model assume that asset returns are normally distributed, when in fact many assets (such as most fixed-income assets) have returns that are not so distributed (generally because they have "fat tails").²¹³ Applying such a model to such data will systematically underestimate extreme negative outcomes. If risk managers, or perhaps the firm's directors, were unaware of this, this failure could be negligent or, if the problem with the model were widely known, even grossly negligent. If, however, the firm's risk managers were fully aware of the problem and used the model anyway, it could be that this was a perfectly sensible choice. For instance, there may have been no more sophisticated models available that corrected for the problem or, if there were, such models may have had their own limitations and the risk managers determined that it was better to use the first model, even given its limitations, than to use any of the others given their limitations. In an appropriate case, therefore, it may be possible to criticize a firm's risk-management practices in a meaningful way, but doing so requires a sophisticated understanding of risk management.

3. Academic Criticisms of *Citigroup*

As noted above, academic criticism of Chancellor Chandler's opinion in *Citigroup* has generally not concerned the Chancellor's application of

every day.

212. See, e.g., *Citigroup*, 964 A.2d at 126 ("When investments turn out poorly, it is possible that the decision-maker evaluated the deal correctly but got 'unlucky' in that a huge loss—the probability of which was very small—actually happened."); Bainbridge, *supra* note 4, at 988 ("An institution that suffered a large loss because a certain risk materialized did not necessarily have a risk management failure.").

213. See JORION, *supra* note 57, at 87, 262–66; SAUNDERS & CORNETT, *supra* note 11, at 270 n.10.

the scien-ter-based *Caremark* standard to the facts of the case. Rather, the criticisms have included (1) an argument that Chancellor Chandler wrongly held that all *Caremark* claims based on risk-management failures must necessarily fail; (2) an argument that Chancellor Chandler missed an opportunity to set legal standards for risk management, or at the very least, to encourage better risk-management practices; and (3) an argument that *Caremark* should be adapted in various ways to apply to alleged failures by the board to oversee risk-management practices.

a. Cognizability of Oversight Claims for Risk-Management Failures

Eric J. Pan suggests that Chancellor Chandler actually held that, as a matter of law, *Caremark* claims cannot be predicated on failures by the board to monitor risk.²¹⁴ To be sure, the Chancellor emphasized the novelty of such applications of *Caremark*,²¹⁵ and there is language in the opinion suggesting that he is highly dubious about such claims. For example, he says that “[t]o the extent the Court allows shareholder plaintiffs to succeed on a theory that a director is liable for a failure to monitor business risk, the Court risks undermining the well settled policy of Delaware law” by inviting courts to review business judgments on their merits.²¹⁶ He even goes so far as to say that “[o]versight duties under Delaware law are not designed to subject directors . . . to *personal liability* for failure to predict the future and to properly evaluate business risk.”²¹⁷ Such language, taken in isolation, can plausibly be read to mean that oversight claims based on alleged failures by the board to properly monitor risk are noncognizable at law.

Nevertheless, such a reading overinterprets these passages.²¹⁸ Chancellor Chandler never says that oversight claims based on failures to monitor risk fail as a matter of law or even that, for some other reason, they necessarily must fail. On the contrary, his opinion presupposes just the opposite. He expressly says that Delaware law “place[s] an extremely high burden on a plaintiff to state a claim for personal director liability for a

214. Pan, *supra* note 33, at 738 (“Chancellor Chandler rejected the plaintiffs’ claims on two grounds,” one of which was that “a board should not be held liable for failure to monitor business risk.”).

215. *Citigroup*, 964 A.2d at 123 (“Plaintiffs’ theory . . . is a bit of a twist on the traditional *Caremark* claim [because it is] based on defendants’ alleged failure to properly monitor Citigroup’s *business risk* . . .”).

216. *Id.* at 126.

217. *Id.* at 131.

218. See Bainbridge, *supra* note 4, at 979 (noting that, although Chancellor Chandler “seem[ed] tempted simply to exclude [risk-management] claims from *Caremark*’s ambit,” nevertheless “he stated that it might be possible for plaintiff to state such a claim ‘under some set of facts’” (citation omitted)).

failure to see the extent of a company's business risk,"²¹⁹ which, while emphasizing the difficulty, presupposes the possibility of carrying that burden. Again, Chancellor Chandler suggests that, although the plaintiffs in *Citigroup* failed to do so, "it may be possible for a plaintiff [alleging an oversight breach related to a failure to monitor risk] to meet the burden under some set of facts."²²⁰ Moreover, had Chancellor Chandler held that such claims fail as a matter of law, his application of *Caremark* principles to the facts of the case would have been superfluous.²²¹ Hence, under *Caremark*, oversight claims predicated on alleged failures to monitor the firm's risk-management practices are not precluded; rather, like all oversight claims, they are just very difficult for plaintiffs to win, primarily because of the scienter requirement common to all oversight claims.

b. Judicial Setting of Risk-Management Standards

Based on what appears to be a misunderstanding of Chancellor Chandler's opinion, Pan further argues that the result in *Citigroup* must be mistaken because "it seems fantastic that the duty to monitor . . . incentivizes boards to take no responsibility for the business results of the company—a complete disregard for the principle that the corporation shall be managed by or under the direction of the board."²²² He believes, rather, that Delaware courts "should begin speaking out about the importance of a board's duty to monitor and to back up their exhortations by expanding the scope and application of the duty to monitor in future cases,"²²³ although he never says specifically how the duty should be expanded. J. Robert Brown goes further. He thinks that the Court of Chancery should "define meaningful [risk-management] responsibilities by the board" and thus squandered an excellent opportunity to do so in *Citigroup*.²²⁴ Like Pan, however, Brown offers no specifics on how the Court of Chancery should go about defining the board's risk-management responsibilities.

As explained above, the holding in *Caremark* turned on the plaintiffs' failure to adequately plead scienter. Regardless of what the Delaware courts may hold concerning what a board should do to monitor risk, as long

219. *Citigroup*, 964 A.2d at 125.

220. *Id.* at 126.

221. *See id.* at 126–31.

222. Pan, *supra* note 33, at 739. Simply as a matter of logic, at worst, *Citigroup* gives boards no incentive to monitor risk; it does not give boards an incentive "to take no responsibility" for such matters. A legal rule may fail to encourage certain conduct without thereby discouraging it.

223. *Id.* at 740.

224. Brown, *supra* note 35.

as the scienter requirement of *Caremark* remains, plaintiffs will very likely lose *Caremark* suits based on failures of the board to monitor risk, just as plaintiffs routinely lose *Caremark* suits based on failures of the board to detect and prevent illegality and fraud by corporate employees.²²⁵ This is true because conscious wrongdoing is rare, and proof of conscious wrongdoing is even rarer. Neither Pan nor Brown expressly say that the Delaware Supreme Court should overrule *Stone* to eliminate scienter as an element of an oversight claim, but any meaningful expansion of oversight liability in connection with risk-management failures would plainly require that this be done. This aspect of the problem will be more fully addressed in Part IV.A. For now, however, the issue is whether, leaving the scienter issue aside, Delaware courts could reasonably be expected to determine risk-management standards for major financial firms.

That question answers itself. Some of the critics of the *Citigroup* case seem to think that the boards of major financial firms like Citigroup did virtually nothing to monitor the risks their firms were running. In fact, as is apparent from their securities filings, such firms, Citigroup included, maintain a vast risk-management apparatus, including hundreds of professionals and immense information technology resources, and the boards of such firms regularly receive reports from risk managers. At best, therefore, the question is not whether such firms should manage their risks, but how they ought to do so. As the very limited and vastly simplified review of risk-measurement models in Part II.B shows, however, creating and implementing risk-management systems requires substantial financial, statistical, and mathematical sophistication. Indeed, even an introductory course in risk management is generally an upper-level MBA course and presupposes significant knowledge of finance and statistics. That is why banks like Citigroup must employ literally hundreds of individuals involved in risk management, many of whom have advanced degrees in finance, statistics, and mathematics. It is absurd to think that a single lawyer, even one who happens to be a vice chancellor or chancellor on the Delaware Court of Chancery, could possibly decide what kinds of risk-

225. Plaintiffs have repeatedly lost on oversight claims. *See, e.g.*, *Stone v. Ritter*, 911 A.2d 362, 373 (Del. 2006) (discussing oversight claims related to criminal conduct by corporate employees); *Desimone v. Barrows*, 924 A.2d 908, 940 (Del. Ch. 2007) (same); *David B. Shaev Profit Sharing Account v. Armstrong*, No. 1449-N, 2006 Del. Ch. LEXIS 33, at *19–20 (Del. Ch. Feb. 13, 2006) (discussing oversight claims related to facilitating fraudulent transactions by third parties); *Guttman v. Huang*, 823 A.2d 492, 506–07 (Del. Ch. 2003) (discussing oversight claims related to accounting irregularities); *In re Caremark Int'l Inc. Derivative Litig.*, 698 A.2d 959, 971 (Del. Ch. 1996) (discussing oversight claims related to criminal conduct by employees). *But see* *Fin. Corp. v. Araneta*, No. 489-N, 2006 Del. Ch. LEXIS 215, at *76–77 (Del. Ch. Dec. 21, 2006) (discussing oversight claims involving looting of a close corporation).

measurement models a financial firm should use, how they should be implemented, how they should be updated (for financial firms are constantly updating and refining their models),²²⁶ how the proprietary data sets on which they are sometimes based should be collected, and what risk tolerances they should be used to calculate.

In practice, even the primary regulator of such firms does not attempt to do this. That is, although the regulations of the Federal Reserve require bank holding companies to employ various risk-based measures in computing their regulatory capital,²²⁷ the Federal Reserve regularly allows such firms to design and implement their own models and measurement approaches,²²⁸ which are invariably more sophisticated than those otherwise required by the Federal Reserve. This is why Alan Greenspan has written that “regulators . . . know far less than private-sector risk managers” and “take their cues from private-sector practitioners,” so that any new risk-management regulations “will reflect the private sector’s already revised market practices.”²²⁹ If the Federal Reserve—with a full-time staff of thousands—cannot dictate good risk-management practices, *a fortiori* a lone judge, even one in Delaware, cannot do so either.

Indeed, if a court attempted to specify the substantive risk-management responsibilities of the board, it could surely do no more than announce some generalized, almost platitudinous standard—for instance, that boards should adopt “reasonable” risk-management practices consistent with the firm’s business and other relevant factors. Such vacuous guidance from a court would not improve actual risk-management practices at all. Furthermore, if a court did such a thing, then in subsequent suits based on alleged failures by the board to comply with the duty the court had announced, the issue would become whether the board’s actual risk-management practices were reasonable in the circumstances. The court would then be required to determine this issue—which would involve

226. See, e.g., 2007 CITIGROUP REPORT, *supra* note 174, at 48 (noting how methodology of estimating the fair value of CDOs is “subject to continuing refinement”).

227. Capital Adequacy Guidelines for Bank Holding Companies: Risk-Based Measure, 12 C.F.R. pt. 225 app. A (2010).

228. See Capital Adequacy Guidelines for Bank Holding Companies: Internal-Ratings-Based and Advanced Measurement Approaches, 12 C.F.R. pt. 225 app. G (2010); SAUNDERS & CORNETT, *supra* note 11, at 596–99 (discussing how the Federal Reserve allows bank holding companies meeting certain criteria to use their own internal risk-measurement models and risk-management processes to comply with regulatory requirements).

229. GREENSPAN, *supra* note 2, at 524. Greenspan does not mention it, but the simple fact that there are so many more private-sector professionals than regulators and these private-sector professionals are so much better compensated than the regulators practically guarantees that the private-sector professionals will always know much more than the regulators.

reviewing for reasonability all the substantive risk-management decisions of the firm, decisions that the court is unequipped to make. There would likely be a battle of experts, but in practice the result would almost always be that the court would defer to the judgments of the board as to which risk-management models and systems were reasonable, and so, undoubtedly after protracted litigation, the directors would prevail unless they had not implemented the kinds of systems they themselves thought they should—unless, that is, they had consciously disregarded their duties. That, of course, is the current law under *Caremark*.

c. Adapting *Caremark* to the Risk-Management Context

Acknowledging that Chancellor Chandler's opinion in *Citigroup* leaves open the possibility that, in an appropriate case, a plaintiff could prevail on the theory that the board failed to properly monitor the risks the company was running,²³⁰ Stephen Bainbridge argues that risk-management and legal-compliance systems are not fundamentally different and that "there are no reasonable grounds in the *Caremark* opinion for limiting the obligations created therein to legal and accounting compliance."²³¹ Nevertheless, he also thinks that risk management differs in degree from legal-compliance and internal accounting controls and that "[s]ome of these differences are significant enough to justify being factored into the *Caremark* analysis."²³²

First, Bainbridge argues that "best practices with respect to enterprise risk management are still evolving" and will "vary from firm to firm."²³³ Hence, "courts should be extremely leery about rendering opinions perceived as creating a roadmap for approaching risk management" because in so doing, "the evolutionary market processes by which optimal best practices emerge may be aborted."²³⁴ This is true, but it seems to overlook the key point about *Caremark*—namely, that directors are liable under *Caremark* only for knowingly failing to do what they themselves think they should do, not for failing to do what a court thinks they should do. That is, even with respect to legal-compliance and accounting controls, there is no duty under *Caremark* for a board to implement any particular kind of system; so long as the board implements some system, the court

230. Bainbridge, *supra* note 4, at 979 ("[Chancellor Chandler] stated that it might be possible for plaintiff to state [a *Caremark* claim predicated on failures to monitor risk management] 'under some set of facts.'" (citation omitted)).

231. *Id.*

232. *Id.* at 981.

233. *Id.* at 982.

234. *Id.*

will inquire only whether the board knew it was failing to use the system to monitor the business.²³⁵ Unless the Delaware Supreme Court overrules *Stone* and eliminates the scienter requirement in a wholesale revision of its principles of directorial oversight, there is no danger of courts dictating to businesses what kinds of risk-management systems they should have. There is no place for such an outcome within *Caremark*.

Noting that obtaining more than the risk-free return requires that the firm bear some risk,²³⁶ Bainbridge also argues that “[r]isk management necessarily overlaps with risk taking” because both concern managing the firm’s activities to obtain the highest possible return for a given level of risk.²³⁷ He thus agrees with Chancellor Chandler that since “Citigroup was in the business of taking on and managing investment and other business risks[,] [t]o impose oversight liability on directors for failure to monitor ‘excessive’ risk would involve courts in conducting hindsight evaluations of decisions at the heart of the business judgment of directors.”²³⁸ In other words, because risk-management decisions are much more closely tied to decisions trading off risk and return than are legal-compliance or accounting-control decisions,²³⁹ and because decisions trading off risk and return are the core decisions protected by the business judgment rule, *Caremark* claims should not be permitted to be predicated on risk-management failures “absent egregious, firm-specific facts.”²⁴⁰

As noted above, the design and implementation of legal-compliance and accounting-control systems also involve risk and return decisions, for such systems enhance shareholder value only to the point that the marginal benefits of investing in a more elaborate system (such benefits coming in the form of reduced losses) exceed the costs.²⁴¹ Nevertheless, Bainbridge is surely right that, in this context, risk-management systems and legal-

235. *Stone v. Ritter*, 911 A.2d 362, 370 (Del. 2006) (holding that directors will be liable only if “(a) the directors utterly failed to implement any reporting or information system or controls [to monitor the business], or (b) having implemented such a system or controls, [the directors] consciously failed to monitor or oversee its operations thus disabling themselves from being informed of risks or problems requiring their attention”). See also Miller, *supra* note 22, at 1156–57 (discussing the holdings in *Caremark* and *Stone*); *supra* Part III.A.

236. Bainbridge, *supra* note 4, at 983 (stating that “shareholders . . . prefer high return projects,” and that since “risk and return are directly proportional . . . implementing that preference necessarily entails choosing risky projects”).

237. *Id.* at 984.

238. *In re Citigroup Inc. S’holder Derivative Litig.*, 964 A.2d 106, 131 (Del. Ch. 2009).

239. See Bainbridge, *supra* note 4, at 986 (“[S]ubstantive analysis of board decisions with respect to the nature, scope, and content of risk management programs are themselves business decisions of the sort protected by the business judgment rule.”).

240. *Id.* at 990.

241. See Miller, *supra* note 26, at 940.

compliance and accounting controls do differ significantly in degree. That said, the point again seems largely irrelevant in light of *Caremark*'s scienter requirement, for if courts were going to pass on the adequacy of both a company's legal-compliance system and its risk-management system, the latter would indeed pose a much more serious danger than the former of a court usurping the business judgment of the board. Under *Caremark*, however, the court does not pass on the adequacy of the system the board has created, whether it be a legal-compliance system, an accounting-control system, or indeed a risk-management system.²⁴² Provided that the board has implemented some system, the only issue is whether the board consciously disregarded its duty to use the system it created. This standard is as directly applicable to risk-management systems as to legal-compliance or accounting-control systems, and it poses as little danger—more accurately, none whatsoever—of a court usurping the board's business judgment. Because of *Caremark*'s scienter requirement, the question before the court is always whether the board thought it was doing its duty, not whether its understanding of its duty was correct. Hence, the nature of that underlying duty—whether it be to monitor the corporation's compliance with law, accounting controls, risk-management systems, or anything else—is irrelevant.

IV. THE INEFFICIENCY OF EXPANDING OVERSIGHT LIABILITY FOR RISK-MANAGEMENT FAILURES

We have primarily seen three things so far: (1) risk management is a difficult, quantitatively intensive business that relies essentially on business judgments, mostly concerning how past market conditions will continue into the future; (2) *Caremark* claims predicated on alleged failures of the board to properly monitor the firm's risk-management practices are likely to fail because of the scienter requirement of such claims; and (3) nevertheless, there has been significant desire, among both public officials and academics, to expand *Caremark* liability to encourage directors at major financial firms to improve risk-management practices to curtail perceived excessive risk taking. This part argues that expanding oversight liability beyond its current *Caremark* form is fundamentally

242. Bainbridge, of course, is aware of this and mentions it elsewhere in his article. See Bainbridge, *supra* note 4, at 985 ("This requirement follows inexorably from *Caremark*'s dictum that 'the level of detail that is appropriate for [the requisite] information system is a question of business judgment.' Courts are not to second-guess a board's determination that the company's risk management programs are adequate." (alteration in original) (footnote omitted)). He does not, however, connect the point with regard to the business judgment involved in accounting-control or legal-compliance systems with his concerns about courts potentially invading these business judgments of the board.

misguided. This argument has two main parts—one legal and one economic. Legally, any meaningful expansion of *Caremark* liability would amount to a revolution in Delaware law tantamount to repealing the business judgment rule, a result that would be so obviously inefficient as to be incontrovertibly out of the question. Economically, even apart from the inefficiencies involved in repealing the business judgment rule, the desired expansion of *Caremark* to control so-called excessive risk taking would be misguided because the kinds of excessive risk taking that expanded oversight liability are not the kinds of excessive risk taking that may have contributed to the financial crisis.

A. EXPANDING *CAREMARK* WOULD REPEAL THE BUSINESS JUDGMENT RULE

To see why expanding the oversight liability of boards beyond the *Caremark* doctrine would effectively repeal the business judgment rule, consider what changes in the law would be needed to effect such an expansion. First, the Delaware courts would have to eliminate the scienter element from oversight liability claims. That is, assuming Delaware courts adopted some substantive standard for director oversight such as gross or simple negligence, if liability still required scienter, then boards would still almost always prevail. Under *Caremark* and *Stone*, directors win oversight cases so easily because plaintiffs cannot prove scienter, not because they cannot prove that directors failed to meet some substantive standard of care. Hence, introducing a substantive standard of care, whatever it may be, but leaving the scienter element in place will do nothing to help plaintiffs win cases. This would be especially true when the alleged failure of oversight concerns risk-management failures, for, as explained above, directors are very unlikely to have the detailed knowledge necessary to identify failures of such practices; hence, in most cases in which a board would have failed to properly oversee the firm's risk-management practices (whatever "properly" might mean in such cases), the directors would not have consciously disregarded their duties. Even in the rare case in which directors were consciously disregarding their duties, they would still likely prevail, for proving scienter would be next to impossible absent the proverbial smoking-gun email or similar proof of a subjective state of mind.²⁴³ Hence, the first requirement of any meaningful expansion of oversight liability is an elimination of *Caremark*'s scienter element.

So, although it is very unlikely, assume that the Delaware Supreme

243. Miller, *supra* note 22, at 1161.

Court overrules *Caremark* and *Stone* and holds that a board may be liable in oversight merely for failing adequately—in some sense of the word—to monitor the business of the company, in particular for failing adequately to oversee the company’s risk-management practices. That is, assume that there could be a pure duty-of-care claim in oversight, a claim that would not require for liability any subjective state of mind on the part of the defendant directors. Would directors then have added incentives to monitor the risk-management practices of the company? In general, they would not. The reason is that section 102(b)(7) of the Delaware General Corporation Law authorizes Delaware corporations to include in their certificates of incorporation provisions eliminating the personal liability of directors for breaches of the duty of care.²⁴⁴ Under *Malpiede v. Townson*, a claim for damages against directors sounding only in the duty of care is immediately dismissible if the corporation has a section 102(b)(7) provision in its certificate of incorporation.²⁴⁵ As is well known, the Delaware General Assembly enacted this provision²⁴⁶ in response to the Delaware Supreme Court’s decision in *Smith v. Van Gorkom*,²⁴⁷ and for many years now “all

244. More precisely, section 102(b)(7) provides that a Delaware corporation’s certificate of incorporation may contain

[a] provision eliminating or limiting the personal liability of a director to the corporation or its stockholders for monetary damages for breach of fiduciary duty as a director, provided that such provision shall not eliminate or limit the liability of a director: (i) For any breach of the director’s duty of loyalty to the corporation or its stockholders; (ii) for acts or omissions not in good faith or which involve intentional misconduct or a knowing violation of law; (iii) under § 174 of [the Delaware General Corporation Law]; or (iv) for any transaction from which the director derived an improper personal benefit.

DEL. CODE ANN. tit. 8, § 102(b)(7) (2001). Even when shielded by a 102(b)(7) provision, therefore, directors may be liable for oversight claims under *Caremark* and *Stone*, for such claims, since they involve scienter, are “acts or omissions not in good faith.” *Id.* If, however, the Delaware Supreme Court eliminated the scienter element and allowed pure duty of care claims in oversight, a section 102(b)(7) provision would clearly eliminate the directors’ personal liability in damages for breaching their duty of care. Of course, a section 102(b)(7) provision does not eliminate the duty of care; rather, it eliminates the directors’ personal liability for breaches of that duty. Hence, even with a section 102(b)(7) provision, equitable relief is still available to plaintiffs. *See generally* Veasey, Finkelstein & Bigler, *supra* note 163 (discussing section 102(b)(7) exculpation). In the context of oversight liability, however, equitable relief is likely irrelevant. Although it is conceivable that, before a loss occurs, a plaintiff could sue the board alleging that it was not properly monitoring the company and seeking an injunction that the board comply with its fiduciary duties in the future, nevertheless by far the most likely scenario is that, after a loss occurs, the plaintiff would sue alleging that the board in the past failed in its duty to monitor and thus should be liable in damages. As explained in the text, a section 102(b)(7) provision would make the latter kind of claim immediately dismissible.

245. *See Malpiede v. Townson*, 780 A.2d 1075, 1093 (Del. 2001) (holding that a complaint that unambiguously and solely asserts only a due care claim is dismissible once the corporation’s section 102(b)(7) provision is invoked); Miller, *supra* note 22, at 1161–62.

246. Act of June 18, 1986, ch. 289, 65 Del. Laws 544.

247. *Smith v. Van Gorkom*, 488 A.2d 858, 874–75 (Del. 1985). Criticisms of the opinion are now standard. *See generally* Fischel, *supra* note 162 (arguing that the court in *Van Gorkom* erred in holding that the business judgment rule did not protect the actions of the company’s directors); Veasey,

well-advised Delaware corporations—and thus virtually all public companies incorporated in Delaware—have such provisions in their certificates.²⁴⁸

Thus, although eliminating the scienter requirement of oversight liability is a necessary condition for meaningfully expanding oversight liability, it is not a sufficient one. Under current law, most complaints sounding in oversight are dismissible because, as in *Citigroup*, plaintiffs cannot adequately plead scienter. If *Stone* were overruled and scienter were no longer required, most oversight complaints would still be dismissible because, unless the plaintiffs could adequately plead scienter (and thus a violation of the duty of loyalty), the complaint would sound only in the duty of care. Thus, the director defendants would be entitled to a dismissal of the complaint if they invoke the corporation's section 102(b)(7) provision. Hence, the Delaware courts could create whatever duties to oversee risk-management practices that critics of *Citigroup* might like, but if director defendants violate them—whether negligently or even grossly negligently—the directors would be exculpated from liability under the corporation's section 102(b)(7) provision. Because of such provisions, the effective standard of liability would still be conscious disregard of duty, and plaintiffs would still virtually always lose.

Assume, as is even more unlikely, that after the Delaware Supreme Court overrules *Stone* to eliminate scienter as an element of oversight liability, the Delaware General Assembly amends or repeals section 102(b)(7) in order to prevent boards from being exculpated against failures to oversee the risk-management practices of the company. Since the General Assembly enacted section 102(b)(7) precisely to quell the perceived crisis generated by *Van Gorkom*,²⁴⁹ its amendment or repeal could generate a new crisis; nevertheless, director oversight suits based solely on alleged breaches of the duty of care could then, at least, get to

Finkelstein & Bigler, *supra* note 163 (discussing the Delaware legislature's passage of section 102(b)(7) in response to the decision of the Delaware Supreme Court in *Van Gorkom*). All of the states (but not the District of Columbia) have enacted provisions more or less analogous to section 102(b)(7). See MARK A. SARGENT & DENNIS R. HONABACH, D & O LIABILITY HANDBOOK (2009) (citing and discussing the laws of various states).

248. See Miller, *supra* note 22, at 1162.

249. See Veasey, Finkelstein & Bigler, *supra* note 163, at 400–04 (discussing the enactment of section 102(b)(7) exculpation). *But cf.* Roberta Romano, *What Went Wrong with Directors' and Officers' Liability Insurance?*, 14 DEL. J. CORP. L. 1, 23 (1989) (noting that, although there was a disruption in the directors' and officers' liability insurance market in 1984–1985, *Van Gorkom* could not have caused the crisis because the case “was decided in 1985 after the insurance crisis was well under way”).

trial.²⁵⁰ The essence of such claims would be that the directors had not done what they should have done to monitor the risks that the firm was running. Since, by hypothesis, this standard of conduct includes no scienter requirement, it must perforce be some merely objective standard, such as negligence or gross negligence. Whatever it may be, however, when the court attempts to implement the standard, it will be inquiring into the substantive merits of the company's risk-management practices.

As mentioned above, even with the corporate information and reporting systems at issue in conventional *Caremark* claims, if the court reviewed the adequacy of such systems on the merits, it would be involved in passing judgment on the merits of business judgments, and thus second-guessing the board's business decisions.²⁵¹ As Chancellor Allen stated in *Caremark*, "the level of detail that is appropriate for such an information system is a question of business judgment"²⁵² because it involves trading off the costs and benefits of alternative systems. Nevertheless, perhaps a limited exception in the oversight context to the principle that courts will not review business decisions on the merits could be defended by analogy with the board's procedural duty of care under *Van Gorkom* to be informed of the material facts reasonably available before actively making a business decision.²⁵³ Obviously, the question of which information and reporting systems are cost justified is vastly more complex than the question of which information is cost justified for the board to gather before making a particular, discrete business decision, but at least both questions concern which information the board should have before it.

If, however, the court were to start reviewing the company's risk-management systems on the merits, even if the review were made under a gross negligence standard, the level of interference by the court with the business decisions of the board (and, for that matter, the decisions of more junior employees as well) would be immense. There are two main reasons for this. First, imagine that the court concentrates on the kinds of risk-management decisions most characteristically made by directors in contradistinction from those made by more junior managers—for instance,

250. Miller, *supra* note 22, at 1161–62.

251. In another article, I argue that "whether an information and reporting system is reasonable is itself a substantive business decision." Miller, *supra* note 26, at 940. This is because, "assuming that a *reasonable* system is one that maximizes shareholder value in the long run, a system will be reasonable if the benefits of the system, in the form of improved decision-making by the board, exceed the costs of designing and implementing it." *Id.* Further, "Measuring these costs and benefits and balancing the one against the other is exactly the kind of highly uncertain judgment at issue in business decisions." *Id.*

252. *In re Caremark Int'l Inc. Derivative Litig.*, 698 A.2d 959, 970 (Del. Ch. 1996).

253. *Smith v. Van Gorkom*, 488 A.2d 858, 872–73 (Del. 1985).

the decision concerning the firm's risk tolerance as opposed to the decision of which financial models to adopt to measure risk. That is, if the board were willing to tolerate a \$10 million VAR with 95 percent confidence, perhaps the court would find that this was grossly negligent and the board should have tolerated only a \$5 million VAR with 99 percent confidence. If it did this, the court would be saying, in effect, that the board was wrong to pursue a higher-return, higher-risk business strategy and should have preferred a lower-return, lower-risk strategy. This, of course, would instantly go to the heart of the board's business judgment,²⁵⁴ for if decisions about the risk-return profile of the firm are not paradigm business judgments, nothing is. Another way to see this is to notice that the overall risk tolerance of the firm is not some discrete decision, but affects virtually every investment and trade that the firm makes. For, as explained in Part II.B, contemporary risk measures aggregate the risk across all the business units of the firm. Thus, criticizing the risk tolerance set by the board is not to criticize any one investment or trade the firm made, but the portfolio of investments and positions resulting from its whole investment and trading strategy. A board ordered to comply with the court's understanding of the correct risk-return profile for the corporation would have an immense job reconfiguring the firm's portfolio—selling some assets and buying others—to achieve a new portfolio at the level of risk that the court had decreed was proper.

Second, virtually every aspect of the firm's business creates risks of some kind or other, and so risk managers, in effect, review all aspects of the firm's business for risk. More than that, they are generally involved in every transaction the firm undertakes. For example, Citigroup notes in its annual report that, "for corporate clients and investment banking activities across the organization," all transactions by the firm must be approved not only by a responsible officer from the relevant business unit but also by a credit officer in the bank's credit risk-management division.²⁵⁵ Therefore, if courts undertook a meaningful review of risk-management practices, they could review every trading decision involving counterparty credit risk the bank undertakes, determining whether the risk managers' assessment of the credit risk to the bank was, say, grossly negligent on the merits. Furthermore, risk managers are even involved in such back-room activities

254. It might also make virtually no sense from a financial point of view. *See infra* Part IV.B.

255. 2007 CITIGROUP REPORT, *supra* note 174, at 54 ("For corporate clients and investment banking activities across the organization, the credit process is grounded in a series of fundamental policies, including . . . [that] a minimum of two authorized-credit-officer signatures are required on extensions of credit (one from a sponsoring credit officer in the business and one from a credit officer in credit risk management).").

as trade processing and settlement, accounting, record keeping, intrafirm communications, computing, and information technology. Risks arising from failures of such support and processing systems are usually called “operational risk,”²⁵⁶ and managing such operational risks is also part of risk management more generally.²⁵⁷ This means that, if a court were to review a firm’s risk-management systems for gross negligence, it would in effect be reviewing virtually every business decision the firm ever makes for gross negligence.

Put another way, shareholders only sue when actions or omissions by the company work out poorly—that is, cause losses. Leaving aside those rare cases in which officers or employees intend to harm the corporation, actions or omissions that work out poorly were actually intended by the corporate agents undertaking them to benefit the company. If they in fact worked out poorly, then *ex ante* there must have been a risk, whether or not recognized at the time, that the action or omission would cause losses to the company. Hence, whenever any action or omission leads to losses for the company, the situation can be characterized as a failure of risk management: plaintiffs can say either that the company wrongfully failed to appreciate the risk that it was running or that, appreciating the risk, the company should not have run it or did not properly insulate itself against it. Thus, if courts review risk-management policies on the merits, they will be empowered to review every business decision the company makes. This would effectively repeal the business judgment rule.

Perhaps because of a failure to appreciate fully what actual risk management entails, some of the critics of *Citigroup* have hoped to find a

256. *E.g.*, GLEASON, *supra* note 52, at 15 (describing operational risk as arising from the “high levels of technical and operational support to deliver information and analysis to [a financial firm’s] traders, brokers, and lenders for structuring and pricing deals . . . including computing and [other] technology”); MARRISON, *supra* note 67, at 360 (defining “operating risk” to include “losses from the back-office operations of processing trades and information,” including from “[i]ncorrectly entering trades,” “[l]osing information on trades,” “failure of computer systems,” and “failures by vendors to correctly perform outsourced processing functions”); SAUNDERS & CORNETT, *supra* note 11, at 168 (defining “operational risk” as the “risk that existing technology, auditing, monitoring, and other support systems may malfunction or break down”). *See also* 2007 CITIGROUP REPORT, *supra* note 174, at 38 (describing operational risks to which the company is exposed, including “the risk of fraud by employees and outsiders, clerical and record-keeping errors, integration of numerous acquired businesses, and computer/telecommunications systems malfunctions” as well as “disruptions of its operating systems arising from events that are wholly or partially beyond its control (for example, natural disasters, acts of terrorism, epidemics, computer viruses, and electrical/telecommunications outages)”).

257. GLEASON, *supra* note 52, at 15–17; MARRISON, *supra* note 67, at 361–71 (discussing operational risk-management techniques); 2007 CITIGROUP REPORT, *supra* note 174, at 64 (describing various aspects of Citigroup’s operational risk-management process).

way to preserve the business judgment rule while expanding oversight liability to allow courts to pass on the substantive merits of a company's risk-management practices. Thus, Brown seemingly approves of the business judgment rule, saying that, if "the board considered the relative risks and made a decision to go forward, the decision almost certainly would have fallen under the business judgment rule and almost certainly would have been insulated from liability,"²⁵⁸ but he nevertheless criticizes Chancellor Chandler's opinion in *Citigroup* because the Chancellor failed "to define the types of information that ought to be reported to the board in order for it to fulfill its oversight function."²⁵⁹ Likewise, Pan agrees that "whether the [Citigroup] board exercised good business judgment . . . is not for a court to decide."²⁶⁰ He does think, however, that the Delaware courts should "[s]trengthen[] the fiduciary duty to monitor"²⁶¹ to ensure that "boards . . . make the effort to collect the right type of information about the corporation."²⁶² With risk management, it cannot be both ways. Once the court begins evaluating the merits of the risk-management practices selected by the board, it would be evaluating the substantive aspects of virtually every business decision the company makes. The court would not only be exercising business judgment, but also exercising it on a scale wide enough to encompass the entire risk-return profile of the company, and deep enough to reach virtually every investment or trade the company makes. If there are cogent reasons for having a business judgment rule, those reasons should foreclose any substantive review by courts of a firm's risk-management practices.

B. THE VARIOUS MEANINGS OF "EXCESSIVE RISK"

Part IV.A argued that expanding oversight liability would effectively repeal the business judgment rule. This section argues that, even apart from the policy justifications underlying the business judgment rule, holding directors personally liable for supposed failures in risk management would not help address the kind of excessive risk taking that may have contributed

258. J. Robert Brown, *Delaware's Top Five Worst Shareholder Decisions for 2009 (#2)*: In re Citigroup, THE RACETOTHEBOTTOM.ORG (Jan. 6, 2010, 6:00 AM), <http://www.theracetothetbottom.org/preemption-of-delaware-law/delawares-top-five-worst-shareholder-decisions-for-2009-2-in.html> [hereinafter Brown, *Shareholder Decisions*]. But perhaps Brown, whose evaluation of Delaware's business judgment jurisprudence is not entirely favorable (he opines that "Delaware cannot be trusted to define fiduciary obligations for directors"), may question even this. See Brown, *supra* note 35.

259. Brown, *Shareholder Decisions*, *supra* note 258.

260. Pan, *supra* note 33, at 739.

261. *Id.* at 740.

262. *Id.*

to the financial crisis. The reason is that risk can be said to be “excessive” in several quite different senses, and when these senses are distinguished, it turns out that the sense in which excessive risk taking may have been a cause of the financial crisis and the sense in which expanding directorial oversight might conceivably reduce excessive risk taking are not the same. Thus, even if director oversight liability were expanded, as some people think desirable, the result would not reduce the probability of a future financial crisis like the one of 2007–2008.

Accordingly, this section distinguishes four ways that the risks a financial firm is running can be excessive, including (1) risks in excess of return-adjusted systematic risk, (2) risks in excess of the risk tolerance of the firm as set by the board of directors, (3) risks in excess of the risk tolerances of shareholders, and (4) excessive systemic risk, that is, risk in excess of what can be socially justified. In fact, only the last of these was arguably a cause of the financial crisis, but only the second is even conceivably a proper predicate for oversight liability.

1. Risks in Excess of Return-Adjusted Systematic Risk

As discussed in Part II.A, it is elementary in financial theory that, because investors are generally risk averse, attaining more than the risk-free rate of return entails bearing some risk. In particular, to the extent that the market is efficient, the rate of return on an investment will reflect the systematic risk associated with that investment, since the unsystematic risk is capable of being diversified away.²⁶³ The market will thus assign to investments with the same level of systematic risk the same market risk premium in discounting their expected cash flows; such investments will thus have the same expected rate of return. Hence, if a firm made an investment with an expected rate of return below that which the market demands for the systematic riskiness of the investment, the firm would not be earning enough return for the risk it was bearing. Equivalently, the firm would be bearing more risk than needed in order to earn the return it expects. Thus, this is one way a firm might be bearing excessive risk: the riskiness of its investments exceeds the risk level that the market will tolerate for the expected returns on the investments. In other words, the firm is bearing risk in excess of the return-adjusted systematic risk of its investments.

263. See, e.g., DAMODARAN, *supra* note 39, at 29–35 (discussing separation of systematic and nonsystematic risk); DAMODARAN, *supra* note 36, at 60–77 (same); ROSS, WESTERFIELD & JAFFE, *supra* note 44, at 243–314 (same).

Calling such investments excessively risky, however, is an odd way to describe the situation. Rather than saying the risk on such investments is too high, it would be clearer and more natural to say that returns on such investments are too low. They are simply bad investments. Detecting and preventing such investments is not the proper job of risk managers. It is, rather, the job of the traders in the firm. It is a trader's job to make profitable investments, and investments such as these are not profitable, even if the expected return is achieved. For, given the risk assumed, a greater return would have been available in the market, and the spread between the available return and the return of the investment actually made would be an economic loss to the firm. Avoiding such losing investments is thus the primary responsibility of the firm's traders, not its risk managers. Risk managers can help in this regard, of course, because these investments, by having risk levels above return-adjusted systematic risk, necessarily increase the firm's overall risk level relative to its return. They are thus prime candidates to be identified by risk managers as investments that ought to be shed, but this is because they are bad investments from the start, not because they are good investments that simply happen to be too risky (in some other sense) for the firm.

Furthermore, expanding oversight liability to try to reduce such excessive risks as these is obviously misguided. One may as well say—for it amounts to exactly the same thing—that we should expand oversight liability so that directors are encouraged to monitor junior employees so that such employees may make better business decisions. This is just to convert oversight liability into a general purpose way for courts to review business decisions at all levels of the business. Risks excessive in this sense, therefore, are not the proper concern of oversight liability.

2. Risks in Excess of the Risk Tolerance of the Firm as Set by the Board

Imagine that the board fixes a certain risk tolerance for the firm, for example, by setting a \$10 million daily VAR with 95 percent confidence. Imagine further that the firm's risk managers then miscalculate the firm's VAR, either because they are using flawed financial models or flawed data, because they are misapplying the models (for instance, because the models are built on statistical assumptions about the data that are known to be false), or for some other identifiable reason that competent risk managers would recognize as involving a mistake. If, as a result of such mistakes, the firm's VAR (as properly calculated) frequently exceeds the level set by the board, then the firm would be running excessive risks—that is, risks exceeding those that the board had determined the firm should run. Clearly,

analogous points could be made about risk tolerances used as targets in other financial models as well—for example, about tolerable levels of credit risk, interest rate risk, and so on. If a firm was running risks excessive in this sense, then clearly this would be a failure of the firm's risk-management systems.

Should the board have oversight liability for such failures? Surely if the directors knew that the firm's risk-management systems were failing in this way and did nothing about it, then they should be liable in oversight. This merely applies the *Caremark* standard, including its scienter requirement, and the efficiency rationale underlying *Caremark* would certainly apply here. But what if the board did not know the firm's risk-management systems were failing but nevertheless was negligent or grossly negligent in not detecting the failures? For one thing, as discussed in Part II.B, because of the difficulty and complexity of risk-management questions, it is very unlikely that even a director who is an expert in risk management could cost-effectively detect serious errors in the company's risk-management practices. Given the large number of models major financial firms use and the complexity of these models, a director would likely never have sufficient knowledge to discover all but the most blatant and egregious errors—precisely the kinds of errors that are least likely to occur in practice. Hence, as a factual matter, there would be very few cases in which, even under a negligence or gross negligence standard, directors would be liable in oversight for failing to detect errors in the firm's risk-management practices. That point aside, however, as discussed in Part IV.A, expanding oversight liability by eliminating the scienter requirement would effectively repeal the business judgment rule. Assuming that the usual efficiency rationales for the business judgment rule are compelling, expanding oversight liability for directors who fail, even negligently or grossly negligently, to detect failures in the company's risk-management practices that result in the company running risks in excess of the risk tolerances set by the board would be inefficient.

More importantly, however, such an expansion of liability would likely do nothing to prevent another financial crisis, for the large losses suffered in the crisis by major financial institutions almost certainly did not arise from risk-management failures of the kind contemplated here. That is, it is not that risk managers at banks were miscalculating the risks that the banks were running, with the result that the banks were taking on more risk than their directors thought.²⁶⁴ There was no such allegation, for example,

264. According to the 2001 CGFS Survey, the stress test most often performed by the surveyed

in the *Citigroup* case. Rather, in the most likely interpretation, in the years leading up to the financial crisis, major financial firms were knowingly, and in exchange for an attractive expected return, taking on very small risks of very large losses,²⁶⁵ and, unfortunately, these risks materialized. As noted above, even after the financial crisis, major banks knowingly take on nontrivial risks of bankruptcy.²⁶⁶ Discussing the bubble in the U.S. residential real estate market leading up to the financial crisis, Judge Posner, although intentionally simplifying the matter, explains the basic idea very clearly:

In gauging the risk of calamity, the key probabilities [the banks] had to consider were that the rise in housing prices was a bubble and that if it burst house prices would fall by at least 20 percent. If both events came to pass, insolvency would loom. Suppose the best guess was that there was a 10 percent probability that the price rise was a bubble and the same probability that if it was a bubble house prices would fall by at least 20 percent. Then the probability that house prices would fall by at least 20 percent was only 1 percent (0.1×0.1), and so disaster would be unlikely to occur for many years, and so the risk of disaster would have seemed worth running.²⁶⁷

It is no failure of risk management if the firm takes on a one-year VAR sufficient to bankrupt the company with a 99.5 percent confidence (that is, a one in two hundred annual chance of bankruptcy), and then almost goes bankrupt. The model predicts that such losses will occur occasionally—indeed, at least once in two hundred years. Without further evidence, there is no reason to think anything was wrong with the model or its application if those losses strike sooner rather than later.²⁶⁸

banks involved equity market crashes. 2001 CGFS SURVEY, *supra* note 74, at 12. One reason for this was that, at the time the survey was performed, such markets were at or near historical highs. *Id.* at 5. Similarly, in the 2005 CGFS Survey, when real estate prices were historically high in what we now know were the early stages of the real estate bubble, the surveyed banks were stress testing for a real estate market crash. 2005 CGFS SURVEY, *supra* note 74, at 10 (“Management requests for one-off tests related mainly to real estate.”). It is reasonable to assume, therefore, that at least some of the major financial firms were, as Judge Posner suggests, aware of the risks they were running. *See infra* note 267 and accompanying text.

265. *See* POSNER, *supra* note 1, at 208–09 (describing how Robert Rubin, a director at Citigroup, had successfully convinced the board to increase the risk tolerance of the bank).

266. *See supra* note 210 and accompanying text.

267. POSNER, *supra* note 1, at 78–79.

268. *See* Miller, *supra* note 173, at 134 (“But didn’t the banks’ leveraging up increase the risk of bankruptcy? Of course it did, but all businesses take on a small but positive risk of bankruptcy. . . . [Bankers] were taking perfectly rational risks that—because they were in fact in a bubble—came to pass contrary to people’s reasonable expectations and had disastrous effects.”).

3. Risks in Excess of the Risk Tolerances of Shareholders

Shareholders might say the firm was running excessive risks in the sense that the firm was making more risky investments than the shareholders may have preferred—that is, taking risks in excess of the shareholders' risk tolerances. Such shareholders might say that they would have preferred lower risks, albeit with lower expected returns. It should be clear, however, that directors should have no liability, whether in oversight or otherwise, for excessive risk taking in this sense.

Shareholders who prefer less risky investments than the board is causing the corporation to make have a simple remedy: they can sell their shares and invest elsewhere. More generally, shareholders rightly care about the level of risk that they themselves are bearing—that is, the level of risk in their own portfolios, not the risk being borne by corporations in which they are invested, which is a quite different thing. Shareholders can adjust the risk level in their own portfolios at will by buying or selling treasury securities—they can decrease the risk (and the expected return) of their portfolio by selling riskier assets and buying more treasuries, and they can increase the risk (and the expected return) of their portfolio by selling treasuries and buying riskier assets.²⁶⁹ In this way, investors can get exactly the risk-return combinations they want. Moreover, because of the low transaction costs of buying and selling financial assets, achieving the desired combination is quite easy. Value-maximizing shareholders should want corporations to take on the highest expected-return investments available. If all companies do that, investors will get the highest risk-adjusted return possible on their portfolios. They certainly ought to have no cause of action against a board of directors that makes that happy result possible for them.

4. Socially Inefficient Risk Taking and Excessive Systemic Risk

Finally, there is one more way that the risks that financial firms bore in the years leading up to the financial crisis may be said to be excessive: the risks may have been excessive not from the point of view of the banks or their shareholders, but from the point of view of society. For example, assume that during the period leading up to the financial crisis, Citigroup was leveraging up, increasing the risk it was bearing, but also increasing its expected return.²⁷⁰ Assume too that Citigroup correctly computed the risks and the expected returns on the investments that it was making, so that the

269. See ELTON ET AL., *supra* note 133.

270. This was actually the case. See POSNER, *supra* note 1, at 208–09.

investments were rational to make at the time—that is, had positive net present value for Citigroup and had an aggregate risk level that, while higher than it was in the past, was within the board’s revised risk tolerance. In no sense, therefore, were the risks that Citigroup was bearing excessive from the point of view of the company or its shareholders.

But now assume that Citigroup’s investment strategy works out badly, and Citigroup suffers large losses. Assume further that the losses are so large that Citigroup becomes insolvent or even bankrupt. If that happens, then other financial firms to which Citigroup owes money will also suffer losses, and financial markets in which Citigroup participated may be disrupted (for instance, if Citigroup starts selling large amounts of securities in order to attempt to pay its debts).²⁷¹ Conceivably, other financial firms could then fail, thus starting a chain reaction in which one firm after another fails, and various markets cease to function. Such a scenario, however, would be a disaster not just for the shareholders and creditors of the firms involved but for virtually everyone in society. A large number of simultaneous bank failures sharply reduces the money supply and the supply of credit in the economy,²⁷² and has a significant adverse effect on consumer and business confidence, including the public’s confidence in other banks that may otherwise be perfectly sound. Thus, a cascade of bank failures imposes costs largely on third parties who are not connected with the banks that ran the underlying risks and that would enjoy virtually all the benefits if the underlying investments had worked out as expected. In other words, many of the costs that result if the banks’ investments work out poorly would be externalized,²⁷³ and the banks

271. If, as is likely, the investment strategy that Citigroup follows is also followed by similar financial institutions, then many events that would adversely affect Citigroup’s financial position would also adversely affect the financial position of these other institutions. In other words, the risks of their failure will become positively correlated. This effect is generally known as convergence. *See generally* Timothy F. Geithner, President and Chief Exec. Officer, Fed. Reserve Bank of N.Y., Remarks Before the Conference on Systemic Financial Crises at the Federal Reserve Bank of Chicago: Changes in the Structure of the U.S. Financial System and Implications for Systemic Risk (Oct. 1, 2004), *available at* http://www.ny.frb.org/newsevents/speeches_archive/2004/gei041001.html (discussing convergence of the kinds of financial transactions undertaken by banks and nonbank financial firms). *See also* POSNER, *supra* note 1, at 106–15.

272. *See* MILTON FRIEDMAN & ANNA JACOBSON SCHWARTZ, *A MONETARY HISTORY OF THE UNITED STATES: 1867–1960*, at 351–52 (1971) (arguing that bank failures during the Great Depression sharply reduced the money supply).

273. There is an important subtlety here. What matters is not that the financial firms are externalizing some of their costs, but that the externalized cost (that is, the cost of the increased cost of capital, multiplied by the probability that such increase will occur) is greater than the cost would be if internalized (that is, the cost the financial firm would bear, in the form of reduced profits, if it pursued a lower-risk, lower-return business strategy). This is Ronald Coase’s point that what matters is not who “causes” a loss but who is the more efficient bearer of the cost. *See* R.H. Coase, *The Problem of Social*

naturally would not take account of such costs in computing the risk and return on their investments. Once these costs are taken into account, it is possible that investment decisions that had positive net present value for the banks had negative net present value for society as a whole. As Judge Posner states,

The essential point is the difference between a 1 percent probability that a firm will go broke [which may be entirely justified by the expected return the firm hopes to obtain], because of risky lending, and a 1 percent probability of a depression because the leading financial firms have a correlated 1 percent risk of going broke. The toleration of the risk is rational for each firm, irrational for society.²⁷⁴

This problem is generally discussed under the rubric of *systemic risk*.²⁷⁵ Although there is significant confusion about the meaning of this term,²⁷⁶ the key idea is that, because large financial firms are generally intertwined through a web of financial relationships and through their participation in capital markets, if the activities of one firm cause it to fail or to become insolvent, then other firms may fail or become insolvent and certain capital markets may become illiquid or otherwise cease to function. This can start a chain reaction that creates costs for financial firms and other persons in society,²⁷⁷ most importantly in the form of an increased

Cost, 3 J.L. & ECON. 1, 2, 15–16 (1960); DAVID D. FRIEDMAN, LAW'S ORDER 36–46 (2000); RICHARD A. POSNER, ECONOMIC ANALYSIS OF LAW 50–55 (7th ed. 2007). In this case, the costs of a chain reaction of financial firm failures is generally immense, though the probability of such a scenario coming to pass is low and difficult to estimate (obviously, the probability is less than the probability that the firm itself fails). This makes comparing the externalized cost, when borne by society, and the internalized cost, when borne by the firm itself, difficult to compare.

274. POSNER, *supra* note 1, at 112.

275. Systemic risk is different from systematic risk. The latter is a component of the risk of investments (the nondiversifiable component, sometimes called market risk) and is borne by the person making the investment. See DAMODARAN, *supra* note 39, at 29–35 (discussing systematic and nonsystematic risk); DAMODARAN, *supra* note 36, at 60–77 (discussing risk analysis in finance and differentiating between diversified and nondiversified risk); ROSS, WESTERFIELD & JAFFE, *supra* note 44, at 243–314 (discussing risk measurement and the role of diversification in reducing risk).

276. E.g., Steven L. Schwarcz, *Systemic Risk*, 97 GEO. L.J. 193, 196 (2008) (“There is . . . a great deal of confusion about what types of risk are truly ‘systemic’ . . . and what types of systemic risk should be regulated.”). Greenspan has observed, “It is generally agreed that systemic risk represents a propensity for some sort of financial system disruption,” but “one observer might use the term ‘market failure’ to describe what another would deem to have been a market outcome that was natural and healthy, even if harsh.” *Id.* (alteration in original omitted) (quoting Alan Greenspan, Remarks at a Conference on Risk Measurement and Systemic Risk, Board of Governors of the Federal Reserve System (Nov. 16, 1995)). See also *Systemic Risk: Examining Regulators’ Ability to Respond to Threats to the Financial System: Hearing Before the H. Comm. on Fin. Servs.*, 110th Cong. (2007).

277. See Schwarcz, *supra* note 276, at 202–04 (arguing that a proper understanding of systemic risk should consider not only the relationships among financial firms, but also the effects of such firms on capital markets).

cost of capital.²⁷⁸ In Steven Schwarcz's definition, "systemic risk" is "the risk that (i) an economic shock . . . triggers . . . either (X) the failure of a chain of markets or institutions or (Y) a chain of significant losses to financial institutions, (ii) resulting in increases in the cost of capital or decreases in its availability."²⁷⁹ We can ask, therefore, whether the activities of financial firms leading up to the financial crisis in 2007–2008 were generating excessive systemic risk.

At least in the context of this Article, however, this is a somewhat misleading way of putting the question. The phenomenon discussed under the rubric of systemic risk concerns activities with costs that are partly externalized. These externalized costs are greater than they would be if they were internalized. The result is that the firms undertaking these activities engage in more than the socially optimal level of such activities. It is not that such activities are too risky in the financial sense of that word; it is that, from the point of view of society, the activities have negative present value. They are simply inefficient. The analogy is to a factory that externalizes some of its costs onto its downstream neighbors by venting pollutants into a waterway that have a certain probability of causing cancer, when it would be cheaper (for society, though not for the factory owner) to pay to haul the pollutants away.²⁸⁰ These activities could be described as creating excessive environmental risk. It would be more natural, however, to say—as we generally do—that such activities are just inefficient.

Now, if this scenario occurred in the years leading up to the financial crisis, it would be a genuine problem. It is, moreover, an example of a well-understood type of collective action problem—a tragedy of the commons.²⁸¹ Specifically, there is an activity that individual actors find to be in their private interest and thus want to pursue, but because a significant fraction of the costs of this activity are being externalized, society as a whole is made worse off by the activity.²⁸² In such cases, it can be entirely proper—indeed economically efficient—for government to

278. *Id.* at 198–99 (“Increases in the cost of capital, or decreases in its availability, are the most serious direct consequences of a systemic failure.”). See also *Hedge Fund Operations: Hearing Before the H. Comm. on Banking and Fin. Servs.*, 105th Cong. (1998) (statement of William J. McDonough, President of the Federal Reserve Bank of New York) (noting that increases in the cost of capital were the most important adverse consequences of the failure of Long Term Capital Management).

279. Schwarcz, *supra* note 276, at 204.

280. Again, the qualification in the text about it being cheaper to haul away the pollution than to vent it into the waterway is important.

281. Garrett Hardin, *The Tragedy of the Commons*, 162 *SCIENCE* 1243, 1243–46 (1968).

282. See, e.g., FRIEDMAN, *supra* note 273, at 28–46; Schwarcz, *supra* note 276, at 198, 206 (describing the problem of systemic risk in terms of a tragedy of the commons).

regulate such activities, perhaps even to prohibit them.²⁸³ Indeed, this is one of the most important purposes of the capital regulations imposed on banks, such as the Basel Capital Accords,²⁸⁴ and the Federal Reserve's regulations implementing them.²⁸⁵ Such regulations are designed to ensure that banks maintain adequate capital to absorb even unexpected losses, and thus protect against the risk of bank insolvencies.²⁸⁶

Directors could be subject to oversight liability for failing to ensure that the corporation complies with legal capital regulations, just as they might be subject to liability for failing to ensure that the corporation complies with any other legal or regulatory requirement. This is simply traditional *Caremark* liability. Apart from questions of illegality, however, oversight liability for failures in risk management related to activities generating excessive systemic risk is entirely misguided. The investments complained of have net present value for the company and include risks within the risk tolerance set by the board. In making such investments or in allowing them to be made, the directors were acting in the best interests of the company (that is, they maximized value for shareholders), though not in the best interests of society as a whole. The shareholders thus have nothing to complain of, even if society generally does.

Indeed, allowing shareholders to bring oversight suits on the basis of the board not preventing the corporation from making investments with socially excessive risks would be perverse. If such suits were allowed, there would be two possibilities: either (1) the subject investments turn out well and make money for the company, in which case the shareholders enjoy the benefits and will not sue, or (2) the investments turn out poorly and the company suffers losses (and perhaps society does as well as if systemic risks materialize), in which case the shareholders will sue and recover such

283. But this is not always the case. All solutions, including government intervention, have costs, and the mere fact that an unregulated market creates costs does not imply that these costs are greater than the costs associated with regulating the market to prevent them. The fact that a market is imperfect does not imply a need for regulation; only when the benefits of regulating exceed its costs is regulation efficient. See Coase, *supra* note 273, at 17–19.

284. The Basel II accord is memorialized in BASEL COMM. ON BANKING SUPERVISION, INTERNATIONAL CONVERGENCE OF CAPITAL MEASUREMENT AND CAPITAL STANDARDS: A REVISED FRAMEWORK (2006), available at <http://www.bis.org/publ/bcbs128.pdf>, and BASEL COMM. ON BANKING SUPERVISION, THE NEW BASEL CAPITAL ACCORD (2001), available at <http://www.bis.org/publ/bcbsca03.pdf>.

285. See 12 C.F.R. pt. 225 app. A, G (2010) (containing the most important implementing regulations of the Federal Reserve); Joint Press Release, Fed. Reserve Bd. et al., Banking Agencies Announce Revised Plan for Implementation of Basel II Framework (Sept. 30, 2005), available at <http://www.federalreserve.gov/boarddocs/press/bcreg/2005/20050930/default.htm>.

286. See SAUNDERS & CORNETT, *supra* note 11, at 586–619.

losses from the directors. The losses sought to be recovered, however, will not be the losses to society generally arising from a materializing systemic risk, but the losses the firm itself bore. Furthermore, under such a rule, the shareholders always win, for either they profit ex post from an investment that was irrational for society ex ante, or else they are indemnified by the directors when the investments lose value. Value-maximizing shareholders would thus want their directors to make as many such investments as possible, for such investments have upside risk for the shareholders but no downside risk. The upshot would be that, when a financial corporation seeks to benefit itself and its shareholders at the expense of society generally, it would be allowed to do so, provided that the attempt is successful. But if it fails, the directors will compensate not the public that was harmed, but the shareholders who were not. Such a rule of law makes no sense at all.

In other words, the interest of society here is to prevent financial institutions from taking socially excessive but privately profitable risks. It makes sense to penalize directors for causing their firms to engage in the relevant transactions, but it makes no sense to appoint shareholders as the enforcers, much less the beneficiaries, of such penalties. Shareholders, as the primary beneficiaries of the activities to be suppressed, have quite the wrong incentives as enforcers.

The problem is a classic collective action problem, and the standard solution to such problems is the one society has already adopted: regulation by government authority. Perhaps the financial crisis has shown that the capital regulation of banks ought to be strengthened. The Dodd-Frank Act, for example, attempts to do exactly this.²⁸⁷ Furthermore, there is certainly a legitimate role for oversight liability when directors fail to properly monitor the business to make sure that it is complying with such regulation, but that is merely traditional *Caremark* liability for failures in legal compliance, not for failures in risk management per se. What is lacking is an intelligible place for oversight liability claims related to actions that,

287. In particular, the Act requires the Federal Reserve to establish enhanced risk-based capital, leverage, and liquidity requirements, and overall risk-management requirements (among other requirements) for certain systemically important financial firms, including bank holding companies with \$50 billion or more in assets. See Dodd-Frank Wall Street Reform and Consumer Protection Act of 2010, Pub. L. No. 111-203, 124 Stat. 1376 (to be codified in scattered sections of 12 and 15 U.S.C.). See also DAVIS POLK & WARDWELL LLP, *supra* note 11, at 7–11 (discussing new capital and other standards for systemically important firms); DELOITTE CTR. FOR FIN. SERVS., THE SOUND OF RUMBLING THUNDER: ASSESSING THE IMPACT OF U.S. FINANCIAL REGULATORY REFORM 3–4 (2010), available at http://www.deloitte.com/assets/Dcom-UnitedStates/Local%20Assets/Documents/FSI/us_fsi_financialreform_rumblingsound_july2110.pdf (same).

while perhaps describable as failures of risk management, do not harm shareholders.

V. CONCLUDING OBSERVATIONS

The purpose of this Article was to determine what role director oversight liability can efficiently play in improving risk-management practices at large financial institutions. Prior treatments of this topic have generally not been informed by a realistic understanding of the financial modeling that risk managers at major financial firms do and the kinds of business judgments that underlie their application. Lacking this realistic understanding, some prior analyses have treated risk management as if it were a discipline akin to, say, automobile maintenance—the kind of thing that professionals know how to do correctly and, if done correctly, will almost inevitably achieve the desired result. Thus, some people seem to assume that, if only financial firms had proper risk-management procedures, they would never incur large losses and the financial crisis of 2007–2008 could have been averted. This is a form of intellectual superstition. It attributes a power to foresee the future to a technical subdiscipline of applied financial theory.

Risk management has nothing like this magical power. The financial models used in risk management (like all financial models) attempt to predict the future on the general assumption that the future will resemble the past in relevant respects, which it often does not. The design and application of such models thus involves the exercise of significant business judgment, and failures of risk management should generally be viewed as failures of business judgment. That businesspeople, acting in good faith and on an informed basis, sometimes—indeed often—make business decisions that work out badly, even catastrophically, is perfectly familiar. Nothing changes if we call such people risk managers.

Once this is clearly understood, there can be little doubt that expanding oversight liability to improve risk-management practices at financial firms would be tremendously inefficient. The key limitation on oversight liability under current Delaware law is the scienter requirement. Expanding oversight liability in a meaningful way means eliminating that requirement, and eliminating that requirement means that courts would be reviewing the substance of risk-management decisions on the merits. Since risk-management decisions are business judgments, courts would be reviewing business judgments on the merits.

Worse, since virtually any business decision with a poor result can be

characterized as a risk-management failure, expanding oversight liability for risk-management failures would authorize the court to review on the merits virtually any business decision made at any level of the firm's business. If the business judgment rule is efficient, expanding oversight liability for risk-management failures would destroy this efficiency. That this has not been previously recognized is likely because the talismanic phrase "risk management" has obscured the nature of the decisions that courts would be reviewing. No one would suggest that better business decisions would be made if we expanded the board's oversight liability for the business decisions of junior employees. If we relabel certain business decisions as risk-management decisions, however, then suddenly the matter looks different. It looks different, however, only if we forget that risk management is just an aspect of all business decisionmaking—an aspect that has been christened with a technical-sounding name.

The primary motivation for such an expansion lay in the perception that excessive risk taking at major financial firms was a primary cause of the financial crisis. But, because of the difficulty of the subject and the great variation in circumstances from firm to firm, directors of such firms are much better placed than anyone else (generally including even regulators) to design and implement risk-management systems. This is why, subject to certain restrictions, their primary regulator, the Federal Reserve, allows them to do just that. Hence, the argument goes, if we want improved risk-management practices, it would make sense to expand oversight liability to give directors stronger incentives to monitor those practices. There are many problems with that argument, but a key one is that the term excessive risk is being used equivocally. There are several ways that risk can be excessive, and the way that excessive risk taking represents a genuine failure of risk-management systems (namely, risk taking in excess of the risk tolerance of the firm as set by the board) is not the way that excessive risk taking may have been a cause of the financial crisis (namely, activities that are inefficient because they externalize costs in the form of systemic risk). Hence, even if expanding oversight liability for risk-management failures did not otherwise involve the tremendous inefficiency of effectively repealing the business judgment rule, such an expansion would still not be well calculated to address the problem with risk taking that may have contributed to the financial crisis.

Finally, there is one more misunderstanding about risk that lurks in much of the popular and even legal discussion of risk taking in connection with the financial crisis. In general usage, though not in financial theory,

risk has strong negative connotations.²⁸⁸ Politicians, for example, will deride the proposals of their opponents as “risky schemes.”²⁸⁹ For lawyers, moreover, talk of risk, and especially talk of excessive risk, tends to conjure up the idea of unreasonable risks as that term is used in negligence doctrine in the law of torts. Indeed, Franklin Gevurtz draws this comparison explicitly when, in a discussion of oversight liability in connection with the *Citigroup* case, he says that “[i]mposing liability to pay the damages resulting from unreasonable risks is a conventional tool in the law to deter creating such risks” and notes that “this concept is a pillar of tort law.”²⁹⁰ That is correct, but business judgment law is very different from tort law,²⁹¹ and the notion of risk in torts is very different from the notion of risk in finance. From an economic point of view, tortious risks are said to be unreasonable when the actor who is creating the risk has a precaution available to him, the cost of which is less than the expected loss resulting from his other actions.²⁹² From a financial point of view, the actor is failing to make a socially desirable investment—an investment with positive net present value for society—and tort liability realigns the actor’s incentives so that the actor does what maximizes value for society. Tort liability is thus about preventing losses that are efficient to prevent.

Risk in the financial sense is not similar at all. In the financial sense, risk is not about whether an investment has positive net present value; it is about how the possible returns are distributed around the expected return. We avoid financial risks not because taking them would not have positive net present value. Indeed, riskier investments usually have greater net present value. Rather, we avoid risks and sacrifice value because we are

288. See, e.g., DAMODARAN, *supra* note 39, at 27 (“Risk, for most of us, refers to the likelihood that in life’s games of chance, we will receive outcomes that we will not like. . . . Thus, risk is perceived almost entirely in negative terms. . . . [But in finance the] definition of risk is both different and broader [and] includes not only the bad outcomes . . . but also good outcomes.”).

289. E.g., James Dao & Frank Bruni, *The 2000 Campaign: The Infighting; Jabs Fly in Presidential Ring over Social Security’s Future*, N.Y. TIMES, May 4, 2000, at A1 (noting how Al Gore referred to George W. Bush’s proposals related to Social Security as “risky”).

290. Franklin A. Gevurtz, *The Role of Corporate Law in Preventing a Financial Crisis: Reflections on In re Citigroup Inc. Shareholder Derivative Litigation*, 23 PAC. MCGEORGE GLOBAL BUS. & DEV. L.J. 113, 127 (2010).

291. Allen, Jacobs & Strine, *supra* note 162, at 453 (discussing differences between the common law of torts and Delaware business judgment law and stating that “in corporate cases, Delaware courts have chosen a definition of gross negligence that is even more difficult for a plaintiff to establish than the gross negligence standard normally applied in American tort or criminal cases”).

292. See *United States v. Carroll Towing Co.*, 159 F.2d 169, 173 (2d Cir. 1947) (articulating the classic “Hand formula,” under which a person is negligent if the cost to the person of avoiding a potential harm is less than the product of the probability of the harm occurring and the magnitude of the harm); POSNER, *supra* note 273, 167–71.

risk averse and are willing to trade some value for reduced variability in returns. Risk in the financial sense is thus a different concept used for different purposes. Risk management, unlike tort liability, is not about preventing losses. It is about smoothing returns, and the question of how smooth returns should be is a business decision. Risk management is the responsibility of boards, not courts.

