



American Finance Association

Tax Incentives to Hedge

Author(s): John R. Graham and Clifford W. Smith, Jr.

Source: *The Journal of Finance*, Vol. 54, No. 6 (Dec., 1999), pp. 2241-2262

Published by: Blackwell Publishing for the American Finance Association

Stable URL: <http://www.jstor.org/stable/797993>

Accessed: 24/11/2009 12:37

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=black>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Blackwell Publishing and American Finance Association are collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Finance*.

<http://www.jstor.org>

Tax Incentives to Hedge

JOHN R. GRAHAM and CLIFFORD W. SMITH, JR.*

ABSTRACT

For corporations facing tax-function convexity, hedging lowers expected tax liabilities, thereby providing an incentive to hedge. We use simulation methods to investigate convexity induced by tax-code provisions. On average, the tax function is convex (although in approximately 25 percent of cases it is concave). Carrybacks and carryforwards increase the range of income with incentives to hedge; other tax-code provisions have minor impacts. Among firms facing convex tax functions, average tax savings from a five percent reduction in the volatility of taxable income are about 5.4 percent of expected tax liabilities; in extreme cases, these savings exceed 40 percent.

IF A FIRM FACES A CONVEX TAX FUNCTION, then hedging that reduces the volatility of taxable income reduces the firm's expected tax liability.¹ For a firm facing some form of tax progressivity, when taxable income is low, its effective marginal tax rate will be low; but when income is high, its tax rate will be high. If such a firm hedges, the tax increase in circumstances where income would have been low is smaller than the tax reduction in circumstances where income would have been high, thus lowering expected taxes.

In examining the importance of various motives to hedge, researchers have employed two general methods: survey and regression analysis. There are potential problems with each. For example, some managers might view the information requested in a survey as proprietary or might hesitate to admit that they employ derivatives to speculate rather than hedge. To employ regression methods, researchers must generate proxies for tax status in addition to other firm characteristics. To the extent that a tax variable (e.g., whether the firm has a net operating loss carryforward) proxies for factors such as financial distress as well as tax circumstances, analysts face an identification problem in interpreting the results.

* Graham is at the Fuqua School of Business, Duke University, and Smith is at the Wm. E. Simon Graduate School of Business Administration, University of Rochester. We thank Steve Huddart, Robert McDonald, Cathy Schrand, David Shimko, René Stulz, two anonymous referees, and seminar participants at the University of North Carolina, Financial Management Association, Southern Finance Association, and the American Finance Association meetings for valuable comments and discussions.

¹ Financial economists have identified other nonmutually exclusive incentives to hedge in addition to tax-related incentives. See Mayers and Smith (1982, 1987), Main (1983), Stulz (1984), Smith and Stulz (1985), and Froot, Scharfstein, and Stein (1993). Smith (1995) provides a review of this literature.

Rather than employing survey or regression analysis, we use simulation methods. Our approach allows us to map a firm's effective tax function to determine the extent of its convexity. Thus, we examine a necessary condition for the firm to have tax-based incentives to hedge. We simulate a given reduction in the volatility of taxable income to examine the extent to which such hedging changes a firm's expected tax liability.

Though little used in corporate finance, the simulation methods we employ offer a potentially powerful tool for addressing this question. The basic structure of the tax code is well defined, yet aspects of the code—especially carryback and carryforward provisions—induce interactions that are difficult to analyze with analytical methods. Our simulations allow us to examine this institutional richness and investigate the extent to which the structure of the tax code induces material convexities in the effective corporate tax function. We analyze the distribution of convexities, thereby allowing a more reasoned assessment of the potential importance of this motive for hedging.

We frame our discussion in terms of the reduction in expected tax liabilities associated with a given reduction in the volatility of taxable income due to increased hedging. But, equivalently, we could discuss the expected increase in taxes from a given increase in volatility due to reduced hedging. (In untabulated analysis we examine this latter case and find no qualitative differences from the results presented below.)

From our analysis of more than 80,000 COMPUSTAT firm-year observations, we find that in approximately 50 percent of the cases, corporations face convex effective tax functions and thus have tax-based incentives to hedge. In approximately 25 percent of the cases, firms face linear tax functions. The remaining firms face concave effective tax functions (which provide a tax-based disincentive to hedge). Of the cases with convex tax functions, roughly one-quarter of the firms have potential tax savings from hedging that appear material—in extreme cases exceeding 40 percent of the expected tax liability. For the remaining firms, the tax savings are fairly small. Thus, our analysis suggests that the distribution of potential tax savings from hedging is quite skewed.

Firms are most likely to face convex tax functions when (1) their expected taxable incomes are near the kink in the statutory tax schedule (i.e., taxable income near zero), (2) their incomes are volatile, and (3) their incomes exhibit negative serial correlation (hence the firm is more likely to shift between profits and losses).

Our methods also allow us to decompose the basic structure of the tax code to examine the incremental impact of statutory progressivity, net operating loss (NOL) carrybacks and carryforwards, investment tax credits (ITCs), the alternative minimum tax (AMT), and uncertainty in taxable income. We find that much of the convexity is induced by the asymmetric treatment of profits and losses in the tax code. Carryback and carryforward provisions effectively allow firms to smooth their losses, thereby reducing tax-function curvature at its most convex points and making the function convex over a

broader range of taxable income. In contrast, the alternative minimum tax and investment tax credits have only a modest effect on the convexity of the tax function.

In Section I, we review the theory of tax-based hedging incentives and model the basic structure of the tax code to highlight the impact of uncertainty on expected tax liabilities. This allows us to identify specific provisions of the tax code that can lead to convex tax functions. We report the results of our simulation of the change in expected tax liabilities from hedging taxable income for a large cross section of COMPUSTAT firms in Section II. We relate the tax savings from hedging to characteristics of the firm and summarize these relations using regression coefficients—these results should provide useful proxies for researchers interested in tax convexity. In Section III, we briefly summarize our results.

I. Hedging and the Firm's Expected Tax Liability

If a firm's effective tax function is linear (the firm faces a constant effective marginal tax rate), the firm's expected tax liability is unaffected by the volatility of taxable income. However, if the firm faces some form of progressivity, then hedging that reduces the volatility of taxable income reduces the firm's expected tax liability. In this section we examine tax functions for a representative firm; this allows us to identify potential sources of tax-function convexity and to illustrate how this firm can reduce its expected tax liability by hedging.

A. Core Tax Structure—Statutory Progressivity

In Figure 1 we graph a firm's tax liability conditional on various levels of taxable income. The *core tax structure* pictured in Panel A ignores carrybacks, carryforwards, ITCs, and the AMT; the slope of the tax function is the period- t marginal tax rate. In this case, the firm's tax liability depends only on its current taxable income. As Panel A demonstrates, the major source of tax-function convexity is the asymmetric treatment of gains and losses—a tax rate of zero for losses but positive tax rates on profits. The close-up in Panel B shows the statutory progressivity of taxable income: the slope of the function varies with a 15 percent tax rate for income between 0 and \$100,000 and a 34 percent rate thereafter. (Assuming statutory progressivity over just the first \$100,000 is a reasonable approximation in many years. However, in some years marginal tax rates vary over several different ranges, thereby inducing regions of local convexity as well as concavity. We use actual statutory schedules in all our analysis other than Figure 1.)

To illustrate how convexity in the tax schedule can provide an incentive to hedge, consider the core tax structure pictured in Panel A of Figure 1. Assuming two equally likely outcomes—low taxable income with no associated tax liability (shown as point L) or high taxable income and a \$1 million tax




Figure 1. The federal income tax function conditional on the level of taxable income. The corporate tax liability is conditional on the level of taxable income assuming there are no tax-loss carrybacks or carryforwards, no investment tax credits, and no alternative minimum tax ("core tax structure" in Panel A) as well as when all of these features exist simultaneously ("extended tax structure" in Panel A). The results are for a hypothetical firm with taxable income in year t as shown on the horizontal axis and earnings in years $t - 3$ through $t + 18$, excluding year t , of \$310,500 (which is the median taxable income for all COMPUSTAT firms over the period 1980 to 1994). The investment tax credit is assumed to be \$1,300,000 (half the average investment tax credits (ITCs) for all COMPUSTAT firms in 1984) in year t and zero in all other years. The alternative minimum tax assumes that \$100,000 of noncash items are added to the tax base in each year and that the alternative minimum marginal tax rate is 20 percent (although it can be reduced to 15 percent using ITCs or to two percent using tax-loss carryforwards or carrybacks). The slope of the tax function represents the marginal tax rate, which is assumed to have a maximum value of 34 percent. The tax liability for the core tax structure is the tax bill in year t . The tax liability for the extended tax structure is the same number *minus* the contemporaneous effect of credits, losses, and AMT, as well as the present value tax benefit (using a ten percent discount rate) associated with carrying back or forward any losses or credits not used contemporaneously. Panel B of Figure 1 (Panel C) presents a close-up of the core (extended) structure tax function from Panel A of Figure 1 and highlights how the marginal tax rate can vary over different regions. The L (H) points in Panel A represent the low (high) outcome in which taxable earnings can assume one of two equiprobable values. In this case the expected tax liability is \$500,000 but could be reduced to zero if the volatility in earnings is eliminated.

(point H)—the firm's expected tax liability is \$500,000. However, if the firm could hedge and completely eliminate the volatility of taxable income, both its taxable income and tax bill would be zero. Thus, through hedging, this firm reduces its expected tax liability by \$500,000.

In Figure 2, we illustrate the effect of uncertainty about taxable income on the firm's expected tax liability. To model uncertainty, we calculate an expected tax liability at each \$100,000 increment along the horizontal axis by averaging the liabilities from 50 simulated income streams at each increment. For each simulation, the innovation in taxable income for years $t - 3$ through $t + 18$, excluding year t , is drawn from a normal distribution with a mean of \$310,500 (the median amount for all COMPUSTAT firms, 1980 to 1994); income for year t is drawn from a distribution with a mean as shown along the horizontal axis. In the low-volatility case, we calculate the standard deviation using a coefficient of variation of 0.09 (the 10th percentile of all COMPUSTAT firm-years, using a rolling average of historical data to calculate the coefficient for each firm-year observation during 1980–1994) applied to expected pretax income of \$310,500; the high-volatility case uses a coefficient of variation of 5.32 (the 90th percentile).²

Panel A of Figure 2 graphs the tax function for the core tax structure with uncertain taxable income. Given the option-like aspects of the core tax structure, the general results reported in Panel A are not surprising, especially for expected taxable income centered near the kink in the tax function: the greater the uncertainty, the higher the expected tax liability.³ To better un-

² To explore how reasonable it is to examine a range of income for a given level of volatility, as we do in Figures 1 and 2, the distribution of income within three volatility categories is shown in the following table. This table shows the percentage of firms with taxable income (U.S.\$ millions) for various income ranges. The "Low vol." category contains firms with volatility in the lowest 25 percent of all firms when they are ranked by volatility. "Med. vol." contains firms with volatility between the 37.5th and 62.5th percentiles of all firms. "High vol." applies to firms in the upper 25 percent of all volatility-ranked firms. Volatility is measured by the firm's absolute coefficient of variation of income.

	<-4	-4 to -2	-2 to 0	0 to 2	2 to 4	4 to 6	6 to 8	8 to 10	>10
Low vol.	1.8%	1.1%	1.8%	3.8%	4.8%	4.1%	3.5%	3.2%	76.0%
Med. vol.	19.8%	6.6%	20.4%	16.4%	7.1%	3.8%	3.0%	2.0%	20.8%
High vol.	19.7%	8.4%	34.6%	17.0%	4.6%	2.6%	1.6%	1.2%	10.2%

The distribution of earnings covers the range depicted in the figures and no particular earnings level should be thought of as obscure or unrepresentative. Therefore, the hypothetical cases we examine in Figures 1 and 2 can be considered reasonable for a typical firm. (Note that managers could use firm-specific probabilities analogous to those shown in the table as weights when determining whether they expect to be in a convex region of their tax function.)

³ Each curve in Figure 2 holds volatility constant as it maps expected taxable income into expected tax liability. We do this to isolate the effect of varying expected taxable income. We explicitly address the potential lack of independence of the first two moments of taxable income (as well as serial correlation) in Section II. Note that the firm-by-firm analysis in Section II does not hold volatility constant over various ranges of income, but instead allows volatility to vary according to firm-specific parameters.




Figure 2. The federal income tax function with uncertain taxable income. The relation between expected taxable income and expected tax liability when income is uncertain but there are no tax-loss carrybacks or carryforwards, no investment tax credits, and no alternative minimum tax (Panel A: core tax structure) and when all these features exist simultaneously (Panel B: extended tax structure). The results are for a hypothetical firm with an expected taxable income in year t as shown on the horizontal axis and expected income in years $t - 3$ through $t + 18$, excluding year t , of \$310,500 (which is the median taxable income for all COMPUSTAT firms during 1980 to 1994). The uncertainty in taxable income is modeled by drawing an innovation from a normal distribution with mean equal to the expected level of income and a standard deviation corresponding to a coefficient of variation of 0.09 (low volatility), 5.32 (high volatility), and 0.93 (medium volatility) which are the 10th percentile, 90th percentile, and median for the absolute coefficient of variation for all COMPUSTAT firms during 1980 to 1994, respectively. The investment tax credit is assumed to be \$1,300,000 (one-half the average investment tax credits (ITCs) for all COMPUSTAT firms in 1984) in year t and zero in all other

derstand the specific values reported, consider a one-period example of how uncertainty affects a firm's expected marginal tax rate.⁴ Assume that if a firm reports positive income, its marginal tax rate is 34 percent. Suppose that a firm makes \$100,000 in the current year but is uncertain about future income: next year it expects either to make \$310,000 with probability 0.50 or lose \$110,000 with probability 0.50. In the case of the \$110,000 loss, in one year the firm will obtain a refund of the current tax payments; consequently, its year- t marginal tax rate is $0.34 - 0.34/1.1 = 0.031$ (assuming a 10 percent discount rate). The firm's year- t tax rate is 34 percent in the \$310,000 case because it does not receive a refund. Thus, the firm has an expected year- t marginal tax rate of $0.5(0.031) + 0.5(0.34) = 0.186$ —this is the slope of the tax function. Overall, uncertainty about taxable income produces a flatter tax function for several reasons: (1) the probability of future losses that can be carried back to offset current taxes paid (as just illustrated), (2) the probability that the firm previously experienced losses that can be carried forward to offset current-period income, and (3) the possibility that expected-loss firms will realize positive earnings in some states (or in the future) and thus have a nonzero current-period effective tax rate. Our simulations allow us to capture this interaction between historical income and expected future income.

B. Extended Tax Structure

Other provisions in the tax code make this analysis materially more complex. First, we return to the case of certain taxable income. The *extended tax structure* in Panel A of Figure 1 reflects the impact of NOL carrybacks and carryforwards, ITCs, and the AMT. We initially assume that in years $t - 3$ through $t + 18$ (excluding year t) the firm has income of \$310,500. Various

⁴ See Graham (1996) for evidence on how the interaction between income uncertainty and tax code provisions affects the distribution of effective tax rates across firms.

years. The alternative minimum tax assumes that \$100,000 of noncash items are added to the tax base in each year and that the alternative minimum tax rate is 20 percent (although it can be reduced to 15 percent using ITCs or to two percent using tax-loss carryforwards or carrybacks). The slope of the tax function represents the marginal tax rate, which is assumed to have a maximum value of 34 percent. The expected tax liability for the core structure is the expected tax bill in year t . The tax liability for the extended tax structure case is the same number *minus* the contemporaneous effect of credits, losses and AMT, as well as the present value of the tax benefit (using a 10 percent discount rate) associated with carrying back or forward any losses or credits not used contemporaneously. The expected liabilities in the graphs are smoothed by using a 12-observation moving average, where each observation is based on the expected income on the horizontal axis and 11 other earnings figures, each incrementally \$100,000 smaller.

levels of year- t taxable income are depicted on the horizontal axis. The ITC is assumed to be \$1,300,000 in year t and zero in all other years. Our AMT analysis assumes that \$100,000 of noncash items (such as accelerated depreciation) is added back to the tax base in each year and that the AMT rate is 20 percent.⁵ For this extended tax structure, the tax liability is the tax obligation for the level of year- t income on the horizontal axis minus any tax credits, plus any AMT liabilities, minus the present value of the tax benefit (employing a 10 percent discount rate) associated with carrying back or forward any losses or credits.

Overall, this extended structure appears broadly convex—although less so than the core tax structure. The extended tax-code provisions reduce the effective convexity (relative to the core tax structure) arising from the asymmetric tax treatment of profits and losses. Although this extended tax structure still suggests a tax-based incentive to hedge, there are nonconvex regions. Panel C of Figure 1 provides a detailed view of this tax function to illustrate how the marginal tax rate varies. For example, suppose that a firm experiences a loss (or has tax credits) in year t that is sufficient to obtain a refund for all taxes paid in year $t - 3$ and a portion of taxes paid in year $t - 2$. In this case, the firm will receive a refund on taxes paid in the following order: (1) $t - 3$ at 34 percent, (2) $t - 3$ at 15 percent, and (3) $t - 2$ at 34 percent. This causes the slope of the firm's tax function to alternate between 15 percent and 34 percent over this range of income. Thus, local nonconvexities occur in this extended structure because in carrying losses and credits back, all taxes in a given year must be refunded before moving to the next year.

Tax-loss and tax-credit carryforwards can also cause the slope of the tax function to alternate between (the present value of) a low tax rate and a high tax rate. For example, the firm depicted in Figure 1 will use its \$1.3 million of ITCs to completely shelter income from $t - 3$ to the present (for current income up to nearly \$4 million). Assume that the firm carries forward enough ITCs to offset all taxes in $t + 1$ that are levied at a 15 percent rate, plus the taxes on one dollar that will be taxed at 34 percent. To calculate the firm's year- t marginal tax rate, we determine the present

⁵ The Tax Reform Act of 1986 introduced the alternative minimum tax in response to some well-publicized cases in which profitable firms paid no federal income taxes. The alternative minimum tax is essentially a new tax system that is separate but parallel to the traditional system. The AMT system has a larger base upon which taxes are determined because it limits the use of some tax-preference items such as accelerated depreciation. The AMT system also has its own tax rate. Currently, the rate is 20 percent but can be reduced to an effective two percent of the pretax AMT base by using the net operating loss carryforwards (or foreign tax credits) or, until 1991, to an effective 15 percent by using investment tax credit carryforwards. Taxes owed are the maximum of taxes determined from the traditional tax system and those from the AMT method. If, because of adding back certain tax-preference items, the AMT causes additional taxes to be paid, the additional amount can be carried forward indefinitely as a tax credit.

value of taxes owed from earning one more dollar of income in year t . If the firm earns one more dollar in period t , it will use ITCs to offset the taxes on this extra dollar but have insufficient ITCs left to carry forward to offset the dollar taxed at 34 percent in $t + 1$; hence, the slope of its year- t tax function is the discounted value of 0.34. In contrast, had the firm been earning one more dollar in the first place, the marginal tax rate at this higher level of income would be the present value of 0.15. In general, the existence of carryforwards leads to a tax function that varies between the present value of a high tax rate and the present value of a low tax rate until income is sufficiently large to exhaust the carryforwards completely.

In Panel B of Figure 2, we illustrate how uncertainty about taxable income affects hedging incentives in the extended tax code. As described above, we simulate uncertain taxable income streams to calculate an expected tax liability at each \$100,000 increment along the horizontal axis. We use a coefficient of variation of 0.93 (the median for all COMPUSTAT firms, 1980 to 1994) in the medium-volatility case. The simulations define the tax liability base and treat ITCs and the AMT as described above.

In Panel B, the low-volatility case using the extended tax structure is like the core tax structure, just shifted down to reflect the impact of the investment tax credit. However, the introduction of these additional provisions of the tax code creates a tax function that is less convex near zero, but with the convexity spread over a wider range. The range of convexity is further widened in the high-volatility case, although this high-volatility expected tax function is not strictly convex. Thus, volatility interacts with the level of taxable income and the provisions of the extended tax code to produce complicated effects on the shape of the tax function. This analysis suggests that the standard approach of using dummy variables to capture tax incentives induced by provisions such as ITCs or carryforwards is too simple. Because these provisions broaden the range of convexity but reduce the degree of curvature near the kink, they reduce incentives to hedge for some firms, yet increase them for other firms. Such effects are likely to be more complex than can be captured with dummy variables.

II. Potential Tax Savings from Hedging

In Section I we illustrate that the effective corporate tax function is broadly convex, although in certain situations the tax function is linear or even concave. We now examine the prevalence of convexity for a cross section of COMPUSTAT firms and calculate the magnitude of potential savings for selected hedging cases. Our measure of taxable income is taxable accounting earnings before extraordinary items and discontinued operations as reported by COMPUSTAT. Thus, our proxy for taxable income is derived from

firms' financial statements—not their tax returns. Although we would prefer to use tax-return rather than financial-statement data, firm-specific tax-return data are not publicly available.⁶

A. Simulation Methods

We assume that taxable income follows a random walk with drift and that hedging reduces the volatility of taxable income. We derive firm-specific estimates for the drift and volatility of the innovation on a rolling historical basis. For example, when determining the convexity of a firm's tax function in 1990, we use firm-specific data from 1973 to 1989 to estimate its drift and volatility parameters. Therefore, the predicted level of taxable income in 1990 is the actual value from 1989 plus a single innovation; the distribution of the innovation is estimated using all prior data. Analogously, when simulating this firm's tax function for 1991, we use data from 1973 to 1990. Any shifts in drift or volatility, whether endogenous (such as changes in operating decisions) or exogenous (such as the Tax Reform Act of 1986), thus are averaged with the other observations.

Employing the firm-specific estimates of drift and volatility, we first estimate the initial expected taxable income and the tax liability base (expected tax liability plus the present value of the tax benefits associated with tax-loss carrybacks and carryforwards, ITCs, and the AMT) for each firm in each year, assuming no reduction in volatility. This involves forecasting taxable income as a random walk through year $t + 18$ (to account for the effect of the carryforward period on the tax bill in year t), as well as considering taxable income in years $t - 3$, $t - 2$, and $t - 1$ (to incorporate the carryback feature of the tax code). Next, to investigate the potential benefit of hedging to reduce volatility, we repeat the procedure, changing only the volatility of year- t taxable income.

We compare the year- t tax liability base to that of the reduced-volatility case to determine the expected tax change associated with hedging. Because the tax liability base includes the present value of the benefits arising from tax-loss carrybacks and carryforwards, ITCs, and the AMT, our comparison does more than simply contrast static year- t tax liabilities between the hedging and no-hedging cases. For example, hedging might reduce the expected year- t tax liability and at the same time reduce the expected tax-loss that will be available to carry forward to offset future tax liabilities, all of which is captured in the tax liability base. Thus, our comparisons capture the dynamic element of year- t hedging that occurs through the provisions of the extended tax code. (Note, however, that our analysis does not reflect any corporate tax savings that would occur from increased interest tax deductions if hedging increases debt capacity; Stulz (1996)).

⁶ Some evidence of whether tax-return and financial-statement data differ materially is provided by Collins, Kemsley, and Shackelford (1997) in their analysis of the wholesale trade industry. Their evidence suggests that the distribution of taxable income from financial statements is generally similar to the distribution from tax returns.

We perform this procedure 50 times (starting in each instance with a new stream of random innovations, and hence a new forecast of taxable income) for each firm in each year. The expected tax savings is the average of the 50 simulations. We repeat the entire process three times, assuming volatility reductions of one, three, and five percent. (To help assess the reasonableness of these assumed volatility reductions, note that in Guay's (1999) examination of firms that initiate derivatives use, he reports an average reduction in stock return volatility of seven percent.)

B. Expected Tax Savings

Based on an analysis of 84,200 firm-year COMPUSTAT observations spanning from 1980 to 1994, we estimate that on average a one percent reduction in volatility reduces a firm's year- t expected tax liability by \$11,582; a three percent reduction reduces the expected liability by \$34,521; and a five percent reduction, by \$57,184 (see of Panel A of Table I). These potential tax savings represent 0.7 percent, 1.9 percent, and 3.1 percent, respectively, of the tax liability base. Because these averages include observations for which the effective tax function is concave or linear, they understate the expected savings. Indeed, approximately 25 percent of all firm-year observations are associated with concave effective tax functions and 25 percent are essentially linear.

For firms with convex tax functions, the mean savings are \$24,666, \$73,782, and \$122,718 for volatility reductions of one, three, and five percent, respectively, implying savings of 1.1, 3.3, and 5.4 percent of the tax liability base (see Panel B of Table I). The distribution is skewed; potential tax savings are nearly 50 percent and worth millions of dollars for firms in the 99th percentile of all firms facing a convex tax function.

Because we measure taxable income with accounting earnings derived from financial statements, our reported numbers potentially overstate the actual tax-statement volatility and savings. For example, the tax code usually recognizes restructuring expenses when they occur, while Generally Accepted Accounting Principles dictate that restructuring expenses are estimated and taken on financial statements in the year in which the event is initiated. We attempt to limit this problem in our primary specification by measuring income prior to extraordinary and discontinued items. However, to examine the extent that our approach fails to control the problem, we also estimate the benefits of hedging based on operating earnings (sales minus cost of goods sold minus selling, general, and administrative expenses)—that is, prior to “special items” and nonoperating income and expenses. Our qualitative results do not change for this alternative measure of income, although the specific magnitude of the potential tax savings is slightly smaller (\$120,192 versus \$122,718 expected tax savings among firms with convex tax functions for a five percent volatility reduction).

We also examine the tax savings due to hedging by year, specifically to check whether our results are affected by the reduction in statutory tax rates associated with the Tax Reform Act of 1986. In principle, the reduction

Table I
The Distribution of Tax Savings from Hedging

The estimated dollar amount and percentage tax savings from hedging are shown for various percentiles from the distribution across all firms. The savings are shown if the volatility of taxable earnings is reduced by one percent, three percent, and five percent. The data are from 84,200 firm-year observations for the years 1980 to 1994. The *tax liability base* is the expected contemporaneous taxable liability plus the present value of tax benefits associated with the following features of the extended tax structure: tax-loss carrybacks and carryforwards, the investment tax credit, and the alternative minimum tax.

	Percentage of Tax Liability Base for Volatility Reductions of			Dollar Amounts for Volatility Reductions of		
	1%	3%	5%	1%	3%	5%
Panel A: Tax Savings for All Observations						
Mean	0.7%	1.9%	3.1%	\$11,582	\$34,521	\$57,184
5th percentile	-0.1	-0.1	-0.1	-8,930	-27,210	-45,650
25th percentile	-0.0	-0.0	-0.0	-15	-47	-84
Median	0.0	0.0	0.0	133	406	685
75th percentile	0.9	2.6	4.3	2,704	8,143	13,670
95th percentile	3.0	9.1	15.1	46,423	139,251	231,476
99th percentile	7.3	22.7	39.2	282,033	841,483	1,397,311
Panel B: Tax Savings for 55,059 Observations of Firm-Years with Convex Tax Functions (Positive Tax Savings from Hedging)						
Mean	1.1%	3.3%	5.4%	\$24,666	\$73,782	\$122,718
5th percentile	0.0	0.0	0.0	25	76	128
25th percentile	0.0	0.0	0.0	310	943	1,588
Median	0.5	1.5	2.5	1,650	4,976	8,348
75th percentile	1.4	4.1	6.8	8,098	24,373	40,784
95th percentile	3.7	11.2	18.7	93,308	279,412	464,135
99th percentile	8.9	28.3	48.9	422,894	1,264,305	2,113,006

in statutory rates could induce a concavity in tax functions following this tax code change: firms might prefer not to hedge if hedging would reduce the likelihood of carrying back post-tax-reform losses to offset prior taxes paid at higher rates. Overall, the effects from the reduction in statutory rates appear to be material in only one year; in 1988, the first year in which the reduced tax rates were completely implemented, the average dollar savings from a five percent reduction in the volatility of income was only \$15,000.

C. The Impact of Individual Tax Code Provisions

Our simulation methods allow us to analyze the impact of individual tax-code provisions on the incentives to hedge. In effect, we can turn tax provisions on and off to identify the incremental impacts of carryforwards, carrybacks, ITCs, and the AMT.

Table II
Incremental Impact of Tax Code Provisions on the Convexity
of the Effective Tax Function

The difference in expected tax savings from a five percent volatility reduction in taxable income when various provisions (carrybacks, carryforwards, investment tax credits) are added one at a time to a “no features” core tax structure. The columns report four measures of expected tax savings: the average change in dollar savings ($\Delta\$$ Savings, in 000’s); the average change as a percentage of the tax base ($\Delta\%$ Savings); the change in the percentage of firms with convex effective tax functions ($\Delta\%$ Convex); and the change in the percentage of firms with concave effective tax functions ($\Delta\%$ Concave). The alternative minimum tax (AMT) row shows the incremental effect of adding the AMT to a tax code that already includes carrybacks, carryforwards, and the investment tax credits (ITCs). The All provisions row reports the effect of simultaneously adding all four features of the tax code (i.e., creating an extended tax code) to a “no features” core tax structure. To be considered “convex” (“concave”) in this table, savings from hedging must be at least 0.5% (–0.5%). The data are from 84,200 observations on COMPUSTAT firms for the years 1980 to 1994.

Tax Provision	$\Delta\$$ Savings	$\Delta\%$ Savings	$\Delta\%$ Convex	$\Delta\%$ Concave
Carryforwards	–7.4	–1.5	–11.7	0.4
Carrybacks	–0.3	–4.5	–11.6	3.0
ITC	–0.2	0.0	0.0	0.0
AMT	–9.7	–0.1	2.7	0.2
All provisions	–3.7	–1.9	–7.1	1.0

In Table II, we report several measures of the incremental effect of the various provisions: the change in dollar savings, the change in savings as a percentage of the tax base, and changes in the percentage of firms with convex and concave effective tax functions. Tax-loss carryforwards and carrybacks each reduce the general convexity of the effective tax function; in each case the percentage of companies with concave functions rises and the tax savings fall, measured either in dollars or as a percentage of the tax base. Figures 1 and 2 suggest that the extended tax code provisions reduce the benefits from hedging for firms near the kink, and they increase the benefits for firms away from the kink. The numbers in Table II indicate that for carryforwards and carrybacks, the former effect dominates the latter.

Carryforwards and carrybacks also introduce concavities into some tax schedules. To examine this effect in more detail, we perform additional untabulated simulations. We decompose the impact of carrybacks and carryforwards into two parts: (1) their interaction with the current tax liability results in a tax function that is everywhere convex, and (2) the present value of tax benefits from using losses to offset future or past taxable income, which can be concave over portions of its range. The convexity of the first part represents the curvature in the core tax schedule, spread over a broader range by the carryback and carryforward features. In the second part, there are both convex and concave components. The convex effect results from limitations on the use of tax-code features (for example, although losses can be carried forward, discounting reduces the value of a one dollar current-

period loss below a dollar), as well as the possibility that a firm might not be able to use the current loss to reduce future tax liabilities. The concave component of the second part occurs because, as the expected income increases, the probability of experiencing a loss declines; thus, the expected value of incremental benefits from using carrybacks or carryforwards also declines, producing a concave region. If this concave component dominates, the effective tax function is concave.

The numbers in Table II suggest that investment tax credits have little impact on the convexity of the effective tax function. We believe that this is a noteworthy result. The logic of DeAngelo and Masulis (1980) suggests that limitations on the use of ITCs should introduce convexity into the effective tax function; if the firm's tax liability falls below the amount necessary to use the firm's current ITCs fully, the unused ITCs must be carried forward, reducing their present value and increasing the firm's effective tax liability. Researchers examining corporate hedging policy have used the presence of ITCs as an instrument to identify a tax-based incentive to hedge (e.g., see Nance, Smith, and Smithson (1993)). Yet, if the incremental impact of ITCs on convexity is negligible, then any association between ITCs and hedging is more likely to reflect other considerations. (For instance, ITCs might proxy for the mix of growth options versus assets in place in a firm's investment opportunity set.)

The incremental effect of the alternative minimum tax depends on the benchmark. If one examines adding the AMT to a tax function without carrybacks, carryforwards, or ITCs, the incremental effect is virtually zero. This is because the primary effect of the AMT on convexity arises from its interaction with these other tax provisions.⁷ In Table II, the AMT row reports the incremental effect of adding the AMT to a tax code that already includes tax-loss carrybacks, carryforwards, and the ITC. In the aggregate, the AMT makes the effective tax function less convex—both the dollar and percentage reduction in expected taxes are reduced when the AMT is added. Interestingly, both the fraction of firms with convex tax functions and the fraction with concave functions increase. The effect of the AMT on tax function convexity is complex because of its interaction with other provisions of the tax code; it can lead to tax rates of two through 20 percent over some ranges of income but tax rates from the usual statutory schedule over other ranges.

The last row of Table II adds all the provisions; collectively, they make the effective tax function less convex. Given the other findings in Table II, results when all of the tax-code provisions are included are unsurprising; each result falls within the range of the incremental effects of the individual pro-

⁷ COMPUSTAT does not provide information about the tax-favored items that the AMT adds back to the traditional tax base. Consequently, our analysis uses the traditional tax base to study the effect the AMT has on the shape of the tax function by either (1) changing the tax rate to which a firm is subject for given ranges of income, and/or (2) stipulating limitations on the effect of the other provisions of the tax code.

Table III
The Relation between Tax Savings from
Hedging and Firm Characteristics

The table shows the relation between various firm characteristics and the expected tax savings that can result from reducing the volatility of pretax earnings. The data are grouped into quartiles based on the percentage tax savings due to a five percent reduction in volatility. The mean value within each quartile is shown for various firm characteristics. Market value is the book value of debt plus the market value of common equity. Debt-to-value is the book value of debt divided by the market value of the firm. Level of earnings is the dollar value (in millions) of pre-tax earnings. Earnings volatility is the absolute coefficient of variation for pre-tax earnings. Serial correlation of earnings is the first-order autocorrelation coefficient for pre-tax earnings. Percent with positive earnings is the percentage of observations for which pre-tax earnings are positive. Percent with ITC measures the percentage of firms with positive investment tax credits on their books, and percent with NOL measures the percentage of firms with positive net operating loss carryforwards on their books. The data are from 84,200 firm-year observations for the years 1980 to 1994.

	Smallest Quartile	Quartile 2	Quartile 3	Largest Quartile
Mean percentage savings	-3.0%	0.0%	1.2%	11.7%
Mean dollar savings	-\$48,430	-\$9,770	\$169,560	\$128,400
Mean value for:				
Market value	1381.5	3607.3	2095.9	615.5
Debt-to-value ratio	0.176	0.156	0.185	0.195
Level of earnings	39.77	194.32	65.93	-3.28
Earnings volatility	1.77	0.90	8.89	13.54
Serial correlation of earnings	0.61	0.73	0.46	0.41
Percent earnings greater than zero	52%	97%	92%	41%
Percent with ITC	3.9%	14.1%	3.7%	0.6%
Percent with NOL	35.3%	6.7%	21.9%	42.6%

visions. Nonetheless, the numbers suggest that the interactions are sufficiently complex that the specific magnitudes would be difficult to forecast based only on the reported incremental effects.

D. Tax Savings and Firm Characteristics

To derive a better understanding of how these potential tax savings vary across firms, we subdivide the sample into quartiles based on the percentage tax savings from hedging. The numbers in Table III indicate that the average tax savings for the smallest quartile are negative—there is a tax disincentive to hedge. For this set of firms, a five percent reduction in the volatility of taxable income results in a three percent tax increase or an average tax increase of \$48,430. For the second quartile, the average percentage tax savings are virtually zero or an average dollar increase of \$9,770; the third quartile savings are 1.2 percent, or \$169,560; for the largest quartile, the savings are 11.7 percent, or \$128,400. This distribution is quite skewed. For example, untabulated results indicate that in the 90 to 95th

percentiles the averages are 11.5 percent or \$134,410, in the 95 to 100th percentiles, 28.9 percent or \$45,986. (Note that by grouping on percentage savings, some firms with high percentage savings have expected tax liability bases that are near zero; thus small dollar savings can represent large percentage savings.)

The remaining data in Table III help identify firm characteristics associated with larger tax-related incentives to hedge. The firms in the second quartile have few tax-based incentives to hedge compared to firms in the third and fourth quartiles. On average, the second-quartile firms have higher market values, lower leverage, higher levels of income, less frequently reported losses, lower income volatility, more ITCs, and fewer NOL carryforwards. In general, the firms in quartile 2 are large, profitable firms that are located on the linear (high income) portions of their tax functions. An untabulated analysis of observations in the 95th and 99th percentiles confirms these findings: firms with the largest tax incentive to hedge are small-to-medium sized, with earnings that are expected to be near zero, frequently shifting between losses and profits.

We perform regression analysis using the expected percentage tax savings from a five percent reduction in the volatility of taxable income (*TAX SAVE*(5%)) as the dependent variable. We include the following explanatory variables (with variable names and *t*-statistics in parenthesis): indicator variables identifying taxable income between $-\$500,000$ and zero (*TI(NEG)*; 40.3) or between zero and $\$500,000$ (*TI(POS)*; 11.3), income volatility measured with the absolute coefficient of variation (*VOL*; 10.9), first-order serial correlation in income (*RHO*; -41.1), a dummy variable indicating the existence of ITCs (*ITC*; -9.5), a dummy indicating the existence of NOL carryforwards (*NOL*; 36.5), and the NOL dummy interacted with the small loss (*NOL*·*TI(NEG)*; -18.3) and small gain (*NOL*·*TI(POS)*; -8.1) indicator variables:

$$\begin{aligned} \text{TAX SAVE}(5\%) = & 4.88 + 7.15 \cdot \text{TI(NEG)} + 1.60 \cdot \text{TI(POS)} + 0.019 \cdot \text{VOL} \\ & - 5.50 \cdot \text{RHO} - 1.28 \cdot \text{ITC} + \text{NOL} \cdot (3.29 - 4.77 \cdot \text{TI(NEG)} \\ & - 1.93 \cdot \text{TI(POS)}). \end{aligned} \quad (1)$$

The regression is estimated using 84,200 firm-year observations; the adjusted R^2 is 8.2 percent. In addition to describing the relation between firm characteristics and the tax incentive to hedge, these estimated coefficients can be used by researchers who are investigating corporate hedging but do not have simulated tax functions available.

The estimated coefficients indicate that expected tax savings increase if income is near zero (close to the kink in the tax function), as volatility increases (given the option-like structure of the tax code), and as serial correlation becomes more negative (because the firm is more likely to switch between profits and losses). The ITC dummy has a negative sign. This is inconsistent with the prediction that is derived from the DeAngelo and Masulis (1980) analysis and (like the results in Table II) suggests that an ob-

served association between ITCs and hedging is driven by other than tax motives. Consistent with the DeAngelo–Masulis analysis, the existence of NOL carryforwards increases the incentive to hedge; however, this effect is attenuated for firms with small profits, and more than offset (indicating a concavity in the tax function) for firms with small losses. To understand this last result, consider a firm with a NOL carryforward but with its distribution of expected profits centered on a small loss—in many circumstances it will be unable to utilize the NOLs. By hedging, the firm tightens its distribution of profits around its mean, further reducing the expected benefit provided by its carryforwards. Thus, one effect of NOL carryforwards is to provide a disincentive to hedge for a firm with expected losses.

To keep the interpretation of the regression analysis manageable, we limit the interaction of the various explanatory variables; yet, one expects additional interaction effects. For example, income volatility might provide a material tax incentive to hedge for firms with income near zero, yet little incentive for firms whose income is high. Likewise, the carryback feature is unimportant to a firm that loses money in each year, but quite important—has its largest effect on convexity—for a firm with volatile taxable earnings, expected taxable earnings near zero, and a negative serial correlation coefficient. In the Appendix, we analyze interactive effects of volatility, serial correlation, and the expected level of income. The results show that, with these additional interactions considered, tax benefits of hedging provided by these income characteristics are consistent with the regression coefficients: serial correlation and level of expected earnings reduce the benefits; volatility increases them.

Finally, note that our simulation methods allow us to focus on tax benefits from hedging isolated from other factors that might affect the incentive to hedge. It is possible that nontax factors provide incentives that work either with or against the factors identified here. For example, we have noted that firms with negative taxable income have tax incentives to hedge. Yet, if these firms are near financial distress, the option value of equity might reduce incentives to hedge.

E. Industry Effects

We investigate whether the incentive to hedge is particularly large in certain industries. We define industries by grouping the data into 75 different two-digit SIC codes (as identified by COMPUSTAT). As our metric for large benefits, we identify industries that have twice as many observations in the upper one percent tail of the benefits distribution as they do in the overall sample. Seven industries meet this standard: agricultural production and livestock (SIC 2), agricultural services (SIC 7), metal mining (SIC 10), communication (SIC 48), motion pictures (SIC 78), health services (SIC 80), and educational services (SIC 82). Because not all of these industries are categorized as engaging in pervasive corporate hedging, it may be difficult for some firms to achieve the estimated potential gains from hedging for tax purposes.

III. Conclusions

We have attempted to provide a measure of the convexity in corporate tax functions that is more direct than past measures, thereby providing more precise estimates of the potential tax savings from hedging. We examine how uncertainty about future taxable income interacts with major provisions of the tax code, including statutory progressivity, tax-loss carrybacks and carryforwards, investment tax credits, and the alternative minimum tax. Using data from COMPUSTAT, we simulate the tax savings from reducing the volatility of taxable income. We find that: (1) Among firms facing a convex tax function, the average tax savings from a five percent reduction in volatility are \$122,718 or about five percent of the tax liability base (defined as the sum of the current tax liability plus the present value of the tax benefits associated with tax-loss carrybacks and carryforwards, the ITCs, and the AMT). (2) The distribution of potential tax savings is quite skewed. Although for approximately 75 percent of firms there is little tax-based incentive to hedge, in extreme cases the savings appear substantial. For example, among firms facing convex tax functions, our estimates suggest that those in the upper one percent of all observations could save over \$2,000,000 annually if they reduce the volatility of taxable income by five percent. (3) The asymmetric treatment of profits and losses drives much of the observed convexity. (4) NOL carryforwards and carrybacks smooth the tax function, reducing its most extreme curvature, but broadening the range of convexity. (5) The AMT introduces a modest increase in convexity; ITCs have little impact.

There are five points to remember when judging the importance of the magnitude of the potential tax benefits that we have identified. First, for firms with convex effective tax functions, the tax savings of hedging are not mutually exclusive from the hedging benefits of controlling underinvestment problems, increased debt capacity, or reducing compensating differentials for risk-bearing by various classes of the firm's claimholders. Thus, the total benefit of hedging is the sum across the various motives. With independent instruments, different hedging targets can be achieved. Therefore, in principle, with the appropriate choice of hedging instruments, a firm can simultaneously manage the impact on its value, reported income, and taxable income.

Second, our methods simulate the change in expected taxes from reducing the volatility of taxable income only in year t . Most firms that hedge in one year also hedge in subsequent years. Our reported numbers thus potentially understate the tax savings from the adoption of an ongoing hedging policy.

Third, we have simulated specific provisions of the tax code, yet other facets may provide material incentives to hedge. Firms do not file a single consolidated tax form; different tax jurisdictions imply that firms face more convexity than our methods recognize. The tax code also contains special deductions related to specific expenses and sources of income that create convexities in special circumstances. Moreover, hedging that reduces cash-flow volatility potentially increases debt capacity, providing larger interest deductions. Our methods do not fully recognize this texture in the tax structure and thus potentially understate tax-related hedging incentives.

Fourth, to determine an optimal hedging program, this tax-related benefit of hedging must be compared to the cost of hedging. For forwards, futures, options, and swaps, this cost has three major components: (1) out-of-pocket costs such as brokerage fees in futures markets, (2) the implicit cost of the bid-ask spread, and (3) the costs of the internal control systems to run the hedging program (including the opportunity cost of management's time). The first two costs have fallen with the growth of the derivatives markets. For example, bid-ask spreads in dollar interest rate swap markets are now generally less than five basis points. There is little evidence on the control-system costs. But following the derivatives debacles at Gibson Greetings, Proctor and Gamble, and Barings Bank, there is more oversight at the level of the corporate board, and companies have been devoting more resources to ensure that hedging programs are better controlled.

Fifth, although our discussion is framed in terms of reducing the volatility of taxable income by using hedging instruments, the identified tax savings can be achieved by other means. For example, a firm could shift income from a "high profit" year into a "loss" year to achieve tax savings arising from tax code convexity. Similarly, firms can change other policies—for example, the choice between leasing or buying assets to shift income and reduce expected tax liabilities.

This analysis represents something of an intermediate good. We have focused on a necessary condition for hedging to reduce expected tax liabilities—the convexity of the firm's effective tax schedule. But in estimating the potential tax savings from hedging, we simply assume that the firm can reduce the volatility of current taxable income by one, three, or five percent. For some firms, these reductions are conservative while for others they are undoubtedly infeasible. It is important to understand the joint distribution of tax-function convexity and potential volatility reduction. For example, suppose that the average convexity paired with the average volatility reduction implies a material reduction in expected tax liabilities. Yet, this calculation might either dramatically overestimate or underestimate the potential tax saving from hedging if these factors are negatively or positively correlated. For example, if those firms with the greatest convexities realize little volatility reduction from hedging, while firms with large feasible reductions in the volatility of taxable income face linear tax schedules, then hedging would provide few tax savings.

Appendix

To gain a better understanding of how the interactions between the level, volatility, and serial correlation of taxable income affect tax savings across firms, we subdivide the sample. We divide firms into groups based on expected taxable earnings, the variability of taxable earnings, and the serial correlation in earnings (see Figure A1). For expected earnings, Group 1 contains all firms with negative expected taxable earnings—approximately 32 percent of the sample population. Group 2 contains firms with expected taxable earnings in the progressive region of the tax schedule (in most years,




Figure A1. Tax savings from hedging conditional on earnings characteristics: Expected earnings, absolute coefficient of variation, and first-order serial correlation. Tax savings expressed as a percentage of the sum of the expected tax bill in year t plus the present value of the benefit from reducing past and future tax bills due to the extended provisions of the tax code, based on a three percent reduction in the volatility of taxable earnings. The results are shown conditional on groupings of expected taxable earnings, the absolute coefficient of variation, and first-order serial correlation. Each characteristic is grouped as follows: (1) expected value: earnings less than zero (approximately 32 percent of the sample population), earnings in the progressive region of the statutory tax schedule (which is earnings between zero and \$100,000 in most years; approximately two percent of the sample population), and evenly split between two groups for the remaining data, (2) absolute coefficient of variation, positive and negative earnings separately: quartiles, and (3) correlation coefficient: less than zero (approximately one percent of the sample population), and evenly split between three groups for the remainder of the data. The data represent 84,200 firm-year observations spanning the years 1980 to 1994.

earnings between zero and \$100,000). This group contains approximately two percent of the sample population. Groups 3 and 4 subdivide the remaining firms. Scanning across the four groups, the largest percentage tax savings occur in Groups 1 and 2 (those nearer the kink, or statutory-progressive region, of the tax schedule).⁸

Figure A1 further subdivides firms into quartiles based on the absolute coefficient of variation for earnings. For positive earnings firms, expected tax savings rise almost uniformly with the coefficient of variation; the higher the volatility of the firm's taxable earnings, the greater the reduction in expected taxes. The same general pattern holds for loss firms, although it is not monotonic.

The firms in Figure A1 are further grouped by the serial correlation coefficient of taxable earnings. Quartile 4 contains all firms with negative estimated serial correlation coefficients—approximately one percent of the sample. Quartiles 1 through 3 contain the remaining firms. The tax savings from hedging are greater for firms with lower serial correlation coefficients. These are the firms that are more likely to alternate between profits and losses. Because firms that alternate between profits and losses from year to year also have expected earnings distributions that range over the kink, it is not surprising that these firms have relatively large tax incentives to hedge. Firms that alternate between profits and losses also are able to utilize the features of the tax code.

Comparing across the histograms in Figure A1, it is evident that the convexity induced by these characteristics of taxable earnings are not independent (see footnote 3). However, the general message holds that lower expected value, higher volatility, and lower serial correlation of taxable earnings increase the tax savings from hedging.

REFERENCES

- Collins, Julie, Deen Kemsley, and Doug Shackelford, 1997, Transfer pricing and the persistent zero taxable income of foreign-controlled U.S. corporations, *Journal of the American Taxation Association* 19, 68–83.
- DeAngelo, Harry, and Ronald W. Masulis, 1980, Optimal capital structure under corporate and personal taxation, *Journal of Financial Economics* 8, 3–30.
- Dolde, Walter, 1993, The trajectory of corporate financial risk management, *Journal of Applied Corporate Finance* 6, 33–41.
- Froot, Kenneth, David Scharfstein, and Jeremy Stein, 1993, Risk management: Coordinating corporate investment and financing policies, *Journal of Finance* 48, 1629–1658.
- Graham, John R., 1996, Debt and the marginal tax rate, *Journal of Financial Economics* 41, 41–73.
- Guay, Wayne R., 1999, The impact of derivatives on firm risk: An empirical examination of new derivative users, *Journal of Accounting and Economics* 26, 319–351.
- Main, Brian, 1983, Corporate insurance purchases and taxes, *Journal of Risk and Insurance* 50, 197–223.

⁸ The across-group relative magnitudes of the savings depicted in Figure A1 are relatively unchanged if we deflate by market value of the firm, rather than by tax liability base.

- Mayers, David, and Clifford Smith, 1982, On the corporate demand for insurance, *Journal of Business* 55, 281–296.
- Mayers, David, and Clifford Smith, 1987, Corporate insurance and the underinvestment problem, *Journal of Risk and Insurance* 54, 45–54.
- Nance, Deana, Clifford Smith, and Charles Smithson, 1993, On the determinants of corporate hedging, *Journal of Finance* 48, 267–284.
- Smith, Clifford, 1995, Corporate risk management: Theory and practice, *Journal of Derivatives* 2, 21–30.
- Smith, Clifford, and René Stulz, 1985, The determinants of firms' hedging policies, *Journal of Financial Quantitative Analysis* 20, 391–405.
- Stulz, René, 1984, Optimal hedging policies, *Journal of Financial and Quantitative Analysis* 19, 127–140.
- Stulz, René, 1996, Rethinking risk management, *Journal of Applied Corporate Finance* 9, 8–24.