

Taxing Risk*

Narayana Kocherlakota

President
Federal Reserve Bank of Minneapolis

Economic Club of Minnesota

Minneapolis, Minnesota

May 10, 2010

*This topic is discussed in greater depth in "Taxing Risk and the Optimal Regulation of Financial Institutions," Economic Policy Paper 10-3 (May 2010), Federal Reserve Bank of Minneapolis.

In the mid-2000s, we—as investors, home buyers, and bank lenders—collectively bet that house prices would not fall by 30 percent in most major metropolitan areas in three years. We were wrong. This mismatch between our expectations and our realizations was the ultimate source of the financial crisis of 2007-09.

The Congress of the United States is currently considering legislation to restructure financial regulation. However, no matter how well-written or how well-intentioned the legislation may be, no law can completely eliminate the kinds of collective investor and regulator mistakes that lead to financial crises. These mistakes have taken place periodically for centuries. They will certainly do so again. And once these crises happen, there are strong economic forces that lead policymakers—for the best of reasons—to bail out financial firms. In other words, no legislation can completely eliminate bailouts. Any new financial regulatory structure must keep this reality in mind.

My theme today is that, although bailouts are inevitable, their magnitude can be limited by taxes on financial institutions. I arrive at this conclusion about the usefulness of taxes by thinking through an analogy that I'll develop at some length. I will argue that, knowing bailouts are inevitable, financial institutions fail to internalize all the risks that their investment decisions impose on society. Economists would say that bailouts thereby create a risk "externality." There is nearly a century of economic thought about how to deal with externalities of various sorts—and the usual answer is through taxation. I will suggest that the logic that argues for taxation to deal with other externalities is exactly applicable in this case as well. As always, any views I share today are my own, and not necessarily those of others in the Federal Reserve System.

My case for taxing banks is distinct from one often heard in popular discourse. This latter logic usually runs: Taxpayers put X dollars into the banking sector, and the banking sector should repay all of that money. This argument is, I think, fundamentally grounded in a desire for revenge: Some big banks—

perhaps now gone forever—took our money, and so all big banks must pay. Taxes are seen as punishment—a means of exacting retribution from the guilty (them) to compensate the innocent (us).

My story is different. At least some big banks did make socially undesirable choices. But — in large part—they were led to make those choices by incentives within the tax and regulatory system. Parts of these incentives were shaped by the ultimately correct expectation that some bailouts would take place in the event of a financial crisis. These government guarantees—no matter how implicit they might have been—created an incentive for financial institutions to make socially undesirable choices. Taxation is a useful way to correct this incentive.

Earlier, I claimed that bailouts of financial institutions are certain to occur in financial crises. Why do I say this? There are many forces at play, but I believe that the strongest has to do with the very nature of financial intermediation. Investors in financial institutions always want the ability to pull out their funds quickly. For this reason, financial institutions' liabilities often take the form of short-term debt and deposits. But such short-term financing instruments are intrinsically prone to self-fulfilling crises of confidence that economists term "runs."

Imagine that Bank X needs \$100 billion of one-day loans to survive. This means that for a given lender to be willing to make a \$1-billion, one-day loan to Bank X, that lender has to believe that Bank X will get another \$99 billion in one-day loans. Then, Bank X may fail simply because every possible lender believes correctly that no lender is willing to lend to Bank X. Such a crisis of confidence can occur regardless of the true condition of Bank X.

This story is hardly a new one. It's exactly why we have deposit insurance: to prevent runs by reassuring bank depositors that their money is safe. But the story has huge consequences for how governments operate. In a financial crisis, there is a tremendous sense of uncertainty. There are some truly insolvent financial firms out there—but no one knows for sure which ones they are. And during a

crisis, the panic in the air means that any institution—even one with solid fundamentals—may be subjected to a run if its investors lose confidence in its solvency.

In such an atmosphere, contagion effects become extremely powerful. Even a slight loss by one short-term creditor can lead all short-term lenders to rush to the safety of Treasury bills. Such flight would endanger the survival of key financial institutions, even if they are fundamentally sound. Governments cannot risk such systemic collapse, and so during times of crisis, they end up providing debt guarantees for all financial institutions. Thus, policymakers inevitably resort to bailouts even when they have explicitly resolved, in the strongest possible terms, to let firms fail.

Many observers of the events of September 2008 have emphasized the need for better resolution mechanisms. Different people mean different things by this, but most want to impose losses on debt holders. I'm not opposed to faster resolutions of bankruptcies. But I do not believe that better resolution mechanisms will end bailouts. Indeed, I'm led to make a prediction. No matter what mechanisms we legislate now to impose losses on creditors, Congress, or some agency acting on Congress' behalf, will block them when we next face a financial crisis. And Congress will do so for a very good reason: to forestall a run on the key players in the financial system.

So, that's my first point: Bailouts are inevitable. Let me move to my next point: Bailouts create inefficiencies in the allocation of real investment. Here's what I mean. Financial institutions make investments that are, by their very nature, risky—that is, their returns are not certain. They finance these investments, at least in part, using debt and deposits.

Now, imagine for a moment that we live in a world without bailouts, so that the government does not provide debt guarantees or deposit insurance. If a financial institution decided to increase the risk level of its investment portfolio, its debt holders and depositors would face a greater risk of loss. By way of compensation for that greater risk, they'd demand a higher yield. As a result, in the absence of

government guarantees, financial institutions would find it more costly to obtain debt financing for highly risky investments than for less risky ones. This effect, on the margin, would curb a firm's appetite for risk. It would have an especially powerful effect on highly leveraged financial institutions, because high debt-to-asset levels mean higher risk of being unable to fulfill debt obligations.

But now return to the real world, with deposit insurance and debt guarantees, and the inevitability of government bailouts. Even if they only kick in during financial crises, these guarantees change this natural market relationship between risk and cost. The depositors and debt holders are now partially insulated from increases in investment risk, and so do not demand a sufficiently high yield from riskier firms. Financial institutions take on too much risk, because they are no longer deterred from doing so by the high costs of debt finance. And this missing deterrence is especially relevant for firms that are highly leveraged, because they should be paying out especially high yields on their debts.

In this way, the expectation of bailouts leads to too much capital being allocated toward overly risky ventures. These misallocations of capital don't create the collective mistakes in predictions that generate financial crises. But the misallocations do mean that society loses a lot from those mistakes—a lot more than is efficient.

What kind of policy would be useful in correcting this inefficiency? In what follows, I will offer an analogy from a completely different arena of public policy that can help us think through this key question.

Consider a factory that creates air pollution as a byproduct of operation. When the firm that owns the factory chooses to produce more output, it incurs various *private* costs: more raw materials, more labor, and so on. But the production increase also generates more pollution that will be absorbed by the surrounding community. The pollution is a *social* cost of production not paid for, or "internalized," by the firm that generates it. Economists refer to such costs as "externalities."

This same distinction between private and social costs applies to financial institutions that are facing debt guarantees. Such guarantees imply that some portion of the risk produced by a firm's investment decisions is absorbed by taxpayers. In making decisions about what to invest in, the firm ignores that portion of risk. It is a social cost of the project that the private firm does not internalize. Just like the pollution, the risk borne by taxpayers is an externality—what I will call a “risk externality.”

This analogy is useful because economists know a lot about how to deal with externalities. We can exploit their years of research to address the problem of financial regulation when government bailouts are inevitable. In particular, that long history of thought says that the best way to correct externalities is by providing the right kinds of incentives through appropriate taxes.

Let me be more specific. Again, let's think about the firm with a polluting factory. Many of its choices affect the amount of pollution produced, including the amount of time that the firm runs the factory during the workweek, the sorts of antipollution technology employed, and the kind of energy used to run the factory. Now, the government could regulate the firm's pollution levels by controlling each and every one of these choices. However, to do so, the government has to choose how to trade off these three (and other) factors against one another.

Its trade-off decisions will be influenced by both pollution considerations and cost factors. If antipollution technology is cheap, the government may simply require the firm to invest in that. But if antipollution technology is expensive, the government may require the firm to switch to using natural gas instead of coal. Making these trade-offs requires a tremendous amount of firm-specific information and firm-specific cost minimization. To put it mildly, historical evidence suggests that governments are not very good at such micromanagement of factory-level operation; that's why we have private markets.

The solution to this difficulty is to regulate the *amount* of pollution produced by the firm, rather than *how* the firm produces that pollution. The central problem here is that pollution has a social cost

that the firm does not internalize when choosing its level of production. From society's point of view, the firm will overproduce pollution. However, the firm *will* choose the socially efficient level of pollution if it is required to pay for—or internalize—the social cost of the pollution.

More concretely, suppose that the firm is told, before choosing its level of production, that the government will measure the amount of pollution that the firm generates and charge the firm a tax that is exactly equal to the social cost of that quantity of pollution. This policy generates a tax *schedule* that translates the amount of pollution generated into an amount paid by the firm. If the firm knows that it faces this tax schedule, its costs of production will include the social cost of pollution, along with labor, materials, energy, and the like. In this way, what was external to the firm becomes internal. As a result, the firm will choose the socially efficient level of production. Just as importantly, it will automatically choose to produce that pollution—and its other more beneficial outputs—in a cost-minimizing fashion. The government does not need to solve the firm's cost-minimization problem.

These lessons about pollution regulation translate directly into lessons about financial regulation. As in the pollution case, a financial institution should be taxed for the amount of risk it produces that is borne by taxpayers. The firm will then choose the socially optimal level of risk.

Here's my preferred policy. The firm is told that the government will estimate the expected, discounted value of bailouts that the financial institution (or any of its stakeholders) will receive in the future. I say "expected" because the amount of the bailout is uncertain (and indeed is likely to be zero much of the time). I say "discounted" because the bailout may be received next year or in 30 years, and we need to discount accordingly. Clearly, this estimate will depend on many firm choices and attributes, including its leverage ratio, the maturity structure of its liabilities, the risk characteristics of its investment portfolio, and its incentive compensation schemes. For example, the expected bailout will be higher for firms with highly risky investments than for firms with less risky portfolios.

Having done this calculation, the government then charges the firm a tax that is exactly equal to the expected discounted value of the firm's bailouts. Just as in the pollution example, this measurement-plus-taxation policy confronts the firm with a tax schedule that translates its choices into a cost paid by the firm. The tax amount exactly equals the extra cost borne by the taxpayers because of bailouts, appropriately adjusted for risk and the time value of money. Knowing that it faces this tax schedule, the firm no longer has an incentive to undertake inefficiently risky investments. Its investment choices will be socially efficient. It is useful to tax a financial institution producing a risk externality, just as it is useful to tax a firm producing a pollution externality. The purpose of the tax in both instances is to ensure that the firm pays the full costs—private and social—of its production decisions.

I emphasized that the pollution tax corrects the pollution externality without creating any new inefficiencies for the firm. The risk tax has the same property. Policymakers are considering a host of regulatory responses to the events of the past two and a half years, including higher capital requirements, leverage caps, and restrictions on incentive compensation. All of these potential changes are good in that they serve to lower the amount of risk-taking by financial institutions. However, they may also create new kinds of inefficiency for the targeted firms. For example, imposing new restrictions on incentive compensation may hamper a firm's ability to motivate its employees. In contrast, my proposed risk tax, like the pollution tax, corrects the risk externality without creating any new inefficiencies.

The proposed tax does require bank supervisors to calculate the expected present value of future bailout payments. These calculations are likely to be complex in a number of ways. Moreover, the calculations could well be controversial. Financial institutions that follow highly risky strategies get especially high profits when those strategies are working. Thus, supervisors would be required to levy

high risk taxes on exactly those institutions that appear extremely successful. For these reasons, it would be useful to develop an objective way to compute the required tax using market information.

Here's what I have in mind. Suppose that, for every relevant financial institution, the government issues a "rescue bond." The rescue bond pays a variable coupon equal to $1/1000$ of the transfers made from the taxpayer to the institution or its stakeholders. (I pick $1/1000$ out of the air; any fixed fraction will do.) Much of the time, this coupon will be zero, because bailouts aren't necessary and so the firm will not receive transfers. However, just like the institution's stakeholders, the owners of the rescue bond will occasionally receive a large payment. In a well-functioning market, the price of this bond is exactly equal to the $1/1000$ of the expected discounted value of the transfers to the firm and its stakeholders. Thus, the government should charge the financial firm a tax equal to 1000 times the price of the bond.

Notice that this approach could be used for a wide variety of financial institutions, including nonbanks. In principle, the government need not figure out in advance exactly which are systemically important and which are not. Instead, it could simply issue a rescue bond for every institution. Then, the market itself could reveal how systemically important each institution is through the price of its rescue bonds. Of course, markets are not always perfect, and it would be inappropriate to rely only on market measures to compute the appropriate taxes. However, the prices of rescue bonds would contain valuable information that should be an important input into the supervisory process.

As I mentioned at the beginning, Congress is in the process of considering changes to the financial regulatory system. In December, the House passed the Wall Street Reform and Consumer Protection Act. The Senate is currently deliberating the Restoring American Financial Stability Act. There is much to like in both pieces of legislation. However, neither piece of legislation incorporates the kind of risk tax that I have described to you. The Senate bill proposes no new taxes on financial

institutions, unless some fail. In that event, taxes could be levied on surviving large financial institutions, regardless of whether or not they had actually engaged in excessive risk-taking. The House bill has a new risk-based assessment on large banks and hedge funds. Such a risk-adjusted tax should have desirable incentive effects on the targeted firms. However, the tax will end once it has raised \$150 billion. This cap is problematic, because once the tax is ended, so too will its desirable incentive effects.

Why do the bills fail to include new levies of the kind that I propose? In my view, both bills significantly understate the extreme economic forces that lead to bailouts during financial crises. Indeed, the opening language of the Senate bill actually declares that it will end taxpayer bailouts. This objective is laudable. But it is not achievable – and thinking that it is can lead to poor choices about the structure of financial regulation.

To wrap up: Bailouts will inevitably happen during financial crises to prevent runs and systemic collapse. We need to structure financial regulation so as to limit the size and occurrence of these bailouts. How should we best design such regulations? The social distortion we face is that debt guarantees create a risk externality, because financial institutions do not bear the full costs of their investment choices. Financial regulation should be designed so as to best control that externality. As is true with any externality, the risk externality can be eliminated with a well-designed tax system. Figuring out the right tax may be complicated, but the task can be eased using appropriate information from financial markets.