

Stress Tests Disclosure: Theory, Practice, and New Perspectives*

Itay Goldstein and Yaron Leitner

1 Introduction

From the early days of stress testing banks in the wake of the financial crisis of 2008, disclosure has been a key issue of discussion. Although banking regulation and supervision were not very transparent over the years leading up to the financial crisis, the crisis created pressure not only to examine banks more closely and in a forward-looking way (which is done with stress tests) but also to reveal the results of the examination publicly. Initially, disclosure was constrained; for example, the first stress tests in Europe did not reveal much information. However, over time, an increasingly more frequent view among regulators and central bankers took hold, according to which they should follow the demand for more transparency and disclose much of the information revealed in the course of stress tests.

There are many controversies surrounding this view among practitioners, academics, and regulators. To give a flavor of the debate, in a *Wall Street Journal* article from March 5, 2012, former governor of the Federal Reserve Board Daniel Tarullo is quoted as saying that “the disclosure of stress test results allows investors and other counterparties to better understand the profiles of each institution.” On the other side, the Clearing House Association is quoted as saying that this disclosure “could have unanticipated and potentially unwarranted and negative consequences to covered companies and U.S. financial markets.” Such debates about the desirability of the disclosure of stress test results continued over the years.

Another issue, which has gained momentum recently and on which there is disagreement, is the transparency of the models that regulators use to conduct stress tests. Under the current policy, the Federal Reserve does not fully disclose its stress test models. Until 2018, the Federal Reserve only provided a broad description of its models (together with the publication of the results). Starting in 2019, the Federal Reserve moved to provide more details about its models, such as certain equations and key variables that influence the result. The Federal Reserve will also now illustrate how its models work on actual loans held by Comprehensive Capital Analysis and Review (CCAR) firms and on some hypothetical loan portfolios. In Europe, model disclosure has been more prevalent in recent years.

Economic theory suggests that public disclosure is not a panacea and that there are trade-offs that need to be considered when deciding how much information should be disclosed

* We thank Kathryn Judge, Til Schuermann, and other participants at the International Monetary Fund (IMF) “Rethinking Financial Stability: The FSAP at 20” conference for very helpful comments. We also thank Sophia Hua for research assistance.

publicly. Empirical evidence from various broad studies in economics and finance provide support for the possible downsides of disclosure. In an early policy paper commissioned by the Committee on Capital Markets Regulation after the introduction of stress tests, Goldstein and Sapra (2014) review various theories highlighting the costs and benefits of greater disclosure and link them to the context of stress tests. They conclude that although the benefits of disclosure are obvious, one has to be mindful of the costs and take steps to design disclosure policy to mitigate their impact.

As this chapter will suggest, we think that it is important for regulators to acknowledge that the disclosure of stress tests results and disclosure of the models behind the tests involve a trade-off and by no means can be viewed as a panacea. We believe it is important to think of a disclosure scheme that accounts for the costs and benefits and how they may evolve over time as the conditions in the economy and the financial system change. We caution against the current trend whereby increased stated policies for greater disclosure of results might just lead to weaker tests and where the resulting complacency in the financial system is not fully understood and accounted for.

The remainder of this chapter is organized as follows: In Section 2, we review the main costs and benefits of the disclosure of stress test results and their implications, as discussed by Goldstein and Sapra (2014). In Section 3, we review the implications of the unified self-contained frameworks, which were developed later on, dedicated to the analysis of disclosure of stress tests results. We focus on the framework of Goldstein and Leitner (2018) and other related works. In Section 4, we discuss the relation between implications for information disclosure and implications for stress test design and how the two are intertwined. In Section 5, we review the evolution of practice on the disclosure of stress test results over the years and caution against the link between greater disclosure and weaker tests. In Section 6, we discuss the disclosure of the models behind the tests. We conclude in Section 7.

2 Costs and Benefits of Disclosure

We review here the costs and benefits of disclosure discussed by Goldstein and Sapra (2014), their implications for stress tests, and some empirical evidence (see also Leitner, 2014). We also discuss some recent research that expands on the early work. We postpone discussion on recent work designed to build unified self-contained frameworks to analyze the disclosure of stress test results and the disclosure of the test models to Section 3 and Section 6.

A widely used argument in favor of disclosure of the test results is that it helps discipline banks. The idea is simple. If investors do not have information about the bank, the price of the bank's securities will not reflect the bank's true risk, and so banks may be induced to take more risk than is socially desired. However, if investors have information about the bank, prices will reflect the true risk, and risky banks will thus find it harder to raise money. This would induce banks to take less risk to begin with. More generally, the benefit of disclosure is that it allows market participants to make more informed decisions. This allows for better allocation of funds and also has a disciplinary effect on banks because they can be rewarded according to their actions. In a general corporate-finance context, various studies show that greater disclosure benefits firms by reducing information asymmetries and constraining managerial misbehavior (e.g., Leuz and Verrecchia, 2000; Greenstone et al., 2006).

Another benefit of disclosure is in allowing greater discipline and accountability of the regulators themselves. For example, disclosing information could make it easier for the regulator to commit to enforcement actions that are optimal *ex ante* but are suboptimal *ex post*. A recent empirical article by Kleymenova and Tomy (2019) provides evidence along these lines: when regulators are required to disclose enforcement actions, they are more likely to issue enforcement actions, as well as to rely on publicly observable signals to issue enforcement orders, suggesting a response to the increased public scrutiny of their actions.

However, the academic literature has also highlighted some problems with disclosure. First, public information might crowd out other sources of information. One form of crowding out is that the regulator may obtain less information from banks. A bank that has a temporary liquidity problem may benefit from government help, but it may be reluctant to reveal information to the regulator for fear of being stigmatized by the market. A bank may also conceal information for fear that it is revealed to competitors. Hence, greater disclosure of bank-specific information via stress tests might imply that banks will be less forthcoming in the information they provide to regulators, and this will limit the effectiveness of the stress tests to begin with. These issues are discussed by Prescott (2008) in the context of the disclosure of traditional bank supervisory information and by Leitner (2012) in the context of the disclosure of complex transactions that banks enter into.

Another form of crowding out of other information sources, which has been discussed much more widely in the general disclosure literature, is that the disclosure of public information can lead to a reduction in the private information that is produced and brought into prices by market investors. To the extent that market prices contain important information about banks' securities, which can help guide decisions made by creditors and others, this crowding out of market information is damaging and might even outweigh the direct positive effect of having more information disclosed via stress tests. A recent review of this literature is provided by Goldstein and Yang (2017). A general lesson is that the effect of information disclosure can vary significantly by the type of information that is being disclosed (Bond and Goldstein, 2015; Goldstein and Yang, 2019). It is always a good idea to disclose information about issues that decision-makers already know about and do not need to learn about from the price because this will encourage market participants to focus their attention on other things that decision-makers may still want to learn about, making prices more informative in a way that helps efficient decision-making. But revealing information, which might not be so precise, about things that the market may have a comparative advantage on might backfire. This will divert the attention of traders away from these dimensions and will make the price a less reliable source of information about them.

Second, as highlighted by Morris and Shin (2002), in the presence of strategic complementarities and coordination problems among bank counterparties, disclosing information about bank fundamentals might become a focal point, leading investors to put excessive weight on it and neglect their own private information. This is because public disclosure not only provides information on fundamentals but also helps investors guess what other investors will do. So even if the regulator is not much more well informed than private investors, these investors may end up overreacting to the regulator's announcement. This is a bad outcome from a social point of view, and it also undermines market discipline because it breaks the link between the bank's financial health and follow-up actions by investors

(e.g., whether they extend credit to the bank). The implication of this is that disclosure is beneficial only if the regulator's information is sufficiently precise.

Third, greater disclosure might encourage bank managers to focus on short-term goals that are not necessarily aligned with long-term value creation. This insight arises from models such as that by Gigler et al. (2014), who show that frequent disclosure can backfire by diverting the incentives of managers. The point goes back to Stein (1989), who argued that managers in public corporations with frequent disclosure and public share prices often try to achieve short-term measurable gains at the expense of long-term value.

Goldstein and Sapra (2014) suggest that these downsides of disclosure arise mostly when bank-specific information is being disclosed and not so much for disclosure of aggregate outcomes. However, these downsides can be remedied in the case of bank-specific information disclosure to alleviate the potential damages. Here are some remedies. First, when disclosing negative results that might trigger a coordination failure among market participants, it would help if the negative results are followed up by a corrective action. Second, it may be better not to release public information on dimensions where the market may have a comparative advantage in providing information. Third, when the information being disclosed is very precise, the negative effects are diminished. Fourth, banks' investment choices should be carefully monitored to ensure that disclosure does not distort their incentives.

An important issue to keep in mind when evaluating the costs and benefits of disclosure is that their relative importance likely depends on the general environment we are in and in particular on whether we are dealing with stress tests in normal times or during a crisis. For example, the disciplinary role of stress tests is more relevant during normal times, when the main purpose is to ensure that banks do not take too much risk. On the other hand, the concern about public information getting too much weight due to a coordination problem, as in Morris and Shin (2002), is likely more acute during a crisis, when the objective is to avoid panic. Although the various theories provide insights along these lines, more work is needed to get a better understanding of how the purpose of a stress test changes along the cycle and how this affects the balance of the costs and benefits of disclosure.

In general contexts (i.e., outside the context of stress tests), several empirical articles have provided evidence for the downsides of disclosure mentioned here. Hertzberg et al. (2011) show how, in the presence of strategic complementarities among lenders, the revelation of a signal publicly causes them to put excessive weight on it. Jayaraman and Wu (2019) show that greater disclosure acts to crowd out information from financial markets and that this acts to reduce the ability of managers to make informed investment decisions. Agarwal et al. (2018) and Kraft et al. (2018) show that greater and more frequent disclosure lead to myopia by corporate managers, who choose to focus on short-term goals rather than long-term innovation and growth.

Following these discussions, recent work has attempted to find evidence of the negative consequences of the disclosure of stress test results using data from US stress tests between 2009 and 2015. Studies by Flannery et al. (2017) and Fernandes et al. (2017) do not find evidence consistent with negative effects. In particular, they do not find evidence for a reduction in other sources of information production or distortion in banks' investment decisions. However, as they note, their sample is limited to a few years that did not exhibit any major events, so it is not clear how these results will generalize to broader and more heterogeneous samples. Moreover, it is difficult to isolate the particular effect of disclosure

from these samples, given that there is no clear counterfactual of what would have happened without disclosure.

Finally, a more recent argument that was brought forward against transparency in banks is developed by Dang et al. (2017). They argue that opacity on the asset side of banks is critical for the liquidity-provision role they perform on their liability side, allowing depositors to benefit from money-like assets. Clearly, the disclosure of stress test results would go against this objective of banks. Empirically, a recent article by Chen et al. (2019) provides evidence in support of this view. They show that banks that are more transparent and provide more precise information about upcoming losses suffer from a decreased ability to provide liquidity to depositors without relying on deposit insurance.

3 Unified Frameworks for Analyzing Stress Tests Information Disclosure

So far, we have discussed various costs and benefits of disclosure in isolation, as they were developed separately in the general disclosure literature and later tied to the context of stress tests. Following the introduction of stress tests, a few articles developed models to evaluate the question of stress test information disclosure using a self-contained framework in which costs and benefits arise in the model and unified policy prescriptions can be provided within the model. We describe a theory by Goldstein and Leitner (2018) that builds on the effects that disclosure will have on the risk-sharing arrangements attained by banks in interbank markets and financial markets more generally. We then describe a few other related works.

The idea that banks engage in risk-sharing arrangements among themselves, which is at the basis of the Goldstein and Leitner (2018) model, is strongly rooted in banking literature and practice (e.g., see the model of Allen and Gale, 2000). Going back to Hirshleifer (1971), it is known that more information might be harmful because it reduces risk-sharing opportunities for economic agents. Intuitively, if the realization of shocks is known, there is no room for mutually beneficial risk-sharing arrangements that ensure against these shocks. However, there are also forces that could make some disclosure desirable even when the objective is risk sharing. Indeed, during the 2008 financial crisis, interbank markets were not performing well, and there was a sense that some disclosure was necessary to prevent a breakdown in financial activity. Although market breakdowns are known to occur when market participants have asymmetric information (as in Akerlof, 1970), they can also occur when market participants share the same information, but the aggregate endowment is expected to be low (see Leitner, 2005), preventing risk-sharing arrangements from taking place. The model by Goldstein and Leitner (2018) captures these forces. We now provide a brief and simple description of it.

Suppose that the value of the bank's assets is $\theta + \varepsilon$, where θ and ε are independent random variables. The regulator learns θ during the stress test. We can think of θ as the regulator's forecast for the value of the bank's assets and of ε as additional noise that is unobservable at the time of the test and that has a mean of zero. Everyone knows the distribution of the random variables θ or ε . No one knows the realized value of ε , and only the regulator knows the realized value of θ . The regulator's objective is to minimize expected losses in the banking system. Suppose that the bank will suffer a loss if it ends up with value below 1 (e.g., because this will trigger a run on the bank). Banks can potentially guarantee that

their values do not fall below 1 by selling assets for a price that equals the expected value. This is an example of a risk-sharing arrangement. However, the ability to get into such an arrangement depends on the overall conditions in the system (e.g., whether expected valuations are above 1) and on the disclosure policy employed by the regulator.

In the basic case in which the realization of the fundamental θ for individual banks is known only to the regulator, Goldstein and Leitner (2018) show that the optimal-disclosure policy depends on investors' prior beliefs regarding θ . If $E(\theta) \geq 1$, the regulator does not need to disclose any information because banks will be able to raise at least \$1 by selling their assets. In this case, disclosure will be harmful because of the Hirshleifer effect: banks whose fundamental is revealed to be below 1 will not be able to participate in risk-sharing arrangements and will not be able to guarantee themselves against losses. If instead $E(\theta) < 1$, then without disclosure, banks will not be able to raise \$1. In this case, full disclosure will lead to a better outcome than no disclosure because at least some banks (those with $\theta \geq 1$) will be able to protect themselves. However, partial disclosure will lead to even better outcomes. Specifically, it is optimal to separate banks into two groups. The first group includes all the banks with $\theta \geq 1$, as well as some banks with $\theta < 1$, such that the group's average θ equals the critical level 1. The second group includes all the other banks. The benefit of this partial disclosure is that it allows some of the banks with values below the critical level ($\theta < 1$) to share risk with banks that have values above the critical level (i.e., banks with $\theta > 1$). In some cases, the regulator can implement this partial disclosure via a simple cutoff rule: the regulator sets a cutoff below 1 and reveals whether the bank's θ is above or below the cutoff.

These results are in the spirit of the Bayesian persuasion literature (Kamenica and Gentzkow, 2011). Essentially, the regulator wants to implement an outcome in which the bank can protect itself against bad outcomes by selling its assets for a sufficiently high price. But because the regulator cannot force investors to purchase assets for more than what they think the assets are worth, the regulator needs to "give up" on some banks with low θ so that on average the remaining banks are above the threshold. Linking to the Bayesian persuasion literature, then, one can think of stress test disclosure rules as an exercise in information design: How should regulators design the disclosure of information in order to achieve their goal (which, in this case, is minimizing expected losses in the banking system)?

The model of Goldstein and Leitner (2018) then goes on to consider more complicated environments, in which not only the regulator knows the fundamental θ of an individual bank, but also the bank knows. Moreover, as is realistically the case, the regulator cannot force a bank to sell its assets. This case could lead to adverse selection because a bank that knows that the value of its asset is high will not agree to sell for only \$1. A strong bank (with high θ) would like to protect itself against its value falling below 1—but not at any price. In other words, each bank has a reservation price, a minimum price at which it is willing to sell in order to obtain insurance. Because the market anticipates this adverse-selection problem, they will not be willing to purchase at a price of \$1, and the partial-disclosure scheme will not work.

Goldstein and Leitner (2018) show that in this case, the regulator needs to disclose more information to maximize its objective function. This is expected: banks of higher type would not agree to be pooled with a group that has a much lower average, and disclosure will thus act to separate banks into more groups, effectively providing more information. The most

realistic case, obtained in the model under some conditions, features optimal disclosure that separates banks into three groups based on their fundamental θ by defining two thresholds. Banks with low θ (below the lowest threshold) are excluded from risk sharing (i.e., they cannot protect themselves against value falling below the threshold), and it does not matter whether the regulator reveals or does not reveal their θ . Banks in the middle (between the two thresholds) are pooled together, and they all sell for the same price (same risk-sharing agreement). Because of the Hirshleifer effect, the regulator must not reveal the exact θ for banks in this group. The rest of the banks (those with θ above the second threshold) will sell for higher prices. For this group of banks, the purpose of disclosure is to ensure that they are willing to sell. If these banks are roughly the same, the regulator does not need to reveal their exact θ because even the highest bank in this group will be willing to sell at a price that reflects the group average. However, if banks in this group are very different from one another in terms of θ , the regulator will need to reveal more information because otherwise, a bank with a very high θ will prefer to take the risk that $\theta + \varepsilon < 1$, rather than get insurance by selling its asset as a very low price. There are even cases in which the only way to implement the optimal-disclosure policy is via full disclosure.

Summarizing the forces of the Goldstein and Leitner (2018) model, we see that the Hirshleifer effect is a force that pushes toward no disclosure, but if the regulator wants to guarantee that bank value is above some threshold, and the expected average value based on market priors is below the threshold, some disclosure is necessary. Adverse-selection problems will push for more disclosure. Adverse selection could be an issue if the bank knows the information that the regulator has and the regulator cannot force the bank to take corrective actions, such as raising capital, in order to protect itself. As these forces strengthen, full disclosure can emerge as the optimal outcome under this model.

The model has interesting implications for understanding how optimal disclosure may change with overall conditions in the financial system. Consider the case in which banks know their types, which are also revealed to the regulator in the course of the stress test. During booms, when there is no threat of bank runs and when every bank has sufficient capital (i.e., $\theta > 1$), the only relevant force is adverse selection, which could necessitate full disclosure. But if some banks are not in good condition (i.e., $\theta < 1$), but the overall state of the banking system is good (i.e., $E(\theta) > 1$), the Hirshleifer effect kicks in, which pushes toward less disclosure (i.e., some banks must be pooled together, but if there is sufficient adverse selection, the regulator may still need to disclose full information regarding the strongest banks). Finally, during a crisis, when the banking system as a whole is undercapitalized (which corresponds to $E(\theta) < 1$), some disclosure may also be necessary to separate some of the weak banks. When designing disclosure policies, regulators have to keep in mind how conditions might change and how this will affect how much they want to disclose.

Outside the model described so far, there can be other forces that push for more disclosure. For example, in a follow-up work, Orlov et al. (2018) develop a model in which the regulator learns from the stress test about systemic risks. In their model, the regulator wants to implement different outcomes across different banks, based on the expected performance of their risky assets in the bad state of the economy. In an optimal arrangement, banks with less fragile risky assets should hold more of them. The optimal-disclosure regime they derive out of this model also entails the separation of banks into different groups, guided by the intention to implement different outcomes for them.

Although this framework considers only the disclosure regime as a tool available to the regulator, governments can use other tools to alleviate fragility, such as directly intervening and injecting funds into banks. With this in mind, Faria-e-Castro et al. (2017) provide a model in which the optimal level of disclosure depends on the government's fiscal capacity, namely, its ability to inject money into banks and the cost of doing so. They show that governments that have more fiscal capacity can run more aggressive disclosure policies (i.e., they disclose more).

In their model, the benefit of disclosure is that it reduces adverse selection, allowing banks to raise capital to finance investment opportunities. The cost of disclosure is related to the Hirshleifer effect—it can induce runs on weak banks. Consistent with our previous discussion, the optimal amount of disclosure depends on the relative magnitude of each of these two forces. Now suppose the government can prevent runs by providing deposit insurance, but to finance this type of intervention, it will need to raise taxes later on, which is costly from a social point of view. Faria-e-Castro et al. (2017) show that if this cost is lower (i.e., if the government has more fiscal capacity), it is optimal to disclose more information. The intuition is simple. If the cost for the government to prevent runs is low, the government does not need to worry about runs, so the dominant force in its disclosure policy is reducing adverse selection.

This result can help explain the striking difference between the stress tests implemented in the United States and in Europe following the 2008–2009 financial crisis. Indeed, the first round of stress tests in the United States was considered to be very revealing, and for this it is always noted as a great success, whereas in Europe, the initial stress tests were weak and did not provide much information. The explanation here is that Europe could not afford to have highly informative tests because of its lack of fiscal capacity to intervene following bad results.

In Chapter 12 of this volume, Judge raises a similar point. Regulators may be hesitant in producing information unless they have the tools in hand to contain the fallout their findings might trigger. As a policy implication, she discusses the importance of designing safety nets, but she also discusses postcrisis reforms that scaled back the authority of financial regulators in the United States to provide crisis-time support. The Goldstein and Leitner (2018) model suggests that in some cases, a well-designed disclosure policy could supplement and possibly reduce the cost of having such safety nets because the central bank could rely not only on its own ability to inject funds into banks but also on cross-subsidies across banks that are pooled together. Reforms that weaken the regulator's authority to support banks during a crisis could potentially make disclosure policy even more relevant.

Finally, the frameworks just described start from the premise that regulators can disclose information according to some rules that are set in advance. This raises questions about commitment ability because *ex post*, the regulator may want to deviate from these rules. For example, in Goldstein and Leitner (2018), when disclosure is according to a cutoff rule, the regulator would like to say that θ is above the cutoff, even when it is below. This reflects the general idea that the regulator may want to hide bad news from the market. In the previous example, this is news about the health of an individual bank, but it could also be news about the health of the overall financial system: if investors did not know $E(\theta)$, the regulator might want to say that $E(\theta)$ is above the threshold, even when it is below. If investors expect this to happen, the stress test would lose its credibility. This issue is discussed by Bouvard et al. (2015).

4 Stress Test Design versus Information Disclosure

The Goldstein and Leitner (2018) article and the related literature discussed in the previous section are framed as analyses of optimal-disclosure schemes, by which regulators decide *ex ante* what rules to follow in presenting the information *ex post*. However, it is important to note that similar results can be obtained when regulators are committed to disclosing everything but have flexibility *ex ante* in how to design the test. In fact, the problem of stress test design is equivalent to the problem of information-disclosure design.

To see this point, suppose that according to the Goldstein and Leitner (2018) model, the optimal-disclosure rule entails no disclosure at all; that is, the circumstances are such that it is optimal not to reveal any information about the realization of the bank's type. For example, this would be the case in the model when the average fundamental of banks $E(\theta)$ is above 1 and the type is only revealed to the regulator and not to the bank itself. Analyzing a model of information design would suggest that after banks' types are revealed in the test, it would be optimal to assign all banks the same score, such that no information about their types is disclosed publicly. What we argue is that the same outcome could be achieved by designing a very weak test that does not identify meaningful differences between banks thus, even if the results of the test are fully revealed to the public, these results do not really differentiate among banks, and effectively, there is no disclosure.

Similarly, consider the case in which the optimal-disclosure rule in the model is to disclose everything, that is, to provide as much information as possible differentiating across the banks. This would be the case in the model if there are big differences across banks and they are aware of their types. Thinking about stress test design, such a solution can be obtained by designing a very strong test that will go to the root of the risks banks face, enabling sharp differentiation. In similar spirit, any disclosure rule prescribed by the model that features some form of partial disclosure (this is, of course, the most common outcome) can be achieved in a model of stress test design by planning a test that will highlight exactly the dimensions that are optimal to be revealed publicly.

Thinking about the problem through the lens of stress test design rather than information-disclosure design is also a way to address commitment problems of the type discussed previously. With information disclosure, one can ask whether the government has the ability to commit to follow *ex post* the *ex ante* optimal disclosure rule. A related concern is that political and legal constraints would make it difficult to lump together banks of different types *ex post* just to adhere to the *ex ante* optimal rule. Such concerns do not arise when we think about the problem as a problem of stress test design. The decision on stress test design is made *ex ante*, and there is no need to commit to it. Once the information is discovered during the stress test, it will be revealed publicly, but how informative it is will depend on the design decision that was made *ex ante*.

Overall, we can see here that the disclosure question cannot be fully separated from the question of the design of the stress tests. The two are intertwined. Ultimately, we care about what market participants will learn, and there are different ways to control this: either through disclosure or through the power of the test. This is an important observation, and we will come back to it when discussing the current practice in regulatory circles on stress test disclosure issues.

5 Current Thinking about Stress Test Information Disclosure in Regulatory Circles

Initially, following the introduction of stress tests, the views about the disclosure of their results were rather mixed among regulators; however, it seems that over the years, a growing consensus emerged that increased transparency of stress test results is desirable. Statements from policymakers reveal their belief in and commitment to greater disclosure, and the amount of information being revealed in both the United States and Europe has clearly increased.

So what kind of information is being provided following stress tests these days? In the United States, the Federal Reserve discloses various capital ratios and other accounting numbers, such as projected losses, revenue, and net income before taxes. The Federal Reserve also discloses loan-loss rates broken down by type of loan for each bank holding company (BHC) participating in the stress tests for both the adverse and severely adverse scenarios. Also disclosed are the Federal Reserve's decisions on capital plans for each BHC and the basis for the decisions. The European Central Bank's disclosure for European banks undergoing stress tests is much more detailed than that of the Federal Reserve. It is also becoming gradually more transparent over time, reaching 17,200 data points disclosed per bank in 2018. This is a striking development, given that in the first stress tests conducted in Europe, the results were kept rather confidential, and only aggregate numbers were provided. The information in Europe is also disclosed in a more user-friendly way than in the United States to help it be absorbed quickly into the financial market. In contrast, the Bank of England provides much less information for the stress tests it conducts.

The general view for greater transparency is clearly conveyed by the Federal Reserve and the European Central Bank in various policy papers and speeches. For example, for the Federal Reserve, the general philosophy is described in a supervisory staff report by Clark and Ryu (2013), who write: "Given a widely held view among supervisors and most third-party observers that the public disclosure of stress testing results enhances available information and supports market discipline, it will continue and it is perhaps even likely to be expanded over time." They then go on to dismiss several of the concerns raised in the academic literature and reviewed previously regarding the possible downsides of a greater-disclosure regime.

Although we agree that disclosure clearly has its benefits, as discussed earlier, we also think it is unrealistic to view disclosure as a panacea and dismiss the possible downsides. Ultimately, there is a trade-off, and theoretical frameworks help clarify the forces and provide policy prescriptions for given environments. For example, it is hard to ignore the fact that meaningful and precise disclosure will, in some cases, act to isolate banks with bad shocks from the market and disrupt the ability of the market to provide risk-sharing opportunities across banks. The question is whether regulators are prepared to accept such outcomes.

It should be noted that regulators have shown concern about such issues in various situations over the years. For example, one of the first uses of the Troubled Asset Relief Program (TARP) funds was providing capital to nine major financial institutions as part of the Capital Purchase Program, a program designed to infuse capital into healthy banks. During the audit,

former Federal Reserve chair Ben Bernanke told the special inspector general for TARP that “there were differences in the nine banks in terms of strength and weakness, but that the selection was generalized in order to avoid stigmatizing any one bank as being a weak bank and creating panic” (Office of the Special Inspector General for the Troubled Asset Relief Program [SIGTARP], 2009). So, the question is whether the Federal Reserve or other policy institutions are prepared to ignore these kinds of considerations when they arise in the future and bear the costs that full disclosure entails.

Another important question to ask when evaluating the current regime of greater transparency is whether the stated goals and public statements about increased transparency are indicative of much valuable information being revealed. Consider the discussion in the previous section on stress test design versus information disclosure. Recall that one way to achieve less actual transparency is to design very weak and uninformative tests and then provide very detailed disclosure of their results, so effectively, no critical information is being revealed. Is it possible that the Federal Reserve, while committing itself to disclosing the results of the tests it conducts, might choose to weaken the tests, such that the negative consequences of greater disclosure are not incurred?

By now, there is some empirical research on the actual informativeness (as opposed to the stated transparency) of stress tests; albeit most of it still comes from earlier years. Although it is difficult to provide conclusive evidence on the informativeness of the tests, these studies do indicate that stress tests provided new valuable information to markets. Flannery et al. (2017) analyze stress tests in the United States between 2009 and 2015. They measure the market response to stress test disclosure using average absolute cumulative abnormal returns (CAR), to avoid positive and negative effects canceling out, and using abnormal trading volume. Absolute CAR is statistically significant around most disclosure dates. Moreover, they show a spillover effect of disclosure from tested banks to untested bank. Abnormal trading volume also spikes around disclosure. They conclude that valuable information is being provided. Similar results are obtained by Fernandes et al. (2017) on a similar data set. They provide evidence that information asymmetry (as proxied by bid–ask spread) increases with the announcement of a stress test and decreases following the release of the results. Importantly, they show that the market response to stress tests disclosures is attenuated after the initial rounds. This raises the possibility that stress tests are not revealing as much information more recently.

Related studies have been conducted on European data, especially following the 2011 European stress test, in which there was a jump in the scale of disclosure (including up to 3,400 data points for each of the 90 participating banks). Petrella and Resti (2013) find that the test result is not fully anticipated by the market, and hence stress tests play an active role in mitigating bank opacity. Ellahie (2015) shows that the detailed data on exposure to credit and sovereign risks disclosed in the 2010 and 2011 EU bank stress tests have significant predictive power for the cross-section of bank equity and credit returns (change in credit default swap [CDS] spread). High sovereign risk exposure implies a subsequent decline in equity return and a widening of the CDS spread.

The concern, however, is that more recently, stress tests have become much less informative, especially in the United States. Inspecting the latest rounds of stress tests in the United States, it is striking that all participating banks are predicted to do so well under unbelievably bad scenarios for the economy. Is this a reflection of the outstanding resilience of the banking

sector or just an indication of weak tests being performed without revealing much critical information? In Chapter 9 of this volume, Sarin and Summers provide a detailed discussion, with references to other sources, of the flaws in the current system of stress tests in the United States. In particular, they highlight that market measures of banks' risks and fragility point to a much gloomier picture than the one provided by the regulatory stress tests. Hence, they call for an overhaul of the current methodologies and for greater reliance on market information when conducting stress tests. Glasserman and Tangirala (2016) provide related evidence in their analysis of the evolution of the information content from stress tests. They argue that the results have become predictable, with the information content declining over time.

Overall, the issue to pay attention to is that commitment to disclosure might come together with weaker tests. Although weak tests could be part of an optimal design, it would be troubling if the stress tests were claiming to identify all risks when, in fact, they were not designed to do so. First, current regulations are based to a large extent on stress test results, so capital disbursements and other measures loosening the positions of banks are triggered once stress test results suggest that risks are low. Second, although market information still seems to depict a different picture than the one portrayed by stress tests, there is a risk that the disclosure provided by the stress tests is affecting, or will affect in the future, the kind of information provided by the financial market and will make the whole system much more complacent. Third, weak tests could reduce the credibility of future tests in the next crisis because the market may lose confidence in the regulator's ability to assess risk. This issue is discussed by Judge in Chapter 12 of this volume. It is also related to the theoretical model by Morrison and White (2013), where crises spread across banks as a result of the common inference about the quality of the regulatory regime.

6 Disclosing the Model

As noted earlier, an important issue related to the disclosure of stress test results is the transparency about the models used by regulators to conduct the stress tests. Like in the case of disclosing the test results, we think there are costs and benefits in disclosing details about the underlying model, and one needs to have a framework in mind to evaluate the trade-off. It is also important to make a distinction between revealing the models to the public versus revealing them to the institution being tested (i.e., the bank).

In terms of disclosing models to the public, the benefits that come to mind are the following. First, disclosure enhances the credibility of the stress test by providing the public with information on the fundamental soundness of the models and their alignment with best modeling practices. Second, it helps the public understand and interpret the results of the stress test. Third, disclosure facilitates comments on the models from the public, including academic experts. These comments could lead to improvements, particularly in the data most useful to understanding the risks of particular loan types.

A possible downside of disclosing models to the public is that it could make it hard to maintain flexibility. The regulator will need to explain or consult with the public every time it wants to make changes in its models. The regulator may be afraid to experiment with new models that might contain mistakes, and this could lead to a situation in which the regulator sticks with old models that are "correct" but do not apply to the new environment. Frame

et al. (2015) suggest that this force might have played an important role in the failure of the supervisory stress tests conducted by the Office of Federal Housing Enterprise Oversight (OFHEO) for Fannie Mae and Freddie Mac. This is an interesting experience to learn from because this is a previous attempt in the United States to tie stress test results to capital requirements, which clearly did not go well.

When it comes to disclosing models to the banks, the biggest disadvantage is that when the regulator reveals the model, banks can game the test. For example, they can use the models to make modifications to their businesses that change the results of the stress test without changing the risks they face. In the presence of such behavior, the stress test could give a misleading picture of the actual vulnerabilities faced by banks. This might become a game focused on passing the test rather than truly increasing resilience. In addition, if a specific asset is perceived as more advantageous to hold based on the regulator's model, disclosing the model could lead to correlated asset holdings across banks, further increasing systemic fragility. Finally, model disclosure could induce banks to use models similar to the regulator's model, rather than building their own capacity to identify, measure, and manage risk. As argued by Schuermann (2014), this "model monoculture," in which all banks have similar internal stress testing models, might miss key idiosyncratic risks faced by the banks.

On the positive side, revealing the model to the bank can help it understand the capital implications of changes to its business activities, such as acquiring or selling a portfolio of assets. Banks have constantly complained in the past about model secrecy, claiming that even their best efforts to prepare for a test could result in unexpected and costly failure, which leads to lack of trust and feeling of arbitrariness. This might have negative real implications, given evidence that regulatory uncertainty causes banks to reduce lending (see Gissler et al., 2016).

As noted earlier, it is useful to have a unified framework for assessing the trade-off involved in the disclosure of the underlying models. Leitner and Williams (2019) recently provide such a framework to evaluate the issue of sharing the models with the bank. In their setting, the bank has better capacity than the regulator to identify and measure risk, but there is a conflict of interest between the bank and the regulator: the bank wants to take more risk than is socially desirable. The outcome of this is that if the regulator discloses the stress test models, the bank games the test by overinvesting in assets for which it knows that the regulator's models underestimate the risk. Essentially, when the regulator reveals a model that shows that an asset is not very risky, it gives a "green light" for the bank to invest, so the bank may end up investing even when it knows that the asset is harmful to financial stability. Not revealing the regulator's model can mitigate this problem, but it opens the door to a new problem: the bank sometimes does not invest in risky assets even though it knows the assets are good from a social point of view. The reason is that although the bank knows that an investment is socially desirable, it may be worried that the regulator's model does not measure risks accurately, which could result in the bank failing the test. Revealing the test mitigates this problem. Hence, there is a trade-off. One implication of this is that revealing the test is preferred to not revealing it if the bank is too concerned about failing the test (e.g., if the bank's private cost of failing the test is sufficiently high). The bank's private cost of failing the test could represent the private cost for bank managers when the bank fails the test (e.g., the manager can get fired), the potential market reaction, the cost of altering the bank's capital plan, and so forth.

Leitner and Williams (2019) show that in some cases, the regulator can alleviate the bank's concern, thereby solving the underinvestment problem, by making the test easier. Specifically, the regulator can reduce the threshold for passing the test. If so, not revealing the test is preferred. However, this does not work if banks are very different from one another and the regulator must apply the same passing threshold to all banks. In this case, the regulator cannot calibrate the passing threshold to induce socially desirable investment by everyone because making the test easier to alleviate concerns by banks that are too cautious (i.e., those with a high private cost of failing the test) could induce the "reckless" banks (those with a low cost of failing the test) to take excessive risks. The implication of this is that if banks are very different from one another and the regulator must apply the same passing threshold to everyone, revealing the test is preferred to not revealing it.

Leitner and Williams (2019) also analyze the more complicated case in which the regulator can reveal only partial information (e.g., whether estimated losses from an asset, according to the regulator's model, are above or below some threshold). One of the results is that in some cases, some disclosure is optimal even if the regulator can set the passing threshold optimally and is not obliged to apply the same threshold to all banks.

This result reflects the general idea that in some cases, it is optimal to combine more than one regulatory tool to achieve a desired outcome. In Leitner and Williams (2019), the regulator wants to prevent gaming while still maintaining socially desirable investment, and under some circumstances, it is optimal to do so by combining partial disclosure with adjustment of the test difficulty (or capital requirements). Intuitively, each of these tools has costs and benefits, and a combination of the two tools helps minimize the total cost (i.e., maximize the total net benefit). Earlier, we saw that during crisis, when the regulator's objective is to prevent runs or ensure that each bank has a sufficient level of capital, the regulator can achieve this goal by combining safety nets and partial disclosure that pools together banks of different levels of stress.

Finally, a widely expressed concern is that disclosing the models could increase correlations in asset holdings among the banks subject to the stress tests. Leitner and Williams (2019) suggest that a similar effect could prevail if the regulator reveals the outcome of applying its models to hypothetical loan portfolios (e.g., under the new policy changes of the Federal Reserve from February 2019). The reason is that these portfolios could serve as a benchmark portfolio in which too many banks invest. In other words, a bank will overinvest in portfolios, for which there is less uncertainty as to how the regulator will measure risk, and it will underinvest in idiosyncratic investments, for which it is unclear how the regulator's models will work.

7 Conclusion

Stress tests have been proposed following the financial crisis as a tool to monitor bank risks in a forward-looking way. Unlike traditional regulatory monitoring and interventions in banks, which are typically kept confidential, the idea with stress tests is to provide more disclosure and transparency. This has become more prominent over the years because the amount of information being provided is now larger. In addition, regulators have taken steps to increase the transparency about the underlying models they are using.

We argue that disclosures of stress test results and of the underlying models involve trade-offs. We described various costs and benefits and also unified frameworks to evaluate

such trade-offs. These forces represent issues that regulators have been concerned about in various situations and are likely to continue to play a role in the future. Hence, it is important to acknowledge that disclosure is not a panacea but, rather, that the disclosure regime has to be thought of carefully, taking into account the various forces and how they change over time.

We caution that if regulators do not acknowledge the trade-offs with disclosure, then commitment to full disclosure could just lead to weak tests which is something that might start happening already. If the system ends up producing such weak tests, while portraying them as strong and informative and basing regulatory actions on them, we might get to a point of extreme complacency, which is dangerous for future financial stability.

References

- Agarwal, Vikas, Rahul Vashishtha, and Mohan Venkatachalam (2018), "Mutual fund transparency and corporate myopia," *Review of Financial Studies*, 31(5), 1966–2003.
- Akerlof, George A. (1970), "The market for 'lemons': Quality uncertainty and the market mechanism," *Quarterly Journal of Economics*, 84(3), 488–500.
- Allen, Franklin, and Douglas Gale (2000), "Financial contagion," *Journal of Political Economy*, 108(1), 1–33.
- Bond, Philip, and Itay Goldstein (2015), "Government intervention and information aggregation by prices," *Journal of Finance*, 70(6), 2777–2811.
- Bouvard, Matthieu, Pierre Chaigneau, and Adolfo De Motta (2015), "Transparency in the financial system: Rollover risk and crises," *Journal of Finance*, 70(4), 1805–1837.
- Chen, Qi, Itay Goldstein, Zeqiong Huang, and Rahul Vashishtha (2019), "Bank transparency and deposit flows," Working Paper, available at www.bde.es/f/webpi/SES/seminars/2019/Fich/sie20190710.pdf.
- Clark, Tim, and Lisa Ryu (2013), "CCAR and stress testing as complementary supervisory tools," Board of Governors of the Federal Reserve System Supervisory Staff Report.
- Dang, Tri Vi, Gary Gorton, Bengt Holmstrom, and Guillermo Ordóñez (2017), "Banks as secret keepers," *American Economic Review*, 107(4), 1005–1029.
- Ellahie, Atif (2013), "Information content of mandated bank stress test disclosures," Working Paper, available at <https://doi.org/10.2139/ssrn.2685919>.
- Faria-e-Castro, Miguel, Joseba Martinez, and Thomas Philippon (2017), "Runs versus lemons: Information disclosure and fiscal capacity," *Review of Economic Studies*, 84(4), 1683–1707.
- Fernandes, Marcelo, Deniz Igan, and Marcelo Pinheiro (2017), "March madness in Wall Street: (What) does the market learn from stress tests?" *Journal of Banking and Finance*, forthcoming.
- Flannery, Mark, Beverly Hirtle, and Anna Kovner (2017), "Evaluating the information in the Federal Reserve stress tests," *Journal of Financial Intermediation*, 29(C), 1–18.
- Frame, Scott, Kristopher Gerardi, and Paul Willen (2015), "The failure of supervisory stress testing: Fannie Mae, Freddie Mac, and OFHEO," Federal Reserve Bank of Boston Working Paper No. 15–4.
- Gigler, Frank, Chandra Kanodia, Haresh Sapra, and Raghu Venugopalan (2014), "How frequent financial reporting can cause managerial short-termism: An analysis of the costs and benefits of increasing reporting frequency," *Journal of Accounting Research*, 52(2), 357–387.
- Gissler, Stefan, Jeremy Oldfather, and Doriana Ruffino (2016), "Lending on hold: Regulatory uncertainty and bank lending standards," *Journal of Monetary Economics*, 81, 89–101.
- Glasserman, Paul, and Gowtham Tangirala (2016), "Are the Federal Reserve's stress test results predictable?" *Journal of Alternative Investments*, 18(4), 82–97.
- Goldstein, Itay, and Yaron Leitner (2018), "Stress tests and information disclosure," *Journal of Economic Theory*, 177(C), 34–69.
- Goldstein, Itay, and Haresh Sapra (2014), "Should banks' stress test results be disclosed? An analysis of the costs and benefits," *Foundations and Trends in Finance*, 8(1), 1–54.

- Goldstein, Itay, and Liyan Yang (2017), "Information disclosure in financial markets," *Annual Review of Financial Economics*, 9, 101–125.
- Goldstein, Itay, and Liyan Yang (2019), "Good disclosure, bad disclosure," *Journal of Financial Economics*, 131(1), 118–138.
- Greenstone, Michael, Paul Oyer, and Annette Vissing-Jorgensen (2006), "Mandated disclosure, stock returns, and the 1964 Securities Acts Amendments," *Quarterly Journal of Economics*, 121(2), 399–460.
- Hertzberg, Andrew, Jose Liberti, and Daniel Paravisini (2011), "Public information and coordination: Evidence from a credit registry expansion," *Journal of Finance*, 66(2), 379–412.
- Hirshleifer, Jack (1971), "The private and social value of information and the reward to inventive activity," *American Economic Review*, 61(4), 561–574.
- Jayaraman, Sudarshan, and Johana Wu (2019), "Is silence golden? Real effects of mandatory disclosure," *Review of Financial Studies*, forthcoming.
- Kamenica, Emir, and Matthew Gentzkow (2011), "Bayesian persuasion," *American Economic Review*, 101(6), 2590–2615.
- Kleymenova, Anya, and Rimmy Tomy (2019), "Observing enforcement: Evidence from banking," Chicago Booth Research Paper No. 19-05.
- Kraft, Arthur, Rahul Vashishtha, and Mohan Venkatachalam (2018), "Frequent financial reporting and managerial myopia," *The Accounting Review*, 93(2), 249–275.
- Leitner, Yaron (2005), "Financial networks: Contagion, commitment, and private sector bailouts," *Journal of Finance*, 60(6), 2925–2953.
- Leitner, Yaron (2012), "Inducing agents to report hidden trades: A theory of an intermediary," *Review of Finance*, 16(4), 1013–1042.
- Leitner, Yaron (2014), "Should regulators reveal information about banks?" *Federal Reserve Bank of Philadelphia Business Review*, Q3, 1–8.
- Leitner, Yaron, and Basil Williams (2019), "Model secrecy and stress tests," Working Paper, available at www.basilwilliams.org/leitner.williams.secrecy.pdf.
- Leuz, Christian, and Robert Verrecchia (2000), "The economic consequences of increased disclosure," *Journal of Accounting Research*, 38, 91–124.
- Morris, Stephen, and Hyun Song Shin (2002), "Social value of public information," *American Economic Review*, 92(5), 1521–1534.
- Morrison, Alan, and Lucy White (2013), "Reputational contagion and optimal regulatory forbearance," *Journal of Financial Economics*, 110(3), 642–658.
- Office of the Special Inspector General for the Troubled Asset Relief Program (2009, October), "Emergency capital injections provided to support the viability of Bank of America, other major banks, and the U.S. financial system," Report 10-001.
- Orlov, Dmitry, Andrzej Skrzypacz, and Pavel Zryumov (2018), "Design of macro-prudential stress tests," Working Paper, available at <https://econpapers.repec.org/paper/redsed018/913.htm>.
- Petrella, Giovanni, and Andrea Resti (2013), "Supervisors as information producers: Do stress tests reduce bank opaqueness?" *Journal of Banking and Finance*, 37(12), 5406–5420.
- Prescott, Edward Simpson (2008), "Should bank supervisors disclose information about their banks?" *Federal Reserve Bank of Richmond Economic Review*, 94(1), 1–16.
- Schuermann, Til (2014), "Stress testing banks," *International Journal of Forecasting*, 30(3), 717–728.
- Stein, Jeremy (1989), "Efficient capital markets, inefficient firms," *Journal of Finance*, 44(4), 1335–1350.