Designing Stress Scenarios*

Cecilia Parlatore† Thomas Philippon‡

May 20, 2021

Abstract

We develop a tractable framework to study the optimal design of stress scenarios. A risk-averse principal (e.g, a manager, a regulator) seeks to learn about the exposures of a group of agents (e.g., traders, banks) to a set of risk factors. The principal asks the agents to report their outcomes (e.g., credit losses) under a variety of scenarios that she designs. She can then take remedial actions (e.g., mandate reductions in risk exposures). The principal's program has two parts. For a given set of scenarios, we show how to apply a Kalman filter to solve the learning problem. The optimal design is then a function of what she wants to learn and how she intends to intervene if she uncovers excessive exposures. The choice of optimal scenarios depends on the principal's priors about risk exposures, the cost of ex-post interventions, and the potential correlation of exposures across and within agents.

JEL Classification: G2, D82, D83

Keywords: stress test, information, bank regulation, filtering, learning

*We thank our discussants Itay Goldstein, Florian Heider, Til Schuermann, and Jing Zeng, as well as Mitchel Berlin, Thomas Eisenbach, Piero Gottardi, Anna Kovner, Ben Lester, Igor Livshits, Tony Saunders, Chester Spatt, Pierre-Olivier Weill, and Basil Williams for their comments. We would also like to thank seminar participants at the NBER Summer Institute, AFA, EFA, SED, NYU, FRB of New York, FRB of Philadelphia, FRB of Boston, University of Wisconsin, Boston College, and the Stress Testing Conference. Abhishek Bhardwaj, Ki Beom Lee, and Luke Min provided excellent research assistance.

†New York University, Stern School of Business, CEPR and NBER. Email: cparlato@stern.nyu.edu

‡New York University, Stern School of Business, CEPR and NBER. Email: tphilipp@stern.nyu.edu

1 Introduction

Stress tests are ubiquitous in risk management and financial supervision. Risk officers use stress tests to set and monitor risk limits within their organizations, and financial regulators around the world use stress tests to assess the health of financial institutions. To give just a few examples: financial firms use stress tests to complement their statistical risk management tools (e.g., Value at Risk); asset managers stress test their portfolios; trading venues stress tests their counter-party exposures; regulators mandate large scale stress tests for banks and insurance companies and use the results to enforce capital requirements and validate dividend policies.[1]

Despite the growing importance of stress testing and the amount of resources devoted to them, there is little theoretical guidance on exactly how one should design the stress scenarios. A theoretical literature has focused on the trade-offs involved in the disclosure of supervisory information (see Goldstein and Sapra, 2014 for a review), which range from concerns about the reputation of the regulator (Shapiro and Skeie, 2015) to the importance of having a fiscal backstop (Faria-e-Castro et al., 2017). Though these papers provide insights as to what to do with the results of asset quality reviews, they do not model *stress testing.* They are silent about the design of forward-looking scenarios, which are the hallmarks of stress testing.

The goal of our paper is to start filling this void. Stress tests are used for risk management. Risk management always contains (at least) two parts: risk discovery (learning) and risk mitigation (intervention). Stress tests are thus always embedded in a risk management framework. Stress tests belong to the risk discovery phase but we cannot think about the design of the test without analyzing the remedial actions that will be taken based on the results of the test. We therefore model the risk discovery stage and the risk mitigation stage. In the particular case of stress tests used for capital adequacy in banking there is often a quasi-mechanical link between the test and the remedial actions since banks are required to maintain a pre-specified level of capital buffer across all scenarios. We do not assume such a mechanical mapping. We assume instead that regulators choose optimal actions conditional on the results of the test. This gives them complete freedom to design the most informative scenarios.[2]

[1] Central banks in the United States, Europe, England, Brazil, Chile, Singapore, China, Australia, and New Zealand, as well as the International Monetary Fund in Japan, have recently used stress tests to evaluate the banking sector's solvency and guide banking regulation.

[2] Our results shed light on risk exposures and capital ratios, but we do not specify a mechanical link between passing the test and having a particular level of equity. To do so one would need to take a stand on many

We model stress testing as a learning mechanism. We consider a principal and a potentially large number of agents. The agents can be traders within a financial firm, or they can be financial firms within a financial system. The principal can be a regulator designing supervisory tests, or a risk officer running an internal stress test. For concreteness we will use the supervisory stress testing example in much of the paper. Banks are exposed to a set of risk factors, but their exposures to aggregate factors are unknown to the regulator. By exposure we mean the relevant elasticity that determines the loss of a position under a given scenario. An exposure is therefore not the same as the book or market value of a position. Banks and regulators usually agree on the nominal size of positions and on the market value of liquid portfolios. They can disagree about the value of illiquid positions, and in all cases, liquid or not, the impact of a scenario on the loss on that position needs to be estimated. What we call "exposure" combines the position (measured with near certainty) with its value under stress scenarios (computed with error). The regulator is risk averse and worries about the financial system experiencing large losses in some states of the world. The regulator then designs a set of scenarios and asks the banks to report what their losses would be under these scenarios. From their responses, the regulator learns about the underlying exposures of the banks. Based on this information, the regulator can decide to intervene, i.e., she can ask a set of banks to reduce their exposures to some factors.

Our main insight comes from writing the learning problem as a Kalman filter. The filter gives us a mapping from prior beliefs and test results into posterior beliefs. The precision of the mapping depends on the scenarios in the stress test. We can then formulate the regulator's problem as an information acquisition problem in which the regulator chooses the precision of her signals. By

features of financial regulations that are not central to our analysis. For instance, imagine that a bank needs the same level of ex-ante equity to satisfy a 9% capital requirement after scenario 1 or a 7% requirement after scenario 2 (presumably because scenario 2 embodies a higher degree of stress). As far as ex-ante capital adequacy is concerned, these two regulations are equivalent. The law sometimes mandates one of these numbers, in which case our model can shed light on the other, but in general the "level" of the ratio and the "degree" of the stress are not independent and a model is likely to pin down a combination as opposed to a particular value for each. A scenario used for "pass/fail" also needs to be plausible in a way that a learning scenario does not have to be. When stress test results are mechanically linked to capital requirements the choice of scenarios can be used to increase the effective requirement (if one assumes that, for some reason, the baseline requirement is too weak) or to implement counter-cyclical ratios (keeping the level of stress constant as the economy improves leads to larger assumed shocks). Finally, passing the stress test means that regulators deem the institution safe and sound even in the stress scenario, and thus that lender-of-last-resort policies would be appropriate.

explicitly considering the structure of the signals generated by the stress tests, we can map the feasible set of precision choices to the primitive parameters of the model, such as the priors of the regulator regarding the banks' exposures. If, for instance, the regulator is worried about a particular risk factor, we can derive the stress test that maximizes learning about the exposures to this risk factor.

Will the regulator focus a particular risk factor or will she try and learn about several factors at the same time? We show that the answer depends on her prior beliefs about the banks' risk exposures and on the marginal cost of intervention. The regulator can choose to mandate a broad risk reduction. This is likely to involve a lot of unnecessary changes and disruptions, but it does not require much information. If the cost of intervention is high, this strategy is not efficient and the regulator will want to learn in order to avoid unnecessary interventions. She can learn by choosing a more extreme scenario along a particular risk dimension, but extreme scenarios lead to less precise answers. The reduction in overall information quality depends on the prior distribution of the risk exposures. Whether or not there is specialization in learning depends therefore on the intervention cost function and on the regulator's prior beliefs.

More generally, the costs of intervention and the prior beliefs of the regulator are central in determining the optimal scenario design. The effect of intervention costs on the optimal scenario is not monotone. On the one hand, a higher intervention cost along a particular risk dimension makes accurate interventions more appealing. Through this channel, learning about the exposures to factors with higher intervention costs is more valuable. On the other hand, learning about the risk exposures is valuable partly because it allows the regulator to intervene to reduce the risk exposures. This mechanism makes the regulator's intervention less responsive to the information in the stress test when intervention costs are higher, and hence, make learning less valuable. When the intervention cost is low enough, the first effect dominates and higher intervention costs along a risk dimension increase the weight of that risk factor in the optimal stress scenario. As the intervention cost increases, the second effect becomes stronger. When the intervention cost is large enough, the second effect dominates and the weight of the risk factor in the optimal stress scenario decreases with the cost to intervene along that risk dimension.

The regulator's expected risk exposure to a risk factor also has opposing effects on the optimal stress scenario. First, a higher expected risk exposure increases the value of intervening along that risk dimension and makes accurate interventions more valuable. This effect pushes the regulator towards more learning (and stressing) about the exposure to risk factors with higher expected

risk exposures. Second, the responsiveness of the regulator's belief, and hence of her intervention policy, to the information contained in the stress test results depends on the level of the prior mean exposure. When the prior mean exposure to a particular factor is low (high), the regulator's beliefs about it are more (less) sensitive to new information.This effect makes learning about factors with higher expected risk exposures less valuable and dominates the first effect when the expected risk exposure is high. Hence, the weight of a factor in the stress scenario is hump-shaped with respect to its expected risk exposure. In the limit, when the expected risk exposure to a factor is high enough, the regulator's intervention policy becomes almost independent of the signal about it and the regulator chooses not to stress that risk factor at all.

The regulator's prior uncertainty, either about the risk exposures or about the risk factors themselves, also shapes the optimal stress scenario. Higher uncertainty about the risk exposure to a particular factor increase the regulator's willingness to intervene along that dimension. Therefore, the regulator finds it optimal to learn more about exposures to risk factors about which she is more uncertain about. On the other hand, higher uncertainty about the risk factor itself makes the regulator less willing to learn about the risk exposures to it. High uncertainty about a risk factor makes the regulator's intervention policy less sensitive to new information about the exposure to that factor. In this case, the regulator puts less weight in the stress scenario on risk factors about which she is more uncertain.

Our model also sheds light on the role of factor with correlated risk exposures, within or across banks. When the exposures to two factors within a bank are correlated, learning about the exposure to one factor also provides information about the exposure to the other factor. Hence, the regulator stresses more the factors with correlated exposures and may focus only on these factors if the correlation is high enough. However, due to the convexity of the information acquisition cost, the regulator's specialization is usually incomplete and she tends to put weight on all factors. Similarly, the regulator optimally puts more weight on systemic factors in stress scenarios. When different banks have correlated exposures to a (systemic) factor, learning about one bank's exposure to this factor provides information about the other banks' exposures. The higher the expected correlation between the exposures to the systemic factor, the more precise the information about the banks' exposure to it and the less costly it becomes for the regulator to learn about it. Hence, the regulator optimally stresses systemic factors more and may choose to specialize and stress only these factors when they are systemic enough (factors with exposures that are sufficiently correlated).

For expositional purposes we develop our insights in the context of a small number of banks and few independent macroeconomic factors. However, our framework can accommodate correlated macro factors, including non-linear combinations of factors, arbitrary and general correlation structures for risk exposures across and within banks, and different preferences for the regulator. Moreover, the linear structure allows it to be scalable and very easy to extend to a large number of banks, scenarios and risk factors and makes it easy to use as a building block in more complex stress test models.

Literature Review

Most of the literature on stress tests focuses on banking. Several recent papers study specifically the trade-offs involved in disclosing stress test results. Goldstein and Leitner (2018) focus on the Hirshleifer (1971) effect: revealing too much information destroys risk-sharing opportunities between risk neutral investors and (effectively) risk averse bankers. These risk-sharing arrangements also play an important role in Allen and Gale (2000). Shapiro and Skeie (2015) study the reputation concerns of a regulator when there is a trade-off between moral hazard and runs. Faria-e-Castro et al. (2017) study a model of optimal disclosure where the government trades off Lemon market costs with bank run costs, and show that a fiscal backstop allows government to run more informative stress tests. Schuermann (2012) analyzes the design and governance (scenario design, models and projection, and disclosure) for more effective stress test exercises. Schuermann (2016) particularly determines how stress testing in crisis times can be adapted to normal times in order to insure adequate lending capacity and other key financial services. Orlov et al. (2017) look at the optimal disclosure policy when it is jointly determined with capital requirements, while Gick and Pausch (2014) and Inostroza and Pavan (2017) do so in the context of Bayesian persuasion. As argued by Goldstein and Leitner (2020), stress test design and disclosure policy are connected. We complement this strand of papers by focusing on the stress scenario design.

While most of the existing literature on stress testing, theoretical and empirical, analyzes the disclosure of stress test results, Leitner and Williams (2018) focus on the disclosure of the regulator's risk modeling. They examine the trade-offs involved in disclosing the model the regulator uses to perform the stress test to banks. However, none of these papers consider the optimal scenario design, which is the focus of our paper.

Most empirical papers on stress tests focus on the information content at the time of disclosure, using an event study methodology to determine whether stress tests provide valuable information to investors. Petrella and Resti (2013) assess the impact of the 2011 European stress test exercise. For the 51 banks with publicly traded equity, they find that the publication of the detailed results provided valuable information to market participants. Similarly, Donald et al. (2014) evaluate the 2009 U.S. stress test conducted on 19 bank holding companies and find significant abnormal stock returns for banks with capital shortfalls. Candelon and Sy (2015), Bird et al. (2015), and Fernandes et al. (2015) also find significant average cumulative abnormal returns for stress tested BHCs around many of the stress test disclosure dates. Flannery et al. (2016) find that U.S. stress tests contain significant new information about assessed BHCs. Using a sample of large banks with publicly traded equity, the authors find significant average abnormal returns around many of the stress test disclosures dates. They also find that stress tests provide relatively more information about riskier and more highly leveraged bank holding companies. Glasserman and Tangirala (2016) evaluate one aspect of the relevance of scenario choices. They show that the results of U.S. stress tests are somewhat predictable, in the sense that rankings according to projected stress losses in 2013 and 2014 are correlated. Similarly, the rankings across scenarios in a given year are also correlated. They argue that regulators should experiment with more diverse scenarios, so that it is not always the same banks that project the higher losses. Acharya et al. (2014) compare the capital shortfalls from the stress tests with the capital shortfalls predicted using the systemic risk model of Acharya et al. (2016) based on equity market data. Camara et al. (2016) study the quality of the 2014 EBA stress tests using the actual micro data from the tests.

Finally, our paper is related to the large theoretical literature on information acquisition following Verrecchia (1982), Kyle (1989), and especially Van Nieuwerburgh and Veldkamp (2010). In this class of models, the cost of acquiring information pins down the set of feasible precisions and determines whether there are signals are complement or substitutes. Vives (2008) and Veldkamp (2009) provide a comprehensive review of this literature. These papers take the information processing constraint on the signal precisions as given. In contrast, our paper focuses on the design of the signals that the regulator receives and endogeneizes the information processing constraint.

The rest of the paper is organized as follows. Section 2 describes the environment. Section 3 describes how the regulator learns from stress test. Sections 4 and 5 characterize the optimal intervention policy and the optimal stress scenarios, respectively. Section 6 discusses the practical

implications of our analysis. Section 7 concludes.

2 Environment

We consider the problem of a principal who wants to learn about the risk exposures of a set of agents in order to take remedial actions. The model has several natural interpretations. The principal could be a chief risk officer and the agents could be traders in her firm. The remedial actions could be hedging or downsizing the traders' positions. Alternatively, the principal could be a regulator and the agents could be a set of banks. The remedial actions could be hedging, reducing new deal flows, selling non-performing assets, or raising capital.

To be concrete we use the regulator/banks metaphor when describing the model. The regulator elicits information from the banks in the form of stress tests. In our model, a stress test is a technology used by regulators to ask questions about profits and losses under hypothetical scenarios. The banks cannot evade the questions and have to answer to the best of their abilities. Banks in our model can only lie by omission: they do not have to volunteer information, but they have to provide estimates of their losses under various scenarios.

2.1 Banks and Risks

There is one regulator overseeing N banks indexed by $i \in [1, .., N]$ exposed to systematic and idiosyncratic risks. The macro-economy is described by a vector of J systematic factors. We denote by s_j the value of factor j. The macroeconomic state of the economy is

$$\mathrm{s} = \begin{bmatrix} s_1 \\ \vdots \\ s_J \end{bmatrix}.$$

The risks of bank i are captured by a vector of J exposures

$$\mathrm{x}_i = \begin{bmatrix} x_{i,1} \\ \vdots \\ x_{i,J} \end{bmatrix},$$

where $x_{i,j}$ represents the exposure of bank i to factor j. We use the term "exposure" to denote the relevant elasticity that determines losses under a given realization of the macroeconomic state.

An exposure is therefore not the same as the nominal value of a position. In many cases (e.g., a commercial loan) the size of the position is unambiguous but the impact of a realization of the macro state on the loss on that position needs to be estimated. What we call "exposure" combines the position (measured with near certainty) with its value in a particular macro state (computed with error).

The losses of bank i in state s are given by

$$y_i\left(\mathrm{s}\right) = \mathrm{s} \cdot \mathrm{x_i} + \eta_i = \sum_{j=1}^{J} x_{i,j} s_j + \eta_i, \tag{1}$$

where η_i is a random idiosyncratic (i.e., bank specific) shock. We normalize the baseline state to $s = 0$ and $\mathbb{E}\left[s\right] = 0$, so that the realizations of the macro state should be interpreted as deviations from the baseline. Our model has only one period so $y_i\left(\mathrm{s}\right)$ should be interpreted as the cumulative losses in state s. The net worth of bank i is then given by

$$w_i\left(\mathrm{s}\right) = \bar{w}_i - y_i\left(\mathrm{s}\right), \tag{2}$$

where $\bar{w}_i$ is the mean level of net worth. Given Equation (1) and Equation (2), the aggregate net worth of the banking system is

$$W\left(\mathrm{s}\right) = \sum_{i=1}^{N} w_i = \bar{W} - \mathrm{s} \cdot \bar{\mathrm{x}} + \sum_{i=1}^{N} \eta_i, \tag{3}$$

where $\overline{W}$ and $\bar{\mathrm{x}}$ are the sum of the corresponding variables across the N banks in the economy.

In keeping with actual stress tests, we usually interpret the systematic factors as innovations to the relevant macro variables. There remains, however, a tension between this interpretation and modern macroeconomic theory. In actual stress tests regulators use traditional macroeconomic variables such as GDP, unemployment, and house prices. In DSGE models, on the other hand, these macro variables would themselves be functions of underlying structural shocks such as productivity, beliefs, risk aversion, etc.[3] Formally, let ϵ^s be the structural shocks and H the

[3]The typical DSGE model contains Euler equations, production functions and resource constraints that lead to a set of equations such as

$$\mathbf{A}\mathbf{s}_t = \mathbf{C} + \mathbf{B}\mathbf{s}_{t-1} + \mathbf{D}\mathbb{E}_t\mathbf{s}_{t+1} + \mathbf{F}\epsilon_t^s.$$

When we solve the model we invert the mapping to obtain a VAR representation

$$\mathbf{s}_t = \mathbf{J} + \mathbf{Q}\mathbf{s}_{t-1} + \mathbf{H}\epsilon_t^s.$$

solution matrix of the DSGE model, so that $\mathrm{s} = H\epsilon^s$. In a fully specified model, banks' losses would also be functions of the structural shocks: $y_i(\epsilon^{\mathrm{s}}) = \tilde{\mathrm{x}}_{\mathrm{i}}'\epsilon^{\mathrm{s}} + \eta_i$, where $\tilde{\mathrm{x}}_i$ are structural exposures. This equation is equivalent to (1) when H is invertible. In that case we can write $\epsilon^s = H^{-1}\mathrm{s}$ and define $\mathrm{x}_i = H'^{-1}\tilde{\mathrm{x}}_i$, and we obtain $y_i(\mathrm{s}) = \mathrm{x}_i'\mathrm{s} + \eta_i$.

In theory the regulator could supply the structural shocks ϵ^s and ask for estimated losses. In practice regulators supply directly the macro variables s. This reflects the fundamental issue of model ambiguity. Even if H is invertible, models for H would likely differ across banks as well as between banks and regulators. By contrast, a handful of macro-economic variables (GDP, credit spreads, house and stock prices, etc.) are well-understood by all participants and capture much of the macro-economic dynamics that matter for expected losses. This is why stress tests are written in terms of s and not ϵ^s. In most of our applications we will assume that H is invertible and that the regulator feels confident about estimating H^{-1}. In that case there is no real difference between estimating x_i or $\tilde{\mathrm{x}}_i$ and we can assume that the factors are independently distributed.

2.2 Regulator's Preferences and Interventions

We model the regulator as a risk manager. As in Acharya et al. (2016), we assume that the regulator has preferences $U(W)$ over the total net worth of the banking system W. This specification makes sense when there is a relatively efficient way to relocate assets and equity across banks, e.g. when healthy banks can take over failed ones.[4] If the regulator believes that the risks in the system are too high she can intervene to force the banks to take remedial actions. In our context, an intervention is any requirement imposed on a bank to decrease its risk exposures. The regulator could force the bank to raise capital. She could ask the banks to sell assets or divest a particular line of business, thereby reducing several exposures at the same time. But she could also use targeted interventions such as imposing higher regulatory ratios against specific types of loans or against specific borrowers.

The most broad based intervention would be to ask banks to increase their capital by Δ. The

In our simple framework we have $\mathbf{J} = \mathbf{Q} = 0$ since we normalize the baseline scenario to 0 and we have only one period.

[4]More generally, we could have $U([w_i]_{1..N})$. This would capture the case where the idiosyncratic failure of bank i matters regardless of the health of the banking sector as a whole. As in the systemic risk literature, we impose the restriction that only W matters. As a result, a financial crisis only happens when the financial system as a whole is under-capitalized.

most granular description of intervention is at the bank×exposure level. If the regulator takes action $\mathrm{a}_i = \{a_{i,j}\}_{j=1..J}$ on bank i, the exposure of bank i to factor j becomes $(1 - a_{i,j})\, x_{i,j}$. To shorten the notation, we define the cell-by-cell multiplication operator $\circ$ as

$$(\mathbf{1} - \mathrm{a}_i) \circ \mathrm{x}_i \equiv [(1 - a_{i,1})\, x_{i,1}, \ldots, (1 - a_{i,J})\, x_{i,J}]\,.$$

Given an intervention policy $(\Delta, \mathbf{a})$, the total net worth of the financial system is $W = \bar{W} + \Delta - \overline{\eta} - \mathrm{s}\cdot \left(\sum_{i=1}^{N} (\mathbf{1} - \mathrm{a}_i) \circ \mathrm{x}_i\right)$. Interventions are costly. There are direct costs born by the regulators and the banks, as well as indirect costs from the disruption of valuable activities. We denote by $\kappa(\Delta)$ be the cost of requiring banks to increase their capital by Δ and by $\Phi(\mathbf{a})$ be the cost of action $\mathbf{a} = \{\mathrm{a}_i\}_{i=1..N}$.

Let $\mathscr{S}$ denote the information set of the regulator at the time when she chooses her intervention policy. The regulator's problem is then to choose an intervention policy to maximize her expected utility given by

$$\mathbb{E}\left[U\left(\bar{W} + \Delta - \overline{\eta} - \mathrm{s}\cdot \left(\sum_{i=1}^{N} (\mathbf{1} - \mathrm{a}_i) \circ \mathrm{x}_i\right)\right) \middle| \mathscr{S}\right] - \Phi(\mathbf{a}) - \kappa(\Delta)\,,$$

where $U(W) = W - \frac{\gamma}{2} W^2$ and $\Phi(\mathbf{a}) = \frac{1}{2} \sum_i \sum_j \phi_j a_{i,j}^2$.

The regulator will choose the optimal intervention policy maximize her expected utility. Conditional on the regulator information set $\mathscr{S}$, the optimal increase in capital required by the regulator is given by

$$\kappa'(\Delta) = \mathbb{E}\left[U'(W) \,|\, \mathscr{S}\right]. \tag{4}$$

The regulator will require an increase in capital such that the marginal cost of raising aggregate banking capital equals the expected marginal utility of doing so. The optimal targeted intervention in bank i along dimension j is given by

$$\frac{\partial \Phi(\mathbf{a})}{\partial a_{i,j}} = \mathbb{E}\left[x_{ij} s_j U'(W) \,|\, \mathscr{S}\right]. \tag{5}$$

In an interior solution, the first order condition (FOC) equates the marginal cost of an intervention to its expected marginal benefit. The expected marginal benefit of reducing the risk exposure to factor j in bank i depends on the covariance between the marginal social utility $U'(W)$ and the contribution of factor j to bank i's losses, $x_{ij} s_j$. Risk reduction is more valuable when we expect U' to be large when $x_{ij} s_j$ is positive.

Targeted interventions might require more information than non-targeted ones. For instance, the regulator could mandate a uniform reduction is all risky activities: this would not require much information, but this strategy would be costly since it would impact socially valuable activities. In that case a capital requirement might be more efficient. On the other hand, if the regulator had access to adequate information, she could mandate reductions only for the activities that create significant systematic risk. We think of stress tests as a way of eliciting information to determine what these activities are. The same logic applies to interventions across banks. Asking all banks to sell assets is simple but it might lead to fire sale prices. With better information the regulator might be able to ask only a subset of banks to reduce their exposures.

2.3 Prior beliefs and stress tests

The banks' risk exposures to the macro factors are unknown to the regulator and the banks. The regulator has prior beliefs over the distribution of exposures within banks and across banks. These prior beliefs come from historical experiences and the regulator's own risk models. We stack the banks' exposures in one large $NJ \times 1$ vector as follows

$$\mathbf{x} \equiv \begin{bmatrix} \mathrm{x}_1 \\ \vdots \\ \mathrm{x}_N \end{bmatrix},$$

and we summarize the regulator's prior over the vector of exposures $\mathbf{x}$ as

$$\mathbf{x} \sim N\left(\overline{\mathbf{x}}, \Sigma_x\right),$$

where the $NJ \times 1$ vector of unconditional means and the $NJ \times NJ$ covariance matrix are, respectively,

$$\overline{\mathbf{x}} = \begin{pmatrix} \overline{\mathrm{x}}_1 \\ \vdots \\ \overline{\mathrm{x}}_N \end{pmatrix} \quad \text{and } \Sigma_x = \begin{bmatrix} \Sigma_x^1 & \Sigma_x^{1,2} & \cdots & \Sigma_x^{1,N} \\ \Sigma_x^{1,2} & \Sigma_x^2 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \Sigma_x^{(N-1),N} \\ \Sigma_x^{1,N} & \cdots & \Sigma_x^{(N-1),N} & \Sigma_x^N \end{bmatrix}$$

with $\Sigma_x^i = \mathbb{V}\mathrm{ar}\left(\mathrm{x}_i\right)$ for all i and $\Sigma_x^{i,j} = \mathbb{C}\mathrm{ov}\left(\mathrm{x}_i, \mathrm{x}_j\right)$ for all $i \neq j$. If Σ_x^i is diagonal the regulator expects the exposures of bank i to the different factors to be independent of each other. If $\Sigma_x^{i,j} = 0$,

the regulator's prior is that the risk exposures of banks i and j are independent. In almost all empirically relevant cases the covariance matrices are not diagonal.

To learn about the banks' risk exposures, the regulator asks the banks to estimate and report their losses under a particular realization of the macroeconomic state. This choice of macroeconomic state is a scenario $\hat{s}$.

Definition 1. A scenario $\hat{\mathrm{s}}$ is a realization of the vector of states s.

A scenario $\hat{s}$ is a row-vector of size J that represents an aggregate state of the economy. We entertain two interpretations of the size of the state space, J. The simplest way is to think of J as exogenously given. There might be a limited number of macroeconomic variables (GDP, unemployment, house prices) that everyone agrees need to be included in the test. The other way to think about J is as a large number capturing the set of all possible risk factors and in any given tests many have zero loadings. A non-zero weight is then a statement about whether that risk factor is included in the particular stress test. Our model can then shed light on which risk factors should be used.

Given our normalization of the baseline state to $s = 0$, a scenario close to 0 is a scenario close to the baseline of the economy. A scenario $\hat{s}$ in which element $\hat{s}_j$ is large, represents a large deviation from the baseline along the dimension of factor j. The larger $|\hat{s}_j|$, the more extreme the scenario along dimension j. When designing a stress test, the regulator specifies a *set* of scenarios for which the banks need to report their losses.

Definition 2. A stress test is a collection of M scenarios $\{\hat{\mathrm{s}}_m\}_{m=1}^{M}$ presented by the regulator, and a collection of estimated losses $\{\hat{y}_{i,m}\}_{i=1..N}^{m=1..M}$ reported by the banks.

For each scenario m, each bank i estimates and reports its net losses $\hat{y}_{i,m}$ given the input parameters in scenario $\hat{\mathrm{s}}_m$.

2.4 Stress test results

Banks use imperfect models to predict their losses under the stress test scenarios. More precisely, we assume that bank i estimates its losses under scenario $\hat{\mathrm{s}}$ as

$$\hat{y}_i\left(\hat{\mathrm{s}}, M\right) = \hat{\mathrm{s}} \cdot \mathrm{x}_i + \hat{\epsilon}_i\left(\hat{\mathrm{s}}, M\right), \tag{6}$$

where the error term $\hat{\epsilon}_i(\hat{s}, M)$ plays an important role in our analysis. We use the following specification

$$\hat{\epsilon}_i(\hat{\mathrm{s}}, M) = \alpha_i(M)\,\epsilon_{i,0} + \tilde{\beta}_i\left(\left\|\hat{s}^{(m)}\right\|\right)\epsilon_{i,1}, \tag{7}$$

where $\alpha_i(M)$ is a positive scalar increasing in the number of scenarios, $\tilde{\beta}_i\left(\left\|\hat{s}^{(m)}\right\|\right)$ is a scalar function, and $\epsilon_{i,0}$ and $\epsilon_{i,1}$ are independent $N(0,1)$ random scalars. We assume $\tilde{\beta}_i(z) = \beta_i\left(z^{\frac{1}{2}} + z^{1+\theta}\right)$, with $\theta > 0$ which guarantees that the regulator can never learn an exposure perfectly.[5] The error terms $\epsilon_{i,0}$ and $\epsilon_{i,1}$ can be correlated across banks.

There are two sources of noise in a bank's report, $\epsilon_{i,0}$ and $\epsilon_{i,1}$, capturing different kinds of measurement errors and model uncertainty. The error term $\alpha_i(M)\,\epsilon_{i,0}$ captures bank level uncertainty that is present even in the baseline scenario. The assumption of an increasing $\alpha_i(M)$ captures capacity constraints either on the bank side or on the regulator side. On the bank side it captures the fairly uncontroversial idea that if a regulator asks too many questions, the average quality of the responses decreases. On the regulator side it also captures the difficulty of designing meaningful scenarios. Macroeconomic stress tests rely on few scenarios and a limited set of risk factors. For example, the Fed currently uses 28 risk factors and a couple of scenarios for the US economy and the rest of the world. The relevant value for M clearly depends on the context. In non-supervisory settings, stress tests of trading books routinely involve tens of thousands of risk factors and scenarios. The second source of noise is modulated by the norm of the scenario and it captures bank level risks that are increasingly more difficult to forecast in more extreme scenarios. These include extreme credit losses as well as potential operational, liquidity and funding risks. Finally, note that in our model the banks are the ones reporting their losses. Often the regulator also has models to estimate the risk exposures. We assume that these models are embedded in the priors of the regulator, but one could also easily add them as additional signals.[6]

The results of the stress test for one bank can be summarized in the $M \times 1$ vector

$$\hat{y}_i\left(\hat{S}\right) = \hat{S}\mathrm{x}_i + \hat{\varepsilon}_i,$$

where $\hat{y}_i\left(\hat{S}\right)$ represents the results of the test for bank i, the $M \times J$ matrix $\hat{S}$ gather the scenarios in the stress test, and the errors in bank i's reported losses are gathered in the $M \times 1$ vector $\hat{\varepsilon}_i$,

[5] Our analysis is robust to other specifications for the error term as long as $\mathbb{V}ar\left(\hat{\epsilon}_i(\hat{\mathrm{s}}, M)\right)$ is continuous at $\hat{s} = 0$ and $\lim_{\hat{s}_j^{(m)} \to \infty} \frac{\left(\hat{s}_j^{(m)}\right)^2}{\mathbb{V}ar\left(\hat{\epsilon}_i(\hat{\mathrm{s}}, M)\right)} = 0$.

[6] The regulator does not need to know the banks' models. The analysis remains unchanged as long as the regulator's beliefs over the model used by the banks are consistent with Equation (6).

Scenario Design	**Stress Testing**	**Intervention**
The regulator chooses stress scenarios	Banks report stress test results	The regulator chooses targetted interventions to reduce banks' risk exposures
$\hat{S}$	$\hat{\mathbf{y}}\left(\hat{S}\right)$	$\mathbf{a}^{\star}\left(\hat{\mathbf{y}}\right)$

Figure 1: Timeline

i.e.,

$$\hat{y}_i\left(\hat{S}\right) = \begin{bmatrix} \hat{y}_i\left(\hat{\mathrm{s}}_1, M\right) \\ \vdots \\ \hat{y}_i\left(\hat{\mathrm{s}}_M, M\right) \end{bmatrix}, \quad \hat{S} \equiv \begin{bmatrix} \hat{\mathrm{s}}_1' \\ \vdots \\ \hat{\mathrm{s}}_M' \end{bmatrix}, \quad \text{and} \quad \hat{\varepsilon}_i = \begin{bmatrix} \hat{\epsilon}_i\left(\hat{\mathrm{s}}_1, M\right) \\ \vdots \\ \hat{\epsilon}_i\left(\hat{\mathrm{s}}_M, M\right) \end{bmatrix}.$$

The structure of the bank's internal models are the outcome of learning about risk exposures from historical data. Therefore, the mistakes a bank makes in computing its expected revenues may be correlated *across scenarios*. The variance-covariance matrix of the errors made by bank i in computing its stress test results is given by

$$\mathbf{\Sigma}_{\hat{\varepsilon}}^{i} \equiv \mathbb{V}ar\left[\hat{\varepsilon}_i\right].$$

Differences in $\mathbf{\Sigma}_{\hat{\epsilon}}^{i}$ across banks reflect differences in information (priors), in the amount or quality of data available to each bank, or in the bank's information processing capacity. We assume that x_i and $\hat{\varepsilon}_i$ are independent, but we allow banks to make correlated mistakes.

2.5 Timing

To summarize, there are three stages in our model: the scenario design stage, the stress testing stage, and the intervention stage. First, the regulator chooses stress scenarios taking into account that the scenario choices will affect the information in the stress test results submitted by the banks. Then, the regulator elicits information from the banks in the form on stress tests. The banks' stress test results consist of projected losses for each bank under each scenario chosen by the regulator. Finally, the regulator chooses her targeted interventions after observing the stress test results. Figure (1) shows the timeline of the model.

3 Learning from stress tests

The information set of the regulator is a crucial input in choosing how to intervene and it depends on the regulator's prior beliefs and on the information she acquires. The main friction in out model is that the banks' exposures to the macro factors are not known by the banks nor by the regulator. However, the banks have imperfect (noisy) models to project their losses in a given state. We model stress tests as a mechanism for the regulator to elicit this information from the banks and acquire information. The results of the stress tests performed by the regulator allow her to learn about banks' risk exposures $\{\mathrm{x}_i\}_{i=1}^{N}$ and shape the regulator's optimal intervention policy.

3.1 A Kalman Filter

By implementing stress tests, the regulator elicits information from the banks which she uses to learn about their risk exposures. More specifically, the stress test results can be interpreted as signals about the banks' risk exposures by defining the error terms in the signals appropriately.

Let us briefly summarize all the sources of risks and information in our model:

1. Prior correlation among exposures within and across banks: These correlations are summarized by the $NJ \times NJ$ covariance matrix Σ_x. The covariance matrix Σ_x is independent of the scenarios and reflect the fundamental common exposures in the banking system as perceived by the regulator.

2. Correlation among estimation errors made by bank i across the M scenarios: These correlations are measured by the $M \times M$ matrix $\mathbf{\Sigma}^i_{\hat{\varepsilon}} \equiv \mathbb{V}ar\left[\hat{\boldsymbol{\varepsilon}}_i\right]$ where $\hat{\varepsilon}_i$ is an $M \times 1$ vector of errors that depends on the underlying scenarios: element m of $\hat{\varepsilon}_i$ is given by$\hat{\boldsymbol{\varepsilon}}_{i,m} = \alpha_i\left(M\right)\epsilon_{i,0} + \tilde{\beta}_i\left(\left\|\hat{s}^{(m)}\right\|\right) \cdot \epsilon_{i,1}$. This part crucially depends on the scenarios and captures the idea that the consequences of extreme scenarios are harder to estimate.

3. Reported losses for each bank i for each scenario in the stress test: $\hat{y}_i\left(\hat{S}\right)$ is an $M \times 1$ vector that summarizes the results of the stress test for bank i. These results depend on the scenarios by design.

For N banks, we respectively stack the reported losses of the banks in the stress test and the

mistakes in them in the $NM \times 1$ vectors

$$\hat{\mathbf{y}} \equiv \begin{bmatrix} \left[\hat{y}_1\left(\hat{S}\right)\right] \\ \vdots \\ \left[\hat{y}_N\left(\hat{S}\right)\right] \end{bmatrix} \quad \text{and} \quad \hat{\boldsymbol{\varepsilon}} = \begin{bmatrix} [\hat{\varepsilon}_1] \\ \vdots \\ [\hat{\varepsilon}_N] \end{bmatrix}.$$

Then, the information available to the regulator after seeing the results of the stress test can then be summarized in the following state space representation

$$\hat{\mathbf{y}} = \hat{\mathbf{S}}\mathbf{x} + \hat{\boldsymbol{\varepsilon}}, \tag{8}$$

where $\hat{\mathbf{S}} \equiv \left(\mathbf{I}_N \otimes \hat{S}\right)$ simply repeats $\hat{S}$ on its diagonal, and $\hat{\boldsymbol{\varepsilon}} \sim N\left(0, \boldsymbol{\Sigma}_{\hat{\varepsilon}}\right)$. Remember that the regulator observes $\hat{\mathbf{y}}$ and wants to learn about $\mathbf{x}$. Expressing the stress test as in equation (8) allows us to apply the Kalman filter and to obtain a full characterization of the posteriors beliefs of the regulator.

Proposition 1. *After observing the results $\hat{\mathbf{y}}$ of the stress test, the posterior beliefs of the regulator regarding the banks' risk exposures are*

$$\mathbf{x}|\,\hat{\mathbf{y}} \sim N\left(\hat{\mathbf{x}}, \hat{\Sigma}_{\mathbf{x}}\right) \ ,$$

where the posterior mean $\hat{\mathbf{x}}$, the Kalman gain K, and the residual covariance matrix $\hat{\Sigma}_{\mathbf{x}}$ are given by

$$\hat{\mathbf{x}} = \left(\mathbf{I}_{NJ} - K\hat{\mathbf{S}}\right)\bar{\mathbf{x}} + K\hat{\mathbf{y}}\,, \tag{9}$$

$$K = \Sigma_{\mathbf{x}}\hat{\mathbf{S}}'\left(\hat{\mathbf{S}}\Sigma_{\mathbf{x}}\hat{\mathbf{S}}' + \boldsymbol{\Sigma}_{\hat{\varepsilon}}\right)^{-1}\,, \tag{10}$$

$$\hat{\Sigma}_{\mathbf{x}} = \Sigma_{\mathbf{x}} - K\hat{\mathbf{S}}\Sigma_{\mathbf{x}}\,. \tag{11}$$

The proof of Proposition 1 is a direct application of the standard Kalman filter. The Kalman gain K is an $NJ \times MN$ matrix. A few special cases can give some intuition. With one bank ($N = 1$), then $K_{j,m}$ is a measure of the amount of information about the exposure to risk factor j contained in the results from scenario m. With one scenario ($J = 1$) and uncorrelated exposures among banks, $K_{j,m}$ also measures the reduction in uncertainty about bank j's exposure to the risk factor.

The posterior covariance matrix $\hat{\Sigma}_{\mathbf{x}}$ plays a critical role in our analysis. $\hat{\Sigma}_{\mathbf{x}}$ measures the residual uncertainty that persists after observing the results of the stress test. The true exposures

are distributed around $\hat{\mathbf{x}}$ with covariance $\hat{\Sigma}_{\mathbf{x}}$. The goal of the stress test is to reduce this residual uncertainty as much as possible, along dimensions that depend on the objective function and on the priors of the regulator.

In the standard state-space representation in Equation (8), the stress test scenarios determine the structure of the signals observed by the regulator. The scenarios also determine the precision of the banks' reported losses in Equation (7). Increasing $|s_j|$ in a scenario makes the results more informative about exposures to factor j, but extreme scenarios reduce the precision of the banks' estimates and the noise might spill over to the measurement of other exposures. On the other hand, the regulator can improve her learning by taking into account the fact that true exposures are correlated across positions and across banks.

When designing the scenarios, the regulator must anticipate how she will interpret and use the results of the test. The extent to which learning takes place is captured by the *expected distribution of the posterior mean*, given by

$$\hat{\mathbf{x}} \sim N\left(\overline{\mathbf{x}}, \Sigma_{\hat{\mathbf{x}}}\right), \tag{12}$$

where the ex-ante variance of the posterior mean is given by

$$\Sigma_{\hat{\mathbf{x}}} \equiv \Sigma_{\mathbf{x}} - \hat{\Sigma}_{\mathbf{x}} = K\left(\mathbf{I}_N \otimes \hat{S}\right)\Sigma_{\mathbf{x}}.$$

The matrix $\Sigma_{\hat{\mathbf{x}}}$ represents the expected amount of learning from stress test $\hat{S}$. If the stress test is pure noise, then $K = 0$, $\hat{\Sigma}_{\mathbf{x}} = \Sigma_{\mathbf{x}}$ and the regulator learns nothing, i.e., $\Sigma_{\hat{\mathbf{x}}} = 0$. If the test is fully informative, then $\hat{\Sigma}_{\mathbf{x}} = 0$ and the regulator learn exactly all the exposures, i.e., $\Sigma_{\hat{\mathbf{x}}} = \Sigma_{\mathbf{x}}$. One goal of the regulator is to maximize $\Sigma_{\hat{\mathbf{x}}} = K\left(\mathbf{I}_N \otimes \hat{S}\right)\Sigma_{\mathbf{x}}$ – or equivalently to minimize the residual uncertainty $\hat{\Sigma}_{\mathbf{x}}$ – to design a more parsimonious and accurate intervention policy. When choosing how much to learn, the regulator takes into account that the Kalman gain K is itself a function of the scenarios, given by equation (10).

Remark. A goal of the regulator is to maximize the amount of learning $\Sigma_{\hat{\mathbf{x}}} = K\left(\mathbf{I}_N \otimes \hat{S}\right)\Sigma_{\mathbf{x}}$ to intervene more accurately. Therefore, the regulator's scenario choice will maximize $\omega'\Sigma_{\hat{\mathbf{x}}}\omega$ for some weight vector $\omega \in [0,1]^{NJ}$, with $\sum_h \omega_h = 1$. Choosing stress scenarios effectively determines which risk exposures the regulator wants to learn more about. Without ex-post interventions, the regulator does not care about learning.

A Simple Example: Scenario Choice and Learning

Consider the case of one bank ($N = 1$), one scenario ($M = 1$), and two risk factors ($J = 2$). The Kalman gain in this case is a 2×1 vector $K = \begin{bmatrix} k_1 \\ k_2 \end{bmatrix}$. To simplify the notation, we omit the argument M and the bank-specific subscript i, and we denote $\sigma_1^2 \equiv \Sigma_{x,11}$, $\rho\sigma_1\sigma_2 \equiv \Sigma_{x,12}$ and $\sigma_\epsilon^2(\hat{s}) \equiv \mathbb{V}\mathrm{ar}\left[\hat{\epsilon}_i(\hat{\mathrm{s}}, 1)\right]$.

The posterior mean is then

$$\hat{\mathrm{x}} = \overline{\mathrm{x}} + K\left(\hat{y} - \hat{s}'\overline{\mathrm{x}}\right)$$

and the amount of learning is given by

$$\Sigma_{\hat{x}} = \begin{bmatrix} k_1\hat{s}_1 & k_1\hat{s}_2 \\ k_2\hat{s}_1 & k_2\hat{s}_2 \end{bmatrix} \begin{bmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{bmatrix},$$

where the Kalman gain is given by

$$k_1 = \frac{\sigma_1^2\hat{s}_1 + \rho\sigma_1\sigma_2\hat{s}_2}{\sigma_1^2\hat{s}_1^2 + 2\rho\sigma_1\sigma_2\hat{s}_1\hat{s}_2 + \sigma_2^2\hat{s}_2^2 + \sigma_{\hat{\epsilon}}^2(\hat{s})} \tag{13}$$

$$k_2 = \frac{\sigma_2^2\hat{s}_2 + \rho\sigma_1\sigma_2\hat{s}_1}{\sigma_1^2\hat{s}_1^2 + 2\rho\sigma_1\sigma_2\hat{s}_1\hat{s}_2 + \sigma_2^2\hat{s}_2^2 + \sigma_{\hat{\epsilon}}^2(\hat{s})} \tag{14}$$

These equations for the components of the Kalman gain show how the scenario choice affects the regulator's learning. The scenario choice $\hat{s}$ affects the Kalman gain in two ways: indirectly, through the amount of noise in the stress test result, $\sigma_{\hat{\epsilon}}^2(\hat{s})$, and directly, by determining the correlation between the stress test result and unknown risk exposures.

From the expressions above, it is easy to see that $\frac{\partial k_j}{\partial \sigma_{\hat{\epsilon}}^2} < 0$ and $\frac{\partial \Sigma_{\hat{x}}(j,j)}{\partial \sigma_{\hat{\epsilon}}^2} \leq 0$, which simply state that more noisy signals (higher $\sigma_{\hat{\epsilon}}^2(\hat{s})$) reduce the weight the regulator puts on the stress test results when updating her beliefs about the risk exposures and leave her with a higher residual uncertainty.

The only way in which the regulator can learn is by deviating from the baseline $s = 0$ since $\Sigma_{\hat{x}}(j,j) = 0$ when $\hat{s} = 0$. However, making the scenario more extreme, by increasing $\hat{s}_j^2$, increases the noise in the signal since, given our assumptions on $\hat{\epsilon}(\hat{s})$, we have

$$\sigma_{\hat{\epsilon}}^2(\hat{s}) = \alpha^2 + b^2\left(\left(\hat{s}_1^2 + \hat{s}_2^2\right)^{\frac{1}{2}} + \left(\hat{s}_1^2 + \hat{s}_2^2\right)^{(1+\theta)}\right),$$

which is increasing in $\hat{s}_j^2$. Therefore, more extreme scenarios reduce the amount of learning from stress tests by increasing the amount of noise in the stress test results.

The direct effect of $\hat{s}_i$ on $\Sigma_{\hat{x}}(i,i)$ depends on the prior covariance matrix of the regulator. For illustration purposes, suppose that $\rho = 0$.[7] When the regulator's priors are uncorrelated among risk exposures, we have

$$\Sigma_{\hat{x}}(j,j) = k_j \hat{s}_j \sigma_j^2 \quad \forall j,$$

where

$$k_j \hat{s}_j|_{\rho=0} = 1 - \frac{\sigma_h^2 \hat{s}_h^2 + \sigma_{\hat{\epsilon}}^2}{\sigma_1^2 \hat{s}_1^2 + \sigma_2^2 \hat{s}_2^2 + \sigma_{\hat{\epsilon}}^2(\hat{s})} \quad \forall j, h = 1, 2,\ j \neq h, \tag{15}$$

which is increasing in $\hat{s}_j$ keeping the amount of noise in the signal constant $\sigma_{\hat{\epsilon}}^2(\hat{s})$. Therefore, keeping the amount of noise fixed, putting more weight on factor j in the scenario leads to the regulator learning more about the bank's exposure to factor j and her putting more weight on the stress test results to update her beliefs about it. This clarifies the partial effect of the scenario choice holding constant the noise.

Lemma 1. *When $\rho = 0$, making factor j more extreme in the scenario, i.e., increasing s_j^2, has a non-monotone effect on the amount learned about the bank's exposure to factor j, i.e.,*

$$\lim_{\hat{s}_j^2 \to 0} \frac{d\Sigma_{\hat{x}}(j,j)}{d\hat{s}_j^2} > 0 \quad \textit{and} \quad \lim_{\hat{s}_j^2 \to \infty} \frac{d\Sigma_{\hat{x}}(j,j)}{d\hat{s}_j^2} = -\theta\sigma_j^2 < 0$$

The overall effect of an increase in $\hat{s}_1$ on the amount of learning about x_1 depends on the trade off between the increase in noise and the increase in correlation between the stress test result and x_1 that are generated by increasing $\hat{s}_1$. When the departure from the baseline is small, i.e., $\hat{s}_j^2$ is small, the direct effect dominates and stressing factor j more leads to more learning about factor j. When the scenario is very extreme along dimension j, i.e., $\hat{s}_j^2$ is large, the increase in noise dominates and stressing factor j more leads to less learning about factor j. The assumption that $\theta > 0$ guarantees that $\Sigma_{\hat{x}}(j,j)$ is eventually decreasing in $\hat{s}_j^2$.

The effect of the scenario choice on learning is a lot more complex when the risk exposures are correlated. In particular, when $\sigma_{\hat{\epsilon}}^2$ is sharply increasing in s_1 and ρ is large, it might be more efficient for the planner to learn about x_1 by increasing $\hat{s}_2$ instead of $\hat{s}_1$.

[7] The intuition extends for the case in which $\rho > 0$.

3.2 Scenario choice as information precision choice

The regulator cares about the scenarios $\hat{S}$ only through the distribution of $\hat{y}$. The stress scenarios $\hat{S}$ determine the structure of the signals about the banks' risk exposures contained in the stress test results $\hat{y}$. Given the Gaussian structure of the random variables and the linearity of the signals, this signal structure only depends on the amount of residual uncertainty that the regulator faces after observing $\hat{y}$, which is measured by the posterior precision of the exposures, $\hat{\Sigma}_{\mathbf{x}}$. Every set of stress scenarios $\hat{S}$ has a unique posterior covariance matrix associated with it. Therefore, choosing a set of scenarios $\hat{S}$ is equivalent to choosing a posterior covariance matrix $\hat{\Sigma}_{\mathbf{x}} \in \mathbf{\Sigma}$ (or alternatively, choosing how much to learn $\Sigma_{\hat{\mathbf{x}}}$), where the set $\mathbf{\Sigma}$ is given by Equations (10) and (11). The shape of the feasibility set $\mathbf{\Sigma}$ is determined only by the regulator's priors Σ_x and the mistakes in the banks' models, $\mathbf{\Sigma}_{\hat{\varepsilon}}$.

Proposition 2. *Choosing stress scenarios $\hat{S}$ is equivalent to choosing a posterior covariance matrix $\hat{\Sigma}_{\mathbf{x}} \in \mathbf{\Sigma}$. When $M < N$, the stress scenario design problem can be further reduced to choosing posterior variances about the risk exposures of any M banks.*

The Kalman filter maps scenarios M to the elements in the posterior covariance matrix$\hat{\Sigma}_x$, which is ultimately what the regulator cares about. More specifically, the Kalman filter implies $\frac{JN(JN-1)}{2}$ equations mapping the $J \times M$ elements in the scenario matrix $\hat{S}$ to the elements of $\hat{\Sigma}_x$. Therefore, when there are fewer scenarios than banks, it is equivalent for the regulator to choose the stress scenarios or $J \times M$ elements of the posterior covariance matrix. In particular, the regulator can focus on choosing the posterior variances about the risk exposures of M banks.

Using Proposition 2 we can write the regulator's scenario choice problem as an information precision choice problem, where the feasible set from which the regulator chooses is determined by the Kalman filter. This set, restricts how much information can be acquired by the regulator and it plays the role of a capacity constraints in models of information acquisition. Figures (2) shows the feasible set of posterior variances $\left\{\hat{\Sigma}_{x,11}, \hat{\Sigma}_{x,22}\right\}$ in a model with one representative bank and two risk factors, and for different values of prior correlations among risk exposures.

As illustrated by the example above, choosing a more extreme scenario has two effects on the amount of information that the regulator can acquire. On one hand, a higher value of $\hat{s}_i$ increases the weight the bank's stress test results put on the bank's exposure to factor i. On the other hand, more extreme scenarios are harder to predict and the stress test results become more noisy, i.e., the variance of the error term, $\mathbf{\Sigma}_{\hat{\varepsilon}}$, is larger. For scenarios that are close to the benchmark, small

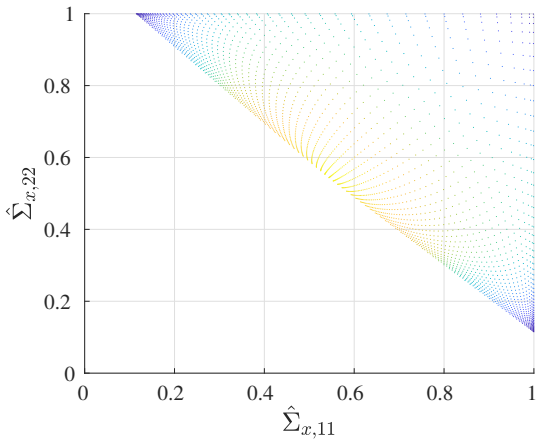

(a) Prior correlation in exposures $\Sigma_{x,12} = 0$

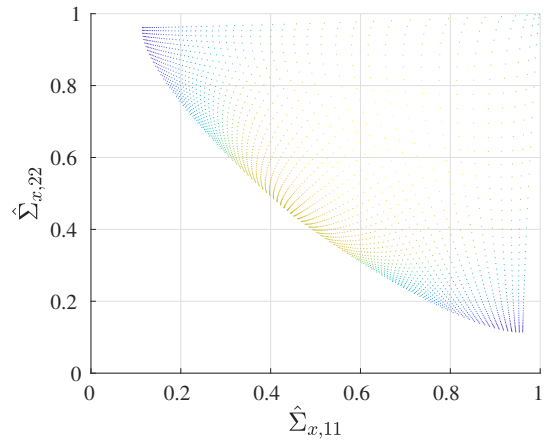

(b) Prior correlation in exposures $\Sigma_{x,12} = 0.2$

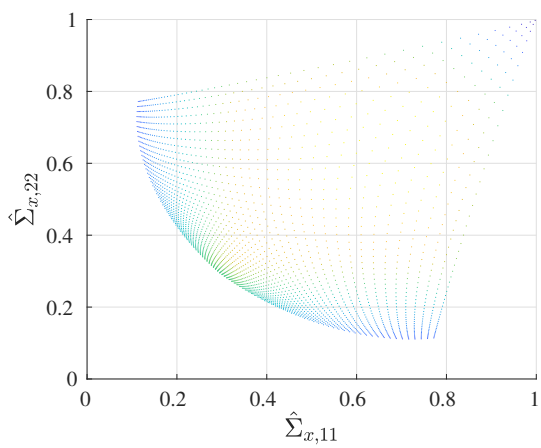

(c) Prior correlation in exposures $\Sigma_{x,12} = 0.5$

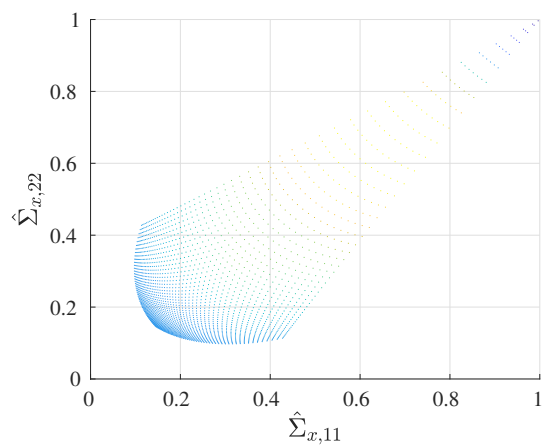

(d) Prior correlation in exposures $\Sigma_{x,12} = 0.8$

Figure 2: Feasible set of posterior variances, $\left(\hat{\Sigma}_{x,11}, \hat{\Sigma}_{x,22}\right)$ when there are two factors and one representative bank for different values of the regulator's prior correlation among the bank's risk exposures to factors 1 and 2.

Note: Figures 1 illustrates the set of feasible posterior variances, $\boldsymbol{\Sigma}$ for different values of prior correlations among risk exposures. The parameters used are $N = 1$, $J = 2$, $\gamma = 1$, $\phi_1 = 1$, $\bar{x}$=[1, 1], $\Sigma_x = \boldsymbol{I}_J$, $\alpha_1(M) = 0.8M$, $\beta_1 = 1$, $M = 1$, $\mathbb{E}\left[\epsilon_{1,0}^2\right] = 2$ and $\mathbb{E}\left[\epsilon_{1,1}^2\right] = 1$.

values of $(\hat{s}_1, \hat{s}_2)$, the first effect dominates and making the scenario more extreme in one dimension translates into lower posterior variances (acquiring more information) in that dimension. For more extreme scenarios, the second effect dominates and moving away form the benchmark reduces the amount of information the regulator can get from the stress test. These countervailing effects limit how much the regulator can learn form the stress test, as can be seen in Figures 2.

The prior correlation between the risk exposures also determines the shape of the feasible set. When the regulator's prior is such that risk exposures are correlated, the cost of increasing the value of $\hat{s}_i$ in terms of decreasing the information about risk exposure j is lower. Hence, the regulator can learn more from the stress tests and reduce the posterior variances. At the same time, the regulator cannot learn about the bank's exposure to factor 1 without learning about the bank's exposure to factor 2. Hence, as it can be seen from panels a, b, c and d in Figure 2, the boundary of set of feasible posterior precisions, $\mathbf{\Sigma}$, becomes more convex. When the correlation is high, the boundary of the feasible set slopes up in the tails. This is in line with the discussion in the example: when s_1 is already large, it becomes more efficient to learn about x_1 by increasing s_2 instead of s_1.

4 Taking action

The regulator values the information provided by the stress test because it allows her to intervene in a more accurate and parsimonious manner. How the regulator's actions depend on the stress test results will determine how valuable it is for the regulator to learn and, therefore, is a key determinant of the optimal scenario choice.

Suppose that the regulator has two ways of intervening in the banking sector. It can mandate an increase of capital of Δ or target interventions at the bank-risk level by requiring each bank i to reduce their exposure to factor j in $a_{i,j}$. Recall that, in this case, the aggregate banking wealth is given by

$$W(\mathbf{a}, \Delta) = \bar{W} + \Delta - \overline{\eta} - \sum_{i=1}^{N} \sum_{j=1}^{J} s_j \left(1 - a_{i,j}\right) x_{i,j}.$$

The regulator takes action after observing the results of the test. Therefore, the regulator's expected utility is

$$V(\mathscr{S}) = \max_{\Delta, \mathbf{a}} \mathbb{E}\left[U\left(W(\mathbf{a}, \Delta)\right) | \mathscr{S}\right] - \Phi(\mathbf{a}) - \kappa(\Delta),$$

where $\mathscr{S}$ denotes the total information from the stress test, which includes the matrix of scenarios and the actual outcomes:

$$\mathscr{S} \equiv \left\{\hat{S}, \hat{\mathbf{y}}\right\}.$$

Since $\Phi(\mathbf{a}) = \frac{1}{2}\sum_{i=1}^{N}\sum_{j=1}^{J}\phi_j a_{i,j}^2$ then we have that the optimal intervention policies for the regulator are given by

$$\kappa'(\Delta) = \mathbb{E}\left[U'(W) \mid \mathscr{S}\right]$$

and

$$\phi_j a_{i,j} = \mathbb{E}\left[s_j x_{i,j} U'(W) \mid \mathscr{S}\right].$$

These expressions are the same as the ones in Equations. (4) and (5). While the FOC for Δ equates the marginal cost of raising aggregate banking capital to the expected marginal utility, the FOC for the targeted interventions also depends on the covariance between scenario-weighted exposures and marginal utility.

Lemma 2. *With linear quadratic preferences, the capital requirement only depends on first order expected losses and solves*

$$\kappa'(\Delta^\star) = 1 - \gamma\left(\bar{W} + \Delta^\star + \mathbb{E}[\mathrm{s}] \cdot \left(\sum_{i=1}^{N}(\mathbf{1} - \mathrm{a}_i) \circ \hat{\mathrm{x}}_i\right)\right).$$

This Lemma is a reflection of the certainty equivalent. It highlights the limits of linear quadratic preferences, but it also emphasizes the fact that broad capital requirements are less likely to depend on the details of the stress test results. Since we define s as deviation from a central scenario we have $\mathbb{E}[s_j] = 0$ unconditionally, but we may want to consider cases where the planner happens to have information about the expected value of s. In this case, we can also interpret the capital requirement as capital adequacy under a severe but realistic scenario so that $\mathbb{E}[\mathrm{s}] > 0$.

Lemma 3. *With linear quadratic preferences, the targeted actions are given by*

$$a_j^\star = \left(\phi_j \mathbb{I}_N + \gamma\mathbb{E}\left[x^j x^{j\prime}\middle|\mathscr{S}\right]\mathbb{E}\left[s_j^2\right]\right)^{-1}\gamma\mathbb{E}\left[x^j x^{j\prime}\middle|\mathscr{S}\right]\mathbf{1}_{N\times 1}\mathbb{E}\left[s_j^2\right] \quad \forall j = 1, ..., J, \tag{16}$$

where $x^j = [x_{1,j}, ..., x_{N,j}]'$.

The optimal intervention along dimension j therefore depends on the posterior means of the banks' exposures to factor j $\hat{x}^j \equiv \mathbb{E}[x^j \mid \mathscr{S}]$ as well as on the residual uncertainty along this dimension$\hat{\Sigma}_x^j \equiv \mathbb{V}ar[x^j \mid \mathscr{S}]$, since $\mathbb{E}[x^j x^{j\prime} \mid \mathscr{S}] = \hat{\Sigma}_x^j + (\hat{x}^j \hat{x}^{j\prime})$. The residual uncertainty is known

in advance since the evolution of the covariance matrix is deterministic, but the posterior mean depends on the random realization the test itself, since $\hat{\mathrm{x}} = \overline{\mathrm{x}} + K\left(\hat{y} - \hat{s}'\overline{\mathrm{x}}\right)$.

Proposition 3. *The regulator intervenes more in bank i along dimension j when she is more risk averse, when factor j is more volatile, when $\mathbb{E}\left[x_{i,j}^2 \,|\mathscr{S}\right]$ is high, and when other banks have correlated exposures.*

A risk averse regulator will intervene more to reduce risk along dimension j in bank i when she perceives a higher risk, which has five components: risk aversion γ; the variance of risk factor j, $\mathbb{E}\left[s_j^2\right]$; the banks own expected exposure $\hat{x}_{i,j}$; the residual uncertainty and the correlation with other banks' exposures, captured by the diagonal and non-diagonal elements of $\hat{\Sigma}_{x,}^{j}$, respectively.

Once we account for the optimal actions the interim utility of the planner is

$$V\left(\mathscr{S}\right) = \mathbb{E}\left[U\left(W\left(\mathbf{a}^{\star}, \Delta^{\star}\right)\right) \,|\mathscr{S}\right] - \Phi\left(\mathbf{a}^{\star}\right) - \kappa\left(\Delta^{\star}\right)$$

which depends on the scenarios in the stress test, $\hat{S}$, and on the stress test results, $\hat{y}$.

5 Optimal scenario design

The regulator will choose the stress test scenarios $\hat{S}$ taking into account their impact on her ex-post interventions. Therefore, the regulator's problem can be written as

$$\max_{\hat{S}\in\bar{\mathbb{R}}} \mathbb{E}_{\hat{y}}\left[V\left(\mathscr{S}\right) \,|\hat{S}\right], \tag{17}$$

Depending on the particular application, we could incorporate a cost of creating additional scenarios for the regulator: choosing M scenarios for the stress test could have a cost $\mathscr{C}\left(M\right)$. In that case the objective function would simply be $\mathbb{E}_{\hat{y}}\left[V\left(\mathscr{S}\right) \,|\hat{S}\right] - \mathscr{C}\left(M\right)$ and the regulator would also choose the number of scenarios to include in the stress test.

Given that $\hat{y}$ is normally distributed and the macro factors are independently distributed with mean zero, we can integrate the indirect value function $\mathbb{E}_{\hat{y}}\left[V\left(\mathscr{S}\right) \,|\hat{S}\right]$ and express it as function of the covariance matrices along each risk dimension j. To see this, note that the regulator's objective depends on the result of the stress test $\hat{\mathbf{y}}$only through $\hat{\mathbf{x}}$. Moreover, the optimal targeted interventions along dimension j only depend on $\hat{\mathbf{x}}$ through $\hat{x}^j$ and on the posterior variance through the posterior covariance matrix of the banks' exposures to factor j $\hat{\Sigma}_x^j$. The following proposition shows how to express the regulator's scenario design problem as an information choice problem.

Proposition 4. *The stress scenario design problem is equivalent to choosing the posterior covariance matrix $\hat{\Sigma}_x$ to minimize the expected targeted intervention cost. Formally, the regulator's scenario design problem is*

$$\min_{\hat{\Sigma}_x \in \mathbf{\Sigma}} \mathbb{E}_{\hat{\mathbf{x}}} \left[\sum_{j=1}^{J} \sum_{i=1}^{N} \phi_j a_{i,j}^{\star} \left(\hat{x}^j, \hat{\Sigma}_x^j \right) \right], \tag{18}$$

where $a_{i,j}^{\star}$ is given by Equation 16 and $\mathbf{\Sigma}$ is the set of feasible posterior covariance matrices implied by the Kalman filter.

Since ex-post interventions are costly for the regulator, she chooses posterior variances to minimize the total amount of expected interventions weighted by the average marginal cost of intervening. Even though the objective function is separable in the dimension of the factors, the precision choices along the factor dimensions are not independent of each other. The set $\mathbf{\Sigma}$ restricts the amount of information about the exposures to different factors that the regulator can choose. For example, when the regulator's prior is that risk exposures to different factors are uncorrelated, increasing the precision of the stress test results as a signal of the exposure to one factor comes at the expense of receiving less precise information about all other factors to which the bank is exposed.

To highlight the intuition behind the regulator's scenario choice, in the remainder of the paper, we will focus on the case in which the regulator finds it optimal to choose one scenario and $M = 1$. Given Propositions 2 and 4 we will restrict our attention to the case in which there is one bank for most of our analysis. In this case, the regulator solves

$$\min_{\left\{\hat{\Sigma}_{x,jj}\right\}_{j=1}^{N} \in \overline{\mathbf{\Sigma}}} \mathbb{E}_{\hat{\mathbf{x}}} \left[\sum_{j=1}^{J} \phi_j a_j^{\star} \left(\hat{x}_j, \hat{\Sigma}_{x,jj} \right) \right], \tag{19}$$

where $\overline{\mathbf{\Sigma}}$ is the set of feasible posterior variances implied by the Kalman filter.

From Proposition Equation (16) we know that the optimal intervention to reduce the exposure to factor j in bank i, $a_j^{\star}\left(\hat{x}_j, \hat{\Sigma}_{x,jj}\right)$, is increasing in the uncertainty about the bank's exposure to factor j, $\hat{\Sigma}_{x,j}$, and in the bank's expected exposure to that factor, $\hat{x}_j$. Moreover, since $\hat{x}_j \sim N\left(\bar{x}_j, \left(\Sigma_{x,jj} - \hat{\Sigma}_{x,jj}\right)\right)$ we can write

$$\hat{x}_j = \left(\Sigma_{x,jj} - \hat{\Sigma}_{x,jj}\right)^{\frac{1}{2}} z_j + \bar{x}_j$$

where $z^j \sim N(0,1)$. The derivative of the expected intervention in dimension j with respect to

$\hat{\Sigma}_x^j$ is

$$\phi_j \frac{d\mathbb{E}_{\hat{x}}\left[a_j^\star\left(\hat{x}_j, \hat{\Sigma}_{x,jj}\right)\right]}{d\hat{\Sigma}_{x,jj}} = \phi_j \mathbb{E}_z\left[\frac{\partial a_j^\star\left(\hat{x}_j, \hat{\Sigma}_{x,jj}\right)}{\partial \hat{\Sigma}_{x,jj}} + \frac{\partial a_j^\star\left(\hat{x}_j, \hat{\Sigma}_{x,jj}\right)}{\partial \hat{x}_j}\frac{\partial \hat{x}_j}{\partial \hat{\Sigma}_{x,jj}}\right]. \tag{20}$$

The first term inside the expectation on the right-hand side of Equation (20) represents the change in the targeted intervention policy due to precautionary motives: the regulator intervenes more aggressively along dimensions she is more uncertain about. This effect shifts the optimal intervention policy in the same way for all realizations of the stress test results: a decrease in $\hat{\Sigma}_x^j$ shifts the optimal policy function to the left.

The second term captures the value of intervening more accurately and it depends on the sensitivity of the targeted intervention to the ex-post expected exposures $\hat{x}_j$, which given by $a_j^{\star\prime}(\hat{x}_j)$, and on how the new information from the stress test changes this posterior mean, which is determined by $\frac{\partial \hat{x}_j}{\partial \hat{\Sigma}_{x,jj}}$. Suppose that the regulator's prior beliefs were very precise. In this case, only very extreme realizations of $\hat{\mathbf{y}}$ would move the regulator's priors and the sensitivity of the intervention policy to new information would be very low, i.e., $\frac{\partial \hat{x}_j}{\partial \hat{\Sigma}_{x,jj}}$ would be small. In this case, increasing the precision of the stress test along dimension j would not improve the accuracy of the intervention much and the value of learning along dimension j would be low. Similarly, if the intervention policy was not very responsive to information, reducing the residual uncertainty of the regulator along dimension j would not have a large effect on the intervention policy and the value of information would be low.

The regulator is always better off with more precise information. What limits her information choice is the feasibility set $\overline{\mathbf{\Sigma}}$, which is determined by the loss model used by the bank and by the regulator's beliefs. While the regulator will always choose posterior variances on the frontier of $\overline{\mathbf{\Sigma}}$, the exact choice of information depends on the marginal benefits of learning across the different dimensions.

If the solution to the regulator's problem is interior, the regulator chooses to diversify her learning and designs a stress test that provides her with information along multiple dimensions. However, the regulator may prefer to specialize her learning and focus only on reducing her uncertainty about a subset of the bank's risk exposures. In the remainder of this section we provide numerical examples to illustrate how the relative intervention cost and the regulator's prior beliefs about the bank's risk exposures determine her optimal scenario choice.

Specialization in learning

Whether the regulator chooses to diversify or specialize her learning about the bank's risk exposures depends on the convexity of the intervention cost function relative to the risk aversion of the regulator, and on the shape of the feasibility set for posterior precisions. More specifically, the regulator will choose to specialize or diversify her learning depending on the convexity of the indifference curves relative to the convexity of $\overline{\mathbf{\Sigma}}$. The ratio $\frac{\phi_j}{\gamma}$ determines the sensitivity of the ex-post intervention to new information and, through it, the shape of the regulator's indifference curves over the posterior precisions. The prior correlation between the exposures determines the concavity of the set $\overline{\mathbf{\Sigma}}$.

Figures (3) show the optimal scenario and optimal posterior variance as a function of $\frac{\phi}{\gamma}$ when there are two symmetric risk factors and one bank whose priors over the risk exposures are uncorrelated. Intuitively, for low values of $\frac{\phi}{\gamma}$ the intervention action $a_j^\star$ is very sensitive to the precision of the information in the stress test about the exposure to factor j, measured by $\hat{\Sigma}_{x,jj}$. In this case, the regulator can reduce her intervention cost significantly by intervening accurately along both risk dimensions and chooses to diversify her learning, i.e., the optimal scenario puts weight on both risk factors. As $\frac{\phi}{\gamma}$ increases, the sensitivity of the optimal intervention to the precision of the stress test result decreases. Now, the regulator needs very precise information to improve the accuracy of her intervention enough to reduce her intervention costs significantly. In this case, the regulator chooses to specialize her learning and learn only about one risk factor (only one factor has positive weight in the optimal scenario). When $\frac{\phi}{\gamma}$ is high enough, the regulator finds it optimal to intervene accurately in one dimension and having a blanket intervention in the other rather than intervening with low accuracy in both dimensions.

Intervention costs

The effect of intervention costs on the optimal scenario is not monotone. When the intervention costs are low, the regulator can intervene substantially ex-post without creating significant dead-weight losses. In this case, intervening inaccurately is not too costly and the regulator cares little about learning about that factor. When the intervention costs are intermediate, interventions are sensitive to the information of the stress test and their cost is significant. In this case, the regulator values learning more since more precise information allows her to avoid wasteful interventions. Finally, when the intervention costs are very high, the ex-post interventions are small irrespective

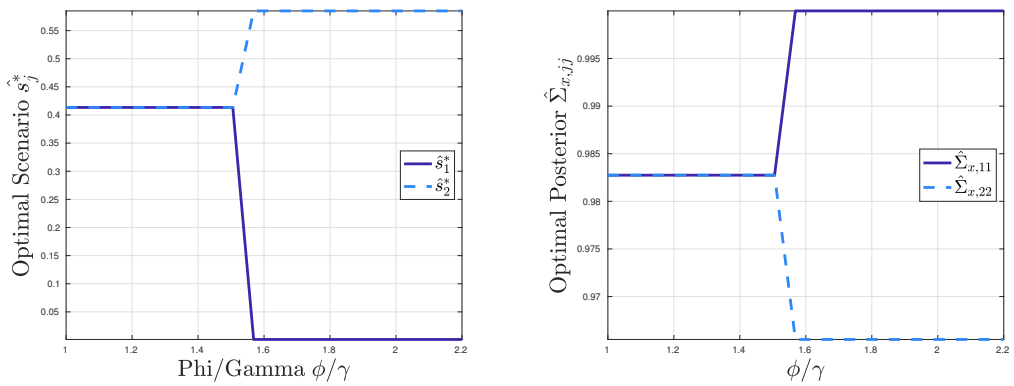

(a) Optimal scenario choice

(b) Optimal posterior

Figure 3: Optimal scenario and information choice as a function the intervention cost to reduce the exposure to factor 2, ϕ_2.

Note: Figure 3 illustrates the regulator's optimal choice of scenario and the implied posterior variance as a function of the intervention cost along the dimension of factor 2. The parameters used are $N = 1$, $J = 2$, $\gamma = 0.3$, $\bar{x}$=[1, 1], $\Sigma_x = \boldsymbol{I}_J$, $\alpha_1(M) = 0.8M$, $\beta_1 = 1$, $M = 1$, and $\theta = 2$, $\mathbb{E}\left[s_k^2\right] = 1$, $\mathbb{E}\left[\epsilon_{1,0}^2\right] = 1$ and $\mathbb{E}\left[\epsilon_{1,1}^2\right] = 1$.

of what the regulator might learn. In this case, the sensitivity of the ex-post intervention with respect to the new information is low and learning is less valuable for the regulator.

Figure (4) illustrates the optimal scenario choice and the implied posterior precisions as the cost of intervening to reduce the bank's exposure to factor 2, ϕ_2, increases when there are two factors and one bank. As it can be seen in the first panel of the figure, the weight of factor 2 in the optimal scenario increases and then decreases with ϕ_2. At first when intervening to reduce the exposure to factor 2 is slightly higher than the cost of intervening along the dimension of factor 1, the regulator finds it optimal to learn more about factor 2 at the expense of decreasing her learning about the exposure to factor 1 to minimize her aggregate intervention costs. However, as ϕ_2 continues to increase, the sensitivity of $a_2^\star$ to the information in the stress test decreases making information about the exposure to factor 2 less effective in decreasing the aggregate intervention cost than information about the exposure to factor 1. This leads to a decrease in the weight of factor 2 and an increase in the weight of factor 1 in the stress test. In the limit, when ϕ_2 goes to infinity and becomes prohibitive, the regulator chooses not to learn about the exposure to factor

2 at all, since her intervention policy will be independent of any information about this exposure.

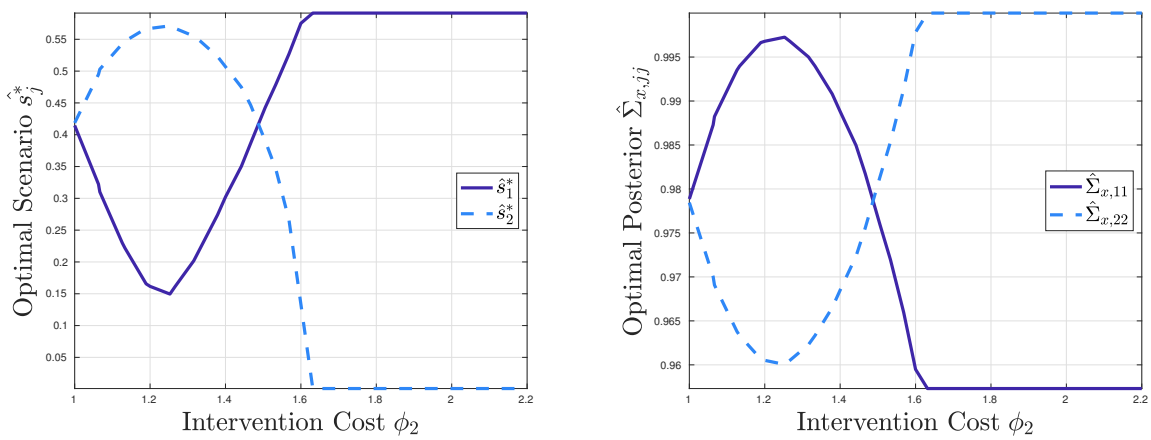

(a) Optimal scenario choice

(b) Optimal posterior

Figure 4: Optimal scenario and information choice as a function the intervention cost to reduce the exposure to factor 2, ϕ_2.

Note: Figure 4 illustrates the regulator's optimal choice of scenario and the implied posterior variance as a function of the intervention cost along the dimension of factor 2. The parameters used are $N = 1$, $J = 2$, $\gamma = 0.3$, $\phi_1 = 1$, $\bar{x}$=[1, 1], $\Sigma_x = \boldsymbol{I}_J$, $\alpha_1(M) = 0.8M$, $\beta_1 = 1$, $M = 1$, and $\theta = 2$, $\mathbb{E}\left[s_k^2\right] = 1$, $\mathbb{E}\left[\epsilon_{1,0}^2\right] = 1$ and $\mathbb{E}\left[\epsilon_{1,1}^2\right] = 2$.

Figure (5) shows the expected optimal intervention as a function of the intervention cost to reduce the exposure to factor 2. The expected intervention along the dimension of factor 2 is decreasing and convex in the intervention cost ϕ_2. When the intervention cost ϕ_2 is small, the expected intervention to reduce the exposure to factor 2 decreases for two reasons. First, the regulator chooses to intervene less because intervening is more expensive. Second, more precise information about the bank's exposure to factor 2 allows the regulator to intervene more accurately and less wastefully. This second effect is not present when the intervention cost ϕ_2 is high, and therefore, the decline in the expected intervention cost is lower.

Expected risk exposures

The prior mean of the exposure of the bank to each factor affects how much the regulator wants to learn about each risk exposure. Similarly to the comparative statics exercise with respect to

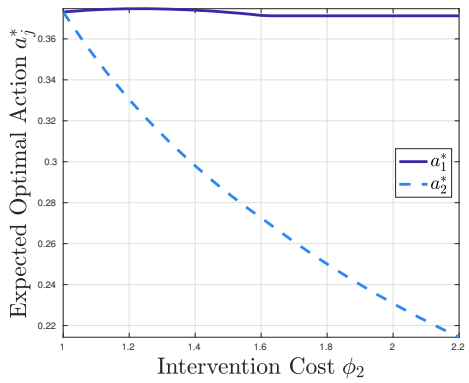

Figure 5: Optimal expected ex-post interventions as a function the intervention cost to reduce the exposure to factor 2, ϕ_2.

Note: Figure 5 illustrates the regulator's optimal expected ex-post interventions to reduce the bank's exposure to factors 1 and 2 as a function of the intervention cost along the dimension of factor 2. The parameters used are $N = 1$, $J = 2$, $\gamma = 0.3$, $\phi_1 = 1$, $\bar{x}$=[1, 1], $\Sigma_x = \boldsymbol{I}_J$, $\alpha_1(M) = 0.8M$, $\beta_1 = 1$, $M = 1$, and $\theta = 2$, $\mathbb{E}\left[s_k^2\right] = 1$, $\mathbb{E}\left[\epsilon_{1,0}^2\right] = 1$ and $\mathbb{E}\left[\epsilon_{1,1}^2\right] = 2$.

intervention costs, the relation between the prior mean exposure and the optimal scenario is non monotonic. There are two opposing effects of a higher expected risk exposure on the optimal stress scenario. First, the regulator's intervention policy is higher when the expected risk exposure to a factor is higher. This makes intervening accurately along this dimension more valuable and increases the value of learning about the exposures to factors with high expected exposures. Moreover, the regulator's beliefs, and hence her intervention policy, are less sensitive to the new information in the stress test when the expected risk exposure is higher. For example, if the prior mean exposure to factor 1 is high, the regulator's posterior mean is anchored around this value, i.e., the posterior mean is likely to be large regardless of the information revealed in the stress test results. In this case, the regulator expects to intervene substantially along dimension 1 for most stress test results. Therefore, new information about the exposure to factor 1 is not very valuable, as it does not change the regulator's ex-post behavior substantially. In this case, the weight of factor 1 in the stress test scenario decreases with the expected prior mean exposure to factor 1.

When the expected risk exposure is low, the first effect dominates and the weight of a factor in the stress test is increasing (decreasing) in the prior mean exposure to it. When the prior risk exposure is high enough, the regulator finds it optimal to set factor 1 to zero in the stress test and

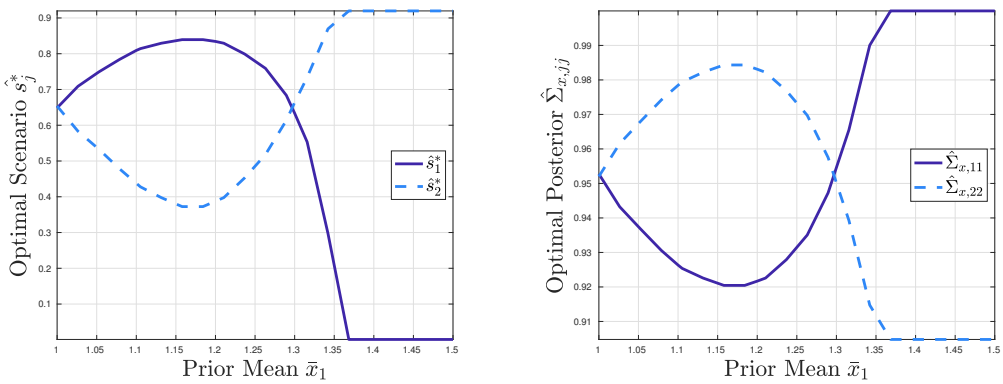

(a) Optimal scenario choice

(b) Optimal posterior variances

Figure 6: Optimal scenario and information choice as a function of the prior mean $\bar{x}_1$.

Note: Figure 6 illustrates the regulator's optimal choice of scenario and the implied posterior variance as a function of the prior mean of the risk exposure to factor 1. The parameters used are $N = 1$, $J = 2$, $\gamma = 1$, $\phi_1 = 1$, $\bar{x}_2$=1, $\Sigma_x = \boldsymbol{I}_J$, $\alpha_1(M) = 1.3M$, $\beta_1 = 1$, $M = 1$, and $\theta = 3$, $\mathbb{E}\left[s_k^2\right] = 1$, $\mathbb{E}\left[\epsilon_{1,0}^2\right] = 1.2$ and $\mathbb{E}\left[\epsilon_{1,1}^2\right] = 1$.

learn only about factor 2. This does not imply that the regulator stops caring about the bank's exposure to risk factor 1. However, the regulator will choose to reduce the risk generated by the bank's exposure to factor 1 by intervening more heavily ex-post and not learning about it at all.

Figure 7 shows the optimal expected ex-post interventions as a function of the prior mean exposure to factor 1 for the case in which there are two risk factors and one bank. Panel (a) in Figure 6 shows the weights of risk factors 1 and 2 in the optimal scenario as $\bar{x}_1$ changes. Panel (b) in the same figure shows the posterior precisions for the bank's risk exposures as a function of the prior risk exposure to factor 1.

Uncertainty

There are two dimensions of uncertainty that affect the regulator's choice of stress scenario: uncertainty about the risk exposures and uncertainty about the realization of the risk factors. As one would expect, everything else equal, the regulator wants to intervene more along dimension about which she is more uncertain about due to precautionary motives. Moreover, when the regulator is more uncertain about the bank's exposure to a risk factor j, she puts more weight on

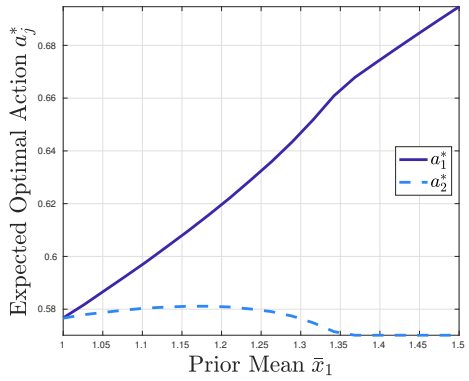

Figure 7: Optimal expected ex-post interventions as a function of the prior mean $\bar{x}_1$.

Note: Figure 6 illustrates the regulator's the regulator's optimal expected ex-post interventions to reduce the bank's exposure to factors 1 and 2 as a function of the prior mean exposure to factor 1.The parameters used are $N = 1$, $J = 2$, $\gamma = 1$, $\phi_1 = 1$, $\bar{x}_2$=1, $\Sigma_x = \boldsymbol{I}_J$, $\alpha_1(M) = 1.3M$, $\beta_1 = 1$, $M = 1$, and $\theta = 3$, $\mathbb{E}\left[s_k^2\right] = 1$, $\mathbb{E}\left[\epsilon_{1,0}^2\right] = 1.2$ and $\mathbb{E}\left[\epsilon_{1,1}^2\right] = 1$.

the stress test results to update her beliefs. This implies that the regulators targeted intervention along dimension j is more responsive to the information contained in the stress test results and information is more valuable.

Figure 8 shows the effect of uncertainty over the optimal scenario choice when there are two factors and one bank. Panels (a) and (b) respectively show the weights of risk factors 1 and 2 in the optimal scenario and the posterior precisions for the bank's risk exposures as a function of the prior uncertainty about the bank's exposure to factor 2, $\Sigma_{x,22}$. Note that the posterior variance about the exposure to factor 2 mechanically increases in $\Sigma_{x,22}$, even when it is optimal to only learn about factor 2.

As discussed above, when the regulator is more uncertain about the bank's exposure to a risk factor, she expects to intervene more and her actions are more responsive to the information contained in the stress test results. This implies that the weight of factor 2 is increasing in the regulator's prior uncertainty about the bank's exposure to factor 2. When the regulator's prior uncertainty about the exposure to factor 2 is large enough, the regulator chooses to learn only about this exposure and only puts weight on factor 2 in the optimal scenario. Panels (a) and (b) in Figure 8 show respectively the weights of risk factors 1 and 2 in the optimal scenario and the posterior precisions for the bank's risk exposures as a function of the prior uncertainty about the

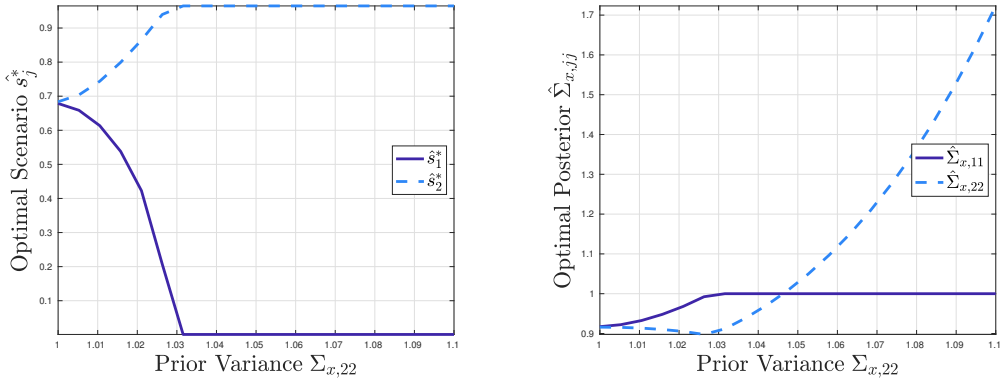

(a) Optimal scenario choice

(b) Optimal posterior

Figure 8: Optimal scenario and information choice as a function of the regulator's prior uncertainty of the exposure to factor 2, $\Sigma_{x,22}$.

Note: Figure 8 illustrates the regulator's optimal choice of scenario and the implied posterior variance as a function of the regulator's prior variance of the exposure to factor 2, $\Sigma_{x,22}$. The parameters used are $N = 1$, $J = 2$, $\gamma = 1$, $\phi_j = 1$, $\bar{x}$=[1, 1], $\Sigma_x = \boldsymbol{I}_J$, $\alpha_1(M) = 0.6M$, $M = 1$, $B_1 = \mathbb{I}_{2\times 2}$, and $C_1 = \mathbb{I}_{2\times 2}$

bank's exposure to factor 2, $\Sigma_{x,22}$. The effect of an increase in uncertainty about the risk factors is analogous. Figure A.1 in the Appendix shows this case.

Correlated risks

The correlation among risk exposures, within and across banks, plays a crucial role in determining how much the regulator can learn from a stress test. As it can be seen from the feasible set of posterior precisions, $\boldsymbol{\Sigma}$, in Figures (2), the regulator faces a trade-off between learning about one risk exposure or the other. When the prior correlation among risk exposures is zero, scenarios that provide a lot of information about the bank's risk exposure to factor 1 contain very little information about the bank's other risk exposures. For example, in this case, an extreme scenario that puts all the weight on risk factor 1 and no weight on the other risk factors contains no information about the bank's risk exposure the remaining factors at all. When the correlation between risk exposures is non-zero, this trade-off is attenuated as signals about one risk exposure always contain some information about the other.

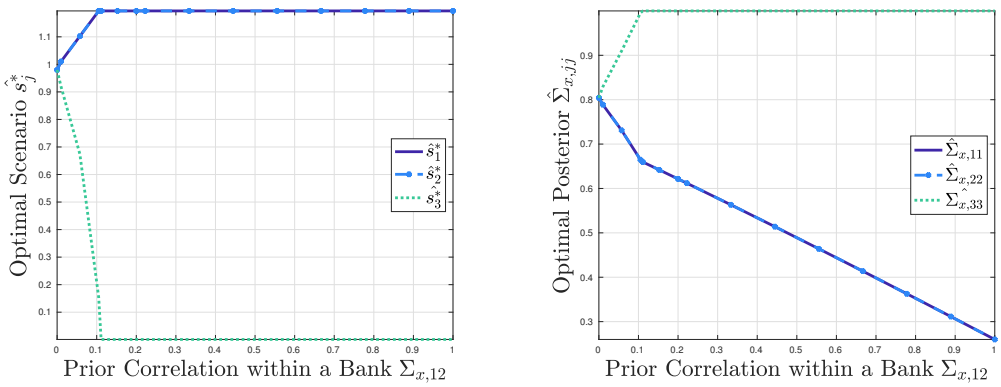

(a) Optimal scenario choice

(b) Optimal posterior variance

Figure 9: Optimal scenario and information choice as a function of the prior correlation in risk exposures, $\Sigma_{x,12}$.

Note: Figure 9 illustrates the regulator's optimal choice of scenario and the implied posterior variance as a function of prior correlation in risk exposures. The parameters used are $N = 1$, $J = 3$, $\gamma = 0.3$, $\phi_1 = 1$, $\bar{x}$=[1, 1, 1], $\Sigma_x = \boldsymbol{I}_J$, $\alpha_1(M) = 0.8M$, $\beta_1 = 1$, $M = 1$, and $\theta = 2$, $\mathbb{E}\left[s_k^2\right] = 1$, $\mathbb{E}\left[\epsilon_{1,0}^2\right] = 1$ and $\mathbb{E}\left[\epsilon_{1,1}^2\right] = 1$

Panel (a) in Figure 9 plots the weight on risk factor 1 in an optimal scenario as a function of the regulator's prior correlation among risk exposures. As one can see in this figure, the correlation in risk exposures does not affect the scenario choice. However, as panel (b) shows, the amount of information that the regulator can get from the same scenario changes as the prior correlation in risk exposures varies. As bank's risk exposures become more correlated in the bank's prior, the more informative the stress test.

When there are more than two factors, the scenario choice will be affected by the prior correlation among the bank's risk exposures to factors 1 and 2. As this correlation increases, the regulator will put more weight on the factors to which the bank's exposures are correlated and decrease the weight of the remaining factor. When the exposures of the bank to factors 1 and 2 are very correlated, the regulator finds it optimal to learn only about these two risk exposures and chooses not to deviate from the benchmark for factor 3.

Similarly, when there are multiple banks, the exposures to the risk factors are likely to be correlated among banks, even when the factors are orthogonal. Figure 10 shows the optimal scenario weights as a function of the correlation between the exposure to risk factor 1 for banks 1

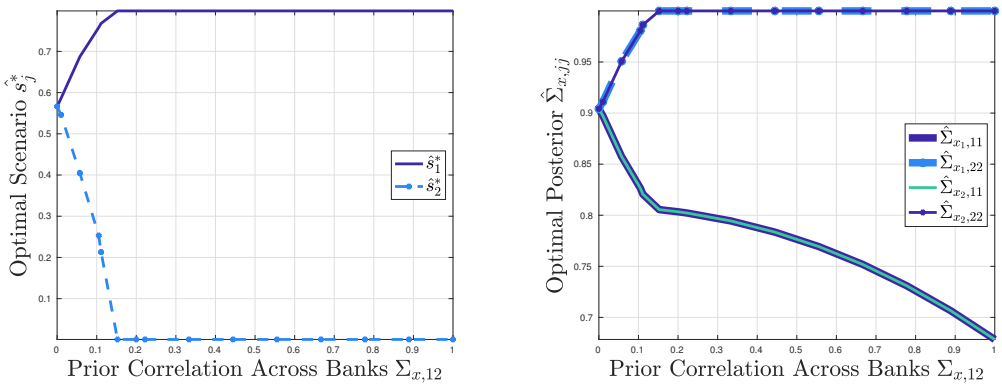

(a) Optimal scenario choice

(b) Optimal posterior variance

Figure 10: Optimal scenario and information choice as a function of the prior correlation between the bank's risk exposures to factor 1, $\Sigma_{x,12}^{11}$.

Note: Figure 10 illustrates the regulator's optimal choice of scenario and the implied posterior variance as a function of prior correlation in risk exposures. The parameters used are $N = 2$, $J = 2$, $\gamma = 0.3$, $\phi_j = 1$, $\bar{x}_2$=1, $\Sigma_x = \boldsymbol{I}_J$, $\alpha_1(M) = 0.1M$, $\beta_i = 1$, $M = 1$, and $\theta = 3$, $\mathbb{E}\left[s_k^2\right] = 1$, $\mathbb{E}\left[\epsilon_{j,0}^2\right] = \mathbb{E}\left[\epsilon_{j,1}^2\right] = 1$ and $\mathbb{E}\left[\epsilon_{1,0}\epsilon_{2,0}\right] = \mathbb{E}\left[\epsilon_{1,1}\epsilon_{2,1}\right] = 1$.

and 2, when there are two banks, two risk factors and one scenario.

When the exposures to factor 1 are correlated across banks, the stress test results of bank i contain information about bank i's exposures to both risk factors and about bank j's exposure to factor 1, as long as the weight of factor 1 in the stress test's scenario is non-zero. Figure 10 shows that, in this case, learning about factor 1 becomes more valuable for the regulator to the point that the stress scenario only puts weight on factor 1 when the risk exposures of the banks to factor 1 are very correlated. In this case, the regulator does not learn about the risk exposures to factor 2 and only chooses to intervene ex-post to reduce the losses associated with this factor. The posterior variance of the risk exposures to factor 2 remain constant and equal to the prior. On the other hand, the posterior variance of the exposure of the banks to factor 1 decreases mechanically with $\Sigma_{x,11}^{12}$ since the the stress test of bank j becomes more informative of bank i's exposure to factor 1 as this correlation increases.

6 Practical implications

Our comparative static exercises above shed light on the optimal stress scenario design in the presence of systemic factor, in times of distress, over time, and its relation to capital requirements.

Systemic factors Stress tests are widely used as a risk management tool. Regulatory stress testing in particular focuses on assessing the resilience of the financial system as a whole. In this context, systemic risk factors that lead to correlated losses among banks are of particular interest. Our analysis on correlated exposures among banks suggests that the optimal stress scenario would put relatively more weight on systemic factors. Moreover, if the correlation of the banks' exposures to some factor is high enough, the optimal stress scenario may put weight only on these systemic factors.

Scenario design in distress times Distress times are often associated with increased uncertainty about the evolution of macroeconomic variables. From our analysis above, it follows that the optimal stress test design in our model calls for stressing more uncertain factors more. Moreover, our model implies that an across the board increase in uncertainty can lead to the optimal stress scenario putting more weight on fewer factors.[8] This implies a negative correlation between the severity of the stress scenarios and the state of the economy.

Evolution of stress scenarios Changes in the information set of the regulator will lead to different scenario choices. The more the regulator knows about the banks' exposures to a particular factor, the less she will choose to stress that factor in the optimal stress scenario. Since the regulator learns more about the exposures to factors that are stressed more in the optimal scenario, one could imagine the optimal evolution of stress scenarios putting weight on different factors each time and resembling experimentation. This changes however, may not be driven by changes in the regulator's objective nor by changes in the expected evolution of risk factors but they may simply reflect the evolution of the regulator's information about the banks' losses.

Capital requirements In our setup, the capital increase Δ set by the regulator is not contingent on the result of the stress test. This is a consequence of our choice to model the scenarios as deviations from a baseline in which $\mathbb{E}[s_j] = 0$. However, our model does not imply that capital requirements are independent of stress test results. Even though we do not model capital

[8] Figure (A.2) in the Appendix shows the optimal scenario weights as a function of uncertainty about symmetric risk factors for the case in which there are two factors and one bank.

requirements explicitly, one can think of capital requirements being met by reducing a bank's exposure to certain risk assets and limiting the bank's losses, which is exactly what our targeted interventions achieve. The uncontingent capital increase can be interpreted as implementing a minimum capital requirement while the contingent targeted interventions can be interpreted as setting a stress capital buffer.

7 Conclusion

Despite the growing importance of stress testing for financial regulation and risk management, economists still lack a theory of the design of stress scenarios. We model stress testing as a learning mechanism and show how to map the scenario choice problem into an information acquisition choice problem. In this framework, we derive optimal scenarios and characterize how the design of these optimal scenarios depends on the cost of interventions, the prior beliefs of the regulator, the precision of regulatory information, the uncertainty about the risk factors, and the presence of systemic risk factors.

Our approach is consistent with the general principles of current policies implement in various jurisdiction, but it has the advantage that our optimal scenarios are not arbitrary. For example, the current policy on stress scenario design in the U.S. allows for the stress scenarios to "follow either a recession approach, a probabilistic approach, or an approach based on historical experiences."[9] These concepts are somewhat vague and have generated much discussion among banks and regulators. Some commentators argue that scenarios should be predictable while others advocate a flexible design to accommodate emerging risks and changing exposures. Our learning approach shows how to incorporate this goals in the design of the stress scenarios.

Recently, the 2020 supervisory stress tests in the U.S. have been scrutinized for not being, in hindsight, stressful enough, and featuring a severely adverse scenario that was a milder economic catastrophe than the one experienced during the COVID-19 pandemic.[10] These criticisms seem to be grounded in the belief that stress scenarios are predictors of the next crisis. Should this be the case? Should the severely adverse scenario be the worst-case scenario? Or are there other considerations that guide the stress scenario choice? Without pondering the goal of the stress tests and the role of the stress scenarios it is impossible to have a serious discussion of what the

[9] See 12 CFR Part 252 Appendix A.

[10] See, for example, https://www.bloomberg.com/opinion/articles/2020-06-29/banks-are-managing-their-stress.

different scenarios are meant to capture and how they should be chosen. Our paper is a first step in this direction.

References

Acharya, V., R. Engle, and D. Pierret (2014). Testing macroprudential stress tests: The risk of regulatory risk weights. *Journal of Monetary Economics 65*, 36 – 53. Carnegie-Rochester-NYU Conference Series on Public Policy "A Century of Money, Banking, and Financial Instability" held at the Tepper School of Business, Carnegie Mellon University on November 15-16, 2013.

Acharya, V. V., L. H. Pedersen, T. Philippon, and M. Richardson (2016). Measuring systemic risk. *The Review of Financial Studies*.

Allen, F. and D. Gale (2000). Financial contagion. *Journal of Political Economy 108*(1), 1–33.

Bird, A., S. Karolyi, T. Ruchti, and A. Sudbury (2015). Bank regulator bias and the efficacy of stress test disclosures. Working Paper.

Camara, B., P. Pessarossi, and T. Philippon (2016). Backtesting european stress tests. NBER Working Paper.

Candelon, B. and A. Sy (2015). How do markets react to stress tests? IMF Working Paper.

Donald, M., S. Peristiani, and V. Savino (2014). The information value of the stress test. *Journal of Money Credit and Banking 46*(7), 1479–1500.

Faria-e-Castro, M., J. Martinez, and T. Philippon (2017). Runs versus lemons: Information disclosure and fiscal capacity. *Review of Economic Studies 84(4)*, 1683–1707.

Fernandes, M., D. Igan, and M. Pinheiro (2015). March madness in wall street: (what) does the market learn from stress tests? IMF Working Paper.

Flannery, M., B. Hirtle, and A. Kovner (2016, March). Evaluating the information in the federal reserve stress tests. FRBNY Staff Report.

Gick, W. and T. Pausch (2014). Bayesian persuasion by stress test disclosure. Technical report, Deutsche Bundesbank.

Glasserman, P. and G. Tangirala (2016). Are the federal reserve's stress test results predictable? *Journal of Alternative Investments 18*(4).

Goldstein, I. and Y. Leitner (2018). Stress tests and information disclosure. *Journal of Economc Theory 177*(13-26), 34–69.

Goldstein, I. and Y. Leitner (2020). Stress tests disclosure: Theory, practice, and new perspectives. *Handbook of Financial Stress Testing Edited by Farmer, Kleinnijenhuis, Wetzer, and Schuermann, Cambridge University Press.*

Goldstein, I. and H. Sapra (2014, March). Should banks' stress test results be disclosed? an analysis of the costs and benefits. *Foundations and Trends(R) in Finance 8*(1), 1–54.

Hirshleifer, J. (1971). The private and social value of information and the reward to inventive activity. *The American Economic Review 61*(4), pp. 561–574.

Inostroza, N. and A. Pavan (2017). Persuation in global games with an application to stess testing. Technical report, Nortwestern University working paper.

Kyle, A. S. (1989). Informed speculation with imperfect competition. *Review of Economic Studies 56*, 317–356.

Leitner, Y. and B. Williams (2018). Model secrecy and stress tests. *working paper*.

Orlov, D., P. Zryumov, and A. Skrzypacz (2017). Design of macro-prudential stress tests. *working paper*.

Petrella, G. and A. Resti (2013). Supervisors as information producers: Do stress tests reduce bank opaqueness? *Journal of Banking & Finance 37*(12), 5406–5420.

Schuermann, T. (2012). Stress testing banks. Wharton Working Paper.

Schuermann, T. (2016). Stress testing in wartime and in peacetime. Working Paper.

Shapiro, J. and D. Skeie (2015). Information management in banking crises. *Review of Financial Studies 28*(8), 2322–2363.

Van Nieuwerburgh, S. and L. Veldkamp (2010). Information acquisition and under-diversificatio. *The Review of Economic Studies 77(2)*, 779–805.

Veldkamp, L. B. A. B. A. (2009). Information Choice in Macroeconomics and Finance. *Manuscript, New York University*.

Verrecchia, R. E. (1982). Information acquisition in a noisy rational expectations economy. *Econometrica: Journal of the Econometric Society*, 1415–1430.

Vives, X. (2008). *Information and Learning in Markets: the Impact of Market Microstructure*. Princeton University Press.

Appendix

The Appendix contains some auxiliary calculation for formulas in the text. It needs to be completed.

A Proofs

A.1 Learning from stress tests

Proof of Lemma 1

When $N = 1$, $M = 1$, and $J = 2$, the Kalman gain is given by

$$K = \begin{bmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{bmatrix} \begin{bmatrix} \hat{s}_1 \\ \hat{s}_2 \end{bmatrix} \left(\begin{bmatrix} \hat{s}_1 & \hat{s}_2 \end{bmatrix} \left(\begin{bmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{bmatrix} \right) \begin{bmatrix} \hat{s}_1 \\ \hat{s}_2 \end{bmatrix} + \sigma_{\hat{\varepsilon}}^2 (\hat{\mathrm{s}}) \right)^{-1},$$

which implies

$$k_1 = \frac{\sigma_1^2 \hat{s}_1 + \rho\sigma_1\sigma_2\hat{s}_2}{\sigma_1^2\hat{s}_1^2 + 2\rho\sigma_1\sigma_2\hat{s}_1\hat{s}_2 + \sigma_2^2\hat{s}_2^2 + \sigma_{\hat{\epsilon}}^2(\hat{s})}, \quad \text{and}$$
$$k_2 = \frac{\sigma_2^2 \hat{s}_2 + \rho\sigma_1\sigma_2\hat{s}_1}{\sigma_1^2\hat{s}_1^2 + 2\rho\sigma_1\sigma_2\hat{s}_1\hat{s}_2 + \sigma_2^2\hat{s}_2^2 + \sigma_{\hat{\epsilon}}^2(\hat{s})}.$$

Moreover,

$$\Sigma_{\hat{x}}(j,j) = k_j\hat{s}_j \left(\sigma_j^2 + \frac{\hat{s}_h}{\hat{s}_j}\rho\sigma_1\sigma_2 \right),$$

where

$$k_j\hat{s}_j = \frac{\sigma_j^2\hat{s}_j^2 + \rho\sigma_1\sigma_2\hat{s}_2\hat{s}_1}{\sigma_1^2\hat{s}_1^2 + 2\rho\sigma_1\sigma_2\hat{s}_1\hat{s}_2 + \sigma_2^2\hat{s}_2^2 + \sigma_{\hat{\epsilon}}^2(\hat{s})} = 1 - \frac{\rho\sigma_1\sigma_2\hat{s}_1\hat{s}_2 + \sigma_h^2\hat{s}_h^2 + \sigma_{\hat{\epsilon}}^2}{\sigma_1^2\hat{s}_1^2 + 2\rho\sigma_1\sigma_2\hat{s}_1\hat{s}_2 + \sigma_2^2\hat{s}_2^2 + \sigma_{\hat{\epsilon}}^2(\hat{s})}.$$

Note that

$$\frac{\partial k_j}{\partial \sigma_{\hat{\epsilon}}^2(\hat{s})} = -\frac{k_j}{\sigma_1^2\hat{s}_1^2 + 2\rho\sigma_1\sigma_2\hat{s}_1\hat{s}_2 + \sigma_2^2\hat{s}_2^2 + \sigma_{\hat{\epsilon}}^2(\hat{s})} \leq 0$$

and

$$\frac{\partial \Sigma_{\hat{x}}(j,j)}{\partial \sigma_{\hat{\epsilon}}^2(\hat{s})} = -\frac{\Sigma_{\hat{x}}(j,j)}{\sigma_1^2\hat{s}_1^2 + 2\rho\sigma_1\sigma_2\hat{s}_1\hat{s}_2 + \sigma_2^2\hat{s}_2^2 + \sigma_{\hat{\epsilon}}^2(\hat{s})} \leq 0$$

Since

$$\hat{\epsilon}(\hat{s}) = \alpha\epsilon_0 + \beta\left(\|\hat{\mathrm{s}}\|^{\frac{1}{2}} + \|\hat{\mathrm{s}}\|^{1+\theta} \right)\epsilon_1,$$

we have

$$\sigma_{\hat{\epsilon}}^2(\hat{s}) = \alpha^2 + \beta^2\left(\left(\hat{s}_1^2 + \hat{s}_2^2\right)^{\frac{1}{2}} + \left(\hat{s}_1^2 + \hat{s}_2^2\right)^{1+\theta} \right),$$

which is increasing in $|s_j|$ for $j = 1, 2$. Therefore, more extreme scenarios decrease the amount of learning. The effect of an increase in noise on the amount of learning is negligible close to the baseline, i.e.,

$$\lim_{|\hat{s}_j|\to 0} \frac{\partial \Sigma_{\hat{x}}(j,j)}{\partial \sigma^2_{\hat{\epsilon}}(\hat{s})} \frac{\partial \sigma^2_{\hat{\epsilon}}(\hat{s})}{\partial |\hat{s}_j|} = -\frac{\Sigma_{\hat{x}}(j,j)\,\beta^2 \left(\left(\hat{s}_1^2 + \hat{s}_2^2\right)^{-\frac{1}{2}} + (1+\theta)\left(\hat{s}_1^2 + \hat{s}_2^2\right)^{\theta} 2 \right) |\hat{s}_j|}{\sigma_1^2 \hat{s}_1^2 + 2\rho\sigma_1\sigma_2\hat{s}_1\hat{s}_2 + \sigma_2^2\hat{s}_2^2 + \sigma^2_{\hat{\epsilon}}(\hat{s})} = 0 \quad \forall j = 1, 2. \tag{A.1}$$

Moreover, the direct effect of a more extreme scenario on the amount of learning is given by

$$\frac{\partial \Sigma_{\hat{x}}(j,j)}{\partial |\hat{s}_j|} = \frac{\partial (k_j \hat{s}_j)}{\partial |\hat{s}_j|}\sigma_j^2 + \frac{\partial k_j}{\partial |\hat{s}_j|}\hat{s}_h \rho\sigma_1\sigma_2 \quad \forall j, h = 1, 2,\ j \neq h.$$

When $\rho = 0$,we have

$$\frac{\partial \Sigma_{\hat{x}}(j,j)}{\partial |\hat{s}_j|} = \frac{\partial (k_j \hat{s}_j)}{\partial |\hat{s}_j|}\sigma_j^2 = \frac{\sigma_h^2 \hat{s}_h^2 + \sigma^2_{\hat{\epsilon}}}{\left(\sigma_1^2\hat{s}_1^2 + \sigma_2^2\hat{s}_2^2 + \sigma^2_{\hat{\epsilon}}(\hat{s})\right)^2}\sigma_j^4 2\,|\hat{s}_j| \geq 0.$$

Then, using Equation A.1 we have that

$$\Sigma_{\hat{x}}(j,j) = \frac{\sigma_j^2 \hat{s}_j^2}{\sigma_1^2\hat{s}_1^2 + \sigma_2^2\hat{s}_2^2 + \alpha^2 + \beta^2\left(\hat{s}_1^2 + \hat{s}_2^2\right)^{1+\theta}}$$

$$\begin{aligned}
\frac{d\Sigma_{\hat{x}}(j,j)}{d\hat{s}_j^2} &= \frac{\sigma_j^2\left(\sigma_h^2\hat{s}_h^2 + \alpha^2 + \beta^2\left(\hat{s}_1^2 + \hat{s}_2^2\right)^{1+\theta}\right) - \sigma_j^2\hat{s}_j^2\beta^2(1+\theta)\left(\hat{s}_1^2 + \hat{s}_2^2\right)^{\theta}}{\left(\sigma_1^2\hat{s}_1^2 + \sigma_2^2\hat{s}_2^2 + \alpha^2 + \beta^2\left(\hat{s}_1^2 + \hat{s}_2^2\right)^{1+\theta}\right)^2} \\
&= \frac{\sigma_j^2\left(\frac{\sigma_h^2\hat{s}_h^2 + \alpha^2}{\beta^2\left(\hat{s}_1^2 + \hat{s}_2^2\right)^{1+\theta}} + 1\right) - \frac{\sigma_j^2\hat{s}_j^2(1+\theta)}{\left(\hat{s}_1^2 + \hat{s}_2^2\right)}}{\left(\frac{\sigma_1^2\hat{s}_1^2 + \sigma_2^2\hat{s}_2^2 + \alpha^2}{\beta^2\left(\hat{s}_1^2 + \hat{s}_2^2\right)^{1+\theta}} + 1\right)^2}
\end{aligned}$$

$$\lim_{\hat{s}_j\to 0}\frac{d\Sigma_{\hat{x}}(j,j)}{d\hat{s}_j^2} > 0 \quad \text{and} \quad \lim_{\hat{s}_j\to\infty}\frac{d\Sigma_{\hat{x}}(j,j)}{d\hat{s}_j^2} = -\theta\sigma_j^2 < 0$$

When $\rho > 0$, using Equation (15) and the definition of $\sigma^2_{\hat{\epsilon}}(\hat{s})$, we have that

$$k_1\hat{s}_1 = \frac{\sigma_1^2\hat{s}_1^2 + \rho\sigma_1\sigma_2\hat{s}_2\hat{s}_1}{\sigma_1^2\hat{s}_1^2 + 2\rho\sigma_1\sigma_2\hat{s}_1\hat{s}_2 + \sigma_2^2\hat{s}_2^2 + \alpha^2 + \beta^2\left(\hat{s}_1^2 + \hat{s}_2^2\right)^{1+\theta}}$$

and

$$\Sigma_{\hat{x}}(j,j) = \frac{\sigma_j^2\hat{s}_j^2 + \rho\sigma_1\sigma_2\hat{s}_2\hat{s}_1 + \rho\sigma_j^2\hat{s}_j\hat{s}_h + (\rho\sigma_1\sigma_2)^2\hat{s}_h^2}{\sigma_1^2\hat{s}_1^2 + 2\rho\sigma_1\sigma_2\hat{s}_1\hat{s}_2 + \sigma_2^2\hat{s}_2^2 + \alpha^2 + \beta^2\left(\hat{s}_1^2 + \hat{s}_2^2\right)^{1+\theta}}.$$

Proof of Proposition 2

The Kalman filter in Equations (10) and (11) imposes restrictions on the set of posterior variances that can be attained. More specifically, the Kalman filter maps the $J \times M$ elements in the stress scenarios $\hat{S}$ to the the $\frac{NJ(NJ+1)}{2}$ elements of the posterior precision $\hat{\Sigma}_{\mathbf{x}}$ from the set $\mathbf{\Sigma}$. Moreover, the regulator's choice will always be on the frontier of the feasibility set, given by $\overline{\mathbf{\Sigma}} \equiv \left\{ K\left(\hat{S}_\omega\right) \hat{S}'_\omega \Sigma_x \quad \text{for all} \quad \omega \in [0,1]^{N\times J} \quad \text{with} \quad \sum_{h=1}^{N\times J} \omega_h \right\}$, where

$$\hat{S}_\omega \equiv \arg\max_{\hat{S}} \omega' K\left(\hat{S}\right) \hat{S}' \Sigma_x \omega.$$

Given our assumptions on $\hat{\epsilon}_i\left(\hat{\mathrm{s}}, M\right)$, $\hat{S}_\omega$ is unique. Hence, since the objective function of the regulator depends on $\hat{S}$ only through the posterior variance, the regulator's scenario choice problem can be cast in term of choosing $\hat{\Sigma}_{\mathbf{x}}$.

When $M < N$,as long as all risk dimensions are spanned, choosing the $J \times M$ elements in $\hat{S}$ is equivalent to choosing $J \times M$ elements of the posterior precision $\hat{\Sigma}_x$. Without loss of generality, one can focus on the posterior variances of the risk exposures of M banks from the set $\overline{\mathbf{\Sigma}}_M \equiv \left\{ K\left(\hat{S}_\omega\right) \hat{S}'_\omega \Sigma_x \quad \text{for all} \quad \omega \in [0,1]^{N\times J} \quad \text{with} \quad \sum_{m=1}^{M} \sum_{j=1}^{J} \omega_{(j-1)J+m} \right\}$.

A.2 Taking action

Proof of Lemma 2

Under linear quadratic preferences, the first order condition that characterizes the optimal capital requirement is

$$\kappa'\left(\Delta^\star\right) = 1 - \gamma \left(\bar{W} + \Delta^\star + \mathbb{E}\left[\mathrm{s}\right] \cdot \left(\sum_{i=1}^{N} \left(\mathbf{1} - \mathrm{a}_i\right) \circ \hat{\mathrm{x}}_i \right) \right).$$

Proof of Lemma 3

Under linear quadratic preferences, the first order condition that characterizes the regulator's optimal targeted intervention policy is

$$\phi_j a^\star_{i,j}\left(\hat{y}\right) = \mathbb{E}\left[x_{i,j} s_j \left(1 - \gamma W\right) \middle| \hat{y}\right] \quad \forall i = 1, ..., N, \quad \forall j = 1, ..., J,$$

where $W = \sum_{i=1}^{N} \left(\sum_{j=1}^{J} \left(1 - a^\star_{i,j}\left(\hat{y}\right)\right) x_{i,j} s_j + \eta_i - d_i \right)$. This is the same as

$$\phi_j a^\star_{i,j}\left(\hat{y}\right) = \mathbb{E}\left[x_{i,j} s_j \left(1 + \gamma \sum_{n=1}^{N} \sum_{h=1}^{J} \left(1 - a^\star_{n,h}\left(\hat{y}\right)\right) x_{n,h} s_h + \gamma \sum_{i=1}^{N} d_i \right) \middle| \hat{y} \right].$$

Since $\mathbb{E}\left[s_j\right] = 0$ for all j, we have

$$\phi_j a^\star_{i,j}\left(\hat{y}\right) = \gamma \sum_{n=1}^{N} \left(1 - a^\star_{n,j}\left(\hat{y}\right)\right) \mathbb{E}\left[x_{i,j} x_{n,j} \middle| \hat{y}\right] \mathbb{E}\left[s_j^2\right].$$

Let $a_j^\star = \left(a_{1,j}^\star, \ldots, a_{N,j}^\star\right)'$. Then, the FOCs can be written as

$$\phi_j a_j^\star = -\gamma \begin{bmatrix} \mathbb{E}\left[x_{1,j}x_{1,j}|\,\hat{y}\right] & \mathbb{E}\left[x_{1,j}x_{2,j}|\,\hat{y}\right] & \ldots & \mathbb{E}\left[x_{1,j}x_{N,j}|\,\hat{y}\right] \\ \mathbb{E}\left[x_{2,j}x_{1,j}|\,\hat{y}\right] & \ddots & \cdots & \vdots \\ \vdots & \cdots & \cdots & \vdots \\ \mathbb{E}\left[x_{N,j}x_{1,j}|\,\hat{y}\right] & \cdots & \cdots & \mathbb{E}\left[x_{N,j}x_{N,j}|\,\hat{y}\right] \end{bmatrix} \mathbb{E}\left[s_j^2\right] a_j^\star + \begin{bmatrix} \gamma \sum_{n=1}^N \mathbb{E}\left[x_{1,j}x_{n,j}|\,\hat{y}\right] \mathbb{E}\left[s_j^2\right] \\ \gamma \sum_{n=1}^N \mathbb{E}\left[x_{2,j}x_{n,j}|\,\hat{y}\right] \mathbb{E}\left[s_j^2\right] \\ \vdots \\ \gamma \sum_{n=1}^N \mathbb{E}\left[x_{N,j}x_{n,j}|\,\hat{y}\right] \mathbb{E}\left[s_j^2\right] \end{bmatrix},$$

which is the same as

$$\phi_j a_j^\star = -\left(\hat{\Sigma}_{x,j} + \hat{x}_j\hat{x}_j'\right)\gamma\mathbb{E}\left[s_j^2\right] a_j^\star + \left(\hat{\Sigma}_{x,j} + \hat{x}_j\hat{x}_j'\right)\mathbf{1}_{N\times 1}\gamma\mathbb{E}\left[s_j^2\right],$$

where $\hat{x}_j = (\hat{x}_{i,j}, \ldots, \hat{x}_{N,j})'$. Then,

$$a_j^\star = \left(\phi_j \mathbb{I}_N + \gamma\left(\hat{\Sigma}_{x,j} + \hat{x}_j\hat{x}_j'\right)\mathbb{E}\left[s_j^2\right]\right)^{-1}\gamma\left(\hat{\Sigma}_{x,j} + \hat{x}_j\hat{x}_j'\right)\mathbf{1}_{N\times 1}\mathbb{E}\left[s_j^2\right] \quad \forall j = 1, \ldots, J \tag{A.2}$$

Proof of Proposition 3

It follows directly from the system in Equation (A.2).

A.3 Optimal scenario choice

Proof of Proposition 4

When utility is linear quadratic and intervention costs are linear, the objective function is

$$O = \mathbb{E}_{\hat{\mathbf{x}}}\left[\mathbb{E}_{s,\eta,x}\left[U\left(W\left(\mathbf{a}^*, \Delta^*\right)\right)|\hat{\mathbf{x}}, \hat{\Sigma}_x\right] - \Phi\left(\mathbf{a}^*\right)\right] = \mathbb{E}_{\hat{\mathbf{x}}}\left[\mathbb{E}_{s,\eta,x}\left[W\left(\mathbf{a}^*, \Delta^*\right) + \frac{\gamma}{2}\left(W\left(\mathbf{a}^*, \Delta^*\right)\right)^2|\hat{\mathbf{x}}, \hat{\Sigma}_x\right] - \Phi\left(\mathbf{a}^*\right)\right]$$

$$\begin{aligned}
O &= \bar{W} - \frac{\gamma}{2}\bar{W}^2 + \mathbb{E}_{\hat{\mathbf{x}}}\left[\Delta\right] - \frac{\gamma}{2}\sigma_\eta^2 - \mathbb{E}_{\hat{\mathbf{x}}}\left[\mathbb{E}_{s,x}\left[\frac{\gamma}{2}\left(\Delta + \sum_{i=1}^N\sum_{j=1}^J\left(1 - a_{i,j}^\star\left(\hat{\mathbf{x}}\right)\right)x_{i,j}s_j\right)^2|\hat{\mathbf{x}}, \hat{\Sigma}_x\right] + \frac{1}{2}\sum_j \phi_j\left(a_{i,j}^\star\left(\hat{\mathbf{x}}\right)\right)^2\right] \\
&= \bar{W} - \frac{\gamma}{2}\bar{W}^2 + \mathbb{E}_{\hat{\mathbf{x}}}\left[\Delta\right] - \frac{\gamma}{2}\sigma_\eta^2 - \mathbb{E}_{\hat{\mathbf{x}}}\left[\mathbb{E}_{s,x}\left[\frac{\gamma}{2}\left(\Delta + \sum_{i=1}^N\sum_{j=1}^J\left(1 - a_{i,j}^\star\left(\hat{\mathbf{x}}\right)\right)x_{i,j}s_j\right)^2|\hat{\mathbf{x}}, \hat{\Sigma}_x\right] + \frac{1}{2}\sum_j \phi_j\left(a_{i,j}^\star\left(\hat{\mathbf{x}}\right)\right)^2\right] \\
&= \bar{W} - \frac{\gamma}{2}\bar{W}^2 + \mathbb{E}_{\hat{\mathbf{x}}}\left[\Delta - \frac{\gamma}{2}\Delta^2\right] - \frac{\gamma}{2}\sigma_\eta^2 - \mathbb{E}_{\hat{\mathbf{x}}}\left[\frac{1}{2}\sum_{j=1}^J\sum_{i=1}^N\left(1 - a_{i,j}^\star\left(\hat{\mathbf{x}}\right)\right)\phi_{i,j}\left(a_{i,j}^\star\left(\hat{\mathbf{x}}\right)\right) + \frac{1}{2}\sum_j \phi_j\left(a_{i,j}^\star\left(\hat{\mathbf{x}}\right)\right)^2\right] \\
&= \bar{W} - \frac{\gamma}{2}\bar{W}^2 + \mathbb{E}_{\hat{\mathbf{x}}}\left[\Delta - \frac{\gamma}{2}\Delta^2\right] - \frac{\gamma}{2}\sigma_\eta^2 - \mathbb{E}_{\hat{\mathbf{x}}}\left[\frac{1}{2}\sum_{j=1}^J\sum_{i=1}^N \phi_{i,j}a_{i,j}^\star\left(\hat{\mathbf{x}}\right)\right]
\end{aligned}$$

where we used that

$$\phi_j a_{i,j}^\star\left(\hat{\mathbf{x}}\right) = \gamma\sum_{k=1}^N\left(1 - a_{k,j}^\star\left(\hat{\mathbf{x}}\right)\right)\mathbb{E}\left[x_{i,j}x_{k,j}|\hat{\mathbf{x}}, \hat{\Sigma}_x\right]\mathbb{E}\left[s_j^2\right].$$

Moreover, since $\Delta^\star$ is given by

$$\kappa'\left(\Delta^\star\right) = 1 - \gamma\left(\bar{W} + \Delta^\star\right) \ ,$$

it is independent of the information revealed by the stress test.

Hence, the regulator's scenario choice problem is equivalent to

$$\min_{\hat{\Sigma}_x \in \mathbf{\Sigma}} \mathbb{E}_z \left[\frac{1}{2} \sum_{j=1}^{J} \sum_{i=1}^{N} \phi_j a_{i,j}^\star \left(\overline{\mathbf{x}} + \left(\Sigma_x - \hat{\Sigma}_x \right)^{\frac{1}{2}} \boldsymbol{z} \right) \right],$$

where

$$a_j^\star = \left(\phi_j \mathbb{I}_N + \gamma \mathbb{E}\left[x^j x^{j\prime} \middle| \hat{y} \right] \mathbb{E}\left[s_j^2 \right] \right)^{-1} \gamma \mathbb{E}\left[x^j x^{j\prime} \middle| \hat{y} \right] \mathbf{1}_{N\times 1} \mathbb{E}\left[s_j^2 \right] \quad \forall j = 1, \ldots, J$$

and $\hat{x}^j = [x_{1,j}, ..., x_{N,j}]'$.

B Additional comparative statics

B.1 Uncertainty about risk factors

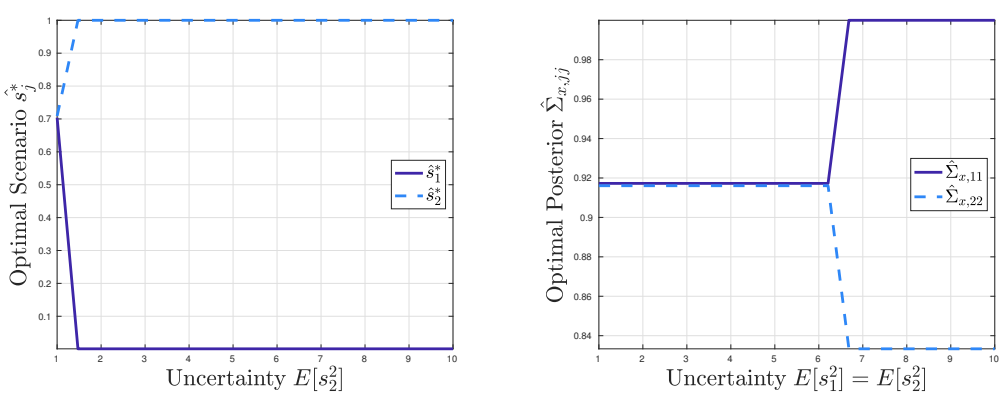

(a) Optimal scenario choice (b) Optimal posterior

Figure A.1: Optimal scenario and information choice as a function of the uncertainty about risk factor 2, $\mathbb{E}\left[s_2^2\right]$..

Note: Figure A.1 illustrates the regulator's optimal choice of scenario and the implied posterior variance as a function of the uncertainty about risk factor 2, $\mathbb{E}\left[s_2^2\right]$. The parameters used are $N = 1$, $J = 2$, $\gamma = 0.3$, $\phi_i = 1$, $\bar{x}$=[1, 1], $\Sigma_x = \boldsymbol{I}_J$, $\alpha_1(M) = 0.8M$, $\beta_1 = 1$, $M = 1$, and $\theta = 1$, $\mathbb{E}\left[s_k^2\right] = 1$, $\mathbb{E}\left[\epsilon_{1,0}^2\right] = 1$ and $\mathbb{E}\left[\epsilon_{1,1}^2\right] = 1$

If one risk factor has a very low variance and will stay close to the the baseline, then it is less valuable to learn about the exposures to it and to intervene to reduce them. In this case, the factor's weight on the expected losses will be small and uncertainty about the exposure to it is less costly. However, if the variance of a risk factor is large it has the potential to be an important driver of bank losses depending on the risk exposures to it. In this case, the regulator has more incentives to learn and intervene along the dimension of this factor to curve its potential impact on losses. Therefore, the regulator will stress a risk factor more in the optimal scenario the highest the uncertainty about it. Figures A.1 show the optimal scenario choice as a function of the uncertainty about risk factor 2, $\mathbb{E}\left[s_2^2\right]$.

B.2 Increased overall uncertainty

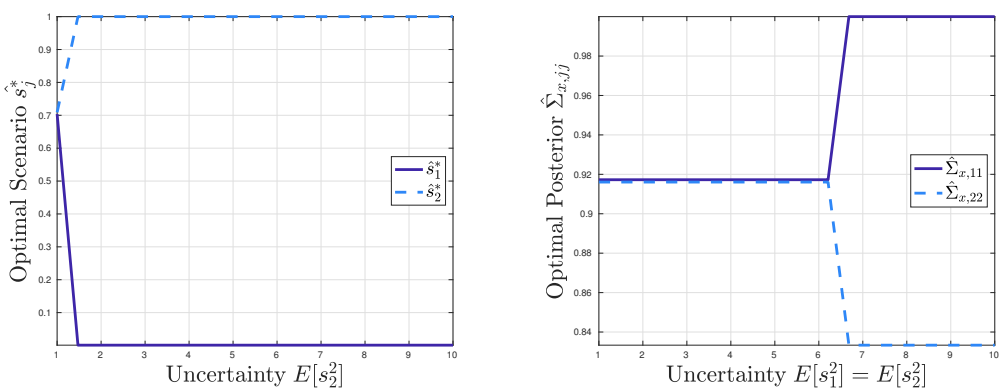

(a) Optimal scenario choice (b) Optimal posterior

Figure A.2: Optimal scenario and information choice as a function of the uncertainty about risk factor 2, $\mathbb{E}\left[s_2^2\right]$..

Note: Figure A.2 illustrates the regulator's optimal choice of scenario and the implied posterior variance as a function of the uncertainty about risk factor 2, $\mathbb{E}\left[s_2^2\right]$. The parameters used are $N = 1$, $J = 2$, $\gamma = 0.3$, $\phi_i = 1$, $\bar{x}$=[1, 1], $\Sigma_x = \boldsymbol{I}_J$, $\alpha_1(M) = 0.8M$, $\beta_1 = 1$, $M = 1$, and $\theta = 1$, $\mathbb{E}\left[s_k^2\right] = 1$, $\mathbb{E}\left[\epsilon_{1,0}^2\right] = 1$ and $\mathbb{E}\left[\epsilon_{1,1}^2\right] = 1$