

Stress Tests and Bank Portfolio Choice*

Basil Williams

New York University

March 2017

Preliminary and Incomplete

Abstract

How informative should bank stress tests be? I use Bayesian persuasion to formalize stress tests and show that regulators can reduce the likelihood of a bank run by performing tests which are only partially informative. Optimal stress tests give just enough failing grades to keep passing grades credible enough to avoid runs. The worse the state of the banking system, the more stringent stress tests must be to prevent runs. I find that optimal stress tests, by reducing the probability of runs, reduce the optimal level of banks' liquidity cushions. I also examine the impact of anticipated stress tests on banks' ex ante incentive to invest in risky versus safe assets.

JEL Classification:

Keywords:

*I thank my advisor S. Viswanathan and Felipe Varas for invaluable feedback and encouragement. I also thank Adriano Rampini for many helpful comments. Address: New York University, Department of Economics, 19 W 4th St., New York, NY 10003. Email: basil.williams@nyu.edu.

1 Introduction

It is widely presumed that investors are better off with more information than less. Much regulation, such as the recent Dodd-Frank Act, is structured in order to help investors become more informed. In the spirit of maximal information disclosure, many presume that stress tests designed to disclose the solvency of banks should be as informative as possible. For example, a recent article in the *Wall Street Journal* predicts that stress tests performed by the European Central Bank may resolve investor uncertainty about European banks, but with a caveat:

If, that is, the tests are credible... The last tests under the sole supervision of the European Banking Authority were risible; several banks passed them only to require bailouts later.¹

The quote laments the inaccuracy of the stress tests, implying that more informative tests are better for investors.

This paper questions the presumption that more informative stress tests improve investor welfare. I show that partially informative stress tests may in fact be optimal. Stress tests only need to be credible enough to prevent runs, so even stress tests with some likelihood of issuing a passing grade to bad banks may be optimal.

I explore the structure of optimal stress tests using a model of bank runs which is most similar to that presented by Allen and Gale (1998). A representative bank issues demand deposit contracts to consumers in exchange for deposits. Banks must choose the promised deposit return as well as the proportion of deposits allocated to safe liquid assets and risky illiquid assets. I assume that at the time the consumers deposit their funds at the bank, they are uncertain about their preferred timing of future consumption and that they are intertemporally risk averse. The representative bank offers a deposit contract which acts as intertemporal insurance by smoothing consumption across future dates. At the second date,

¹Davies, Paul, (2014, Oct. 22) The real reason to stress about Europe's bank test, *Wall Street Journal*

depositors learn whether they prefer to consume early or late, and they also receive a public signal about the future return of the risky asset. If this signal indicates a low enough future return, then patient investors prefer to withdraw early and the bank is forced to liquidate all assets, which is costly.

One important difference between my model and that of Allen and Gale is that the public signal about the risky asset does not resolve all future uncertainty. The remaining uncertainty creates scope for additional disclosure via a stress test, which I formalize using the Bayesian persuasion framework of Kamenica and Gentzkow (2011). Under this framework, stress tests are conceptualized as a distribution of signals, conditional on the true future return of the risky asset. Both the structure and result of the stress tests are public information, and the test is costless to perform, so neither information asymmetries nor auditing costs exist as frictions in the model.

Despite the frictionless nature of the tests, I still find that partially informative tests are optimal. If the original public signal is low enough to induce a run, regulators may improve consumer welfare by performing a test which reduces the likelihood of runs. It is optimal to design the test so that all good banks and some bad banks pass, but some bad banks fail. In particular, it is necessary to fail only just enough banks to maintain the credibility of the passing grade, so that a passing grade induces a high enough posterior belief in the risky asset's return that patient investors do not run. A test of this form improves welfare, because without it, the original public signal was such that a run would have occurred with certainty, so the bank would have been forced to liquidate the risky asset at a discount. Under the optimal test, however, passing grades result in no-run equilibria, so the probability of costly liquidation is reduced.

I also find that the lower the initial public confidence in the risky assets, the more stringent the optimal stress tests must be. In particular, in times of very low depositor confidence in the bank, bad banks should receive passing grades with less and less likelihood. Intuitively, low public confidence in the banking sector requires a very strong signal of quality in order

to prevent runs. The harder a test is to pass, the stronger the signal of quality given by a passing grade, and the resulting improved confidence is then sufficient to prevent runs.

The paper also examines the relationship between well-designed stress tests and the bank's optimal portfolio and contract, showing that stress tests and liquidity buffers may be seen as policy substitutes for preventing runs. Specifically, because optimal stress tests reduce the likelihood of runs, the need for a liquidity cushion for run prevention is not as high as it would be otherwise. In addition, well-designed stress tests enable banks to offer a higher interest rate to early withdrawers. The reason is that without stress tests, raising the promised deposit return to early withdrawers may raise the probability of runs by making early withdrawal more attractive, whereas optimal stress tests reduce that risk by lowering the likelihood of runs, so the optimal deposit return rises.

My model builds on the theory of bank runs, which began most notably with the seminal paper of Diamond and Dybvig (1983), who showed how banks subject to runs can attract deposits by offering intertemporal risk sharing. Allen and Gale (1998) build on their model by showing how runs may be induced by exogenous signals about fundamentals, and allow banks to decide how much to invest in safe vs risky assets. My model is closely related to the third section of that paper in which the risky asset may be liquidated at a discount, but unlike their paper I assume an exogenous, linear liquidation technology in order to simplify my analysis. The most crucial difference between my paper and theirs is that in my model, uncertainty is not fully resolved until the final date; this creates scope for informative stress tests, which their paper does not consider.

I am aware of three papers which apply the Bayesian persuasion framework of Kamenica and Gentzkow (2011) to the optimal design of bank stress tests: Gick and Pausch (2012), Goldstein and Leitner (2013), and Faria-e-Castro, Philippon, Martinez (2015). The most significant feature which distinguishes my paper from these papers is my examination of the impact of anticipated stress tests on the bank's asset portfolio, deposit return, and liquidity cushion, which they do not consider. Several other differences exist. Faria-e-Castro, et. al

focus on providing a joint theory of disclosure and government fiscal backstops. Gick and Pausch assume an exogenous binary form for investor utility, which depends by assumption on the binary quality of the bank and a binary action of the investor. In my model, however, the utility of the investors is determined endogenously from the payoff of the demand deposit contract offered by the bank and the late investor's decision to withdraw early. They also assume an exogenous form for the regulator's utility which depends directly on the mass of investors which take a prudent action, whereas my regulator's objective is simply the endogenous utility of the investors.

The main value of disclosure in Goldstein and Leitner (2013) is that banks can engage in risk-sharing arrangements by pooling bad banks with good banks, in order to protect against a fall in future capital. In my model, bad banks may pool with good banks, but the value of pooling is so that endogenous runs occur less often. Their paper assumes that the bank's return on cash is a discontinuous function, whereas the discontinuity in my model is explicitly modeled as the discount suffered from early liquidation during an endogenous run. They examine a case both with and without bank private information, and they find that private information may necessitate more than two disclosure scores; in my model, there is no private information, so two disclosure scores are sufficient.

Parlatore (2015) considers the effect of bank transparency on depositors' incentives to run and on the bank's ex ante contract and portfolio decision. She explores the effects of varying levels of transparency, which is given by a single parameter controlling the probability that a depositor's signal matches the true underlying state. In contrast, I consider both the level and the form of transparency, given by a flexible structure in which the probability of receiving an accurate signal may differ across banks of different qualities. In her model, increasing transparency monotonically decreases welfare because it increases the likelihood of runs, so zero transparency is always an optimal level. However, I show that if prior depositor beliefs are bad enough, some transparency may improve welfare by identifying good banks and deterring unnecessary runs. Bouvard, Chaigneau, and de Motta (2015) consider optimal

bank transparency, but focus on the commitment problem of a regulator who strategically interacts with depositors, ignoring the bank's incentives.

Section 2 presents the environment and solves for the optimal consumption profile given some particular portfolio and contract terms. Section 3 finds the structure of optimal stress tests and explains how they can improve consumer welfare. Section 4 examines the impact of anticipated stress tests on the bank's ex ante portfolio decision and contract terms, and Section 5 concludes.

2 Environment

The environment consists of two periods, or three dates: 0, 1, and 2. In this environment, there is a consumption good, and two types of assets which transform date 0 consumption to date 2 consumption. The first type of asset is safe and liquid: it gives a certain return of one at date 2 and may be liquidated at date 1 without cost. The second type of asset is risky and illiquid: it gives a risky return R at date 2 and may be liquidated at date 1 only by selling it for a fraction $\theta \in (0, 1)$ of its expected return. I assume the date 2 return on the risky asset takes one of two values: in a good state, the return is $G > 0$, and in a bad state, the return is $B = 0$. I assume that at date 0, the expected return on the risky asset is greater than 1, creating an incentive to invest in the risky asset.

There is a continuum of ex ante identical consumers, with measure normalized to 1. At date 0, consumers have total endowment E of the consumption good, and are uncertain about their future consumption preferences; some will be impatient, preferring to consume at date 1, and the rest will be patient, being indifferent between date 1 and date 2 consumption. At date 0, consumers know only the probabilities of being patient or impatient, which I assume are equal. At date 1, consumers learn whether they are patient or not, and also receive a public signal $p \sim \text{unif}(0, 1)$, which indicates the probability that the risky asset gives a good return G at date 2.

In the spirit of Epstein and Zin (1989), I distinguish between the consumer's desire to smooth consumption across dates and across states. In particular, I assume that at any particular date, consumers are risk neutral, but they desire to smooth consumption across dates. Let s_t denote the state at date t , which may include the public signal p , the impatience of each consumer, or the actual realized return R on the risky illiquid asset. Let $c_t(s_t)$ be a consumer's date t consumption, contingent on the state s_t , and let $\beta_i(s_1) \in \{0, 1\}$ be the discount factor applied to date 2 consumption by consumer i . At date 1, an impatient investor's utility is given by $u(c_1(s_1))$, and a patient investor's utility is given by $u(c_1(s_1)) + u(E[c_2(s_2)|s_1])$, where $u(c) = \min[c, k]$. At date 0, before any uncertainty is resolved, a consumer's utility is given by

$$E[u(c_1(s_1)) + \beta_i(s_1)u(E[c_2(s_2)|s_1])].$$

These preferences are broadly similar to Epstein-Zin preferences, where here the expectation of date 2 consumption is analagous to a risk-neutral Epstein-Zin certainty equivalent operator, and if the output of this operator is denoted by z , then the function $W(c, z) = u(c) + \beta u(z)$ is analagous to the Epstein-Zin intertemporal aggregator. The concavity of u creates incentives to smooth consumption between dates 1 and 2. So consumers will prefer consumption profiles that offer similar levels of consumption to patient and impatient investors. That is, they desire to insure against intertemporal risk, which creates a role for banks.

Banks may make investments on behalf of consumers. They pool consumers' assets, and thereby offer insurance to consumers against uncertain liquidity demands. Free entry forces banks to compete, so they offer deposit contracts which maximize the utility of consumers, and banks fund those contracts by investing deposits into the safe liquid asset and the risky illiquid asset. I assume that the demand deposit contract takes the following form: consumers who withdraw at date 1 receive fixed return $r > 0$ if the bank can afford to pay it; if not, the bank is forced to liquidate all assets and distribute them equally among early withdrawers.

Late withdrawers simply receive the remainder of the banks' assets at date 2, which is zero if the bank is forced to liquidate and distribute all assets at date 1.

I can summarize the timing as follows. At date 0, consumers invest their deposits with banks. Banks invest those deposits in the two assets, allocating $L \geq 0$ to the liquid safe asset and $I \geq 0$ to the illiquid risky asset, and promising early withdrawers a fixed return $r > 0$ if affordable. At date 1, consumers privately learn whether they are patient or impatient and receive a public signal $p \sim \text{unif}[0, 1]$ indicating the probability that the risky asset will yield a good return G at date 2. Contingent on this information, patient consumers decide whether to withdraw early or late, and all early withdrawers approach the bank requesting return r . If the bank can afford to pay r to all early withdrawers, it does so, and then at date 2 the consumers receive the bank's remaining assets. If the bank cannot afford to pay r to all early withdrawers, it is forced to liquidate all assets and distribute them across all early withdrawers.

2.1 Optimal Consumption Profile

At date 0, the bank receives deposits E from consumers and must decide how to invest those deposits and what consumption profile to offer consumers. Because the bank competes with other banks for the business of consumers, it chooses its portfolio and offered contract to maximize the expected utility $E[u(c_{j1}(p)) + \beta_j u(E[c_{j2}(R, p)|p])]$ of consumers, where $j = 1$ denotes an impatient consumer who prefers to consume at date 1, and $j = 2$ denotes a patient consumer who is indifferent between early and late consumption, so then $\beta_1 = 0$ and $\beta_2 = 1$.

It is helpful to consider the bank's problem using backward induction. At date 1, the bank has already selected L liquid assets and I illiquid assets, and has promised return r . In addition, consumers privately know their patience level, and all have belief p that the risky asset will yield a good return G . Given this information, what is the highest aggregate consumption a bank can offer consumers? First, the consumption profile must be feasible,

so early withdrawers cannot receive any more consumption than the total value of the early liquidated assets. Denoting the amount of liquidated liquid and illiquid assets by $l \in \{0, 1\}$ and $i \in \{0, 1\}$, respectively, and the total mass of patient consumers who withdraw early (run) by α , I have

$$c_{11} + \alpha c_{21} \leq lL + i\theta E[R|p]I. \quad (1)$$

Similarly, the late withdrawing investors' consumption cannot exceed the total remaining assets left over at date 2, which consists of the unliquidated assets plus the unwithdrawn liquidated assets. Taking expectations gives:

$$\begin{aligned} (1 - \alpha)E[c_{22}|p] &\leq \left((1 - l)L + (1 - i)E[R|p]I \right) + \left(lL + i\theta E[R|p]I - c_{11} - \alpha c_{21} \right) \\ &= L + \left(i\theta + (1 - i) \right) E[R|p]I - c_{11} - \alpha c_{21}. \end{aligned} \quad (2)$$

The deposit contract says that early withdrawers can receive no more than the promised fixed return $r > 0$, and that if the bank cannot afford to pay this, it is forced to liquidate all of its assets and distribute them equally among early withdrawers. So the deposit contract gives the following two conditions:

$$\max[c_{11}, c_{21}] \leq r \quad (3)$$

If $\min[c_{11}, c_{21}] < r$, then $i = l = 1$ and

$$c_{11} + \alpha c_{21} = L + \theta E[R|p]I. \quad (4)$$

Next, note that because investor patience is privately known, all early withdrawers must be treated equally, so

$$\text{if } \alpha > 0, \text{ then } c_{11} = c_{21}. \quad (5)$$

Lastly, I consider the patient investor's incentive to run, which occurs only when he expects

to do at least as well by running as by waiting.

$$\max[u(c_{11}), u(c_{21})] > u(E[c_{22}|p]) \implies \alpha = 1 \quad (6)$$

$$\max[u(c_{11}), u(c_{21})] < u(E[c_{22}|p]) \implies \alpha = 0 \quad (7)$$

At date 1, then, the bank's problem is to maximize aggregate utility

$$u(c_{11}(p)) + \left(\alpha(p)u(c_{21}(p)) + (1 - \alpha(p))u(E[c_{22}(R, p)|p]) \right)$$

subject to constraints (1) through (7). Given L , I , r , and p , the optimal consumption profile is the solution to this problem, and is given in the following proposition.

Proposition 1 (Optimal Consumption Profile). *Given L , I , r , and p , let p^* be the lowest belief p that does not induce a run:*

$$p^* = \inf\{p \in \mathbb{R} : u(r) \leq u(L - r + E[R|p]I)\}.$$

(i) *If $p \geq p^*$, then*

- *The liquid asset is liquidated, the illiquid asset is not, and all patient investors withdraw late: $l = 1$, $i = 0$, and $\alpha(p) = 0$.*
- *Impatient investors receive the fixed promised payment, and patient investors receive the remaining assets: $c_{11}(p) = r$ and $c_{22}(R, p) = L - r + RI$.*

(ii) *If $p < p^*$, then*

- *Both assets are liquidated and all patient investors withdraw early: $l = i = 1$ and $\alpha(p) = 1$.*
- *All investors receive an equal portion of the total assets, which portion is strictly less than fixed return r : $c_{11}(p) = c_{22}(p) = (L + \theta E[R|p]I)/2 < r$.*

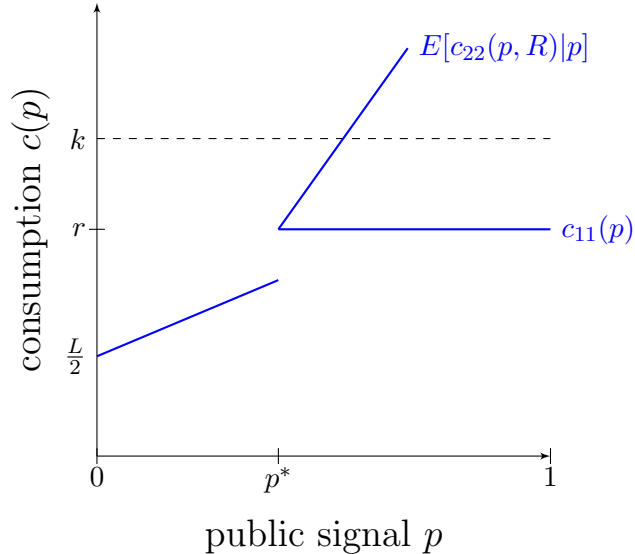


Figure 1: Consumption profiles of early and late withdrawers as a function of public signal p , where p indicates the probability of a good return G on the risky illiquid asset.

I define the parameter $\lambda \equiv L - r$ to be the quantity of the safe liquid asset reserved for the late withdrawers, and I interpret this quantity λ as the bank's liquidity cushion.

3 Stress Tests

Suppose that at date 1, after the public signal p is revealed but before the consumers privately learn their type, a regulator can perform a stress test on the risky asset to reveal additional information about its future return. The testing procedure and results of the test are public information; the bank cannot falsify either. However, the regulator can strategically choose the testing procedure in order to maximize the welfare of the consumers. The result of the stress test is a signal $\sigma \in \{g, b\}$ taking either a good value g or a bad value b . The probability of receiving either signal depends on the underlying true future return $R = \{G, 0\}$ of the risky illiquid asset. Stress tests are formalized as a pair of conditional distributions $\pi(\sigma|R = G)$ and $\pi(\sigma|R = 0)$, which are public information and chosen by the regulator.

By choosing the stress test $\pi(\sigma|R)$, the regulator induces not only a distribution $\mu(\sigma)$

over the signal outcomes σ , where

$$\mu(\sigma) = \pi(\sigma|R = G)p + \pi(\sigma|R = 0)(1 - p),$$

but also posterior beliefs p_σ about the probability of a high return G , conditional on the revealed signal σ , where

$$p_\sigma = \frac{\pi(\sigma|R = G)p}{\mu(\sigma)}.$$

Choosing the stress test $\pi(\sigma|R)$ is equivalent to choosing distribution $d(p_\sigma) = \mu(\sigma)$ over posterior beliefs p_σ , provided that the mean posterior belief equals the prior belief p :

$$p = p_g d(p_g) + p_b d(p_b).$$

Given liquid holdings L , illiquid holdings I , fixed return r for early withdrawers, let $U(p|L, I, r)$ be the date 1 average consumer utility under the optimal consumption profile $\{c_{11}(p), c_{21}(p), c_{22}(R, p), \alpha(p)\}$:

$$U(p|L, I, r) = \frac{1}{2}u(c_{11}(p)) + \frac{1}{2}\left(\alpha(p)u(c_{21}(p)) + (1 - \alpha(p))u(E[c_{22}(R, p)|p])\right).$$

The regulator's disclosure problem is as follows:

$$V(p|L, I, r) \equiv \max_{p_g, p_b} d(p_b)U(p_b|L, I, r) + d(p_g)U(p_g|L, I, r)$$

s.t.

$$p_b, p_g, d(p_b), d(p_g) \in [0, 1]$$

$$d(p_b) + d(p_g) = 1$$

$$p_b \leq p_g$$

$$p = p_b d(p_b) + p_g d(p_g).$$

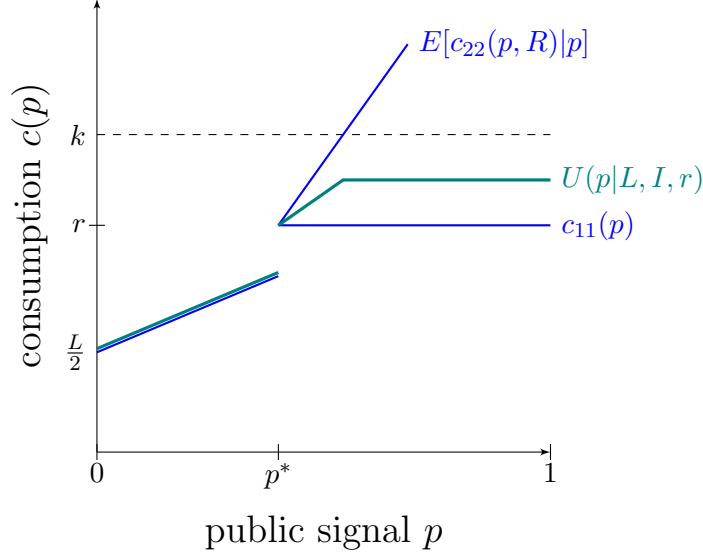


Figure 2: Average utility of consumers at date 1.

Proposition 2. Let $p^{**}(L, I, r)$ be the belief at which the late withdrawer's expected date 2 consumption equals the kink k : $L - r + E[R|p^{**}]I = k$.

1. Suppose the portfolio (L, I) and fixed early return r are such that (a) without a stress test, the public signal p would induce a run: $p < p^*(L, I, r)$; and (b) there exists some feasible belief at which runs are avoided: $p^*(L, I, r) \leq 1$.

Then an optimal stress test (p_b, p_g) must satisfy:

$$p_b = 0, \quad p_g \in [p^*, p^{**}],$$

which implies

$$\pi(g|R = G) = 1, \quad \pi(g|R = 0) = \frac{p(1 - p_g)}{p_g(1 - p)}.$$

2. Otherwise, an uninformative test is weakly optimal:

$$p_b = p_g = p, \quad \pi(g|R = G) = \pi(g|R = 0)$$

The proposition explains that in bad times, where $p < p^*$, the regulator can improve

consumer welfare by performing partially informative stress tests. Without the tests, depositors would run with certainty, forcing the bank to liquidate its risky assets at a significant discount. However, with partially informative stress tests, there is a positive probability that the bank will pass the test and patient investors will refrain from running. Stress tests therefore create value by reducing the likelihood of costly liquidation.

Condition (1b) stipulates that in order for stress tests to be valuable at all, there must exist some feasible belief at which investors would avoid runs. The condition is equivalent to $p^* \leq 1$, where $p^*(L, I, r, \lambda) = (u(r) - \lambda)/(GI)$. Intuitively, if the return r promised to early withdrawers is too high, the liquidity cushion λ too low, or the risky payoff GI in a good state too low, then patient investors understand that even in the best case scenario they will do worse by waiting than by withdrawing early, so they run no matter what. As a consequence, costly liquidation is unavoidable, so informative stress tests offer no value.

In explaining optimal stress tests, I focus on the test where $p_g = p^*$, although the proposition states that a range of posterior beliefs slightly above the run cutoff p^* could also be optimal. This is simply a consequence of using a piecewise linear utility function, which makes tractable the ex ante portfolio problem explored in the next section. For a strictly concave utility function, $p_g = p^*$ is uniquely optimal, so I focus on $p_g = p^*$ because it is optimal in either case.

Now consider the form of the optimal stress test. If conditions (1a) and (1b) hold, then the optimal test passes a good bank with certainty but passes a bad bank with a lower, but positive, probability. The idea here is that stress tests should be informative, but not too harsh. Because all good banks pass with certainty, depositors know that failing grades can only come from bad banks, so the posterior belief from a failing grade is as low as possible: $p_b = 0$. However, because optimal stress tests can produce passing grades for both good and bad banks, a passing grade is less clear, so it's possible for a passing grade to give depositors less than full confidence: $p_g < 1$.

Because runs induce costly liquidation, it is optimal to construct a test which induces

runs with the lowest likelihood possible. This is a different objective from designing a test which produces failed grades with the lowest likelihood. Such a test would simply pass every bank, but would be completely uninformative, and if prior beliefs p were below the run cutoff p^* , then runs would occur with certainty. Therefore, it is necessary to fail banks with some likelihood in order for a passing grade to have any credibility. The optimal test gives a failing grade with just enough likelihood so that passing grades are credible enough to prevent runs.

The following corollary shows how the stringency of the stress tests varies with the public signal p .

Corollary 1. *Under optimal stress tests, as public prior beliefs p decrease,*

1. $\pi(g|R = 0)$ decreases, so stress tests become more stringent, and
2. $d(p_b)$ increases, so failing grades are given with greater frequency.

The corollary highlights a natural intuition, which is that the less confidence consumers have in the economy, the more credible a passing grade must be in order to prevent runs. Credibility is achieved by decreasing the likelihood $\pi(g|R = 0)$ that bad banks receive a passing grade, as shown in part (1) of the Corollary. Part (2) highlights the fact that the total likelihood of failing grades increases as well, which results both from the increased stringency of the tests, as well as the lower prior likelihood p that the bank is good.

4 Liquidity Cushion and Portfolio Choice

I now examine how stress tests affect the bank's optimal portfolio L and I as well as promised early return r and liquidity cushion λ . Without loss of generality, I can assume that $r \leq L$, because otherwise the bank would be forced to liquidate the illiquid risky asset in every state, which is clearly not optimal.

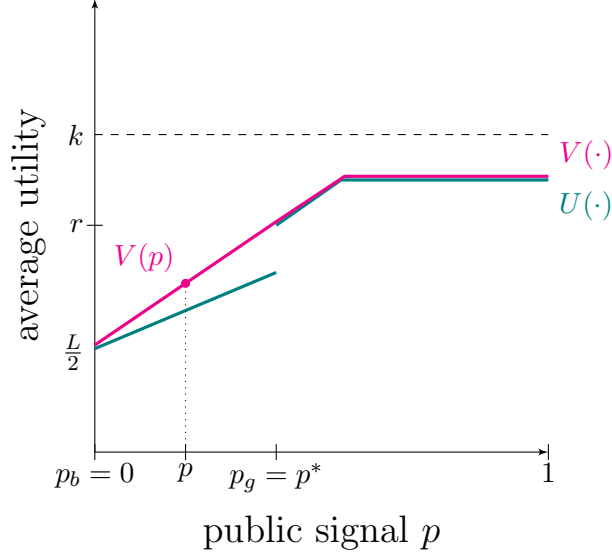


Figure 3: Average utility of consumers at date 1.

4.1 No stress tests

Suppose the bank does not anticipate stress tests. Then at date 1, the average current and future utility of consumers is captured by the function $U(p|L, I, r, \lambda)$. So by backward induction, the bank chooses its investments, early return, and liquidity cushion to maximize the expected average date 1 utility $U(p|L, I, r, \lambda)$, where the expectation is taken over realizations of the public signal p . I write the bank's problem with the objective function U expanded to display the run and no-run cases as follows:

$$\max_{L, I, r, \lambda} \int_0^{\hat{p}(L, I, r, \lambda)} 2u\left(\frac{L + \theta E[R|p]I}{2}\right) dp + \int_{\hat{p}(L, I, r, \lambda)}^1 \left(u(r) + u(\lambda + E[R|p]I)\right) dp$$

s.t.

$$L + I \leq E \tag{8}$$

$$r + \lambda \leq L \tag{9}$$

$$L \geq 0, \quad I \geq 0, \quad r \geq 0, \quad \lambda \geq 0. \tag{10}$$

where

$$\hat{p} \equiv \max[0, \min[1, p^*]] \quad (11)$$

and

$$p^*(L, I, r, \lambda) = \frac{u(r) - \lambda}{GI}. \quad (12)$$

Constraint (8) is simply the bank's budget constraint: its total investments cannot exceed consumer deposits E because the bank has no other source of funding. Constraint (9) highlights the two uses of the safe liquid asset. The first is the fixed return r to early withdrawers if beliefs are optimistic enough to avoid runs. It is optimal to pay early withdrawers with the safe liquid asset when possible because it may be liquidated at no cost. The remaining portion of the safe liquid asset is stored as a liquidity buffer λ . Constraint (12) defines the minimum belief p^* at which patient depositors will refrain from running, and is the value of p at which the patient investor's expected consumption from waiting $\lambda + E[R|p]I$ equals the utility $u(r)$ of withdrawing early. Because the cutoff belief p^* may be greater than 1 or less than 0, the function defined in constraint (11) maps the cutoff belief p^* to a feasible probability value $\hat{p} \in [0, 1]$ so that the objective function avoids integrating over infeasible values of p .

The utility function $u(c) = \min[c, k]$ is piecewise linear with a kink at k , so the total deposits E offered to the bank must fall in a particular range in order for the concavity of u to be relevant. In particular, I make the following two assumption about the amount of deposits E , relative to the kink k .

Assumption 1. $\frac{E}{2} < k$

This assumption places an upper bound on the total wealth of the consumers. If it were violated, then the bank could offer consumers the certain utility of k by simply letting $L = E$,

$I = 0$, and $r = \lambda = E/2 \geq k$. If this were the case, then the deposit contract would have no risk because nothing is invested in the risky asset and because early and late withdrawers receive equal consumption $E/2 \geq k$, so the maximal utility of k is guaranteed. Assumption 1 rules out this possibility by capping the total deposits. This creates an incentive to invest in the risky asset in order to benefit from its higher expected return $E[R] > 1$.

Assumption 2. $\sqrt{2GE} > k$

This assumption places a lower bound on the total wealth of consumers. Otherwise, investors only operate on the increasing linear portion of the utility curve, and then are effectively risk neutral, so invest all their assets in the risky asset. Assumption 2 prevents this corner solution by guaranteeing investors have enough wealth for the kink to be relevant, inducing risk aversion.

Although there are feasible portfolios (L, I) and contracts (r, λ) which never induce runs, the following proposition asserts that these are never optimal. Denote optimal values with a subscript n if the bank anticipates no stress tests.

Proposition 3 (Optimal Runs). *If the bank does not anticipate stress tests, then the optimal portfolio and contract $(L_n, I_n, r_n, \lambda_n)$ induces runs with some positive probability: $p^*(L_n, I_n, r_n, \lambda_n) > 0$.*

For runs to occur with any positive probability, the promised early return r must exceed the liquidity cushion λ . It is never optimal for the early return r to be strictly less than the liquidity cushion λ , because then the late investor is guaranteed to consume more than the early investor, so the contract does not smooth consumption between dates 1 and 2 sufficiently. A better contract would, for a fixed safe liquid investment L , raise r and lower λ so that early and late investors would consume more similarly.

It is also never optimal for the early return r to equal the liquidity cushion λ , but for a more subtle reason. For such a contract, there is an intertemporal benefit of raising r and lowering λ so that $r > \lambda$, but there is also the cost induced by a run. However, for r close

to λ , runs are unlikely, because only very low signals p can cause expected late consumption $\lambda + E[R|p]I$ to fall below the early return r . Not only are runs very unlikely for r close to λ , but when they do occur they are not very costly because the risky illiquid asset is not worth very much anyway. Because of this, the benefit from intertemporal smoothing dominates the likelihood and cost of a run, so the bank can improve on $r = \lambda$ by raising the early return r while lowering the liquidity cushion λ . As a consequence, some probability of runs occur under the optimal portfolio-contract.

For the remainder of this section, I adopt an additional assumption in order to simplify the analysis.

Assumption 3. $\frac{GE}{1+G} < k$

The assumption guarantees that there is no portfolio-contract (L, I, r, λ) that satisfies both $r \geq k$ and $p^* \leq 1$. This permits me to assume that if there exists some $p \in [0, 1]$ such that patient depositors refrain from running, then the amount r promised to early investors must be strictly less than the kink k .

Proposition 4 (Optimal Liquidity Cushion, No Disclosure). *Given endowment E and kink k , (i) there exists a good return $G > 0$ satisfying Assumption 2 and liquidation discount $\theta \in [0, 1]$ such it is optimal to hold a positive liquidity cushion: $\lambda_n > 0$; and (ii) there exists a good return $G > 0$ satisfying Assumption 2 and liquidation discount $\theta \in [0, 1]$ such that it is not optimal to hold a liquidity cushion: $\lambda_n = 0$.*

The proposition states that the optimal presence or absence of a liquidity buffer λ depends on the good return G and the liquidation discount θ . Numerical simulations indicate that a low enough good return G or a low enough liquidation discount θ imply that it is optimal to hold a positive liquidity buffer: $\lambda_n > 0$. The intuition is that if the good return G is low, then only very high signals p will assure late investors of consumption exceeding the early return r , so costly runs occur with high frequency without reserving some of the safe asset as a liquidity buffer. Alternatively, if the liquidation discount θ is low enough, then

regardless of the likelihood of a run, they are very costly when they do occur, so lowering their likelihood with a liquidity cushion $\lambda_n > 0$ is optimal. Alternatively, if the good return G or liquidation discount θ are high enough, then runs occur with low likelihood or are not very costly, so no liquidity buffer is needed, and $\lambda_n = 0$.

4.2 Optimal Stress Tests

If the bank anticipates optimal stress tests, then the date 1 average current and future utility of depositors is given by $V(p|L, I, r, \lambda)$, and the bank maximizes the expected value of V where the expectation is taken over the public signal p :

$$\max_{L, I, r, \lambda} \int_0^1 V(p|L, I, r, \lambda) dp$$

s.t.

$$L + I \leq E \tag{8}$$

$$r + \lambda \leq L \tag{9}$$

$$L \geq 0, \quad I \geq 0, \quad r \geq 0, \quad \lambda \geq 0. \tag{10}$$

Theorem 1 (No Liquidity Cushion Under Optimal Disclosure). *Suppose that the bank anticipates an optimal stress test at date 1. Then it is optimal for the bank to reserve no liquidity buffer: $\lambda_o = 0$.*

Provided Assumption 3 is satisfied, the result of the theorem holds, independent of the good return G or the liquidation cost θ . This is in contrast to the case of no disclosure, in which the optimal liquidity buffer may be positive given sufficiently low G or θ . Intuitively, optimal stress tests lower the likelihood of a run, so the value of a liquidity buffer is decreased. In addition, under optimal stress tests, runs occur only when investors are certain the risky illiquid asset is worthless ($p_b = 0$), in which case there is no loss from being forced to

liquidate it early. This effectively eliminates the cost of runs, which is the only reason for holding liquidity buffers in the first place.

Note that although this result technically requires that the risky asset returns nothing in the bad state ($B = 0$), the intuition is similar if the bad state yields a positive return ($B > 0$). In that case, optimal stress tests lower, rather than eliminate, the cost of runs by inducing depositors to run only when they are certain the risky asset will yield the lowest possible return $B > 0$, which implies the liquidation cost $(1 - \theta)BI > 0$ is also the lowest possible value. Therefore, optimal stress tests should lower the optimal size λ_o of the liquidity cushion because the cost of runs have decreased.

Theorem 2. *If without stress tests, no liquidity cushion is optimal ($\lambda_n = 0$), then under optimal stress tests, the return r_o to early withdrawers—and therefore the investment L_o in the safe liquid asset—should rise: $L_o = r_o > L_n = r_n$.*

If holding a positive liquidity cushion is not optimal with or without stress tests, then the liquid asset must be used entirely to fund the return r to early withdrawers, so $L_n = r_n$ and $L_o = r_o$. A low early return r may allocate too much consumption to the late withdrawers and not enough to the early withdrawers, making the ex ante consumption profile too risky for agents who do not yet know their time preferences. So the value of raising r (and therefore L) is that it smoothes consumption across dates 1 and 2. The cost of raising r (and therefore L) is twofold: (1) depositor beliefs p must clear a higher bar p^* to avoid running, so runs are more likely, and (2) less deposits are allocated to the high return illiquid asset. Under optimal stress tests, the cost of runs is eliminated, so the only remaining cost of raising the early return r is the foregone return R on the illiquid asset. Therefore, optimal stress tests should result in a higher return r_n to early withdrawers and therefore a greater investment L_n in the safe liquid asset.

5 Conclusion

The paper presents a model in which stress tests may be optimally designed in order to reduce the probability of a run. Runs are costly because they force banks to sell illiquid assets at a discount, reducing the resources available to withdrawing depositors. It is optimal for stress tests to pass all good banks, but to fail just enough bad banks to maintain the credibility of a passing grade. In particular, passing grades must be credible enough to prevent runs. The less confidence depositors have in the bank, the more stringent stress tests must be to prevent runs. By reducing the probability of runs, optimal stress tests reduce the value of liquidity cushions, and therefore the optimal level of liquidity cushions. In addition, the reduced probability of runs permit greater consumption smoothing by allowing a higher return to early withdrawers.

References

- Allen, Franklin and Douglas Gale, 1998, Optimal financial crises, *The Journal of Finance* 53(4): 1245-1284.
- Bouvard, Mattieu, Pierre Chaigneau, and Adolfo de Motta, 2015, *The Journal of Finance* 70(4).
- Diamond, Douglas and Philip Dybvig, 1983, Bank runs, deposit insurance, and liquidity, *The Journal of Political Economy* 91(3): 401-419.
- Epstein, Larry and Stanley Zin, 1989, Substitution, risk aversion, and the temporal behavior of consumption and asset returns: a theoretical framework, *Econometrica* 57(4): 937-969.
- Faria-e-Castro, Miguel, Joseba Martinez, and Thomas Philippon, 2015, Runs versus lemons: information disclosure and fiscal capacity, Working Paper, NBER.

Gick, Wolfgang and Thilo Pausch, 2012, Persuasion by stress testing, Working Paper, Harvard University.

Goldstein, Itay and Yaron Leitner, 2013, Stress tests and information disclosure, Working Paper, University of Pennsylvania.

Kamenica, Emir and Matthew Gentzkow, 2011, Bayesian persuasion, *American Economic Review* 101: 2590-2615.

Parlato, Cecilia, 2015, Transparency and Bank Runs, Working Paper, New York University.

Appendix

Proof of Proposition 1 The piecewise linearity of the utility function, while improving tractability of the date 0 portfolio problem, makes the date 1 consumption problem more complicated than in Allen and Gale (1998). Throughout the proof, I abbreviate c_{11} by c_1 . I also assume that $r \leq L$, because otherwise costly liquidation occurs for all $p \in [0, 1]$, which is clearly not optimal.

First note that for any (L, I, r, p) , $l = 1$ weakly dominates $l = 0$. To see this, observe that for any $\{c_1, c_{21}, \alpha, c_{22}, i, l = 0\}$ satisfying constraints (1) through (7), by (4) I must have $c_1 = r$. So then if l is changed so $l = 1$, the only changed condition is (1), which is clearly satisfied. Therefore, any consumption allocation attainable by $l = 0$ is attainable by $l = 1$, so $l = 1$ is weakly optimal.

I next prove two lemmas which will be useful in completing the proof of the proposition.

Lemma 1. *If $l = 1$ and $i = 0$, then*

1. $c_1 = r$

2. $L + pGI - r \geq u(r)$

3. $\alpha = 0$

4. $E[c_{22}] = L + pGI - r$.

5. Aggregate utility is $u(r) + u(L + pGI - r)$.

Proof: By (4), $c_1 = r$ so item 1 holds. If item 2 is strict, then $\alpha > 0$ requires (by (5)) $c_{21} = r$ and (by (7)) $u(E[c_{22}]) \leq u(r)$, which is clearly improved on by $\alpha = 0$, $u(E[c_{22}]) = u(L + pGI - r) > u(r)$. If item 2 holds with equality, then by (2) raising α will weakly decrease $E[c_{22}]$, so $\alpha = 0$ is optimal. This shows that item 2 implies items 3 through 5.

Now suppose item 2 is violated. Then by (2) raising α will strictly decrease $E[c_{22}]$, so for all $\alpha \in [0, 1]$, $u(E[c_{22}]) \leq u(L + pGI - r) < u(r)$, and so by (6), $\alpha = 1$. This gives, by (5), $c_{21} = r$, and (1) implies $2r \leq L$. But this contradicts my assumption that item 2 is violated. Therefore, item 2 must hold, and so must items 3 through 5. \square

Define $W \equiv L + \theta pGI$ to be the total wealth if all assets are liquidated.

Lemma 2. *If $l = 1$ and $i = 1$, then*

1. *If $W - r \geq u(r)$, then aggregate utility is $u(r) + u(W - r)$, and $\alpha = 0$, $c_1 = r$, and $E[c_{22}] = W - r$ is an optimal allocation.*
2. *If $W - r < u(r)$, then aggregate utility is $2u(W/2)$, $\alpha = 1$, and $c_1 = c_{21} = W/2 < r$.*

Proof: First note that if $c_1 < r$, then $\alpha = 1$. To see this, observe that (4) and (2) imply $(1 - \alpha)E[c_{22}] = 0$. If $\alpha < 1$, then by (6) I have $u(c_1) \leq u(E[c_{22}])$ and $E[c_{22}] = 0$, which is a contradiction.

(Part 1) Suppose $r \geq k$. Then letting $c_1 = r$, $\alpha = 0$, and $E[c_{22}] = W - c_1 = W - r \geq u(r)$ implies $u(E[c_{22}]) = u(W - r) \geq u(r)$, so this is a feasible allocation. Also, the aggregate utility under this allocation is $u(r) + u(W - r) = 2k$, which is the unconstrained maximum utility, so this allocation must be an optimal one.

Suppose $r < k$. If $c_1 < r$, then as noted above, $\alpha = 1$, and (4) implies that $r > c_1 = c_{21} = W/2$. This contradicts the assumption of Part 1, so $c_1 = r$. Given this, if $W - r > r$, then $\alpha > 0$ requires by (5) $c_{21} = r$ and by (7) $u(E[c_{22}]) \leq u(r)$, which is clearly improved on

by $\alpha = 0$, $u(E[c_{22}]) = u(W - r) > u(r)$. And if $W - r = r$, then raising α will leave $E[c_{22}]$ unchanged, so $\alpha = 0$ is an optimal value.

This gives $c_1 = r$, $\alpha = 0$, and $E[c_{22}] = W - r$ as an optimal allocation, with aggregate utility of $u(r) + u(W - r)$.

(Part 2) If $c_1 = r$, then by (2) raising α will strictly decrease $E[c_{22}]$, so for all $\alpha \in [0, 1]$, $u(E[c_{22}]) \leq u(W - r) < u(r)$, and so by (6), $\alpha = 1$. This gives, by (5), $c_{21} = c_1$ and by (1), $2c_1 = 2c_{21} \leq W < u(r) + r \leq 2r$, which is a contradiction. So $c_1 < r$, $\alpha = 1$, $c_1 = c_{21}$ and (4) gives $c_1 = W/2$. \square

Armed with Lemmas 1 and 2, I am ready to prove Proposition 1. I have already shown above that $l = 1$ weakly dominates $l = 0$.

Part (i): If $i = 0$, then Lemma 1 indicates that aggregate utility is $u(r) + u(L - r + rGI)$. If $i = 1$, then Lemma 2 indicates that if $W - r \geq u(r)$, then aggregate utility is $u(r) + u(W - r) \leq u(r) + u(L - r + pGI)$. If $W - r < u(r)$, then $W < r + u(r) \leq 2r$, and by Lemma 2 aggregate utility is $2u(W/2) \leq 2u(r) \leq u(r) + u(L - r + pGI)$. So clearly $i = 0$ is optimal, and by Lemma 1, I have $c_1 = r$, $\alpha = 0$, $E[c_{22}] = L - r + pGI$.

Part (ii): If $p < p^*$, then $L - r + pGI < u(r)$, so by Part 2 of Lemma 1 I must have $i = 1$ is optimal. Also, this implies $W - r = L + \theta pGI - r < L + pGI - r < u(r)$, so by Lemma 2 I have $\alpha = 1$, $c_1 = c_{22} = W/2$, and the proposition is proved. \square

Proof of Proposition 2 Part 1: Denote the left limit of $U(p)$ by $U_-(p)$, and note that the jump discontinuity at p^* implies $U_-(p^*) < U(p^*)$. If $p_g = p^*$ and $p_b = 0$, then the constraints of the disclosure problem imply $d(p_g) = p/p^*$ and $d(p_b) = 1 - p/p^*$, so the objective function is $(1 - p/p^*)U(0) + (p/p^*)U(p^*) > (1 - p/p^*)U(0) + (p/p^*)U_-(p^*) = U(0 \cdot (1 - p/p^*) + p^*(p/p^*)) = U(p)$, where the second to last equality follows from the linearity of U over $[0, p^*)$. So I have shown that $(p_g = p^*, p_b = 0)$ is strictly better than an uninformative test $p_g = p_b = p$. Next, note that the constraints imply that $p_b \leq p \leq p_g$, with $p < p^*$. If $p_g < p^*$, then by the linearity of U over $[0, p^*)$, the objective function takes value equal to $U(p)$, which I have shown is not optimal. So the optimal test must

have $p_g \geq p^*$ and $p_b \leq p < p^*$. The constraints imply that $d(p_b) = (p_g - p)/(p_g - p_b)$ and $d(p_g) = (p - p_b)/(p_g - p_b)$. Substituting into the objective function and differentiating with respect to p_g gives

$$FOC(p_g) : \frac{U'(p_g)(p_g - p_b) - (U(p_g) - U(p_b))}{(p_g - p_b)^2},$$

which has the same sign as $U'(p_g) - (U(p_g) - U(p_b))/(p_g - p_b)$, or

$$u'(L - r + p_g GI)GI - \frac{u(r) + u(L - r + p_g GI) - 2u((L + \theta p_b GI)/2)}{(p_g - p_b)}. \quad (13)$$

Because $U(\cdot)$ is strictly increasing below p^* and weakly increasing thereafter, and $p_b < p < p^* \leq p_g$, the second term must be positive. If $p_g > p^{**}$, then the first term is zero, so (13) is negative. Now suppose $p_g \in [p^*, p^{**}]$. If $r \geq k$, then the first term is zero, so (13) is negative. If $r < k$, then (13) reduces to $GI - GI(p_g - \theta p_b)/(p_g - p_b) \leq 0$, with strict inequality if $p_b > 0$ and equality otherwise. So for any $p_b < p$ and $p_g \geq p^*$, the disclosure objective function is strictly decreasing in p_g for $p_g > p^{**}$ and strictly decreasing in p_g for $p_g \in [p^*, p^{**}]$ if p_b is nonzero. So for any $p_b < p$, an optimal p_g must satisfy $p_g \in [p^*, p^{**}]$.

Now differentiate the objective with respect to p_b to get

$$FOC(p_b) : \frac{U'(p_b)(p_g - p_b) - (U(p_g) - U(p_b))}{(p_g - p_b)^2},$$

and without loss of generality set $p_g = p^*$, which I have shown is optimal for any $p_b < p$. The $FOC(p_b)$ has the same sign as $U'(p_b) - (U(p^*) - U(p_b))/(p^* - p_b) < U'(p_b) - (U_-(p^*) - U(p_b))/(p^* - p_b) = 0$, where the equality follows from the linearity of U over $[0, p^*)$. So $FOC(p_b)$ is strictly negative, and therefore $p_b = 0$ is uniquely optimal.

Using Bayes rule, I can back out the stress test $\pi(\sigma|R)$. Note that $0 = p_b = (1 - \pi(g|G))p/\mu(b)$ implies $\pi(g|G) = 1$, and $p_g = \pi(g|G)p/(\pi(g|G)p + \pi(g|0)(1 - p))$ implies $\pi(g|0) = p(1 - p_g)/(p_g(1 - p))$.

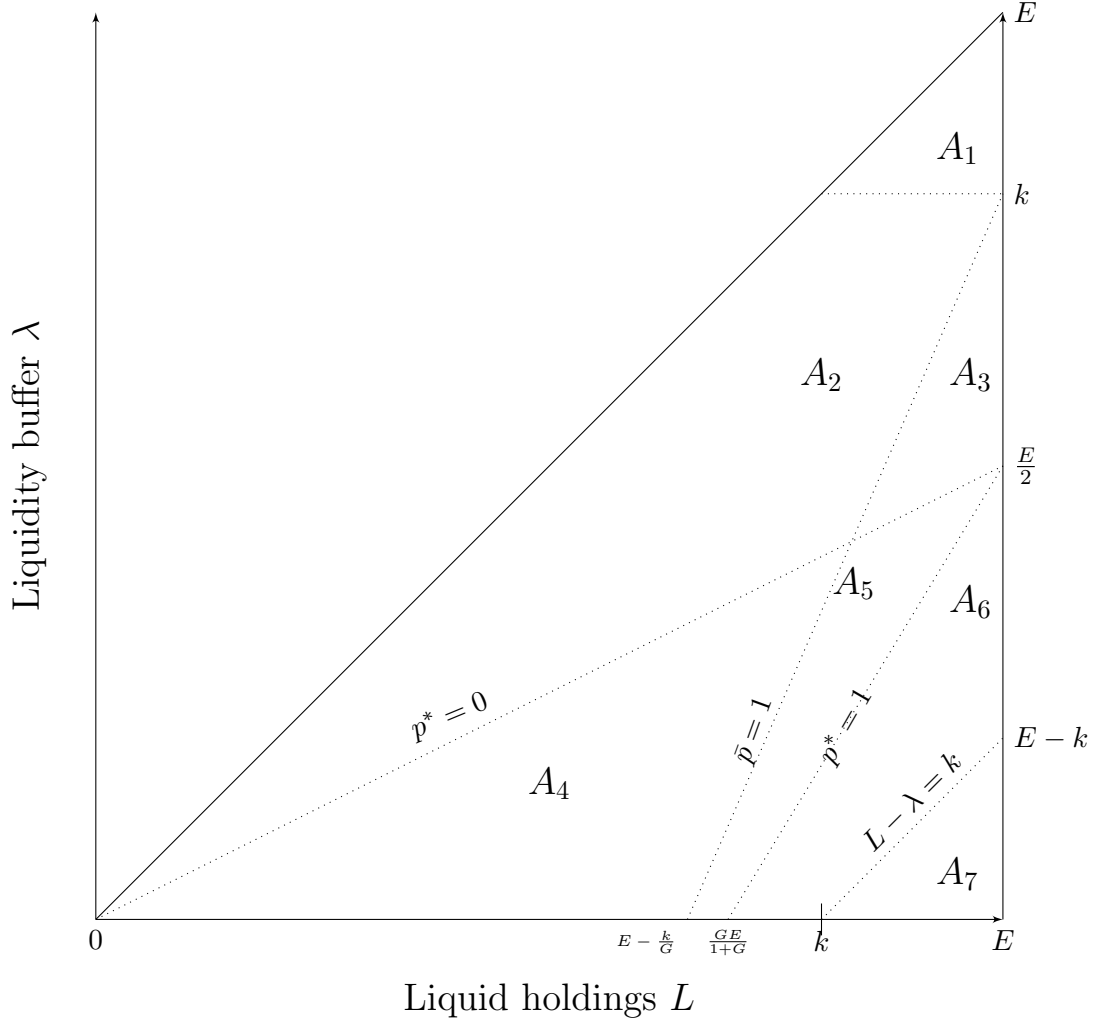


Figure 4: Differentiable regions of the portfolio-contract $L - \lambda$ space.

Part 2: If $p^* > 1$ or if $p^* \leq 0$, then $U(\cdot)$ is weakly concave over $[0, 1]$. So for any $d(p_b)$, $d(p_g)$, p_b , and p_g satisfying the constraints, $d(p_b)U(p_b) + d(p_g)U(p_g) \leq U(d(p_b)p_b + d(p_g)p_g) = U(p)$. Since $U(p)$ can be achieved by an uninformative signal $p_b = p_g = p$, it must be weakly optimal. \square

Proof of Corollary 3 *Part 1:* The optimal posterior belief p_g is independent of p , and $\pi(g|0) = p(1 - p_g)/(p_g(1 - p))$, which decreases as p decreases. *Part 2:* Since $p_b = 0$, the constraint $p = d(p_b)p_b + d(p_g)p_g$ implies $d(p_g) = p/p_g$. Therefore, $d(p_b) = 1 - d(p_g) = 1 - p/p_g$, which decreases as p decreases. \square

Proof of Proposition 3 Because the objective function is weakly increasing in I and λ ,

without loss of generality, Equations (8) and (9) bind. So the choice variables can be reduced to $L \in [0, E]$ and $\lambda \in [0, L]$. Denote the objective function under no stress tests by $O_n(L, \lambda)$. The objective function is differentiable only in certain regions, with kinks identified by the dotted lines in Figure 4. The figure is drawn assuming $k < E$, but this is not necessary for the proofs. If $k \geq E$, then regions A_1 and A_7 are simply eliminated. Where O_n is differentiable, I have

$$\begin{aligned} \frac{\partial O_n}{\partial L} = & \int_0^{\hat{p}} u' \left(\frac{L + \theta pGI}{2} \right) (1 - \theta pG) dp + \int_{\hat{p}}^1 \left[u'(L - \lambda) - u'(\lambda + pGI)pG \right] dp \\ & - \frac{\partial \hat{p}}{\partial L} \left[u(L - \lambda) + u(\lambda + \hat{p}GI) - 2u \left(\frac{L + \theta \hat{p}GI}{2} \right) \right] \end{aligned} \quad (14)$$

$$\begin{aligned} \frac{\partial O_n}{\partial \lambda} = & \int_{\hat{p}}^1 \left[-u'(L - \lambda) + u'(\lambda + pGI) \right] dp \\ & - \frac{\partial \hat{p}}{\partial \lambda} \left[u(L - \lambda) + u(\lambda + \hat{p}GI) - 2u \left(\frac{L + \theta \hat{p}GI}{2} \right) \right] \end{aligned} \quad (15)$$

The strategy is to show that the optimum must be in region A_4 , and then to show that it does not lie on the upper boundary $p^* = 0$ of A_4 . First consider region A_1 , where $\lambda > k$. In this region, $p^* < 0$, so $\hat{p} = 0$, and $L - \lambda < k$. Therefore, $\partial \hat{p} / \partial \lambda = 0$, so $\partial O_n / \partial \lambda = -1$, and in region A_1 , the optimum must lie along $\lambda = k$. In region A_3 , $\bar{p} > 1$, $\lambda < k$, and $p^* < 0$, so $\partial O_n / \partial L = 1 - G/2 < 0$, which implies the optimum must lie along $\bar{p} = 1$. In region A_2 , $p^* < 0$, $\lambda < k$, and $\bar{p} < 1$, so $\partial O_n / \partial \lambda = -1 + \bar{p} < 0$, which implies the optimum in A_2 must lie along $\bar{p} = 1$ or $p^* = 0$. Also, in region A_2 , $\partial O_n / \partial L = 1 - G\bar{p}^2/2$. Let $\lambda(L)$ be the value of λ which solves $\bar{p} = 1$ for a given L . Then $dO_n(L, \lambda(L))/dL = \partial O_n(L, \lambda(L))/\partial L + (\partial O_n(L, \lambda(L))/\partial \lambda)\lambda'(L) = 1 - G\bar{p}^2/2 + (-1 + \bar{p})\lambda'(L) = 1 - G/2 + 0 < 0$, so O_n is decreasing along $\bar{p} = 1$ as L increases, which implies the optimum in A_2 must lie on $p^* = 0$. Because A_2 includes the optima of A_1 and A_3 , the optimum of (A_1, A_2, A_3) together must lie on $p^* = 0$, with $\bar{p} \leq 1$.

I next establish a lemma which simplifies the computation of (14) and (15) when $p^* \geq 0$.

Lemma 3. *Suppose $p^* \geq 0$. Then the consumption $c_{11}(p)$ under a run ($p < p^*$) does not exceed the kink k .*

Proof. If $p < p^*$, then $c_{11}(p) = (L + \theta pGI)/2 < (L + \theta \hat{p}GI)/2$. In regions A_6 and A_7 , $p^* > 1$, so $\hat{p} = 1$, and also $L > k$. Then $L + \theta \hat{p}GI = L + \theta G(E - L)$. If $\theta G \leq 1$, then $L + \theta G(E - L) \leq E < 2k$. If $\theta G > 1$, then $L + \theta G(E - L) = \theta GE - (\theta G - 1)L < \theta GE - (\theta G - 1)k < \theta GE - (\theta G - 1)(E - k/G) = \theta k + E - k/G < \theta k + k < 2k$. In regions A_4 and A_5 , where $p^* < 1$ and $L - \lambda < k$, I have $L + \theta \hat{p}GI = L + \theta p^*GI < L + p^*GI = L + L - 2\lambda = 2(L - \lambda) < 2k$. So in regions A_4 through A_7 , $c_{11}(p) = (L + \theta pGI)/2 < k$, and the lemma is proved. \square

Now consider the region (A_4, A_5, A_6, A_7) , where $p^* \geq 0$. In regions A_6 and A_7 , I have $\bar{p} > 1$ and $p^* > 1$, so $\hat{p} = 1$ and then $\partial \hat{p} / \partial \lambda = 0$. Therefore, (15) is 0, the objective function is constant in λ , so an optimal point of (A_6, A_7) must lie on the upper boundary $p^* = 1$. In the interior of A_5 , I have $p^* < 1$ but $\bar{p} > 1$, so (15) is $\int_{p^*}^1 [-1 + 1] - (-2/GI)(1 - \theta)(L - 2\lambda) > 0$; therefore $\partial O_n / \partial \lambda > 0$, and the optimum must lie on the upper boundaries $\bar{p} = 1$ or $p^* = 0$. Now consider (14) in A_5 , which is $\int_0^* (1)(1 - \theta pG)dp + \int_{p^*}^1 (1 - pG)dp - (E - 2\lambda)/(GI^2)(1 - \theta)(L - 2\lambda)$. Along $p^* = 0$, I have $0 = p^* = (L - 2\lambda)/GI$, which implies $L = 2\lambda$, so along $p^* = 0$, (14) reduces to $1 - G/2 < 0$ and (15) reduces to 0. Let $\lambda(L) = L/2$, so that for any L , $\lambda(L)$ solves $p^* = 0$. Then along $p^* = 0$, I have $dO_n(L, \lambda(L))/dL = \partial O_n / \partial L + \partial O_n / \partial \lambda \cdot \lambda'(L) = 1 - G/2 + 0 \cdot 1/2 < 0$, so O_n is decreasing along $p^* = 0$ as L increases, and the optimum of A_5 must lie along $\bar{p} = 1$. Because A_5 contains an optimal point of (A_6, A_7) , I must have that the optimal point of (A_5, A_6, A_7) lies along $\bar{p} = 1$. Because A_4 contains the optimal point of (A_1, A_2, A_3) along $p^* = 0$ and also contains the optimal point of (A_5, A_6, A_7) along $\bar{p} = 1$, the global optimum must lie in A_4 .

In the region A_4 , where $p^* \in [0, 1]$ and $\bar{p} \leq 1$, (14) and (15) may be written as

$$\frac{\partial O_n}{\partial L} = 1 - \frac{G}{2}\bar{p}^2 + \frac{G}{2}(1 - \theta)p^{*2} - \frac{E - 2\lambda}{GI^2}(1 - \theta)(L - 2\lambda) \quad (16)$$

$$\frac{\partial O_n}{\partial \lambda} = \bar{p} - 1 + \frac{2}{GI}(1 - \theta)(L - 2\lambda). \quad (17)$$

Along $p^* = 0$, $L - 2\lambda = 0$, so (16) reduces to $1 - \bar{p}^2 G/2$ and (17) reduces to $\bar{p} - 1 < 0$. Therefore, along any point of the interior of $p^* = 0$, lowering λ is an improvement. Now check the endpoints of $p^* = 0$ by examining how O_n changes along $p^* = 0$. Letting $\lambda(L) = L/2$ as above so that $\lambda(L)$ solves $p^* = 0$ for any L , observe that $dO_n(L, \lambda(L))/dL = \partial O_n/\partial L + \partial O_n/\partial \lambda \cdot \lambda'(L) = 1 - \bar{p}^2 G/2 + (\bar{p} - 1)/2$. At the right-most endpoint of $p^* = 0$ in A_4 , $\bar{p} = 1$, so $dO_n(L, \lambda(L))/dL = 1 - G/2 < 0$, so moving slightly to the interior of $p^* = 0$ is an improvement. At the left-most endpoint of $p^* = 0$, $L = \lambda = 0$, so $\bar{p} = k/(GE)$, so then Assumption 2 implies $\partial O_n/\partial L = 1 - \bar{p}^2 G/2 = 1 - (k/(GE))^2 G/2 > 1 - (\sqrt{2G}/G)^2 G/2 = 0$. Therefore, raising L is an improvement over $L = \lambda = 0$, so the left-most endpoint of $p^* = 0$ cannot be optimal, and therefore the optimal point of A_4 must satisfy $p^* > 0$. The optimal point of A_4 is the global optimum, so the Proposition is proved. \square

Proof of Proposition 4 In region A_4 , the objective function O_n is strictly concave in λ , which can be seen by noting that (17) is strictly decreasing in λ . Therefore, the optimum must lie on the set of points where (17) is equal to zero, which is a positive-sloped straight line intersecting $(L = (GE - k)/(G + 2(1 - \theta)), \lambda = 0)$, as well as the upper right corner of A_4 . At upper-right corner of A_4 , I have $p^* = 0$ (so $L - 2\lambda = 0$) and $\bar{p} = 1$, so (16) reduces to $1 - G/2 < 0$. Letting $\lambda(L)$ denote the value λ that sets (17) equal to zero, I have $dO_n(L, \lambda(L))/dL = \partial O_n/\partial L + \partial O_n/\partial \lambda \cdot \lambda'(L) = \partial O_n/\partial L + 0 \cdot \lambda'(L) = \partial O_n/\partial L$ along $\lambda(L)$. At the upper right corner of A_4 , $\partial O_n/\partial L = 1 - G/2 < 0$, so the optimum cannot lie on the upper right corner of A_4 .

Next, I show that (16) is strictly decreasing in L . Because \bar{p} is increasing in L , the sum of the first two terms are decreasing in L . The second two terms can be combined as $-(1 - \theta)p^{*2}G((E - 2\lambda)/(L - 2\lambda) - 1/2)$, and differentiating with respect to L gives $-(E - 2\lambda)^2 p^*/(GI^2(L - 2\lambda)) < 0$, so the sum of the last two terms is also strictly decreasing in L . Therefore, O_n is strictly concave in L . I have shown above that along $\lambda(L)$, $dO_n(L, \lambda(L))/dL = \partial O_n/\partial L$, so by the L -concavity of O_n , the optimal λ is nonzero if and only if $\partial O_n/\partial L > 0$ at the intersection of the graph of $\lambda(L)$ and $\lambda = 0$, which occurs at the

point $(L = (GE - k)/(G + 2(1 - \theta)), \lambda = 0)$. Evaluating (16) at this point, and multiplying by $2(G + 2(1 - \theta))^2$ to eliminate fractions indicates that (16) has the same sign as

$$(1 - \theta)E[EG^2 + 4(1 - \theta)(EG + k) + kG] + k^2[-(G - 2\theta)^2 - 2G + 7\theta - 3]. \quad (18)$$

If $\theta = 1$, then (18) reduces to $-k^2G(G - 2)$, which is negative for $G > 2$, so by the continuity of (18) there exists a neighborhood of $(\theta = 1, G = 2)$ which contains a point $\theta < 1, G > 2$ in which (18) is negative, so that $\lambda = 0$ is optimal. On the other hand, if $(\theta, G) = (0, 2)$, then (18) reduces to $12E^2 + k(6E - 11k) > 12E^2 + E/2(6E - 11k) > 12E^2 + E/2(6E - 22E) = 4E^2 > 0$, where the two inequalities follow from Assumptions 1 and 2 respectively. Therefore, there exists a neighborhood of $(\theta, G) = (0, 2)$ that contains a point with $\theta > 0, G > 2$ so that $\lambda > 0$ is optimal. \square

Proof of Theorem 1 I first show that the optimum must lie in A_4 . Denote the objective function under optimal stress tests by $O_o(L, \lambda)$. By Proposition 2, an uninformative test is weakly optimal if $p^* \leq 0$ or if $p^* > 1$, so $O_n = O_o$ in all but regions A_4 and A_5 , and the optimum must lie in A_4 or A_5 . Also note that if $p \geq p^*$, then $V(p) = U(p)$, but if $p < p^*$, then $V(p) = 2u((L + pGI)/2)$, which is simply $U(p)$ with $\theta = 1$. Therefore, the derivative of O_o in regions A_4 and A_5 may be obtained by simply substituting $\theta = 1$ into (14) and (15). For region A_5 , then, (15) is equal to zero, so an optimal point of A_5 must lie on the upper boundaries $\bar{p} = 1$ or $p^* = 0$. Let $\lambda(L) = L/2$, so that for any L , $\lambda(L)$ solves $p^* = 0$. Then along $p^* = 0$, I have $dO_n(L, \lambda(L))/dL = \partial O_n/\partial L + \partial O_n/\partial \lambda \cdot \lambda'(L) = 1 - G/2 + 0 \cdot 1/2 < 0$, so O_n is decreasing along $p^* = 0$ as L increases, and the optimum of A_5 must lie along $\bar{p} = 1$. This boundary is shared by A_4 , so the optimum must lie in A_4 .

In the region A_4 , the λ derivative of O_o is given by (17) with $\theta = 1$, which reduces to $\bar{p} - 1$. This value is strictly negative everywhere but the boundary $\bar{p} = 1$, so the optimal point of A_4 must lie on either the lower boundary $\lambda = 0$ or the right boundary $\bar{p} = 1$. However, the L derivative of O_o is given by (16) with $\theta = 1$, which reduces to $1 - \bar{p}^2G/2$. This value is

strictly negative on the boundary $\bar{p} = 1$, so the optimum cannot lie on the boundary $\bar{p} = 1$, and must lie on the lower boundary $\lambda = 0$. \square

Proof of Theorem 2 If $\lambda_n = 0$, then the optimal L_n without stress tests lies along the lower boundary where the L -derivative equals zero. Therefore, rearranging the last two terms of (16), I have that at L_n ,

$$0 = \frac{\partial O_n}{\partial L} \Big|_{L=L_n} = 1 - \frac{G}{2}\bar{p}^2 + \frac{G}{2}(1-\theta)p^{*2} - \frac{E-2\lambda}{GI^2}(1-\theta)(L-2\lambda) \quad (19)$$

$$= 1 - \frac{G}{2}\bar{p}^2 - (1-\theta)p^{*2}G\left(\frac{E-2\lambda}{L-2\lambda} - \frac{1}{2}\right) \quad (20)$$

$$< 1 - \frac{G}{2}\bar{p}^2 = \frac{\partial O_o}{\partial L} \Big|_{L=L_n}. \quad (21)$$

By the L -concavity of O_o , I must have $L_o > L_n$, and therefore $r_o = r_o + \lambda_o = L_o > L_n = r_n + \lambda_n = r_n$. \square