

**RUNS, PANICS AND BUBBLES:
DIAMOND DYBVIK AND MORRIS SHIN RECONSIDERED**

By

Eric Smith and Martin Shubik

August 2012

COWLES FOUNDATION DISCUSSION PAPER NO. 1869



**COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281**

<http://cowles.econ.yale.edu/>

Runs, Panics and Bubbles : Diamond Dybvig and Morris Shin reconsidered

Eric Smith

Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501

Martin Shubik

Cowles Foundation, Yale University, 30 Hillhouse Ave.,

New Haven, CT 06511

external faculty, Santa Fe Institute

1399 Hyde Park Road, Santa Fe, NM 87501

August 24, 2012

Abstract

The basic two-noncooperative-equilibrium-point model of Diamond and Dybvig is considered along with the work of Morris and Shin utilizing the possibility of outside noise to select a unique equilibrium point. Both of these approaches are essentially nondynamic. We add an explicit replicator dynamic from evolutionary game theory to provide for a sensitivity analysis that encompasses both models and contains the results of both depending on parameter settings.

JEL Classification: C73, D84, E59

Key Words: Multiple equilibria, runs, replicator dynamics, sensitivity analysis

1 The Diamond-Dybvig model and multiple equilibria

We doubt that there is a single universal explanation of financial bubbles and panics. There are just too many plausible possibilities available. A concatenation of many different special circumstances may suffice. Nevertheless Diamond and Dybvig [1] (DD) in a seminal article presented a basic structure simple enough to be illustrated in the form of a two-stage game with two well defined noncooperative equilibria, one reflecting a panic and the other an optimally functioning economy with the choice for panic not exercised. Their general framework has served as a model for a variety of questions concerning coordinated equilibria and endogenous risk, and provided ground for comparison of different equilibrium concepts.

The essential features of multiple-equilibrium models of the DD type is reviewed in an article by Diamond [2]. There are three decision points $T = 0, 1, 2$, and an investment opportunity with increasing returns to the time invested. An individual investing at $T = 0$ does not know whether she will need to consume at $T = 1$ or at $T = 2$. This can be modeled by assuming that at time $T = 1$ each agent is informed of the outcome of a lottery which with probability θ declares that she needs to consume at time $T = 1$ and with probability $(1 - \theta)$ that she will need to consume at $T = 2$. Banks issue demand deposits and can act as proxy investors for individuals; because it pools deposits the bank's needs for early termination of the investment are less severe than those faced by any individual. Therefore, under optimal conditions, the bank's demand deposits remain liquid assets even though its investments are not.

The model of demand-deposit contracts is that a bank repays an amount r_1 per unit deposited for money held for one period, and an amount r_2 for money held two periods, where $r_1 < r_2$. In the example of Ref [2], agents consume their withdrawals, labeled c_1 at $T = 1$ and c_2 at $T = 2$. If their consumption utility is $U(c) = (1 - 1/c)$, they are risk averse. Continuing with the example, all agents are endowed with one unit which they deposit, $(\theta, 1 - \theta)$ are set to $(1/4, 3/4)$, and two assets are considered. One (interpreted as the underlying investment opportunity) is illiquid; if liquidated in period $T = 1$ pays only 1, while if held until period $T = 2$ it pays 2. The

resulting expected utility of an individual holding the asset is

$$\frac{(1 - 1/1)}{4} + \frac{3(1 - 1/2)}{4} = 0.375.$$

A more liquid asset, in this case one created by bank deposits, might pay $(r_1, r_2) = (1.28, 1.813)$ in the two periods, giving an expected utility for a one-unit deposit of

$$\frac{(1 - 1/1.28)}{4} + \frac{3(1 - 1/1.813)}{4} = 0.391 > 0.375,$$

although the total amount paid out is less for the second asset than the first. Ref. [2] describes a variety of situations in which either consumers or investors might be sufficiently risk averse to prefer the second asset to the first despite its lower total payout.

However, as long as $r_1 > 1$, a second equilibrium also exists in which agents who do not need to withdraw at $T = 1$ do so anyway, making the bank insolvent, ensuring zero payout to any individual withdrawing at $T = 2$ (and thereby ensuring that there will be no such waiting individuals at this Nash equilibrium), and generating a new lottery to be paid either r_1 or nothing for the agents who withdraw at $T = 1$. The second equilibrium is the bank run.

The original Diamond-Dybvig model is ingenious in capturing the essential trade-off between “created” liquidity and co-created endogenous risk. However, it is fundamentally an equilibrium analysis, leaving many aspects of the solution concept implicit, and not considering the dynamics of panics. In particular, it makes multiple uses of “expectations” in the descriptive application of the model, which require further elaboration. As is standard in games with multiple equilibria, the choice of any investor to run or not to run is a best response “*if all are expected to do this*” [2]. Yet in the crucial question whether it can be rational for investors to deposit in a bank subject to runs, rationality is defended “*provided that the probability of a run is small enough*” [1]. The proposition that expectations can be rational but time-inconsistent (choosing to deposit and later choosing to run) requires the introduction of an additional (delayed) signal, and hence the refinement of the problem to an extensive-form game and a solution by correlated equilibrium. When these features were added several years later in a stimulating article by Morris and Shin [3], their conclusion was that under a general hash of uncertainty only one equilibrium survives for a given state of the world θ .

To consider the relation of these conclusions and also provide some context in dynamics, we first consider the role of interpretation of solution concepts and then analyze a simplified version of the Diamond-Dybvig model at three levels: in the original Nash equilibrium of the strategic form with implicit expectations, in a correlated subgame-perfect equilibrium of the corresponding extensive form, and finally using a replicator dynamic to identify subgame fixed points and their stability.

1.1 Durable intuition, formal claims that depend sensitively on modeling paradigm

The extent to which formal solution methods capture, ratify, or distort intuitions about what constitute important economic problems can be complicated. Together with bank runs, other problems such as currency attacks [3], or more general processes involving threat, illustrate this difficulty. The intuition derived from the empirical descriptive tradition – that outcomes depend on mutual confidence – may in the end be more compelling than conclusions of models attempting to formalize rational choice. Confidence is a concept that grows out of information and dynamics. In models, it may be omitted entirely, approximated with additional layers of variables and operators reflecting belief [4], or entailed in specific dynamical mechanistic models for learning or updating actions.

The original DD model [1] and its subsequent exposition [2] invoke but do not model expectations. As is known from the substantial literature on *folk theorems* for repeated games [5, 6, 7], expectation or implicit prior agreement can be a powerful tool for defending a very wide range of outcomes to a game as non-cooperative equilibria, but this very flexibility calls into question what is thereby learned from “solving” the game. The virtue of *not* formalizing expectations, however, is that it can allow the equilibrium analysis to address a rather subtle intuition that, against a background of contracts and trades that is mostly orderly, multiple equilibria do nonetheless exist and are in some sense indeterminate, and that this indeterminacy is an important source of endogenous risk. Using equilibrium as a base case (and to establish notation) we first repeat the standard two-equilibrium solution for a simplified bank-run model in strategic form.

We then address the problem of *contingently* rational expectations by expanding the strategic form into an extensive form in which knowledge about

the moves by nature can support subgame-perfect, non-cooperative, correlated equilibria. Our information model is one instance of a result of Morris and Shin [3] (MS) which appears to refute the intuition from the DD equilibrium analysis. MS show that in the presence of small imprecisions in agents' signals of the population state that underlie the correlated equilibrium, no strategy can outperform a simple threshold strategy, which then selects between the coordinated and uncoordinated equilibria, removing indeterminacy and its associated risk. A key feature of this result is that the amount of imprecision may be arbitrarily small, and its form may vary widely, and the MS theorem will still hold, even though a singular limit with no imprecision remains indeterminate as in the DD analysis. Such regular limits with removable singularities at the limit points are not uncommon in rational-expectations equilibria, but they raise a caution about the interpretation of theorems. As we have discussed [8] in regard to one of the better-known cases – the Hahn paradox, or what establishes the salvage value of a fiat money in long-term circulation – singular limits may signal the inherent *fragility* of rational-expectations strategies, and in real situations these fragilities can be exposed in explicit models of dynamics.

Therefore, in a third treatment, we re-formulate the process of identifying the non-cooperative rational-expectations equilibrium of MS as the dynamical convergence toward an Evolutionary Stable State [9] (ESS) in a replicator dynamic. The simple structure of the model we will adopt ensures that, in a suitable low-fluctuation limit, the ESS correspond to the Nash rational-expectations equilibria. While a dynamical search process will recover the MS decision rule for sufficiently large imprecisions and some prior distributions over population states, the limit of small imprecision leads to a regime shift for many (very reasonable) cases, in which the DD instability between coordinated equilibria and bank runs is restored. Because a replicator dynamic¹ requires the specification of an explicit update rule, within the context of any particular model class we obtain a description of behavior both at *and away from* fixed points. These permit us to go beyond the mere identification and description of equilibria [11], to a full treatment of the dynamics in multiple-equilibrium systems [12, 13, 14, 15, 16, 17], including metastable residence in different equilibria and punctuated escapes between equilibria.

¹The replicator dynamic is equivalent to Bayesian learning from repeated samples [10], so this formulation captures a wide range of cognitive as well as population-level processes for convergence to equilibria.

All of our results are carried out with highly generic modeling assumptions.

Our analysis upholds some qualitative claims from both DD and MS, but it does not support any of these as general claims, instead arguing that each captures a robust feature of process in limited contexts. Where context matters, our dynamical models also generally produce quantitative differences from the estimates of equilibrium properties provided by either DD or a constructive implementation of the MS decision mechanism. From DD, we recover the existence of two *relevant* steady states corresponding roughly to their high-liquidity coordinated equilibrium and their bank run. Banks can provide both liquidity and risks of bank runs, and agents can be rational to deposit in banks as long as the likelihood of runs is sufficiently rare. However, our coordinated equilibrium may not be close to the General Equilibrium that they propose as the Pareto-superior outcome. Finally – but *only* as a result of a full process model – we present a framework to determine how frequent runs should be. From MS, we recover the intuitive result that in conditions of high imprecision, a simple threshold-based decision rule can be robust and optimal, and may remove the indeterminacy of multiple equilibria. In many such cases the price system in the resulting non-cooperative equilibrium differs significantly from the DD General Equilibrium prices; and is by that measure inefficient. However, we add to the MS rational-expectations result a formalization of the concept of fragility: when rational expectations are replaced by generic and reasonable constructive decision rules, singular limits give way to transitions between dynamical regimes with unique or with multiple equilibria. At such transitions the appropriate description shifts from an emphasis on particular equilibria to an emphasis on distributions of outcomes and their temporal behaviors.

1.2 Basic framework: banks as creators of liquidity and risk of runs, and the claim of two equilibria

We use a simplified version of the Diamond-Dybvig banking model [1] to study multiplicity and indeterminacy of equilibria in all three descriptions. We simplify the forms of utility and replace several of their continuous-valued decision variables with discrete (binary) variables where no generality is lost in the analysis of the core problem.² The essential features of their model

²The simplification of the utility consists of omitting explicit time-discounting structure which increases the appeal of the narrative particular to bank runs but can be replaced

which we retain are these:

1. **Illiquid production:** A single good may be invested in production for either a short term (one period) or a long term (two periods). If withdrawn in the short term it returns only the original input. If withdrawn after the long term it yields a multiplier $R > 1$ per unit of original input. Each agent α in a continuum $\alpha \in [0, 1]$ is given an endowment of one unit of the input. Some investors will learn after investing that they must withdraw on the short term. The important feature of the model is that Arrow-Debreu securities do not exist to cover the risk of being someone who must withdraw early;³ the act of investing then makes the goods-supply illiquid because of the waiting time required to achieve the multiplier R .
2. **Banks as proxy investors:** Banks exist which can accept demand-deposits from agents and can invest on behalf of the agents. Deposit contracts that promise payouts $r > 1$ to agents who withdraw after one period, *subject to availability of funds*, can provide some or all of the liquidity that Arrow-Debreu securities would have provided (had they existed), but promised payouts $r > 1$ leave the banks susceptible to runs, in which some agents cannot be repaid. Following Ref. [1], we suppose that in cases where a fraction $> 1/r$ of agents attempt to withdraw, making the banks insolvent, a random ordering of the agents is generated, and the agents are fully-paid r in that order until the initial deposits are exhausted; the remaining agents are paid nothing.
3. **A natural risk of early withdrawal:** A natural process partitions the agents into a fraction $\theta \in [0, 1]$ which we term *Type-1*, who must withdraw after one period, and a remainder termed *Type-2*, who have the option to leave their deposits in the bank for both periods, but who may also withdraw after one period. Agents will be randomly

by general conditions on the saturation of a single utility of payouts. The continuous variables removed include amounts of the endowment invested, and amounts withdrawn, which in the original model [1] took binary boundary solutions anyway.

³Ref. [1] characterize the surprise signal that one must withdraw early as “private information” to justify the lack of Arrow-Debreu securities within a General Equilibrium framework. The model interpretation extends, however, to many other reasons for incomplete contracts.

partitioned by nature between the two types, and only learn their individual type after the payout contract r has been declared and deposits have been made. Early withdrawal by Type-2 agents to the point of insolvency constitutes a bank run.

The main result of Diamond and Dybvig is that, for $r \in (1, R)$, and a variety of situations that produce a measure of Type-1 agents $\theta < 1/r$ (few enough agents are forced by nature to withdraw early that insolvency is not assured), two Nash equilibria can exist for the joint decision to deposit/not-deposit and for the agents who learn they are Type-2 to run/not-run. In the high-yielding equilibrium all agents deposit, and no Type-2 agents run. In the low-yielding equilibrium all agents deposit and all Type-2 agents run. Whether the low-yielding equilibrium is a best response compared to not-depositing depends on the likelihood of runs, for which a generating process is not explicitly given in Ref. [1]. If θ takes a definite value $\langle \theta \rangle$ (which is a regular limit of θ sampled from progressively narrower distributions with mean $\langle \theta \rangle$), the payout $r = r_{\langle \theta \rangle} \equiv R / (1 + (R - 1) \langle \theta \rangle)$ delivers an assured return which is the risk-minimizing solution.⁴ At this solution, on the interval $r \in [1, r_{\langle \theta \rangle}]$ the social welfare, defined as the sum of expected utilities of payout, is monotone increasing for the high-yielding equilibrium, and monotone-decreasing for the bank-run equilibrium, since the mean of the latter is always 1 and the variance increases with r . The payout $r = 1$ gives no incentive to run, but is identical to each agent's investing individually, and so provides no liquidity-incentive to deposit. In this way the provision of liquidity and the risk associated with bank runs are intrinsically coupled (prior to the discussion of deposit insurance, which we do not enter here). The emphasis in Ref. [1] was on the ability of banking to recover the liquidity of the General Equilibrium contract in one of the Nash equilibria.

⁴The more general case, maximizing *ex ante* expected utility for agents who are required to deposit, maximizes the function

$$\langle \theta \rangle u(r) + (1 - \theta) u\left(\frac{(1 - \langle \theta \rangle r) R}{(1 - \langle \theta \rangle)}\right)$$

over r . The result is the condition reported in Ref's. [1, 2], that

$$u'(r) = Ru' \left(\frac{(1 - \langle \theta \rangle r) R}{(1 - \langle \theta \rangle)} \right).$$

1.3 Subgame-perfect equilibrium of an extensive-form game, and threshold strategies as a kind of regulator

In an explicitly-described extensive-form game, the incentive and risk associated with bank deposits and bank runs are captured as properties of *subgame-perfect* non-cooperative equilibria. On top of extensive-form formulations of multiple-equilibrium problems of the kind posed by Diamond and Dybvig, Morris and Shin [3] add the explicit observation of the population state – how many Type-1 and how many Type-2 agents – by the Type-2 agents as a part of the game specification, which enables them to consider properties of *correlated* subgame-perfect equilibria. The important additional feature added by Morris and Shin is imperfect observation of the population state, which has the crucial effect of further partitioning the Type-2 agents among (continuous) values of the observed signal.⁵ This smoothing permits them to prove that no strategy for the wait/run subgame dominates a simple threshold strategy, and that furthermore, the use of any such threshold strategy removes the degeneracy of the Diamond-Dybvig equilibria, yielding a unique correlated subgame-perfect equilibrium for any θ . We characterize their smoothing as a *regulator*; the width in the dispersion of observations of the population state does not matter to their formal proof of existence and uniqueness of equilibrium, but in reasonable abstractions of dynamics it will matter, and as the regulator becomes weak, the multiple equilibria and non-determinism of Diamond and Dybvig will re-emerge.

Fig. 1 shows an extensive-form game for the bank-run model, with the following sequence of moves:

1. In the first (continuous-valued) move, agents negotiate a payout r for early withdrawal from the bank, and make deposits.⁶ The bank is a

⁵In Ref. [3] imperfect observation is discussed in a wider context of formalizing information and beliefs. Here our concern is with the mechanistic consequences of dispersing population responses, which may be given comparable analyses in many contexts of inference, reinforcement learning, or population updating.

⁶Here for brevity we pass over how the negotiation is performed. Because a unique equilibrium value will be computable for the cases we consider, at the equilibrium it will be sufficient for all agents to declare it independently. In a larger discussion of convergence toward equilibrium, many mechanisms could be specified, along lines similar to those we describe below for the selection of threshold strategies, and the convergence of these mechanisms to deposit- or no-deposit-equilibria could be pursued in a more elaborate




Figure 1: An extensive form game for bank runs. In the first ply, the agents move, declaring a payout for early withdrawals $r \in [1, R]$. In the second ply, nature moves, selecting a fraction of Type-1 agents $\theta \in [0, 1]$, and randomly assigning agents within the two groups. (Filled wedges indicate continuous decision variables.) In the third ply, which is the sub-game, agents of Type-2 make a binary decision to run or wait.

strategic dummy, which will accept any value the agents declare, but enforces it in later rounds, along with implementing sorting protocols in cases of insolvency when only some agents can be paid. For simplicity we have all agents deposit as a rule of the game; this is a property of the noncooperative equilibrium for the cases we will consider, and was the boundary solution derived in Ref. [1].

2. In the second (continuous-valued) move, nature samples a measure θ to make Type-1 agents, from a distribution $\rho(\theta)$, which is common knowledge. The agents are randomly assigned types and informed privately of their types.
3. In the third (binary) move, agents either withdraw at the first period or do not withdraw. All Type-1 agents are required to withdraw early, and Type-2 agents withdraw according to Morris-Shin threshold strategies. We regard withdrawals as simultaneous, so the choices by Type-2

dynamical model.

agents are made non-cooperatively. While agents know the distribution $\rho(\theta)$ from which population-states are sampled, and each knows his own type, they do not know the exact measure θ . To obtain a signal corresponding to the Morris-Shin imperfect information, we have each Type-2 agent sample an integer number K agents from the population, and use the fraction $k/K \equiv \tilde{\theta}$ they find of Type-1 agents as a sample estimator for θ . The values $\tilde{\theta}$ will be binomially distributed with mean θ and variance $\theta(1-\theta)/K$. Agents also choose thresholds θ^* , withdraw early if $\tilde{\theta} < \theta^*$, and wait if $\tilde{\theta} \geq \theta^*$. (We describe below different mechanisms by which values for θ^* are chosen.) The magnitude of K will determine the resulting robustness or fragility of the threshold strategies.

1.3.1 Bayesian rational-expectations correlated equilibrium

For any value of r chosen in the first round of the extensive form, Type-2 agents may identify Bayesian rational-expectations threshold values θ^* as shown next. These define the equilibrium of the wait/run subgame, and the equilibrium of the outer move is then obtained by maximizing *ex ante* expected utility over r .

The conditional probabilities for sample estimators $\tilde{\theta}$ given actual population states θ are

$$p(\tilde{\theta} | \theta) = \theta^k (1-\theta)^{K-k} \binom{K}{k} \Big|_{\tilde{\theta}=k/K}. \quad (1)$$

From these and the common-knowledge distribution $\rho(\theta)$, each agent may form a Bayesian posterior distribution for the actual population state,

$$\rho(\theta | \tilde{\theta}) = \frac{\rho(\theta) p(\tilde{\theta} | \theta)}{\int_0^1 d\theta' \rho(\theta') p(\tilde{\theta} | \theta')}. \quad (2)$$

From any actual population state θ , because we assume a continuum of agents, the measure of agents who will withdraw early, using any common value of the threshold θ^* , will be given by

$$f_{\theta}^{(\theta^*)} \equiv \theta + (1-\theta) \sum_{\tilde{\theta} \geq \theta^*}^1 p(\tilde{\theta} | \theta). \quad (3)$$

For a population of such agents, the solvency threshold $\underline{\theta}$, beyond which early-withdrawing agents are uncertain to be paid and late-withdrawing agents are paid nothing, is given by $f_{\theta}^{(\theta^*)}\big|_{\theta=\underline{\theta}} = 1/r$.

A Type-2 agent who samples $\tilde{\theta}$, in a background population using threshold θ^* , will receive a Bayesian-posterior expected utility if his own threshold tells him to wait at that $\tilde{\theta}$, of

$$\langle u | \tilde{\theta} \rangle_{\text{wait}} = \int_0^{\underline{\theta}} d\theta \rho(\theta | \tilde{\theta}) u\left(\frac{(1 - f_{\theta}^{(\theta^*)}r)R}{1 - f_{\theta}^{(\theta^*)}}\right) + \int_{\underline{\theta}}^1 d\theta \rho(\theta | \tilde{\theta}) u(0). \quad (4)$$

Alternatively, if his own threshold tells him to run at that $\tilde{\theta}$, his expected payoff will be

$$\langle u | \tilde{\theta} \rangle_{\text{run}} = \int_0^{\underline{\theta}} d\theta \rho(\theta | \tilde{\theta}) u(r) + \int_{\underline{\theta}}^1 d\theta \rho(\theta | \tilde{\theta}) \left[\frac{1}{f_{\theta}^{(\theta^*)}r} u(r) + \left(1 - \frac{1}{f_{\theta}^{(\theta^*)}r}\right) u(0) \right]. \quad (5)$$

The Bayesian-Nash equilibrium value for θ^* is the value of $\tilde{\theta}$ for which the difference in expected utility,

$$\langle u | \tilde{\theta} \rangle_{\text{wait}} - \langle u | \tilde{\theta} \rangle_{\text{run}} = \int_0^1 d\theta \rho(\theta | \tilde{\theta}) \frac{u\left(\max\left(\frac{(1 - f_{\theta}^{(\theta^*)}r)R}{1 - f_{\theta}^{(\theta^*)}}, 0\right)\right) - u(r)}{\max(1, f_{\theta}^{(\theta^*)}r)}, \quad (6)$$

equals zero. The proof of existence of such equilibria is sketched in App. A. For large K and a sufficiently smooth prior distribution $\rho(\theta)$, the posterior $\rho(\theta | \tilde{\theta})$ is dominated by the distribution over sample estimators and produces a Nash θ^* for any r , which is asymptotically (in large K) independent of the prior. Note that, for such a narrow $\rho(\theta | \tilde{\theta})$, we expect that the Nash threshold will fall close to the value where the wait- and run-payoffs which are the arguments to the integral are equal, but with $f_{\theta}^{(\theta^*)}$ in place of θ^* – that is, it will produce $(r, f_{\theta}^{(\theta^*)})$ pairs near the risk-minimizing General Equilibrium contract.

The outer optimization of r then follows from properties of the utility, and uniqueness follows from the monotone-decreasing utility of the payout terms in Eq. (6) with r .

1.4 Dynamics and the fragility of the regulator; how bistability re-emerges in context

The rational-expectations correlated equilibrium is defined as a best-response of a single agent to a population already using the equilibrium strategy. It is not a constructive solution, but often it can be approximated dynamically (and in that sense, “constructed”) by embedding small sub-populations of agents with variant strategies in a background with a hypothesized strategy, and then using a replicator dynamic in which the fitness of each type is its expected payout to update the population numbers, thereby shifting the hypothesized strategy. The “fundamental theorem of evolutionary game theory” [18] states that the rest points of the replicator dynamic converge to the Nash equilibria as populations become homogeneous.⁷ The replicator dynamic provides a general way to introduce population variation into the analysis of strategies such as the Morris-Shin threshold strategies, and in particular to study the ways they may become “fragile” to dynamics that sample neighborhoods of the rational-expectations fixed point.

Here we demonstrate the behavior of a replicator dynamic on threshold strategies for a utility of the form $u(r) = r^\alpha$ for $0 \leq r \leq 1$ and $u(r) = 2 - r^{-\alpha}$ for $r \leq 1 \leq R$, and three distributions $\rho(\theta)$, shown in Fig. 2. The essential properties of the utility for this particular problem are finiteness of $u(0) = 0$, which makes strategies with nonzero risk of zero payout rationalizable, and $ru'(r)$ monotone decreasing for $r \geq 1$. The latter property is necessary for utilities with a single payout (from either time period) to produce solutions with interior maxima, as shown in App. A.

We introduce a distribution $\rho^*(\theta^*)$ over sub-populations of agents using different strategies θ^* . In examples we take this distribution to be binomial (on 15 elements), because the approximately Gaussian form of both ρ^* and the sample-estimator density $p(\tilde{\theta} | \theta)$ from Eq. (1) leads to a convolution which is also approximately Gaussian, and hence a population behavior

⁷Because the replicator dynamic with utility as fitness is equivalent to Bayesian updating [10], the dynamics described here apply equally well to reinforcement learning by Bayesian updating or to literal population-dynamic processes.

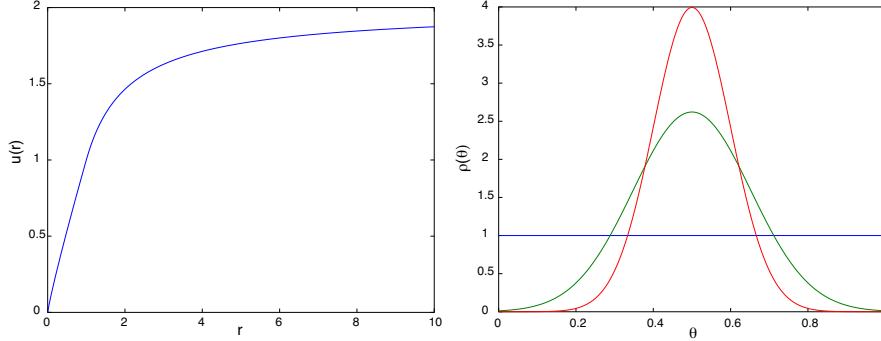


Figure 2: Left panel: Functional form $u(r) = r^\alpha$ for $0 \leq r \leq 1$ and $u(r) = 2 - r^{-\alpha}$ for $r \leq 1 \leq R$ for $\alpha = 0.9$ and $R = 10$. Right panel: three densities $\rho(\theta)$ to be considered, all with $\langle \rho \rangle = 0.5$. Blue is uniform on $\theta \in [0, 1]$; green is Gaussian standard deviation ≈ 0.152 , red is Gaussian with standard deviation = 0.1.

comparable to the behavior of a homogeneous population of agents using the mean threshold

$$\langle \theta^* \rangle = \int d\theta^* \rho(\theta^*) \theta^* \quad (7)$$

with a slightly increased variance. The aggregate measure of agents who withdraw early in a population with measure θ of Type-1 agents (randomly assigned) is simply the integral

$$\int d\theta^* \rho(\theta^*) f_\theta^{(\theta^*)} = f_\theta. \quad (8)$$

The expected utility within each sub-population relative to the aggregate, averaged over the actual distribution $\rho(\theta)$, is then given by

$$\begin{aligned} \langle u | \theta^* \rangle - \langle u \rangle &= \int_0^\theta d\theta \rho(\theta) \left(f_\theta - f_\theta^{(\theta^*)} \right) \left[u \left(\frac{(1 - f_\theta r) R}{1 - f_\theta} \right) - u(r) \right] \\ &\quad - \left(\frac{u(r) - u(0)}{r} \right) \int_\theta^1 d\theta \rho(\theta) \left(1 - \frac{f_\theta^{(\theta^*)}}{f_\theta} \right). \end{aligned} \quad (9)$$

The resulting shift in the aggregate threshold value $\langle \theta^* \rangle$, if this population is repeatedly subjected to a replicator dynamic with sub-population numbers

increased or decreased in proportion to relative utilities [18], is given by

$$\frac{d\langle\theta^*\rangle}{dt} \equiv \frac{d}{dt} \int d\theta^* \rho(\theta^*) \theta^* = \int d\theta^* \rho(\theta^*) [\langle u | \theta^* \rangle - \langle u \rangle] \theta^*. \quad (10)$$

Here we use such a replicator dynamic to identify the fixed points θ^* as a function of payout r in the subgame, and then simply maximize expected utility along the subgame-perfect (r, θ^*) contour in the distribution $\rho(\theta)$, to define an equilibrium for the extensive-form game.

1.4.1 The uniform distribution: uniqueness of ESS and the corresponding subgame-perfect correlated equilibrium

The properties of such a mixed population, and their behavior under a replicator dynamic, for a payout $r = 1.2$ in the subgame, are shown for the uniform distribution $\rho(\theta)$ in Fig. 3.⁸ The functional dependence of the aggregate measure of early-withdrawing agents f_θ closely resembles that of each sub-population $f_\theta^{(\theta^*)}$ (left panel). The relative utilities (9) for $\theta^*r \ll 1$ (middle panel) are maximized near the population mean $\langle\theta^*\rangle$, as agents whose thresholds are lower miss out on yield R for a fraction of the samples θ , while those whose thresholds are higher often receive zero payout by waiting to withdraw, while a bank run by the majority of agents makes the bank insolvent before period-two. For $\theta^*r \geq 1$, Type-1 agents alone suffice to make the bank insolvent before most Type-2 agents run, and the lowest thresholds fare best from rare low signals $\tilde{\theta}$. The resulting velocity (10) for the population-mean (right panel), shows a unique stable fixed point for $\langle\theta^*\rangle$ at this (and indeed at all) r , which is the behavior that would be expected from the theorem of Ref. [3] for the rational-expectations equilibrium.

Fig. 4 shows the expected utility

$$\langle u \rangle = u(r) + \int_0^1 d\theta \rho(\theta) \max \left(1 - f_\theta, 1 - \frac{1}{r} \right) \left[u \left(\max \left(\frac{(1 - f_\theta^{(\theta^*)} r) R}{1 - f_\theta^{(\theta^*)}}, 0 \right) \right) - u(r) \right] \quad (11)$$

⁸We use the value $r = 1.2$, which is larger than the equilibrium payout value $r \approx 1.1$, because its interesting transitions occur in the center of the θ range where they are easier to view in graphs. We also use a distribution $\rho^*(\theta^*)$ in the first two panels of Fig. 3, to illustrate population structure and the qualitative behavior of relative utility, which is 10x wider than the distribution used to identify the replicator-dynamic velocity and to closely approximate the Nash fixed points in the third panel and later figures.




Figure 3: Identification of the fixed-point threshold θ^* for a given r in the wait/run subgame. Left panel: population structure. 15 thresholds parametrize contours $f_\theta^{(\theta^*)}$ in sub-populations (fine curves), with (binomial) weight $\rho^*(\theta^*)$ shown shaded. Aggregate population response f_θ (heavy line) remains nearly that for a Gaussian sample estimator, but with a slightly increased half-width. Half-way point for Type-2 agents who run, and the solvency limit $f_\theta = 1/r$, are shown as horizontal lines. Middle panel: utilities for each sub-population relative to the mean are graphed (height and color scale) as a function of the population-mean $\langle\theta^*\rangle$, and the sub-population offset $(\theta^* - \langle\theta^*\rangle)$. Right panel: the velocity of the population-mean $d\langle\theta^*\rangle/dt$ if this population structure is updated using a replicator dynamic. The unique, stable zero crossing is the ESS which converges to the Nash best-response as $\rho^*(\theta^*)$ is made narrow.

along the subgame-perfect (r, θ^*) contour (left panel), and the nearly-linear dependence of r on θ^* along this contour (right panel). The equilibrium value r for any θ^* grows more slowly with decreasing θ^* than a naïve estimate

$$r^* \equiv \frac{R}{1 + (R-1)f_{\theta^*}^{(\theta^*)}} = \frac{R}{1 + (R-1)(1 + \theta^*)/2}, \quad (12)$$

which places the inflection point of the curve $f_\theta^{(\theta^*)}$ at the point where the wait- and run-payoffs $R(1 - f_\theta^{(\theta^*)}r) / (1 - f_\theta^{(\theta^*)}) \Big|_{\theta=\theta^*}$ and r are equal. App. B compares the efficiency of this equilibrium to that of the risk-minimizing competitive equilibrium of Sec. 1.2. The smaller-than-expected r required to produce runs at small θ^* reflects asymmetry between a fixed *marginal* loss of running too early, when $f_\theta \approx \theta$, and the *absolute* loss $\sim (1 - \theta^*)$ from running too late, when $f_\theta \approx 1$, which increases with decreasing θ^* .




Figure 4: Left panel: expected utility (11) in the uniform distribution $\rho(\theta)$, along the subgame-perfect (r, θ^*) contour given by the rest points of the replicator dynamic. Marker (+) indicates the interior maximum which is the equilibrium of the extensive-form game. Utility level at $r = 1$ (red line), corresponding to the no-deposit outcome, shown for reference. Right panel: locus of best-response pairs (r, θ^*) (blue) compared to the naïve expectation (12) (green).

1.4.2 Non-uniform $\rho(\theta)$ and the restoration of bistability

The uniqueness and stability of threshold equilibria, exhibited for uniform $\rho(\theta)$, can fail for quite ordinary non-uniform $\rho(\theta)$ which increase the frequency of losses from waiting during a run, relative to gains from avoiding too-early sacrifice of yield R . We illustrate this by comparing the replicator subgame rest points for three Gaussian $\rho(\theta)$ with mean $\langle \theta \rangle = 1/2$ and different variances, shown in Fig. 2. Standard deviations ∞ (the uniform distribution), 0.152 (critical), and 0.1 (bistable) are shown.

Fig. 5 shows the sign of the velocity for $d\langle \theta^* \rangle / dt$ from Eq. (10) as a function of r and $\langle \theta^* \rangle$ for the three cases. Stable $\langle \theta^* \rangle$ result at any r when positive velocities fall to the left of negative velocities in the plot. For the uniform distribution fixed points are stable at all r . The Gaussian with standard deviation = 0.152 shows a classic critical point, where the threshold $\langle \theta^* \rangle$ is marginally stable and infinitely sensitive to $r \approx 1.225$. The Gaussian with standard deviation = 0.1 shows the characteristic “spinodal” behavior associated with bistability and hysteresis. For $r < 1.12$ a single high-yielding equilibrium exists. For $1.12 < r < 1.24$, a high-yielding equilibrium co-exists

with a low-yielding equilibrium corresponding to the Diamond-Dybvig bank run; between these two is an unstable fixed point. The basin of attraction of the bank-run equilibrium increases with increasing r , until for $r > 1.24$ it is the only equilibrium.

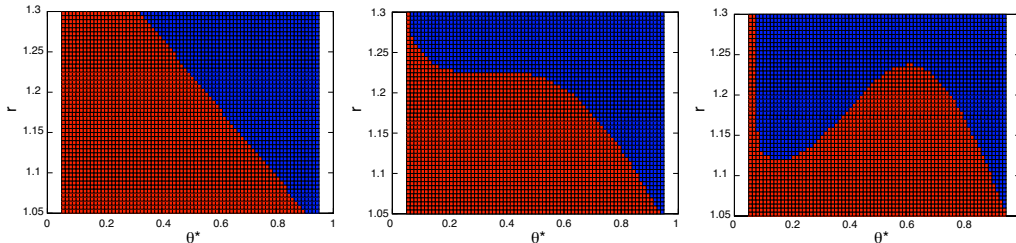


Figure 5: Signs of the evolutionary velocity $d\langle\theta^*\rangle/dt$ for the three prior distributions shown in Fig. 2. Each horizontal cross-section (single r value) corresponds to a curve like that in the right panel of Fig. 3. Red is positive velocity; blue is negative. Left panel: the uniform distribution in which each r has a unique stable fixed point θ^* . Middle panel: Gaussian $\rho(\theta)$ with standard deviation ≈ 0.152 which is the critical value for onset of bistability. Right panel: Gaussian $\rho(\theta)$ with standard deviation = 0.1 showing two stable fixed points and one unstable fixed point over the range $r \in [1.12, 1.24]$.

Is the region of bistability relevant to equilibrium for the full extensive-form game? That is, does the original intuition developed by Diamond-Dybvig for equilibrium, that agents can be rational to deposit in banks susceptible to runs, emerge in our treatment of threshold equilibria? Fig. 6 shows the optimization problem for expected utility (11), along the high-yielding subgame-perfect (r, θ^*) contour, where this contour exists. An interior maximizer of $\langle u \rangle$ does indeed fall within the domain of bistability, at $r \approx 1.1162$ (left panel), and the velocity profile $d\langle\theta^*\rangle/dt$ (middle panel) is robustly that for a bistable solution in an open neighborhood of this equilibrium. Thus, within mild and generic assumptions about testing-and-reward dynamics to identify equilibria, whether unique stable equilibria, or multiple equilibria with risk, are recovered depends on the distribution of population states in relation to the distribution of sample-estimators.

Finally, we note that the dynamical treatment adds the last feature required for a consistent analysis of the multiplicity or stability of equilibria: In the stochastic version of a replicator dynamic appropriate to real processes

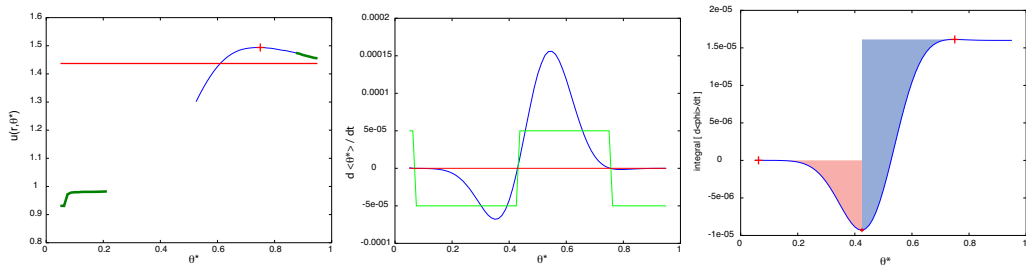


Figure 6: Left panel: utility along high (blue) and low (green-heavy) (r, θ^*) fixed-point contours in the wait/run subgame. Solutions drop out where the hysteretic threshold is crossed, and one or the other branch ceases to exist. The region with no stable thresholds covers much of the support of the prior distribution of $\rho(\theta)$. The utility-maximizing r value (marker) lies along the high-threshold contour, *within* the region of bistability. Middle panel: evolutionary velocities $d\langle\theta^*\rangle/dt$ (blue) at the subgame-perfect equilibrium r value, to be contrasted with the third panel of Fig. 3 for the uniform distribution. Red is zero (for reference), and green shows the sign of the velocity, which is a cross section at $r = 1.1162$ of panel three of Fig. 5. Right panel: integral of the evolutionary velocity which defines a “fitness landscape”. Time integrals of the potential, which correspond to appropriate velocity-weighted red and blue shaded regions, are proportional to the log-escape rates in a stochastic evolutionary dynamic; the difference of these areas determines the log-ratio of occupancy of the high and low equilibria (markers).

(such as finite-population sampling) that continuously generate variation in parameters such as θ^* , the integral of the velocity field defines a *potential* (often called a “fitness landscape”⁹, shown in the right panel of Fig. 6) for transitions between the two fitness maxima. It is a basic result for escape processes in Markov chains [23, 24, 25], including stochastic evolutionary games [26, 16, 17], that for one-dimensional processes the log-escape rates are proportional (with a minus sign) to the time-integrals of this potential between the equilibria and the minimum (indicated schematically by shading in the figure). The log-ratios of escape rates equal the log-ratios of occupancy of the two equilibria in the long run, which are thus proportional to the dif-

⁹This usage follows Sewall Wright [19, 20, 21, 22].

ference between the two integrated areas. The larger integrated area for the high-yielding equilibrium in Fig. 6 will cause most time to be spent at this point, and the ratio of time spent in the high-yielding versus the bank-run equilibrium will grow exponentially as the standard deviation of the sampling distribution $\rho^*(\theta^*)$ is taken to zero. The theory of escapes using this formalism has been applied to conventional contracts [12, 13, 14, 15], and this formula for the relative occupancy of states has been used in a context similar to annealing to define an approach to equilibrium refinement [11].

In a modeling approach that requires commitment to particular, constructive, dynamic mechanisms of this kind for setting θ^* , we are able to characterize not only the equilibria but also the non-equilibrium behavior that determines the frequency of runs, while retaining some generality of results. More fundamentally, we may shift the focus of description between the monostable and bistable regimes. In the monostable regime, the fixed point and its neighborhood govern almost all observables, in keeping with an emphasis on equilibrium. In the bistable regime the relevant objects of description may become the distributions over equilibria and transitions between them, which only in some cases reduce naturally to criteria for equilibrium refinement.

1.5 Why do threshold strategies become fragile in this way?

The purpose of our analysis is to show how apparent anomalous sensitivities in paradigms or even theorems can reflect important context-dependence in the formalization of the relevant concepts, and in some cases this context can be made explicit and even quantified. One anomalous sensitivity is the mismatch between the apparent demonstration of multiple equilibria and risk in the Diamond-Dybvig bank-run model, contrasted with the general proof of existence of unique equilibria in models of similar type by Morris and Shin, with a slightly more explicit strategic treatment and the new element of signal imprecision. A second sensitivity is the apparent removable singularity of threshold models: any dispersion in signals implies a unique equilibrium, but exact signaling of the population state to all agents removes the conditions needed for the theorem. The key element of the Morris-Shin construction is the possibility for *a part* of a population of identical agents to run at the optimal threshold. The equilibrium at such a threshold sacrifices efficiency

compared to the competitive equilibrium because overall production is reduced, but makes well-defined the marginal value of shifts in the threshold. However, the limit as signals are made arbitrarily precise is one where the sensitivity of the population state to threshold changes becomes infinite – equivalent to arbitrarily defining the value of a unit step function to be $1/2$ at its transition. A theorem that makes essential use of partial runs, in a limit where the measure of partial runs becomes infinitely sensitive to the optimization parameter, does not seem a sufficiently strong result to invalidate the intuition that multiple equilibria exist and carry risk in realistic situations. We now explain how the replicator dynamic formalizes this notion of a “sufficiently strong” regularization mechanism, and why it recovers multiple equilibria, not only at the singular limit, but in an open domain of smooth partial-run transitions.

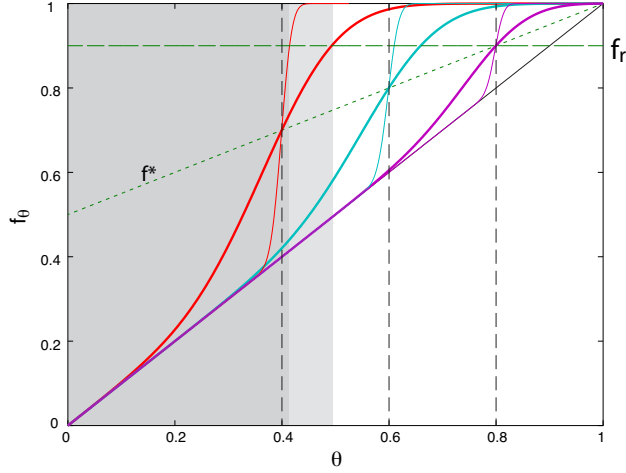


Figure 7: The transition from robust to fragile threshold strategies in the replicator dynamic. Here r is chosen so that at the inflection point $f_r \equiv f_\theta^{(\theta^*)}|_{\theta=\theta^*} = 0.9$, the wait- and run-payouts are equal, $R(1 - f_r r) / (1 - f_r) = r$. Three population-aggregate thresholds $\langle \theta^* \rangle = 0.4, 0.6, 0.8$ are illustrated, with the last value approximating the Nash threshold at this r . Broad transitions $f_\theta^{(\theta^*)}$ (at $K = 20$) are shown as heavy solid colored lines, and narrow transitions $f_\theta^{(\theta^*)}$ (at $K = 1000$) are the corresponding light colored lines. About the background $\langle \theta^* \rangle = 0.4$, the range of θ values where agents are better to wait than to run are shaded, and the θ values where they are better to run are unshaded. For the gradual transition at $K = 20$ (light shading), the range where waiting is advantageous extends almost 0.1 above $\langle \theta^* \rangle$ toward the equilibrium value $\theta^* = 0.8$, favoring best-responses larger than $\langle \theta^* \rangle$ unless the underlying distribution $\rho(\theta)$ is strongly increasing. For the sharp transition at $K = 1000$ (dark shading), the range where waiting is advantageous extends only ~ 0.01 above $\langle \theta^* \rangle$, and the disadvantage of waiting during a run covers almost the entire interval up to $\theta^* = 0.8$. Therefore even small positive $d\rho(\theta)/d\theta$ will create negative velocities in the replicator dynamic near $\langle \theta^* \rangle = 0.4$.

Threshold strategies function as regulators by balancing the marginal loss of expected utility from choosing too low a threshold relative to the aggregate population and sacrificing the multiplier R , against the alternative loss from choosing too high a threshold and receiving payout 0 in period two when a run has left the bank insolvent. The approximate magnitudes of these effects can be estimated from the behavior of $f_{\theta}^{(\theta^*)}$ shown in Eq. (3), that $f_{\theta}^{(\theta^*)} \rightarrow \theta$ for $\theta < \theta^*$, and $f_{\theta}^{(\theta^*)} \rightarrow 1$ for $\theta > \theta^*$. These changes occur in a narrow range of θ when $p(\tilde{\theta} | \theta)$ is narrow.

Agents who run at $\theta^* < \langle \theta^* \rangle$ receive $u(r)$ rather than $\approx u(R(1 - \theta^*r) / (1 - \theta^*))$ for a measure of events $\sim \rho(\theta^*) (\langle \theta^* \rangle - \theta^*)$. Agents who run at $\theta^* > \langle \theta^* \rangle$ receive $u(0)$ rather than a lottery over $u(0)$ and $u(r)$ with mean payout 1, for a measure of events $\sim \rho(\theta^*) (\theta^* - \langle \theta^* \rangle)$. Since the former loss is monotone decreasing in θ^* while the latter is fixed, in a uniform distribution there will be a unique fixed point where they are equal.

If, however, we replace the uniform distribution by a more concentrated distribution such as a Gaussian shown in Fig. 2, the measure $\rho(\theta)_{\theta > \langle \theta^* \rangle}$ of loss events from running late exceeds the measure $\rho(\theta)_{\theta < \langle \theta^* \rangle}$ from running early where $d\rho/d\theta > 0$, and the converse where $d\rho/d\theta < 0$. With sufficient curvature in $\rho(\theta)$, the interior fixed point splits into two fixed points, but the latter is unstable because the best response to any aggregate value $\langle \theta^* \rangle$ at this lower fixed point is always a still-lower θ^* , which protects agents from waiting during bank runs.

For wide $p(\tilde{\theta} | \theta)$ and “robust” regulators, the transition region of $f_{\theta}^{(\theta^*)}$ is wide and the marginal utility is determined by the small quantity $df_{\theta}^{(\theta^*)}/d\theta^*$. For narrow $p(\tilde{\theta} | \theta)$ and “fragile” regulators, the transition is abrupt and marginal utility is dominated by the two extreme behaviors described above. The transition from robust to fragile is illustrated in Fig. 7.

2 On Models and Theories

In this essay we claim no more than having taken a next reasonable step in extending two noncooperative game theory analyses of the potential for “rational bubbles”. But both of these stimulating papers have adhered to the basis of equilibrium analysis. We have taken their work a step further in the consideration of a specific dynamical context given by a replicator

equation and we have provided a dynamic sensitivity analysis. However when reading books such as Kindleberger [27] or old classics such as MacKay [28] or a description of instability such as that supplied by Minsky [29], one realizes that in spite of the attractiveness of our replicator story added to the two equilibrium analyses these three models only go a small way to offering a general dynamic theory of bubble behavior. We are in complete accord with the view that one should push the limits of economic and game theoretic analysis of *homo ludens* as far as possible. But when confronted with items such as incomplete knowledge of the rules of the game and a multiplicity of socio-psychological phenomena to account for we are still far from understanding context and dynamics.

A Proof of existence of interior threshold strategies

Lemma: If the utility satisfies $ru'(r)$ monotone decreasing on $r \geq 1$, then the expected utility (11) along the subgame-perfect (r, θ^*) contour has an interior maximum.

Proof: Existence follows from the signs of derivatives of $\langle u \rangle$ at $r = 1$ and $r = R$.

For $r = 1$, there is never an incentive to run, while for all $\theta < 1$ there is an incentive to wait, so the subgame-perfect $\theta^*|_{r \rightarrow 1} > 1$ and $f_\theta = \theta$ everywhere. Then the expected utility becomes

$$\langle u \rangle|_{r=1} = \langle \theta \rangle u(1) + (1 - \langle \theta \rangle) u(R), \quad (13)$$

and the marginal expected utility may be computed and shown to be

$$\left. \frac{d\langle u \rangle}{dr} \right|_{r=1} = \langle \theta \rangle [u'(1) - Ru'(R)]. \quad (14)$$

As long as $R > 1$ and $ru'(r)$ is monotone decreasing on $r \geq 1$, this quantity is always positive.¹⁰

¹⁰This restriction on u says that the utility increases more slowly than the logarithm on $1 < r < R$. Combining this requirement with a lower bound $u'(0) > -\infty$ motivates the form of utility shown in Fig. 2. In the original model of Ref. [1] the same effect was achieved by discounting period-two payouts in the utility relative to period-one payouts,

Conversely, if $r = R$, The payout $R(1 - f_\theta r) / (1 - f_\theta) < R$ for all θ , so there is never an incentive to wait, making $\theta^*|_{r \rightarrow R} < 0$ and $f_\theta \rightarrow 1$ everywhere. The utility therefore becomes

$$\langle u \rangle|_{r=R} = \left(1 - \frac{1}{R}\right) u(0) + \frac{1}{R} u(R), \quad (15)$$

and the marginal expected utility becomes

$$\left. \frac{d\langle u \rangle}{dr} \right|_{r=R} = -\frac{[u(R) - u(0)] - Ru'(R)}{R^2}. \quad (16)$$

For any concave utility, $[u(r) - u(0)] - ru'(r) > 0$, so the derivative (16) is negative and $\langle u \rangle$ has an interior maximum.

B Efficiency of threshold equilibria in relation to competitive equilibria

To illustrate the inefficiency of the threshold equilibria without reference to the detailed form of the utility, we consider the risk-minimizing equilibrium. Readers of Ref. [2] will recognize that the competitive equilibrium for utilities of our form with $\alpha = 1$ are recovered in the following results by replacing $R \rightarrow \sqrt{R}$, and the same efficiency arguments go through.

In the risk-minimizing competitive equilibrium of Sec. 1.2, Type-2 agents never ran and the payout to those waiting until period-two to withdraw, in any population state θ , was

$$r_{\text{Competitive}} = \frac{R}{1 + (R - 1)\theta}. \quad (17)$$

In contrast, for the threshold strategy used here, the inflection point of the function $f_\theta^{(\theta^*)}$ occurs at $\theta = \theta^*$, where $f = (1 + \theta^*)/2$ (shown as the line with slope 1/2 in the left panel of Fig. 3). When this inflection point closely

but apart from some narrative appeal to the case of bank runs, the method used is not critical to the result. The reason such a strong requirement exists on utility, however, is that early withdrawal reduces period-two payout proportionally, to $(1 - f_\theta r) R$, amplifying the second marginal utility in Eq. (14) by R .

approximates the subgame-perfect threshold θ^* , as it does in the numerical examples, the payout to agents who withdraw early is

$$r_{\text{Threshold}} = \frac{R}{1 + (R - 1)(1 + \theta^*)/2}. \quad (18)$$

For the sake of comparison, considering a narrow distribution where $\theta^* \approx \langle \theta \rangle \approx \theta$ for almost-all samples θ , the threshold equilibrium is less efficient than the competitive equilibrium by

$$\frac{1}{r_{\text{Threshold}}} - \frac{1}{r_{\text{Competitive}}} = \left(1 - \frac{1}{R}\right) \frac{1 - \langle \theta \rangle}{2}. \quad (19)$$

One half of Type-2 agents withdraw “needlessly” to implement the mechanism of regulation, reducing their own payouts because they decrease the total production relative to a competitive equilibrium.

References

- [1] Douglas W. Diamond and Philip H. Dybvig. Bank runs, deposit insurance, and liquidity. *J. Pol. Econ.*, 91:401–419, 1983.
- [2] Douglas W. Diamond. Banks and liquidity creation: a simple exposition of the Diamond-Dybvig model. *Economic Quarterly*, 93:189, 2007.
- [3] Stephen Morris and Hyun Song Shin. Unique equilibrium in a model fo self-fulfilling currency attacks. *Amer. Econ. Rev.*, 88:587–597, 1998.
- [4] Atsushi Kajii and Stephen Morris. The robustness of equilibria to incomplete information. *Econometrica*, 65:1283–1309, 1997.
- [5] Drew Fudenberg and Eric Maskin. The folk theorem in repeated games with discounting or with incomplete information. *Econometrica*, 54:533–554, 1991.
- [6] A Rubinstein and A. Wolinsky. Remarks on infinitely repeated extensive-form games. *Games Econ. Behavior*, 9:110–115, 1995.
- [7] Drew Fudenberg and Jean Tirole. *Game Theory*. MIT press, Cambridge, Mass, 1991.

- [8] Eric Smith and Martin Shubik. Endogenizing the provision of money: costs of commodity and fiat monies in relation to the valuation of trade. *J. Math. Econ.*, 47:508–530, 2011.
- [9] John Maynard Smith and George R. Price. The logic of animal conflict. *Nature*, 246:15–18, 1973.
- [10] Cosma Rohilla Shalizi. Bayesian learning, evolutionary search, and information theory. 2008.
- [11] Michihiro Kandori, George J. Mailath, and Rafael Rob. Learning, mutation, and long run equilibria in games. *Econometrica*, 61:29–56, 1993.
- [12] H. Peyton Young. The evolution of conventions. *Econometrica*, 61:57–84, 1993.
- [13] H. Peyton Young. The economics of conventions. *J. Econ. Perspectives*, 10:105–122, 1996.
- [14] Conventional contracts. Young, h. Peyton. *Rev. Econ. Studies*, 65:776–792, 1998.
- [15] H. Peyton Young and Mary A. Burke. Competition and custom in economic contracts: a case study of Illinois agriculture. *Amer. Econ. Rev.*, 91:559–573, 2001.
- [16] Eric Smith. Symmetry and collective fluctuations in evolutionary games. submitted, 2012.
- [17] Eric Smith and Supriya Krishnamurthy. Bridging scales with generating functions, and the large-deviations analysis of stochastic evolutionary game theory. submitted, 2012.
- [18] Josef Hofbauer and Karl Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge U. Press, New York, 1998.
- [19] Sewall Wright. *Evolution and the genetics of populations: genetics and biometric foundations vol. 1*. U. Chicago Press, Chicago, IL, 1984.
- [20] Sewall Wright. *Evolution and the genetics of populations: genetics and biometric foundations vol. 2 (The theory of gene frequencies)*. U. Chicago Press, Chicago, IL, 1984.

- [21] Sewall Wright. *Evolution and the genetics of populations: genetics and biometric foundations vol. 3 (Experimental results and evolutionary deductions)*. U. Chicago Press, Chicago, IL, 1984.
- [22] Sewall Wright. *Evolution and the genetics of populations: genetics and biometric foundations vol. 4 (Variability within and among natural populations)*. U. Chicago Press, Chicago, IL, 1984.
- [23] M. I. Freidlin and A. D. Wentzell. *Random perturbations in dynamical systems*. Springer, New York, second edition, 1998.
- [24] Robert S. Maier and D. L. Stein. Escape problem for irreversible systems. *Phys. Rev. E*, 48:931–938, 1993.
- [25] Eric Smith. Large-deviation principles, stochastic effective actions, path entropies, and the structure and meaning of thermodynamic descriptions. *Rep. Prog. Phys.*, 74:046601, 2011.
- [26] Josef Hofbauer and Karl Sigmund. Evolutionary game dynamics. *Bull. Am. Math. Soc.*, 40:479–519, 2003.
- [27] Charles Kindleberger. *Manias, Panics, and Crashes: a History of Financial Crises*. Wiley, New York, 1978.
- [28] Charles MacKay. *Extraordinary population delusions and the madness of crowds*. Harmony Books, New York, (1841) 1980. with a foreword by Andrew Tobias.
- [29] Hyman Minsky. *Stabilizing an unstable economy*. McGraw Hill, New York, 2008.