

Multiple Equilibria

Philip H. Dybvig¹
Washington University in Saint Louis
dybvig@wustl.edu

1 October, 2023
based on the lecture presented 8 December, 2022

¹I am grateful for helpful comments from Dilip Abreu, Ben Bernanke, Doug Diamond, Weiting Hu, Gerry Jaynes, Todd Keister, Shu Li, Manju Puri, Karl Shell, and Chester Spatt. This written version is based on the lecture Dybvig (2022). Dybvig (2023) is another version of the lecture, intended for academic economists rather than a general audience. Copyright The Nobel Foundation 2022

Bank failures did considerable economic damage during the Great Depression. This empirical result was documented by Ben Bernanke (Bernanke (1983)), and the theory by Doug Diamond and myself (Diamond and Dybvig (1983)) provided a plausible mechanism for this result. The 2022 Sveriges Riksbank Prize in Economic Science was awarded to the three of us for “research on banks and financial crises,” focusing on these two papers and Doug’s related paper on delegated monitoring (Diamond (1984)). My lecture will focus on (having) multiple equilibria, an important feature of the model in my Prize paper with Doug. Multiple equilibria had also been a theme in my prior research, and I want to give credit to my co-authors in those prior papers for their contributions to my thinking.

As Doug mentioned,² our model has an equilibrium with bank runs and an equilibrium without bank runs. There are prior examples of multiple equilibria sprinkled through the literature in economics. However, I think that at the time we wrote our paper, having multiple equilibria was largely viewed as a defect. In the introductory macro course, we were told that if you don’t have a unique equilibrium, it’s not a real economic model, because you can’t do prediction. Basically, they wanted to say that we need to have enough equations to determine equilibrium or else something is missing. Like a lot of important-sounding assertions, this assertion doesn’t stand up to close examination. When you’re a first-year graduate student, you need something for a foundation, and it would be too confusing to have all the assumptions up for debate when first learning the material. Mostly, we take assertions in first-year classes as absolute truth. Some simplifications are corrected later, but in other cases we carry these ideas through our careers.

I have no conflicts to report. Actually, I like it when the medical school people give this notice because it’s important. I’ve seen papers by people who had big conflicts that they should have reported, and they didn’t. If a scholar is presenting a paper on the financial crisis and was working as an employee or consultant to a firm that blew up, we need to know that, and perhaps we won’t take it seriously when the scholar has a model in which no market participants did anything wrong or de-emphasizes the firm’s role in the crisis. For these reasons, I’m getting into the habit of including this

²References to what Doug or Ben mentioned refer to their Nobel Lectures, presented before mine: Diamond (2022, 2023) and Bernanke (2022, 2023).




Figure 1: Stephen A. Ross

disclosure.

My talk is dedicated to Steve Ross (figure 1), who was advisor to Doug and me, and was just the best advisor ever. I can't conceive of how I would have developed as a scholar (or a person) without his beneficial influence. He was a great financial scholar who set in place the basic framework and even the name of agency problems in finance and economics (Ross (1973)), and with John Cox he also introduced the concept of risk-neutral probabilities (martingale pricing) that is widely used by academics and practitioners (Cox and Ross (1976)). He also made many other important contributions, including his Arbitrage Pricing Theory (e.g. Ross (1976a) and Chen, Roll, and Ross (1986)), which is the foundation of factor pricing models, the Fundamental Theorem of Asset Pricing (Dybvig and Ross (1987), building on Ross (1976b,1978)), binomial option pricing (Cox, Ross, and Rubinstein (1979)), and widely-used models and theorems about the term structure of interest rates (e.g. Cox, Ingersoll, and Ross (1981, 1985a, and 1985b) and Dybvig, Ingersoll, and Ross (1996)). Everything I have done in my research was impacted positively by what I learned from Steve.

I will start by talking about a couple of common definitions of equilibrium in economics, followed by a discussion of the role of multiple equilibria in the Prize paper with Doug. Then I will talk about some roots of my thinking,

with a focus on multiple equilibria in my prior papers with Chester Spatt and Gerry Jaynes. I will also discuss the sunspots model of Dave Cass and Karl Shell (Cass and Shell (1983)), which we referred to in the Prize paper to answer the puzzle of why people might invest in a bank when a run is expected. (The answer to the puzzle is that the run could happen with small probability, as a function of some commonly observed noise.) Then, I want to talk about a couple of ideas for future research.

Okay, what is equilibrium? I suspect that the intuition about equilibrium in economics comes from physics. In Newtonian physics, if you have a ball on a surface and subject to a uniform gravitational force from below, there is an equilibrium if the surface is locally flat.³ If, where it is flat, it's a strict local minimum, then it's going to be a stable equilibrium in the sense that it takes energy to move away from the equilibrium, and if you move a small enough distance away, it's going to stay nearby. If the surface is locally flat, but goes down in at least one direction, then the equilibrium is unstable, because moving downhill does not take any energy and once you move away a little bit, the slope keeps you moving away. While considering stability of equilibrium seems out of fashion in economics, it is consistent with what we do in the Prize paper, which focuses on the stable equilibria and ignores an unstable equilibrium with randomization.

Two popular definitions of equilibrium in economics, still in use now, are competitive equilibrium and Nash equilibrium. There are many, many other definitions of equilibrium we use in Economics, and different definitions are appropriate in different contexts, but let's just consider these two classical definitions. Competitive equilibrium makes intuitive sense when each agent is small compared to the economy and does not behave strategically, for example, when each agent cares only about the price and not about what choices other agents make given the price. Competitive agents take prices as given and they look at the prices in the economy and then ask how much to demand at that price. An equilibrium is a point at which supply equals demand, i.e., agents' excess demands add up to zero.

Another definition, the primary definition used in this Lecture, is Nash equi-

³To make the claims in the text rigorous, we have to have some regularity, such as the elevation of the surface being twice differentiable in its horizontal position.

librium, which Doug mentioned. Nash equilibrium was named after the great mathematician John Nash, who also won the Economics Prize. In Nash equilibrium, each agent makes a choice that is optimal given the choices of the other agents. When we model a single agent, we write down an optimization problem with choice variables, an objective function, and constraints. The objective function could be to maximize profit or the net present value of profits over time, or alternatively to maximize utility. Utility is a way of representing preferences. We often assume in economics that agents act as if they are maximizing a utility function, that is, that the utility function is the objective for the agent's optimization problem. What do you do when multiple agents are interacting? In this case we need some definition of equilibrium to describe the interactions, because what's optimal for me depends on what the other agents do. What should be assumed about that? The Nash equilibrium assumption is simple: each agent's choice is the optimal reaction to the other agents' choices. Implicitly, this says that if everyone knows what equilibrium everyone is playing, nobody would want to switch. Put another way, they are doing what they want to do if they assume the same thing about what everyone is doing. So that's the definition of Nash equilibria. And when we say that there are multiple equilibria, we are just saying that in the competitive model, there exists more than one vector of prices that make supply equal to demand, or that under Nash equilibrium, there exists more than one list of choices, one for each agent, that is an equilibrium, in the sense that each agent's choice is the best response to the other agents' choices.

At the time Doug and I were working on the Prize paper, game theory and economics were further apart than they are now, and game theorists at that time were less hostile on average than economists to the existence of multiple equilibria. In particular, it was an essential feature of coordination games. A classic example of a coordination game is that you are meeting someone in New York but you haven't agreed on where to meet (before cell phones, so we cannot communicate in real time). I can choose to go to Grand Central or to the Empire State Building, and you have the same choice. It is a Nash equilibrium for both of us to go to Grand Central, it is another Nash equilibrium for both of us to go to the Empire State Building, and there is also an unstable equilibrium in random strategies. In economics,

there is also a classical result (Sonnenschein (1972)⁴) that there can be many competitive equilibria in the standard model,⁵ but I also think this result was viewed as a defect in the model. Our Prize paper helped economists to understand that having multiple equilibria can be an important conclusion of the analysis, something that is an important part of the economics that can be interesting.

Doug has already talked about our model. You can have an equilibrium in which everybody takes their money out whether or not they need it. We call that equilibrium a bank run. It turns out to be bad for everybody. If you think everybody else is trying to take out their money, then you realize the bank is going to run out of money before you get there if you wait. So, you will also try to take out all your money.

Alternatively, there is an equilibrium in which agents take out their money only if they really need it. It could be that somebody down the street is selling a classic car, which I think is super cool, and I really want my money out now to buy the car, and I'm going to take my money out to do this. Or it could be that I have some kind of health expense. Or, I might have a friend who needs money, or it could be that I feel a sudden need to take a trip and get away from things. Whatever the reason that I want the money, it's a private reason so I can't necessarily document to the bank that I really prefer to get the money now, but nonetheless it's valuable for me to have the option to take out my money if I want it. That is the feature of liquidity that Doug talked about, and why it is valuable to consumers.

Figure 2 has a picture of Doug. Doug showed a much cooler picture of the two of us. I'm jealous, but I'm grateful too because I like seeing that picture. I want to talk a little bit about Doug. He and I have worked together a lot. I can't remember ever having a real fight. We would have little arguments and the arguments were about both of us wanting the research to be as good as possible, the paper to be as good as possible. One of us would

⁴See also Sonnenschein (1973), Mantel (1974), and Debreu (1974)

⁵The papers were posed in terms of excess demand function that are arbitrary up to continuity, homogeneity of degree 1, and Walras' law. If we look at all but one good (the other determined by Walras' law) and all but one price (taking the other as numeraire), we just have continuity. A standard result in analysis says that the set of zeros of a continuous function can be any closed set.



Figure 2: Douglas W. Diamond

say “Let’s assume this,” and the other one say “No, we can’t assume that, or the model will be impossible to solve. Let’s assume this other thing.” The first one would say, “No, we can’t assume that; it would throw away all the economics.” We would go back and forth, but there was never any unpleasantness. We both wanted everything to be as good as possible. Just a superior co-author.

Back to the paper. The idea is that you may want your money out, and you may want to get it out early, depending on the private shock that you can’t document to the bank. This is the formalization of the notion that you like liquidity. When you have something that looks like bank deposits, which allow you to get your money out whenever you want, you have liquidity that you value because you don’t really know, as Doug put it, how long you want your money in, or as I am putting it, an opportunity may come up that makes you value using the money more than you normally do. So, liquidity is valuable for the depositor, who is a customer of the bank. But bank assets are illiquid; they pay a lot more if you leave them in. Let’s say that the bank made a loan to a developer who has a half-completed shopping center. Well, the money is in this half-completed shopping center. You can’t sell a half-completed shopping center, not for very much money. Even if you write the contract with the borrower that gives the bank the option of calling the loan to get money back, the borrower is not going to agree to giving back

very much money in this contingency, because the borrower is not going to have money to return. Doug and I thought bank assets were illiquid due to information asymmetry. If somebody comes to you to sell bank assets, you know they've spent a lot of time researching loans and they know a lot about what they are worth, and they're going to try to sell the ones that are lemons, that are terrible, consistent with George Akerlof's Prize for the lemons model (Akerlof (1970)). It's exactly this kind of adverse selection that makes bank assets illiquid. In our model, we made a simpler assumption. We assumed the illiquidity was in the technology, like in the shopping center. The reason we did that is because the nature of the bank assets' illiquidity is not the most important part of the model. In general, we worked really hard to make the model as simple as possible, and I think most economic researchers think it's simple. Having a simple model is really beneficial because people have additional degrees of freedom that they can still handle. Scholars have added many other features to the model and they can still solve it. It's a little bit like a crystal or a poem. It is also similar to a gedankenexperiment (thought experiment) in physics like the one in Einstein, Podolsky, and Rosen (1935) with entanglement, as described by the Physics Laureates this year. I don't think anyone complained that they did not model explicitly how the detectors worked or how the particles are created in the first place.⁶ Just as they were looking for the simplest setting in which to discuss entanglement, we want to have the simplest possible setting to explain how valuable liquidity provision by banks makes them subject to strategic runs.

Bank assets are illiquid, which means you don't get much back if you liquidate them before maturity, but demand deposits are liquid in response to depositor preferences, and this is a liquidity mismatch. We also assume the bank assets are riskless. That is a feature, not a bug! Some people say they can't believe Diamond and Dybvig (1983) assumes bank assets are riskless given that we know bank assets are risky! It's a feature because everybody knows that the banks can fail if bank assets are risky. That's obvious. What we're saying is that even if the assets are riskless, and there are enough assets in the bank to pay off all the depositors who should be getting money, then there can still be problems. You can still have bank failures. Of course, if the assets are risky, it doesn't take away the possibility of multiple equilibria, in some

⁶Someone could complain that their example was not rich enough to create results violating Bell's (1964) inequality, but to be fair, that is probably only with hindsight.

of which can fail due to a run.

Because of liquidity mismatch, there are multiple equilibria. There is a good equilibrium in which people withdraw only when the money is needed. There is a bad equilibrium (the bank run) in which everyone takes their money out whether they need it or not. There is also another mixed strategy equilibrium in between in which people randomize, but it is unstable and its economics are not really much different than the run equilibrium, so we needn't discuss it further.

We know from Ben's research in his lecture that bank failures in the economy did tremendous damage in the Great Depression, consistent with our model. This raises the natural policy question of what can government do to prevent the bank run equilibrium? One way you can prevent people from running on the bank is that you can have deposit insurance so that people are guaranteed to get their money back whether or not there is a run. Even if everyone runs on the bank and the assets are gone, the government-back deposit insurance fund will pay them off. That's the policy we think is probably most effective. In a setting in which bank assets are risky, you also need for there to be monitoring of the banks so they don't take too much risk. That is outside of our model.

A second possible policy choice we talked about is a discount window at the central bank, where banks can borrow money if they have trouble meeting withdrawals. This seems less reliable than deposit insurance, since, as Doug said, the central bank may suffer from a credibility problem.⁷ If people are unsure whether the central bank will always fund the banks that are in trouble, say because the central bank decides the fund shortage is due to fraud and not just excess withdrawals, then lending by the discount window is not going to be an effective way assuring depositors their claims are safe, and they may still run on the bank.

The third policy we talked about is suspension of convertibility, which is

⁷To be fair, an insurance fund can also be subject to a credibility problem. Given current institutions, lack of credibility of the insurer is more likely to be due to soundness of the insurance fund's credit or its government backing, rather than a decision that the bank is not worthy of funding, since deposit insurance payoffs benefits depositors but not necessarily bank owners.

shutting off withdrawals, a so-called bank holiday. Bank holidays stop the run by preventing anyone from taking money out, which I thought of as completely bad, since it eliminates the benefit of liquidity provision by banks and it may induce a run by depositors wanting to get out before the bank is closed. I found Ben's comments in his lecture interesting, since I didn't know anything good about bank holidays until I heard his comment that the bank holiday was used as a reset after the Great Depression and was actually beneficial. But as a policy for preventing runs, it's not a good policy to say "Don't worry: if there is a run, we are going to keep people from taking out their money by shutting down the bank." Indeed, depositors may try to withdraw their money in anticipation of a bank holiday.

It has been said, and with a lot of nuance, that several markets in the 2008 financial crisis look just like our bad equilibrium, and both Ben and Doug mentioned some other crises that have this feature, too. In the 2008 crisis, there were runs on the repo market (the market for repurchase agreements, a big shadow banking market), money market mutual funds, some traditional banks, and bank commercial paper. Those are some of the markets for which the 2008 crisis looked pretty much just like the bad equilibrium in our model.

Doug's lecture mentioned briefly our policy piece Diamond and Dybvig (1986). I stand by most of what we say in this piece. I think the description of money market funds is accurate, but a little embarrassing. It says that ignoring some extreme times, you can think of money market funds as being very safe because they normally have both liquid assets and liquid liabilities. However, the extraordinary times are the interesting times, and the assets became illiquid during the financial crisis.

Here is a quote from the paper that is a pretty good description (from 1986!) of the 2008 financial crisis:

Proposals to move toward 100% reserve banking would prevent banks from fulfilling their primary function of creating liquidity. Since banks are an important part of the infrastructure in the economy, this is at best a risky move and at worst could reduce stability because new firms that move in to fill the vacuum left by banks may inherit the problem of runs.

In modern terminology, this quote talks about shadow banking. You can think of this as a discussion of the repo market in the 2008 crisis. The quote mentions 100% reserve banking because that was the theme of the conference where it was presented, but the quote applies more generally to policies that restrict banks from performing their function of providing liquidity. It says if you restrict banks too much and keep them from doing their jobs, then some other institutions (like the repo market) can arise that will replace banks and their function. But it's risky because bank liquidity may not be surplus liquidity (and it probably isn't), and because the loss of liquidity might cause instability because the new firms that move in to fill the vacuum left by banks may have runs. We now have a name – shadow banks – for the firms that move in to replace the functions of banks, and we saw the instability of shadow banks in the 2008 financial crisis. In particular, the run on the repo market was a big part of the contraction of the economy.

Another insight from Diamond and Dybvig (1986) is the reminder that when we say deposit insurance removes the incentive for depositors to run, that means full insurance, not capped insurance. In the classic failure of Continental Illinois Bank, many uninsured deposits ran, creating a big disruption and problems for the regulators, but the regulators insured most of the deposits ex post, incurring the cost of full insurance but getting only part of the benefit. As stated in Diamond and Dybvig (1986), "...there is a much stronger case for 100% insurance than for limiting insurance, especially if the alternative is for regulators to retain their discretion to in fact insure most 'uninsured deposits.' " In spite of a few things I would say differently, I still like our policy piece Diamond and Dybvig (1986), and I recommend it to you if you are doing research in this area. I think most of it stands up.

I already discussed multiple equilibria in our Prize paper: a good equilibrium and the bad equilibrium. Doug talked about it and I talked a little about it that basically if you don't need the money out, you don't run in the good equilibrium, but you do run in the bad equilibrium. I'm going to move on now to talk about multiple equilibria in some prior joint papers I wrote with other co-authors.

The first prior paper I want to discuss is Dybvig and Spatt (1983a). Chester (see Figure 3) was my good buddy in graduate school. He was a year ahead of me and showed me the ropes. In this paper, we model adoption of a tech-




Figure 3: Chester S. Spatt

nology, for example, telephones or a standard unit of measurement. Whether you want a phone or not depends on who else has a phone. If everyone else has a phone, you probably want to pay the money to buy a phone because all your friends will have a phone and you can talk to them. If nobody else has a phone, you don't want a phone because you have to pay for it but you don't get any benefits. This is very similar to bank run models, in that you like to do what everybody else is doing. In the interesting case when the technology is useful, there is a bad equilibrium (with nobody adopting) like the run equilibrium and a good equilibrium (with a set of people adopting) like the no-run equilibrium. We also have the policy prescription of a guarantee of a minimum number of adopters which, like deposit insurance in the banking model, is costless in equilibrium and eliminates the bad equilibrium. I should note that late in the review process we learned about Rohlfs (1974), which has a model encompassing ours but not the insurance suggestion, and instead suggests a discount for early adopters that is costly in equilibrium.

Chester and I had another paper, Dybvig and Spatt (1983b), that had multiple equilibria. The paper is about consumer information and product quality. In the story you have a restaurant that is producing food and higher quality is more expensive to produce, and the question is how much customers will be willing to pay, which depends on what quality they expect. One equilibrium is for the restaurant to offer the lowest possible quality all the

time, and charge customers the (low) price they are willing to pay for that. Expect junk and you get junk and so you don't pay very much. But there are also equilibria where larger quality is offered, and how much larger can be offered depends on the quality of the consumer's information. This paper was written before the internet and online reviews. At that time, if you had a restaurant on a highway and it takes a long time for a typical person to return or maybe they never return, then people don't have very good information. That is why you have an alternative way of certifying quality for restaurants on the highway, and that is why there is a larger fraction of chain restaurants on the highway than in town. Quality of chain restaurants is assured by a different mechanism, the quality of the chain rather than the individual restaurant. This is a case where the important question is about multiple equilibria. The paper was never published. We were silly. The paper was actually accepted at the Review of Economic Studies and the editor, Oliver Hart, wanted us to take out a little section that we thought was really important (because it contained the comparative statics with respect to consumer information) and we said no, we are going elsewhere. Instead, we should have said "Can't we leave it in, please? It is important and indeed it is what makes the paper unique." Maybe Oliver would have asked us to take out something else, and that would have been okay. If we had consulted a senior colleague before acting, we would not have made such a silly mistake. That's the way it goes.

Gerry Jaynes and I also had prior papers with multiple equilibria, Dybvig and Jaynes (1979, 1980). Gerry (see Figure 4) was my professor at Penn and coincidentally he moved to Yale at the same time I moved to Yale with my advisor, Steve Ross. These are both labor papers with a search flavor and have almost the same assumptions. In the models, firms have to pay not only the wages that they pay to the workers, but they also have to pay training costs. The training costs are firm-specific, but if anything the impact of the friction would be bigger with general training that can be transferred. Workers have all submitted applications to all the firms that pay more than their reservation wage and what they are paid if they are employed, and if there is an opening, the firm picks randomly from the workers who applied. This is what I mean by "search flavor;" the random matching is simpler than in a typical search model in which both sides wait for a match. The first paper, Dybvig and Jaynes (1979), considers equilibria in which all firms pay the same wage, with the result that there can be unemployment very similar




Figure 4: Gerald D. Jaynes

to what Keynes described. In our model, unemployment is defined (as is customary) as having people who have no job but are willing to work at less than the prevailing wage. In an equilibrium with unemployment, there are many firms offering the same wage that is higher than the competitive wage would be, but other firms cannot enter and make a profit, because when you enter the people in your firm are going to have applications in at the firms that pay more and you are going to incur higher training costs than other firms. This is an equilibrium for any wage in some range above the competitive wage. In this range, the wage savings from offering the workers' reservation wage is less than the increase in training costs. This model anticipated the emphasis on multiple equilibria in macro models by Farmer (2008, 2012, 2013). In our paper, we connect our model to Keynes (1936) using the following quote:

... the postulates of the classical theory are applicable to a special case only and not to the general case, the situation which it assumes being a limiting point of the possible positions of equilibrium.

That is consistent with our model. The competitive solution is always an equilibrium of our model, and it has either full employment or less than

full employment with firms paying the workers' reservation wage, but there are also equilibria with a higher common wage and unemployment. We never published this or our other paper because we moved in very different directions in our careers.⁸ A couple of years after this, Gerry was on a commission studying the status of black Americans, and I moved more into Finance, but I still think these are interesting papers. I want to give a shout out to Gerry.

The other paper with Gerry, Dybvig and Jaynes (1980), has very similar assumptions. The difference is that we do not restrict attention to equilibria in which firms all pay the same wage. If my firm offers a higher wage than yours and that costs me more, but I'll also have a lower quit rate and in equilibrium, the shape of the distribution must make firms just indifferent about which wage to pay. Given the search flavor of the model, I think of this paper as anticipating Burdett and Mortensen (1998). We share their central assumption that workers search while employed as well as when unemployed, but our random matching does not have "structural unemployment" because firms and workers are never both waiting to be matched. Looking at both papers together makes clear that wage dispersion comes from the assumption about "searching while employed" more than the general search frictions, since wage dispersion is there in our random matching model that does not make firms search for a match. However, our simplified form of search does imply an interesting difference in our results. Our model's solution includes equilibria with mass points in the distribution of wages, while their model cannot. The reason is that a firm sitting at a mass point in their model can reduce the structural unemployment and benefit from increasing the wage a little bit above the mass point. That accelerates arrival of agents in their model (since there can now be random matching with the agents at the same wage before increasing the wage slightly). In our model, that is not necessarily useful, because the firm chooses immediately one of the workers with a lower wage or unemployed to fill the position, so there is no structural unemployment and no advantage to having more applicants. Our model can have a discrete set of wages, a continuum of wages, or a hybrid.

⁸Gerry points out, and I had forgotten, that we did submit the 1980 paper to the QJE, but we didn't want to make the changes that were requested. Probably, we should have consulted a senior colleague, who would have counseled us to make the changes.



Figure 5: Karl Shell

I also want to give a shout out to Karl Shell, who was one of my mentors. He and another professor of mine, Dave Cass, had a nice paper (Cass and Shell (1983)), that was cited in our Prize paper because it answers a possible puzzle in the model. If agents expect the run equilibrium, why would they deposit money in the first place? If you know there is going to be a run equilibrium, you would be better off not depositing. Well, the answer is in the sunspots model. Sunspots are supposed to be an exogenous shock, observable by everybody, that is not directly economically relevant. In reality, we know sunspots can disrupt communications, and in extreme cases, could bring down the power grid. But that's not the idea. The idea is supposed to be something that's commonly observable, but gives us a source of uncertainty on which to coordinate, but does not need to impact payoffs except through the strategic coordination.⁹ The sunspots model says that if there are multiple equilibria and some randomness everyone observed, it is also an equilibrium to select one of the equilibria based on the commonly observed randomness. It could be that there is a very small fraction of the time when the number and size of sunspots are really big (or really small or whatever) and cause the bank run equilibrium, in which case it may still pay

⁹Like having riskless assets in the Prize paper, this is an “even if” proposition: sunspots may matter even if they are irrelevant for payoffs except for the equilibrium selection. Of course, commonly observed randomness that affects payoffs a little or a lot can also play the same role.

to put your money in the bank. I also wanted to mention that Karl and his students have written lots of interesting papers extending the Prize paper.¹⁰ That is part of the reason why so many people are so familiar with our work. Along similar lines, I should also mention Neil Wallace. He wrote some early papers extending the Prize paper, and his students have also contributed to the literature.¹¹

I want to mention a couple of ideas for further work. Manju Puri and co-authors (Iyer and Puri (2012), Iyer, Puri, and Ryan (2016), and Martin, Puri, and Ufier (forthcoming)) have studied actual bank runs. Based on casual observation, the simple equilibria in our model or the dynamic models I have seen don't look so much like what they find in actual bank runs. I don't think that invalidates our basic messages, since liquidity mismatch, instability, and the usefulness of deposit insurance will remain. However, we would probably learn something from trying to model practice more closely. One idea is to use another equilibrium concept, the concept of rationalizability of David Pearce (Pearce (1984)). He was a student when I was teaching at Princeton. Doug Bernheim also has a similar solution concept (Bernheim (1984)). One of the things about Nash equilibrium that is a little unsatisfying is it is unclear how people can know which Nash equilibrium everybody else is thinking about. Rationalizability says that all I can impose is that everybody in the economy has some reasonable idea that is consistent with the reasonable ideas other people could have. That will give you a larger set of equilibria and maybe could help us to explain the empirical results of Manju and her co-authors.

Figure 6 shows some other mentors (besides Steve Ross, Gerry Jaynes, and Karl Shell). Michael McGill was my undergraduate professor whose graduate economics course I took without any prerequisites. Since it was half Hamiltonians and related concepts from physics, and I was a physics major, it worked out. He also helped me to figure out where to apply for graduate school. These others (Bill Brainard, Don Brown, Sharon Oster, Bob Polak, and Al Klevorick) were all mentors from graduate school. All were very helpful in my career. Of these, Bill Brainard was most influential, and my breadth in research follows Bill (and Steve Ross).

¹⁰For example, see Peck and Shell (2003,2010), Antinolfi and Prasad (2008), Ennis and Keister (2016), Gu (2011), and Keister (2016).

¹¹See, for example, Wallace (1988, 1990, 1996), Alonso (1996), Andolfatto, Nosal, and Sultanum (2017), Cavalcanti and Monteiro (2016), and Sultanum (2014).



William Brainard



Donald Brown




Alvin Klevorick



Michael J. P. Magill



Sharon Oster



Robert Pollak

Figure 6: Some other mentors

It has been a lot of fun thinking about the Prize paper and some of its roots, with a focus on multiple equilibria in some of my earlier papers. I would like to thank my colleagues, co-authors, and mentors. It is really great that I was able to interact with you.

References

Akerlof, George A., 1970, The Market for ‘Lemons’: Quality Uncertainty and the Market Mechanism, *Quarterly Journal of Economics* 84, 488-500.

Alonso, Irasema, 1996, On avoiding bank runs, *Journal of Monetary Economics* 37, 73-87.

Andolfatto, David, Ed Nosal, and Bruno Sultanum, 2017, Preventing bank runs, *Theoretical Economics* 12, 1003-28.

Antinolfi, Gaetano, and Suraj Prasad, 2008, Commitment, banks, and markets, *Journal of Monetary Economics* 55, 265-77.

Bell, John S., 1964, On the Einstein Podolsky Rosen Paradox, *Physics* 1, 195-200.

Bernanke, Ben S., 1983, Nonmonetary effects of the financial crisis in the propagation of the Great Depression, *American Economic Review* 73, 257-276.

Bernanke, Ben S., 2022, Economic Sciences Prize Lecture (video), available online at <https://www.nobelprize.org/prizes/economic-sciences/2022/bernanke/>.

Bernanke, Ben S., 2023, Economic Sciences Prize Lecture (written version), Brookings Institution working paper.

Bernheim, B. Douglas, 1984, Rationalizable Strategic Behavior, *Econometrica* 4, 1007-28.

Burdett, Kenneth, and Dale T. Mortensen, 1998, Wage Differentials, Employer Size, and Unemployment, *International Economic Review* 39, 257-273.

Cass, David, and Shell, Karl, 1983, Do sunspots matter? *Journal of Political Economy* 91, 193–227.

Cavalcanti, Ricardo do O., and Paulo K. Monteiro, 2016, Enriching information to prevent bank runs, *Economic Theory* 62, 477-94.

Chen, Nai-fu, Richard Roll, and Stephen A. Ross, 1986, Economic Forces and the Stock Market, *Journal of Business* 59, 383-403.

Cox, John C., Jonathan E. Ingersoll, Jr., and Stephen A. Ross, 1981, A Re-Examination of Traditional Hypotheses about the Term Structure of Interest Rates, *Journal of Finance* 36, 769-799.

Cox, John C., Jonathan E. Ingersoll, Jr., and Stephen A. Ross, 1985a, An Intertemporal General Equilibrium Model of Asset Prices, *Econometrica* 53, 363-384.

Cox, John C., Jonathan E. Ingersoll, Jr., and Stephen A. Ross, 1985b, A Theory of the Term Structure of Interest Rates, *Econometrica* 53, 385-407.

Cox, John C., and Stephen A. Ross, 1976, The Valuation of Options for Alternative Stochastic Processes, *Journal of Financial Economics* 3, 145-166.

Cox, John C., Stephen A. Ross, and Mark Rubinstein, 1979, Option Pricing: A Simplified Approach, *Journal of Financial Economics* 7, 229-263.

Debreu, Gérard, 1974, Excess-demand functions, *Journal of Mathematical Economics* 1, 15–21.

Diamond, Douglas W., 1984, Financial intermediation and delegated monitoring, *Review of Economic Studies* 51, 393–414.

Diamond, Douglas W., 2022, Economic Sciences Prize Lecture (video), available online at <https://www.nobelprize.org/prizes/economic-sciences/2022/diamond/>.

Diamond, Douglas W., and Philip H. Dybvig, 1983, Bank Runs, Deposit Insurance, and Liquidity, *Journal of Political Economy* 91, 401-19. Reprinted in *Federal Reserve Bank of Minneapolis Quarterly Review* Winter 2000, 14-

23, <http://minneapolisfed.org/research/qr/qr2412.pdf>.

Diamond, Douglas W., and Philip H. Dybvig, 1986, Banking Theory, Deposit Insurance, and Bank Regulation, *Journal of Business* 59, 55-68.

Dybvig, Philip H., 2022, Nobel Lecture: Multiple Equilibria (video), available online at <https://www.nobelprize.org/prizes/economic-sciences/2022/dybvig/>

Dybvig, Philip H., 2023. "Nobel Lecture: Multiple Equilibria." *Journal of Political Economy*, forthcoming (version for academic economists).

Dybvig, Philip H., Jonathan E. Ingersoll, Jr., and Stephen A. Ross, 1996, Long Forward and Zero-Coupon Rates Can Never Fall, *Journal of Business* 69, 1-25.

Dybvig, Philip H., and Gerald Jaynes, 1979, Microfoundations of Wage Rigidity and Unemployment, Yale University working paper.

Dybvig, Philip H., and Gerald Jaynes, 1980, Output Supply, Employment, and Intra-Industry Wage Dispersion, Cowles Foundation Discussion Paper 546, Yale University.

Dybvig, Philip H., and Stephen A. Ross, 1987, Arbitrage, a contribution to *The New Palgrave: a Dictionary of Economics* 1, New York: Stockton Press, 100-106.

Dybvig, Philip H., and Chester S. Spatt, 1983a, Adoption Externalities as Public Goods, *Journal of Public Economics* 20, 231-47.

Dybvig, Philip H., and Chester S. Spatt, 1983b, Does it Pay to Maintain a Reputation? *Consumer Information and Product Quality*, Yale working paper.

Einstein, A., B. Podolsky, and N. Rosen, 1935, *Physical Review* 47, vol. 10, 777-80.

Ennis, Huberto M., and Todd Keister, 2009, Bank Runs and Institutions: The Perils of Intervention, *American Economic Review* 99, 1588-607.

Farmer, Roger, 2008, Aggregate Demand and Supply, *International Journal of Economic Theory* 4, 77-93.

Farmer, Roger, 2012, Confidence, Crashes and Animal Spirits, *Economic Journal* 122, 155-172.

Farmer, Roger, 2013, Animal Spirits, Financial Crises and Persistent Unemployment, *Economic Journal*, 123, 317-340.

Gu, Chao, 2011, Herding and bank runs, *Journal of Economic Theory* 146, 163-88.

Iyer, Rajkamal, and Manju Puri, 2012, Understanding Bank Runs: The Importance of Depositor-Bank Relationships and Networks, *American Economic Review* 2012, 1414-45.

Iyer, Rakjamal, Manju Puri, and Nick Ryan, 2016, A Tale of Two Runs: Depositor Response to Bank Solvency Risk, *Journal of Finance*, 2687-726

Keister, Todd, 2016, Bailouts and financial fragility, *Review of Economic Studies* 83, 704-36.

Keynes, John Maynard, 1936, *The General Theory of Employment, Interest, and Money*, London: Macmillan.

Mantel, Rolf (1974). On the characterization of aggregate excess-demand, *Journal of Economic Theory*. 7, 348–353.

Martin, Christopher, Manju Puri, and Alexander Ufier, forthcoming, Deposit Inflows and Outflows in Failing Banks: The Role of Deposit Insurance, *Journal of Finance*

Pearce, David, 1984, Rationalizable Strategic Behavior and the Problem of Perfection, *Econometrica* 4, 1029-50.

Peck, James, and Karl Shell, 2003, Equilibrium Bank Runs, *Journal of Political Economy* 111, 103-23.

Peck, James, and Karl Shell, 2010, Could making banks hold only liquid assets induce bank runs? *Journal of Monetary Economics* 57, 420-27.

Rohlfis, Jeffrey, 1974, A theory of interdependent demand for a communications service, *Bell Journal of Economics* Spring 1974, 16-37.

Ross, Stephen A., 1973, The Economic Theory of Agency: The Principal's Problem, *American Economic Review* 63, 134-139.

Ross, Stephen A., 1976a, The arbitrage theory of capital asset pricing, *Journal of Economic Theory* 13, 341-60.

Ross, Stephen A., 1976b, Return, risk, and arbitrage, in Irwin Friend and James Bicksler, *Risk and Return in Finance*, Cambridge, Ma.: Ballinger.

Ross, Stephen A., 1978, A Simple Approach to the Valuation of Risky Streams, *Journal of Business* 51, 453-75.

Sonnenschein, Hugo, 1972, Market Excess Demand Functions, *Econometrica* 40, 549-563.

Sonnenschein, Hugo, 1973, Do Walras' identity and continuity characterize the class of community excess-demand functions? *Journal of Economic Theory*. 6 (4): 345-354.

Sultanum, Bruno, 2014, Optimal Diamond-Dybvig mechanism in large economies with aggregate uncertainty, *Journal of Economic Dynamics and Control* 40, 950192.

Wallace, Neil, 1988, Another attempt to explain an illiquid banking system: The Diamond-Dybvig model with sequential service taken seriously, *Federal Reserve Bank of Minneapolis Quarterly Review* 12, 3-16.

Wallace, Neil, 1990, A banking model in which partial suspension is best, *Federal Reserve Bank of Minneapolis Quarterly Review* 14, 11-23.

Wallace, Neil, 1996, Narrow banking meets the Diamond-Dybvig model, *Federal Reserve Bank of Minneapolis Quarterly Review* 20, 3-13.