

Nonparametric Estimation of the Diamond-Dybvig Banking Model

Bruno Sultanum

The Diamond and Dybvig (1983) model has been extensively used to explain episodes of runs against financial institutions. In the model, depositors face uncertainty about whether they would prefer to consume in an early or late period. Because there are costs associated with an early liquidation of investments, depositors can benefit from an insurance contract with respect to their preference shock. The optimal insurance will transfer resources from those depositors who prefer to consume in the late period, and therefore get a better return in their investments, to those who prefer to consume in the early period. Such transfers, however, cannot be contingent on the depositor preference because these are not observed—the contract must be incentive compatible so they reveal their true preferences in equilibrium. As Diamond and Dybvig (1983) argue, an arrangement that has this property is a bank contract. The bank promises the efficient transfer in the early period to any depositor who claims the resources. In one equilibrium, only those who actually have preference for early consumption claim early payments. However, there is also an equilibrium where depositors fear that every other depositor, including those with preference for later consumption, will claim early payments. As a result, depositors fear no resources will be left at the bank for consumption late and all of them have incentives to claim payments in the early period, generating a self-fulfilling bank run.

■ I thank Huberto Ennis, Todd Keister, Neil Wallace, Thomas A. Lubik, Zhu Wang and Quang Vuong for helpful comments. The views expressed in this paper are those of the author and should not necessarily be interpreted as those of the Federal Reserve Bank of Richmond or the Federal Reserve System.

Even though this argument is intuitive, whether or not the model generates runs under a well-designed bank contract depends on the particular specifications of the environment. For example, if there is no aggregate uncertainty, in the sense that the number of early consumers in the economy is known, a simple suspension scheme is able to prevent runs from happening. This observation was made by Diamond and Dybvig (1983), but they highlight that this is only true without aggregate uncertainty.¹ Later, Wallace (1988) pointed out the importance of *sequential service* in order to generate bank runs in the Diamond-Dybvig model. Sequential service is a constraint that payments must be done in sequence, as depositors arrive at the bank, so payments to one depositor cannot be made contingent on future withdrawal demand. However, as Green and Lin (2000, 2003) later point out, aggregate uncertainty and sequential service alone are not enough to generate bank runs in the Diamond-Dybvig model. As they pose it, the theory as it stood was incomplete. A huge theoretical literature has followed trying to understand what other ingredients are necessary for the existence of bank runs.

The approach to address this question and understand whether the Diamond-Dybvig model actually generates bank runs or not has been to build examples where bank runs do exist—see, for instance, Peck and Shell (2003), Ennis and Keister (2009b), and Sultanum (2014). The examples of bank runs are built with particular distributions of liquidity needs and other primitives, such as preferences. However, for the model to explain observed runs, we need its primitives to be consistent with empirical observations. In particular, the distribution of liquidity needs in the model should be consistent with the empirical one. Without this consistency between data and model, the model would still be “incomplete” as a theory to explain historical run episodes. Hence, developing tools to estimate the Diamond-Dybvig model is an important step in understanding how bank runs actually work and how to prevent them.

Moreover, the advantages of estimating the Diamond-Dybvig model go beyond the positive aspect of explaining empirical observations. There are also normative advantages. One way policymakers can use this tool is to estimate the model for different markets and institutions (possibly also making it state contingent) and then test for which markets and institutions a run equilibrium exists. That is, the model can guide policymakers to when and where runs are a possibility, allowing timely policy measures to be taken before a run ever happens.

¹ As pointed out later by Ennis and Keister (2009a), this result also relies heavily on the bank’s ability to commit to a suspension scheme.

In this paper, I construct a structural estimator for the distribution of liquidity needs in the version of the Diamond-Dybvig model studied in Sultanum (2014). The data requirement for the estimator is that the econometrician observes the total amount withdrawn. This assumption serves two purposes. First, it is a very weak data requirement, which is always welcome since more detailed data may only be available to regulators (sometimes not even to them). Second, in the model, a depositor either withdraws his entire deposit or nothing. In practice, people can withdraw only part of their money or can withdraw money from multiple accounts at the same bank. So the “amount withdrawn” is a clear and well-defined measure both in theory and in the model. It is not clear how to match the observation of partial withdrawals or withdrawals from different accounts in the data to the model.

What makes this problem difficult is that aggregate payments are observed, but the preferences that define the liquidity needs are not observed. Therefore, in order to estimate the distribution of liquidity needs, one must establish a map between payments and preferences. However, because payments embed an insurance against the preference shock risk, how much is paid for any given realization of preference shocks is endogenous. In particular, it depends on the distribution of liquidity needs.

There is a large literature that studies whether past run episodes against financial institutions were due to coordination failure, as in Diamond and Dybvig (1983), or not. This literature focuses on indirect tests of theoretical frameworks. That is, it tests some of the implications of the theory rather than estimates a particular model and tests whether it generates runs or not. Two recent examples are Foley-Fisher et al. (2015) and Schmidt et al. (2016). Foley-Fisher et al. (2015) develop a model to study runs against extendible funding agreement-backed notes (XFABN) issued by life insurers. Their theory suggests an instrument for the strategic complementarity among investors that is (plausibly) exogenous to variations in fundamentals. So, to test whether self-fulfilling runs played a role in the withdrawals from XFABN or not, the authors test the correlation between the instrument implied by the theory and the observed withdrawals. Schmidt et al. (2016) study runs against money market mutual funds. The authors develop a model and test its different predictions. For example, they test whether outflows from sophisticated investors in reaction to worse fundamentals are greater than from unsophisticated ones, where sophistication is defined by the quality of the information the investor has access to.

I see the nonstructural approach that has been used as complementary to a fully structural one. It sheds light on whether self-fulfilling

bank runs exist or not. In fact, in my estimation procedure I assume that the econometrician knows whether past runs were due to coordination failure or not. However, without the estimation of a structural model, it is hard to test particular theories—such as Diamond-Dybvig. Once the theory is tested, then we can use it to make predictions and/or policy recommendations.

The econometric method I use in this paper builds on those developed for estimation of auctions. Specifically, it builds on Guerre et al.'s (2000) idea of using the equilibrium conditions of the model to map observable to unobservable variables. Since its publication, Guerre et al. (2000) has spurred a huge empirical and theoretical literature. Some of the more recent theoretical examples are Campo et al. (2011), which allows bidders to be risk-averse; Krasnokutskaya (2011), which considers bidders' unobserved heterogeneity; and Kastl (2011), which proposes an estimation method for auctions with discrete bids. On the empirical side we have, for example, Cassola et al. (2013), which uses the extension in Kastl (2011) to study liquidity demand from European banks during the 2007 financial crisis; and Hortaçsu and Kastl (2012), which quantifies the dealers' advantage from observing customers' orders using data on Canadian Treasury auctions.

Even though a lot can be done using, and improving on, the estimation procedure I discuss in this paper, as the literature on the estimation of auctions has shown, the goal of this paper is not to fully investigate all the properties and possible extensions of a particular estimator. The goal here is to provide an illustrative framework that future researchers can build on in order to estimate bank-run models in different settings. I believe that a full investigation is only worth it with a particular dataset and institutional framework in mind. For this reason, a lot of the discussion here is abstract and details on data and applications are left for future research.

The paper is organized as follows. Section 1 describes the model, the equilibrium concept, and provides a characterization of the solution. Section 2 describes the data requirement, discusses identification, and provides a nonparametric estimator and a numerical example of the procedure. Section 3 discusses how the model can be used to test for the existence of bank-run equilibria. Section 4 discusses practical difficulties and challenges associated with estimating the model. Section 5 concludes.

1. THE MODEL

The model builds on Sultanum (2014), which is an extension of Peck and Shell (2003) with a continuum of agents. The advantage of this

setting is that the optimal bank contract can be easily characterized by a second-order differential equation, which will be used in the proposed estimation procedure.

Environment


There are three periods, zero, one, and two, and a unit measure of agents called depositors. In period zero, each depositor is endowed with one unit of wealth, which they can invest to consume in periods one and two (agents do not consume in period zero). The investment technology is as follows. Investments in period zero pay gross return 1, if liquidated in period one, and gross return $R > 1$, if liquidated in period two. Depositors are identical in period zero. In period one, each depositor receives a preference shock. The preference shock turns them into one of two types: *patient* or *impatient*. The utility of a type impatient depositor is $u(c_1)$, while that of a type patient depositor is $u(c_1 + c_2)$, where (c_1, c_2) is consumption in periods one and two, respectively. The utility function u is the constant relative risk-aversion (CRRA) utility function, $u(c) = \frac{c^{1-\gamma}-1}{1-\gamma}$, and I assume that the risk-aversion parameter γ is greater than 2.

Let $\alpha \in [0, 1]$ denote the fraction of depositors of type impatient. The value of α is assumed to be a random variable with cumulative distribution function F and density function f , which satisfies $f(\alpha) > 0$ for all α in the support $[0, 1]$. The density f is also assumed to be continuous and differentiable in the support. Conditional on the realization of α , the event that a depositor is of type impatient is i.i.d. across agents and has Bernoulli distribution with parameter α . Throughout the text, I refer to α as the aggregate liquidity need in the economy.

One can think that the event where more than, say, 90 percent of depositors are type impatient has probability zero. This could be formalized by allowing the support for the aggregate liquidity need to differ from $[0, 1]$. That is, in general, F could have support $[\alpha_l, \alpha_h]$ with $0 \leq \alpha_l < \alpha_h \leq 1$. In this case, α_l and α_h would also have to be estimated. The approach I describe in this paper can be extended to address this case, which I believe to be of interest. However, as previously stated, going through all the details and extensions of the estimation procedure is beyond the goal of this paper.

Sequence of actions and bank contracts

Depositors face risk in the form of preference shocks (be patient or impatient). As a result, an insurance arrangement is desirable in order to improve depositors' *ex-ante* welfare. Following Peck and Shell

Figure 1 Sequence of Actions

(2003) and, more closely, Sultanum (2014), I focus on a form bank contract where resources are deposited in a bank and depositors can withdraw resources if they want to. That is, in period zero, all resources are deposited in the bank. In the beginning of period one, each depositor observes his own type (which is private information). No one observes the realization of α . Then, agents simultaneously decide whether to withdraw resources from the bank or not. The bank serves the withdrawal requests of the individuals in a random sequence, which the literature refers to as the sequential service constraint (see Wallace [1988] for details on sequential service). A depositor's position in the queue is uniformly distributed among depositors who decide to withdraw resources from the bank in period one.² After all withdrawal payments are made, what is left in the bank pays a gross return of R from period one to period two. In period two, the bank distributes the amount left to those who did not withdraw in period one. Figure 1, extracted from Sultanum (2014), depicts the sequence of actions.

A bank contract tells how much a depositor who withdraws in period one receives as a function of his queue position and how much a depositor who waits until period two to withdraw receives as a function of the number of withdrawals in period one. We can formalize it as a pair of continuous functions, $m = (c_1, c_2)$, where $c_1 : [0, 1] \rightarrow \mathbb{R}_+$, and $c_2 : [0, 1] \rightarrow \mathbb{R}_+$. The function $c_1(z)$ gives the payment to a depositor who withdraws in period one and has position z in the queue. The function $c_2(\bar{z})$ gives the payment to a depositor who waits until period two to withdraw when the fraction of people who withdrew in period one is $\bar{z} \in [0, 1]$. The continuity on m is without loss of generality with

² Although depositors arrive at the bank in sequence, the rate of arrivals of depositors at the bank cannot be measured by the bank and, therefore, cannot be used as a factor to determine payments. This assumption can be rationalized by assuming that the time interval in which agents arrive varies proportionally to the number of agents visiting the bank.

respect to finding the constrained optimal outcome, which we define later.

Feasibility of a contract requires that payments must not be greater than the resources available. I impose that the total amount paid must exactly equal the resources available. This requirement is without loss of generality because utility functions are strictly increasing. The feasibility conditions can be written in terms of the functions c_1 and c_2 as

$$c_2(\bar{z}) = \frac{1 - \int_0^{\bar{z}} c_1(z) dz}{1 - \bar{z}} R \quad \text{for all } \bar{z} \in (0, 1), \quad \text{and} \quad \int_0^1 c_1(z) dz = 1. \quad (2.1)$$

A bank contract m and the sequence of actions induce a Bayesian game where each player has only two types, either *patient* or *impatient*, and two actions, either withdraw in period one or period two. A strategy profile is a function s that maps types $\theta \in \{\textit{patient}, \textit{impatient}\}$ into probability measures over the periods of withdrawal, $\{\textit{period one}, \textit{period two}\}$. I consider only symmetric Bayesian Nash equilibria of this game, where symmetric means that players of the same type use the same strategy.

It is important to note that the game is simultaneous. That is, when a depositor is deciding whether to withdraw or not, he does not observe the withdrawal decisions of other depositors. One could think that, in practice, people have at least some idea (or signal) of other depositors' actions. For instance, one could see whether or not there is a line in front of the bank, as beautifully illustrated in the Frank Capra movie *It's a Wonderful Life*. Of course, whether this signal is available or not depends on the setting. These days, when many withdrawal decisions are done online or by phone, such as in mutual funds and other shadow banks, it seems reasonable to assume that depositors do not have much information on other depositors' actions prior to their withdrawal decision. For simplicity, I do not allow depositors to observe any other depositors' actions or obtain any signal that is informative of such actions.

The optimal bank contract

The bank problem is to design a contract $m = (c_1, c_2)$ that maximizes ex-ante welfare of depositors. This assumption can be justified by an extension of the model where a competitive bank sector has banks competing to attract depositors from other banks. To keep the exposition simple, however, I follow the literature and directly assume that the goal of the bank is to maximize depositors' welfare.

The outcome that maximizes ex-ante welfare of depositors must be such that only impatient depositors consume in period one, while all the patient depositors consume in period two. This is the case because the return from period one to period two, R , is strictly greater than 1. Therefore, I am interested in bank contracts that have an equilibrium in which only impatient depositors withdraw in period one. I call such equilibrium a no-run equilibrium, and, when a bank contract has a no-run equilibrium, I call it an incentive-compatible bank contract.

When a depositor observes his type, he uses Bayes' rule to update his belief over the distribution of α . Let $f_p(\alpha) = (1 - \alpha)f(\alpha)/\int_0^1(1 - z)f(z)dz$ be the density of α conditional on the depositor being of type patient. Note that impatient depositors withdraw in period one because they derive no utility from period-two consumption. Therefore, in order to verify that a contract is incentive-compatible, we just need to verify that a patient depositor is better off withdrawing in period two when the other patient depositors are withdrawing in period two.

A feasible bank contract $m = (c_1, c_2)$ is incentive compatible if, and only if, it satisfies

$$\int_0^1 \int_0^\alpha \frac{u(c_1(z))}{\alpha} dz f_p(\alpha) d\alpha \leq \int_0^1 u(c_2(\alpha)) f_p(\alpha) d\alpha. \quad (2.2)$$

The left-hand side of the above inequality is the expected utility of a patient depositor if he withdraws in period one, and the right-hand side of the inequality is his expected utility if he withdraws in period two—all conditional on the other depositors withdrawing in period one only if they are impatient types.

When depositors are playing the no-run equilibrium, the *ex-ante* welfare associated with a bank contract $m = (c_1, c_2)$ is

$$W(m) = \int_0^1 \left[\int_0^\alpha u(c_1(z)) dz + (1 - \alpha)u(c_2(\alpha)) \right] f(\alpha) d\alpha. \quad (2.3)$$

A bank contract is said to be optimal if it achieves the maximum of $W(m)$ among all feasible and incentive compatible bank contracts $m = (c_1, c_2)$.

Let us assume for a moment that the incentive-compatibility constraint does not bind in this problem. Then, using the same approach as in Sultanum (2014), we can show that an optimal bank contract $m = (c_1, c_2)$ always exists and $w(\alpha) = \int_0^\alpha c_1(z) dz$ is the unique solution to the second-order differential equation

$$w''(\alpha)u''(w'(\alpha)) = h(\alpha) \left[u'(w'(\alpha)) - Ru' \left(\frac{1 - w(\alpha)}{1 - \alpha} R \right) \right] \quad (2.4)$$

with boundary conditions $w(0) = 0$ and $w(1) = 1$, where $h(\alpha) = \frac{f(\alpha)}{1-F(\alpha)}$.³ Therefore, to solve the model it suffices to first solve the differential equation (2.4), then recover c_1 and c_2 using that $c_1(\alpha) = w'(\alpha)$ and $c_2(\alpha) = \frac{1-w(\alpha)}{1-\alpha}R$.

Equation (2.4) differs from the one in Sultanum (2014) because the Lagrange multiplier of the incentive-compatibility constraint shows up in their characterization, but it does not show up here. In the present setting, because the utility of the patient types is the same as the impatient, the incentive-compatibility constraint of agents does not bind. There are two steps to show this result. The first one is to note that, in any solution of equation (2.4), we must have $c_1(\alpha) = w'(\alpha) \leq c_2(\alpha) = \frac{1-w(\alpha)}{1-\alpha}R$ for all α . Otherwise, the boundary condition would not be satisfied. The second step is to show that this inequality in consumption implies that the period-two distribution of consumption stochastically dominates the period-one distribution of consumption when other patient types withdraw only in period two. That is, patient depositors are better off choosing period-two consumption when they believe that other patient depositors are also waiting to consume in period two. Therefore, we can conclude that the bank contract is incentive compatible.

2. ESTIMATION

The primitives of this economy are given by the risk-aversion parameter, γ , the return, R , and the distribution of the liquidity needs, F . In this section, we establish an estimator for the distribution of the liquidity needs under the assumption that we know γ and R or that they can be identified separately.

The assumption that the return, R , is known seems natural since one can observe market returns from bank balance sheets. The

³ It is easier to solve the differential equation (2.4) in terms of a system of differential equations, where the marginal utility of period-one consumption, $m_1(\alpha) = u'(w'(\alpha))$, and period-two consumption, $c_2(\alpha) = \frac{1-w(\alpha)}{1-\alpha}R$, are the main variables. That is,

$$\begin{aligned} m_1'(\alpha) &= h(\alpha)[m_1'(\alpha) - Rc_2(\alpha)^{-\gamma}] \\ c_2'(\alpha) &= \frac{1}{1-\alpha}[c_2(\alpha) - Rm_1(\alpha)^{-1/\gamma}] \end{aligned}$$

with boundary conditions $c_2(0) = R$ and $c_2(1) = R/m_1(1)^{1/\gamma}$. By picking the initial $c_1(0)$, one can target the final condition $c_2(1) = R/m_1(1)^{1/\gamma}$. That is, combining the fact that the solution is continuous in the initial condition and that the solutions cannot cross, one can use the intermediate value theorem to argue that an initial $c_1(0)$ such that the boundary condition is satisfied must exist. Moreover, one can show that in such a solution $c_1(1) = c_2(1) = 0$.

assumption that risk-aversion is known, however, deserves justification. This assumption is made for tractability since identifying risk-aversion and distributions together is challenging in this, and also other, settings. For example, Campo et al. (2011) establish that the risk-aversion parameter cannot be identified in first-price auctions together with the distribution of valuations. One could impose additional parametric assumptions in order to identify both the risk-aversion parameter and the distribution of liquidity needs. I consider such analysis interesting but leave it for future research.

There are also two structural assumptions that are necessary for our estimation procedure. Namely, that the bank contract is optimal, as described in the previous section, and that depositors play the no-run equilibrium, where only impatient types withdraw in period one. Alternatively, we could have assumed that we can separately identify the periods in which the no-run equilibrium is played. This is equivalent to saying that, at least ex post, we know whether a bank run happened or not.

In terms of observed data, we assume we have N independent instances of our economy, and in each one we observe only how many total early payments were made as a fraction of the total resources. That is, we can observe a sequence of realizations $\{w_n\}_n$ that are independent of each other. The sample can be interpreted either as a sample over time of the same bank or a sample with N identical banks. In either case, it is important that $w_n = w(\alpha_n)$, where $\{\alpha_n\}_n$ are independent and identically distributed according to F , and $w(\alpha)$ solves (2.4). In the next subsection I show that these data contain enough information to identify F .

I would like to emphasize that this is a very weak data requirement. Only total outflows from the financial institution being studied are necessary. One could try to improve upon the estimation procedure I discuss here by having additional data available, for example, by having microdata on individual depositors. Additional data would also allow for extensions of the model where more primitives of the economy could be identified. But, as I show, just data on outflows already provide a lot of information, allowing us to identify the distribution of liquidity needs.

Identification

A crucial problem in structural estimation is whether the observed data are enough to identify the primitives of the model. In the context of our model, the assumption is that we observe total withdrawals. Let the distribution of total withdrawals be denoted by G . So the question

is whether we can identify the distribution of liquidity needs, F , from the distribution of total withdrawals, G .

In order to answer this question, we use the solution condition of the model to relate G and F . If the map between these two distributions is unique, then the model is identified. So let us look at these conditions. First, the differential equation in (2.4) implies that w is strictly increasing and, therefore, the inverse of w exists. Moreover, because w takes value in the $[0, 1]$ interval, G has also support $[0, 1]$ and we have that $G(\tilde{w}) = \mathbb{P}[w(\alpha) \leq \tilde{w}] = F(w^{-1}(\tilde{w}))$. Because w and F are differentiable, we also know that G is differentiable (since it is the composition of differentiable functions) and it satisfies $g(w(\alpha))w'(\alpha) = f(\alpha)$. We can now use these conditions to rewrite the differential equation (2.4) in terms of G . We get that


$$w''(\alpha)u''(w'(\alpha)) = h_G(w(\alpha))w'(\alpha) \left[u'(w'(\alpha)) - Ru' \left(\frac{1-w(\alpha)}{1-\alpha} R \right) \right] \quad (3.1)$$

with boundary conditions $w(0) = 0$ and $w(1) = 1$, where $h_G(w) = \frac{g(w)}{1-G(w)}$.

Since we know u and R , for each G we can solve the differential equation (3.1) for w . Once we have w , we can recover F using $F(\alpha) = G(w(\alpha))$. Note that this procedure identifies F . To see this, assume that two distributions, F_1 and F_2 , generate the same G . That is, $F_1(\alpha) = G(w_1(\alpha))$ and $F_2(\alpha) = G(w_2(\alpha))$, where w_1 and w_2 are solutions to the differential equation (2.4) associated with F_1 and F_2 , respectively. If that is the case, then w_1 and w_2 would both have to solve (3.1). But one can show that equation (3.1) admits only one solution, which implies that $w_1 = w_2$ and $F_1 = F_2$. Therefore, we can conclude that the distribution of total withdrawals, G , combined with the first-order condition that characterizes the optimal bank contract, contains enough information to identify the distribution of aggregate liquidity needs F .

Estimation steps and numerical example

I propose an indirect nonparametric estimation of F . This estimation has three steps. First, we estimate the distribution of total withdrawals G . Call this estimator \hat{G} . Then we solve the differential equation (3.1) where G is replaced with \hat{G} . Call the solution to this differential equation \hat{w} . Finally, the estimator of the cumulative distribution of liquidity needs is $\hat{F}(\alpha) = \hat{G}(\hat{w}(\alpha))$, and its density estimator is $\hat{f}(\alpha) = \hat{g}(\hat{w}(\alpha))\hat{w}'(\alpha)$.

Figure 2 Distributions

The problem of estimating G is a standard nonparametric estimation problem for a continuous distribution over a compact support. One must choose a bandwidth $h > 0$ and a Kernel function $k : \mathbb{R} \rightarrow \mathbb{R}_+$, where $\int k(u)du = 1$. Then we have that


$$\hat{g}(w) = \frac{1}{Nh} \sum_n k\left(\frac{w - w_n}{h}\right) \quad \text{and} \quad \hat{G}(w) = \int_0^w \hat{g}(\tilde{w})d\tilde{w}. \quad (3.2)$$

There are different ways of choosing the bandwidth h and the Kernel k . I refer the interested reader to Pagan and Ullah (1999) for a full discussion.

To illustrate how this estimator works in practice, below I simulate the model and use our procedure to estimate the underlying distribution F . I consider a specification of the model with $\gamma = 3.0$, $R = 1.15$, and F is a normal distribution with mean $\mu = 0.5$, standard deviation $\sigma = 0.25$, and truncated between 0 and 1.

The numerical exercise is performed with the following steps. I first solve the model for w using equation (2.4). Then I draw 500 observations α_n from F and generate the sample $\{w_n\}_n$ using $w_n = w(\alpha_n)$ for $n = 1, \dots, 500$. For the nonparametric estimator of G , I choose the bandwidth $h = 1.06\hat{\sigma}_w N^{-1/5}$, where $\hat{\sigma}_w$ is the standard deviation of the sample $\{w_n\}_n$, and the Kernel function

$$k(u) = 1(|u| \leq 1) \frac{35}{32}(1 - u^2)^3.$$

Figure 3 Bank Contract

After estimating \hat{G} , I obtain \hat{w} using equation (3.1). Then I compute \hat{F} and \hat{f} .

This procedure allows us to estimate two interesting objects. The first one is the distribution of liquidity needs, which is characterized by its cumulative distribution \hat{F} and the associated density \hat{f} . Another interesting outcome of this procedure is the estimation of the bank contract itself. Hence, we can see how the bank contract $m = (c_1, c_2)$ compares to the estimated one $\hat{m} = (\hat{c}_1, \hat{c}_2)$.

Figures 2 and 3 depict the estimation outcomes. Figure 2 depicts the distribution of aggregate liquidity needs. We can see that the estimates, \hat{F} and \hat{f} , provide a good approximation of the true cumulative distribution F and its density f . Figure 3 shows the graph with the true bank contract $m = (c_1, c_2)$ and the estimated one $\hat{m} = (\hat{c}_1, \hat{c}_2)$ for comparison. The contracts are extremely close.

3. TESTING FOR THE EXISTENCE OF BANK RUNS

A bank-run equilibrium is defined as an equilibrium where depositors withdraw not because they have liquidity needs but because they believe all other depositors are withdrawing. When estimating a bank-run model, an important question the econometrician could have in mind is whether a bank-run equilibrium exists or not. In this section we discuss possible econometric tests to address this question.

Consider a more general formulation of our model where the support of F is any interval $[\alpha_l, \alpha_h] \subset [0, 1]$. I focus on this generalization because, in the previous formulation of the model, a bank-run equilibrium always exists if $\alpha_h = 1$, so the question of whether a bank run exists would be uninteresting. In the context of this generalized model, the condition for the existence of a bank-run equilibrium is that

$$\int_0^{\alpha_h} u(c_1(\alpha))d\alpha \geq u(c_2(\alpha_h)). \quad (4.1)$$

The left-hand side of the inequality is the expected utility of a *patient* depositor if he decides to withdraw in period one, while the right-hand side of the inequality is his expected utility if he decides to withdraw in period two—all conditional on every other depositor withdrawing in the first period.

Define the propensity to run as

$$P = \int_0^{\alpha_h} u(c_1(\alpha))d\alpha - u(c_2(\alpha_h)). \quad (4.2)$$

For an econometrician who has the prior that the model is capable of explaining observed bank runs, the null hypothesis is that $P \geq 0$. The advantage of formulating the problem in this way is that this is a hypothesis that can be empirically tested. That is, define the statistic \hat{P} as

$$\hat{P} = \int_0^{\hat{\alpha}_h} u(\hat{c}_1(\alpha))d\alpha - u(\hat{c}_2(\hat{\alpha}_h)). \quad (4.3)$$

Traditional econometric methods can be applied to derive the distribution of \hat{P} , build confidence intervals, and, ultimately, test the hypothesis that $P \geq 0$. That is, the properties of \hat{P} can be used to test whether the model generates bank runs or not.

The test I propose here is essentially different from what is tested in most empirical literature on bank runs. While the focus of the existing empirical literature is to test whether or not past episodes of bank runs were due to coordination failure by estimating the model, the test I am proposing can reveal whether bank runs can happen or not in equilibrium.

Testing for the possibility of a bank run in the model is helpful in two ways. First, this allows us to test the theory itself. Second, if the theory is successful in explaining past run episodes, it can be used to also inform policymakers of which markets and institutions are vulnerable to runs prior to a run happening, when measures can still be taken to prevent them.

The model can also extend to be state contingent, so the propensity to run can be a function $P(\theta)$, where θ contains relevant information

such as economic growth and unemployment. This would allow us to predict under which conditions runs are more likely to happen.

4. CHALLENGES AND PRACTICAL PROBLEMS

Many assumptions are necessary to use the estimation procedure discussed above. Below I discuss some of the difficulties an econometrician would face when taking the model to a particular dataset.

The data requirement for the estimator is very weak, an econometrician only has to observe the early withdrawals in the Diamond-Dybvig model. One issue that arises, however, is that it is not clear how to match early withdrawals in the model with the data. Each application requires the econometrician to define what in Diamond-Dybvig is labeled as *early* versus *late* dates. For the XFABN studied by Foley-Fisher et al. (2015), for example, the answer seems natural. These notes feature specific dates when investors have the option to extend their notes. In other settings, however, the answer may be more challenging.

Before taking the model to the data, an econometrician also has to decide what exactly is the unit of observation the Diamond-Dybvig model represents. Does it represent the entire financial sector? Or does it represent particular financial institutions? If financial institutions have access to a complete set of liquidity contracts, then liquidity demand that is idiosyncratic to one of them does not matter because they would insure against using the available liquidity contracts. In this case, only liquidity demand in the banking sector as a whole matters for allocations. However, if financial institutions do not have access to a complete set of liquidity contracts, then each one should be considered in isolation as a unit of observation.

Another difficulty the model suggests is that, similar to estimation of auctions, combined identification of risk aversion and distribution is challenging. In the estimation procedure I propose here, I assume that the risk-aversion parameter of depositors' utility is known (or it could be separately identified). Once the risk-aversion parameter is known, the econometrician can use the second-order differential equation that characterizes the optimal contract to pin down the distribution of liquidity needs in the economy from the distribution of total withdrawals. However, the map between the two is only unique because the risk-aversion parameter is known. That is, just information on the distribution of total withdrawals would not be enough for the econometrician to identify risk aversion and distribution of aggregate liquidity needs.

A problem, similar to the one created by having to identify the risk-aversion parameter, would also arise if the econometrician has to

identify the Lagrange multiplier of the incentive-compatibility constraint. In the version of the Diamond-Dybvig model I study, the incentive-compatibility constraint of the patient depositors does not bind because the utility function of patient and impatient depositors is the same. In the original Diamond-Dybvig model, however, the utility of a patient depositor is ρ times the utility of an impatient depositor. Hence, the preferences I use are a particular case of Diamond and Dybvig (1983) where ρ equals one. Under the general formulation used in Diamond and Dybvig (1983), the incentive compatibility can bind and the solution to a second-order differential equation that characterizes the optimal contract would depend on the Lagrange multiplier associated with this constraint. Diamond and Dybvig (1983) assume that ρR is greater than one. If we assume the same in our model, the incentive compatibility does not bind for the same reason it does not bind when ρ equals one. However, the econometrician would still have to identify the preference parameter ρ .

Another crucial assumption I make is that the econometrician can identify periods when depositors played the run and no-run equilibrium. This can be challenging in practice for many reasons. In particular, there seems to be a lot of disagreement among economists, after an episode of high demand for liquidity, over whether such episode was caused by fundamental liquidity demand or by a self-fulfilling run. If it is the former, the econometrician should keep this observation in the sample; if it is the latter, he should exclude it. However, if the econometrician eliminates observations with high liquidity demand because he mistakenly identifies those as runs, he would create a sample selection problem and bias the estimator. That is because his sample would be $w_n = (\alpha_n)$, but the α_n would not be drawn from F because he is excluding with some probability observations of high α .

An econometrician would have a similar problem if depositors are more likely to run when the realization of the aggregate liquidity demand is high. Imagine a situation where depositors use as a coordination device a “sunspot” variable x that is correlated with the aggregate liquidity need α . In this case, if the econometrician excludes observations where there is a run, his sample would suffer selection issues and the estimator would be biased. The issue is the same as before, the α_n would not be drawn from F because he is excluding with some probability observations α_n that correlate with the realization of the sunspot variable that leads depositors to run.

5. CONCLUSION

Green and Lin (2000) have called for a complete theory of bank runs—a theory that explains why bank runs happen and also why society is unable to design mechanisms to prevent such bad outcomes. However, a scientific theory is only complete once it is consistent with empirical observations. Thus, particular examples that generate bank runs, as the literature has provided, are important steps toward a complete theory of bank runs, but they are not the final step.

In order to move closer to this final goal, in this paper I attempt to illustrate how the theory can be taken to the data by providing an approach to estimate the version of the Diamond-Dybvig model proposed by Peck and Shell (2003) and extended by Sultanum (2014). The estimator builds on the literature that studies the estimation of auctions. In particular, it builds on the indirect nonparametric approach to estimate first-price auctions proposed by Guerre et al. (2000). I believe this exercise can provide us with a laboratory to think about issues relating to the estimation of the model.

The exercise highlights many challenges we have to handle in order to successfully take this model to the data. However, I believe the main message of this paper is very positive. Many of these challenges have been faced by economists in different fields, and we can borrow many of the tools they have developed.

REFERENCES

- Campo, Sandra, Emmanuel Guerre, Isabelle Perrigne, and Quang Vuong. 2011. "Semiparametric Estimation of First-Price Auctions with Risk-Averse Bidders." *Review of Economic Studies* 78 (January): 112–47.
- Cassola, Nuno, Ali Hortaçsu, and Jakub Kastl. 2013. "The 2007 Subprime Market Crisis Through the Lens of European Central Bank Auctions for Short-Term Funds." *Econometrica* 81 (July): 1309–45.
- Diamond, Douglas W., and Philip H. Dybvig. 1983. "Bank Runs, Deposit Insurance, and Liquidity." *Journal of Political Economy* 91 (June): 401–19.
- Ennis, Huberto M., and Todd Keister. 2009a. "Bank Runs and Institutions: The Perils of Intervention." *American Economic Review* 99 (September): 1588–1607.
- Ennis, Huberto M., and Todd Keister. 2009b. "Run Equilibria in the Green-Lin Model of Financial Intermediation." *Journal of Economic Theory* 144 (September): 1996–2020.
- Foley-Fisher, Nathan C., Borghan Narajabad, and Stephane H. Verani. 2015. "Self-fulfilling Runs: Evidence from the U.S. Life Insurance Industry." Board of Governors of the Federal Reserve System Finance and Economics Discussion Series 2015-032 (March).
- Green, Edward J., and Ping Lin. 2000. "Diamond and Dybvig's Classic Theory of Financial Intermediation: What's Missing?" Federal Reserve Bank of Minneapolis *Quarterly Review* 24 (Winter): 3–13.
- Green, Edward J., and Ping Lin. 2003. "Implementing Efficient Allocations in a Model of Financial Intermediation." *Journal of Economic Theory* 109 (March): 1–23.
- Guerre, Emmanuel, Isabelle Perrigne, and Quang Vuong. 2000. "Optimal Nonparametric Estimation of First-Price Auctions." *Econometrica* 68 (May): 525–74.
- Hortaçsu, Ali, and Jakub Kastl. 2012. "Valuing Dealers' Informational Advantage: A Study of Canadian Treasury Auctions." *Econometrica* 80 (November): 2511–42.

- Kastl, Jakub. 2011. "Discrete Bids and Empirical Inference in Divisible Good Auctions." *Review of Economic Studies* 78 (July): 974–1014.
- Krasnokutskaya, Elena. 2011. "Identification and Estimation of Auction Models with Unobserved Heterogeneity." *Review of Economic Studies* 78 (January): 293–327.
- Pagan, Adrian, and Aman Ullah. 1999. *Nonparametric Econometrics*. Cambridge: Cambridge University Press.
- Peck, James, and Karl Shell. 2003. "Equilibrium Bank Runs." *Journal of Political Economy* 111 (February): 103–23.
- Schmidt, Lawrence, Allan Timmermann, and Russ Wermers. 2016. "Runs on Money Market Mutual Funds." *American Economic Review* 106 (September): 2625–57.
- Sultanum, Bruno. 2014. "Optimal Diamond-Dybvig Mechanism in Large Economies with Aggregate Uncertainty." *Journal of Economic Dynamics and Control* 40 (March): 95–102.
- Wallace, Neil. 1988. "Another Attempt to Explain an Illiquid Banking System: The Diamond and Dybvig Model with Sequential Service Taken Seriously." Federal Reserve Bank of Minneapolis *Quarterly Review* 12 (Fall): 3–16.