Discussion of the Papers:

"Sharing the Risk of Settlement Failure" by Fujiki, et al.; "Risks in Interlinked Settlement Systems: How to Measure the Impact of Settlement Delay in Italian RTGS System" by Impenna and Masi; and "Systemic Risk, Interbank Relations and Liquidity Provision by the Central Bank" by Freixas, et al.

William C. Hunter*

Federal Reserve Bank of Chicago

The papers presented in this session address interesting and important questions regarding the design and operation of payments systems and each makes a contribution to the literature. In my discussion, I will first comment on the two theory papers and then turn to the empirical paper dealing with the costs of settlement delay.

Let me start with the elegant and technically demanding paper by Fujiki, Green, and Yamazaki. First, let me say that this paper is very promising and will be important in setting a new standard in the modeling of payments systems risk and related issues. Informal discussions of the costs of settlement failure often focus on the indirect effects on parties who are not direct participants in a transaction. For example, the "cascades of defaults" story has the failure of bank A to make its payments to B adversely affect a bank C that has no direct dealings with bank A, as long as C has transactions with B. So a question arises whether it would be better to simply have all transactions bilateral, with no involvement of third (fourth, fifth, etc.) parties.

I interpret the Fujiki, et. al. paper as seeking to justify involvement of parties in the settlement system who have no direct relation to the transacting parties. The rationale for such involvement is that these "extraneous" parties can share the risk of payments failure (with appropriate compensation paid to take on this insurance role), resulting in a Pareto-improved outcome.

The key feature of the model in the paper involves the specification of preferences of the agents. As structured, the specification implies that in the absence of risk agent 4 would not trade with agents 1, 2, or 3. The reason is that agent 4 places low value on agent 3's endowment and agent 3 places low value on agent 4's endowment. In contrast, agents 1, 2, and 3, place high value on the endowment of each other and hence are willing to trade with each other. So, in the case

* I am indebted to David Marshall and James Moser for helpful discussions of the papers presented in this session. Errors of interpretation are my own. The views expressed are my own and not necessarily those of the Federal Reserve Bank of Chicago or the Federal Reserve System.

of certainty, agent 4 is somewhat a fifth wheel. The authors introduce the possibility of payments failure into the model by having the transfer from agent 2 to agent 3 fail with positive probability. In this case, agent 3 would like to receive an insurance payment from agent 4 conditional on the failure of the agent 2 to agent 3 transfer. To get this "insurance policy," agent 3 must make an unconditional transfer to agent 4 (the "insurance premium"). The paper rigorously proves that the equilibrium where this holds is Pareto optimal.

The questions I have regarding this major result of the paper revolve around the interpretation the authors give to agent 4. They interpret agent 4 as a government or central bank or something like a government guarantee (like the guaranteed finality in FEDWIRE, for example). While this is a plausible interpretation, it does raise the question of how dependent the insurance result is on the particular preference specifications assumed for the agents in the model. As the model is specified, agent 4 has his own preferences and endowment. However, one would think that a true central bank or system guarantor would have preferences which are a function (directly or indirectly) of all agent payments (that is, the payments among agents 1, 2, and 3). In the banking interpretation of the model, the reason is that bank failure is presumed to have systemic or system-wide consequences that the central bank does not like. So, an obvious question is how would things change if bank failure or agent payment transfer failure is directly affected by agent 4's (the central bank's) utility? A related question relates to the reasonableness of interpreting agent 4 as a "clearing house coalition" that stands ready to provide liquidity in the event of a payment failure by one of its members. The appendix to the paper suggests that this is reasonable. I am not totally convinced of this. In particular, a model based on a clearing house coalition would be more symmetric than the set-up in this paper. Each agent would have a positive probability of direct transfer failure, but would stand ready to provide funds in the event that a transfer fails to which it is not a direct party. I think that this "clearing house coalition" could be modeled using a strategy analogous to this paper, and would provide a theoretical rationale for the development of private clearing arrangements with risk sharing. These arrangements are quite common and were even more common in the era prior to Federal Reserve guarantees. It should be quite easy to show that where all agents' utility involve all others' consumption and where all agents are risk averse, their willingness to share in the risk of failure (losses) will be proportional to their individual degrees of risk aversion and not necessarily the volume of transactions. Another question I have involves the role of moral hazard in the model presented in the paper. Specifically, it would appear that moral hazard considerations would come into play since the role of agent 4 is to guarantee the system.

The paper by Freixas, et al. attempts to model the incentives of clearinghouse participants to default in a game-theoretic context. In the model, each bank receives deposits, which can be prematurely liquidated (at a gross return below unity) or can be held to maturity (which pays a gross return exceeding unity). In addition, agents can forego the banking sector and "store" their wealth, for a gross return of unity. Some of the depositors need liquid assets early (the "early diers," in the language of the paper). The remaining depositors ("late diers") move to a new location where they present a check on their original deposits to the bank in their new location. The new bank gives the late dier funds in exchange for the check, and then presents the check to the original bank. Now, the original bank could have prematurely liquidated its entire holdings, in which case this check is worthless. If every bank did this, we would be in the "gridlock equilibrium" described in the paper. This would be gridlock, because no payments would be made. In contrast, the original bank could have prematurely liquidated nothing, in which case our "late dier" could write a check for his entire deposits, and this check would be honored in full. If every bank did this, we would be in the "credit-line equilibrium" described in the paper. There is a continuum of possible equilibria between these two poles.

The paper provides a welfare ranking of the various equilibriums. The credit-line equilibrium dominates autarky (in which agents forego the banking system and use only storage), which in turn dominates the gridlock equilibrium. Unfortunately, the credit line equilibrium is unstable. That is, it is subject to coordination failure. Therefore, the policy question is how to set up a bank transfer mechanism that minimizes the danger of gridlock? Unfortunately, I think the paper says too little about how to implement the credit line equilibrium. While there is some comment on the role of real time gross settlement systems and deposit insurance, it seems to me that the paper is somewhat incomplete along this dimension.

It should be noted that the paper focuses on the problem of "strategic default": The gridlock equilibrium in the paper arises because banks choose not to honor checks drawn on their accounts. I am not sure that this type of strategic default is really what we are afraid of as central bankers. I would think that in a dynamic version of the game described in the paper, reputation effects (or simple legal punishments for default) would extinguish the "gridlock." Rather, the problem central bankers fear most is that resulting from an exogenous payment system failure which reduces the volume of credit transactions (with the accompanying welfare losses). I believe that a version of this model could be constructed that formally addresses this issue.

The second part of the paper is a multi-bank extension of an earlier Freixas/Parigi paper recently published in the Journal of Financial Intermediation. In this part, there is an exogenous bank-specific shock that can induce bank failure. The bank's depositors have information that their bank will fail, but this information is private. So, they have an incentive to write checks against their failed bank. This possibility induces all banks to not honor all checks, resulting in gridlock. That is, the credit-line equilibrium cannot be implemented. As is obvious, if the shock realization is publicly observable ex post, the central bank can implement the credit line equilibrium by guaranteeing overdrafts to solvent institutions ("collateralized overdrafts").

A key feature of this second part (and of the earlier paper) is that inefficiencies arise if bad banks are not closed early. In other words, we want to liquidate banks that cannot honor checks before the checks are actually written. To do so in this model, we need the depositors (who know whether their bank is good or bad) to run all bad banks rather than writing checks on them as the checks would allow one bank's failure to spread to other banks. Is this realistic? In a world with deposit insurance, depositors, no matter how knowledgeable, can not be counted on to discipline bad banks by running. It is up to the deposit insurer to do so. Thus, this part of the paper is more relevant to bank closure policy. It is not clear what this part of the paper contributes to payments system design, per se.

The final paper I want to discuss (by Impenna and Masi) examines risks in interlinked settlement systems. The paper attempts to measure the maximum cost that a disruption of the Italian payments would have, taking into consideration the fact that two linked settlement systems exist: the large-value payments system BIREL, which is a real-time gross settlement system, and the securities settlement system, SSS, which involves net settlement. The paper does not describe in detail the exact nature of the linkages. This inhibits the reader's ability to carefully judge the severity of the potential problems which could arise out of the interrelationships. For example, the claim that settlement delays in the netting system (SSS) could result in a settlement delay in the RTGS system, potentially mitigating the supposed advantages of RTGS, are difficult to verify.

The paper constructs a stylized model of the "maximum cost of settlement delay" due to a delay in the SSS link. The paper simply takes the average daily net outflows of the cash leg of SSS, and multiplies that by an estimate of the probability of delay. This, in a sense, gives the "expected shortfall" due to a delay, under the assumption that the amount of the shortfall is uncorrelated with the likelihood of a delay in payments. However, it is not clear that this is a "maximum cost." It looks more to be an "average cost," under this independence assumption.

The key equation in the paper defines the so-called probability of settlement delay. The underpinnings of this equation are not well spelled-out in the paper. This probability is increasing in the variability of daily cash outflows of SSS. Also, this probability is decreasing in the turnover ratio. So, less reserves imply a higher probability of delay. The model also has the probability of delay increasing in the quantity, nf/NDB, the ratio of the number of daily settlement failures to the number of negative daily balances in the whole system. This is very opaque. Why model the probability of settlement delay as in equation (3) of the paper? Why choose the functional form given in equation (3), where the elasticities of the probability of delay with respect to CV and nf/NBD are unity and the elasticity of this probability with respect to V is minus one? How do we know that p in equation (3) is actually a true probability measure (i.e., between zero and one)? It is very important that one gets these elasticities right, since the key result of the paper is that the probability is relatively small in the case of the Italian payments system. The quantitative estimates of this probability (given in Table 3) will obviously be affected by the elasticity assumptions embedded in equation (3) of the paper.

The contribution of this paper lies in its attempt to quantify the costs of settlement failure, using actual daily data. However, I am not totally convinced that the framework chosen by the authors is the most robust one for estimating this cost. At best, the model seems to give an estimate of the "average cost of an average delay." However, the concern of regulators is not with "average costs," but with "tail events"; that is, low probability events that trigger a breakdown in the payments system, with potential large-scale real economic effects. In other words, our concern is with the maximum cost that can occur (with some small, but non-zero, probability) over some reasonable time period. The reason we are concerned with tail events is our belief that the cost of systemic events is not linear in the magnitude of the event. Rather, we have an (unmodelled) belief that the social cost of systemic crises is convex in the size of the crisis. So, learning about the probability of small problems does not tell us much about the true social cost of systemic failure. Perhaps the authors will focus their attention on this issue in their future research.

In closing, I want to note that each of these papers offers unique insights into payments systems operations and risks. I encourage the authors to continue their important work and the readers of this volume to study these papers carefully.