

# Sequential Defaulting in Financial Networks

Pál András Papp

ETH Zürich, Switzerland

apapp@ethz.ch

Roger Wattenhofer

ETH Zürich, Switzerland

wattenhofer@ethz.ch

---

## Abstract

We consider financial networks, where banks are connected by contracts such as debts or credit default swaps. We study the clearing problem in these systems: we want to know which banks end up in a default, and what portion of their liabilities can these defaulting banks fulfill. We analyze these networks in a sequential model where banks announce their default one at a time, and the system evolves in a step-by-step manner.

We first consider the *reversible model* of these systems, where banks may return from a default. We show that the stabilization time in this model can heavily depend on the ordering of announcements. However, we also show that there are systems where for any choice of ordering, the process lasts for an exponential number of steps before an eventual stabilization. We also show that finding the ordering with the smallest (or largest) number of banks ending up in default is an NP-hard problem. Furthermore, we prove that defaulting early can be an advantageous strategy for banks in some cases, and in general, finding the best time for a default announcement is NP-hard. Finally, we discuss how changing some properties of this setting affects the stabilization time of the process, and then use these techniques to devise a *monotone model* of the systems, which ensures that every network stabilizes eventually.

**2012 ACM Subject Classification** Applied computing → Economics; Theory of computation → Market equilibria; Theory of computation → Network games

**Keywords and phrases** Financial Network, Sequential Defaulting, Credit Default Swap, Clearing Problem, Stabilization Time

**Digital Object Identifier** 10.4230/LIPIcs.ITCS.2021.57

**Related Version** An archive version is available at <https://arxiv.org/abs/2011.10485>.

## 1 Introduction

The world's financial system is a highly complex network where banks and other financial institutions are interconnected by various kinds of contracts. These connections create a strong interdependence between the banks: if one of them goes bankrupt, then this also affects others, causing a cascading effect through the network. Such ripple effects also had an important role in the financial crisis of 2008, and hence there is an increasing interest in the network-based properties of these systems.

One fundamental question in these networks is the so-called *clearing problem*: given a network of banks and contracts, we need to decide which of the banks can fulfill their payment obligations, and which of the banks cannot, and thus have to report a default. This question is of high interest both for financial authorities and for the banks involved.

With two simple kinds of contracts, one can already build a financial network model that captures a wide range of phenomena in real-life financial systems. Previous work has mostly focused on the equilibrium states in these models, i.e. the fixed final states where the recovery rates of banks are consistent with their current assets and liabilities. However, in practice, most events in a financial system happen gradually, one after another: a single



© Pál András Papp and Roger Wattenhofer;  
licensed under Creative Commons License CC-BY  
12th Innovations in Theoretical Computer Science Conference (ITCS 2021).

Editor: James R. Lee; Article No. 57; pp. 57:1–57:29



Leibniz International Proceedings in Informatics  
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

46 bank announces a default, which might prompt another bank to reevaluate its situation, and  
 47 also need to call being in default. This sequential development is an inherent part of the way  
 48 financial networks behave, and as such, it is crucial to understand.

49 In particular, there is a range of natural questions that only arise if we study how  
 50 the system develops in a step-by-step fashion. Can we reach every equilibrium state in a  
 51 sequential manner? How does the ordering of default announcements influence the final  
 52 outcome? Is there an optimal strategy of timing the announcements, either from a financial  
 53 authority's or a single bank's perspective? How long can the sequential process last, and in  
 54 particular, is it guaranteed to always stabilize eventually?

55 In this paper we analyze the development of financial systems in a sequential model,  
 56 where banks update their situation one after another. We first study the *reversible model*,  
 57 which is a natural sequential setting in such networks. We analyze this model from three  
 58 main perspectives:

- 59 ■ *Stabilization time:* We show that a system can easily keep running infinitely in this  
 60 model. Moreover, the time of stabilization heavily depends on the ordering of default  
 61 announcements. We also present a more complex system that does stabilize eventually,  
 62 but only after exponentially many steps.
- 63 ■ *Globally best solution:* We show that finding the ordering which results in the smallest  
 64 (or largest) number of defaulting banks in the final state is NP-hard.
- 65 ■ *Defaulting strategies:* We study the best defaulting strategy of a single bank, and show  
 66 that surprisingly, a bank may achieve the best outcome by announcing its default as early  
 67 as possible. We also prove that in general, finding the best time to report a default is  
 68 NP-hard.

69 Moreover, since the possibly infinite runtime is the most unrealistic aspect of this model,  
 70 we analyze the reasons behind this phenomenon, and we discuss how it can be avoided in  
 71 our sequential model.

- 72 ■ *Monotone sequential model:* We show that with two minor changes to the setting (a more  
 73 sophisticated update rule and a slightly different handling of defaulting banks), we can  
 74 develop a monotone model variant where the recovery rate of banks can only decrease,  
 75 and the system is always guaranteed to stabilize after quadratically many steps. We also  
 76 compare this setting to the reversible model in terms of defaulting strategies.

## 77 **2 Related Work**

78 The network-based analysis of financial systems has been rapidly gaining attention in the last  
 79 decade. Most studies are based on the early financial network model of Eisenberg and Noe  
 80 [11], which only assumes simple debt contracts between the banks. The propagation of shocks  
 81 has been analyzed in many variants of this base model over the last decade [9, 5, 4, 1, 13, 15];  
 82 in particular, the model has been extended by default costs [20], cross-ownership relations  
 83 [24, 12] or game-theoretic aspects [6].

84 However, the common ground in these model variants is that they can only describe *long*  
 85 *positions* between the banks: a better outcome for one bank always means a better (or the  
 86 same) outcome for other banks. This already allows us to capture how the default of a single  
 87 bank can cause a ripple effect in the system, but it also ensures that there is always an  
 88 equilibrium which is simultaneously best for all banks [11, 20]. As such, long positions cannot  
 89 represent e.g. the opposing interests of banks in real-world financial systems. In particular,

90 banks in practice often have *short positions* on each other when a worse situation for one  
91 bank is more favorable to another bank, mostly due to various kinds of financial derivatives.

92 The recent work of Schuldenzucker et. al. [22] presents a more refined model where the  
93 network also contains credit default swaps (CDSs) besides regular debt contracts. CDSs are  
94 financial derivatives that essentially allow banks to bet on the default of another bank in the  
95 system; they have played a dominant role in the financial crisis of 2008 [14], and have been  
96 thoroughly studied in the financial literature [10, 16]. While CDSs are still a rather simple  
97 kind of derivative, they already allow us to model short positions in the network; as such,  
98 their introduction to the system leads to remarkably richer behavior. In our paper, we also  
99 assume these two kinds of contracts in the network.

100 The work of [22, 23] discusses various properties of this new model: they show that  
101 systems may have multiple solutions (equilibrium states) in this model, and finding a solution  
102 is PPAD-complete. They also show that with default costs, these systems might not have a  
103 solution at all, and deciding whether a solution exists becomes NP-hard. The work of [19]  
104 studies a range of objective functions for selecting the best solution in this model, showing  
105 that the best equilibrium is not efficiently approximable to a  $n^c$  factor for any  $c < \frac{1}{4}$ . The  
106 work of [18] analyzes the model from a game-theoretical perspective, discussing how the  
107 removal or modification of contracts can lead to more favorable equilibria for the acting  
108 banks, and showing that such operations can lead to game-theoretical dilemmas.

109 However, all these results only analyze the model in terms of equilibrium states. This  
110 is indeed important when the market is hit by a large shock, and a central authority has  
111 to analyze the whole system, identify its equilibria, and possibly select one of them to  
112 artificially implement. However, apart from these rare occasions, the network mostly evolves  
113 sequentially, with banks announcing defaults in a step-by-step manner. For an understanding  
114 of real-world networks, it is also essential to study this gradually developing behavior of the  
115 process besides the equilibrial outcomes.

116 Sequential models of financial networks have already been studied in several papers;  
117 however, most of them consider some variant of the debt-only model with long positions  
118 [7, 2]. The paper of [22] notes that sequential clearing in their model would be dependent  
119 on the order of defaults, but does not investigate this direction any further. At the other  
120 end of the scale, the work of [3] introduces a very general sequential model (where payment  
121 obligations can be a function of all banks and all previous time steps), with a specific focus  
122 on expressing concrete real-world examples in this setting. As such, to our knowledge, there  
123 is no survey that considers a simple network model with both long and short positions, and  
124 analyzes the step-by-step development of financial systems in this model.

125 Finally, we point out that the clearing problem indeed has a high relevance in practice,  
126 e.g. when financial authorities conduct stress tests to analyze the sensitivity of real-world  
127 networks. One concrete example for a study of this problem is the European Central Bank's  
128 stress test framework [8].

## 129 **3 Model Definition**

### 130 **3.1 Banks and contracts**

131 Our financial system model consists of a set of *banks* (or *nodes*)  $B$ . We denote individual  
132 banks by  $u$ ,  $v$  or  $w$ , and the number of banks by  $n = |B|$ . Banks are connected by two kinds  
133 of contracts that both describe a specific payment obligation from a debtor bank  $u$  to a  
134 creditor bank  $v$ . The amount of payment obligation is called the *weight* of the contract.

135 The simpler kind of connection is a *simple debt* contract, which obliges the debtor  $u$  to

## 57:4 Sequential Defaulting in Financial Networks

136 pay a specific amount  $\delta$  to the creditor  $v$ . This liability is unconditional, i.e.  $u$  owes this  
137 amount to  $v$  in any case.

138 Besides debts, banks can also enter into *conditional debt contracts* where the payment  
139 obligation depends on some external event in the system. One of the most frequent forms  
140 of such a conditional debt is a credit default swap (*CDS*), which obliges  $u$  to pay a specific  
141 amount to  $v$  in case a specific third bank  $w$  (the *reference entity*) is in default. More  
142 specifically, if  $w$  can only fulfill a  $r_w$  portion of its payment obligations (known as the  
143 recovery rate of  $w$ ), then a CDS of weight  $\delta$  implies a payment obligation of  $\delta \cdot (1 - r_w)$  from  
144  $u$  to  $v$ . For simplicity, we assume that all conditional debt contracts are CDSs.

145 In practice, CDS contracts can, for example, be used by a bank as an insurance policy  
146 against the default of its debtors. If  $v$  suspects that its debtor  $w$  might not be able to fulfill  
147 its payment obligation, then  $v$  can enter into a CDS contract (as creditor) in reference to  
148  $w$ ; if  $w$  goes into default and is indeed unable to pay, then  $v$  receives some payment on this  
149 CDS instead. However, banks may also enter into CDSs for other reasons, e.g. speculative  
150 bets about future developments in the market. As a sanity assumption, we assume that no  
151 bank can enter into a contract with itself or in reference to itself.


152 Besides the contracts between banks, a financial system is described by the amount of  
153 funds (in financial terms: *external assets*) owned by each bank, denoted by  $e_v$  for a specific  
154 bank  $v$ . The external assets and the incoming payments describe the total amount of assets  
155 available to  $v$ , while the outgoing contracts describe the total amount of payment obligations  
156 of  $v$ . If  $v$  is not able to fulfill all these obligations from its assets, then we say that  $v$  is *in*  
157 *default*. If  $v$  is in default, then the fraction of liabilities that  $v$  is able to pay is the *recovery*  
158 *rate* of  $v$ , denoted by  $r_v$ . Note that  $r_v \in [0, 1]$ , and  $v$  is in default if  $r_v < 1$ . We represent the  
159 recovery rates of all banks in a recovery rate vector  $r \in [0, 1]^B$ .

160 For an example, consider the financial system in Figure 1a with 3 banks. The banks have  
161 external assets of  $e_u = 2$ ,  $e_v = 1$  and  $e_w = 0$ . Bank  $u$  has a debt of weight 2 towards both  $v$   
162 and  $w$ , and there is a CDS of weight 2 from  $w$  to  $v$ , with  $u$  as the reference entity. In this  
163 network,  $u$  has a total payment obligation of 4, but only has assets of 2, so  $u$  is in default,  
164 with a recovery rate of  $r_u = \frac{2}{4} = \frac{1}{2}$ . Bank  $u$  must use its funds of 2 to pay 1 unit of money  
165 to both  $w$  and  $v$ , proportionally to its obligations. Since  $r_u = \frac{1}{2}$ , the CDS from  $w$  to  $v$  will  
166 induce a payment obligation of  $2 \cdot (1 - r_u) = 1$ . The payment of 1 coming from  $u$  allows  
167  $w$  to fulfill this obligation to  $v$ , thus narrowly avoiding default (hence  $r_w = 1$ ). Finally,  $v$   
168 receives 1 unit from both  $u$  and  $w$ , has funds of 1 itself, and no payment obligations, so it  
169 has a positive equity of 3, and  $r_v = 1$ .

170 For convenience, we will use a simplified version of this notation in our figures: we only  
171 show the weight  $\delta$  of a contract when  $\delta \neq 1$ , and we only show the external assets of  $v$   
172 explicitly if  $e_v \neq 0$ . We also write  $e_v = \infty$  to conveniently indicate that  $v$  can pay its  
173 liabilities in any case.

174 We also note that many of our constructions in the paper contain banks that have the  
175 exact same amount of assets and liabilities, like  $w$  in this example. This is a somewhat  
176 artificial ‘edge case’ that still ensures  $r_w = 1$ . However, this is only for the sake of simplicity;  
177 we could avoid these edge cases by providing more assets to the banks in question.

178 Finally, we point out that contracts in a real-world financial system are often results of  
179 an earlier transaction between the banks, i.e. the creditor  $v$  previously offering a loan to  
180 the debtor  $u$ . We assume that such earlier payments are implicitly represented in  $e_u$ , and  
181 as such, the external assets and the contracts are together sufficient to describe the current  
182 state of the system.



■ **Figure 1** Two example systems on 3 banks. External assets are shown in rectangles besides the banks. Simple debts are denoted by blue arrows from debtor to creditor, while CDSs are denoted by light brown arrows from debtor to creditor, with the payment obligation shown beside the arrow.

### 3.2 Assets, liabilities and equilibria

We now formally define the liabilities and assets of banks in our systems. Note that due to the conditional debts, the payment obligations in a network are always a function of the recovery rate vector  $r$ .

Assuming a specific vector  $r$ , the liability  $l_{u,v}$  of a bank  $u$  towards a bank  $v$  is defined as the sum of payment obligation from  $u$  to  $v$  on all contracts, i.e.

$$l_{u,v}(r) = \delta_{u,v} + \sum_{w \in V} \delta_{u,v}^w \cdot (1 - r_w),$$

where  $\delta_{u,v}$  is the weight of the simple debt contract from  $u$  to  $v$  (if this contract exists, and 0 otherwise), and  $\delta_{u,v}^w$  is the weight of the CDS from  $u$  to  $v$  in reference to  $w$  (if it exists, and 0 otherwise). The total *liability* of  $u$  is simply the sum of liabilities to all other banks:  $l_u(r) = \sum_{v \in V} l_{u,v}(r)$ .

However, the actual *payment*  $p_{u,v}$  from  $u$  to  $v$  can be less than  $l_{u,v}$  if  $u$  is in default. If  $u$  is in default, then it has to spend all of its assets to make payments to creditors. Most financial system models assume that in this case,  $u$  has to follow the *principle of proportionality*, i.e. it has to make payments proportionally to the corresponding liabilities. This means that if  $u$  can pay an  $r_u$  portion of its total liabilities, and it has a liability of  $l_{u,v}$  towards  $v$ , then the payment from  $u$  to  $v$  is  $p_{u,v}(r) = r_u \cdot l_{u,v}(r)$ .

On the other hand, we can define the *assets* of a bank  $v$  as the sum of  $v$ 's external assets and its incoming payments in the network; that is,

$$a_v(r) = e_v + \sum_{u \in V} p_{u,v}(r).$$

If  $v$  is in default, then all these assets are used for  $v$ 's payment obligations; otherwise,  $a_v - l_v$  of these assets remain at  $v$ . Note that while both  $a_v(r)$  and  $l_v(r)$  are formally a function of  $r$ , we often simplify this notation to  $a_v$  and  $l_v$  when the recovery rate is clear from the context.

Recall that the recovery rate of  $v$  indicates the portion of payment obligations that  $v$  is able to fulfill. As such, a valid choice of  $r_v$  requires  $r_v = 1$  if we have  $a_v \geq l_v$ , and  $r_v = \frac{a_v}{l_v}$  if  $a_v < l_v$ . For simplicity, let us introduce a separate function  $R$  to denote this dependence on  $a_v$  and  $l_v$ ; that is, we define the function  $R : [0, \infty) \times [0, \infty) \rightarrow [0, 1]$  as

$$R(a, l) = \begin{cases} 1, & \text{if } a \geq l \\ \frac{a}{l}, & \text{otherwise.} \end{cases}$$

We say that a vector  $r \in [0, 1]^B$  is an *equilibrium* (or a *clearing vector*) of the system if for each bank  $v \in B$ , we have  $r_v = R(a_v(r), l_v(r))$ ; that is, if the recovery rate vector

is consistent with the assets and liabilities it generates in the network. Previous work has mostly focused on the analysis of different equilibrium states. Recall that while it is mostly straightforward to find the equilibrium states in our example constructions, the problem is PPAD-hard in general [23].

We have already seen a simple example equilibrium in Figure 1a; for another example that is slightly more challenging to compute, let us consider Figure 1b. Here bank  $u$  is again always able to pay its liabilities, so  $r_u = 1$  in any case. Furthermore, neither  $r_v = 1$  nor  $r_w = 1$  can provide an equilibrium in this network, so both  $v$  and  $w$  must be in default in any solution. Thus any equilibrium must have

$$r_v = \frac{a_v}{l_v} = \frac{1 + 1 - r_w}{3} \quad \text{and} \quad r_w = \frac{a_w}{l_w} = \frac{3 \cdot r_v}{2}.$$

This implies that the only equilibrium is  $r_v = \frac{4}{9}$ ,  $r_w = \frac{2}{3}$ .

### 3.3 Sequential models of defaulting

We have defined the equilibria of the system as the states  $r$  that would fulfill the payment criteria if every bank were to simultaneously update its recovery rate to  $r$ . However, in practice, the announcement of defaults usually happens in a sequential manner, due to different sources of delay in the system: even if it is clear from  $a_v$  and  $l_v$  that a bank  $v$  is only able to fulfill a specific  $r_v$  portion of its liabilities, this might not be immediately known to the creditors of  $v$  (due to incomplete information), or the legal framework may first allow  $v$  to try to obtain further funds before officially having to announce its default. As such, the officially announced recovery rate  $r_v$  might not always equal  $R(a_v, l_v)$ , and  $v$  has to explicitly announce the changes in  $r_v$  in order to make other banks aware of this situation.

Hence in our sequential model, each step of the process will consist of a single bank announcing an *update* to its recovery rate. That is, given the assets  $a_v$  and liabilities  $l_v$  currently available to  $v$ , if the official recovery rate  $r_v$  does not equal  $R(a_v, l_v)$ , then bank  $v$  can (and eventually has to) announce a new official recovery rate of  $r_v := R(a_v, l_v)$ . Since this affects both the payments received by the debtors of  $v$  and the payment obligations on CDSs in reference to  $v$ , it can have various effects on the system, providing new assets and liabilities to some banks; as a result, these banks may also end up with a higher or lower asset/liability balance than their currently announced recovery rate, and thus they will also have to execute a new update at some point.

More formally, we consider discrete time steps  $t = 0, 1, 2, \dots$ . Each step consists of a single bank  $v$  announcing an update to  $r_v$ . That is, if  $v$  has assets  $a_v^{(t-1)}$  and liabilities  $l_v^{(t-1)}$ , but a recovery rate of  $r_v^{(t-1)} \neq R(a_v^{(t-1)}, l_v^{(t-1)})$  at time  $t - 1$ , then we say that  $v$  is *updatable* at time  $t - 1$ . In each time step  $t$ , we select a bank  $v$  that is updatable at time  $t - 1$ , and define the state of the system at time  $t$  by (i) setting  $r_v^{(t)} = R(a_v^{(t-1)}, l_v^{(t-1)})$  for the bank  $v$  that executes the update, (ii) setting  $r_u^{(t)} = r_u^{(t-1)}$  for every other bank  $u \neq v$ , and (iii) calculating  $a_u^{(t)}$  and  $l_u^{(t)}$  for all  $u \in B$  based on this new vector  $r^{(t)}$ .

We assume that initially, each bank  $v$  has  $r_v^{(0)} = 1$ , and we compute  $a_v^{(0)}$  and  $l_v^{(0)}$  accordingly. We say that the sequential process *stabilizes* in round  $t$  if there is no updatable bank in round  $t$ .

## 4 Basic Properties

We begin by discussing some fundamental properties of this sequential setting.

## 255 4.1 Reversibility and infinite cycling

256 One important property of the sequential model is that even if a bank  $v$  goes into default,  
 257 it can easily return from this default later. That is, future updates in the system might  
 258 increase the payment obligation on an incoming CDS of  $v$ , thus increasing  $a_v$  and possibly  
 259 raising  $\frac{a_v}{l_v}$  above 1 again. This is in line with real-world financial systems, where returning  
 260 from a default is also often possible if a bank acquires new assets. Due to this property, we  
 261 also refer to this setting as the *reversible model*.

262 Note that in practice, defaulting banks are often given a limited amount of time to obtain  
 263 new assets and thus reverse a default; however, our sequential setting does not define an  
 264 explicit timing of defaults (only their order), so such rules are not straightforward to include  
 265 in our model. Nonetheless, we point out that many of our example constructions also work if  
 266 we assume that defaults are only reversible for a specific (constant) number of rounds.

267 Another important property is that in a cyclic network topology, our model can easily  
 268 result in an infinite loop of updates. Consider the example in Figure 2, where the default of  $v$   
 269 indirectly provides new assets to  $v$ . Since  $r_v = 1$  initially,  $u$  must first update to  $r_u = 0$ , and  
 270 as a result,  $v$  must update to  $r_v = 0$ . However, this leads to new liabilities in the network,  
 271 providing assets to both  $u$  and (indirectly) to  $v$ , so  $u$  (and then  $v$ ) must update its rate back  
 272 to  $r_u = r_v = 1$ . This returns the system to its initial state, where  $u$  (and  $v$ ) will continue by  
 273 updating their recovery rates to 0 again.

274 If we keep repeating these few steps, then  $u$  and  $v$  alternate between  $r_u = r_v = 0$  and  
 275  $r_u = r_v = 1$  endlessly. Note that the system does have an equilibrium in  $r_u = r_v = \frac{1}{2}$ ;  
 276 however, instead of converging to this state, the banks keep on periodically repeating the  
 277 same few steps. The possibility of such behavior in a sequential setting has already been  
 278 noted in [22] or [3] before. While this looping behavior is certainly undesired, it follows  
 279 straightforwardly from the reversibility of defaults and the existence of cycles in the network  
 280 topology. As such, these situation could also occur in real-world systems, requiring a financial  
 281 authority to intervene and set the system artificially to its equilibrium.


## 282 4.2 Dependence on the order of updates

283 Another key property of the sequential model is that the final outcome becomes dependent  
 284 on the ordering of updates, i.e. whether some banks announce their default earlier or later.


285 We show a simple example of this dependence on the *branching gadget* of Figure 3, which  
 286 has already been used as a building block in the works of [23] and [19]. In this system,  
 287 neither of the two banks  $u$  and  $v$  have any assets initially, so they are unable to fulfill their  
 288 obligations. However, if  $u$  is the first one to report default (updating to a new recovery rate  
 289 of  $a_v^{(0)}/l_v^{(0)} = 0$ ), then this provides 1 unit of new assets to  $v$ , which means that  $v$  does not  
 290 default anymore; the system stabilizes with  $r_u = 0$ ,  $r_v = 1$ . Similarly, if  $v$  is the first one to  
 291 execute an update, then this provides new assets to  $u$ , and the system stabilizes with  $r_u = 1$ ,  
 292  $r_v = 0$ . Thus both banks are strongly motivated to delay their default announcement as long  
 293 as possible, as this might allow them to avoid defaulting entirely.

294 We can also note that there are further equilibrium states where both  $u$  and  $v$  are in  
 295 default, e.g. when  $r_u = \frac{1}{2}$  and  $r_v = \frac{1}{2}$ ; due to its symmetry, one might even argue that  
 296 this is the ‘fair’ equilibrium to implement. However, this equilibrium is not reachable in  
 297 any way through sequential updates; the only possible endstates of the sequential model are  
 298  $(r_u, r_v) = (0, 1)$  and  $(r_u, r_v) = (1, 0)$  as described above.


299 This shows that even in terms of the final outcome, the sequential model can significantly  
 300 differ from the static analysis of the system. This is not due to the presence of fractional



■ **Figure 2** Example system for an infinite loop in the reversible model.



■ **Figure 3** Example system where the outcome depends on the order of announcements.



■ **Figure 4** Example of an equilibrium that is not reachable in the sequential model.

301 recovery rates: we can also easily have equilibria with integer (i.e., 0 or 1) recovery rates that  
 302 is not reachable in a sequential setting. In Figure 4, bank  $u$  is the only node who can execute  
 303 an update, which immediately leads to the unique final state  $r_u = 0, r_v = r_w = 1$ . However,  
 304  $r_u = 1$  with  $r_v = r_w = 0$  also forms an equilibrium in this system, so this phenomenon is  
 305 indeed a result of the sequential nature of our model.

## 5 Results

307 We now move on to a deeper analysis of the model. We mainly focus on the length and  
 308 outcome of the sequential process, and how the ordering of updates affects these properties.

309 Since our proofs will require more complex constructions, we switch to a simpler notation  
 310 in our figures: instead of directly showing the liability  $\delta \cdot (1 - r_w)$  on a CDS, we only label  
 311 the CDS by the weight  $\delta$  and the reference entity  $w$ , or simply by  $w$  when  $\delta = 1$ . Nonetheless,  
 312 recall that each such CDS still denotes a liability of  $\delta \cdot (1 - r_w)$ .

### 5.1 Stabilization time

314 One fundamental question is the number of rounds it takes until the sequential process  
 315 stabilizes, i.e. until no node can execute an update anymore. We first analyze this aspect in  
 316 detail.

317 We have already seen in Figure 2 that even in simple examples, it can easily happen that  
 318 the system does not stabilize at all.

319 ► **Corollary 1.** *There is a system which never stabilizes.*

320 Furthermore, with the appropriate ordering of default announcements, we can also obtain  
 321 any finite value as a stabilization time.

322 ► **Lemma 2.** *For any integer  $k$ , there exists a system and an ordering such that the system  
 323 stabilizes after exactly  $k$  steps.*

324 **Proof.** Consider the system on Figure 5. Similarly to Figure 2, this system allows us to  
 325 produce an arbitrarily long sequence by switching only  $u$  and  $v$  repeatedly. However, when  
 326  $w$  announces a default, then both  $u$  and  $v$  gain enough assets to fulfill their obligations, so  
 327 the system stabilizes after at most 2 more updates.

328 This allows us reach any magnitude of stabilization time, apart from a constant offset.  
 329 We can then simply add  $O(1)$  more independent defaulting nodes to reach the desired value  
 330  $k$ . ◀



331 This already shows that stabilization time can heavily depend on the order of updates. A  
 332 more extreme case of this is when the choice of the first update already decides between two  
 333 very different outcomes for the system.

334 ► **Lemma 3.** *There is a system where depending on the first update, the system either*  
 335 *stabilizes in 1 step, or does not ever stabilize.*

336 **Proof.** Figure 6 is obtained by combining the base ideas of Figures 2 and 3. In this network,  
 337 either bank  $w_1$  or  $w_2$  must execute the first update.

338 If  $w_2$  is the first to announce  $r_{w_2} = 0$ , then  $w_1$  receives a payment of 1, and the system  
 339 immediately stabilizes; no other bank will make an update.

340 However, if we update  $r_{w_1} = 0$  first, then  $w_2$  survives, but on the other hand,  $u$  receives  
 341 no assets at all. In this case, nodes  $u$  and  $v$  are in the same situation as in Figure 2, and  
 342 thus the upper part of the system will never stabilize. ◀

343 Finally, infinite loops are not the only examples of long stabilization: it is also possible  
 344 that the system does stabilize eventually, but for any ordering of updates, this only happens  
 345 after exponentially many steps.

346 ► **Theorem 4.** *There is a system where for any possible ordering, the system eventually*  
 347 *stabilizes, but only after  $2^{\Omega(n)}$  steps.*

348 **Proof sketch.** This proof requires a significantly more complex construction than our previous  
 349 statements. We only outline the main idea of the construction here, and we discuss the  
 350 details in Appendix A.


351 The first step of the proof is to build a *stable bit gadget*, which represents a mutable  
 352 binary variable. The gadget offers a simple interface to set the bit to 0 or 1 through external  
 353 conditions, and otherwise maintains its current value until the next such operation is executed.

354 Besides this, we create gadgets that describe logical states of an abstract process, similarly  
 355 to a finite automaton. We also encode conditional transitions between these state gadgets,  
 356 i.e. ensure that the system can only enter a given logical state if some banks currently have  
 357 a specific recovery rate. This allows us to describe a logical process where the next state of  
 358 the system is always determined by the current state and the current value of some stable  
 359 bit gadgets.


360 Using these tools, we can essentially design a binary counter on  $k = \Omega(n)$  bits, with  $k$   
 361 stable bits representing the bits of the counter. This counter will proceed to count from 0 to  
 362  $2^k - 1$ , and only stabilize after the counting has finished, resulting in a sequence of at least  
 363  $2^k$  steps.

364 The most challenging task is to ensure that in every step of the process, there is only one  
 365 possible update we can execute next: the appropriate next step of the counting procedure.  
 366 To achieve this, we not only need to ensure that some banks become updatable at specific  
 367 times, but we also have to force the banks to indeed execute these updates, by encoding  
 368 them as requirements in the transition conditions of our logical states. This results in a  
 369 heavily restricted construction where there is essentially only one valid ordering of updates:  
 370 the one that corresponds to the step-by-step incrementation of the binary counter. ◀


371 Note that another possible approach for measuring the stabilization time of our systems  
 372 is to consider the number of *defaulting steps*, i.e. to only count the steps when a bank  $v$   
 373 updates from  $r_v = 1$  to  $r_v < 1$ . One can check that our results on stabilization time also  
 374 hold for this alternative metric.



■ **Figure 5** Example system for Lemma 2. Recall that a CDS labeled with  $v$  still describes a payment obligation of  $1 - r_v$ .



■ **Figure 6** Example system for Lemma 3, where stabilization time depends on the choice of the first update.



■ **Figure 7** Example system for Lemma 6, i.e. where the number of defaults depends on the choice of the first update.

375 Finally, as a theoretical curiosity, we point out that our binary variable and state machine  
 376 gadgets in the proof of Theorem 4 demonstrate that we can essentially use financial networks  
 377 as a model of computation. We discuss the expressive power of this model in Appendix C.

378 ► **Theorem 5.** *We can use financial networks to simulate any Turing-machine with a finite*  
 379 *tape.*

## 380 5.2 Outcome with the fewest defaults

381 In case of a larger shock, a financial authority could also be interested in the final state of  
 382 the system, and in particular, the number of banks that end up in default. This can again  
 383 heavily depend on the order of updates; in fact, even a single decision in the ordering can be  
 384 critical from this perspective.

385 ► **Lemma 6.** *Depending on the first update, the number of defaults can be either  $O(1)$  or*  
 386  *$n - O(1)$ .*

387 **Proof.** Consider the system on Figure 7. If  $u$  is the first to report a default with  $r_u = 0$ ,  
 388 then  $v$  receives 1 unit of payment, and thus no other node defaults. On the other hand, if  $v$   
 389 reports a default first, then  $u$  survives, but all the nodes in the lower chain have no incoming  
 390 assets, and thus they all have to report a default eventually. So based on the first update,  
 391 the number of defaults is either 1 or  $n - 3$ . ◀


392 Hence if the authority has some influence over the ordering of updates, e.g. by allowing  
 393 more flexibility to some banks than to others, then it could dramatically reduce the number  
 394 of banks that end up in default. Unfortunately, even if we have complete control over the  
 395 ordering, it is still hard to find the best possible ordering (in terms of the number of defaults  
 396 in the final outcome).

397 ► **Theorem 7.** *It is NP-hard to find the number of defaulting nodes in the best possible*  
 398 *ordering.*

399 **Proof.** We reduce the question to the MAXSAT problem: given a boolean formula in  
 400 conjunctive normal form, the goal of MAXSAT is to find the assignment of variables that  
 401 satisfies the highest possible number of clauses [17].



■ **Figure 8** Clause gadget for the MAX-SAT reduction in Theorem 7.



■ **Figure 9** Clause gadget for the MAX-SAT reduction in Theorem 8.

402 Assume we have a MAXSAT problem on  $k$  variables  $x_1, \dots, x_k$ , and  $m$  clauses. Note  
 403 that in our financial systems, the branching gadget of Figure 3 is a natural candidate for  
 404 representing a boolean variable, since in any sequence, exactly one of  $u$  and  $v$  will eventually  
 405 default. We point out that this gadget has already been used for similar purposes before in  
 406 [23] and [19].

407 Hence for each variable  $x_i$ , we create a separate branching gadget in our system, and  
 408 consider node  $u$  to represent the literal  $x_i$ , and node  $v$  to represent the literal  $\neg x_i$ . That is,  
 409 we will consider  $x_i = \text{TRUE}$  if  $u$  defaults, while we consider  $x_i = \text{FALSE}$  if  $v$  defaults.

410 Furthermore, for each clause of the input formula, we create the clause gadget shown in  
 411 Figure 8, with the CDSs labeled by the banks representing the literals in the clause. For  
 412 example, the gadget in the figure is obtained for the clause  $(x_1 \vee x_3 \vee \neg x_4)$ . If any of the  
 413 banks  $x_1$ ,  $x_3$  or  $\neg x_4$  default, then  $v$  receives enough assets to pay its debt, whereas otherwise,  
 414  $v$  must eventually default.

415 If we aim to avoid as many defaults as possible, then the reasonable ordering strategy  
 416 is to first evaluate all the variable gadgets, and the clause gadgets only afterwards. In this  
 417 case, each bank  $v$  of a clause gadget survives if and only if there is a true literal in the  
 418 corresponding clause. This way the number of defaulting nodes in the final state is always  
 419 exactly  $k$  in the variable gadgets, and at most  $m - \text{OPT}$  in the clause gadgets, where  $\text{OPT}$   
 420 denotes the maximal number of satisfiable clauses in our MAXSAT problem. Thus the  
 421 minimal number of defaulting nodes in the system is altogether  $k + m - \text{OPT}$ . Finding this  
 422 value also allows us to determine  $\text{OPT}$ , which completes our reduction. ◀

423 To analyze the effects of a shock, one might also be interested in the worst possible  
 424 ordering; a similar reduction shows that this is also hard to find.

425 ▶ **Theorem 8.** *It is NP-hard to find the number of defaulting nodes in the worst possible*  
 426 *ordering.*

427 **Proof.** We can apply the same reduction from MAXSAT as before; we only need to slightly  
 428 change the clause gadgets. Consider the clause gadget of Figure 9 for the example clause  
 429  $(x_1 \vee x_3 \vee \neg x_4)$ . To maximize the number of defaulting banks in this system, we can first  
 430 evaluate the variables gadgets, which then allows us to produce an extra default for each  
 431 clause that has a true literal. Thus the maximum number of defaulting nodes is  $k + \text{OPT}$ ,  
 432 which completes our reduction. ◀

### 433 5.3 Individual defaulting strategies

434 It is also natural to consider the effect of the ordering from the perspective of a single bank  
 435  $v$ . More specifically, is  $v$  motivated to immediately report its own default? Can it achieve a  
 436 better outcome for itself by carefully timing its updates?

437 Intuitively, one would expect that banks are motivated to report their default as late as  
 438 possible, in hope of obtaining further assets in the meantime. This is indeed true in many

## 57:12 Sequential Defaulting in Financial Networks

439 cases. For example, in the branching gadget of Figure 3,  $u$  and  $v$  clearly have a short position  
440 in each other, and if either of them can wait long enough such that the other bank reports a  
441 default first, then it obtains new assets from the incoming CDS and thus manages to avoid a  
442 default entirely.

443 However, due to the complex interconnections in a network, it is in fact also possible that  
444  $v$  achieves a better outcome if it reports a default earlier; it is even possible that this is the  
445 only strategy which allows  $v$  to avoid a default in the endstate of the system. We consider  
446 this one of our most surprising results.

447 ► **Theorem 9.** *There exists a system where a bank  $v_1$  can only avoid a default in the final  
448 state of the system if  $v_1$  is the first bank to report a default.*

449 **Proof.** Consider the system in Figure 10, where only  $v_1$  or  $v_2$  can report a default initially,  
450 since no other node has any liabilities.

451 Assume that  $v_1$  is the first to report a default, updating to  $r_{v_1} = 0$ . This influences the  
452 network in two ways:  $v_2$  obtains assets of 1, and  $u_2$  now has a new liability of 1 as a result.

453 Thus the next update can only be executed by  $u_2$ , resulting in  $r_{u_2} = 0$ . On the one hand,  
454 this provides assets to  $u_1$ ; on the other hand, it creates liabilities for  $w_2$ . As a result, the  
455 next update can only be executed by  $w_2$ .

456 When  $w_2$  announces  $r_{w_2} = 0$ , this results in more liabilities for the defaulting  $u_2$ , and  
457 more assets for  $v_1$ . These assets make  $v_1$  the only updatable next node, bringing  $v_1$  back  
458 from its default with  $r_{v_1} = 1$ .

459 When  $v_1$  announces  $r_{v_1} = 1$ , then  $u_2$  loses some of its liabilities, and  $v_2$  loses its assets.  
460 This does not affect  $u_2$ , which remains at  $r_{u_2} = 0$  due to the default of  $w_2$ ; however,  $v_2$  now  
461 also has to report a default. The system finally stabilizes after  $v_2$  updates to  $r_{v_2} = 0$ : the  
462 assets/liabilities of  $v_1$  and  $u_1$  are affected, but neither of them has to make an update. Thus  
463 the final solution has  $r_{v_1} = 1$  and  $r_{v_2} = 0$ .


464 On the other hand, if  $v_2$  is the first to report default, then due to the symmetry of the  
465 system, the final outcome will have  $r_{v_1} = 0$  and  $r_{v_2} = 1$ . Note that in both cases, after the  
466 first update is executed, the remaining steps are already determined, and no alternative  
467 ordering is possible. Hence the only way for  $v_1$  to avoid a default in the final outcome is to  
468 be the first one to report a default. ◀

469 We can also show that in general, it is NP-hard to find the best default-reporting strategy  
470 for a bank. This even holds if the behavior of the rest of the network is ‘predictable’, i.e.  
471 if there is essentially only one ordering that the system can follow. This implies that any  
472 interpretation of this problem, e.g. optimizing a bank’s best-case payoff or worst-case payoff,  
473 is also hard.


474 ► **Theorem 10.** *It is NP-hard to find the time of defaulting that provides the highest payoff  
475 to a specific bank in the final outcome.*

476 **Proof sketch.** The main idea of the proof is to combine the binary counter construction of  
477 Theorem 4 with the MAXSAT reduction. That is, given a binary counter on  $k = \Theta(n)$  bits,  
478 we add a new node  $v$  to the system such that

- 479 (a)  $v$  can choose to default anytime,
- 480 (b) the default of  $v$  terminates the counting process, stabilizing the counter in its current  
481 state,
- 482 (c)  $v$  then comes back from its default, and its assets in the final state are proportional to the  
483 amount of clauses satisfied in a SAT formula, where the value of the variables is derived  
484 from the finalized state of the bits in the counter.



■ **Figure 10** Example where early defaulting is the best strategy, with multiple source and sink nodes for a cleaner topology.



■ **Figure 11** Infinite convergence to an equilibrium state.

485 This means that the counter essentially enumerates all the possible value assignments of the variables, and the best defaulting strategy is obtained if counting is terminated at the assignment that satisfies the highest number of clauses. However, finding this assignment is NP-hard.

489 The details of the construction are discussed in Appendix B. ◀

## 490 6 Achieving Stabilization

491 While the reversible sequential model is realistic from many perspectives, the infinite looping property is clearly not reasonable in real-world systems. As such, it is natural to ask if there is a way to modify the model to avoid this situation, and instead ensure that every financial system stabilizes eventually.

495 In this section, we investigate the causes of this infinite behavior in the sequential model. We first show that we require more sophisticated update rules to avoid a specific kind of infinite behavior, namely when the system converges to an equilibrium. We then discuss liability freezing, a different (but in some sense also realistic) approach of handling defaulting banks in the network. Finally, we show that if we combine these two modifications, we can obtain a monotone sequential model where our systems always stabilize after polynomially many steps.

### 502 6.1 More sophisticated update rules

503 **Convergence to an equilibrium.** Since the addition of conditional debt contracts drastically increases the complexity of the model, it is a natural first assumption that such an infinite pattern can only arise if the system contains a CDS. However, this is not the case: we can also obtain a (slightly different kind of) infinite sequence in systems with only regular debts.

507 Consider the example system in Figure 11. Since bank  $u$  has  $l_u = 2$  and  $a_u = 1$  initially, it can begin by updating its recovery rate to  $r_u = \frac{1}{2}$ . As a result,  $v$  and  $w$  must also announce recovery rates of  $r_v = r_w = \frac{1}{2}$ . With  $a_u = \frac{1}{2}$ , bank  $u$  now has to update to  $r_u = \frac{1}{4}$ , which then gives  $r_v = r_w = \frac{1}{4}$ . Each such round prompts another round of updates, slowly converging to  $r_u = r_v = r_w = 0$ . While this is indeed the only equilibrium of the system, the process takes infinitely many steps to reach this state.

513 **Explicit computation of equilibria.** Such a convergence process can easily occur in any network with cycles; as real-world financial systems are also known to contain cycles [21],

515 we can easily encounter such a situation in practice. In this case, it seems that a financial  
 516 authority (or the banks involved) have no other option than to explicitly compute this  
 517 equilibrium, and set their recovery rates to the appropriate values.

518 Fortunately, it is known that in case of fixed liabilities in the network (i.e. only simple  
 519 debts), this is computationally feasible: there always exists a single maximal solution that is  
 520 simultaneously best for all banks, and this solution can be found in polynomial time [20],  
 521 essentially by repeatedly solving a system of linear equations. Thus an authority could indeed  
 522 find this solution, and banks could directly update to these recovery rates in order to skip  
 523 the convergence steps.

524 This allows us to introduce the notion of *smart updates*: after each updating step, we can  
 525 consider the current liabilities in the network fixed, and we assume that the equilibrium of  
 526 the system is computed under these liabilities (essentially reducing the convergence process  
 527 to a single step). This equilibrium defines a *tentative recovery rate* for each bank  $v$ , denoted  
 528 by  $\bar{r}_v$ . In smart updates, we assume that whenever  $v$  executes an update, it always updates  
 529 to  $r_v := \bar{r}_v$ .

530 In the example of Figure 11, this means that the tentative recovery rates  $\bar{r}_u = \bar{r}_v = \bar{r}_w = 0$   
 531 are already computed initially, and thus any bank executing an update will immediately set  
 532 its recovery rate to 0. This way the process already stabilizes after each bank has executed  
 533 one update. In general, we achieve stabilization in this setting when  $r_v = \bar{r}_v$  for each bank  $v$   
 534 in the network.


535 While the explicit computation of equilibria may seem artificial, in practice, defaulting  
 536 banks are often subject to more thorough supervision by the authorities. As such, it is not  
 537 so unrealistic that the situation of a defaulting bank  $v$  is first analyzed by an authority, and  
 538 this analysis determines the official recovery rate of  $v$ .

539 Also, recall that while equilibria are easy to find in debt-only networks, the introduction  
 540 of CDSs changes this picture entirely. With CDSs, there can easily be multiple equilibria  
 541 that are Pareto-optimal, and finding any of them is already a PPAD-hard problem [23].  
 542 Thus this explicit computation of  $\bar{r}_v$  is only possible for a single step of the process, when  
 543 we consider the current payment obligation on each CDS fixed. As defaults rarely happen  
 544 simultaneously in practice, it can indeed be realistic to assume that we can analyze the  
 545 current (fixed) liabilities in the network after each new update.

546 Finally, note that smart updating is not yet enough to avoid an infinite convergence. In  
 547 the system shown in Figure 12,  $v$  can initially fulfill its obligations, while  $u$  must update to  
 548  $r_u = \frac{1}{2}$ . This creates new liabilities of 2 for  $v$ , leading to the tentative recovery rates  $\bar{r}_v = \frac{2}{3}$   
 549 and thus  $\bar{r}_u = \frac{1}{3}$  after this first step. If  $u$  updates again (to  $r_u = \frac{1}{3}$ ), then the liability on the  
 550 CDS again increases, and thus the next computed equilibrium has an even lower  $\bar{r}_u$ .

551 Each step of this process provides new tentative recovery rates, obtained as  $\bar{r}_v = \frac{2}{5-4 \cdot r_u}$   
 552 and  $\bar{r}_u = \frac{\bar{r}_v}{2} = \frac{1}{5-4 \cdot r_u}$ . This results in an infinite convergence to the equilibrium  $r_v = \frac{1}{2}$ ,  
 553  $r_u = \frac{1}{4}$ . Note that we can observe this behavior regardless of whether  $v$  ever updates its  
 554 recovery rate to the new  $\bar{r}_v$  value; the assets of  $u$  are calculated independently of the recovery  
 555 rate reported by  $v$ .

556 **Optimistic updates.** Another natural variant of smart updates is the *optimistic update* rule.  
 557 To avoid the convergence phenomenon of Figure 11, this setting also assumes that the system  
 558 is analyzed by an authority after each update. However, defaulting and non-defaulting nodes  
 559 are now handled in a different manner in this analysis. More specifically, if a bank  $v$  is not  
 560 in default (it has  $r_v = 1$  currently), then it is given the benefit of a doubt: we assume that it  
 561 can fulfill its obligations, regardless of how many assets it currently has. On the other hand,



■ **Figure 12** Example of infinite convergence in a network, even in case of smart updates. Recall that the label  $4u$  on the CDS describes a payment obligation of  $4 \cdot (1 - r_u)$ .

562 banks in default are handled the same way as in case of smart updates.

563 This distinction can indeed be realistic: if  $v$  is non-defaulting, then  $a_v$  might not even be  
 564 known to other banks, so the creditors of  $v$  have no better option than to assume that they  
 565 will receive all payments from  $v$ . On the other hand, the assets of defaulting banks are under  
 566 more thorough scrutiny in most legal frameworks.

567 Formally, optimistic update means that after each step of the process, we use a modified  
 568 version of the liability network to compute the equilibrium. Whenever there is a contract of  
 569 current weight  $\delta$  from  $u$  to  $v$  with  $r_u = 1$ , then we remove this contract from the network, and  
 570 instead (i) we add a new debt of weight  $\delta$  from  $u$  to an artificial sink node  $s$ , ensuring that  $u$   
 571 still has this liability, and (ii) we increase the value of  $e_v$  by  $\delta$ , ensuring that  $v$  always has  
 572 these assets. In contrast, if  $r_u < 1$ , we do not execute any changes on the outgoing contracts.  
 573 This modified network ensures that until a bank reports a default, its lack of assets does not  
 574 affect its creditors. We then use the same algorithm of [20] to find the equilibrium in this  
 575 modified system, and set the next tentative recovery rates accordingly.

576 Revisiting the system in Figure 12, we see that bank  $u$  can again first update to  $r_u = \frac{1}{2}$ ,  
 577 which results in  $\bar{r}_v = \frac{2}{3}$ . However, with optimistic updates,  $u$  cannot make an update again:  
 578 until  $v$  adjusts its recovery rate to this new value, the tentative recovery rate of  $u$  remains  $\frac{1}{2}$ ,  
 579 since we still expect to get the entire payment from the non-defaulting  $v$ . Note, however,  
 580 that optimistic updating still does not prevent an infinite convergence in this system if, for  
 581 example,  $u$  and  $v$  keep on updating alternately.

## 582 6.2 Liability freezing

583 We have seen that neither smart nor optimistic updating prevents an infinite sequential  
 584 process in itself. For this, we also need to change another aspect of our model, namely how  
 585 the contracts of  $v$  are handled once  $v$  goes into default.

586 Debts are rather simple from this perspective: they describe a previously established  
 587 payment obligation in the network, so there is no incentive to change them if  $v$  defaults.

588 CDSs, however, pose a more complicated question, since they describe payment obligations  
 589 that are dynamically changing. So far, we assumed that even after  $v$  defaults, the payment  
 590 obligations on its CDSs keep changing as the reference entities are updated. Another possible  
 591 approach is to assume *liability freezing*: whenever  $v$  goes into default, the liabilities on any  
 592 incoming or outgoing CDS are fixed at the current value for the rest of the process. That is,  
 593 a CDS with weight  $\delta$  and reference entity  $w$  at time  $t$  is essentially converted into a simple  
 594 debt contract with weight  $\delta \cdot (1 - r_w^{(t)})$ , and this weight does not change in the future, even  
 595 if  $r_w$  is updated.

596 This can be realistic when there is a larger time difference between subsequent defaults:  
 597 by the time the next default happens, the previous bank has already completed the first  
 598 phase of the insolvency process, and its incoming/outgoing payments have been established

599 and fixed. Indirectly, such a framework suggests that if  $v$  defaults, then it is expected to  
 600 immediately ‘cash in’ its incoming debts and fulfill its payment obligations, and not wait for  
 601 a more favorable situation.

602 The main advantage of this approach is that if we combine liability freezing with optimistic  
 603 updates, it provides a *monotone sequential model* where recovery rates can only decrease  
 604 throughout the process. Intuitively, when a bank  $w$  makes an update, then CDSs in reference  
 605 to  $w$  could only provide more assets to a bank  $v$  if we still have  $r_v = 1$ , as otherwise the  
 606 liability on the CDS is already fixed. However, if  $r_v = 1$ , then the optimistic approach  
 607 assumes anyway that  $v$  can pay its liabilities, and thus the update has no effect on other  
 608 banks in the system.

609 This monotonic property ensures that any system stabilizes eventually in this model; on  
 610 the other hand, it also means that once a bank  $v$  announces a default in this model, it has  
 611 no possibility to reverse this default in the future, and its recovery rate can only get smaller  
 612 with further updates.

613 By revisiting Figure 2, we can observe that both liability freezing and optimistic updates  
 614 are crucial ingredients to achieve this monotonicity. Without liability freezing, the system  
 615 loops infinitely if  $u$  and  $v$  make updates in an alternating fashion, both with smart and with  
 616 optimistic updates. On the other hand, if we combine liability freezing with smart updates,  
 617 then  $v$  can still alternate between  $r_v = 0$  and  $r_v = 1$  indefinitely; if  $u$  never makes an update,  
 618 then the liability on the CDS will never be fixed at a specific value.

### 619 6.3 Stabilization in the monotone model

620 We now discuss the main properties of the monotone model. We first show that the model  
 621 indeed ensures an eventual stabilization for any ordering. The key observation for this is  
 622 that the recovery rate of banks can never increase in this model.

623 ► **Theorem 11.** *The recovery rate of a bank can only decrease in the monotone model.*

624 **Proof.** The main idea is to show that for any bank  $v$ ,  $\bar{r}_v$  can only increase if we still have  
 625  $r_v = 1$  currently. This shows that we can never have  $\bar{r}_v > r_v$ , and thus no update can  
 626 increase  $r_v$ .

627 Assume that node  $w$  updates  $r_w$  in a specific step, and assume for contradiction that this  
 628 is the first step that increases  $\bar{r}_v$  for some bank  $v$  with  $r_v < 1$ . This means that the current  
 629 update is still a decrease of  $r_w$ , since we must have  $\bar{r}_w < r_w$ . The update of  $r_w$  can have two  
 630 kinds of effects on the system: it can change the liabilities on CDSs that are in reference to  
 631  $w$ , and it can result in a lower amount of assets for the creditors of  $w$ . We analyze these two  
 632 effects separately.

633 Since the monotone model has liability freezing, the liability on a CDS from  $u$  to  $v$  (in  
 634 reference to  $w$ ) can only change if we currently still have  $r_u = r_v = 1$ . Thus while this  
 635 extra payment may increase  $\bar{r}_v$ , we will still have  $\bar{r}_v \leq r_v$  afterwards. Since the model uses  
 636 optimistic updates and  $r_u = r_v = 1$ , both  $u$  and  $v$  only have debts towards the artificial sink  
 637  $s$  in the input graph of the equilibrium algorithm (which computes the tentative recovery  
 638 rates), so the changes to  $\bar{r}_u$  and  $\bar{r}_v$  do not affect the tentative recovery rate of any other  
 639 node.

640 As for the creditors of  $w$ , we consider two cases. If this is not a defaulting step (we  
 641 already had  $r_w < 1$  before the update), then updating  $r_w$  does not change the liabilities in  
 642 the input graph of the equilibrium algorithm (apart from the case of some non-defaulting  
 643 nodes, as discussed above), so the tentative recovery rates will remain unchanged.



644 On the other hand, if this is a defaulting step, then the outgoing debts of  $w$  will now be  
 645 redirected from  $s$  to the actual creditors of  $w$ . However, this operation can only result in  
 646 less assets for a bank. More specifically, one can observe that any configuration of payments  
 647 in this new graph is also a valid configuration of payments in the original graph before the  
 648 redirection step. Hence if the  $\bar{r}_v$  value of any bank  $v$  increases with this step, then this  
 649 contradicts the fact that the previous  $\bar{r}_v$  was obtained from a maximal equilibrium of the  
 650 system. ◀

651 ▶ **Theorem 12.** *The monotone model allows at most  $n$  defaulting and  $O(n^2)$  updating steps.*

652 **Proof.** Since recovery rates are always decreasing, every bank can default at most once, thus  
 653 the number of defaulting steps is at most  $n$ .

654 For the  $O(n^2)$  upper bound, we show that there are at most  $n$  updating steps between  
 655 any two consecutive defaulting steps. This is rather straightforward: recall from the proof of  
 656 Theorem 11 that if bank  $w$  executes a non-defaulting update, then this can only change the  
 657 value of  $\bar{r}_v$  for banks  $v$  that are not in default. Thus for any bank  $v$  in default,  $\bar{r}_v$  can not  
 658 change between two defaulting steps of the process. This means that any bank can execute  
 659 at most 1 updating step between two consecutive defaulting steps, limiting the number of  
 660 steps in this period to  $n$ . ◀

661 We point out that this upper bound is asymptotically tight: we can easily construct  
 662 a system and an ordering that indeed takes  $\Omega(n^2)$  steps in the monotone model. The  
 663 construction does not even require CDSs in the network; it only contains simple debt  
 664 contracts.

665 ▶ **Lemma 13.** *There is a system with an ordering that lasts for  $\Omega(n)$  defaulting and  $\Omega(n^2)$   
 666 updating steps.*

667 **Proof.** Let  $m$  be a parameter with  $m = \Theta(n)$ , and consider Figure 13. All the banks  $w_1,$   
 668  $\dots, w_m$  will eventually report a default in this system, so the number of defaulting steps is  
 669 indeed  $m = \Omega(n)$ .

670 Let  $w_1, \dots, w_m$  report a default in this order throughout the process. After  $w_i$  has  
 671 reported a default, bank  $v$  can always decrease its recovery rate to a new value of  $r_v = \frac{m-i}{m}$ .  
 672 Finally, after each such update of  $v$ , assume that all the nodes  $u_1, \dots, u_m$  make an update  
 673 step, also announcing a new recovery rate of  $\frac{m-i}{m}$ ; they can indeed all do this due to the  
 674 update executed by  $v$ . This ordering has  $\Omega(m^2) = \Omega(n^2)$  updating steps altogether. ◀


## 675 6.4 Defaulting strategies

676 Finally, we discuss how the monotone model compares to the reversible model in terms of  
 677 defaulting strategies.


678 When finding the globally best ordering, the two models turn out to be very similar. In  
 679 fact, our proofs from Section 5.2 can also be carried over to the monotone model without  
 680 any changes.

681 ▶ **Corollary 14.** *Lemma 6 and Theorems 7 and 8 also hold in the monotone model.*

682 In terms of individual defaulting strategies, the branching gadget again provides a simple  
 683 example where late defaulting is beneficial: by delaying their updates, banks  $u$  and  $v$  can  
 684 again entirely avoid a default.



■ **Figure 13** Example system for  $\Theta(n^2)$  stabilization time in the monotone model.



■ **Figure 14** Example system where early defaulting is the best strategy in the monotone model.

685 However, early defaulting is a more difficult question in this setting. In particular, we  
 686 cannot hope for a result that is analogous to Theorem 9, since once a bank reports a default,  
 687 there is no way to reverse this in the future. Nonetheless, early defaulting can still be a  
 688 beneficial strategy in the monotone model: there are cases when a bank cannot avoid an  
 689 eventual default in any way, but early defaulting can still allow the bank to have a higher  
 690 recovery rate in the final state.

691 ► **Theorem 15.** *There exists a system where a bank  $v$  only obtains its highest possible  
 692 recovery rate in the final state of the system if  $v$  is the first bank to report a default.*

693 **Proof.** Consider the system in Figure 14, where either  $v_1$  or  $v_2$  can execute the first update.  
 694 We analyze the defaulting strategies of bank  $v_1$  in this system.

695 Assume that  $v_1$  is the first to execute a step, announcing  $r_{v_1} = \frac{3}{4}$ . This gives new assets  
 696 to  $v_2$  (resulting in  $\bar{r}_{v_2} = 1$ ), and new liabilities to  $u_2$  (resulting in  $\bar{r}_{u_2} = 0$ ). The next update  
 697 can only be executed by  $u_2$ , setting  $r_{u_2} = 0$ ; at this point, the system stabilizes.

698 On the other hand, assume that  $v_2$  first announces  $r_{v_2} = 0$ . This provides  $\bar{r}_{v_1} = 1$  and  
 699  $\bar{r}_{u_1} = 0$ , so as a next step,  $u_1$  will announce a default. However, this results in new liabilities  
 700 for  $w$ , so as a next step,  $w$  has to update to  $r_w = \frac{1}{3}$ . With this,  $v_1$  only has 2 assets altogether,  
 701 so  $v_1$  must announce  $r_{v_1} = \frac{1}{2}$ . Hence  $v_1$  achieves a lower recovery rate in the final state if it  
 702 is not the first bank to announce a default.

703 Note that with some further modifications, we can also make the example symmetric to  
 704 ensure that both  $v_1$  and  $v_2$  are motivated to be the first one to default. ◀

705 Finally, one might also wonder if the monotone model allows an analogous result to  
 706 Theorem 10, i.e. a hardness result on finding the best defaulting strategy of a single bank.  
 707 However, note that the simple formulation of Theorem 10 was possible due to the fact that  
 708 the proof construction only allowed one possible ordering in the rest of the system.

709 If we were to introduce a similar setting in the monotone model, then the banks could  
 710 always find the best outcome in polynomial time, since the sequence can only last for  $O(n^2)$   
 711 steps. As such, in the monotone model, we can only expect similar hardness results for more  
 712 complex formulations of this problem, such as finding the best defaulting time with respect  
 713 to, e.g., the best-case or worst-case ordering of the remaining banks in the system.

714 ——— **References** ———

- 715 **1** Daron Acemoglu, Vasco M Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. The  
716 network origins of aggregate fluctuations. *Econometrica*, 80(5):1977–2016, 2012.
- 717 **2** Daron Acemoglu, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. Systemic risk and stability in  
718 financial networks. *American Economic Review*, 105(2):564–608, 2015.
- 719 **3** Tathagata Banerjee and Zachary Feinstein. Impact of contingent payments on systemic risk  
720 in financial networks. *Mathematics and Financial Economics*, 13(4):617–636, 2019.
- 721 **4** Marco Bardoscia, Stefano Battiston, Fabio Caccioli, and Guido Caldarelli. Pathways towards  
722 instability in financial networks. *Nature Communications*, 8:14416, 2017.
- 723 **5** Stefano Battiston, Guido Caldarelli, Robert M May, Tarik Roukny, and Joseph E Stiglitz. The  
724 price of complexity in financial networks. *Proceedings of the National Academy of Sciences*,  
725 113(36):10031–10036, 2016.
- 726 **6** Nils Bertschinger, Martin Hoefer, and Daniel Schmand. Strategic Payments in Financial  
727 Networks. In *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*,  
728 volume 151 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 46:1–46:16,  
729 Dagstuhl, Germany, 2020. Schloss Dagstuhl–Leibniz-Zentrum für Informatik.
- 730 **7** Péter Csóka and P Jean-Jacques Herings. Decentralized clearing in financial networks.  
731 *Management Science*, 64(10):4681–4699, 2017.
- 732 **8** Stéphane Dees, Jérôme Henry, and Reiner Martin. Stamp€: stress-test analytics for macro-  
733 prudential purposes in the euro area. *Frankfurt am Main: ECB*, 2017.
- 734 **9** Gabrielle Demange. Contagion in financial networks: a threat index. *Management Science*,  
735 64(2):955–970, 2016.
- 736 **10** Darrell Duffie and Haoxiang Zhu. Does a central clearing counterparty reduce counterparty  
737 risk? *The Review of Asset Pricing Studies*, 1(1):74–95, 2011.
- 738 **11** Larry Eisenberg and Thomas H Noe. Systemic risk in financial systems. *Management Science*,  
739 47(2):236–249, 2001.
- 740 **12** Matthew Elliott, Benjamin Golub, and Matthew O Jackson. Financial networks and contagion.  
741 *American Economic Review*, 104(10):3115–53, 2014.
- 742 **13** Helmut Elsinger, Alfred Lehar, and Martin Summer. Risk assessment for banking systems.  
743 *Management science*, 52(9):1301–1314, 2006.
- 744 **14** Ingo Fender and Jacob Gyntelberg. Overview: global financial crisis spurs unprecedented  
745 policy actions. *BIS Quarterly Review*, 13(4):1–24, 2008.
- 746 **15** Matt V Leduc, Sebastian Poledna, and Stefan Thurner. Systemic risk management in financial  
747 networks with credit default swaps. *Available at SSRN 2713200*, 2017.
- 748 **16** Yee Cheng Loon and Zhaodong Ken Zhong. The impact of central clearing on counterparty  
749 risk, liquidity, and trading: Evidence from the credit default swap market. *Journal of Financial  
750 Economics*, 112(1):91–115, 2014.
- 751 **17** Christos H Papadimitriou. *Computational complexity*. Addison-Wesley, 1994.
- 752 **18** Pál András Papp and Roger Wattenhofer. Network-Aware Strategies in Financial Systems. In  
753 *47th International Colloquium on Automata, Languages, and Programming (ICALP 2020)*,  
754 volume 168 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 91:1–91:17,  
755 Dagstuhl, Germany, 2020. Schloss Dagstuhl–Leibniz-Zentrum für Informatik.
- 756 **19** Pál András Papp and Roger Wattenhofer. Default ambiguity: Finding the best solution to  
757 the clearing problem, 2020. ArXiv preprint arXiv:2002.07741.
- 758 **20** Leonard CG Rogers and Luitgard AM Veraart. Failure and rescue in an interbank network.  
759 *Management Science*, 59(4):882–898, 2013.
- 760 **21** Steffen Schuldenzucker and Sven Seuken. Portfolio compression in financial networks: Incentives  
761 and systemic risk. In *Proceedings of the 21st ACM Conference on Economics and Computation*,  
762 EC '20, page 79, New York, NY, USA, 2020. Association for Computing Machinery.
- 763 **22** Steffen Schuldenzucker, Sven Seuken, and Stefano Battiston. Clearing payments in financial  
764 networks with credit default swaps. In *Proceedings of the 2016 ACM Conference on Economics  
765 and Computation*, EC '16, pages 759–759, New York, NY, USA, 2016. ACM.

## 57:20 Sequential Defaulting in Financial Networks

- 766 **23** Steffen Schuldenzucker, Sven Seuken, and Stefano Battiston. Finding Clearing Payments  
767 in Financial Networks with Credit Default Swaps is PPAD-complete. In *8th Innovations in*  
768 *Theoretical Computer Science Conference (ITCS 2017)*, volume 67 of *Leibniz International*  
769 *Proceedings in Informatics (LIPIcs)*, pages 32:1–32:20, Dagstuhl, Germany, 2017. Schloss  
770 Dagstuhl–Leibniz-Zentrum für Informatik.
- 771 **24** Stefania Vitali, James B. Glattfelder, and Stefano Battiston. The network of global corporate  
772 control. *PloS one*, 6(10):1–6, 2011.

# 773 Appendices

## 774 **A** Binary counter construction

775 In this section, we describe the binary counter construction that proves Theorem 4.


### 776 **A.1** Stable bit gadget

777 One of the basic building blocks of this construction is the so-called *stable bit gadget*, shown  
 778 in Figure 15; we have already applied the base idea of this gadget in the proof of Theorem 9.  
 779 The gadget consists of two nodes  $v_1$  and  $v_2$ , which have an outgoing CDS in reference to  
 780 each other. Note that if some external condition sets  $r_{v_1}$  to 0, then this results in  $r_{v_2} = 0$ ,  
 781 even if we had  $r_{v_2} = 1$  before. Similarly, if we change to  $r_{v_1} = 1$ , then this results in  $r_{v_2} = 1$ ,  
 782 even if we had  $r_{v_2} = 0$  before.


783 The key property of this gadget is that it allows us to ensure that banks  $v_1$  and  $v_2$  remain  
 784 in a specific state. Assume that we initially have  $v_1$  and  $v_2$  in the state  $r_{v_1} = r_{v_2} = 1$ , and  
 785 some event (i.e. the default of an external node) creates another liability for  $v_1$ . This leads  
 786 to  $r_{v_1} = 0$ , and hence  $r_{v_2} = 0$ . However, after this point, even if the extra liabilities for  $v_1$   
 787 are removed, the banks  $v_1$  and  $v_2$  do not return to their initial recovery rate, but remain in  
 788 this new state of  $r_{v_1} = r_{v_2} = 0$  instead.

789 Hence if we add conditional assets and liabilities to both nodes of the gadget, then  
 790 activating these contracts will allow us to flip the state of the bit to 0 or 1 as required, and  
 791 then the gadget will store this state until the next such activation. More specifically, consider  
 792 the extended version of the gadget in Figure 16. The default state of the external nodes  $z_1$ ,  
 793  $z_2$  is  $r_{z_1} = 1$  and  $r_{z_2} = 1$ ; in this case,  $v_1$  and  $v_2$  are in the same situation as in Figure 15, so  
 794 they retain their current recovery rates. However, if we set  $r_{z_1} = 0$  and  $r_{z_2} = 1$ , then this  
 795 allows us to set the gadget to  $r_{v_1} = r_{v_2} = 1$ , regardless of its previous state. Similarly, if we  
 796 have  $r_{z_1} = 1$  and  $r_{z_2} = 0$ , then this allows us to set  $r_{v_1} = r_{v_2} = 0$ , regardless of the previous  
 797 state. Our construction ensures that after each such operation, the external nodes  $z_1, z_2$  are  
 798 returned to their default state  $r_{z_1} = r_{z_2} = 1$ .


799 Thus the gadget essentially acts as a memory cell for storing a single bit, which is  
 800 modifiable through the recovery rates of external banks. Our main construction will use such  
 801 gadgets to store the current bits of the binary counter, and apply these external operations  
 802 to increment the counter to the next value.



■ **Figure 15** Stable bit gadget.



■ **Figure 16** Resettable version of the stable bit gadget (the sink node is split into two for a cleaner topology).



■ **Figure 17** State gadget, obtained as the combination of a condition gadget (left) and a variant of the stable bit gadget (right).

## 803 A.2 States and conditions

804 In order to ensure that the bits are changed in the correct order for the incrementation, we  
 805 create another set of gadgets that capture the current state of the counting process, and  
 806 allows us to control the transitions between these states.

807 First, note that CDSs essentially allow us to describe specific conditions, and ensure that  
 808 an event only happens if these conditions are fulfilled. Assume that we have some nodes  
 809  $z_1, \dots, z_c$  and  $z'_1, \dots, z'_d$ , which are all ‘binary nodes’ in the sense that the system guarantees  
 810 that they always have a recovery rate of either 0 or 1. Let us first analyze the left-hand  
 811 component of the system in Figure 17 separately; we will refer to this building block as  
 812 the *condition gadget*. In this gadget, node  $u$  has incoming CDSs in reference to banks  
 813  $z_1, \dots, z_c$ , and outgoing CDSs in reference to banks  $z'_1, \dots, z'_d$ , and an outgoing debt of weight  
 814  $c$ . Furthermore, we have another node  $w$  with a liability of 1, and an incoming CDS in  
 815 reference to  $u$  which has a weight of  $c + 1$ .

816 The key property of this gadget is that it only allows  $r_w = 0$  if  $r_{z_1} = \dots = r_{z_c} = 0$  and  
 817  $r_{z'_1} = \dots = r_{z'_d} = 1$ . That is, if all these conditions are fulfilled, then  $u$  has  $c$  assets and  $c$   
 818 liabilities, thus  $r_u = 1$  and  $r_w = 0$ . However, if any of the nodes  $z_i$  have  $r_{z_i} = 1$ , then  $u$   
 819 has at most  $c - 1$  assets and  $r_u \leq \frac{c-1}{c}$ . This ensures that  $w$  receives a payment of at least  
 820  $\frac{1}{c} \cdot (c + 1) \geq 1$ , thus avoiding default with  $r_w = 1$ . Similarly, if any of the nodes  $z'_i$  have  
 821  $r_{z'_i} = 0$ , then  $u$  has at least  $c + 1$  liabilities and  $r_u \leq \frac{c}{c+1}$ . This again means that  $w$  gets a  
 822 payment of at least  $\frac{1}{c+1} \cdot (c + 1) = 1$ , thus ensuring  $r_w = 1$  again.

823 Hence this gadget allows us to select a set of binary banks, and ensure that  $w$  only goes  
 824 into default if each of these banks have the desired recovery rate. This allows us to control  
 825 the transitions between a set of states, only permitting entry into a state if a set of conditions  
 826 are fulfilled.

827 We can combine this condition gadget with a stable bit gadget to obtain our *state gadget*  
 828 as shown in Figure 17. The state gadget ensures that a specific set of conditions are fulfilled  
 829 before allowing  $w$  to default. This then sets the stable bit to 0, representing the fact that  
 830 the execution is currently in this state; we can then use bank  $v_2$  as a reference entity in  
 831 the CDSs of other condition gadgets to make some events dependent on the condition that  
 832 we are currently in this state. Once we exit the state, we can use the technique shown in  
 833 Figure 16 to ensure that the state bit is set back to 1 again. Also, note that since the bit is  
 834 stable, once we enter the state, the bit remains set to 0 until it is artificially reset with this  
 835 technique, even if the entry condition of the state becomes false in the meantime (and thus  
 836  $w$  updates back to  $r_w = 1$ ).

837 Hence the state gadgets will allow us to represent the execution of the counting process  
 838 through a set of states and transitions between these states, with the next state always  
 839 depending on our current state and possibly the value of some other stable bits in the

840 system (the bits of the counter, in our current case). By defining the appropriate conditions  
 841 (including a specific previous state) for each state, we can ensure that the only valid ordering  
 842 of the system is an execution that follows these prescribed transitions between the states.

### 843 A.3 Resetting the states

844 In order to guarantee that the system is only in one state at any point in time, we also have  
 845 to ensure that the resetting operations are indeed executed after exiting a state (i.e. that  
 846  $r_{v_1}$  and  $r_{v_2}$  is indeed updated to 1). Due to this, encoding the transition from a state  $s_1$   
 847 to another state  $s_2$  becomes a nontrivial task. The natural approach would be to enforce  
 848 the resetting of  $s_1$  by including these updates in the entry condition of  $s_2$ . However, recall  
 849 that the entry condition of  $s_2$  also requires that  $s_1$  is active (as a preceding state), so this  
 850 makes the entry condition of  $s_2$  contradictory, hence impossible to fulfill. On the other hand,  
 851 after entering state  $s_2$ , there is no straightforward way to verify the resetting of  $s_1$  anymore;  
 852 furthermore, the system already has two active states at once in this case.

853 In order to solve this problem, we take each state  $s$  of our original system design, and  
 854 replace it by three consecutive state gadgets  $\text{ENTRY}_s$ ,  $\text{RESET}_s$  and  $\text{EXIT}_s$ , known as the *entry*  
 855 *phase*, *resetting phase* and *exit phase* of  $s$ . State  $\text{ENTRY}_s$  will have the same entry conditions  
 856 as the original state  $s$  did, and any other state that was previously following state  $s$  will now  
 857 follow after state  $\text{EXIT}_s$ . The entry condition for states  $\text{RESET}_s$  and  $\text{EXIT}_s$  will be that we are  
 858 currently in states  $\text{ENTRY}_s$  and  $\text{RESET}_s$ , respectively. Thus instead of passing through state  
 859  $s$ , the execution will pass through all 3 phases of  $s$  in this predefined order in our modified  
 860 system.

861 The key idea is that the three classes of states will reset each other in a round-robin  
 862 fashion throughout the execution. State  $\text{RESET}_s$  will provide new assets to the nodes  $v_1$  and  
 863  $v_2$  of state  $\text{ENTRY}_s$ , thus resetting their recovery rate to 1. Then state  $\text{EXIT}_s$  will have it as  
 864 an entry condition that the value of the banks  $w$ ,  $v_1$  and  $v_2$  of state  $\text{ENTRY}_s$  are all set to 1.  
 865 This ensures that  $\text{EXIT}_s$  is indeed only reached when all recovery rates in  $\text{ENTRY}_s$  are set  
 866 back to their initial value. This way we can make sure that when the execution leaves  $\text{EXIT}_s$   
 867 and enters the state  $\text{ENTRY}_{s'}$  of the following state  $s'$ , then the state  $s$  is not considered  
 868 active anymore.

869 In order to ensure that  $\text{RESET}_s$  and  $\text{EXIT}_s$  are also reset to inactive, we execute the same  
 870 steps in any two succeeding states for both of them. That is, we ensure that  $\text{EXIT}_s$  will  
 871 provide new assets to reset the stable bit of  $\text{RESET}_s$  (in the same fashion that  $\text{RESET}_s$  does  
 872 this for  $\text{ENTRY}_s$ ), and we ensure that in any original state  $s'$  succeeding  $s$ , the state  $\text{ENTRY}_{s'}$   
 873 has in its entry condition that banks  $w$ ,  $v_1$  and  $v_2$  of state  $\text{RESET}_s$  are all set to 1 (in the  
 874 same fashion that  $\text{EXIT}_s$  does this for  $\text{ENTRY}_s$ ). Similarly, in order to ensure that  $\text{EXIT}_s$  is  
 875 reset to inactive, we take every original state  $s'$  succeeding  $s$ , and in  $\text{ENTRY}_{s'}$  we provide new  
 876 assets to reset the stable bit of  $\text{EXIT}_s$ , while we require in the entry condition of  $\text{RESET}_{s'}$   
 877 that banks  $w$ ,  $v_1$  and  $v_2$  of  $\text{EXIT}_s$  are set to 1.

878 Hence by representing each logical state by three consecutive state gadgets, we can  
 879 ensure that any ordering of updates is indeed forced to reset each state to inactive when  
 880 leaving the state and entering the following one. Note that this method results in a higher  
 881 number of entry conditions for each of our state gadgets, but this has no effect on the overall  
 882 construction. Furthermore, we point out that the reason to have at least three such classes is  
 883 to avoid the situation when a state gadget investigates its own banks in its entry condition,  
 884 which could lead to undesired behavior.

#### 885 A.4 Technical details of managing states

886 While this already describes the general technique of using state gadgets, there are still  
887 several details to discuss for completeness.

888 One such example is the handling of bank  $w$  in the state gadget: note that this bank is  
889 slightly differently from  $v_1$  and  $v_2$  in the sense that it does not receive extra assets to be  
890 reset, but its resetting is still checked as a condition in  $\text{EXIT}_s$ . This is because when state  
891  $\text{RESET}_s$  is reached, then the exit state  $\text{EXIT}_s$  of the previous state  $\hat{s}$  has already been reset  
892 to inactive, which means that the entry conditions of  $\text{ENTRY}_s$  are not fulfilled anymore. This  
893 allows us to update  $r_u$  to a new value of  $r_u \leq \frac{c}{c+1}$ , which then allows us to set  $r_w = 1$ , and  
894 indeed enter  $\text{EXIT}_s$ . Hence by not explicitly resetting  $w$  but still checking in  $\text{EXIT}_s$  that  $w$  is  
895 reset, we can ensure that bank  $u$  of state  $s$  is also reset to its initial state of  $r_u \leq 1$ , and thus  
896 the  $s$  state can only be reactivated if the entry conditions are fulfilled again at some point.

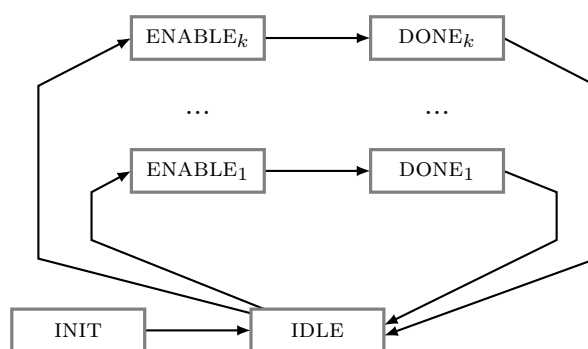
897 Furthermore, note that while our construction mostly requires us to implement logical  
898 ‘and’ relations in the entry conditions of state gadgets, we occasionally also have to implement  
899 a logical ‘or’. One such case is the central state of our binary counter which will have multiple  
900 different preceding states, i.e. there are multiple states that finish by enabling this state as  
901 the next one; to activate this state, we only require one of the preceding states to be active,  
902 and not all of them. In this specific case, it is rather simple to insert this ‘or’ condition into  
903 our state gadget: we simply add an incoming CDS in reference to each of these preceding  
904 states, but we still select the weight of the input CDS of  $w$  based on the original  $c$  value, i.e.  
905 as if there was only one preceding state. This way  $u$  receives a payment of 1 if we are in any of  
906 these preceding states (and the remaining conditions are fulfilled), and since we can not have  
907 two active states at the same time, these extra CDSs will always only result in a payment of  
908 0 or 1 for  $u$  altogether. In a more general setting (e.g. if we want to encode different further  
909 conditions for different predecessor states), we can create a separate transition state for each  
910 such condition, and then use the same method to set these transition states as predecessors.

911 Finally, our analysis has so far assumed that inactive states gadgets have recovery rates  
912 of  $r_u < 1$  and  $r_w = 1$ , and we discussed how we can maintain this invariant throughout the  
913 process. However, we also have to ensure this when initializing the construction; since we  
914 initially begin with  $r_u = 1$ , the node  $w$  of any state gadget could already report a default in  
915 the initial time step, even though the entry conditions of the state are not satisfied.

916 For this initialization step (i.e. achieving  $r_u < 1$  in each state gadget), we introduce  
917 a special node  $y_0$  whose default indicates that the system has already been initialized, i.e.  
918 that  $r_u < 1$  in each state gadget. Then we slightly modify the state gadgets such that the  
919 outgoing debt from node  $w$  is replaced by a CDS in reference to  $y_0$ ; this way no  $w$  can update  
920 before all the banks  $u$  are initialized, but after we set  $r_{y_0} = 0$ , the recovery rate of  $y_0$  will  
921 never change, and thus the outgoing contract of each node  $w$  will behave as a single debt for  
922 the rest of the process.

923 To ensure this behavior, it suffices to add a starter node  $y_1$  with an outgoing CDS in  
924 reference to each  $u$ , and carefully choose  $e_{y_1}$  such that  $y_1$  only goes into default if each bank  
925  $u$  has executed an update. Then we can include  $y_0$  in a stable bit gadget with another node  
926  $y'_0$ , with  $y'_0$  also having an outgoing CDS in reference to  $y_1$ . This system can only begin with  
927 all the state gadget nodes  $u$  executing an update. The slight default of  $u$  then also sends  
928  $y'_0$ , and then  $y_0$  into default; since  $y_0$  has no way to reverse this, it will remain in default  
929 indefinitely. We can then use the default of  $y_0$  as the trigger condition for the first real state  
930 of our counting process.





■ **Figure 18** Illustration of the main states of our binary counter system.

## 931 A.5 Overall construction

932 Given the tools to represent states, bits and conditions, the construction of the binary  
 933 counter becomes straightforward. Let us introduce a parameter  $k$  such that  $k = \Theta(n)$ .  
 934 We create  $k$  distinct stable bit gadgets that represent the  $k$  bits of a counter, which will  
 935 count from 0 to  $2^k - 1$ . We also add a set of state gadgets which control the counting  
 936 process. More specifically, we add an IDLE state to capture the state between two consecutive  
 937 incrementations. Furthermore, for each bit of the counter (i.e. each  $i \in \{1, \dots, k\}$ ), we add  
 938 two states that describe the incrementation of the  $i^{\text{th}}$  bit, and we call them  $\text{ENABLE}_i$  and  
 939  $\text{DONE}_i$ . The states of our process are illustrated in Figure 18.

940 For  $\text{ENABLE}_i$ , the entering condition is that the process is currently in the IDLE state, and  
 941 that this is indeed a valid next incrementation of the counter, i.e. that the bits at positions  
 942  $1, \dots, i - 1$  are all set to 1, and the bit at position  $i$  is set to 0. When entering  $\text{ENABLE}_i$ , we  
 943 ensure that this state sets the bits at positions  $1, \dots, i - 1$  to 0, and it sets the bit at position  
 944  $i$  to 1. The entering condition of the  $\text{DONE}_i$  state is that all of these updates are indeed  
 945 executed, and thus the counter is indeed correctly incremented. From the  $\text{DONE}_i$  state, we  
 946 lead the execution back to the IDLE state without any condition.

947 Thus the financial system indeed has essentially only one possible ordering, aside from  
 948 the fact that we are free to choose the order of updating the bits of the counter in each  
 949 incrementation. In this single ordering, the system works as a binary counter: in the idle  
 950 state, the only valid next step is to always execute the next incrementation on the counter.  
 951 Since the construction consists of only  $O(k)$  different gadgets, each having only  $O(1)$  nodes,  
 952 this indeed allows for a choice of  $k = \Theta(n)$ , and thus the counting process indeed lasts for  
 953 at least  $2^{\Omega(n)}$  steps (note that this also holds if we only count defaulting steps, i.e. when a  
 954 bank  $v$  updates from  $r_v = 1$  to  $r_v < 1$ ). Once all the bits are set to 1, there is no next state  
 955 that the process can enter from the IDLE state, so the system indeed stabilizes eventually.

956 As a technical detail, note that we must also ensure that the process begins in the IDLE  
 957 state. This can be ensured by adding a further state INIT and a further stable bit gadget with  
 958 this state. The state INIT can be entered if this stable bit is set to 1 (and the initialization  
 959 node  $y_0$  has already defaulted), so it will be the only state that the process can enter in the  
 960 beginning. We allow the process to also enter the IDLE state from INIT. However, within the  
 961 INIT state, this stable bit is set to 0, and the system does not provide a way for this stable  
 962 bit to ever be reset to 1; hence the INIT state cannot ever be entered again, and thus it plays  
 963 no role after this point.

964 Furthermore, note that for the simplest implementation, we can consider the counter bits  
 965 to be 1 when the nodes of the stable bit are in default, and 0 when they are not in default:

966 this ensures that the initial state of the system (when every bank has a recovery rate of 1)  
 967 is indeed a valid initialization of our construction. Otherwise (if we want the 0 bits to be  
 968 represented by defaulting bit gadgets), we can use further initial states to ensure that each  
 969 stable bit is initialized to the desired value before the process first enters the IDLE state.

970 Recall that, as discussed before, each state in our description will in fact be split to three  
 971 consecutive states to ensure that resetting is always executed. However, this only increases  
 972 the number of state gadgets in our system by a constant factor, and thus it has no effect on  
 973 our analysis.

## 974 **B Best defaulting strategy for a single bank**

975 In this section we prove the claim of Theorem 10, i.e. that finding the best defaulting strategy  
 976 (the best time to report a default) for a single bank is an NP-hard problem. We combine the  
 977 binary counter construction of Appendix A with the MAXSAT reduction technique to show  
 978 that any efficient algorithm that finds the best time to report a default would also provide  
 979 an efficient solution to MAXSAT.

### 980 **B.1 Overall idea**

981 The main idea of our construction is to create a binary counter system where each counter  
 982 bit represents one of the variables of our input MAXSAT formula. The counting process  
 983 then corresponds to enumerating all the  $2^k$  possible value assignments to the variables.

984 We then add a further node  $v$  to this system that wants to find the best time to report its  
 985 own default. This bank  $v$  will only have an opportunity to report a default in the IDLE state  
 986 of the counter, i.e. exactly once for each of the  $2^k$  possible assignments. When  $v$  reports a  
 987 default, this will immediately stop the counting process, thus fixing the value of the  $k$  stable  
 988 bits to their current value forever. Then for each clause of the formula, we add a clause  
 989 variable that defaults exactly if at least one of the bits corresponding to the literals in the  
 990 clause are set to true (but only after the counting has stopped). Finally, we ensure that for  
 991 each such clause node, bank  $v$  receives a unit of payment through an incoming CDS.

992 This results in a construction where, by choosing a time to report its default,  $v$  can  
 993 essentially select a value assignment to the variables, and the final amount of assets received  
 994 by  $v$  will be determined by the number of satisfied clauses under this assignment. Thus  
 995 selecting the best defaulting time for  $v$  is equivalent to selecting the best assignment for  
 996 MAXSAT, which completes the reductions.

997 Note that the behavior of the binary counter system is completely predictable, since it  
 998 only has essentially one valid ordering of updates, but it is still NP-hard to find the optimal  
 999 defaulting time for  $v$ . This implies similar hardness results in more general systems that  
 1000 have very different orderings: it is still NP-hard to find the best defaulting time if we, for  
 1001 example, assume that the remaining part of the systems follows the ordering that is the  
 1002 most/least beneficial for  $v$ .

1003 Furthermore, we note that in order to simplify our clause gadgets, we can easily extend  
 1004 the binary counter construction by a negated version of each of the  $k$  stable bits, which  
 1005 are similarly set and checked in the  $\text{ENABLE}_i$  and  $\text{DONE}_i$  states. This step essentially  
 1006 provides another counter that is counting backwards from  $2^k - 1$  to 0 simultaneously to our  
 1007 original counter, without having any effect on the magnitude of the number of nodes. More  
 1008 importantly, in our case, it provides a convenient access to the negation of each of variable;  
 1009 with this, we can directly check the value of any literal in the clauses of the formula.

## 1010 B.2 Technical details

1011 Consider the bank  $v$  for which we want to find the best defaulting strategy. In order to  
 1012 ensure that  $v$  can only report a default when the counter is in the IDLE state, we simply add  
 1013 an outgoing CDS of weight 1 to  $v$  (and select  $e_v = 0$ ). Then similarly to the design of a state  
 1014 gadget, we connect  $v$  to a stable bit gadget on banks  $w_1$  and  $w_2$  such that the default of  $v$   
 1015 will lead to the default of both  $w_1$  and  $w_2$  (i.e.  $e_{w_1} = e_{w_2} = 0$ , and  $w_1$  has outgoing CDSs in  
 1016 reference to  $v$  and  $w_2$ , while  $w_2$  has an outgoing CDS in reference to  $w_1$ ). Then  $w_1$  and  $w_2$   
 1017 can never return from this default; this will ensure that even after  $v$  receives extra assets in  
 1018 the future, the counting still does not continue.

1019 More specifically, the outgoing CDS of  $v$  is in reference to bank  $v_2$  of the  $\text{ENTRY}_{\text{IDLE}}$   
 1020 state; since this bank defaults every time when the IDLE state is visited,  $v$  indeed has the  
 1021 opportunity to report a default at each of the  $2^k$  counting phases. Then in the  $\text{EXIT}_{\text{IDLE}}$  state,  
 1022 we add it as a condition that both  $r_v = 1$  and  $r_{w_2} = 1$ ; this ensures that after  $v$  defaults, the  
 1023 counter can never enter the  $\text{EXIT}_{\text{IDLE}}$  state again, so the counting indeed stops at the current  
 1024 value.

1025 Furthermore, for each literal of the formula (i.e. each variable and its negation), we create  
 1026 a node that defaults in this final state if the literal is set to true. That is, given a literal  $\ell_i$   
 1027 (a variable or its negated version), we add a bank  $\ell_i$  representing this literal. This bank  $\ell_i$   
 1028 has an outgoing CDS in reference to  $w_2$ , and an incoming CDS in reference to the stable  
 1029 bit gadget in the counter that represents the negated version of  $\ell_i$ . This implies that (i) the  
 1030 bank  $\ell_i$  can only go into default once  $v$  has reported a default and the counter was stopped,  
 1031 and (ii) in this case, it goes into default exactly if the literal  $\ell_i$  is set to true in the chosen  
 1032 assignment.

1033 Then for each clause  $c_i$  of the formula, we simply add a bank representing  $c_i$ , and draw  
 1034 an outgoing CDS from  $c_i$  for all the banks  $\ell_i$  that correspond to a literal included in the  
 1035 clause. As such,  $c_i$  is in default exactly if at least one of the literals in the clause is true.

1036 Finally, for each such clause node  $c_i$ , we add an incoming CDS to  $v$  in reference to this  
 1037 bank  $c_i$ . With this, the incoming assets of  $v$  equal the number of satisfied clauses under the  
 1038 chosen assignment. Furthermore, we add another incoming CDS in reference to  $w_2$  in order  
 1039 to compensate for the outgoing CDS that allows  $v$  to default. With this, the total payoff of  
 1040  $v$  in the final state (the difference of its assets and liabilities) is indeed equal to the number  
 1041 of satisfied clauses in the formula. Hence the best outcome for  $v$  is indeed the assignment  
 1042 where the maximal possible number of clauses are satisfied, which is NP-hard to find.

1043 Hence whenever  $v$  reports a default, the connected stable bit is set to 0, and the counting  
 1044 stops permanently. After this point, the literal gadgets corresponding to true literals will  
 1045 have to report a default eventually, followed by the appropriate clause gadgets. Hence the  
 1046 system will indeed eventually stabilize in the state where  $v$  receives all the payments for  
 1047 the clause gadgets, so its assets in the final state are indeed defined by the quality of the  
 1048 MAXSAT assignment. This completes our reduction.

## 1049 C Financial networks as a model of computation

1050 We now make a detour to briefly discuss how financial systems can behave as a model of  
 1051 computation. While this is not very relevant for the analysis of real-world financial networks,  
 1052 it is still an interesting aspect of our reversible model from a theoretical perspective.

1053 First of all, note that we have only considered financial networks with finitely many banks;  
 1054 as such, in our base model of these systems, we cannot hope to model a general Turing  
 1055 machine (TM) with an infinitely long tape. Therefore, we only discuss how our networks can

1056 model a Turing machine which only has a finitely long tape (also known as a linear bounded  
 1057 automaton). We point out that generalizing our constructions to the infinite case would not  
 1058 be as straightforward as to simply allow infinitely many banks and contracts in the network.  
 1059 For example, in our binary counter or in the TM simulation design below, this generalization  
 1060 would lead to infinitely many state gadgets, and the initialization of these gadgets would  
 1061 already require infinitely many updates in the beginning, thus not allowing the main part of  
 1062 the process to begin after a finite amount of steps.

1063 Recall that the main tools for simulating a TM have already been introduced in the  
 1064 binary counter construction: state machines were explicitly discussed and used in the counter,  
 1065 and the stable bit gadgets are a natural candidate to simulate the tape cells of a TM over a  
 1066 binary alphabet. Note that for a very simple and crude encoding of the TM, the stable bit  
 1067 gadgets are not even needed: since a TM with a finite tape can only have finitely many valid  
 1068 configurations (in terms of current state, tape content and tape pointer position), we can  
 1069 encode the transitions between these configurations as a finite automaton, and build this  
 1070 state machine using our state gadgets.

1071 However, a much more elegant way of modeling a Turing machine with our systems is to  
 1072 indeed use stable bits as tape cells, and encode an addressing mechanism on the tape. That  
 1073 is, besides our financial subsystem representing the state machine part of the TM, we also  
 1074 create a binary counter which stores the position of the TM pointer of the tape; when the  
 1075 pointer is moved to the left or the right in a transition from one state to another, we simply  
 1076 increment or decrement the value of this counter.

1077 Then in an auxiliary state following the transition, we can use the value of the counter  
 1078 and the content of the tape to copy the content of the currently chosen tape cell to a specific  
 1079 stable bit gadget. That is, for each cell  $c$  of the tape, we have two specific states  $\text{READ}_{c,0}$  and  
 1080  $\text{READ}_{c,1}$  that are only entered as a next step if the counter value currently points to  $c$ ; state  
 1081  $\text{READ}_{c,0}$  is activated if stable bit gadget of  $c$  is currently set to 0, while  $\text{READ}_{c,1}$  is activated  
 1082 if  $c$  is set to 1. We use this extra state to copy the content of the cell to a specific stable  
 1083 bit gadget, and then we only make our next transition in the state machine based on the  
 1084 current logical state and the value of this single stable bit (as in case of a Turing machine).

1085 We can use a similar technique for overwriting the value in the current cell: we add two  
 1086 auxiliary writing states  $\text{WRITE}_{c,0}$  and  $\text{WRITE}_{c,1}$  for each cell  $c$ , and based on the value of the  
 1087 counter and the bit we want to write, only one of these states gets activated. This state then  
 1088 copies the desired bit value to the corresponding stable bit gadget of the tape.

1089 Note that this more sophisticated simulation method still requires a separate state for  
 1090 each cell of the tape. However, for a state machine of  $s$  states and an available tape of length  
 1091  $m$ , this approach can be implemented with  $O(s)$  banks in the state machine,  $O(\log m)$  banks  
 1092 in the counter, and  $O(m)$  banks for the stable bit gadgets and auxiliary states of the tape  
 1093 cells; in contrast to this, the crude approach has a  $O(s \cdot m)$  factor in the total number of  
 1094 banks. Even more importantly, this simulations method allows us keep the banks modeling  
 1095 the state machine and the banks modeling the tape separately, thus providing a much cleaner  
 1096 representation of the Turing machine in our systems.

## 1097 **D A note on further possible sequential models**

1098 In the paper, we have studied two different sequential models of our financial networks: the  
 1099 reversible model (which allows banks to return from a default if they acquire new assets)  
 1100 and the monotone model (which guarantees an eventual stabilization for any ordering in any  
 1101 network). Since both the reversibility of defaults and an eventual stabilization of the network

1102 are realistic properties in real-world financial systems, it is a natural idea to try to obtain an  
1103 even more accurate sequential model by appropriately combining these two settings. For  
1104 example, the financial framework could ensure that the liability freezing rule is applied on a  
1105 bank if, for example, it already has to report a default for the third time. Alternatively, a  
1106 financial authority could actively monitor the network to look for infinite cycling patterns,  
1107 and enforce liability freezing in specifically chosen situations. We leave it to future work to  
1108 explore a more complex (and possibly more realistic) line of models in this direction.

1109 Note that our discussion of sequential models is not exhaustive; there are various further  
1110 changes we can execute to model the sequential process slightly differently. While these  
1111 alternative models may come with some convenient properties, the changes often also  
1112 introduce new undesired side effects into the model.

1113 For a simple example of such an alternative model, assume that the financial industry is  
1114 aware of the heavily dynamic behavior of payment obligations on CDSs in these networks,  
1115 and thus the authorities decide to measure the assets of a bank by estimating an expected  
1116 incoming payment on CDSs. That is, we select a global constant  $\mu \in [0, 1]$ , and given a CDS  
1117 of weight  $\delta$  from  $u$  to  $v$  (in reference to  $w$ ), we define the incoming assets of bank  $v$  on this  
1118 CDS as  $\mu \cdot \delta \cdot r_u$ , regardless of the (dynamically changing and possibly inaccurate) current  
1119 value of  $r_w$ .

1120 One clear disadvantage of this slightly simpler model is that defaults can easily remain  
1121 undetected in the system. E.g. in the branching gadget of Figure 3, a choice of  $\mu = 1$  will  
1122 mean that neither of the two banks ever report a default, since they can both fulfill their  
1123 obligations in a possible best-case situation. On the other hand, the survival of both  $u$  and  $v$   
1124 can never be an equilibrium, since neither of them receives any assets, and thus it remains  
1125 undetected that either  $u$  or  $v$  should be in default in any ‘reasonable’ outcome. Similarly, a  
1126 choice of  $\mu = 0$  will force both  $u$  and  $v$  to report a default with  $r_u = r_v = 0$ , even though  
1127 this is not a realistic outcome in the network. While the choice of  $\mu \in (0, 1)$  seems like a  
1128 reasonable compromise, it can actually lead to both of these problems (both undetected and  
1129 false defaults) in practice, and as such, it does not allow an accurate analysis of the network.