



# Dynamic Interventions for Networked Contagions

Marios Papachristou  
papachristoumarios@gmail.com  
Cornell University  
United States

Siddhartha Banerjee  
sbanerjee@cornell.edu  
Cornell University  
United States

Jon Kleinberg  
kleinberg@cornell.edu  
Cornell University  
United States

## ABSTRACT

We study the problem of designing dynamic intervention policies for minimizing cascading failures in online financial networks, as well as more general demand-supply networks. Formally, we consider a dynamic version of the celebrated Eisenberg-Noe model of financial network liabilities, and use this to study the design of external intervention policies. Our controller has a fixed resource budget in each round, and can use this to minimize the effect of demand/supply shocks in the network. We formulate the optimal intervention problem as a Markov Decision Process, and show how we can leverage the problem structure to efficiently compute optimal intervention policies with continuous interventions, and give approximation algorithms in the case of discrete interventions. Going beyond financial networks, we argue that our model captures dynamic network intervention in a much broader class of dynamic demand/supply settings with networked inter-dependencies. To demonstrate this, we apply our intervention algorithms to a wide variety of Web-related application domains, including ridesharing, online transaction platforms, and financial networks with agent mobility; in each case, we study the relationship between node centrality and intervention strength, as well as fairness properties of the optimal interventions.

## CCS CONCEPTS

• Theory of computation → Design and analysis of algorithms.

## KEYWORDS

dynamic resource allocation, financial contagion, optimal bailouts, markov decision process, fairness

## ACM Reference Format:

Marios Papachristou, Siddhartha Banerjee, and Jon Kleinberg. 2023. Dynamic Interventions for Networked Contagions. In *Proceedings of the ACM*

Supported in part by a Simons Investigator Award, a Vannevar Bush Faculty Fellowship, AFOSR grant FA9550-19-1-0183, a Simons Collaboration grant, and a grant from the MacArthur Foundation. SB gratefully acknowledges funding from the NSF under grants ECCS 1847393 and CNS 1955997. MP was also supported by a LinkedIn PhD Fellowship, and a grant by the A.G. Leventis Foundation. The authors thank C. Caramanis, A. Jalan, E. Pierson, D. Shmoys, and A. Wang for helpful discussions regarding our paper. An extended version of the paper can be found here.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

WWW '23, April 30–May 04, 2023, Austin, TX, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9416-1/23/04...\$15.00

<https://doi.org/10.1145/3543507.3583470>

Web Conference 2023 (WWW '23), April 30–May 04, 2023, Austin, TX, USA.  
ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3543507.3583470>

## 1 INTRODUCTION

**Motivation.** The world consists of interconnected entities that interact with one another in the form of networks. Networks experience shocks due to adverse scenarios, such as (partial) failures of nodes and edges. When exogenous shocks hit networks, such shocks propagate through the edges of the network, causing cascades that may affect a significant population of nodes in the network. This phenomenon is known as *network contagion*, and a particularly interesting category of networks that undergo contagion are *financial networks*. In financial networks, financial entities, such as individuals, businesses, and banks have *liabilities* to one another as well as assets which can be attributed either *internally*, i.e., within the financial network in question, or *externally*, i.e., outside the financial network. It is often the case that the entities within these networks do not have adequate means to pay off their financial obligations, so they become default [13, 18]. A *planner* acts as an external force and is responsible for (optimally) allocating resources, also known as *interventions* (or bailouts) subject to budget constraints so that defaults are averted [1, 9, 22, 23].

When modeling and studying such networked interactions, much of the literature assumes the interactions are *static* [12, 14, 15, 18, 20, 22–24]. However, in many situations, networked interactions evolve, subject to an *uncertain environment*, and where the planner's interventions at some point in time affect the state of the system in future times, resembling the dynamics in queuing network models [2, 21], or epidemics [4, 8, 10, 11, 19]. So far, limited attention has been given to contagion processes that evolve dynamically: First of all, such approaches have either been considered in limited settings of a financial system together with specific strategies for the mitigation of systemic risk that do not fall under an intervention regime [6, 7, 16], or are based on different modeling assumptions and are computationally more intensive [3].

Moving beyond contagion in financial systems, we argue that our models have a much broader scope across a variety of Web-based applications. In particular, with small modifications, our model and algorithms can apply to resource allocation problems arising in ridesharing, ad placement, influence maximization, allocation of computational resources, and contagion in digital financial transaction networks, to name a few. At a high level, our framework can handle settings that can be modeled as a supply/demand network, subject to defaults, and where the demand is proportionally split between neighbors of a node upon its default. Consequently, we hope our ideas can find wider usage in other such settings.


**A Model for Dynamic Interventions in Networks.** We study networks subject to a contagion process over time, with the system's current state and interventions affecting its future state due to

the accumulation of liabilities at nodes. Our model extends the static Eisenberg-Noe model [13] of contagion to a dynamic setting with the accumulation of liabilities. The network is overseen by an external regulator, who has limited resources to try and intervene to stop a contagion. Designing such intervention policies, however, requires leveraging both spatial and temporal information.

We initially study the design of an optimal *fractional* intervention policy, which is a computationally tractable problem (assuming that the system responds *optimally* at every round). Next, we study the allocation of discrete resources, a computationally intractable problem whose static version has been studied in [23]. Here, we design approximate discrete intervention policies that are based on the fractional intervention policy. Note that the concurrent work of [5] studies dynamic interventions under the Eisenberg-Noe model *only* in the case where the rationing of payments among agents is *constant* (i.e., does not vary with time), which is a particular case of our framework.

**Interventions in General Demand-Supply Networks.** While our models are motivated by (and build on) existing models of liabilities and defaults in financial networks, we posit that these models and the corresponding resource allocation problems can extend far beyond this setting to account for many applications on the Web. One place where such a clearing mechanism can be utilized is in *ridesharing* [2] (Uber, Lyft, etc.). The nodes of the network represent *neighborhoods* of a city, the external world represents suburban areas of the city, and the internal liabilities between nodes represent the number of rides requested from a neighborhood  $i$  to a neighborhood  $j$ . The external liabilities (resp. external assets) correspond to rides requested from a neighborhood  $i$  to suburban areas (resp. rides from the suburban areas to a neighborhood  $i$ ). The interventions correspond to additional supply (for example, autonomous vehicles) that can be dispatched to any neighborhood. Assets represent available vehicles, and shocks represent adverse scenarios such as traffic jams. Similarly, in a *computing cluster*, nodes represent compute nodes, and liabilities represent the amount of computing that can be displaced to adjacent nodes; assets represent available computing power at each node, shocks represent failures of computing resources, and interventions represent allocations of backup resources to the existing centers (which can also be utilized, e.g., because of high demand).

More generally, any problem that corresponds to a *supply and demand network that evolves* at which, when the corresponding obligations of the node cannot be met in full, they should be distributed *proportionally* towards the nodes that demand the corresponding resources. A planner seeks to allocate resources in this environment subject to a *budget constraint* can be captured by our framework. Our framework can support the maximization of a variety of objectives. In this paper, we focus on a particular objective function. However, the solutions produced are equivalent to every strictly increasing objective. For discrete interventions (where the problem becomes NP-Hard), we provide approximation guarantees of the computed solution concerning the optimal solution that works under realistic monotonicity assumptions. As a result, our dynamic contagion framework is suitable for resolving resource allocation problems in a variety of domains (both through the lens of fractional and discrete allocations); financial transaction networks (physical



**Figure 1: Schematic representation of our problem setting.**

or on the Web [e.g., Venmo, cryptocurrencies]), ridesharing, high-performance computing, ad-placement, to name a few.

**Our Technical Contributions.** Formally, we study the problem of optimally allocating (fractional and discrete) resources subject to (i) contagion effects and (ii) an uncertain environment, i.e., an environment that experiences financial shocks. We generalize the model of Eisenberg and Noe [13] to discrete time as a *Markov Decision Process* (MDP) and formulate the optimal intervention problem (Sec. 2). The resulting MDP is very high-dimensional; nevertheless, we show how we can leverage the problem structure to compute near-optimal intervention policies for continuous interventions efficiently. Moreover, under discrete interventions, we demonstrate how the planner can use the above continuous intervention policies to derive heuristic control policies with formal approximation guarantees (Sec. 4). In addition, our framework also supports the incorporation of additional fairness constraints (based on generalizations of the Gini Coefficient) for the distribution of the resources such that interventions are more equitable across nodes; surprisingly, we also show that incorporating such objectives has little effect on the welfare objective in our setting. We supplement our theoretical results with experiments (Sec. 6) on real-world data from the Venmo transaction platform, semi-artificial data from mobility patterns, real-world data for ridesharing applications from New York City’s Taxi and Limousine Commission, and synthetic data with core-periphery structure<sup>1</sup>.

## 2 SETTING

**Notation.** We use  $[n]$  to denote the set  $\{1, \dots, n\}$ . For vectors (resp. matrices), we use  $\|x\|_p$  for the  $p$ -norm of  $x$  (resp. for the induced  $p$ -norm); for the Euclidean norm (i.e.,  $p = 2$ ), we omit the subscript.  $\mathbf{0}$  (resp.  $\mathbf{1}$ ) denotes the all zeros (resp. all ones) column vector, and  $\mathbf{1}_S$  represents the indicator column vector of the set  $S$ . We use  $x \wedge y$  (resp.  $x \vee y$ ) as shorthand for the coordinate-wise minimum (resp. maximum) of vectors  $x$  and  $y$ . For a given vector  $x$ , we use the array notation  $x(i : j)$  to denote a sub-vector of  $x$  from  $x_i$  to  $x_j$  (inclusive range). Finally order relations  $\geq, \leq, >, <$  denote coordinate-wise ordering.

**System Model (see Fig. 1).** The model we study is a natural dynamic extension of the Eisenberg-Noe model [13]. The system consists of  $n$  entities  $[n]$ , connected via a dynamic network, where each *directed* edge  $(i, j)$  denotes that entity  $i$  owes a liability to entity  $j$ .

At the start of each (discrete) round  $t$ , new *internal* liabilities  $\ell_{ij}(t) \geq 0$  get generated in the system. Moreover, let  $\tilde{P}(t-1)$  denote the *clearing vector* from round  $t-1$ , i.e., each agent clears  $\tilde{P}_i(t-1) \in [0, P_i(t-1)]$  of its liabilities in round  $t-1$  where  $P_i(t-1)$  are the total liabilities in round  $t-1$ . The rest of the

<sup>1</sup>Code: <https://github.com/papachristoumarinos/dynamic-clearing>

liabilities are getting forwarded in time. In accordance with the proportional clearing rule of the EN model, we have that the same fraction  $\left(1 - \frac{\tilde{P}_i(t-1)}{P_i(t-1)}\right)$  of each of agent  $i$ 's liabilities are cleared. Thus, the total liability between  $i, j \in [n]$  at the start of round  $t$  (i.e., before clearing) is given by

$$p_{ij}(t) = \ell_{ij}(t) + p_{ij}(t-1) \cdot \left(1 - \frac{\tilde{P}_i(t-1)}{P_i(t-1)}\right),$$

In addition, agents also experience external liabilities  $b_i(t) > 0$ , generated at the start of each round  $t$ . Let  $\ell_i(t) = \sum_{j \in [n]} \ell_{ij}(t)$ ; then the total liabilities owed by  $i$  at the start of round  $t$  are  $P_i(t) = b_i(t) + \sum_{j \in [n]} p_{ij}(t) = b_i(t) + \ell_i(t) + (P_i(t-1) - \tilde{P}_i(t-1)) > 0$ . This induces the liability network in round  $t$ , with the (weighted, directed) relative liability matrix  $A(t)$  given by  $a_{ij}(t) = \frac{p_{ij}(t)}{P_i(t)}$ . Now let  $\beta_i(t) = \sum_{j \in [n]} a_{ij}(t)$  denote the *fractional internal liability* for any node  $i$ , and  $\beta(t)$  the vector of these internal liabilities. Throughout the rest of the paper we assume the system has *non-vanishing external liabilities*, i.e., that the following holds for  $A(t)$ :

ASSUMPTION 1 (NON-VANISHING EXTERNAL LIABILITIES).

$$\|\beta(t)\|_\infty = \|A^T(t)\|_1 < 1 \text{ for all } t \in [T]$$

In order for the overall system to clear their liabilities (in particular, given there are always external liabilities), the EN model assumes that each agent  $i$  has some additional “assets”  $c_i$  which can contribute to clearing their liabilities. In the same vein, in a dynamic interaction setting, we assume each agent has external assets (or “revenue streams”), which in each round  $t$  provide an instantaneous revenue  $c(t) = \{c_i(t)\}_{i \in [n]} \geq \mathbf{0}$ . Now, as in the EN model, the clearing vectors  $\tilde{P}(t) \geq \mathbf{0}$  must satisfy the following *zero-input dynamics* constraints

$$\tilde{P}(t) \leq P(t) = b(t) + \ell(t) + P(t-1) - \tilde{P}(t-1) \quad (1a)$$

$$\tilde{P}(t) \leq A^T(t)\tilde{P}(t) + c(t) \quad (1b)$$

We refer to first constraint (Eq. 1a) as the *solvency constraint*, since if for some node  $i$  it holds with equality, it means that this node is able to repay its debts in full. We refer to the second constraint (Eq. 1b) as the *default constraint*, since when it holds with equality for some node  $i \in [n]$ , this means that  $i$  partially repays its debts proportionally to its creditors. Finally, clearing vectors are always non-negative. Note by definition the bounds in Eqs. ?? are non-negative; we can thus compress these constraints as  $\tilde{P}(t) \in [0, P(t) \wedge (A^T(t)\tilde{P}(t) + c(t))]$ .

**Dynamic Control.** Until now, we have assumed that all inputs (internal/external liabilities, and external asset values) are exogenous. We now augment this model with an additional *centralized controller*, who is provided with some resource budget in each round, and can allocate this budget to try and minimize defaults.

**Actions.** Consider a controller (or planner) who has access to a bounded amount of resource  $B \geq 0$  at each round, and seeks to inject this into the network (i.e. allocate a fractional quantity  $Z_i(t) \geq \mathbf{0}$  to each agent  $i$  subject to the constraint  $\mathbf{1}^T Z(t) \leq B$ ). The per-round allocation  $Z(t)$  may be subject to additional bounds  $Z(t) \leq L$  for some given maximum allowed allocation vector  $L$ ; ignoring this is equivalent to letting  $L \geq \mathbf{1}B$ . Therefore, the action space  $\mathcal{Z}$  is

given by  $\mathcal{Z} = \{z \in \mathbb{R}^n : \|z\|_1 \leq B, \mathbf{0} \leq z \leq L\}$ . We always assume the policy function  $Z(t)$  is Markovian.

**State transitions.** Eq. 1 can be modified to incorporate allocations  $Z(t)$  to get

$$\tilde{P}(t) \in [0, P(t) \wedge (A^T(t)\tilde{P}(t) + c(t) + Z(t))] \quad (2)$$

$$\mathbf{1}^T Z(t) \leq B, Z(t) \in [\mathbf{0}, L]$$

Note that in the above equations,  $A(t)$  is implicitly a function of  $P(t)$ , and consequently, this makes the constraints non-linear. In the full paper, we show that the necessary and sufficient condition for the set of allocations satisfying Eq. 2 to be *convex* (in  $\tilde{P}(1:T), Z(1:T)$ ) is that the relative liability matrix  $A(t)$  is constant over time (which makes the dynamics linear; see also the concurrent results of [5]).

Next, given the current state  $P(t)$ , exogenous input  $c(t)$ , and action  $Z(t)$ , a natural ‘maximal’ choice of the clearing vector  $\tilde{P}(t)$  is for it to be the *fixed point* of the system

$$S(t) = \begin{pmatrix} P(t) \\ \tilde{P}(t) \end{pmatrix} \mapsto \underbrace{\begin{pmatrix} P(t) \\ P(t) \wedge (A^T(t)\tilde{P}(t) + c(t) + Z(t)) \end{pmatrix}}_{=\Phi(S(t), Z(t); S(t-1), U(t))}. \quad (3)$$

When the round is clear from the context we will use the abbreviation  $\Phi_t(s, z)$  to denote the mapping with information up to time  $t$  acting on the state action pair  $(s, z)$ , i.e. for all  $z$  we have that  $s(z) = \Phi_t(s(z), z)$ . Under Asm. 1, we can use the Banach fixed-point theorem to assert that this has a unique solution (since  $\Phi_t$  is a contraction with respect to  $s$  for a given  $z$ ; see [18]). We henceforth make the following assumption on the agents’ response

ASSUMPTION 2 (MAXIMAL CLEARING). *In each round  $t$ , the agents maximally clear their liabilities, i.e., with  $\tilde{P}(t)$  as the unique fixed point of  $\tilde{P}(t) = P(t) \wedge (A^T(t)\tilde{P}(t) + c(t) + Z(t))$ .*

The above condition, taken from the EN model, is standard in the finance literature – it imposes a natural requirement that agents try and clear liabilities as soon as possible, subject to the proportional clearing rule. If one wants to maximize flows (or minimize defaults), one may be tempted to think that this is without loss of generality. In the context of dynamic external interventions, however, this is not the case; one can create examples where dropping this assumption leads to higher overall rewards. These settings however are somewhat extreme, and it may be possible to eliminate them via other assumptions. In the full paper we exhibit one such natural regime; finding more general conditions remains a challenging direction for future work.

This is a reasonable assumption since natural agents do not usually have knowledge of their future, have limited memory, and are not usually able to respond in a globally optimal way, given a realization of the sample path. From a mathematical viewpoint, lifting Asm. 2 leads to non-convexities which yields a hard-to-solve non-convex optimization problem to find the globally optimal policy (cf. full paper).

**Exogenous Shocks.** Till now, we have been agnostic in our model description as to the exact nature of the exogenous shocks to the system, i.e., the per-round internal and external liabilities, and external asset payouts. We assume that the environment is

stochastic, i.e. the instantaneous assets and (internal and external) liabilities induce uncertainty in the system in forms of a disturbance. We denote the financial environment at round  $t$  as  $U(t) = (b(t), c(t), \{\ell_{ij}(t)\}_{i,j \in [n]})$ . We assume that the environment is a Markov Chain:

**ASSUMPTION 3 (MARKOVIAN EXOGENOUS SHOCKS).** *The financial environment  $U(t)$  is a Markov Chain, i.e.  $\Pr[U(t) = u(t)|U(t-1) = u(t-1), \dots, U(1) = u(1)] = \Pr[U(t) = u(t)|U(t-1) = u(t-1)]$ .*

The state space of  $U(t)$  is denoted by  $\mathcal{U}$ . It is easy to observe that under Asm. 3 the sequence  $S(t) = (P(t), \tilde{P}(t))$  is a Markov Chain (MC). More specifically, at time  $t$  the only information needed to determine  $p_{ij}(t)$  is the instantaneous liabilities (which is a MC), the action  $Z(t-1)$  and the remaining liabilities from times  $t-1$ , therefore extra information from round 0 up to  $t-2$  is redundant. Since the external liabilities are also an MC and the sum of  $p_{ij}(t)$  only depends on the state of the system at  $t-1$  then  $S(t)$  is a MC based wrt to  $S(t-1)$  and  $Z(t-1)$ . Also Eq. 2 depends only on the state of the system at time  $t-1$  and the calculated maximum liabilities at the start of round  $t$  therefore the optimal clearing vector that occurs on the element-wise minimum of the RHS of the inequalities of Eq. 2 is dependent on the previous state  $S(t-1)$  and the action  $Z(t-1)$ . That defines a transition kernel  $\mathcal{T}, \mathcal{T}((s, z) \rightarrow s') = \mathcal{T}(s'|s, z) = \Pr[S(t) = s'|S(t-1) = s, Z(t-1) = z]$ . We also denote the projection on  $(s, z)$  of the kernel (which is a distribution itself) as  $\mathcal{T}(\cdot|s, z) = \mathcal{T}_{s,z}(\cdot)$ . The MC is also associated with an initial distribution over the state space  $S(0) \sim \pi_0$ . The state space of the MC is denoted by  $\mathcal{S}$ .

### 3 CONTINUOUS INTERVENTIONS

Given the above setting of networked interactions over time with stochastic shocks, we can now formulate the problem of optimal dynamic interventions for maximizing various objectives as a Markov Decision Process (MDP). In this section, we formalize this, and show that when interventions are allowed to be continuous, the MDP can be solved optimally for a wide range of objectives. Continuous actions are however often infeasible in practice, and so in the next section, we turn to the question of designing approximately-optimal controllers given discrete actions.

**Rewards & Objective.** The stochastic reward incurred by a state-action pair at time  $t$  is  $R(t) = \mathbf{1}^T \tilde{P}(t)$ . Note that here we can use any function of  $\tilde{P}(t)$  that is coordinate-wise strictly increasing due to [13, Lemma 4] and get the same solution. Thus our framework allows for the maximization of a large family of reward functions<sup>2</sup>

For simplicity, we have chosen  $R(t) = \mathbf{1}^T \tilde{P}(t)$  since it corresponds to a measure of how much money circulates in and out of the network. The overall objective that is to be maximized is the sum of rewards over a finite horizon  $[T]$ ,

$$\begin{aligned} \max_{\Pi} \quad & \mathbb{E}_{S(0) \sim \pi_0} \left[ \sum_{t=0}^{T-1} R(S(t), Z(t) = \Pi(t, S(t)), U(t)) \right] \\ \text{s.t.} \quad & (2) \quad \forall t \in [T] \end{aligned} \quad (4)$$

<sup>2</sup>Other possible reward functions can be, for example,  $R(v(t), S(t), Z(t), U(t)) = v^T(t) \tilde{P}(t)$  for some function  $v(t) > 0$ . The analysis of the approximation algorithm for these functions is similar to the one followed in [23].

where  $\Pi : [T] \times \mathcal{S} \rightarrow \mathcal{Z}$  is a policy function. We also assume that there are no accumulated debts and interventions from time  $t \leq 0$ . We let  $r(s, z) = \mathbb{E}_{U(t)} [R(S(t) = s, Z(t) = z, U(t))]$ .

**Value Function.** We define the value function  $V^\Pi(t, s)$  as the optimal reward we can collect from time  $t$  onward, starting from state  $s$  and applying policy  $\Pi$ . The value function obeys the HJB equations with respect to actions chosen from the action set, i.e.  $V(t, s) = \max_{z \in \mathcal{Z}} \{r(s, z) + \mathbb{E}_{s' \sim \mathcal{T}_{s,z}} [V(t+1, s')]\}$ . In Fig. 2 we present an example of an allocation scenario under our model for a toy dynamic network with  $T = 2$  periods.

**Efficiently Computing the Optimal Value Function** The above MDP has a very high-dimensional state and action space ( $\mathbb{R}^{2nT}$  and  $\mathbb{R}^{nT}$  respectively), so a priori it is unclear if it can be solved efficiently. Surprisingly, we show below that we can exploit the structure of the problem – in particular, the fact that the random shocks are exogenous (Asm. 3), and the maximal clearing assumption (Asm. 2) – to give a closed-form expression for the value function as an expectation over the exogenous shock vector  $U(1 : T)$ ; moreover, this also allows us to compute it efficiently (and thus find near-optimal policies) via Monte Carlo estimation.

We now proceed to show how to calculate the value function  $V(t, s)$  and the optimal policy  $Z^*(t)$ . First, it is easy to check that given a realization of the random shocks  $u(t : T)$ , the optimal reward (and policy) can be written as a sequence of nested linear programs. More surprisingly, we prove that we can exchange the maximum and expectation operators in the value function due to the structure of our model. Consequently, when the shocks are generated randomly, we get that the value function (Thm. 1) can be approximated by sampling  $N$  sample paths and then, for each sample path  $u(t : T)$  solving a sequence of  $T - t + 1$  linear programs.


Our algorithm (Alg. 1) is comprised of two routines: first, **Compute-Value-Function-Given-Sample-Path** takes as an input a realization  $u(t : T)$  of exogenous shocks, a starting state  $s(t-1) = s$ , and budget constraints  $L$  and  $B$  and solves a sequence of  $T - t + 1$  nested linear programs, where the optimal solution at round  $t$  is fed to calculate the optimal solution at round  $t+1$ . The second routine (**Aggregate**) takes as input a natural number  $N$ , the budget constraints  $L$  and  $B$ , a financial environment  $\mathcal{U}$ , and the starting state  $s$ . The algorithm then samples  $N$  exogenous shock realizations from  $\mathcal{U}$ . Conditioned on any of the sample paths  $u_i(t : T)$  with  $i \in [N]$ , it calls the first routine to compute the sample value function  $\tilde{V}_{u_i}(t, s)$ . It aggregates all solutions estimates  $\bar{V}(t, s)$ :

**THEOREM 1.** *Under Asms. ??, the following are true*

- (1) *The value function  $V(t, s)$  satisfies  $V(t, s) = \mathbb{E}_{U(t:T)} [\max_{z_t, \tilde{p}_t} \{ \mathbf{1}^T \tilde{p}_t + \max_{z_{t+1}, \tilde{p}_{t+1}} \{ \mathbf{1}^T \tilde{p}_{t+1} + \max_{z_{t+2}, \tilde{p}_{t+2}} \{ \mathbf{1}^T \tilde{p}_{t+2} + \dots \} \} \} ]$  and corresponds to solving a sequence of linear programs.*
- (2) *For  $N = \frac{\log(2/\delta)(T-t+1)^2 \Delta^2}{2\epsilon^2}$  samples,  $\epsilon > 0$ , and  $\Delta = \sup_{\mathcal{U}} (\|b\|_1 + \|\ell\|_1)$ , Alg. 1 returns a solution  $\bar{V}(t, s)$  such that  $|\bar{V}(t, s) - V(t, s)| \leq \epsilon$  with probability at least  $1 - \delta$ .*

### 4 DISCRETE INTERVENTIONS

We next focus on the problem of allocating discrete interventions. For the discrete interventions problem, each node can be allocated discrete resources up to some value  $L_j \in \mathbb{N}$ . A simpler version of



**Figure 2: A contagion network over  $T = 2$ , with instantaneous internal liabilities  $\ell$ , external liabilities  $b$ , and external assets  $c$  that are identical over the two rounds. The total budget is  $B = 2$  at each round. The optimal allocation for  $t = 1$  is  $z^*(1) = (2, 0, 0)^T$  in which case all nodes are able to cover their debts, and no liabilities are carried over from  $t = 1$  to  $t = 2$ . Similarly, in  $t = 2$  the optimal intervention vector is  $z^*(2) = (2, 0, 0)^T$  and all liabilities are cleared. The value function equals  $R(1) + R(2) = 10$ .**

---

**Algorithm 1** Dynamic Clearing With Fractional Interventions

---

**Compute-Value-Function-Given-Sample-Path**( $L, B, u(t : T), s$ )

- (1) Given the initial state at round  $t - 1$  calculate  $A(t)$  and  $P(t)$
- (2) For each  $t' \in [t, T]$ 
  - (a) Let  $u(t') = (b(t'), c(t'), \text{vec}(\{\ell_{ij}(t')\}_{i,j \in [n]}))$  be the financial environment.
  - (b) Let  $\tilde{P}^*(t'), Z^*(t')$  be the optimal solution to  $\max_{\tilde{P}(t'), Z(t')} \mathbf{1}^T \tilde{P}(t')$  subject to the dynamics of (2) and the random shocks  $b(t'), c(t'), \{\ell_{ij}(t')\}_{i,j \in [n]}$ .
  - (c) If  $t' < T$ , use  $\tilde{P}^*(t')$  to calculate  $A(t' + 1)$  and  $P(t' + 1)$ .
- (3) Return  $V_{u(t:T)} = \sum_{t' \in [t, T]} \mathbf{1}^T \tilde{P}^*(t')$ .

**Aggregate**( $N, L, B, \mathcal{U}, s$ )

- (1) Sample  $N$  i.i.d. sample paths  $\{u_i(t : T)\}_{i \in [N]} \sim \mathcal{U}$  where a sample path consists a realization of the environment on  $T - t + 1$  periods.
- (2) For every  $i \in [N]$  compute

 $V_{u_i(t:T)} = \text{Compute-Value-Function-Given-Sample-Path}(L, B, u_i(t:T), s)$ 

- (3) Return  $\bar{V}(t, s) = \frac{1}{N} \sum_{i=1}^N V_{u_i(t:T)}$ .
- 

the problem studied in [23] allowed the interventions to get two distinct values  $\{0, L_j\}$ . The analysis in this case is exactly the same with the general case. The total budget is again  $B$  and does not change with time as well. We refer the the action space of this setting with  $\mathcal{Z}_d = \{z \in \mathbb{N}^n : \|z\|_1 \leq B, \mathbf{0} \leq z \leq L\}$ . Note that  $\mathcal{Z}$  defined on Sec. 3 corresponds to the *fractional relaxation* of  $\mathcal{Z}_d$ .

We again seek to find the optimal policy which maximizes the value function at round  $t = 0$ , subject to the dynamics

$$S(t) = \begin{pmatrix} P(t) \\ \bar{P}(t) \end{pmatrix} \mapsto \underbrace{\begin{pmatrix} P(t) \\ P(t) \wedge (A^T(t) \bar{P}(t) + c(t) + Z(t)) \end{pmatrix}}_{=\Psi(S(t), Z(t); S(t-1), U(t))}. \quad (5)$$

Again, when the round is clear from the context, we will use the abbreviation  $\Psi_t(s, z)$  to denote the mapping with information up to time  $t$  acting on the state action pair  $(s, z)$ , i.e., for all  $z$  we have that  $s(z) = \Psi_t(s(z), z)$ . Similarly to  $\Phi_t$ , the operator  $\Psi_t$  is a contraction according to Asm. 1. In [23, Theorem 1] it has been proven that the same problem is NP-Hard for  $T = 1$  by reduction from the Set

Cover problem. Therefore, the problem in question is at least as hard as the combinatorial optimization problem of [23]. Therefore, we seek an approximate policy which yields an *approximate value function*  $V^{SOL}(t, s)$  such that  $V^{SOL}(t, s) \geq (1 - \gamma(t, s)) \cdot V^{OPT}(t, s)$  for some  $\gamma(t, s) \in (0, 1)$ .

For the problem with  $T = 1$  the work of [23] obtains a  $(1 - \beta_{\max})$ -approximation for the problem with a randomized rounding algorithm where  $\beta_{\max}$  is the maximum row sum of the relative liability matrix. Because the matrix  $A$  will be different in both cases, which would correspond to different clearing solutions. However, note that a rounding regime that iteratively rounds the fractional optimal solutions in a backward fashion will *not* yield correct results. The reason is that a suboptimal action at round  $t + 1$  can affect the optimal fractional action at round  $t$ .

---

**Algorithm 2** Randomized Rounding

---

**Sample-Interventions**( $L, B, u(t : T), s$ )

- (1) Until constraints are satisfied and the approximation guarantee is not violated
  - (a) For every agent  $i \in [n]$  sample (independent) interventions

$$z_{d,i}(t : T) \sim \text{Bin}\left(\frac{z_{r,i}^*(t : T)}{L_i}, L_i\right)$$

- (2) Return the value function  $V_{u(t:T)}^{SOL} = \sum_{t' \in [t, T]} \mathbf{1}^T \tilde{P}_d(t')$  given the calculated approximate (SOL) policy after calculating the clearing payments  $\tilde{P}_d(t : T)$ .

**Aggregate-Discrete**( $N, L, B, \mathcal{U}, s$ )

- (1) Sample  $N$  exogenous shocks  $\{u_j(t : T)\}_{j \in [N]} \sim \mathcal{U}$
  - (2) For  $j \in [N]$ 
    - (a) Call **Compute-Value-Function-Given-Sample-Path**( $L, B, u_j(t : T), s$ ) and get the optimal fractional policy  $z_r^*(t : T)$ .
    - (b) Calculate  $V_{u_j(t:T)}^{SOL} = \text{Sample-Interventions}(L, B, u_j(t : T), s)$
  - (3) Return  $\bar{V}^{SOL} = \frac{1}{N} \sum_{j=1}^N V_{u_j(t:T)}^{SOL}$
- 

The first idea on deriving an approximation algorithm for the problem described is adapting the approximation algorithm presented in [23] to the dynamic setting: More specifically let  $t$  be a

fixed round and  $s(t)$  be a fixed state at round  $t$  and let  $u(t : T)$  be a sample path for the random sequence  $U(t : T)$  from time  $t$  onward. Conditioned on the realization of  $u(t : T)$  the planner seeks to solve the following deterministic problem

$$\begin{aligned} \max_{z(t:T)} \quad & V_{u(t:T)}(t, s(t)) = \sum_{t'=t}^T R(s(t'), z(t') = \Pi(t', s(t')), u(t')) \\ \text{s.t.} \quad & s(t') = \Psi_{t'}(s(t'), z(t'), u(t')), z(t') \in \mathcal{Z}_d \quad \forall t' \in [t, T]. \end{aligned}$$

We also consider the fractional relaxation where the decision variables lie in  $\mathcal{Z}$ . Let  $V_{u(t:T)}^{SOI}(t, s(t))$  be the value that the approximation algorithm produces, let  $V_{u(t:T)}^{REL}(t, s(t))$  be the optimal solution of the relaxation (i.e., when the interventions belong to  $\mathcal{Z}$ ) and let  $V_{u(t:T)}^{OPT}(t, s(t))$  be the optimal solution. From optimality we know that  $V_{u(t:T)}^{REL}(t, s(t)) \geq V_{u(t:T)}^{OPT}(t, s(t))$ . Since  $u(t : T)$  is given, the optimal policy for the relaxed program consists of solving  $T - t + 1$  LPs sequentially and finding the pair of the clearing vector and optimal policy. We produce a rounded policy  $SOL$  that corresponds to an allocation vector  $z_d(t : T)$  randomly by rounding the matrix  $z_r^*(t : T)$  of the optimal relaxed policy by rounding all random variables  $z_{d,i}(t')$  as Binomial i.i.d. vectors on  $L_i$  trials with biases  $z_{r,i}^*(t')/L_i$  for all  $t' \in [t, T]$ .

**THEOREM 2.** *Alg. 2 yields the following approximation guarantee (on expectation):  $\mathbb{E}_{u(t:T), z_d(t:T)}[V^{SOI}(t, s(t))] \geq (1 - \gamma) \cdot \mathbb{E}_{u(t:T)}[V^{OPT}(t, s(t))]$ , where  $\gamma = \sup_{u(t:T)} \max_{t' \in [t, T], i \in [n]} \beta_i(t')$  is the worst-case financial connectivity of a deterministic contagion network produced in the environment space  $\mathcal{U}$ .*

Note that in the proof of Thm. 2, there is no dependence on the states created by the rounded outcomes. Therefore, we can compare and lower bound by the fractionally relaxed policy's reward value at each round  $t$ . Specifically, the proof of Thm. 2 lies in the observation that under a realization of the financial environment, a default node on the rounded solution can serve at least its external assets plus the intervention it gets, so, in expectation, it can serve at least its external assets plus the intervention of the fractional solution. For the optimal solution of the fractional relaxation, we know [18, 23] that for every  $S \subseteq [n]$ , the sum of the external assets and the fractional interventions over  $S$  is at least the sum of payments weighted by the “negated” financial connectivities. Moreover, for all the solvent nodes, the amount of payments they can serve is at least the corresponding fractional payment weighted by the node's connectivity. Combining both observations, we get that the value function of the rounded solution is at least a factor of 1 minus the worst financial connectivity across (the remaining) rounds of the optimal solution.

## 5 FAIRNESS IN INTERVENTION POLICIES

We say that an allocation is fair across the nodes if the intervention each node gets “does not differ a lot from its neighbors”, whereas the “neighborhood” of a node can be expressed in terms of the existing financial network or can be expressed in terms of an auxiliary network as in [23]. In this way, we can measure fairness in allocations in various settings. For instance, we can compare the interventions between a node and all other nodes in the network,

interventions between a node and its neighbors on the financial network, and interventions between nodes belonging to different population groups (such as minority groups). All fairness metrics have to be scale-invariant, namely, do not change when the budget provided changes from  $B$  to  $\alpha B$  for some  $\alpha > 0$ .

Driven by the above desiderata, we call an allocation rule  $Z(1 : T)$  in the model *fair* if the allocations of a node do not “differ much” from its neighbors. As a starting point, we consider the Gini Coefficient [17] and generalize it accordingly to our model. In detail, we measure the deviation between a node and its neighbors on a graph sequence  $\{H_t\}_{t \in [T]}$  with weights  $w_{ij}(t) \geq 0$  ( $w_{ii}(t) = 0$  for all  $i \in [n], t \in [T]$ ), with the following measure  $GC(t; H_t) = \frac{\sum_{(i,j) \in E(H_t)} w_{ij}(t) |Z_i(t) - Z_j(t)|}{\sum_{i \in [n]} Z_i(t) (\sum_{j \in [n]} (w_{ij}(t) + w_{ji}(t)))}$ . Note that this mea-

sure of inequality is well defined: when all nodes get the same allocations, it equals zero, and when one node gets all the allocations, it equals one. Examples of the sequence  $\{H_t\}_{t \in [T]}$  include: (i) Setting  $w_{ij}(t) = \mathbf{1}\{i \neq j\}$  yields the standard Gini Coefficient (GC). Such a measure does not capture the topology of the problem and aims to distribute interventions equitably.

(ii) Setting  $w_{ij}(t) = a_{ij}(t)$ , i.e., the weights between the nodes represent the actual relative liabilities between such pairs of nodes. This measure considers spatial interactions and the strength of ties (i.e., the relative importance of liabilities) to distribute the resources. We call this fairness constraint the Spatial Gini Coefficient (SGC). A similar fairness measure has been considered in [23], yet their measure is asymmetric.

(iii) If every node is associated with a sensitive attribute and  $q_i(t) \in [0, 1]$  is the probability of the node having this sensitive attribute, we can, for example, use  $w_{ij}(t) = |q_i(t) - q_j(t)| \cdot \mathbf{1}\{a_{ij}(t) > 0\}$  to put high weights on adjacent pairs which deviate in this attribute. For instance,  $q_i(t)$  can represent the probability of a node belonging to a minority group. Thus the weights would give importance to mitigating inequalities between neighboring groups of minorities and non-minorities. Moreover, if we want to enforce a stronger version of fairness, we can consider  $w_{ij}(t) = |q_i(t) - q_j(t)|$ , which penalizes all deviations in allocations between nodes with high deviations in their sensitive attribute. If more than one sensitive attributes are present, the weights can be modified accordingly to capture the average (or maximum) deviation of the attributes between a pair of nodes. We call this coefficient the Property Gini Coefficient (PGC).

This motivates the definition of a  $g(t)$ -fair allocation to be the allocation policy  $Z(1 : T)$  for which  $GC(t; H_t) \leq g(t)$  for all  $t \in [T]$  for some function  $g(1 : T) \in [0, 1]$  that does *not* depend on the clearing payments and the interventions. In our experiments we use  $g(t) = \text{constant}$ . This corresponds to the following additional linear constraints on the *action space* for an additional set of decision variables  $\omega(t) \in \mathbb{R}^{|E(H_t)|}$ ; for all  $t \in [T]$  we have

$$\begin{aligned} \omega_{ij}(t) &\geq 0 & \forall (i, j) \in E(H_t) \\ \omega_{ij}(t) &\leq |Z_i(t) - Z_j(t)| & \forall (i, j) \in E(H_t) \\ \sum_{(i,j) \in E(H_t)} w_{ij}(t) \omega_{ij}(t) &\leq g(t) \sum_{i,j \in [n]} (w_{ij}(t) + w_{ji}(t)) Z_i(t). \end{aligned}$$

After imposing the fairness constraints, a question one might ask is “How does the optimal value function without the fairness

constraints compares to the optimal value function with fairness constraints for any type of the (generalized) Gini Coefficient?”. For this reason, we define the Price of Fairness (PoF) to be  $\frac{\mathbb{E}[\text{OPT sans fairness}]}{\mathbb{E}[\text{OPT with fairness}]}$ . Since the fairness-constrained setting has additional constraints compared to the no-fairness setting, the price of fairness is always at least 1. In the static version, [23] proves that under discrete allocations there exist instances where the PoF can be unbounded, yet when the allocations are fractional the PoF is always bounded. Subsequently, in the dynamic setting it is easy to observe that the same result holds. Finally, in Sec. 6 we show how incorporating further fairness constraints to the problem affects the distribution of interventions in our datasets with respect to the nodes’ clearing payments and we also give quantitative results regarding the PoF.

## 6 EXPERIMENTS

**Datasets.** We run experiments with a variety of datasets, including synthetic networks, financial networks from various online settings, as well as an application to a non-financial networked resource allocation problem arising in ridesharing.

*Synthetic Liability Network (Stochastic blockmodel).* We generate a network of  $n = 50$  nodes and  $T = 10$  rounds whereas the structural graphs  $G_t$  are drawn i.i.d. from an SBM with Core-periphery structure with 2 blocks of size  $n_{\text{core}} = 10$  and  $n_{\text{periphery}} = 40$  and edge probability matrix  $\begin{pmatrix} 0.6 & 0.35 \\ 0.35 & 0.1 \end{pmatrix}$ . The internal liabilities and the external liabilities are drawn i.i.d. from  $\text{Exp}(1)$ , where the internal liability between  $(i, j)$  at round  $t$  is realized conditioned on the edge  $(i, j)$  existing on  $G_t$ , and the asset vector is set to 0 for every time. *Online Financial Network (Venmo transaction data).* We use publicly available data (<https://github.com/sa7mon/venmo-data>) from public Venmo transactions to form a dynamic transaction network. The dataset is split into three distinct time periods: (i) July 2018 to September 2018 (3.8M transactions) October 2018 (3.2M transactions) January 2019 to February 2019 (167K transactions). For our experiments we used the first period (July 2018 to September 2018) as it was the period with the most transactions. The amounts of the transactions are not provided in the data so we generate random transactions and use the provided topology (see App. B for more information). We have ignored data points for which the sender or the receiver of the transaction were missing.

*Physical Financial Network (Cellphone mobility data).* We extend the (static) version of the data processed from the SafeGraph platforms in [23] to the dynamic setting. Due to space constraints we have deferred the description and creation of the dataset to App. B.

*Non-financial Allocation Network (Extra dispatches in ridesharing).* To demonstrate the applicability of our framework beyond financial networks, our final experiment looks at the problem of creating extra dispatches (for example, using autonomous vehicles) in ridesharing networks to mitigate instantaneous demand-supply imbalances. We use publicly available trip data from the NYC Taxi and Limousine Commission (TLC). The TLC data are split in time periods, with each entry containing the start time of a ride, and its source and destination location IDs (corresponding to zones, e.g. Washington Heights, East Harlem etc.). We build a temporal network for Manhattan: nodes in the graph are rides between zones,

and the rest of the rides (to and from other boroughs) belong to the external network. The time period corresponds to January 2021 and the  $T = 31$  rounds correspond to different days. The edge weights (instantaneous liabilities) are determined by the number of rides requested from a zone to another (or to and from the external network). See App. B for more information.

**Interventions Experiments.** We run experiments in two settings: fractional and discrete interventions. For the fractional intervention setting, we report the payments, cumulative reward, and interventions for the datasets in question; since we solve the problem optimally we do not report any competing method. For the discrete intervention setting, our method is similar in spirit to the randomized-rounding algorithm of [23] for *static* allocations, which has been shown to perform very well in practice, and in fact, outperform existing heuristics even in more limited settings (in particular, binary action settings where the decision is to either to allocate or not). In addition, [23] also presents a greedy hill-climbing algorithm which under specific conditions gives an approximation guarantee. However, this algorithm is designed for the binary setting, and so it is unclear how one can extend the analysis to a dynamic setting and a more general allocation rule.

In Fig. 3 we plot the the average clearing payments, the average cumulative reward and the average *fractional* interventions for the Synthetic Core-periphery data and the Venmo data together with confidence intervals by averaging over 50 samples of the random networks. We plot the clearing payments of the 5 most *important* nodes (in terms of total payments).

Similarly, in Fig. 4 we plot the the average clearing payments, the average cumulative reward and the average *discrete* interventions for the TLC data and the SafeGraph data. The datasets here are deterministically created (see App. B).

**Price of Fairness Studies.** To induce fair allocations (in the fractional case) we run Alg. 1 with the constraints of Sec. 5. In Tab. 1, we run experiments with the fairness constraints where we constrain the Spatial GC and the Standard GC to be at most  $g(t) = 0.5$  at all times for different values of the budget  $B$  (per dataset). We report the PoF for the corresponding experiments in Tab. 1. For the datasets that involve randomness, we take the average over 50 independent runs.


Tab. 1 indicates that all fairness measures achieve a PoF that is very close to 1 in all datasets. This indicates that, generally our allocation algorithm can respect algorithmic fairness constraints with a very small cost on the total welfare.

**Insights from Experiments.** Due to space constraints, we give a chosen subset of data from our experiments (see full paper for more experiments).


In Fig. 5 we present the relationship between the fractional interventions and the clearing payments for the Synthetic Core-periphery dataset and the Venmo dataset. We plot the relationship

| Fairness Constraint<br>Budget $B$                    | Synthetic<br>50 | TLC<br>100 | Venmo<br>50 | SafeGraph<br>500K |
|------------------------------------------------------|-----------------|------------|-------------|-------------------|
| Spatial GC ( $w_{ij}(t) = a_{ij}(t)$ )               | 1.001           | 1.007      | 1.019       | 1.037             |
| Standard GC ( $w_{ij}(t) = \mathbf{1}\{i \neq j\}$ ) | 1.011           | 1.009      | 1.017       | 1.102             |


**Table 1: Price of Fairness. The payments and allocations can be found at Fig. 5. We set  $g(t) = 0.5$ .**



**Figure 3: Fractional interventions in financial networks based on the stochastic blockmodel (a-b) and Venmo data (c-d).**



**Figure 4: Figs (a-b): Interventions (extra dispatches) in ridesharing (January 2021 NYC data; we report 5 busiest neighborhoods). Figs (c-d): Discrete interventions in a financial network (based on SafeGraph Data, December 2020–April 2021).**



**Figure 5: Total payments of nodes vs. total interventions received (under financial networks based on the Stochastic Blockmodel and Venmo transaction data): We use  $L = B \cdot 1$ . SGC = Spatial Gini Coefficient, GC = (Standard) Gini Coefficient. Constraining the GC/SGC to be at most  $g(t) = 0.5$  the  $R^2$  drops since the interventions are distributed “more equally” between the nodes, i.e. “more central” nodes get a lower amount of bailouts when fairness is imposed).**

between the average total interventions and the network characteristics subject to fairness ( $g(t) = 0.5$ ) and no-fairness ( $g(t) = 1$ ) constraints. For the fairness constraints, we consider the Spatial Gini Coefficient (SGC) and the (Standard) Gini Coefficient (GC). We fit an ordinary least squares model to the points and a robust linear model via iteratively reweighted least squares with Huber’s T criterion. For the Venmo data, we observe a good positive correlation ( $R^2 = 0.61$ ) between the interventions and the payments. This is suggestive of the fact that nodes that are generally bailed out are central in the system, so bailing them out benefits the connections they have. When fairness constraints are applied,  $R^2$  decreases since

the fairness constraint balances interventions between adjacent nodes (in the case of SGC) and all nodes (in the case of the Standard GC) which may itself result in nodes that are less important to the contagion process getting high interventions. For the synthetic Core-periphery data, core and periphery nodes are well-separated on the plots’ “left-hand” and “right-hand” sides. We also observe that when we constrain the Standard GC to be at most 0.5 the  $R^2$  drops since the interventions are distributed “more equally” between core and periphery nodes, which results in more central (i.e. core) nodes receiving a lower amount of bailouts.

## REFERENCES

- [1] Dohyun Ahn and Kyoung-Kuk Kim. 2019. Optimal intervention under stress scenarios: A case of the Korean financial system. *Operations Research Letters* 47, 4 (2019), 257–263.
- [2] Siddhartha Banerjee, Daniel Freund, and Thodoris Lykouris. 2021. Pricing and optimization in shared vehicle systems: An approximation framework. *Operations Research* (2021).
- [3] Shane Barratt and Stephen Boyd. 2020. Multi-period liability clearing via convex optimal control. Available at SSRN 3604618 (2020).
- [4] Christian Borgs, Jennifer Chayes, Ayalvadi Ganesh, and Amin Saberi. 2010. How to distribute antidote to control epidemics. *Random Structures & Algorithms* 37, 2 (2010), 204–222.
- [5] Giuseppe Calafiore, Giulia Fracastoro, and Anton V Proskurnikov. 2022. Control of Dynamic Financial Networks. *IEEE Control Systems Letters* (2022).
- [6] Agostino Capponi and Peng-Chu Chen. 2015. Systemic risk mitigation in financial networks. *Journal of Economic Dynamics and Control* 58 (2015), 152–166.
- [7] Hong Chen, Tan Wang, and David D Yao. 2016. *Financial Network and Systemic Risk A Dynamic Model*. Technical Report. working paper.
- [8] Wei Chen, Chi Wang, and Yajun Wang. 2010. Scalable influence maximization for prevalent viral marketing in large-scale social networks. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*. 1029–1038.
- [9] Gabrielle Demange. 2018. Contagion in financial networks: a threat index. *Management Science* 64, 2 (2018), 955–970.
- [10] Kimon Drakopoulos, Asuman Ozdaglar, and John Tsitsiklis. 2014. An efficient curing policy for epidemics on graphs. In *53rd IEEE Conference on Decision and Control*. IEEE, 4447–4454.
- [11] Kimon Drakopoulos, Asuman Ozdaglar, and John N Tsitsiklis. 2017. When is a network epidemic hard to eliminate? *Mathematics of Operations Research* 42, 1 (2017), 1–14.
- [12] Beni Egressy and Roger Wattenhofer. 2021. Bailouts in financial networks. *arXiv preprint arXiv:2106.12315* (2021).
- [13] Larry Eisenberg and Thomas H Noe. 2001. Systemic risk in financial systems. *Management Science* 47, 2 (2001), 236–249.
- [14] Matthew Elliott, Benjamin Golub, and Matthew O Jackson. 2014. Financial networks and contagion. *American Economic Review* 104, 10 (2014), 3115–53.
- [15] Zachary Feinstein. 2019. Obligations with physical delivery in a multilayered financial network. *SIAM Journal on Financial Mathematics* 10, 4 (2019), 877–906.
- [16] Zachary Feinstein and Andreas Sojmark. 2019. A dynamic default contagion model: From Eisenberg-Noe to the mean field. *arXiv preprint arXiv:1912.08695* (2019).
- [17] Corrado Gini. 1921. Measurement of inequality of incomes. *The Economic Journal* 31, 121 (1921), 124–126.
- [18] Paul Glasserman and H Peyton Young. 2015. How likely is contagion in financial networks? *Journal of Banking & Finance* 50 (2015), 383–399.
- [19] Jessica Hoffmann, Matt Jordan, and Constantine Caramanis. 2020. Quarantines as a Targeted Immunization Strategy. *arXiv preprint arXiv:2008.08262* (2020).
- [20] Matthew O Jackson and Agathe Pernoud. 2020. Credit Freezes, Equilibrium Multiplicity, and Optimal Bailouts in Financial Networks. *SSRN* (2020).
- [21] Frank Kelly and Elena Yudovina. 2014. *Stochastic networks*. Vol. 2. Cambridge University Press.
- [22] Arian Klages-Mundt and Andreea Minca. 2022. Optimal intervention in economic networks using influence maximization methods. *European Journal of Operational Research* 300, 3 (2022), 1136–1148.
- [23] Marios Papachristou and Jon Kleinberg. 2022. Allocating Stimulus Checks in Times of Crisis. In *Proceedings of the Web Conference 2022*.
- [24] Pál András Papp and Roger Wattenhofer. 2021. Default ambiguity: finding the best solution to the clearing problem. In *International Conference on Web and Internet Economics*. Springer, 391–409.

## A PROOFS

**Proof of Thm. 1.** *Base Case*  $t' = T$ . Let  $f_T$  be the PDF of  $U(T)$ . We have that

$$\begin{aligned}
V(T, s) &= \max_{z \in \mathcal{Z}, \tilde{p} \in \mathcal{P}} \langle f_T, \mathbf{1}^T \tilde{p} \rangle \\
&= \max_{z \in \mathcal{Z}} \int_{\mathcal{U}} f_T(u) \max_{\tilde{p} \in \mathcal{P}} \mathbf{1}^T \tilde{p}(u) du \\
&= \max_{z \in \mathcal{Z}} \lim_{N \rightarrow \infty} \sum_{i=1}^N f(u_i) \Delta u_i \max_{\tilde{p}(u_i) \in \mathcal{P}} \mathbf{1}^T \tilde{p}(u_i) \\
&= \max_{z \in \mathcal{Z}} \lim_{N \rightarrow \infty} \max_{(p(u_1)^T, \dots, p(u_N)^T)^T \in \mathcal{P}(z^T, \dots, z^T)^T} \sum_{i \in [N], j \in [n]} f(u_i) \Delta u_i \tilde{p}_j(u_i) \\
&= \lim_{N \rightarrow \infty} \max_{(p(u_1)^T, \dots, p(u_N)^T)^T \in \mathcal{P}(z^T, \dots, z^T)^T} \sum_{i \in [N], j \in [n]} f(u_i) \Delta u_i \tilde{p}_j(u_i) \\
&= \int_{\mathcal{U}} f_T(u) \max_{z, \tilde{p}(u)} \mathbf{1}^T \tilde{p}(u) du \\
&= \mathbb{E}_{U(T)} [\max_{z, \tilde{p}} \mathbf{1}^T \tilde{p}].
\end{aligned}$$

The equalities follow from: (i) definition of expectation, (ii) pushing the maximization wrt  $s$  inside since  $z$  is fixed, (iii) definition of Riemannian integral, (iv) the  $N$  optimization problems being decoupled since  $z$  is fixed and thinking of the optimization as a large problem with a state vector of dimension  $2 \times N \times n$ , (v) pushing the optimization inside the limit (rewards are bounded regardless of the value of  $N$  and the corresponding mappings are continuous), (vi) definition of integration, and (vii) definition of the expected value.

*Inductive Hypothesis.* Assume that for  $t' = t + 1$  we have that  $V(t+1, s_t) = \mathbb{E}_{U(t+1:T)} [\max_{z_{t+1}, \tilde{p}_{t+1}} \{\mathbf{1}^T \tilde{p}_{t+1} + \max_{z_{t+2}, \tilde{p}_{t+2}} \{\mathbf{1}^T \tilde{p}_{t+2} + \dots\}\}]$ .

*Inductive Step.* For  $t' = t$  we have that

$$\begin{aligned}
V(t, s_t) &= \max_{z_t} \{r(s_t, z_t) + \mathbb{E}_{s_{t+1} \sim \mathcal{T}(s_t, z_t)} [V(t+1, s_{t+1})]\} \\
&= \max_{z_t} \mathbb{E}_{U(t)} \left[ \max_{\tilde{p}_t} \mathbf{1}^T \tilde{p}_t + \mathbb{E}_{U(t+1:T)} \left[ \max_{z_{t+1}, \tilde{p}_{t+1}} \{\mathbf{1}^T \tilde{p}_{t+1} + \max_{z_{t+2}, \tilde{p}_{t+2}} \{\mathbf{1}^T \tilde{p}_{t+2} + \dots\}\} \right] \right] \\
&= \max_{z_t} \mathbb{E}_{U(t:T)} \left[ \max_{\tilde{p}_t} \{\mathbf{1}^T \tilde{p}_t + \max_{z_{t+1}, \tilde{p}_{t+1}} \{\mathbf{1}^T \tilde{p}_{t+1} + \max_{z_{t+2}, \tilde{p}_{t+2}} \{\mathbf{1}^T \tilde{p}_{t+2} + \dots\}\}\} \right] \\
&= \mathbb{E}_{U(t:T)} \left[ \max_{z_t, \tilde{p}_t} \{\mathbf{1}^T \tilde{p}_t + \max_{z_{t+1}, \tilde{p}_{t+1}} \{\mathbf{1}^T \tilde{p}_{t+1} + \max_{z_{t+2}, \tilde{p}_{t+2}} \{\mathbf{1}^T \tilde{p}_{t+2} + \dots\}\}\} \right].
\end{aligned}$$

The equalities follow from: (i) the HJB equations, (ii) the inductive hypothesis, (iii) the fact that the maximization over  $\tilde{p}_t$  is independent of the sample paths from round  $t+1$  onwards and thus we can reorganize the expectations into one expectation over sample paths  $U(t:T) \sim \mathcal{U}$ , (iv) identically to the base case argument.

At any point, with probability 1, the value functions  $V_{u_i}(t:T)$  are between 0 and  $\sum_t \mathbf{1}^T(b(t) + \ell(t)) \leq (T - t + 1) \cdot \Delta$ , where  $\Delta = \sup_{\mathcal{U}} (\|b\|_1 + \|\ell\|_1)$  since the maximum reward can be achieved when all debts are paid and all nodes are solvent. Thus, by standard Chernoff bounds, one needs to choose  $N = \frac{\log(2/\delta)(T-t+1)^2 \Delta^2}{2\epsilon^2}$  samples to get an  $\epsilon$ -accurate estimation of the actual value function with probability at least  $1 - \delta$ .  $\square$

**Proof of Thm. 2.** *Approximation Guarantee.* Let  $D_d(t')$  and  $R_d(t')$  are the sets of default nodes and solvent nodes under discrete interventions at round  $t'$ . We have that

- (1) If  $i \in D_d(t')$  we have that  $\mathbb{E}_{z_d(t:T)} [\tilde{p}_{d,i}(t') | i \in D_d(t')] \geq c_i(t') + z_{r,i}^*(t')$ .
- (2) If  $i \in R_d(t')$  we have that  $\mathbb{E}_{z_d(t:T)} [\tilde{p}_{d,i}(t') | i \in R_d(t')] = \mathbb{E}_{z_d(t:T)} [p_{d,i}(t') | i \in R_d(t')] = \mathbb{E}_{z_d(t:T)} [\sum_{t'' < t'} (b_i(t'') + \ell_i(t'')) - \sum_{t'' < t'} \tilde{p}_{r,i}(t'')] \geq \mathbb{E}_{z_d(t:T)} [\sum_{t'' < t'} (b_i(t'') + \ell_i(t'')) - \sum_{t'' < t'} \tilde{p}_{r,i}^*(t'')] = \mathbb{E}_{z_d(t:T)} [p_{r,i}(t')] \geq \tilde{p}_{r,i}^*(t') \geq (1 - \max_{i \in [n]} \beta_{r,i}(t')) \tilde{p}_{r,i}^*(t')$ . The statement follows from: (i) definition of solvent node in the rounded solution, (ii) recursively using the definition of  $p_{d,i}(t'')$  for all  $1 \leq t'' < t'$ , (iii) point-wise optimality of the fractional clearing vector, (iv) definition of the solvency constraint for the fractional relaxation, (v) feasibility of the fractional solution, (vi)  $\max_{i \in [n]} \beta_{r,i}(t') > 0$  by Asm. 1.

Moreover, by [23] we know that for every subset  $S \subseteq [n]$  we have that the fractional solution satisfies

$$\left(1 - \max_{i \in S} \beta_{r,i}(t')\right) \sum_{i \in S} \tilde{p}_{r,i}^* \leq \sum_{j \in S} [c_i(t') + z_{r,i}^*(t')]. \quad (6)$$

By letting  $S = D_d(t')$  on Eq. 6 and since  $\max_{i \in D_d(t')} \beta_{r,i}(t') \leq \max_{i \in [n]} \beta_{r,i}(t')$  we have that

$$\left(1 - \max_{i \in [n]} \beta_{r,i}(t')\right) \sum_{i \in D_d(t')} \tilde{p}_{r,i}^* \leq \sum_{j \in D_d(t')} [c_i(t') + z_{r,i}^*(t')].$$

Moreover,  $p_{r,i}(t') \geq \tilde{p}_{r,i}^*(t') \geq (1 - \max_{i \in [n]} \beta_{r,i}(t')) \tilde{p}_{r,i}^*(t')$  where the second inequality is due to feasibility and the last inequality is because we multiply with a quantity that is strictly in  $(0, 1)$ . Therefore we have that the expected reward of the rounded solution at time  $t'$  is at least  $(1 - \max_{i \in [n]} \beta_{r,i}(t'))$  the optimal reward, i.e.  $\mathbb{E}_{z_d(t:T)} [R(s(t'), z(t') = \text{SOL})] \geq (1 - \max_{i \in [n]} \beta_{r,i}(t')) \cdot R(s(t'), z(t') = \text{OPT}) \quad \forall t' \in [t, T]$ .

We sum over  $t' \in [t, T]$  to conclude that  $\mathbb{E}_{z_d(t:T)} [V_{u(t:T)}^{\text{SOL}}(t, s(t))] = \sum_{t' \in [t, T]} \mathbb{E}_{z_d(t:T)} [R(s(t'), z(t') = \text{SOL})] \geq (1 - \max_{t' \in [t, T], i \in [n]} \beta_{r,i}(t')) \cdot V_{u(t:T)}^{\text{OPT}}(t, s(t))$ .

Taking expectation with respect to  $u(t:T) \sim U(t:T)$  and using Hölder's inequality we arrive to the desired result.

$$\begin{aligned}
& \mathbb{E}_{z_d(t:T), u(t:T)} [V^{SOL}(t, s(t))] \\
& \geq \mathbb{E}_{u(t:T)} \left[ \left( 1 - \max_{t' \in [t, T]} \max_{i \in [n]} \beta_i(t') \right) V_{u(t:T)}^{OPT}(t, s(t)) \right] \\
& = \mathbb{E}_{u(t:T)} [V_{u(t:T)}^{OPT}(t, s(t))] - \\
& \quad \mathbb{E}_{u(t:T)} [ |V_{u(t:T)}^{OPT}(t, s(t))| \max_{t' \in [t, T]} \max_{i \in [n]} \beta_i(t') ] \\
& \stackrel{\text{Hölder}}{\geq} \mathbb{E}_{u(t:T)} [V_{u(t:T)}^{OPT}(t, s(t))] - \\
& \quad \mathbb{E}_{u(t:T)} [\|V_{u(t:T)}^{OPT}(t, s(t))\|_1] \sup_{u(t:T)} \max_{t' \in [t, T]} \max_{i \in [n]} \beta_i(t') \\
& \stackrel{V \geq 0}{\equiv} \mathbb{E}_{u(t:T)} [V_{u(t:T)}^{OPT}(t, s(t))] - \\
& \quad \mathbb{E}_{u(t:T)} [V_{u(t:T)}^{OPT}(t, s(t))] \sup_{u(t:T)} \max_{t' \in [t, T]} \max_{i \in [n]} \beta_i(t') \\
& = \left( 1 - \sup_{u(t:T)} \left\{ \max_{t' \in [t, T]} \max_{i \in [n]} \beta_i(t') \right\} \right) \cdot \mathbb{E}_{u(t:T)} [V_{u(t:T)}^{OPT}(t, s(t))].
\end{aligned}$$

□

## B DATASETS

**Ridesharing.** We construct an instance of the dynamic clearing problem as follows:

- (1) We define the network as rides between locations at the same borough (this can be extended to include rides from different boroughs; but here we focus on one borough for clarity of exposition).
- (2) The data is split into non-overlapping frames that correspond to some duration. For exposition clarity we have used 1-day intervals. Again, here we can use smaller intervals (e.g. 5min) to represent demand for rides realistically.
- (3) We define  $\ell_{ij}(t)$  as the total number of rides from location  $i$  to location  $j$  at timestamp  $t$ .
- (4) We define  $b_i(t)$  to be the total number of external (outbound) rides requested from location  $i$  to outside of the borough.
- (5) We define  $c_i(t)$  to be the total number of internal (inbound) rides requested for location  $i$  from outside the borough.

**Venmo Transaction Data.** We construct the dynamic contagion instances as follows:

- (i) The timestamps are grouped in a weekly basis according to which year and week of the year they correspond to.
- (ii) For the whole dataset (corresponding to the July–September 2018 period), we create 2 sets:  $V_1$  and  $V_2$ .  $V_1$  corresponds to the top-100 nodes in terms of the number of incoming transactions, and  $V_2$  corresponds to the top-100 nodes in terms of the number of outgoing transactions. As the vertex set we use  $V = V_1 \cup V_2$ .
- (iii) We count the transactions between nodes of  $V$ , as well as the transactions from  $V$  to the outside system, and from  $V$  to the inside system for each round (week of the year 2018). For nodes with zero outgoing transactions we add one transaction.

(iv) We create random liabilities as follows

$$\begin{aligned}
\ell_{ij}(t) &= \mathbf{1}\{\text{\# of transactions } i \rightarrow j \text{ at round } t \text{ is } > 0\} \\
&\quad \times \text{Gamma}(\text{\# of transactions } i \rightarrow j \text{ at round } t, 1), \\
&\quad \forall i, j \in [n], t \in [T] \\
b_i(t) &= \max\{1, \text{Gamma}(\text{\# of transactions from } i \text{ to outside}, 1)\}, \\
&\quad \forall i \in [n], t \in [T] \\
c_i(t) &= \mathbf{1}\{\text{\# of transactions from outside to } i \text{ is } > 0\} \\
&\quad \times \text{Gamma}(\text{\# of transactions from outside to } i, 1), \\
&\quad \forall i \in [n], t \in [T].
\end{aligned}$$

Note that for  $b_i(t)$  we assert a *positive* value in order for Asm. 1 to hold.

**Cellphone Mobility Data.** We use mobility data from the SafeGraph platform spanning the period of December 2020 to April 2021. The static version of the data (i.e. with one round) has been introduced in [23]. Our paper extends the experiment to the multi-period setting, resulting in five financial networks each one of which for each month of the period of December 2020 to April 2021. SafeGraph offers insights on mobility patterns of people between Census Block Groups (CBGs) and Points of Interest (POIs), such as grocery stores, fitness centers, and religious establishments. We create a sequence of bipartite liability networks with monthly granularity between CBGs and POIs using the mobility data as a proxy to create liabilities. Roughly, given an initial pair of coordinates (latitude and longitude) we identify  $k$ -nearest neighboring CBGs and, subsequently, the POIs that these CBGs interact with. Each POI contains data from the visits of unique mobile devices from the corresponding CBGs, where we assume that each distinct device represents a unique person. SafeGraph also logs data about the dwelling time of devices, namely for how long each device stays in a POI, which we use to classify the visitors to two categories: customers and employees. For the employees, we add a financial liability that corresponds to the average monthly wage of such an employee for the corresponding POI as it is determined by its NAICS code. For the customers, we add edges from the corresponding CBGs to the POIs with value being the average consumption value for the specific POI category (given by the POI's NAICS code) as it is defined by the U.S. Economic Census. For each CBG node, we estimate the average size of households per CBG, the average income level and the percentage of people that belong to a minority group, which we use in order to estimate the assets and liabilities (internal and external) of CBGs. For the interventions of CBGs we use the US Cares Act rules to determine the interventions. Regarding the POI interventions we used data from loans provided in April 2020 as part of the SBA Paycheck Protection Program (PPP) and impute the interventions of the nodes that do not have intervention information. The processed data was *calibrated* similarly to [23] to reflect the specific interaction between POIs and CBGs across the different time spans.

The full paper contains extensive information on how to build the network, and determine  $c_i(t)$ ,  $b_i(t)$ ,  $\ell_{ij}(t)$  and  $L_i$  for each round  $t$ , which has been omitted due to space constraints.