UNIVERSITÀ DEGLI STUDI FIRENZE
DISEI
DIPARTIMENTO DI SCIENZE PER L'ECONOMIA E L'IMPRESA

Working Papers - Economics

Financial Contagion and Financial Lockdowns

Gabriele Camera and Alessandro Gioffré

Working Paper N. 01/2024

DISEI, Università degli Studi di Firenze
Via delle Pandette 9, 50127 Firenze (Italia) www.disei.unifi.it

Financial Contagion and Financial Lockdowns*

Gabriele Camera — Chapman University
Alessandro Gioffré — University of Florence

January 3, 2024

Abstract

Extreme financial shocks often elicit extraordinary policy interventions that preclude financial activity on a large scale, for example as the 1933 U.S. "bank holiday." We study these interventions using a random matching framework where the financial contagion process is explicit and the diffusion of the initial shock can be analytically characterized. The study suggests that there is scope for forced closures of individual firms or even economy-wide financial lockdowns only when firms are financially vulnerable and policy institutions are not well-functioning. Here, ordinary policy alone cannot prevent or sufficiently mitigate contagion, while complementing it with a lockdown or individual closures can do so, and improve social welfare if the initial shock is severe but not widespread.

Keywords: matching models, financial crises, contagion.

JEL codes: C6, D6, E5

1 Introduction

In March 1933 the U.S. financial system was in shambles. The financial crisis that gripped the country since 1929—creating thousand of bankruptcies—took a new ominous turn, when the gold reserves of the Federal Reserve Bank of New York fell below the legal 40% limit of their paper currency liabilities. This pushed the newly elected President F.D. Roosevelt to take an extraordinary

* We thank participants at the TEME 2022 conference in Honor of Jasmina Arifovic, an anonymous Guest Editor and two anonymous Referees for comments. Gabriele Camera, Economic Science Institute, Chapman University, One University dr., Orange, CA 92866; Tel.: 714-628-2806; e-mail: camera@chapman.edu. Alessandro Gioffré, DISEI, University of Florence; e-mail: alessandro.gioffre@unifi.it

step on March 6, 1933: ordering the suspension of all banking transactions until further notice. Presidential proclamation 2039 resulted in a complete "financial lockdown"—euphemistically referred to as a "bank holiday"—which lasted an entire week (Federal Reserve Bank of Boston, 1999).[1]

This extraordinary intervention came at the behest of the president of the Federal Reserve Bank of New York, G. L. Harrison, who reasoned that—given the increasing number of bank insolvencies and financial turmoil—suspending all bank activities "would permit the country to calm down and allow time for the enactment of remedial legislation" (Burns, 1974). This time-out would allow policy steps to be taken to reorganize the banks that, although financially distressed, had sound assets and a reasonable chance to return to solvency, a path forward that was codified in the Emergency Banking Act passed by Congress on March 13.[2]

This paper studies the impact of large-scale extraordinary interventions, such as the U.S. financial lockdown of March 1933, and smaller scale extraordinary interventions that only target firms directly affected by a shock. To study financial contagion we adapt the mathematical technique developed in Camera and Gioffré (2014) to model the diffusion of punishment in indefinitely repeated matching games. That technique forms also the basis of the study in Camera and Gioffré (2021) about the effects of non-medical "lockdown" interventions to contain the diffusion of viral infections. Here we bring to bear those techniques to illustrate how extraordinary policy interventions can be used to stem rapidly spreading financial contagion, and to study their consequences in terms of social welfare.

We use an analytical framework that makes explicit the process of contagion,

[1]The proclamation categorically stated: "No such banking institution or branch shall pay out, export, earmark, or permit the withdrawal or transfer in any manner or by any device whatsoever, of any gold or silver coin or bullion or currency or take any other action which might facilitate the hoarding thereof; nor shall any such banking institution or branch pay out deposits, make loans or discounts, deal in foreign exchange, transfer credits from the United States to any place abroad, or transact any other banking business whatsoever."

[2]These events are not uncommon even in more recent years. A database of financial crises from 1970 through 2012 compiled by the International Monetary Fund finds that in 8 out of 147 financial crises, governments prevented access to bank deposits or altogether ordered the suspension of all banking transactions (Argentina 1989 and 2001-02, Brazil 1990, Ecuador 1999, Panama 1988, Uruguay 2002, Chad 1983, Cyprus 2013); see Laeven and Valencia (2012). Beyond 2012, the Greek government closed banks for about two weeks in the summer of 2015.

and ties it to economic activity. In our model, the economy is populated by a finite number of players interpreted as financial firms, who gain from trading with each other. These firms are infinitely-lived — that is, they cannot go bankrupt — and trade in pairs that are randomly created at the beginning of each period, and dissolve at the end of each period. The economy is subject to an unanticipated financial shock which partitions the population into financially healthy and financially distressed firms, which have a lower payoff from trading than healthy firms. It is assumed that both kinds of firms have always an incentive to trade, and that the financial status of a firm cannot be observed by their counterpart. It is also assumed that the shock can propagate only in meetings between a financially distressed and a healthy firm. Hence, the process of financial contagion is random and operates on a population that is constant.

The diffusion of the initial shock depends on two economic fundamentals: the financial vulnerability of healthy firms, and the effectiveness of institutions tasked with developing ordinary policy interventions that enable economic recovery. Financial vulnerability is modeled as the probability of shock propagation in a meeting between a healthy and a distressed firm. The effectiveness of ordinary policy for economic recovery is modeled as the probability that all financially distressed firms simultaneously return to full financial health. We say that economic fundamentals are weak when firms are financially vulnerable (high probability of shock propagation) and ordinary policy is ineffective (low probability of return to healthy state). We emphasize that the structure of our model is designed to magnify the relevance of policy interventions directed at restoring financial health of distressed firms — relative to a model where firms can go bankrupt — because distressed firms keep operating in the market and hence present a persistent threat of contagion for the financial system. In this scenario, there may be scope for extraordinary policy interventions designed to contain or prevent contagion.

The analysis proceeds as follows. First, we characterize the contagion process following the initial shock, and show how ordinary policy intervention influences it. These policy responses affect the propagation mechanism only probabilistically.

Second, we show how contagion is affected when ordinary policy is complemented by some extraordinary intervention, which is deterministic and designed to immediately stop the progression of financial contagion. A first type of intervention is what we call a "financial lockdown," whereby all financial activity is temporarily suspended, as during a bank holiday. It is assumed that implementing this large-scale intervention can buy time to design a more effective ordinary policy response. Here, a tradeoff exists between contagion mitigation and temporary economic inactivity for all firms. We also study a second type of extraordinary intervention, which we call "individual suspension." Here, only firms that are financially healthy are allowed to trade, so this small-scale intervention completely stops contagion but forces distressed firms out of the market for a prolonged period of time (until they recover their financial health). Here, a tradeoff exists between contagion prevention and prolonged economic inactivity for a subset of firms. Based on this analysis, we calculate firms' expected payoffs and then construct a measure of social welfare that allows us to numerically assess the economic impact of extraordinary interventions.

Our findings can be summarized as follows. Financial contagion is a common phenomenon in economies with weak fundamentals because that weakness allows the shock to spread easily. In that case, even shocks that are limited to a small subset of the economy can rapidly spread. This suggests that there might be scope for extraordinary policy interventions that mitigate or completely prevent contagion primarily in economies where fundamentals are weak, i.e., where firms are financially vulnerable and ordinary policy response is not very effective. To uncover if such an extraordinary response is welfare-improving, we compare welfare under ordinary intervention to cases when a complementary extraordinary intervention is also implemented.

The numerical exercises suggest that extraordinary interventions can indeed be useful in economies with weak fundamentals, but only under certain conditions. These conditions depend on the size and the intensity of the initial shock. Preventing financial activity either on a large or a small scale improves welfare when the initial shock is deep and affects only a limited number of firms. When the shock affects only a limited number of firms, there are many firms that can benefit from extraordinary interventions that can mitigate or prevent contagion. Moreover, when the shock is deep, mitigating contagion implies mitigating large losses. The opposite holds when the initial shock is widespread or is not very significant, which is when the costs of an extraordinary intervention may not be offset by its benefits. Whether a large- or a small scale intervention is more suitable also depends on these two margins, intensity and scale of shock. Since a financial lockdown imposes large short run losses, a small-scale intervention is preferable if very few firms are affected by the shock. Here, many firms benefit while only very few suffer from prolonged inactivity. By contrast, a large-scale intervention is preferred as the scale of the shock grows because too many firms would suffer from prolonged inactivity, while too few would benefit from it.

This analysis contributes to the economics literature that studies financial contagion, an extensive review of which is outside the scope of this paper. Here we briefly identify the main research themes in this literature, and then explain how our work contributes to expand it; we refer the reader to Glasserman and Young (2016), for a more exhaustive review of financial contagion studies.

A main strand of literature studies how the specific links underlying some exogenous network structure affect financial contagion. Examples include Rochet and Tirole (1996), Allen and Gale (2000), Freixas et al. (2000), and Dasgupta (2004), which study the relationship between liability structure of banks and financial fragility. Specifically, Rochet and Tirole (1996) analyzes the impact of monitoring in interbank lending, and Allen and Gale (2000) shows that network connectivity enhances resilience to individual insolvencies. These models assume an unexpected exogenous liquidity shock occurs, as also assumed in our paper. By contrast, Freixas et al. (2000) studies network connectivity effects in a framework where contagion is driven by deposit flow, rather than an exogenous liquidity shock. Dasgupta (2004) uses a global game to study how linkages between banks, in the form of cross-holding of deposits, can become a source of financial contagion.

Other studies investigate the role of exchange frictions in specific network architectures. Gale and Kariv (2007), for example, study how the cost of intermediation affects efficiency in incomplete financial networks.

A second strand of literature adopts a reduced-form approach to model agents' interactions, to study the aggregate effect of their choices on financial networks. Examples include Eisenberg and Noe (2001), which studies a cascade of firms defaults triggered by individual payment shortfalls, by considering a clearing vector that defines the degree of connectivity among firms. This connectivity notion is generalized in Acemoglu et al. (2015), which studies the relationship between network structure optimality and shocks magnitude, finding that a densely connected network enhances financial stability only if shocks are moderate. Gai et al. (2011) considers an epidemiological-like model of financial contagion to assess the role of network concentration and complexity, on financial fragility. They provide numerical experiments to test the effectiveness of policy measures designed to mitigate the fragility of the financial system. Haldane and May (2011) and Elliot et al. (2014) examine nonmonotonic paths of financial contagion, focusing on the double role of connectivity—as a means of risk diversification and as a channel of contagion. They show that shock transmission dominates shock absorption if the degree of connectivity is low, while the opposite effect emerges if the degree of connectivity is high.

Given the above, our analysis offers a unique methodological contribution. Whereas the literature primarily focuses on studying financial contagion in models characterized by an implicit shock propagation process, we work with a model where the shock propagation process is made explicit. We leverage the theory of random matching to characterize the diffusion of financial contagion. This allows us to develop a framework where the transmission of the initial shock is made explicit via a pairwise random matching process.[3] This approach differs from the

[3]In particular, we do not derive expected contagion paths using Monte Carlo simulations — as often done in studies of random networks — but instead use an explicit transition matrix to calculate the expected path of financial contagion. In this manner, we can also calculate expected payoffs in closed-form for any contagion state. A benefit of having an explicit solution is a more precise and faster calculation for any initial financial shock. It also allows us to derive some partial results for the effect of various parameters changes on welfare.

reduced-form approach that is more standard in the literature, where, instead, the propagation of the initial shock depends on the (exogenous) topology of the financial network (e.g., as in Elliot et al., 2014). Seen this way, our study is closer to the strand of literature that focuses on the impact of the network structure on the transmission of shocks. Unlike those studies, we do not focus on the pros and cons of different network structures, but rather work with an environment in which all interactions are decentralized and in pairs. This introduces an explicit transmission component, which allows us to study the evolution of the initial shock, as it spreads from meeting to meeting, something that the existing literature does not consider. We think that such an approach can be helpful to organize one's thoughts about financial contagion phenomena.

In fact, although some of our results may seem intuitive, the proposed analytical approach is instrumental to develop insights that are less obvious. In particular, we show that there generally is no "one-size-fits-all" regulatory/policy response to a severe financial crisis. This finding is especially meaningful considering that developing countries tend to follow the example of the U.S. when it comes to financial regulatory or policy matters. Is this advisable? For example, consider how U.S. regulators dealt with the severe banking crisis of March 2023 ignited by Silicon Valley Bank: the bank was forced closed as soon as it suffered a bank-run and, after declaring it systemically important, regulators allowed it to re-open under new owners, four days later. Should a developing country follow this same plan if faced with a crisis of equivalent severity? Our analysis suggests that superficial policy mimicking might be ill-advised and policymakers should thoughtfully evaluate all of their available options.[4]

To put our contribution in context, it is also helpful to juxtapose the process of financial contagion studied in this paper, to the viral contagion studied in Camera and Gioffré (2021). In both studies, contagion entails the diffusion of an undesirable condition via random pairwise meetings. Hence, the main similarity is that contagion in society can only occur when two individuals directly interact.

[4]We thank an anonymous reviewer for raising this point and for suggesting this consideration about interventions in developing vs. developed countries.

Yet, there are also differences. The first lies in whether the recovery process of an economic agent is asynchronous and autonomous as opposed to synchronous and externally induced. In Camera and Gioffré (2021) an individual can recover their health autonomously, i.e., without some kind of policy intervention, while this is assumed impossible in the present paper where, instead, the process of recovering financial health must be initiated by an external authority, and is synchronous as if it happens, it simultaneously applies to all distressed firms. A second difference is that in the present paper we do not assume that the mere existence of unhealthy firms creates costs for society. In Camera and Gioffré (2021) unhealthy individuals induced societal costs but suffered no personal cost and remained as productive as everyone else. By contrast, in the present paper unhealthy firms do not induce societal cost but must bear a direct cost, from lost productivity. Consequently, financial and health contagion give rise to different trade-offs involving individual and social costs vs benefits. The two studies have also different implications for whether lockdowns can bring about a zero-prevalence state, i.e., a situation where everyone is healthy. Sufficiently prolonged health-related lockdowns can attain a zero-prevalence state without additional medical interventions thanks to the possibility of autonomous recovery – even if re-infection is possible. By contrast, prolonged financial-related lockdown spells cannot induce full recovery; the zero-prevalence state requires additional policy interventions, even if re-infection is impossible.

The paper proceeds as follows. The model is in Section 2. Section 3 analyzes the contagion processes, and Section 4 numerically studies how extreme interventions affect the evolution of the financial shock. Finally, Section 5 constructs payoff functions and numerically studies the welfare impact of the two kinds of extraordinary interventions. Section 6 offers a final discussion.

2 Modeling the economy

Time is discrete and infinite, denoted $t = 0, 1, 2, \ldots$. The economy is composed of a constant population of $N = 2n \geq 4$ infinitely-lived firms, which initially are all financially *healthy*. Firms trade in a decentralized market where they meet at random and in pairs in each period. Each meeting consists of a trading opportunity that generates a positive payoff to both counterparts, while the payoff is zero for a firm that does not trade. It follows that by design, economic activity in the model is beneficial, as it creates surplus, and all meetings automatically lead to a trade. Letting $y_{i,t} \geq 0$ denote the payoff to firm i on date t, the expected payoff at the start of the economy is $\mathbb{E} \sum_{t=0}^{\infty} \delta^t y_{i,t}$ where $\delta \in (0, 1)$ is a common discount factor.

A financial shock takes place in $t = 0$, which affects a proportion k/N of firms, transforming them into (financially) *distressed* firms, i.e., a firm being unable to fully meet its financial obligations, or being technically insolvent. The consequence of the shock is that while the payoff from trade is $y > 0$ for healthy firms, it falls to $\alpha y \in (0, y)$ for a distressed firm. No firm can ever go bankrupt although it can be technically insolvent. The financial shock can thus be characterized according to how many firms it affects (its size k) and how deeply it affects their profitability (its intensity $1 - \alpha$). Since there are no formal bankruptcies, the financial shock does not reduce the number of firms operating in the market.

Table 1: Financial shocks.

		Size k	
		Small	Large
Intensity $1 - \alpha$	Low	Negligible	Substantial
	High	Substantial	Extreme

An extreme shock is both intense and widespread, i.e., it induces a large drop in profitability for a large fraction of firms. Conversely, the shock is negligible. In-between cases correspond to substantial shocks; see Table 1.

Financial contagion happens because distressed firms can transmit the shock to healthy firms by trading with them. It is assumed that if a healthy and a distressed firm trade, the healthy firm becomes distressed with probability $p \in (0, 1]$. The transmissibility of the shock, p, can be interpreted as measuring the financial vulnerability of firms—an economic fundamental. To explain, the shock does not generally propagate after a single business dealing with a distressed firm but only if a healthy firm repeatedly engages with distressed firms: it takes an average of $1/p$ transactions for this to occur, so financial vulnerability increases in p.[5]

We make two additional assumptions about the propagation of the shock.

Assumption 1. *A healthy firm can become distressed independent of its history.*

Assumption 2. *A firm's financial health is unobservable by other firms.*

Recall that financially distressed firms earn a positive profit from trading and cannot be identified. It follows that healthy firms prefer trading and risk financial distress than staying out of the market. This will lead to financial contagion. Since being distressed is an absorbing state it follows that we need some policy intervention to mitigate or prevent contagion, and to restore the financial health of distressed firms. This will be discussed in a later section. Now, we study how the financial shock is transmitted through trading activity.

3 Financial Contagion: Characterization

At the start of each period a matching process partitions the population into pairwise meetings. This means that everyone, healthy and distressed, is assigned to a pair, for the period. Pairs are selected using a uniform random matching process, e.g., as in Diamond (1982). Therefore, given a population N, in each period t, the probability that firm i is assigned to meet any other firm is $\frac{1}{N-1}$. Meetings last just one period, meaning that rematching takes place in each period.

[5] One can think of the propagation mechanism as a delayed payment, not quite a default as that kind of problem would be publicly observable – something that we rule out in the model. A significantly delayed payment might create prolonged balance sheet problems to the creditor and an increase in insolvency risks. In particular, if the missed payment is large, this might lead to liquidity shortages that must be immediately addressed by selling assets with substantial haircuts, or by costly borrowing.

Figure 1: Timeline of events.

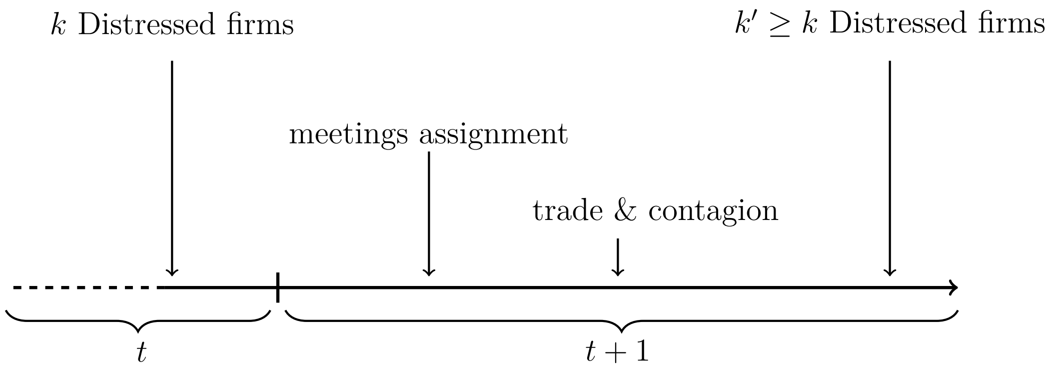

Fig. 1 displays the timeline of events in a period. First, a partition into pairwise meetings is proposed, involving all firms. Then firms meet. The financial shock can be transmitted only in meetings between healthy and distressed firms, i.e., in what we call a *mixed match*, which spreads the financial shock with probability p. As a result of this, at the end of $t+1$ we have $k' \geq k$ distressed firms, which is a random variable. This process repeats itself indefinitely and governs the stochastic progression of the financial shock in the economy.

3.1 How the shock propagates

In this section we describe how the shock propagates by utilizing the matrix-based tools developed in Camera and Gioffré (2014, 2021). Consider the start of a period. Let $k = 1, \ldots, N$ denote the number of distressed firms present in the market. We start by deriving the probability $Q_{kk'}(N)$ that $k' \geq k$ firms are distressed by the *end of the period*. Two sources of randomness affect this probability: the matching process, which determines how many mixed meetings occur between distressed and healthy firms, and the probability p determining if contagion occurs in these meetings.

Lemma 1. *We have:*

$$Q_{kk'}(N) = \sum_{\ell=k'-k}^{\min(k,N-k)} \lambda_{k\ell}(N) \binom{\ell}{k'-k} p^{k'-k}(1-p)^{\ell-(k'-k)}.$$

where

$$\lambda_{k\ell}(N) := \begin{cases} \dfrac{\ell!\binom{k}{\ell}\binom{N-k}{\ell}(k-\ell-1)!!(N-k-\ell-1)!!}{(N-1)!!} & \text{if } \ell \in L_k \\ 0 & \text{if } \ell \notin L_k, \end{cases}$$

and

$$\ell \in L_k := \begin{cases} \{0, 2, 4, \ldots, \min(k, N-k)\} & \text{if } k = \text{even}, \\ \{1, 3, 5, \ldots, \min(k, N-k)\} & \text{if } k = \text{odd}. \end{cases}$$

A formal extended proof of Lemma 1 can be found in Camera and Gioffré (2021, Section 3.1). Here, we provide a shorter proof, and some intuition for the result.

In Lemma 1, the function $\lambda_{k\ell}(N)$ is the probability that, if we have k distressed firms and $N-k$ healthy firms, then there can be $\ell = 0, 1, \ldots, \min(k, N-k)$ pairs composed of one distressed and one healthy firm. These mixed matches are the only meetings were contagion can occur in the model. The notation ℓ accounts for the possibility of having an odd or even number of distressed firms. We also note that we use the notation $n!! = n \cdot (n-2)!!$ where, $0!! = 1$ and $(-1)!! = 1$.

To see why $Q_{kk'}(N)$ takes the form shown above note that the financial shock is transmitted in a mixed meeting with probability $p < 1$. This has two implications. First, not all mixed meetings result in $k' - k \geq 0$ new distressed firms. Hence, we must have *at least* $k' - k$ mixed meetings if we want to have exactly $k' - k$ new distressed firms. This explains why we work with $\ell \geq k' - k$ in the probability $\lambda_{k\ell}(N)$. Second, to ensure that transmission of the shock occurs in *exactly* $k' - k$ of these ℓ meetings we have to account for the binomial probability $\binom{\ell}{k'-k}p^{k'-k}(1-p)^{\ell-(k'-k)}$.

Now note that $Q_{kk'}(N)$ only accounts for situations in which there is at least one distressed firm, i.e., $k' \geq k > 0$. The remaining cases have the obvious properties $Q_{00} = 1$, $Q_{0k'} = 0$ for all $k' > 0$, and $Q_{kk'} = 0$ for all $k > k' \geq 0$. It follows that the evolution of distressed firms during a period is described by the

$(N+1) \times (N+1)$ upper-triangular Markov matrix $\mathcal{Q}$:

$$\mathcal{Q} := \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & \dots & 0 & 0 & 0 \\ 0 & Q_{11} & Q_{12} & 0 & 0 & 0 & \dots & 0 & 0 & 0 \\ 0 & 0 & Q_{22} & Q_{23} & Q_{24} & 0 & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \dots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & 0 & 0 & \dots & Q_{N-2,N-2} & Q_{N-2,N-1} & Q_{N-2,N} \\ 0 & 0 & 0 & 0 & 0 & 0 & \dots & 0 & Q_{N-1,N-1} & Q_{N-1,N} \\ 0 & 0 & 0 & 0 & 0 & 0 & \dots & 0 & 0 & 1 \end{pmatrix}.$$

By means of example, if $N = 4, p = 1/2, k = 2$ we have $\lambda_{20} = 1/3, \lambda_{21} = 0, \lambda_{22} = 2/3$, and $Q_{22}(4) = \lambda_{20} + \lambda_{22}\binom{2}{0}(\frac{1}{2})^2 = \frac{1}{3} + \frac{1}{6}$, $Q_{23}(4) = \lambda_{22}\binom{2}{1}(\frac{1}{2})^2 = \frac{1}{3}$, $Q_{24}(4) = \lambda_{22}\binom{2}{2}(\frac{1}{2})^2 = \frac{1}{6}$, so $\sum_{k'=2}^{4} Q_{2k'}(4) = 1$.

Matrix $\mathcal{Q}$ contains the probabilities $Q_{kk'}(N)$ in row $k+1$ and column $k'+1$, for $k, k' = 0, \dots, N$. Because the number of distressed firms is a finite Markov chain, the square matrix $\mathcal{Q}$ is a Markov matrix. Hence, the transition probabilities taking place after $j \geq 1$ periods are determined by the product $\mathcal{Q}^j$. Matrix $\mathcal{Q}$ is a central element to calculate the evolution of financial contagion, but not the only one. We must incorporate policy interventions to have a complete picture.

3.2 Policy interventions

Financial distress is an absorbing state. Hence, in this section we enrich the model by discussing possible policymakers' responses, using a reduced-form approach. One kind of intervention is designed to restore firms' financial health: we call this an *ordinary policy intervention.* We differentiate it from another kind of policy response that, instead, is solely directed at mitigating or preventing further contagion, which we call *extraordinary policy intervention.*

Intuitively, these two kinds of interventions differ in their effect on the propagation mechanism of contagion and on the associated economic costs. Ordinary policy interventions, such as capital requirements or liquidity support provision, affect the propagation mechanism only probabilistically and, hence, prevent contagion only in expectation. By contrast, extraordinary interventions are deterministic and designed to immediately stop the progress of financial contagion by preventing the operation of distressed firms. These novel forms of intervention create economic costs that are absent from traditional policy responses. Hence, a tradeoff will emerge between the social costs from forced financial inactivity and the social benefit from stopping the progress of contagion relative to traditional policy responses.

Ordinary intervention. The ordinary intervention is a policy designed to attain what we call "full economic recovery." It is assumed this policy is put in place as soon as the shock occurs, and it operates probabilistically. In each period, there is a probability $h \in [0, 1)$ that all distressed firms reacquire their financial health. This probability is iid across periods and independent of trading activity. This reduced-form approach is meant to capture traditional policy interventions designed to restore overall financial health in the economy.[6] Once this policy is in place, full economic recovery is expected in $1/h$ periods. We thus interpret h as measuring the effectiveness of policy institutions' performance—our second economic fundamental. The underlying idea is that more effective (strong) policy institutions devise and implement policies that bring about full recovery faster than less effective (weak) institutions. Given h, we characterize the evolution of distressed firms using the Markov matrix $\mathcal{Q}$. Fixing the economy size N, we have:

Lemma 2. *Let there be $k = 0, 1, \ldots, N$ distressed firms when policy $h \geq 0$ is in place. The expected number of distressed firms after j periods of trade is*

$$\mu_k(j) = (1 - h)^j e_k^\top \mathcal{Q}^j \kappa.$$

The Lemma immediately follows from direct calculation. The properties of $\mu_k(j)$ are discussed in Camera and Gioffré (2014, Theorem 1). The probability that full economic recovery is achieved at any point during the first j periods is

$$h + (1 - h)h + (1 - h)^2 h + \ldots + (1 - h)^{j-1} h \equiv 1 - (1 - h)^j.$$

With the complementary probability $(1 - h)^j$, economic recovery does not take place, in which case we must account for the contagion process (the second part

[6] The probabilistic effectiveness of traditional regulatory interventions reflects the idea that the regulator might be unable to correctly assess the liquidity needs or insolvency risks of every firm, hence might miscalibrate the intervention or intervene with a delay (as during the recent financial crisis). In all these circumstances, some distressed firms might end up transmitting the shock to healthy counterparts, hence contagion may progress further.

of the expression). Matrix $\mathcal{Q}^j$ determines the possible evolution of the number of distressed firms after j periods of trade. We use $\kappa = (0, 1, \ldots, N)$ to denote a column vector that contains all possible numbers of distressed firms in the economy (including the 0 absorbing state). The vector e_k is an $(N+1)$-dimensional column vector with 1 in the $(1+k)^{th}$ position and 0 everywhere else. It is used to pin down the initial state of the economy in terms of the $k \geq 0$ distressed firms present. The transpose of this vector, $e_k^\top$, selects the $(1+k)^{th}$ row of matrix $\mathcal{Q}^j$, i.e., the transition probabilities associated with this initial k state. The non-zero elements of that row are the probabilities to transition from k to $k' = 0, 1, \ldots, \min(2k, N)$ distressed firms j periods from now. Hence, $e_k^\top \mathcal{Q}^j \kappa$ is the expected number of distressed firms after j periods of trade.

Ordinary policy cannot prevent contagion in a meeting. It affects the expected number of distressed firms because it probabilistically enables their financial recovery. Hence, we consider complementary policy responses designed to *contain* contagion, i.e., responses designed to reduce the diffusion of the shock.

A large-scale extraordinary intervention. The first response we consider is called "financial lockdown." It revolves around preventing all financial activities for $T \geq 1$ periods during which policymakers can develop tools designed to permanently reduce the transmissibility p of the shock and, hence, to mitigate contagion. In other words, a financial lockdown is a way to temporarily prevent financial contagion, by temporarily suspending all financial activities, in order to reduce the speed of financial contagion once financial activities restart. This modeling feature reflects the motivation for imposing "bank holidays" in the past. These interventions were justified as a way to prevent financial contagion in the very short run, while giving time to policymakers to put in place measures expected to mitigate or entirely stem contagion when markets reopened (e.g., Silber, 2009).

Given a financial lockdown of duration T, Lemma 2 can be slightly modified to calculate the expected evolution of distressed firms. First, matrix $\mathcal{Q}$ underlies a lower probability of contagion $p' < p$. Second, the contagion process operates only after T periods. Hence, given a population N, $k = 0, 1, \ldots, N$ currently distressed firms, an ordinary policy h, and a financial lockdown of $T < j$ periods, the expected number of distressed firms j periods in the future is

$$\mu_k(j,T) = (1-h)^j e_k^\top \mathcal{Q}^{j-T} \kappa.$$

Here we account for the fact that there can be full economic recovery in every period, including the first T periods, given ordinary policy h. Hence, if there is no recovery during any of the j periods, then the financial lockdown prevents further financial contagion for T periods, and reduces the spread of contagion in the next $j - T$ periods, when $\mathcal{Q}^{j-T}$ determines the evolution of distressed firms.

A small-scale extraordinary intervention. A second response we consider is called "individual suspension." Here, distressed firms are prevented from trading, hence the probability p of transmission in a meeting drops to zero. In other words, individual suspension is a way to completely prevent financial contagion.[7] To model individual suspension, consider that the random matching process partitions the population into meetings at the start of each period. This determines a *proposed* partition into trading pairs, independent of the firms' financial status. Individual suspension forces distressed firms to be inactive, meaning that they cannot join the proposed meeting. As a result, trade takes place only in pairs comprising two healthy firms, because healthy firms with a distressed counterpart are unable to trade. This is an analytically convenient way to maintain tractability, as we can work with a matching process that operates on a population of fixed size N. As a consequence, the volume of trade falls until all distressed firms recover their financial health, as a consequence of ordinary policy.

Hence, given a population N, $k = 0, 1, \ldots, N$ currently distressed firms, an

[7] Determining the financial health of a firm, and forcing it inactive can be complex and costly. We assume policymakers can do so costlessly, i.e., we ignore preventive costs of contagion. Note that in a financial lockdown, the entire market is shut down, and no firm is allowed to operate. Instead, under individual suspension, the distressed firm is placed under the control of the regulator but its operations do not completely cease—they are simply limited to what is necessary to resolve the firm's solvency problems, with the objective to restore investors' confidence and fully resume its operations. In this scenario, although the suspended firm does not operate normally, payments to creditors are generally allowed since the regulator's objective is to prevent contagion of other firms, not to create an illiquidity spiral.

ordinary policy h, and a complementary extraordinary intervention based on individual suspension, the expected number of distressed firms j periods in the future is $(1-h)^j k$. This immediately follows from Lemma 2, since there is no contagion in this case. The advantage of individual suspension, relative to a financial lockdown, is that it allows *some* healthy firms (depending on the realizations of the matching process) to keep trading without any contagion risks. The disadvantage, relative to a financial lockdown, is that distressed firms are expected to remain inactive for a prolonged time, $1/h$ periods.

4 Contagion: numerical experiments

In this section, we numerically study the expected trajectory of distressed firms following a financial shock. We first consider a case where there is only ordinary intervention, and then complement it with extraordinary interventions. In presenting our results, it is convenient to differentiate economies depending on the strength of their two economic fundamentals affecting the contagion process: the financial vulnerability of healthy firms (governed by p), and the performance of policy institutions (governed by h).

We say that the economy has *weak* fundamentals if firms are financially vulnerable (high p) and institutions are not well-functioning so ordinary policy is not very effective (low h). Here, financial shocks are easily transmissible and their effects long-lasting.[8]

[8] By well-functioning policy institutions, we mean institutions that can act quickly and competently to address economic problems. For concreteness, a well-functioning institution has a low level of corruption, takes policy actions with limited delays and with a clear plan, acts transparently, and is accountable to the public (e.g., see Guy, 2010). Financial vulnerability means greater risks of technical insolvency. These kinds of risks, may stem from more than one factor and in particular, an increased *credit risk*, i.e., the risk that the assets on the institution's balance sheet underperform or rapidly deteriorate in value as it happened during the Great Financial Crisis in 2008, an increased *liquidity risk*, i.e., the risk that the institution has not enough liquidity to service a large volume of withdrawals as during the recent collapse of Silicon Valley Bank, and an increase in *funding risks*, i.e., such as interest-rate risk for banks, when a rapid monetary policy tightening increases the cost of the institution's short-term liabilities above the revenue from its long-term assets. Technical insolvency can have negative cascading effects on the financial institutions that are more closely connected to the insolvent bank, which progressively spread to other institutions, and so on. This gives rise to a financial "contagion" process. Contagion may stem from a decline in investors' confidence in other financial institutions — e.g., as during the U.S. banking panics last century — or may trigger a deterioration in

Conversely, the economy has *strong* fundamentals (low p, high h). Intuitively, an economy is characterized by "weak fundamentals" when financial firms exhibit traits that contribute to create financial instability. These factors include, asset valuations that are historically very high, excessive leverage, funding risks associated with exceedingly high levels of short-term debt, market or funding liquidity problems, or an excessive risk appetite by financial institutions.

Table 2: Economic Fundamentals.

		Institutions' performance	
		$h = 0.01$	$h = 0.05$
Financial Vulnerability	$p = 0.01$	Moderate-a	Strong
	$p = 0.05$	Weak	Moderate-b

The two remaining situations capture in-between scenarios, where either firms are financially vulnerable but policy institutions are well-functioning, $(p, h) =$(high,high), or the reverse $(p, h) =$(low,low). Table 2 shows the parameterizations used in the numerical analysis, which is based on a population size $N = 1,000$.

4.1 The impact of ordinary policy

Fig. 2 reports the expected number of distressed firms over 200 periods.

---

the balance sheet of other institutions through a de-leveraging process — e.g., as it happened during the Great Financial Crisis. The model remains agnostic about the underlying causes of financial vulnerability because our objective is to study the mechanics of financial contagion for a given probability p of transmission of the shock, not the underlying causes of the process generating p.

Figure 2: Contagion and Recovery Under Ordinary Policy Interventions.

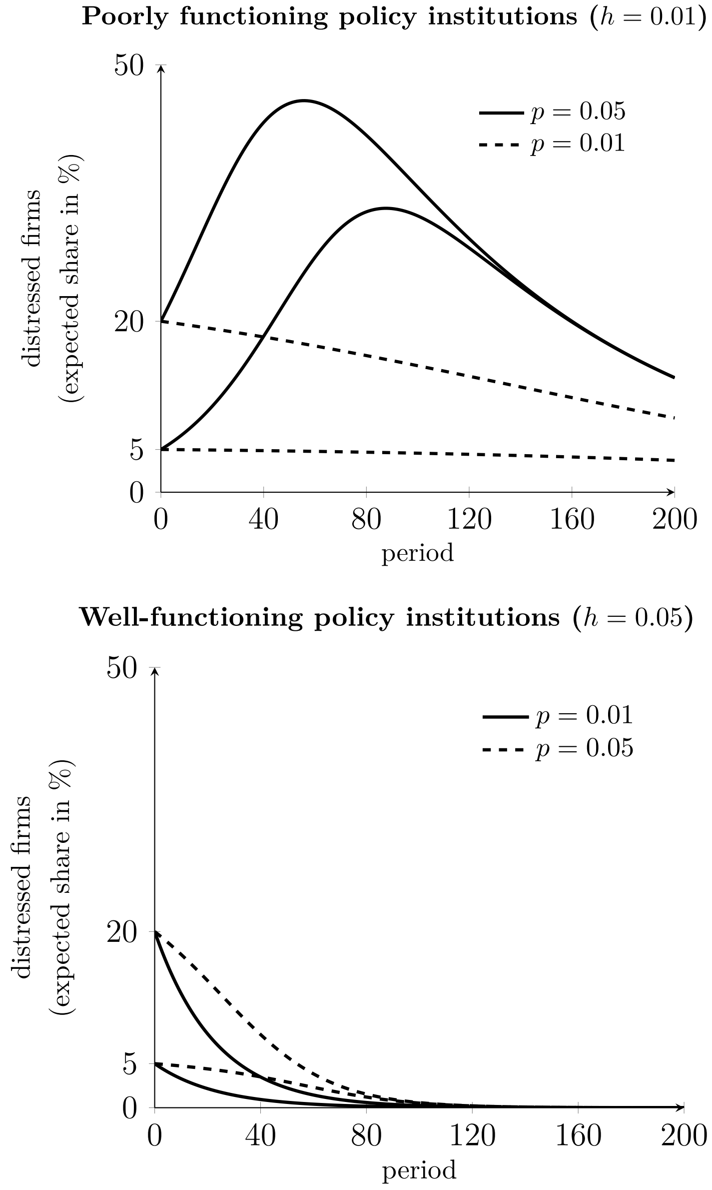

Notes: Population size $N = 1,000$ and transmission rates $p = 0.01, 0.05$. Top panel: $h = 0.05$. Bottom panel: $h = 0.01$. In each panel there are two initial financial shocks affecting, respectively, 5% and 20% of financial firms. The solid lines illustrate an economy with *weak fundamentals* in the top panel ($h = 0.01, p = 0.05$), and *strong fundamentals* in the bottom panel ($h = 0.05, p = 0.01$). The dashed lines identify two intermediate cases: (a) weak institutions but financially healthy firms in the top panel ($h = 0.01, p = 0.01$), and (b) strong institutions but financially vulnerable firms in the bottom panel ($h = 0.05, p = 0.05$).

Each panel considers two kinds of shocks, small and large shocks, affecting 5% and 20% of firms, respectively. The top panel considers $h = 0.01$ and the bottom $h = 0.05$, i.e., poorly vs. well-functioning policy institutions. If we interpret a period as a day, this means that the effective policy intervention $h = 0.05$ is expected to lead to full recovery in about 20 days(vs. 100 for $h = 0.01$). The lines trace the trajectory of the expected number of distressed firms, for the small and the large shock. The solid lines in the top panel show the trajectory when fundamentals are weak, while in the bottom panel they refer to the case of strong fundamentals. The dashed lines refer to the case of moderate fundamentals. There are two main observations.

Observation 1. *When policy institutions are well functioning, ordinary policy in expectation can contain the shock even if firms are financially vulnerable.*

To illustrate this point, consider the solid lines in the bottom panel in Fig. 2 (strong fundamentals). Here the expected number of distressed firms monotonically declines following the shock. The expected share of distressed firms is 1% after about 40 periods for a small shock, and about 70 periods for a large shock. In other words, contagion can be contained even in case of a large shock. This occurs because the transmission rate is low ($p = 0.01$) and the ordinary policy response is quite effective ($h = 0.05$). In this case, even a large shock is not expected to quickly spread to the rest of the market.

The large 20% shock can also be contained when firms are financially vulnerable; see the dashed lines ($p = 0.05$). This is because ordinary policy is effective and can foster a quick recovery in expectation. For $h = 0.05$, we need a much higher transmission rate for the expected number of distressed firms to increase above 20%, for example $p = 0.1$. The reason why firms' financial vulnerability does not matter as much in this case, is rooted in the way ordinary policy operates in our model. Greater financial vulnerability of firms means a greater probability p of transmitting the shock in a meeting, a process that operates locally at the firm level. By contrast, a more effective ordinary policy implies a greater probability h of economic recovery for every distressed firm, a process that operates globally at the financial market level. Hence, in this numerical experiment the extensive margin of policymaking dominates the intensive margin of financial health. This intuition is behind a second observation.

Observation 2. *When policy institutions are poorly functioning, then ordinary policy in expectation cannot contain the shock if firms are financially vulnerable.*

To illustrate this observation, consider the solid lines in the top panel in Fig. 2 (weak fundamentals). For each kind of shock, small and large, the initial shock cannot be contained, as it is expected to spread to a larger segment of firms, as compared to the initial shock. This is due to a transmission rate ($p = 0.05$) and an ineffective ordinary policy response ($h = 0.01$). A shock affecting only 5% of firms, is expected to spread less widely and rapidly than a large 20% shock. Following the small shock, the expected peak in distressed firms occurs after 90 periods and is about 33%, while the peak occurs earlier and is higher, for large shock (60 periods, 45% peak). Note also that each shock is persistent, and the long-run dynamics are independent of the shock's size; after 200 periods, the expected number of distressed firms is still above 5%. The dashed lines reveal that the shocks are expected to be contained when $p = 0.01$. However, as ordinary policy is not very effective, the recovery is very slow even if the shock is small. The 5% shock in the top panel is expected to drag out indefinitely, decaying very slowly, unlike what is reported in the bottom panel.

An insight from this analysis is that there is scope for extraordinary interventions when fundamentals are weak, i.e., when ordinary policy is not very effective and firms are financially vulnerable. In particular, the discussion above suggests that the strength of the institutions tasked with conducting ordinary policy is key because in the model the extensive margin of ordinary policy dominates intensive margin of vulnerability. With an effective ordinary intervention (high h) we do not see contagion dynamics, even if the initial shock is rather widespread and even if the shock can be easily transmitted. The natural question is thus to investigate whether extraordinary interventions can be a helpful complement to ordinary interventions, when these are ineffective at containing the initial shock.

4.2 Complementing ordinary with extraordinary interventions

Here we focus on economies with weak fundamentals, i.e., $p = 0.05$ and $h = 0.01$. Fig. 3 reports the expected number of distressed firms for a large shock affecting 20% of firms. The solid lines illustrate two types of financial lockdowns, short and long ($T = 5, 40$), each of which reduces firms' financial vulnerability by half, from $p = 0.05$ to $p = 0.025$ when markets reopen (square markers). The dashed line illustrates a policy of individual suspension for all firms affected by the initial shock. The dot-dashed line is for the case of no extraordinary intervention, as a comparison.

Figure 3: Contagion and Recovery Under Extraordinary Policy Interventions.

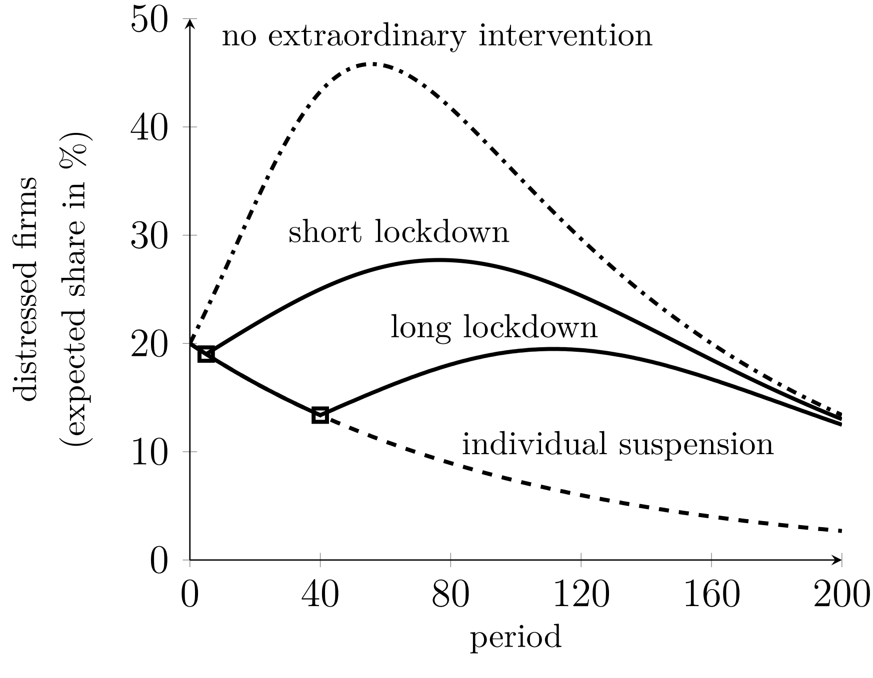

Notes: Population size $N = 1,000$, high initial shock (20%), high transmission rate ($p = 0.05$), and poor effectiveness of ordinary policy ($h = 0.01$). The solid lines illustrate an extraordinary intervention consisting of a financial lockdown of, respectively, 5 and 40 periods, after which the transmission rate is cut in half ($p = 0.025$). The dashed line illustrates an extraordinary intervention consisting of individual suspension of the firms affected by the initial shock.

A lockdown prevents contagion when markets are closed but does not necessarily do so when markets reopen; it depends on the effects of the lockdown on p.

Suspending individual firms, instead, completely prevents financial contagion by isolating problematic firms outside of the market. This reduces market activity for a prolonged period of time, until ordinary policy measures ensure full recovery has taken place. This means that, by design, extraordinary interventions cannot be a perfect substitute of ordinary intervention. Extraordinary interventions in the model are assumed incapable of addressing the factors underlying the shock, which instead can be addressed by ordinary policy intervention (the parameter h). In other words, the extraordinary interventions are modeled as being capable of only partially addressing one of the factors underlying the process of contagion: the firms' vulnerability to financial shocks (transmission rate p). The lockdown partially reduces p once markets reopen, while individual suspension completely reduces it to zero. We summarize the analysis as follows:

Observation 3. *Extraordinary interventions are a useful short-run complement to ordinary policy interventions.*

Consider the solid lines. Here, the lockdown mechanically prevents contagion while markets are closed. In fact, the expected number of distressed firms falls during the closure because ordinary policy intervention operates in the background. Longer periods of inactivity lead to greater declines in the expected number of distressed firms by the time markets reopen. However, the main impact of this extraordinary intervention hinges on the additional measures that policymakers can put in place to reduce the shock transmissibility p once markets reopen. This aspect of the intervention implies that suspending all financial activity for T periods has a positive impact on the expected number of distressed firms also after markets reopen because the intervention directly affects the speed of diffusion of the initial shock. As a result, financial contagion is reduced (shorter lockdown) or entirely prevented (longer lockdown) when markets reopen. This results in an accelerated economic recovery as compared to the no-extraordinary intervention scenario (dot-dash line). Notice, however, that the long-run trajectory of the expected number of distressed firms is similar for shorter and longer lockdowns because long-run dynamics are primarily entirely governed by ordinary policy intervention. It is in this sense, that this kind of large-scale extraordinary intervention can only complement but cannot substitute for ordinary policy intervention.

Now consider the dashed line. Here, financial contagion is prevented because all firms initially affected by the shock can no longer trade. In a way, this policy sets $p = 0$ thus ensuring a monotonic decline in the expected number of distressed firms. Once again, the speed of the decline is governed by the characteristics of ordinary policy intervention. Hence, even in this case this extraordinary intervention can only complement but cannot substitute ordinary policy.

Note that by design we model extraordinary interventions as being policy actions that are limited in scope. They are simply stop-gap measures introduced at the onset of an unexpected crisis, that are meant to prevent or slow down contagion, not to address the underlying crisis factors. In other words, extraordinary interventions cannot per se resolve the sudden crisis. A proper crisis response must necessarily include policy actions designed to address the specific factors underlying the shock (e.g., market illiquidity, regulatory shortcomings, agency problems, etc.). This is what we refer to as "ordinary intervention." In such a scenario, are extraordinary policy actions needed, and if so under what conditions and which one is more effective?

We can summarize the main difference between the two kinds of extraordinary interventions, small and large-scale, as follows. A financial lockdown can limit but not necessarily prevent contagion, and its impact is primarily observed in the short-run. The cost of this intervention is temporary trade inactivity for all firms, distressed and healthy. A famous example is the nationwide "banking holiday" of March 6, 1933, when U.S. President Roosevelt ordered by proclamation the suspension of banking activities. By contrast, individual suspension prevents financial contagion and, therefore, has a long-lasting impact. The cost of this intervention is trade inactivity only of distressed firms, until they can be restructured. As a recent example, on March 10, 2023, Silicon Valley Bank—among the top 20 U.S. banks— suffered a bank run. U.S. regulators immediately suspended its operations, took possession of it, and, after declaring it systemically important, they took the very unusual step of guaranteeing the entirety of its uninsured deposits, and on March 13 the bank was allowed to reopen under new ownership. The two interventions, lockdown and individual suspension, thus not only have different benefits but also different costs and, therefore, should have a differential impact on welfare, which we now investigate.

5 How extraordinary interventions affect welfare

In this section, we study the welfare effects of extraordinary interventions by constructing a measure of social welfare that corresponds to per-capita expected payoffs. We show the existence of a trade-off between the benefits of contagion mitigation and the costs from lost incomes due to trading inactivity. In doing so, we assume away any long-run adverse effects of trading inactivity (e.g., lower productivity, destruction of financial networks and so on).

5.1 Expected payoffs

Consider the expected payoff to a firm after the financial shock occurs. Let the ordinary policy intervention be characterized by h and $k = 0, 1, \dots, N$ denote the current number of distressed firms, with $N - k$ healthy firms.

We start by determining the boundary points of payoffs, i.e., for the aggregate state $k = 0, N$. For $k = 0$ we let

$$w_0 = \frac{y}{1-\delta},$$

denote the payoff for any firm in the economy. This is so, because all firms are healthy. Instead, when $k = N$ all firms are distressed. The payoff to a firm in this case is denoted w_N, which can be recursively written as

$$w_N = \alpha y + h\delta w_0 + (1-h)\delta w_N \qquad \Rightarrow \qquad w_N := \frac{\alpha y + h\delta w_0}{1-(1-h)\delta}.$$

The firm earns αy independent of the meeting. Instead, the continuation payoff depends on whether there is full recovery or not. In the first case (with probability h) the payoff is w_0, discounted by δ. Otherwise, the continuation payoff remains w_N since every firm is distressed.

Now consider $0 < k < N$. It should be clear that w_N is the payoff to a distressed firm, as their income is independent of the aggregate state and only depends on the policy parameter h. Instead, the payoff to a healthy firm depends on the aggregate state k. This is because a healthy firm may become distressed only in meeting with a distressed firm—an event that depends on k. We let the probability of a healthy firm meeting another healthy firm be

$$\sigma_k := \frac{N-k-1}{N-1}.$$

This probability is used to calculate w_k for a healthy firm, when $k = 1, \ldots, N-1$.

Start by noticing that for $k = N-1$ the expected payoff to a generic healthy firm can be recursively defined by

$$w_{N-1} = y + h\delta w_0 + (1-h)\delta[pw_N + (1-p)w_{N-1}].$$

To understand this expression, notice that a healthy firm earns y independent of the meeting. The future payoff instead, depends on whether there is full recovery or not. In the first case (with probability h) the payoff is w_0, which is then discounted by δ. Otherwise, the continuation payoff depends on whether the firm becomes distressed or not. Indeed, this firm meets a distressed firm with probability one. It becomes distressed with probability p, netting continuation payoff w_N, and otherwise the continuation payoff remains w_{N-1}.

For $k = 1, \ldots, N-2$, the expected payoff can also be recursively defined, but is a bit more complicated because we have to keep track of all possible random meetings leading to new distressed firms. We have

$$\begin{aligned} w_k =& y + h\delta w_0 + (1-h)\delta\Big\{\sigma_k \sum_{k'=k}^{N-2} Q_{k,k'}(N-2)w_{k'} \\ &+ (1-\sigma_k)[pw_N + (1-p)\sum_{k'=k-1}^{N-2} Q_{k-1,k'}(N-2)w_{k'+1}]\Big\}. \end{aligned}$$

The first two terms are the same as before. The other terms capture continuation payoffs when there is not full recovery. This continuation payoff depends on whether the firm meets another healthy firm, with probability σ_k. If it does, then it remains healthy (and so does its counterpart) and its continuation payoff is $w_{k'}$. Here, k' depends on the contagion resulting from the random meetings between the remaining $N-2$ firms: k of these are distressed and $N-k-2$ are healthy. This random process is governed by the probabilities $Q_{k,k'}(N-2)$, which are the elements of matrix $\mathcal{Q}$ for a population of $N-2$ firms. The population participating in the random matching is reduced by 2, because we already fixed a match of two firms.

Instead, if the healthy firm meets a distressed firm (with probability $1-\sigma_k$) there are two possibilities. It may become distressed, with probability p, so its continuation payoff is w_N, or may remain healthy, with probability $1-p$. In the latter its continuation payoff is $w_{k'+1}$, where k' depends on the new distressed firms resulting from the random meetings between the remaining $N-2$ firms, $k-1$ distressed and $N-k-1$ healthy. This random process is governed by the probabilities $Q_{k-1,k'}(N-2)$. Since we are fixing a match between a healthy and a distressed firm, we go to $k'+1$ distressed firms to account for the distressed party in the match that we fixed.

The vector of payoffs $\boldsymbol{w} := (w_0, \ldots, w_N)^\top$ can be represented in closed form. To do so, let $\mathcal{A}, \mathcal{B}, \mathcal{I}$ define $(N+1)\times(N+1)$ matrices. $\mathcal{I}$ is a standard identity matrix. $\mathcal{A}$ and $\mathcal{B}$ are probability matrices used to calculate expected continuation payoffs of a healthy firm. $\mathcal{A}$ refers to meetings with healthy firms, and $\mathcal{B}$ to meetings with distressed firms that do not cause contagion; the formal definitions are in the proof of Lemma 3 in Appendix A. We have the following.

Lemma 3. *Given w_N and w_0, and the $(N+1)\times(N+1)$ matrices $\mathcal{A}, \mathcal{B}, \mathcal{I}$ defined in Appendix A, we have:*

$$w_k = e_k^\top \{\mathcal{I} - (1-h)\delta[\mathcal{A} + (1-p)\mathcal{B}]\}^{-1}[(y + h\delta w_0)\mathbf{1} - e_N(1-\alpha)y + (1-h)p(\mathbf{1}-\boldsymbol{\sigma})w_N],$$

for $k = 0, 1, \ldots, N$.

Proof of Lemma 3. See Appendix A. □

Payoffs under a financial lockdown. Using the payoffs calculated above, we can define expected per capita payoffs when a financial lockdown is imposed for T periods. Let $k > 0$ be the number of distressed firms in a period that coincides with the start of the lockdown. The expected lifetime payoff to *any* healthy firm is defined by:

$$\tilde{w}_k = [1-(1-h)^T]\delta^T w_0 + (1-h)^T \delta^T w_k.$$

Since the market is closed for T periods, all payoffs must be discounted by δ^T. There is full economic recovery in those T periods with probability $[1-(1-h)^T]$ in which case every firm will earn w_0 when the market reopens (independent of its current state, healthy or distressed), in period $T+1$. Instead, if there is no economic recovery when markets reopen, then the firm's expected payoff is w_k.

Similarly, the expected lifetime payoff to *any* distressed firm is defined by:

$$\tilde{w}_N = [1-(1-h)^T]\delta^T w_0 + (1-h)^T \delta^T w_N, \qquad \text{for any } k \leq N.$$

The aggregate cost of a financial lockdown corresponds to the income lost during the T periods of the intervention. Since a healthy firm earns income y in a period, and a distressed firm αy, if there are $k > 0$ distressed firms, then lockdown destroys $yT(k\alpha + N - k)$ income.

Payoffs under individual suspension. When distressed firms are suspended from market activity, they earn zero income until they are readmitted to the market—when all firms recover their financial health. Their payoff is thus independent of k and given by

$$\hat{w}_N = \sum_{t=1}^{\infty} \delta^t (1-h)^{t-1} h w_0.$$

Instead, the expected payoff for a healthy firm depends on k and can be recursively defined as we did before. A difference is that the firm earns zero instead of y when it meets a distressed firm (which is inactive). Also, we no longer have financial contagion. Using the definition of $\sigma_k = \dfrac{N-k-1}{N-1}$, we have

$$\hat{w}_k = \sigma_k y + \delta[h w_0 + (1-h)\hat{w}_k].$$

In the aggregate, the cost of individual suspension is the income loss sustained in the economy until there is full recovery. Since a distressed firm earns income αy in a period, and the recovery is expected in $1/h$ periods, when there are $k > 0$ distressed firms, the intervention has an expected aggregate cost of $yk\alpha/h$.

Welfare. We measure social welfare using the *average payoff*, which for the case of no extraordinary intervention is defined as

$$W_k(0) = \frac{k}{N}w_N + \frac{N-k}{N}w_k, \qquad \text{(no extraordinary intervention)}$$

i.e., the weighted average of the payoff to distressed and healthy firms.

For the two complementary extraordinary interventions instead we have

$$W_k(T) = \frac{k}{N}\tilde{w}_N + \frac{N-k}{N}\tilde{w}_k, \qquad \text{(financial lockdown)}$$
$$W_k = \frac{k}{N}\hat{w}_N + \frac{N-k}{N}\hat{w}_k, \qquad \text{(individual suspension).}$$

We use these measures to numerically illustrate the trade-off facing a policy-maker in choosing whether or not to implement some extraordinary intervention. This tradeoff depends on the income lost due to trading inactivity, and the income earned thanks to the reduction in financial contagion.

To understand how the extraordinary interventions differentially affect welfare, recall that ordinary policy intervention operates in the long-run; it takes the form of a probabilistic and simultaneous recovery of financial health by all distressed firms. Extraordinary interventions operate in the short run, motivated by the desire to preclude or reduce financial contagion, when ordinary policy cannot achieve that goal. Now consider that healthy-distressed meetings generate income only if both firms are allowed to trade. It follows that a financial lockdown destroys *all the income* but only for a short time. When markets reopen, every meeting generates income, but the transmission of the financial shock p is reduced but not eliminated. That is to say, financially healthy firms are still exposed to the risk of becoming distressed. By contrast, individual suspension destroys *some income* for a prolonged period of time—the income of distressed firms and of their counterparts, healthy and distressed alike. Hence, not every meeting generates income but contagion completely stops as if $p = 0$.

It should be clear that the best-case scenario for a lockdown is $T = 1$. This is a situation when measures to reduce firms' financial vulnerability can be implemented with the minimum amount of time (one period) and, hence, the least possible income loss. On the other hand, the best-case scenario for individual suspension is $k = 1$, i.e., when the shock is so small that financial contagion can be completely prevented by forcing the minimum amount of firms to sit idle. These considerations inform the welfare analysis that follows.

Remark 1. *The expected benefit of each extraordinary intervention is hump-shaped in the size k of the shock.*

To understand why this is so, consider that limiting contagion is least beneficial when $k = 1$ (smallest size shock) as just one firm can be the vector of contagion. The benefit derived from lowering p or excluding distressed firms is directly linked to the number of mixed meetings (healthy-distressed meetings). Given random matching with uniform probability, the number of mixed meetings is hump-shaped in the size k of the shock, reaching a peak when half of firms are distressed and half are healthy. Notice that extraordinary interventions *cannot* increase welfare—as compared to no intervention—if the initial shock is sufficiently widespread. For $k = N$, there is no benefit because the shock cannot progress further. Consequently, the beneficial effect from an extraordinary intervention must be hump-shaped.

Intuitively, the benefits from limiting contagion are minimal when the shock affects a very small portion of firms; in this case, the probability of contagion is minimally impacted because very few firms can be vector of contagion. The benefits from limiting contagion are also minimal when the shock affects almost the entire economy, in which case the probability of contagion is also minimally impacted because few firms are healthy (and the costs from market shutdowns or exclusion are also large).

Now consider that extraordinary interventions have costs—the income loss associated with temporary suspension of all financial activity (lockdown) or market exclusion of distressed firms (individual suspension).

Remark 2. *The expected economic cost of each extraordinary intervention varies monotonically with the size k of the shock, but in opposite directions. It decreases for lockdowns, and increases for individual suspension.*

To see why, note that a financial lockdown creates income losses that are only temporary. This economic loss falls in the size k of the shock because as more firms are financially distressed, average productivity falls as well. By contrast, individual suspension creates protracted income losses (for $1/h$ periods), which however are suffered only by a subset of firms—those initially exposed to the shock and their random trade partners. This economic loss increases in size of the shock because as more firms are financially distressed, individual suspension gets closer to a financial lockdown that is protracted for $1/h$ periods.

These remarks suggest that the two extraordinary interventions have different welfare consequences depending on the size of the initial shock. They do not imply that extraordinary interventions *necessarily* improve welfare, which generally depends on the economic fundamentals and the severity of the financial shock. This is discussed by means of numerical illustrations, in the section that follows.

5.2 Numerical experiments on welfare

In this section, we report results of a numerical analysis about the welfare consequences of the two types of extraordinary policy interventions, following financial shocks of various sizes and intensities. Given Observations 1-2, we contrast economies with weak, moderate and strong fundamentals; see Table 2.

The numerical experiment involves the entire range of possible initial shocks, ranging from 1% to 100% of firms. We consider a shock of intensity $\alpha = 0.5$, a discount factor $\delta = 0.995$, and, for computational convenience, limit the population of firms to $N = 100$ in these experiments. We then expand the analysis to $\alpha < 0.5$ for the specific case of weak fundamentals. The welfare impact of complementing ordinary with some kind of extraordinary policy is determined by calculating the welfare gain or loss — in percentage — relative to welfare without any kind of extraordinary intervention (only ordinary policy). That is to say, the welfare impact corresponds to the ratio $\dfrac{W_k(T)}{W_k(0)} - 1$ for the case of a financial lockdown lasting T

periods, and $\frac{W_k}{W_k(0)} - 1$ for the case of individual suspension.

Figs. 4-5 illustrate the welfare impact of the extraordinary intervention on the vertical axis. The horizontal axis reports the size of the initial shock. A positive (negative) value identifies a welfare gain (loss) relative to the case of no extraordinary intervention. The results of our analysis can be organized into three main observations.

Observation 4. *There is scope for financial lockdowns in economies with weak fundamentals. The best-case scenario is a high-intensity shock that is not widespread.*

Fig. 4 and Table 3 provide evidence, for a lockdown lasting 5 periods.

Figure 4: Welfare Consequences of a Financial Lockdown.

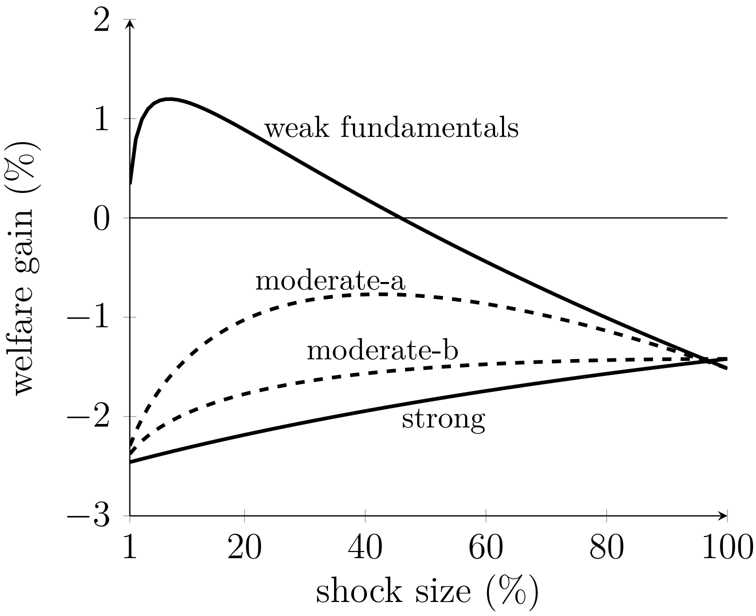

Notes: Horizontal axis: size of initial shock (share of firms initially distressed). Vertical axis: welfare gain relative to no extraordinary intervention. Population size $N = 100$, the lockdown lasts 5 periods and reduces p by half, the income from trading is 0.5 and 1 for, respectively, distressed and healthy firms. Each curve refers to one of the four possible combinations of economic fundamentals described in Table 2.

The lower solid line refers to the welfare impact when economic fundamentals are strong. Here, the lockdown is not welfare-improving even for shocks that affect large segments of the economy. In this case the shock is transmitted slowly (low p) and is not expected to have long-lasting consequences because ordinary policy is effective at bring about full economic recovery (high h). Hence, the probability $p(1-h)$ that a distressed firm transfers the shock to a healthy trade counterpart is already very small—about 1% since $p = 0.01$ and $h = 0.05$. Though the lockdown is assumed (in this numerical experiment) to cut the transmission rate p by half, the probability of contagion in a meeting falls only by about half of a percentage point. It follows that the benefit from limiting contagion, when markets reopen, is dominated by the temporary income loss associated with the 5-period lockdown. The welfare loss declines with the size of the shock because the number of low-productive firms increases, hence the economic cost of temporarily suspending all financial activity falls as well.

Instead, when fundamentals are weak we see a net benefit of suspending all financial activity (higher solid line). Here, contagion spreads rapidly (high p) and the consequences of the shock are expected to lasts longer (low h) so an extraordinary intervention that cuts the transmission rate in half induces a large reduction in the probability of contagion (from about 5% to about 2.5% for $p = 0.05$ and $h = 0.01$). The welfare gain rapidly increases as the shock grows in size, and as rapidly starts to slowly decline—eventually becoming a welfare loss for larger-size shocks. This non-monotonicity in welfare gains hinges on the monotonic decline in the cost of the lockdown with the size of the shock, and hump-shaped benefits from mitigating financial contagion. Such non-monotonicity in welfare is also exhibited when economic fundamentals are neither weak nor strong (dashed lines) although in these cases, the benefit of the intervention is always more than offset by its cost.

An insight is that complementing ordinary policy with a temporary suspension of all financial activity can be welfare-improving when ordinary policy is not very effective, but only if the initial shock is not too widespread, in which case there is little additional benefit or, in fact, a net loss from the additional intervention. Table 3 reinforces this message by showing how, under weak fundamentals, welfare varies with the shock *intensity*, i.e., with the trade profitability parameter α.

Table 3: Percentage welfare gains or losses (weak fundamentals)

	Intensity $1-\alpha$				
Size k/N	0.05	0.25	0.5	0.75	0.95
0.05	-2.14	-0.73	1.15	3.20	4.96
0.25	-2.19	-0.98	0.71	2.65	4.41
0.50	-2.27	-1.39	-0.13	1.35	2.75
0.75	-2.34	-1.74	-0.87	0.18	1.18
0.95	-2.38	-1.98	-1.39	-0.67	0.03

Notes: Population size $N = 100$, the economy has weak fundamentals $(p, h) = (0.05, 0.01)$, the lockdown lasts $T = 5$ periods and reduces p by half. The cells report the percentage welfare gain as we vary size and intensity of the shock, i.e., the proportion of firms initially affected by the shock (k/N), and the productivity decline $1 - \alpha$ experienced by distressed firms.

The cost of a financial lockdown is independent of p and h, while it increases in the lockdown's duration T so welfare will certainly fall if T is sufficiently large, for any kind of shock. In the Table $T = 5$. The cost also falls in the intensity and size of the shock ($1-\alpha$ and k), because the lost income from forced inactivity declines. Hence, in the Table the cost declines as we move down along the main diagonal. Note that the cells below the main diagonal have all negative values, while we see welfare gains in the right upper quadrant. Fixing a row, as we move right in the table, the intensity of the shock increases so the income loss associated with the temporary suspension of all financial activities falls, while the benefit from contagion-mitigation increases. This explains why the last column has only welfare gains: here, distressed firms have the largest productivity declines. Fixing a column, as we move down in the table, the size of the shock increases so both the income loss associated with the temporary suspension of all financial activities as well as the benefit from contagion-mitigation falls. This explains why welfare gains decline or eventually become losses as we move down a column.

This same intuition applies when we consider interventions that target only distressed firms. This kind of intervention can limit economic costs, and hence be preferable to wholesale suspension of all financial activities.

Observation 5. *If the initial financial shock is intense but limited to very few firms, then individual suspension is a welfare superior intervention relative to a financial lockdown.*

Fig. 5 illustrates the welfare consequences of individual suspension, compared to a financial lockdown under weak and strong fundamentals. As for large-scale extraordinary interventions, suspending only distressed firms from trade is not a good idea when fundamentals are strong (dashed line). Instead there is scope for this kind of intervention when fundamentals are weak. In this case we see that individual suspension may be more suitable than a financial lockdown.

Figure 5: Financial Lockdowns vs. Individual Suspension: Welfare Comparison

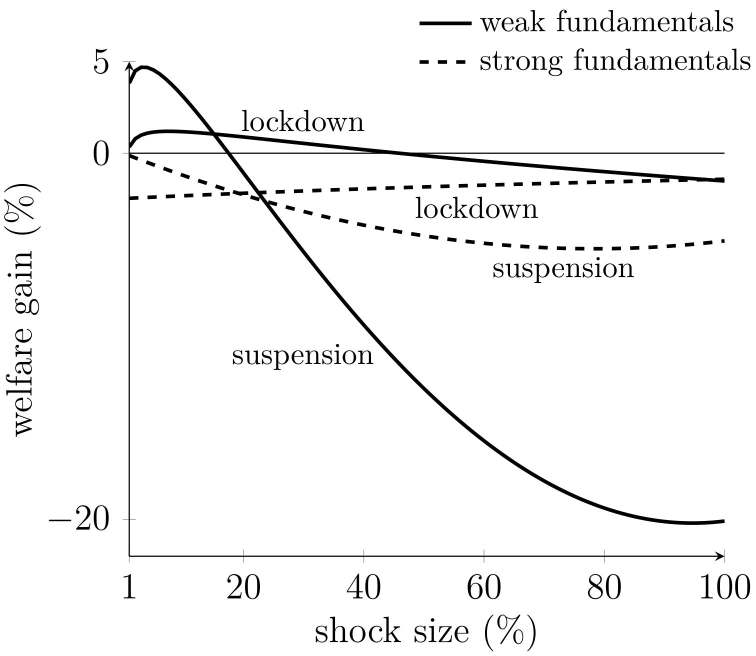

Notes: Horizontal axis: size of initial shock. Vertical axis: welfare gain relative to no extraordinary intervention. Population size $N = 100$, the economy has either weak or strong fundamentals $(p, h) = (0.05, 0.01)$ and $(p, h) = (0.01, 0.05)$, the lockdown lasts 5 periods and reduces p by half, the income from trading is $\alpha = 0.5$ and 1 for distressed and healthy firms.

Recall that the contagion-mitigation benefits of any extraordinary intervention are hump-shaped in the size k of the shock. However, financial lockdowns and individual suspension generate economic losses that move in opposite ways with k: they increases under individual suspension, while they fall under a financial lockdown. This makes individual suspension welfare superior (inferior) relative to the other intervention when the shock is sufficiently limited (widespread). When very few firms are financially distressed, not so much income is lost by excluding them from financial trading activities for a prolonged number of periods.

The results above assume that the policymaker response to a financial shock is based on a correct assessment of the economic fundamentals. In practice, policymakers might be uncertain about fundamentals when the shock hits.[9] In this case, the policymaker might mistakenly take no action when in fact some kind of extraordinary intervention would have been better. As the crisis progresses, the true state of fundamentals might be observed and, if so, there would be a sudden switch in policy; this, for instance, might explain the U.S. government lack of intervention at the beginning of the 2008 financial crisis.

To investigate this possibility, consider the parameters adopted for Fig. 5 and two specific initial shocks, $10\%, 30\%$. The policymaker is uncertain about the fundamentals, but observes the initial shock, and must simply decide if and how to intervene. Suppose that the policymaker attributes probability f that the economy has strong fundamentals, and $1 - f$ that they are weak. We can thus calculate the expected welfare gain (or loss) from implementing either of the two extraordinary policies considered before, and determine whether to implement either one of them, or none at all (hence, simply rely on ordinary policy). Here, the policymaker does not choose policies at random: she selects the one with the highest expected welfare gain. However policy mistakes can occur when the probability assigned to the true state of fundamentals is sufficiently small. This would lead to sudden policy shifts if new information can be learned as the crisis progresses. This additional analysis is summarized in the following:

Observation 6. *Uncertainty about economic fundamentals might cause a policy switch, whereby a financial crisis is characterized by initial policy inaction, followed by a switch to some extraordinary intervention.*

To develop intuition, consider Figure 6. The horizontal axis reports the policymaker's belief that the fundamentals are strong (the probability f). The vertical axis reports the expected welfare gain or loss for each extraordinary policy, based on that belief (dashed lines refer to individual suspension, solid lines to a financial lockdown).

[9]We thank an anonymous Referee for suggesting this additional analysis.

The left panel consider an initial small shock affecting 10% of firms. Here, there is a preferred policy, hence the policymaker should not apply a policy at random. If f is sufficiently small ($f < f_1$) then individual suspension is optimal since the solid line lies below the dashed line. Otherwise, no extraordinary policy should be implemented. In other words, in the case of a small shock, a financial lockdown is never optimal even under uncertainty about fundamentals. Intuitively, if fundamentals are weak, then the lockdown is dominated by individual suspension, while if fundamentals are strong then a lockdown creates a welfare loss; see Fig. 5. What kind of policy mistakes are possible here? Suppose that the true state of the economy is "weak fundamentals" but the policymaker beliefs are $f > f_1$. Here, the policy mistake is to do nothing; the two solid lines in Fig. 5 tell us that some extraordinary intervention would be preferred (suspension being the best). The opposite mistake occurs when $f < f_1$, if fundamentals are in fact strong. Here, the policy mistake is to implement individual suspension, instead of doing nothing; see the dashed lines in Fig. 5.

Figure 6: Policy choices under uncertainty about fundamentals

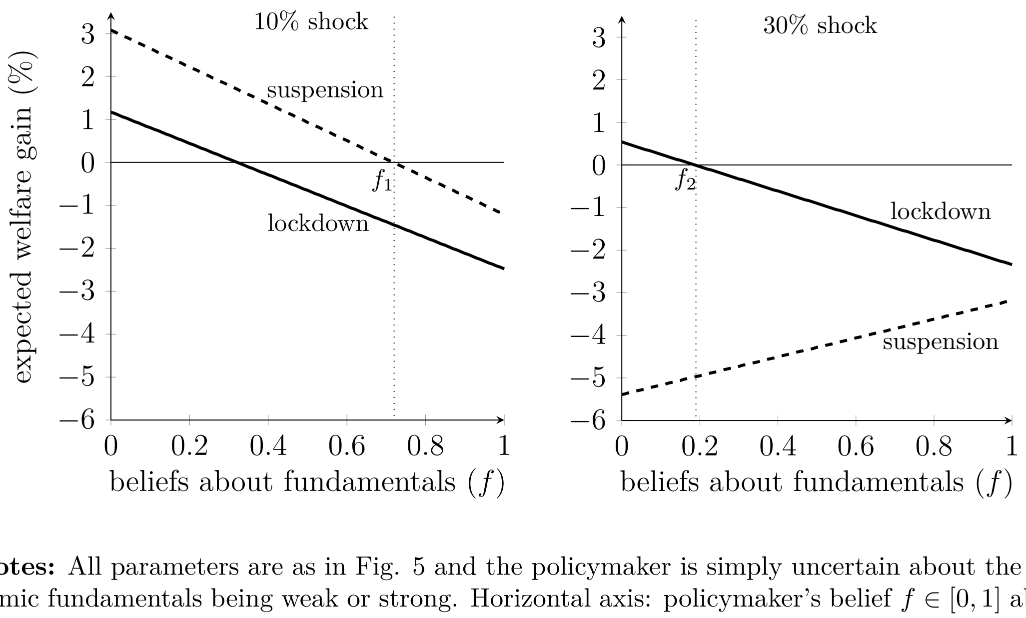

Notes: All parameters are as in Fig. 5 and the policymaker is simply uncertain about the economic fundamentals being weak or strong. Horizontal axis: policymaker's belief $f \in [0, 1]$ about fundamentals being strong. Vertical axis: expected welfare gain relative to no extraordinary intervention. Left panel: policymaker observes a 10% initial shock. Right panel: policymaker observes a 30% initial shock.

The right panel considers a large shock affecting 30% of firms. From Fig. 5, we know that individual suspension is never optimal — for weak and strong fundamentals — but a financial lockdown might be suitable if fundamentals are weak. It follows that if f is sufficiently small ($f < f_2$) then the uncertain policymaker will go for a financial lockdown (correctly so) and, otherwise, will take no extraordinary action. Again, here two types of policy mistakes can occur: a financial lockdown is imposed when in fact it should not (because fundamentals are strong), or there is a failure to act (because fundamentals are indeed weak). Supposing that the policymaker learns the true fundamentals over time, we might observe sudden policy reversals — a sudden financial lockdown for instance — which reminds us of policy actions taken during the great recession of 2008. In other words, it is conceivable that a rational policymaker would not intervene at the onset of a financial crisis and sharply reverse course as the crisis progresses.

Figure 7: Duration of Financial Lockdowns: Welfare Comparisons

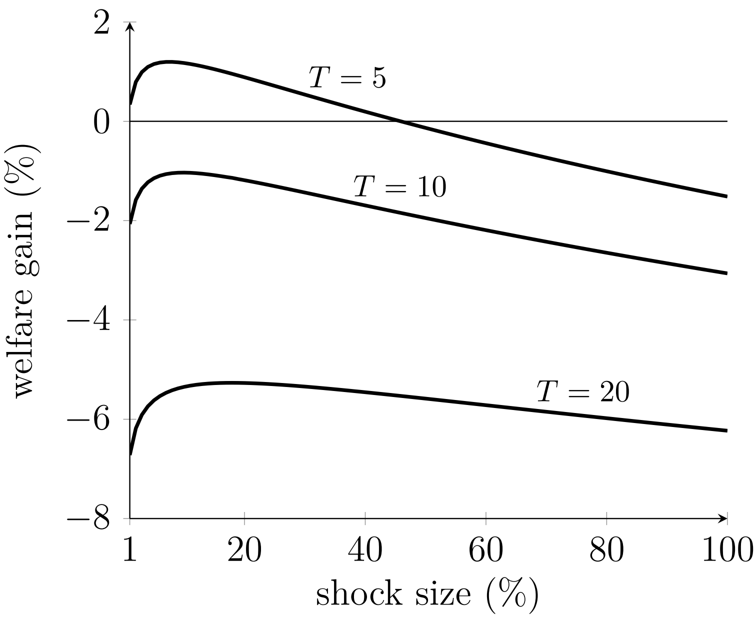

Notes: Horizontal axis: size of initial shock. Vertical axis: welfare gain relative to no lockdown. The population is $N = 100$, the economy has weak fundamentals $(p, h) = (0.05, 0.01)$, the lockdown reduces p by half and lasts alternatively $T = 5, 10, 20$ periods. The income from trading is $\alpha = 0.5$ and 1 for, respectively, distressed and healthy firms.

To conclude this section, we ran numerical experiments to show how welfare depends on the duration T of a lockdown. Fixing the same parameters used in Fig. 5, welfare falls as T increases because the cost of the lockdown is proportional to its duration while the benefit exerted through a reduction of the transmission probability p is not. This can be seen in Fig. 7 reporting welfare gains or losses for lockdowns of durations $T = 5, 10, 20$. Each lockdown cuts by half the probability of transmission p. Hence, when shocks are sufficiently small, in expectation we see a positive welfare impact of a lockdown only for $T = 5$, as the impact of a lockdown is negative as it gets longer, $T = 10, 20$.

6 Discussion

A main insight from this study is that complementing ordinary with extraordinary interventions to respond to a financial shock is *not* necessarily welfare-improving. It may be advisable when economic fundamentals are weak, i.e., when firms are financially vulnerable and the public institutions tasked with crafting and implementing ordinary policy are not very effective in their actions. In this case, there may be scope for *brief* suspensions of all or some financial activities — as in the case of bank holidays — primarily when the financial shock is intense but not widespread. In other words, an extraordinary intervention is not a suitable response to a systemic shock but it can be useful when the initial shock is severe and limited to a small segment of the financial market. Outside of this best-case scenario, the intervention is counterproductive because its contagion-reduction benefits would be more than offset by the economic costs it imposes on society. The analysis also suggests that smaller-scale interventions — which only constrain the activity of firms affected by the shock — are a preferable response.

Recent experience seems consistent with this finding. Consider the severe U.S. banking crisis of March 2023, when regulators confronted a sudden shock that involved only a handful of banks — Silicon Valley Bank and Signature Bank in particular. The U.S., a country with strong fundamentals, chose to implement an extraordinary policy broadly consistent with "individual suspension." The operations of the two distressed banks were suspended soon after the beginning of alarming bank runs. Afraid that these bank runs could trigger a global financial contagion, regulators immediately closed the banks and took possession of them (via newly created banks), at the same time taking steps to facilitate their acquisition from healthy banks. This averted contagion, restored depositors' confidence, and the banks fully reopened under new ownership in a matter of days.[10]

The mechanism behind financial lockdowns in this paper differs from the often-heard behavioral interpretation of the 1933 U.S. national "bank holiday," which somehow allowed depositors — interpreted as players engaged in the coordination game discussed in Diamond and Dybvig (1983) — to tacitly re-coordinate on the good equilibrium and away from the bank run equilibrium. Albeit intriguing, this explanation does not seem consistent with empirical evidence. On the one hand, the local bank holidays of 1932 all failed to achieve such a re-coordination despite the smaller number of depositors involved in those "games" (for a discussion of local bank holidays during previous banking panics see Engemann, 2013). One would imagine that pausing a larger coordination game (via a national bank holiday) should not be more likely to allow successful re-coordination as compared to a smaller scale game (via regional bank holidays). Typically miscoordination increases in the number of players, not the reverse, as strategic uncertainty is more significant. We are also not aware of experimental evidence suggesting that pausing a coordination game where players have coordinated on the bad equilibrium is a sure recipe to foster re-coordination on the good equilibrium. The idea behind our paper follows a different logic. We take the view that the 1933 bank holiday can be interpreted as a technical pause that allowed policymakers to buy time to achieve three goals: (1) better understanding the underlying causes of the banking crisis, (2) designing a suitable policy response, and (3) allowing an

[10]Silicon Valley Bank was the second largest bank failure in the history of the U.S. after Lehman Bros. in 2008. It had the potential to trigger a full blown global financial crisis. On March 8, the bank announced a large loss on the sale of securities and the day after large depositors began to quickly withdraw their funds. On March 10 regulators closed the bank and spun it over to an FDIC institution, newly created. Simultaneously, the Fed created an emergency lending program to absorb the risk of the bank's distressed securities. The bank reopened for business on March 13, under new ownership. Signature bank suffered a bank run on March 10, when its depositors were spooked by SVB's collapse on the previous day. Regulators closed it on March 12, transferring its assets and liabilities also to an FDIC-created bank, and its branches reopened for business on March 20 under new ownership.

effective implementation of this response. This seems indeed consistent with the events in 1933, when the national bank holiday allowed Congress to convene an extraordinary session to implement very drastic emergency legislation.[11]

We believe these considerations can offer useful insights for policymakers. At the same time, we emphasize the exploratory nature of this work and its obvious limitations. For instance, we adopt a reduced form approach to modeling policy, do not consider an explicit process describing how financial firms generate value, and abstract from general equilibrium considerations. As such, the observations that emerge from the numerical experiments offered in this paper should not be taken as general statements directly applicable to field economies, but suggestive of possible correlations between policy actions and outcomes. Indeed, some fruitful extensions of the model suggests themselves. For instance, the model could easily include considerations about the resources needed to build and maintain the infrastructure necessary to monitor and manage the financial health of system participants. It could also be extended to account for the possibility that extraordinary interventions might destroy segments of the financial network, or simply create protracted stigma or confidence problems that would not otherwise occur.

References

Acemoglu, D., Ozdaglar, A., and Tahbaz-Salehi, A. 2015. Systemic risk and stability in financial networks. *American Economic Review* 2, 564-608.

Allen, F., and Gale, D. 2000. Financial contagion. *Journal of political economy* 108.1, 1-33.

Burns, Helen M. 1974. *The American Banking Community and the New Deal Banking Reform, 1933-1935.* Westport: Greenwood Press.

[11] The article in Federal Reserve Bank of Boston (1999, p. 19) notes that "During the holiday, federal officials worked day and night on the emergency legislation." That work was not about fostering coordination of depositors, but rather to deliver the emergency banking act which "gave the President "tremendous World War powers of regulation over transactions in credit, currency, gold and silver, including foreign exchange." It also empowered the Secretary of the Treasury, "to require delivery at the Treasury of all gold and gold certificates held by anybody in the country." We thank an anonymous Referee for suggesting we discuss the differences between the often-heard behavioral interpretation for the 1933 banking holiday, and ours.

Camera, G., Gioffré, A. 2014. A tractable analysis of contagious equilibria. *Journal of Mathematical Economics* 50, 290-300.

Camera, G., Gioffré, A., 2021. The economic impact of lockdowns: A theoretical assessment. *Journal of Mathematical Economics* 97.

Dasgupta, A. 2004. Financial contagion through capital connections: A model of the origin and spread of bank panics. *Journal of the European Economic Association* 2(6), 1049-1084.

Diamond, Peter. 1982. Aggregate Demand Management in Search Equilibrium. *Journal of Political Economy* 90, 881-94.

Diamond, Douglas W. and Philip H. Dybvig. 1983. Bank Runs, Deposit Insurance, and Liquidity. *Journal of Political Economy* 91(3), 349-527.

Eisenberg, L., and Noe, T. H. 2001. Systemic risk in financial systems. *Management Science* 47(2), 236-249.

Elliott, M., Golub, B., and Jackson, M. O. 2014. Financial networks and contagion. *American Economic Review*, 104(10), 3115-53.

Engemann, Kristie M. 1999. *Banking Panics of 1931-33*. Federal Reserve Bank of St. Louis, Federal Reserve History. Accessed at https://www.federalreservehistory.org/essays/banking-panics-1931-33 on July 1, 2023.

Federal Reserve Bank of Boston. 1999. *Closed for the Holiday: The Bank Holiday of 1933*. Accessed at http://www.bos.frb.org/about/pubs/closed.pdf on March 8, 2022.

Freixas, X., Parigi, B. M., and Rochet, J. C. 2000. Systemic risk, interbank relations, and liquidity provision by the central bank. *Journal of money, credit and banking* 611-638.

Gai, P., Haldane, A., and Kapadia, S. 2011. Complexity, concentration and contagion. *Journal of Monetary Economics* 58(5), 453-470.

Gale, D. M. and Kariv, S. 2007. Financial networks. *American Economic Review* 97(2), 99-103.

Glasserman, P. and Young, H.P. 2016. Contagion in financial networks. *Journal of Economic Literature*, 54(3), pp.779-831.

Peters, B. Guy. 2010. *Institutional Design and Good Governance* In *The Good Cause. Theoretical Perspectives on Corruption.*, p. 83-97. Editor(s): Gjalt de Graaf, Patrick von Maravic and Pieter Wagenaar. Publisher: Verlag Barbara Budrich.

Haldane, A. G. and May, R. M., 2011. Systemic risk in banking ecosystems. *Nature* 469(7330), 351-355.

Laeven, L. and F. Valencia. 2012. Systemic Banking Crises Database: An Update. IMF Working Paper # 12/163.

Mortensen, D., and C. Pissarides. 1994. Job Creation and Job Destruction in the Theory of Unemployment. *Review of Economic Studies*, 61(3), 397-415.

Rochet, J. C. and Tirole, J. 1996. Interbank lending and systemic risk. *Journal of Money, credit and Banking* 28(4), 733-762.

Silber, William L. (2009) Why Did FDR's Bank Holiday Succeed? FRBNY *Economic Policy Review*, July, 19-30.

A Appendix

B Proof of Lemma 3

To express w_k for $k = 0, \ldots, N$ in vectorial form, we consider matrix $\mathcal{Q}$ when the population is $N-2$, with a slight abuse in notation. This matrix has $N-1$ rows and columns (corresponding to $k=0$ all the way to $k=N-2$). Let its elements $Q_{kk'}(N-2)$ be denoted $\tilde{Q}_{kk'} \equiv Q_{kk'}(N-2)$, for convenience.

Now construct two $(N+1)\times(N+1)$ matrices, $\mathcal{A}$ and $\mathcal{B}$.

Matrix $\mathcal{A}$ accounts for the continuation payoff components w_k that emerge if the firm is healthy and meets another healthy firm. We have:

$$\mathcal{A} := \begin{pmatrix} \sigma_0 & 0 & 0 & 0 & 0 & \cdots & 0 & 0 & 0 & 0 \\ 0 & \sigma_1\tilde{Q}_{11} & \sigma_1\tilde{Q}_{12} & 0 & 0 & \cdots & 0 & 0 & 0 & 0 \\ 0 & 0 & \sigma_2\tilde{Q}_{22} & \sigma_2\tilde{Q}_{23} & \sigma_2\tilde{Q}_{24} & \cdots & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \sigma_3\tilde{Q}_{33} & \sigma_3\tilde{Q}_{34} & \cdots & 0 & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & 0 & \cdots & \sigma_{N-3}\tilde{Q}_{N-3,N-3} & \sigma_{N-3}\tilde{Q}_{N-3,N-2} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \cdots & 0 & \sigma_{N-2} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \cdots & 0 & 0 & \sigma_{N-1} & 0 \\ 0 & 0 & 0 & 0 & 0 & \cdots & 0 & 0 & 0 & 0 \end{pmatrix}.$$

To understand the non-zero elements of $\mathcal{A}$, note that they account for the probabilities of financial contagion outside of the firm's meeting. That is, the probabilities that, conditional on the state being k today, and conditional on the firm being healthy and meeting another healthy firm, we reach a state $k' \geq k$ in the following round. The positive elements in $\mathcal{A}$ are obtained by multiplying each row $k+1 = 1, \ldots, N-1$ of $\mathcal{Q}$ by σ_k, i.e., the probability of a healthy firm meets another healthy firm when $k = 0, \ldots, N-2$. The two last rows have zero vectors because there cannot be further contagion outside the firm's meeting. In row N, corresponding to the state $k = N-1$, the healthy firm meets a distressed firm, as $\sigma_{N-1} = 0$. In the last row, the firm cannot be healthy, since the state is $k = N$. These last two rows, simply account for w_{N-1} and w_N.

Matrix $\mathcal{B}$, accounts for the continuation payoff components w_k that emerge if the firm is healthy, meets a distressed firm, but does not become distressed (with probability $1-p$). We have:

$$\mathcal{B} := \begin{pmatrix} 1-\sigma_0 & 0 & 0 & 0 & \cdots & 0 & 0 & 0 \\ 0 & 1-\sigma_1 & 0 & 0 & \cdots & 0 & 0 & 0 \\ 0 & 0 & (1-\sigma_2)\tilde{Q}_{11} & (1-\sigma_2)\tilde{Q}_{12} & \cdots & 0 & 0 & 0 \\ 0 & 0 & 0 & (1-\sigma_3)\tilde{Q}_{22} & \cdots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & (1-\sigma_{N-2})\tilde{Q}_{N-3,N-3} & (1-\sigma_{N-2})\tilde{Q}_{N-3,N-2} & 0 \\ 0 & 0 & 0 & 0 & \cdots & 0 & 1-\sigma_{N-1} & 0 \\ 0 & 0 & 0 & 0 & \cdots & 0 & 0 & 1 \end{pmatrix}.$$

As for matrix $\mathcal{A}$, the non-zero elements of $\mathcal{B}$ account for the probabilities of financial contagion *outside* of the firm's meeting. The positive elements in $\mathcal{B}$ are obtained by multiplying each row $k+1 = 1, \ldots, N-1$ of $\mathcal{Q}$ by $1-\sigma_k$, i.e., the probability of a healthy meeting a distressed firm when $k = 0, \ldots, N-2$. The first row is a zero vector because everyone is healthy ($k = 0$), so this row accounts for w_0. The last row, corresponding to the state $k = N$, is also a vector of zeros because everyone is distressed already. So this last row accounts for w_N.

To account for the continuation payoff component w_N that emerges if the firm is healthy, meets a distressed firm, and becomes distressed (with probability p), we define the $(N+1)-$dimensional column vectors

$$\boldsymbol{\sigma} := (\sigma_0, \ldots, \sigma_{N-1}, 0)^\mathsf{T},$$
$$\mathbf{1} := (1, 1, \ldots, 1)^\mathsf{T}.$$

Recall that e_k is the $(N+1)$-dimensional column vector that has 1 in the $(k+1)^{th}$ entry and zero everywhere else. Putting all these elements together, we can define $\boldsymbol{w}$ recursively as:

$$\boldsymbol{w} = (y + h\delta w_0)\mathbf{1} - e_N(1-\alpha)y + (1-h)\delta[\mathcal{A}\boldsymbol{w} + p(\mathbf{1} - \boldsymbol{\sigma})w_N + (1-p)\mathcal{B}\boldsymbol{w})].$$

Here, the element $-e_N(1-\alpha)y$ captures the drop in earnings for the vector component $N+1$, i.e., when $k = N$ since in that case the earnings in a trade meetings are αy.

The equation above can be rewritten as

$$\{\mathcal{I} - (1-h)\delta[\mathcal{A} + (1-p)\mathcal{B}\}\mathbf{w} = (y + h\delta w_0)\mathbf{1} - e_N(1-\alpha)y + (1-h)p\delta(\mathbf{1} - \boldsymbol{\sigma})w_N$$

The result of the lemma immediately follows from multiplying the expression above, first, by $\{\mathcal{I} - (1-h)\delta[\mathcal{A} + (1-p)\mathcal{B}]\}^{-1}$ and, then, by $e_k^\top$.