Crisis Risk and Risk Management

René M. Stulz,
The Ohio State University, NBER, and ECGI

Dice Center WP 2023-10
Fisher College of Business WP 2023-03-010

May 2023

Crisis Risk and Risk Management

René M. Stulz*

May 2023

Preliminary

ABSTRACT

This paper assesses the current state of knowledge about crisis risk and its implications for risk management. Better data that became available since the Global Financial Crisis (GFC) has improved our understanding of crisis risk. These data have been used to show that some types of crises become predictable when one accounts for interactions between risks. Specifically, a financial crisis is much more likely in the years following both high credit growth and high asset valuations. However, some other types of crises do not seem predictable. There is no evidence that the frequency of economic and financial crises is increasing. The existing data show that political crises make economic crises more likely, so that, as suggested by the concept of polycrisis, feedback between non-economic crises and economic crises can be important, but there is no comparable evidence for climate events. Strategies that increase firm operational and financial flexibility appear successful at reducing the adverse impact of crises on firms.

* Fisher College of Business at The Ohio State University, NBER, and ECGI. Prepared for the NYU VRI conference on "Managing Compound Risks in a Polycrisis World." I am grateful for comments from Richard Berner, Harry DeAngelo, Leandro Sanz, and Til Schuermann, as well as from the conference participants.

1. Introduction.

Since the 1990s, much progress has been made in identifying, measuring, and monitoring risks. Armed with these identification, measurement, and monitoring tools, firms can make decisions about what risks to take, what risks to avoid, and what risks to mitigate. We call risk management the process of identifying, measuring, monitoring, and managing risks. Despite this progress, some observers argued both during and after the global financial crisis (GFC) that risk managers failed to anticipate and prepare for the crisis. These observers were generally not convinced by the argument of the Economics Nobel prizewinner Robert Lucas that, because markets are efficient, "The best and only realistic thing you can do in this context is to keep your eyes open and hope for the best."[1] The last few years have heightened this concern about how risk managers should deal with crisis risk as the COVID-19 crisis appeared to be the ultimate black swan event. Most recently, there has been much discussion that we are entering a "new epoch of crisis,"[2] that we are in a polycrisis or may be facing the heightened risk of a polycrisis. In short, a polycrisis is a situation where multiple crises are occurring at close to the same time and they are feeding on each other so that the adverse impact of a polycrisis is worse than the adverse impact of the sum of the individual crises. In this paper, we assess the current state of knowledge about crisis risk and draw some implications for the practice of risk management. This paper is not a review of the literature as excellent reviews already exist.[3]

To understand crisis risk, we need to define what a crisis is. A frequent definition is that a crisis is a tail outcome in a distribution. For instance, for growth of GDP per capita, a crisis could be an outcome that has less than a 1% chance of occurring given the historical distribution of GDP per capita growth. This definition is incomplete. If extreme outcomes are not qualitatively different from less extreme outcomes, there is no good reason to focus on these extreme outcomes separately from focusing on the whole distribution of outcomes. Studying crises makes sense as a separate topic of study if the economy or critical

[1] "In defence of the dismal science," by Robert Lucas, The Economist, August 6, 2009.
[2] "The slow-motion tidal wave consuming our economy," by Richard Bookstaber, The New York Times, March 27, 2023.
[3] See, for instance, Claessens and Kose (2013) and Sufi and Taylor (2021), for reviews of the literature on financial crises.

parts of the economy function differently in a crisis than in non-crisis or normal times. A good analogy is a car engine. The engine can work at low speeds or high speeds. However, if a crucial component of the engine is broken, the engine stops working normally and no longer performs its function. If the engine has a low speed, pressing on the accelerator will solve the problem. However, that approach will be useless if a key component is broken. If a crisis means that the financial markets do not function normally, then the financial policies a firm has put in place that rely on financial markets functioning normally will not work and the firm will need a different set of financial policies to deal with a crisis. We will define a crisis as a tail outcome for the economy such that there are important disruptions in the functioning of the economy or of critical parts of the economy.

Much effort has been spent in economics and finance to construct global databases of crises since the GFC. Global databases are required for statistical analysis because crises are too rare at the country level. Some of these databases rely on quantitative definitions of crises, while others use narrative definitions. We want to focus on crises that affect the whole economy rather than crises that affect specific industries or specific sectors of the economy. For crises that affect the whole economy, there are two broad approaches that have been followed to construct databases. One approach identifies instances where there is a sharp drop in economic growth or consumption. Perhaps the best-known database is the one from Barro and Ursúa (2008) that sets the threshold for a crisis at a cumulative 10% drop in consumption or GDP from the start of the crisis to the trough. Using this threshold, a recent study by Ćorić (2021), that seems to include the largest number of countries, finds that the probability of a country entering an episode of a 10% cumulative drop in real GDP from 1820 to 2016 is 2.32% per year. This definition constitutes a high threshold for an episode to constitute a crisis as the GFC is not a crisis in the U.S. by this definition. We will follow the literature and call these crises catastrophes. An extension of this approach is to focus on recessions instead, but this definition seems too low of a threshold. While it is likely that the economy does not work normally for an extreme drop in consumption or economic growth such as a 10% drop, there are many recessions where the economy does not experience important disruptions. In 19 developed countries for which there is extensive data since 1870, there are over 200 peacetime recessions (Jordà, Schularick, Taylor, 2020).

The second definition refers to financial or banking crises. This definition has the advantage that it focuses directly on situations where the financial system stops working normally. Many different databases have been constructed for financial crises that differ along some dimensions. Sufi and Taylor (2022) review these databases and discuss how they differ. With these databases, the U.S. had no post-WWII financial crisis until the savings and loan crisis of the 1980s and the GFC was the only other financial crisis before the COVID pandemic. Using a sample of 14 countries from 1870 to 2008, Jordà, Schularick, and Taylor (2013) find that one peacetime recession out of four involves a financial crisis. The consensus in the literature is that a recession that involves a financial crisis is deeper than a recession without a financial crisis. Research since the GFC indicates that financial crises are more predictable than previously believed provided one takes into account interactions among risks. In this context, predictability means that the probability that a crisis will occur over some future period changes over time, so that there are times when the probability that a crisis will occur is much larger than at other times. Specifically, a financial crisis is much more likely to occur in the next two to four years if credit growth and asset returns have been high. Neither credit growth nor asset returns separately increase the probability of the occurrence of a financial crisis much.

Crisis risk involves the risk of a type of crisis we have experienced in the past as well as the risk of an unknown crisis type or a type of crisis we know exists but with which we have little experience. The existing quantitative work on crises takes the view that data on crises in the past is useful for forecasting the likelihood of future crises. Statistical tools are of great help in assessing the risk of crises we have experienced historically, but they have three important limitations. First, they require that the probability distribution of crisis occurrence is stationary. As we will see, the predictability literature models changes in the probability of crisis occurrence, but the conditional distribution in that literature is stationary – given the condition of the economy, the probability of occurrence is the same. This stationarity assumption seems more plausible for some risks than others. For instance, it would not make sense that the risk of extreme weather outcomes will be the same in the future than it was in the past. Second, while some types of crises are predictable, some types of crises are closer to black swans. The COVID-19 pandemic is a crisis of the

latter type. It caused a sharp decrease in income worldwide, but it was not predictable. While we knew that there was a risk of a pandemic, the historical frequency of pandemics was extremely low. There are other types of crises with extremely low probabilities. There are types of crises that we cannot even describe because we have never experienced them. Third, it would be a mistake to assume that we have a full catalogue of possible crisis types. To quote Secretary of Defense Rumsfeld, there are unknowns unknowns.[4] Using a terminology more grounded in statistics and economics, there is what Kay and King (2020) call radical uncertainty.

Though much of the focus in developed countries has been on banking crises, the literature also studies currency crises and sovereign crises. All three types of crises can occur at the same time, or one type of crisis can cause another one. Acharya, Drechsler, and Schnabl (2014), for instance, show how interactions between the state and banks can create a feedback effect where a banking crisis weakens the financial situation of the state as the state supports the banking sector and the weakened situation of the state causes further deterioration in the financial situation of banks. The existing crisis literature measures crises at the country level. A crisis in one country can propagate to another country. The phenomenon of propagation is often called contagion. With contagion, a country that is healthy falls into a crisis state because there is a crisis in another country. Contagion can arise because of financial links as well as because of trade links. There is also a school of thought that contagion can simply arise because of animal spirits, somewhat like seeing an accident reduces traffic speed.

The polycrisis concept is generally understood to mean that different types of crises occur all at once and interact. Following Tooze (2021), discussions of polycrises so far have been qualitative. Tooze (2022) argues that we have simultaneously geopolitical crises as well as a climate crisis and an energy crisis. He does not define what makes a crisis and has no quantitative criteria to define crises. To consider the frequency of other types of crises, we consider foreign policy crises and climate events in part because

[4] Kuritzkes and Shuermann (2010) have a useful set of definitions for the various types of risks: "A risk is *known* (**K**) if it can be identified and quantified ex ante. (…) A risk is *unknown* (**u**) if it belongs to a set of risks that can be identified but not meaningfully quantified at present. (…) A risk is *unknowable* (**U**) if the existence of the risk or set of risks is not predictable, let alone quantifiable, ex ante."

Tooze (2022) focuses specifically on these crises but also because data is available. While geopolitical crises are the source of economic catastrophes throughout history, none of the recorded economic catastrophes in the databases we consider have climate-related events at their origin. This is an emerging area of research, but the existing evidence for the U.S. shows that no climate event has had anything close to systemic implications.

The evidence on crisis risk accumulated since the GFC has many lessons for risk management. Both financial institutions and bank regulators tend to calibrate risk models so that financial institutions have enough capital to cope with events that reach far in the tail of the distribution of outcomes (Schuermann, 2020). However, data on the frequency of crises makes it plain that models calibrated that way imply that the financial institution has to have enough capital to cope with most crises that have occurred in history. Hence, a better understanding of crisis risk is important to determine what type of events a firm has to have enough capital to cope with. However, the fact that there is a great deal of uncertainty concerning crises in general and also about the future distribution of crisis risk means that calibrating risk management models to what we know about past crisis risk means that one likely understates the true crisis risk.

The type of crises we discuss affects the whole economy. Commonly, firms use derivatives and insurance contracts to offset losses that could put them in costly financial distress. However, with crises, such tools are of limited use. We cannot all buy insurance or take derivatives positions against crises. Instead of using financial instruments to mitigate the impact of crises, firms can choose to be organized in such a way that they are more resilient. Resilience as an objective for risk management has the advantage that we can all become more resilient and that it helps to deal with known risks as well as unknown risks. A good analogy is to think about which type of car is most appropriate. A race car optimized for speed will be best if the weather and the road are perfect. However, if there is any risk of bad weather or stretches of poor-quality road, a four-wheel SUV will complete the trip when the race car will end up in a ditch. Firms can be built like race cars or like SUVs. They can choose to be built so that a shock throws them off course permanently or so that they can keep going even if a wheel gets stuck in the mud. Firms choose the level of financial and operational flexibility they have. More financial and operational flexibility will make a

firm better able to cope with shocks. A lesson from the COVID-19 crisis is that firms with more financial and operational flexibility performed better. If crises are rare, building resilience to crises may not make sense for shareholders because resilience may be too expensive. In contrast, if crises are relatively frequent, building resilience may be optimal for firms. Resilience involves a trade-off for firms. More resilient firms are likely to be less profitable in normal times, but more profitable in crises. Consequently, how much to invest in resilience depends on the frequency and costs of crises.

The paper proceeds as follows. In Section 2, we introduce various crisis definitions and provide data on the frequency of crises for various definitions. We then turn to what we know about the predictability of crises in Section 3. In Section 4, we examine how crises propagate through contagion and the impact of non-economic crises. In Section 5, we show how existing crisis data and our understanding of crises can help make sense of the phenomenon of polycrisis. In Section 6, we address the implications for risk management of what we have learned about crisis risk. We conclude in Section 7.

2. How frequent are crises?

Economists have devoted considerable efforts to building databases of crises. One goal of these efforts is to improve our understanding of crises and to help assess the conditions that cause crises. However, another important motivation for such databases is to be able to estimate the average return of financial assets over periods of time that include crises. If crises are rare, a country may have high stock returns over an extended period of time simply because it has not had a crisis. As a result, the high stock returns of that country are not indicative of the average return of holding stocks. They are only indicative of the returns of stocks when there is no crisis. These databases generally cover many countries over a long period of time. By definition, crises must be infrequent events. Because they are infrequent at the level of individual countries, it is difficult to reach conclusions about the frequency of crises and about their predictability using the data of only one country. The hope underlying the construction of global crises databases is that there is enough commonality in the distribution of crises across countries that statistical inference can be conducted by pooling countries.

In this section, we consider first databases that consider economic catastrophes, namely sharp decreases in economic production or in consumption. We then turn to financial crises. We end the section by discussing data for other types of crises.

2.1. Catastrophes and recessions.

Barro and Ursúa (2008, 2012) develop a database that measures cumulative declines in output or consumption of at least 10% from peak to trough across a large number of countries. Ćorić (2021) extends the database so that it ends in 2016 for GDP per capita and includes many more countries. For real GDP per capita, the dataset has 495 disasters since 1820.[5] An issue with such datasets is that the number of countries increases over time. A country can enter the dataset because it did not exist before or because its data was not available before. As a result, the dataset starts with one country in 1820 and ends with 163 in 2016. The dataset has a large discrete increase in the number of countries in 1950 as the number of countries increases from 56 to 135 in 1950.

The paper finds fewer disasters using aggregate real GDP than real GDP per capita. For aggregate real GDP, there are 324 disasters and 13,961 observations, so that the probability of entering into an economic disaster is 2.32 percent a year. The distribution of outcomes for disasters is highly skewed. The most likely outcome is a fall in real GDP between 10% and 15%.

Disasters are much more likely in countries that have a low GDP per capita compared to the U.S. These countries are less resilient along many dimensions. As discussed later, poor institutions lead to more economic instability. Further, poorer countries are often more dependent on few export goods, so that they are less diversified. The duration of a disaster follows a skewed distribution with a mode of two years, so that some disasters last over ten years. Most of the disasters involve countries with less than 10 million inhabitants with per capita income that is typically a relatively small fraction of the per capita income of the U.S. In particular, 50% of the disasters are in countries with GDP below 23% of U.S. GDP per capita.

[5] We thank Bruno Ćorić for making the data available to us.

There are only nine disasters where the GDP of the country is larger than U.S. GDP. After WWII, almost all disasters are in countries that were not part of the OECD at its formation. Specifically, there are only four economic disasters in OECD countries.

In Figure 1, we show the percentage of countries that are in a disaster state over time. It is immediately apparent that the 2000s until 2016 were a period with relatively few disasters. In the Figure, we show the percentage of countries in disaster using the whole dataset as well as the percentage of countries in disaster using the 49 countries that are in the dataset in 1949. The countries that are already in the dataset in 1949 have a much lower frequency of being in a disaster than the countries that were added in 1950. This confirms that developed countries are less likely to be in a disaster state post-WWII.

While Figure 1 uses all data available, it is also useful to look at a dataset of 40 countries for which data is available for a long time period. For many of these countries, the data start in 1870. Using that dataset, Barro and Ursúa (2008) show that using cumulative real GDP growth of -10% as the criterion for a disaster, there are 183 catastrophes in their sample. In this dataset, 51 catastrophes are before World War I and 35 are post-World War II. Wars are periods with a large number of countries that fall into a catastrophe: 31 out of 40 in World War I and 35 out of 40 in World War II. The Great Depression has 23 countries falling into a catastrophe.

An alternative approach is to look at recessions. Jordà, Schularick, and Taylor (2020) present data that makes it possible to look at frequencies of recessions since 1870 for 17 advanced economies. The probability that an advanced country enters a recession from 1870 to 2008 is 11%, so slightly more frequently than one year out of ten. The highest probability is 17% for the period 1870 to 1913. The lowest probability is 6% from 1946 to 1972. The period from 1973 to 2008 has a slightly lower probability than average at 9%.

2.2. Financial crises.

Most of the data on financial crises is data on banking crises. A good definition of a financial crisis is that it "is a breakdown of normal financial market activities to such an extent that capital flees, resources

are misallocated, institutions are destabilized, and disorder spills into the real economy, causing job loss, bankruptcies, recession, and social distress and political ferment." (Bruner and Carr, 2023). Whereas catastrophes and recessions are determined based on quantitative criteria alone, classifications of financial crises are often mostly qualitative. With these classifications, a financial crisis is one where the banking system is distressed so that it does not function normally and/or there is a need for large-scale government intervention to prevent failure of the financial system. We discuss the evidence using this approach first and then turn to a more quantitative definition.

At the end of the previous section, we discussed data on the frequency of recessions among developed countries. For that dataset, slightly more than one recession in four involves a financial crisis and recessions with a financial crisis are deeper. The probability that a financial crisis recession starts in a given year is 4%. From 1946 to 1972, there is no financial crisis recession in the dataset. From 1973 to 2008, the probability that a country has a financial crisis recession start in a given year is 3%.

In Figure 2, we show the distribution of the number of countries having a financial crisis from 1870 to 2020 using the Jordà, Knoll, Kuvshinov, Schularick, and Taylor (2019) dataset of 18 developed countries.[6] In this Figure, the long period from World War II to 1974 is striking by its absence of financial crises. Before World War I, it was often the case that in a given year a country or two would have a financial crisis. For instance, in the first ten years of the dataset, crises start in six years out of ten. The same is also the case in the 1920s. However, 8 countries start a crisis in 1931 and 11 start a crisis in 2008. No other year in the dataset has as many crises starting in one year.

Other databases of crises are more encompassing in the countries they consider. One widely used data set is Laeven and Valencia (2020). This database starts in 1970. It defines a banking crisis as a situation with significant signs of financial distress in the banking system, including bank failures, and significant banking policy interventions in response to losses in the banking system. The database's focus is the whole world. The database includes 151 banking crises from 1970 to 2017. It finds that only four countries have

[6] The current version of the data is available at https://www.macrohistory.net/database/ .

a banking crisis start since 2011: Cyprus, Guinea-Bissau, Moldova, and Ukraine. Laeven and Valencia (2020) find only three countries that experience more than two systemic banking crises over the sample period. From 1970 to 2017, the global financial crisis (GFC) of 2008 stands out. No other year has more than three high-income countries in crisis.

A more recent approach in identifying banking crises focuses on a more quantitative methodology. Baron, Verner, and Xiong (2021) create a dataset that is the union of two sets of crises: bank equity crises and bank panic crises. For a bank equity crisis, they require a 30% drop in the value of bank equity and widespread bank failures. They define widespread bank failures as the failure of a top five bank or five more bank failures than typical. For a bank panic crisis, they require a rapid withdrawal of deposits from the banking system. Their dataset includes 46 countries from 1870 to 2016. They find 151 crises that result from a decrease in equity values and 192 crises that involve a panic. A panic can occur without a drop in bank equity values of 30% or more and a drop in bank equity values of 30% or more can occur without a panic. In Figure 3, we show the distribution over time of financial crises that experience a drop in bank equity values of 30% or more and of financial crises that experience a panic. Four crises stand out: 1907, 1914, 1931, and 2008. The frequency of a bank panic is 2.68% for the whole sample period, but 3.91% since 1980. For bank equity crises, the frequency is 2.09% for the whole sample period and 4.36% since 1980. This evidence suggests that the frequency of crises is higher in the recent past than over the whole sample period. It is also interesting to note that large drops in bank equity values without panics seem to occur more frequently after WWII than before.

Financial crises lead to interventions by governments or central banks. These interventions have become increasingly massive. Metrick and Schmelzing (2021) have constructed a database of these interventions. The database covers the years 1257 to 2019. The crises they consider include the crises discussed in this section. They also add to these databases additional candidate crises. Their database includes 902 crises with 1,886 interventions. They argue that the frequency of interventions and the size of interventions have been rising since the late 17th century. They also find that governments have become more aggressive over time and more so in wealthier countries.

2.3. Other crisis types.

The literature on financial crises has also collected information on currency and sovereign crises. Laeven and Valencia (2020) provide data on both types of crises since 1970 for more than 160 countries. They identify 236 currency crises from 1970 to 2017 using the definition of a depreciation that is at least 30% relative to the dollar or 10% higher than the previous year. For their sovereign crisis database, they find 79 episodes. For either type of crisis, there are almost none for high-income countries over the sample period. For low- and middle-income countries, the largest number of sovereign debt crises per year is in the early eighties where there are nine episodes in 1982 and 1983. There is no year after that with more than four episodes. Throughout the sample period, there is no year with less than one episode. For currency crises, there is no year with no episode for low- and middle-income countries. The peak year is in 1993 with more than 18 episodes, but 2015 has more than 10 episodes.

There are many datasets that provide data on losses of various types. However, these datasets do not define what a crisis is. For instance, data on the number of deaths from geophysical, meteorological and climate events by year is available. It is not clear what threshold would have to be crossed for there to be a crisis. It is noteworthy that, using global data, since 1900 the three decades with the lowest number of deaths are 1910, 1990, and 2010.[7] All these decades record less than 100,000 deaths. The highest decade was the 1920s with more than 500,000 deaths. A study for the U.S. on similar events representing disasters at the county level finds 125 county-level disasters from 1920 to 2010, so 1.4 per year on average (Boustan, Kahn, Rhode, and Yanguas, 2020). Climate events show the limitation of classifying crises as tail events. Suppose one were to define a climate crisis event as an extreme temperature given the historical distribution of temperatures. With such a definition, tail events would occur frequently because the probability that an average temperature in some area is the highest ever in that area is relatively large when temperatures keep increasing. However, such an outcome would not be surprising. It also would not correspond with a situation where the economy no longer functions normally.

[7] https://ourworldindata.org/natural-disasters.

In political science, the International Crisis Behavior Project housed at Duke University maintains a detailed global database of international political crises from 1918 to 2019.[8] These crises are "interstate military-security crises"[9] that meet some specific criteria to insure homogeneity of the database. The database includes 496 crises and has been used in much academic work, including the study of Berkman, Jacobsen, and Lee (2011) discussed later. Figure 4 shows the number of countries in crisis from 1918 to 2019. There are more countries now than at the beginning of the sample period. However, despite this increase in the number of countries over time, the number of countries in crisis in 1939 is higher than at any other time.

3. Are crises predictable?

Systemic banking crises or dramatic shocks to the economy are rare events at the country level. The probability of occurrence of such an event is less than 5% per year in databases that cover a large number of countries. Though systemic banking crises are more frequent post-1980 than from the end of World War II to 1980, they are not more frequent than over longer periods of time. Economic catastrophes involving a 10% cumulative GDP drop are almost non-existent in developed countries from the end of World War II to 2019. The global databases have been used to estimate the frequency of crises and to assess whether crises are predictable. In this section, we first examine evidence on whether there are systematic differences in crisis probability across countries. We then discuss the evidence that crisis probabilities differ across time within countries, so that some types of crisis are predictable. Lastly, we examine how the evidence on predictability differs across crisis types.

3.1. Do crisis risks differ across countries?

In the Baron, Verner, and Xiong (2021) database of crises, Argentina has nine crises since 1870. One might conclude that it is not surprising that Argentina has so many crises. However, the U.S. also has nine

[8] https://sites.duke.edu/icbdata/.
[9] See Brecher and Wilkenfeld (1997).

crises in the database. The record number of crises goes to Japan with 12 followed by Italy with 11. In contrast, Indonesia has only two crises and Canada only three. One way to look at this data is to say that crises occur randomly and that if a crisis has a 5% probability of occurring in a given year, we expect to observe seven crises over 146 years. There will be random variation in the number of crises and, because of that random variation, some countries will have fewer crises than seven and others will have more. However, whether a crisis occurs is like the outcome of flipping a weighted coin where the weight is the same for all countries.

Existing empirical research suggests that some country-specific factors influence the probability that a country will experience a financial crisis and some countries can be expected to experience more financial crises than others. Demirgüç-Kunt and Detragiache (1998) examine the determinants of banking crises and find that, in their dataset from 1980 to 1992, banking crises are less likely in countries where law and order are better respected. They also find that countries with deposit insurance are more likely to have banking crises. This evidence suggests that moral hazard may play an important role in having poor economic outcomes develop into crises.

Acemoglu, Johnson, Robinson, and Thaicharoen (2003) conduct a more extensive study focused on the institutional causes of macroeconomic volatility and crises. To measure the quality of institutions, they use an index from political sciences that measures the extent of constitutional constraints on the exercise of arbitrary power by the executive. The score is available for every independent country. Some of the authors of the study have shown in other work that the index is correlated with other measures of the quality of institutions and with financial development. They conclude that the quality of institutions plays a large role in explaining macroeconomic volatility in contrast to macroeconomic variables. For instance, when they do not allow a role for institutions, they find that the size of government consumption is positively associated with macroeconomic volatility. However, when they control for institutions, the size of government consumption is no longer significant. The study investigates the severity of crises. It finds that crises are worse with lower quality of institutions controlling for the level of economic development.

Institutions are not constant, so that changes in institutions could affect macroeconomic volatility and crises. For instance, Funke, Schularick, and Trebesch (2021) show that following the election of populist leaders, the constraints on the executive decrease. They find that economic performance weakens substantially following the election of populist leaders, but in addition their evidence suggests that the risk of crises increases as well. Crises themselves could lead to changes in the political leaders and might affect the probability of subsequent crises. The same authors in Funke, Schularick, and Trebesh (2016) find that crises are followed by an increase in the following of far-right leaders.

3.2. Do crisis risks differ across time?

Does the risk of a crisis vary over time and what does it depend on? Studies before the GFC seemed to focus more on macroeconomic weaknesses as predictors of crises. For instance, Demirgüç-Kunt and Detragiache (1998) find for their sample from 1980 to 1992 that crises are more likely when a country has low growth, high real interest rates, and high inflation. They also find some evidence that a credit boom precedes a banking crisis. However, the consensus before the GFC was that the extent to which crises probabilities changed through time was quite limited.

More recent studies use the type of database we discussed in Section 1. These databases have much longer sample periods and typically more countries. However, in addition, these studies allow for interactions between different types of risks. These interactions appear to make a large difference in the predictability of financial crises. Baron, Xiong, and Ye (2023) study a database of 20 advanced economies from 1870 to 2021. They look at GDP crises defined as a one-year change in growth rate of GDP below the 2nd percentile of historical distributions over the past 50 years. They exclude the World Wars to determine the threshold for GDP growth. With this data, they have 102 crises. Of these crises, 35 are banking crises, 43 result from wars, and 9 from a natural disaster. All nine episodes of natural disaster are in 2020 and correspond to the COVID-19 pandemic. They have 15 crises that are unclassified. The frequency of experiencing a crisis in a particular year is 2.7% from 1870 to 1949 and 1.8% from 1950 to 2021. The severity of crises falls sharply between the two periods as it is -11.5% from 1870 to 1949 and then -5.9%

from 1950 to 2021. They find that when credit expansion over the last year is in the top quintile constructed from past data and past market returns are in the top decile of past three-year returns using past data, the probability of a crisis in the next two to four years is 11.3% compared to the unconditional probability of a GDP crisis in two to four years of 6%.

Baron, Xiong, and Ye (2023) construct a disaster index which is the probability of a crisis in 2 to 4 years. This index depends on last year's credit expansion, last three years' stock returns, and the interaction between credit expansion and stock returns. The interaction between credit expansion and stock returns is the most important variable in explaining GDP crises. Future stock returns are negatively related to the disaster index, so that a high disaster index predicts negative stock returns. The disaster index predicts half of the GDP crises. Not surprisingly, it also predicts GDP crises that do not happen, but the false positive rate is only 13%. It is interesting to note that the disaster index is elevated ahead of the 2008 crisis in the U.S.

Greenwood, Hanson, Shleifer, and Sørensen (2022) focus directly on predicting financial crises rather than GDP crises. They use a sample of 42 countries from 1950 to 2016. For that sample, they find that a country that has rapid credit and asset price growth over the prior three years has a 40% probability of entering in a financial crisis within the next three years. That probability is 7% in normal times when neither credit nor asset price growth are high. They show that credit growth alone has a modest impact in predicting a crisis. For a crisis to have a high probability of occurring, one needs to have both high credit growth and high asset price growth.

The authors construct a Red Zone indicator. A country is in the business Red Zone if nonfinancial business credit growth is in the top quintile over the last three years and if stock market returns over the same window are in the top tercile. If a country is in the business Red Zone, the probability of a crisis in the next three years is 45% in contrast to the unconditional probability of 4%. They also construct a household Red Zone. A country is in the household Red Zone if the growth of household debt over the last three years is in the top quintile and if the growth in house prices is in the top tercile. With the household Red Zone, the probability of a crisis in the next three years is 27%. They find that 64% of the crises have

overheating in businesses or households. The probability of a crisis in the next three years is 68.8% when a country is overheating in both businesses and households. However, it is rare for a country to be overheating in both business and households. The U.S. was in the household Red Zone from 2002 to 2006. However, the predictability of crises result holds when the GFC is not in the sample.

The studies considered in this section are focused on predicting crises over a multi-year window. Other studies focus instead on predicting crises one year ahead. Such studies have a different interpretation. A recent example of such a study is Marfè and Pénasse (2023). They construct a sample of crises where a crisis is a one-year drop in consumption that is greater than two standard deviations of the consumption growth rate in a country. Their sample period is 1876 to 2020 and their sample includes 42 countries. With their definition of a crisis, their sample has 177 crises. The probability of a crisis is 3.7%. The studies discussed earlier in this section predict a crisis when asset prices are high. This study has the opposite prediction, which is that a crisis is more likely when asset prices are low. Baron, Xiong, and Ye (2023) have a straightforward explanation for this difference in results. Large drops in consumption tend to occur after a fall in asset prices. In other words, the financial crisis comes first followed by the economic crisis. As a result, a financial crisis is a good predictor of an economic crisis. In their dataset, the U.S. has three crises: 1921, 1930, and 1932. Given the rarity of crises, a country has a probability of roughly 10% of not experiencing a crisis in 75 years. The best predictors of a crisis are whether a country is in a recession already, whether there is a war or political crisis abroad, whether the country is at war, and whether stock prices are low.

There are two important lessons from this section. First, the research on the predictability of crises has made progress since the GFC in that there is now strong evidence that using interactions between credit conditions and asset valuations helps predict crises. This type of crisis involves a financial crisis. Many crises start as a financial crisis that is followed by a sharp decrease in consumption. The literature does not seem to have much success in predicting crises that do not involve a financial crisis. The types of crises for which there is little evidence on predictability might be shown to be predictable with future research, but for now little can be said about their predictability. Second, the situation of financial markets is helpful in

predicting financial crises, but not in the way one would expect. The comment of Lucas cited at the beginning of this paper implies that crises are not predictable because, if they were, they would happen immediately. His statement is certainly correct in the sense that if it were learned that the stock market would fall by a large amount at a known future date for sure, it would fall right now. More generally, if financial markets were anticipating a higher risk of a crisis, they would value risky assets less today. As a result, we would expect asset values to be low when the risk of a crisis is higher. The opposite is the case. High asset values predict a higher likelihood of a future crisis. Such a result is hard to reconcile with rational expectations, but is consistent with investors having expectations where they extrapolate from the recent past so that they expect more good times ahead (see, for instance, Gennaioli, Shleifer, and Vishny, 2015). Related research by López-de-Salido, Stein, and Zakrajšek (2017) shows elevated credit sentiment in the prior two years predicts a decline in economic activity in the current and following year.

4. Contagion and compound risks.

There is a vast literature in finance and economics that examines how crises travel across countries as well as how different types of crises can interact or feed on each other. In this section, we first discuss this literature. We then turn to some evidence concerning the impact of non-economic shocks. As before, we do not review the literature but focus on what we believe is the most relevant evidence for our discussion.

4.1. Multiple crises.

In Section 1, we addressed the issue of the definition of a crisis and of the frequency of crises. We discussed separately economic crises and financial crises. An economic crisis is one that involves a sharp decline in consumption growth. Such a crisis can occur with or without a financial crisis. The literature has examined the interaction between financial crises and economic crises. Jordà, Schularick, and Taylor (2013) investigate recessions in 14 developed countries from 1870 to 2008. They distinguish between "normal" recessions and "financial" recessions. The "financial" recessions are recessions following financial crises. They find a large difference between normal recessions and financial recessions. After five years, the real

GDP per capita is 5% lower for a financial recession than a normal recession. It follows from this that the interaction of a financial crisis with a recession has a large adverse impact on the severity of the recession. Baron, Wei, and Ye (2023) have a sample of 102 economic crises. In that sample, 35 economic crises are also banking crises. For these crises, GDP per capita falls by 8.8% in the year of the economic crisis. In contrast, 43 economic crises involve war and the drop in GDP per capita is 19.5%. Comparing the results of the two studies, it follows that financial crises make recessions worse, but the worst economic crises are those involving war.

Greenwood, Hanson, Shleifer, and Sørensen (2022) address the issue of whether a higher probability of a financial crisis in foreign countries is associated with a higher probability of a financial crisis in the local country. They construct an index of the fraction of countries that are in a Red Zone. They find that this variable helps predict crises. In particular, a country that is not in a Red Zone when a high fraction of countries is in a Red Zone has a probability of crisis that is close to the one it would have if it were in a Red Zone. When they account for the global variables, they find that the probability of a subsequent crisis in the next three years rose from 31% in 2002 when the US first entered the household Red Zone to 51% in 2006. Similarly, the study of Marfè and Pénasse (2023) finds that a country is more likely to be in crisis next year if other countries are in crisis this year.

Crises can spread across countries through many channels (Pritsker, 2000, reviews these channels). However, for a crisis to impact another country, the two countries cannot be isolated from each other and have to have some types of commonalities. For instance, the literature has shown the existence of contagion through trade (Dornbusch, Park, and Claessens, 2000). If a country experiences a crisis, it will import less, which will affect the exports of the other country. Alternatively, the two countries could have similar investments. As these similar investments lose value, each country will become poorer. With this perspective, globalization should lead to more contagion of crises as countries are more closely connected. Some of this literature emphasizes the impact of changes in rates in the U.S. A sharp increase in rates in the U.S. can have an adverse impact on the rest of the world and hence cause other countries to enter a state of crisis. Many emerging countries were in crisis in the early 1980s and part of the explanation was that these

countries had dollar debts that became more expensive. In the 1990s, there was much discussion that capital flows themselves could destabilize economies and contribute to crises. Karolyi (2003) reviews the existing evidence on this argument and points to reasons to be doubtful about its importance.

Forbes (2012) distinguishes between interdependence and contagion. She defines interdependence as a high correlation across markets. In contrast, contagion involves spillovers from extreme negative events. Not surprisingly, interdependence has increased over time. She finds that countries more vulnerable to contagion have more levered banking systems, greater trade exposure, weaker macroeconomic fundamentals, and greater external liabilities. Bae, Karolyi, and Stulz (2003) examine contagion using stock returns across countries. They show that correlations are not useful to understand the joint occurrence across countries of extreme stock index returns. They use a logistic multinomial model to predict joint extreme returns across countries and find that joint occurrences of extreme returns across countries are more likely when interest rates are high. Their study shows the limitations of statistical tools that are used to model normal times for understanding periods of extreme outcomes.

Though the studies discussed in Section 2 did not make that distinction, the literature also distinguishes between financial, sovereign, and currency crises. Laeven and Valencia (2020) document for their dataset that a country with a banking crisis has a 22.5% probability of having a currency crisis over the next two following years. A country with a currency crisis has an 8.9% probability of having a banking crisis and a 6.8% probability of having a sovereign crisis over the next two years. Lastly, while a country with a sovereign crisis has only a 2.6% probability of having a banking crisis over the next two years, it has a 16.9% probability of having a currency crisis. The connection between sovereign crises and banking crises played an important role in the European crisis that followed the GFC. Countries that spent much to bail out their banks ended up in a weakened state which increased the cost of their sovereign debt (Acharya, Drechsler, and Schnabl, 2014).

When we are considering multiple crises, the crises could all occur at once or start sequentially where one crisis causes the next. It is therefore important to understand in which order different types of crises occur and whether the joint occurrence of multiple crises leads to worse economic outcomes. A classic

study (Kaminsky and Reinhart, 1999) conducts such an analysis distinguishing between banking and currency crises. It starts with a narrative approach for crises from the 1970s to 1995. Their sample has 76 currency crises and 26 banking crises. They find that the typical sequence is one where a country takes measures to liberalize its financial system. This liberalization causes a rapid increase in credit that ends with problems in the banking sector followed by a currency crisis. When there is both a currency crisis and a banking crisis, the currency crisis worsens the banking crisis, creating a vicious spiral. In this study, when there is both a currency crisis and a banking crisis, the crises feed on each other and make each other worse.

4.2. Non-economic shocks and crises.

In Section 1, we report results for the frequency of political crises. We would expect that financial crises would impact a country's politics and in Section 2 we reference some studies that show that financial crises give strength to populist movements. It would seem likely that in some cases political crises cause financial crises or disasters. As discussed in Section 2, a large number of economic disasters are associated with wars. Berkman, Jacobsen, and Lee (2011) use the measure of political crises discussed in Section 2 to show how wars are associated with economic catastrophes. Their dataset has 95 separate wars. They have GDP per capita data for 88 different war-actor countries. In 24 cases, GDP does not fall. In the other cases, it falls by a mean of 24.9% and a median of 19.8%. They match their dataset with the Barro and Ursúa (2008) dataset of economic catastrophes from 1918 to 2006. They have 19 countries with 37 economic catastrophes. The impact of a political crisis on the probability of an economic disaster is moderate as the start of a political crisis increases the probability of an economic catastrophe by 2.8%. The unconditional probability of an economic catastrophe is 2.2% in that dataset. The study does not go on to examine how the probability of an economic catastrophe is affected by different types of political crises. It is important to note, however, that a political crisis that can start with relatively moderate acts may eventually evolve into a war.

The Berkman, Jacobsen, and Lee (2011) study is more focused on how political crises affect stock returns. Specifically, they find that when a political crisis starts, stock market valuations drop and they

rebound when the crisis ends. They propose a measure of the severity of a political crisis. They show that more serious political crises have more adverse effects on stock prices. Relatedly, Miller (2023) finds that democratizations reduce stock market valuations. However, his work also shows that democratizations do not affect economic growth. Hence, poor stock market outcomes do not necessarily imply poor outcomes for the economy. Political conflicts or political change could have implications for redistribution that could be adverse for shareholders but that may not necessarily affect economic growth. Existing evidence shows that correlation between economic growth and stock market returns over long periods is actually negative (Ritter, 2013).

It is not straightforward to interpret the impact of crises on the stock market. Stock prices embed subjective expectations from market participants which may differ from objective expectations. Further, stock prices can fall because investors require more compensation for bearing risk as well as because investors expect lower cash flows. It is therefore not surprising that stock prices can fall sharply without there being a financial crisis or an economic catastrophe.

The other type of non-economic shocks that we have discussed are natural shocks, namely shocks involving earthquakes, floods, hurricanes, and other natural phenomena. The databases of economic catastrophes do not seem to trace these catastrophes to natural shocks.[10] If natural shocks include pandemics, then COVID-19 did cause economic tail events across countries and disasters of the early 1920s might be attributable to the Spanish flu pandemic. Though the magnitude of the COVID-19 shock was extremely large on a quarterly basis, it is much less so when measured on a yearly basis. The COVID-19 shock could have been much more damaging to economic growth, but the public sector response, both from the government and from the Federal Reserve, was overwhelming. The ability of authorities to use policy measures to mitigate the impact of shocks means that shocks become less likely to cause a financial crisis or an economic disaster.

[10] Marfè and Pénasse (2023) include famines as one of their factors. It appears to have no meaningful explanatory power. It is likely, however, that databases that cover emerging countries more completely might lead to different conclusions about the role of famines and floods through history.

Natural shocks have localized effects. For them to affect the economy as a whole, they have to be of a magnitude that has not been experienced in the U.S. A recent study of natural disasters in the U.S. by Boustan, Kahn, Rhode, and Yanguas (2020), where disasters are selected based on a disaster declaration, finds only 125 disasters when measured at the county level from 1920 to 2010 that caused 25 or more deaths. Another recent study investigates the impact of natural disasters on the financial sector. Blickle, Hamerling, and Morgan (2022) examine the impact on banks of FEMA disasters over the last quarter century. They find insignificant or small effects on bank performance. Large banks seem to benefit because loan demand increases. Local banks seem to avoid making mortgage loans to the areas where flood maps understate the risk of floods.

There is one example in the history of the U.S. where a natural shock might have contributed to a financial crisis (Odell and Weidenmier, 2004). This natural shock is the San Francisco earthquake. However, to the extent it contributed to the panic of 1907, it was because of the policy response of European countries rather than because of the shock itself. The earthquake had a localized cost that amounted to less than 2% of U.S. GDP. The costs were estimated at an amount of $350-$500 million, but mostly British insurance companies had claims to pay amounting to up to $200 million. As they started paying these claims, they created an outflow of gold from the UK. The Bank of England reacted by raising interest rates and by pressuring UK joint-stock companies from discounting American bills. These actions led to a tightening of credit in the U.S. that set the stage for the panic of 1907 and one of the sharpest but also shortest recession in U.S. history. However, Bruner and Carr (2023) caution us in putting too much weight on the San Francisco earthquake in their definitive study of the 1907 panic. They emphasize that financial "crises are cascades of shocks and information problems to which bank runs, market crashes, rumors, hoarding, fear, and panic are predictable responses."

5. Polycrisis and compound risks.

We now turn to the concept of polycrisis and how to understand this concept in light of the existing knowledge we have about crises that we discussed in the earlier sections. The concept of polycrisis seems

to have originated with Morin and Kern (1999). In that work, they do not refer to the type of crisis we have discussed, but instead refer to events that play out over longer periods of time. Subsequently, the concept is used in relation to the multiple crises confronting the European Union after the GFC, when the European Union was facing what some have described as both a sovereign debt crisis and a migration crisis. Lawrence, Janzwood, and Homer-Dixon (2022) develop the concept further. In their work, they define a crisis as "a sudden (non-linear) event or series of events that significantly harm, in a relatively short period of time, the well-being of a large number of people." They further emphasize that a polycrisis occurs when different systems are in crisis and these crises interact with each other in a way that the outcome of the polycrisis is worse than the impact of the crises if they did not interact. With this approach, a polycrisis differs from a systemic crisis because a polycrisis always involves crises in multiple systems.

Most recently, Tooze (2022) popularized the concept and argued that it could be helpful in understanding the current world. Subsequently, the World Economic Forum (2023) predicted that the "risk of polycrises accelerates." Tooze (2022) argues that the world faces a number of radical challenges or crises. He makes several points about them. Most importantly, many of them are surrounded by considerable uncertainty. At the time he made these arguments, there was considerable concern about whether Russia might launch some nuclear device against Ukraine. Tooze (2022) rightfully points out, as did Kay and King (2020), that it is hard, if not impossible, to attach a meaningful probability to such events. He also stresses that these risks seem to be materializing all at once and that they reinforce each other. He proposes the device called "Krisenbilder" to describe how these risks interact.

The analysis of Tooze (2022) seems to focus more on risks such that tail outcomes for these risks would have serious adverse effects than on the tail outcomes themselves. For instance, a risk is the risk of lethal new COVID-19 variants. This risk could materialize in a new pandemic, but it could also fail to materialize. Similarly, another risk is the risk of a Eurozone sovereign debt crisis. Again, this risk could or could not materialize. It seems therefore better to think of his approach as one that focuses on some risks whose materialization could have important adverse effects, so that we could end up, for instance, with a large adverse shock to the economy or to population. With this perspective, a polycrisis seems a misnomer, but

it would be a situation where there is a high probability that many of these risks will materialize. These risks are such that they interact with each other, so that materialization of one risk makes it more likely that another risk will materialize. They are also such that materialization of two distinct risks at the same time is worse than materialization of these two risks separately.

The contagion literature focuses on the interaction of risks, but it does so among some very specific economic risks. However, we know from the contagion literature that we discussed before that various causation chains have been documented where for instance a banking crisis leads to a sovereign debt crisis which then makes the banking crisis worse. The polycrisis looks at risks with a much wider lense and considers risks that are much less amenable to be assessed quantitatively. For instance, Tooze (2022) includes in his "Krisenbilder" the risk of nuclear escalation. It is possible that some political scientists might suggest a model that would produce a probability distribution for this risk, but it seems reasonable to conclude that many of the risks he focuses on are such that one will have to live with a considerable amount of uncertainty about their frequency and severity.

Risk managers have used heat maps for a very long time. With such maps, they can represent the likelihood that critical risks will materialize and the severity of these risks. One way to construct such maps is to plot risks in the frequency-severity space. In such a heatmap, a polycrisis as discussed by Tooze (2022) would look like a number of high severity risks that have high likelihoods of materialization. The current time is not the first time that the heatmap would look like this. One might easily think of times in the past where the heatmap would have looked worse.

There are two weaknesses of the heatmap that have to be considered in light of the implications of the concept of polycrisis. First, the heatmap does not tell us how much confidence we can have in the assessment of frequency and severity of a risk. It can only represent risks that can be put on the map. Not all risks can be put on the map because some are not known and because there is too much uncertainty about others. There is a danger with any graphical representation that one concludes that it represents the sum of our knowledge, which can make it dangerous. Second, the heatmap does not show how the risks interact. However, the "Krisenbilder" of Tooze (2022), while it shows interactions, does not show the

magnitudes of these interactions and the likelihood of the risks materializing. A better approach would seem to be to construct conditional heatmaps when risks interact. Specifically, one could construct a heatmap that would result from a shift in the likelihood of the materialization of specific risks.

6. Lessons for risk management.

There are many lessons for risk management from our discussion. We focus on four: 1) calibration of risk models; 2) the fact that some crises are predictable but others not; 3) the fact that the future may not be like the past; and 4) mitigation of crisis risk.

6.1. Risk model calibration.

Especially for financial institutions, it is common to calibrate risk models so that the institution will have enough capital to cover losses far in the tail of the distribution of losses (Schuermann, 2020). Regulators have followed this approach as well as financial institutions in their own practice. These calibrations are often performed using the normal distribution. On a yearly basis, having capital to cover losses at the 99.9% confidence level means that the institution has enough capital to cover losses for 999 years out of a thousand. The Basel II Accord used a 99.9% calibration and calibrations at the 99.95% level were not uncommon for models of economic capital. With the data we have presented, financial crises are not as frequent as what Jamie Dimon, the CEO of JPMorgan Chase, purportedly told his daughter, namely that they occur every five to seven years.[11] However, the data does imply that they occur once every 20 to 25 years. With such a frequency, a risk model calibrated to a one in a thousand years tail outcome should be one where the financial institution can ride out comfortably forty crises in a thousand years. Focusing on crisis outcomes is a good way to figure out whether the outcome from risk models comports with historical experience. The same reasoning applies to the design of stress tests.

[11] Top 60 Jamie Dimon Quotes (2023 Update) - Quotefancy.

6.2. Some crises are predictable, others not.

We saw earlier that we have a better understanding of the predictability of some types of crises than we used to. The key discovery is that financial crises appear to be predicted much better when one accounts for an interaction between two types of risks – credit and valuation risks. One could look at this literature as evidence on the importance of compound risk. If one treats the two risks separately, predictability is low. Skeptics might question whether this predictability is evidence of data mining and whether future crises might fail to be predictable using that literature. An argument against that view is that we have a body of economic theory that supports the predictability that is evidenced in the literature. It is a body of literature that builds on the insights of Kindleberger (1978) and Minsky (1977) that market participants can be too optimistic and that they eventually discover that they were too optimistic, which leads to a reversal in economic activity and a crisis. Another body of theory suggested explicitly that crises could not be predictable because if the market expected a crisis, it would happen immediately.

The crises that seem predictable are those that involve large credit expansions and high asset valuations. These are not the only crises. We saw that the worst economic crises are associated with wars. The most recent crisis was the COVID-19 crisis. It was of a type that was not predictable. We knew that the unconditional probability of a pandemic was extremely low and suddenly we had a massive crisis. Therefore, when it comes to crisis risk, we have to accept that there are types of crises that we understand well and that are predictable, but other types that are not predictable, that we may not even know of, and for which the concept of probability may have little meaning.

6.3. Stationarity.

A concern with the literature that relies on historical data is that it often assumes that the distribution of crisis risk in the historical data will also be the distribution of crisis risk in the future. For some types of crises, this seems a good starting point. However, the financial system changes all the time. The increasingly strong reactions of policymakers to crises may affect the risk of crises in the future. Statistical analysis is not a substitute for economic analysis, but is part of economic analysis.

The assumption that the distribution of a risk is stationary is highly problematic for climate risks. For instance, it is well-established that the probability of record-shattering climate extremes is increasing (Fischer, Sippel, and Knutti, 2021). At the same time, however, it is not clear how climate events can either cause the type of crises we have discussed or how they can make these types of crises worse. In a perverse sense, the economic and financial crises might be good for climate in that economic activity and hence carbon production falls. As seen for the U.S., climate events are unlikely to have major effects on the economy or the financial system unless they are very different from what we have observed historically. For instance, climate shocks that make large parts of a country unlivable would cause the type of catastrophes we have discussed. However, scaling up the magnitudes of existing events we have experienced is unlikely to cause economic or financial crises of the type we have discussed. While Tooze (2022) seems to use the term climate crisis to characterize a state of affairs that will last for decades, this nomenclature does not seem helpful in understanding crisis risk.

6.4. Mitigating crisis risk.

Risk management consists of identifying risks, assessing their importance, and then deciding on whether to take these risks and if so whether to mitigate them. Firms can take derivatives positions and purchase insurance contracts that will offset losses they experience. It makes sense for firms with diversified ownership to enter such financial positions when the losses could impose on them costs of financial distress or increase their cost of doing business. Taking financial positions to protect a firm against financial distress costs arising from crises would be expensive if it were at all feasible. Not all firms could enter such positions since in equilibrium we cannot all hedge with financial instruments against crisis risk. The firm would experience costs in most years. The present value of the costs of financial distress resulting from a crisis might be small given that crises arise infrequently. However, in addition, contracts might not be available and the risks a firm is exposed to may be hard to evaluate or even unkown. Instead, a firm may be better off organizing itself so that it is resilient when risks it is exposed to materialize. Firms have a number of

ways to be more resilient. They can be more financially flexible and more operationally flexible. Financial flexibility and operational flexibility give them more options to respond to adverse development.

Existing evidence of how firms were affected by the COVID-19 crisis is useful to understand the benefits of resilience. The COVID-19 shock impacted firms by reducing their income if they needed physical proximity to their consumers and by making it harder to produce and manage their operations if these activities required physical proximity of workers and managers. For some firms, the COVID-19 shock represents a large drop in revenues while their costs did not fall or fell little. This put them in a situation where they stopped being profitable. Firms could deal with this shock better if they had financial capacity they could use to keep funding their fixed costs. As a result, we would expect firms that had more financial flexibility to be less affected by the COVID-19 shock.

Fahlenbrach, Rageth, and Stulz (2022) examine how the performance of firms in March 2020 was affected by their financial flexibility. They investigate the stock price drop of firms during March 2020 before the announcement of the massive stimulus package on March 23, 2020. They use three proxies for financial flexibility, namely cash to assets, short-term debt to assets, and long-term debt to assets. They compare firms with low financial flexibility to firms with high financial flexibility. A firm with high (low) financial flexibility is at the 75th (25th) percentile of cash to assets and the 25th (75th) percentile of short-term debt to assets and long-term debt to assets. During the period of March 2020 they consider, the mean decrease in the value of common stock across firms is 37.8 percentage points. A firm with high financial flexibility experienced a drop in stock price that is 9.7 percentage points lower or 26% percentage smaller. They find that firms with greater financial flexibility also fared better during the global financial crisis. When they consider firms more sensitive to social distancing, they find that these firms experienced an additional benefit from financial flexibility compared to the average firm. Specifically, firms that were highly exposed to social distancing suffered a lower drop in their stock price of 9 percentage points if they were at the 75th percentile of the cash to assets distribution instead of the 25th percentile of that distribution.

Barry, Campello, Graham, and Ma (2022) investigate how the COVID-19 shock affected firms' business plans depending on how flexible these were financially and operationally. They find that firms

that had more workplace flexibility were better able to cope with the shock, and so were firms with more investment flexibility. Firms with greater financial flexibility saw their employment levels and investment levels less affected by the COVID-19 shock.

The literature we have described focuses on operating companies as opposed to financial companies. However, financial companies also can choose to be in a better position to respond to a crisis by having financial flexibility. For instance, they can have a diversified funding base as well as having excess capital. The CEO of JPMorgan Chase, Jamie Dimon has often talked about his goal for the bank to have a fortress balance sheet. In his list of "JPMorgan Chase Principles and Strategies," the third item is "We will maintain a fortress balance sheet." He further explains that "It is capital and liquidity combined with strong earnings and margins that provide the ability to withstand extreme stress."[12]

Though flexibility means that a firm is better able to cope with the unexpected, flexibility is not free. Firms face a tradeoff between having more flexibility to be better equipped to react to crises and being more profitable in normal times. How a firm will resolve this tradeoff will depend on a number of factors, but the likelihood of experiencing a crisis would be an important consideration. If crises are extremely rare, firms may not find it worthwhile to keep flexibility to cope with them. However, if the probability of experiencing a crisis increases, the value of flexibility increases as well.

Good corporate governance seems particularly important when firms invest in resilience. If investing in resilience is valuable for shareholders, it is still likely to mean that short-term performance will be weaker than if the firm did not invest in resilience. In such a situation, it would be important for the board to resist short-termism pressures. Flexibility also means that a firm has slack, which means that it has resources that could be wasted. Again, good governance would make it less likely that resources would be wasted.

[12] See Chairman and CEO letter to shareholders, Annual Report 2018, for the quotes.

7. Conclusion.

Since the GFC, economists have made progress in collecting data on economic and financial crises and in using that data to study whether crises are predictable. The evidence does not support the view implied by some uses of the polycrisis concept that there is a higher frequency of economic and financial crises. We also considered political crises and natural disaster events. Political crises have not become more frequent, but there is evidence that political crises make economic crises more likely. The evidence on the predictability of crises shows that it is important to consider the interaction of risks. Specifically, neither credit growth nor high valuations alone help predict crises much, but the interaction of these two variables makes a large difference in predicting crises. However, not all crises are caused by the combination of credit growth and high valuations. Historically, the worst economic crises are associated with wars. There is no evidence that crises not associated with the combination of credit growth and high valuations can be predicted. The most recent economic crises in the data are associated with the COVID-19 pandemic and there is no evidence of predictability for crises of this type.

Our current understanding of crises makes it clear that not all crises are predictable and that it would be unreasonable to believe that we know all types of crises that could occur. With this understanding, it is not the case that risk managers can catalogue all risks according to their frequency and severity. As we saw, in some cases the frequency and severity depend on other risks or may be highly uncertain. More research on interactions between non-economic risks and financial and economic risks is called for to improve our understanding of these issues. However, the existing evidence makes it important for firms to think of measures they can take to protect themselves in the event of unforeseen shocks. With such shocks, there is value to firms to invest in resilience. Evidence from the COVID-19 crisis shows that firms that are more operationally and financially flexible are better able to cope with unforeseen events. Unfortunately, resilience is not free. It can decrease the performance of firms in good times or leave them with slack that could be misused. Good corporate governance can decrease the costs of resilience.

References

Acemoglu, Daron, Simon Johnson, James Robinson, and Yunyong Thaicharoen. "Institutional causes, macroeconomic symptoms: volatility, crises and growth." *Journal of monetary economics* 50, no. 1 (2003): 49-123.

Acharya, Viral, Itamar Drechsler, and Philipp Schnabl. "A pyrrhic victory? Bank bailouts and sovereign credit risk." *The Journal of Finance* 69, no. 6 (2014): 2689-2739.

Bae, Kee-Hong, G. Andrew Karolyi, and Reneé M. Stulz. "A new approach to measuring financial contagion." *The Review of Financial Studies* 16, no. 3 (2003): 717-763.

Barry, John W., Murillo Campello, John R. Graham, and Yueran Ma. "Corporate flexibility in a time of crisis." *Journal of Financial Economics* 144, no. 3 (2022): 780-806.

Blickle, Kristian, Sarah Ngo Hamerling, and Donald P. Morgan. "How Bad Are Weather Disasters for Banks?." *FRB of New York Staff Report* 990 (2021).

Baron, Matthew, Emil Verner, and Wei Xiong. "Banking crises without panics." *The Quarterly Journal of Economics* 136, no. 1 (2021): 51-113.

Baron, Matthew, Wei Xiong, and Zhijiang Ye. "Quantifying Dark Matter in Finance: Time-Varying Disaster Risk and Asset Prices." Working Paper, Cornell University, 2023.

Barro, Robert J., and José F. Ursúa. "Macroeconomic crises since 1870." Working paper, National Bureau of Economic Research, 2008.

Barro, Robert J., and José F. Ursúa. "Rare macroeconomic disasters." *Annu. Rev. Econ.* 4, no. 1 (2012): 83-109.

Berkman, Henk, Ben Jacobsen, and John B. Lee. "Time-varying rare disaster risk and stock returns." *Journal of Financial Economics* 101, no. 2 (2011): 313-332.

Boustan, Leah Platt, Matthew E. Kahn, Paul W. Rhode, and Maria Lucia Yanguas. "The effect of natural disasters on economic activity in US counties: A century of data." *Journal of Urban Economics* 118 (2020): 103257.

Brecher, Michael, and Jonathan Wilkenfeld. *A study of Crisis*. University of Michigan Press, 1997.

Bruner, Robert F., and Sean D. Carr. *The Panic of 1907: Heralding a new era of finance, capitalism, and democracy*. Wiley, 2023.

Claessens, Stijn, and M. Ayhan Kose. Financial crises: Explanations, types, and implications. IMF, 2012.

Ćorić, Bruno. "Economic disasters: a new data set." *Finance Research Letters* 39 (2021): 101612.

Demirgüç-Kunt, Asli, and Enrica Detragiache. "The determinants of banking crises in developing and developed countries." *Staff Papers* 45, no. 1 (1998): 81-109.

Dornbusch, Rudiger, Yung Chul Park, and Stijn Claessens. "Contagion: Understanding how it spreads." *The World Bank Research Observer* 15, no. 2 (2000): 177-197.

Fahlenbrach, Rüdiger, Kevin Rageth, and René M. Stulz. “How valuable is financial flexibility when revenue stops? Evidence from the COVID-19 crisis.” *The Review of Financial Studies* 34, no. 11 (2021): 5474-5521.

Fischer, E. M., Sebastian Sippel, and Reto Knutti. "Increasing probability of record-shattering climate extremes." *Nature Climate Change* 11, no. 8 (2021): 689-695.

Forbes, Kristin. "The" Big C": identifying contagion." NBER working paper (2012).

Funke, Manuel, Moritz Schularick, and Christoph Trebesch. “Going to extremes: Politics after financial crises, 1870–2014.” *European Economic Review* 88 (2016): 227-260.

Funke, Manuel, Moritz Schularick, and Christoph Trebesch. “Populist leaders and the economy.” *American Economic Review* (2021), forthcoming.

Gennaioli, Nicola, Andrei Shleifer, and Robert Vishny. "Neglected risks: The psychology of financial crises." *American Economic Review* 105, no. 5 (2015): 310-314.

Greenwood, Robin, Samuel G. Hanson, Andrei Shleifer, and Jakob Ahm Sørensen. "Predictable financial crises." *The Journal of Finance* 77, no. 2 (2022): 863-921.

Jordà, Òscar, Moritz Schularick, and Alan M. Taylor. "When credit bites back." *Journal of money, credit and banking* 45, no. s2 (2013): 3-28.

Jordà, Òscar, Katharina Knoll, Dmitry Kuvshinov, Moritz Schularick, and Alan M. Taylor. “The rate of return on everything, 1870-2015.” *Quarterly Journal of Economics* 134, no. 3 (2019): 1225-1298.

Jordà, Òscar, Moritz Schularick, and Alan M. Taylor. “Disasters everywhere: The costs of business cycles reconsidered.” Working paper, National Bureau of Economic Research, 2020.

Kaminsky, Graciela L., and Carmen M. Reinhart. "The twin crises: the causes of banking and balance-of-payments problems." *American economic review* 89, no. 3 (1999): 473-500.

Karolyi, G. Andrew. "Does international financial contagion really exist?." *International Finance* 6, no. 2 (2003): 179-199.

Kay, John, and Mervyn King. *Radical uncertainty: Decision-making beyond the numbers*. W.W. Norton & Company, 2020.

Kindleberger, Charles P. *Manias, panics, and crashes: A history of financial crises*. Basic Books, 1978.

Kuritzkes, Andrew, and Til Schuermann. “What we know, don’t know, and can’t know about bank risk.” In *The Known, the Unknown, and the Unknowable in Financial Risk Management*, ed. by Richard J. Herring, Neil A. Doherty, and Francis X. Diebold, Princeton University Press, 2010.

Laeven, Luc, and Fabian Valencia. "Systemic Banking Crises Database II." *IMF Economic Review* 68, no. 2 (2020): 307-361.

Lawrence, Michael, Scott Janzwood, and Thomas Homer-Dixon. "What is a Global Polycrisis? And How Is It Different From A Systemic Risk?" *Cascade Institute Discussion Paper* (2022).

López-Salido, David, Jeremy C. Stein, and Egon Zakrajšek. "Credit-market sentiment and the business cycle." *The Quarterly Journal of Economics* 132, no. 3 (2017): 1373-1426.

Marfè, Roberto, and Julien Pénasse. “Measuring macroeconomic tail risk.” Working paper, Collegio Carlo Alberto, 2023.

Metrick, Andrew, and Paul Schmelzing. “Banking-crisis interventions, 1257-2019.” Working paper, 2021.

Minsky, Hyman P. A theory of systemic fragility, pp. 138-152, in E.D. Altman and A.W. Sametz, Eds. *Financial crises: Institutions and markets in a fragile environment.* John Wiley and Sons, 1977.

Morin, Edgar, and Anne Brigitte Kern. *Homeland earth*. Cresskill (1999).

Odell, Kerry A., and Marc D. Weidenmier. "Real shock, monetary aftershock: The 1906 San Francisco earthquake and the panic of 1907." *The Journal of Economic History* 64, no. 4 (2004): 1002-1027.

Pritsker, Matt. “The channels of financial contagion.” 2000.

Ritter, Jay R. "Is Economic Growth Good for Investors?" *Journal of Applied Corporate Finance* 24, no. 3 (2012): 8-18.

Schuermann, Til. "Capital adequacy pre-and postcrisis and the role of stress testing." *Journal of Money, Credit and Banking* 52, no. S1 (2020): 87-105.

Sufi, Amir, and Alan M. Taylor. "Financial crises: A survey." Working paper, National Bureau of Economic Research, 2021.

Tooze, Adam. *Shutdown: How Covid shook the world’s economy*. Viking, 2021.

Tooze, Adam. “Crisis Pictures (Krisenbilder) – mapping the polycrisis.” Chartbook #73, 2022.

World Economic Forum, *The Global Risks Report 2023*.

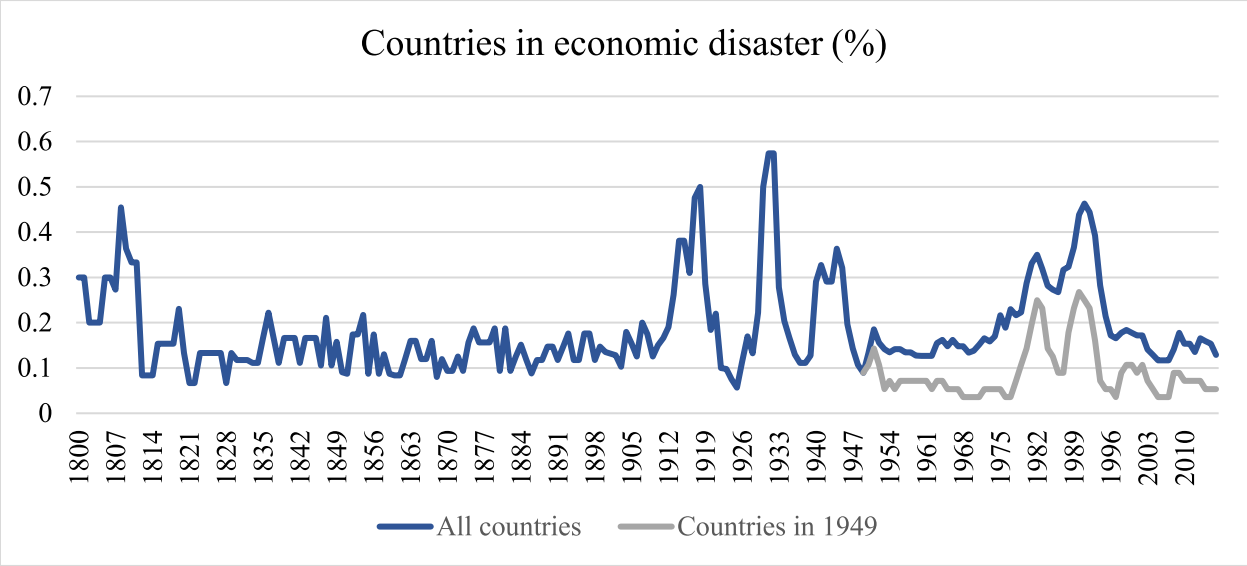

Figure 1. Percentage of countries in economic disaster.

This figure uses the data from Ćorić (2021) that defines an economic disaster as a drop in real GDP of 10% or more. We thank Bruno Ćorić for providing the data. The figure shows results for all countries in the sample each year. It also shows results for the countries in the sample in 1949 for subsequent years.

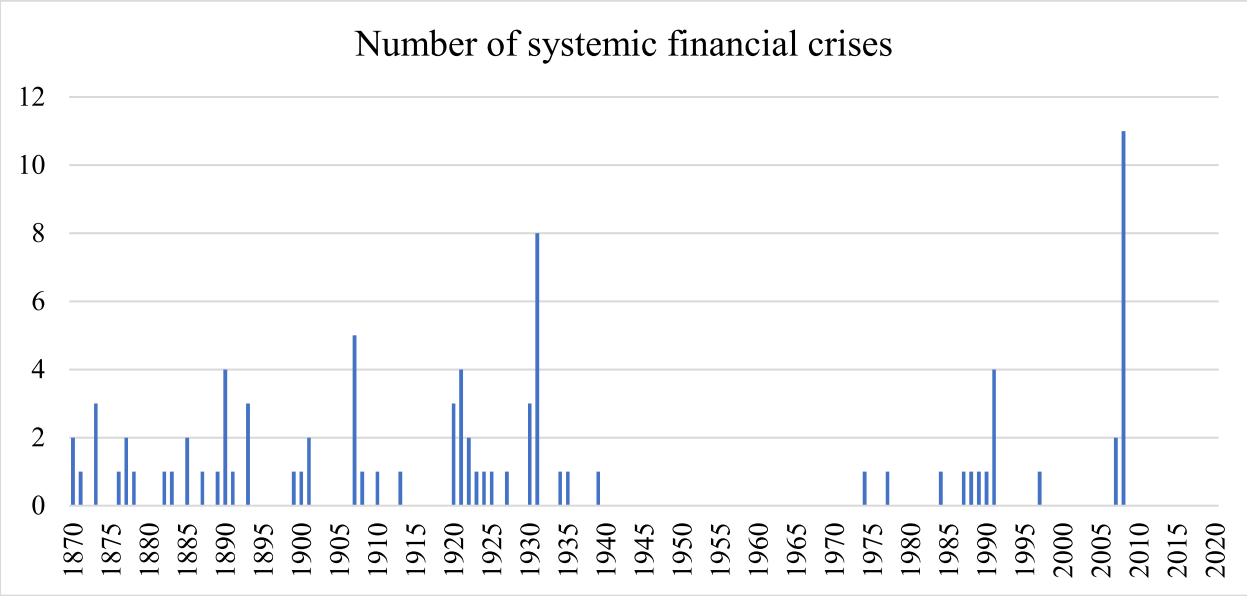

Figure 2. Number of systemic financial crises.

This figure uses the data for eighteen countries from 1870 to 2020 from Jordà, Knoll, Kuvshinov, Schularick, and Taylor (2019) that indicates whether a country is in a systemic crisis for each sample year. The current version of the data is available at https://www.macrohistory.net/database/ .

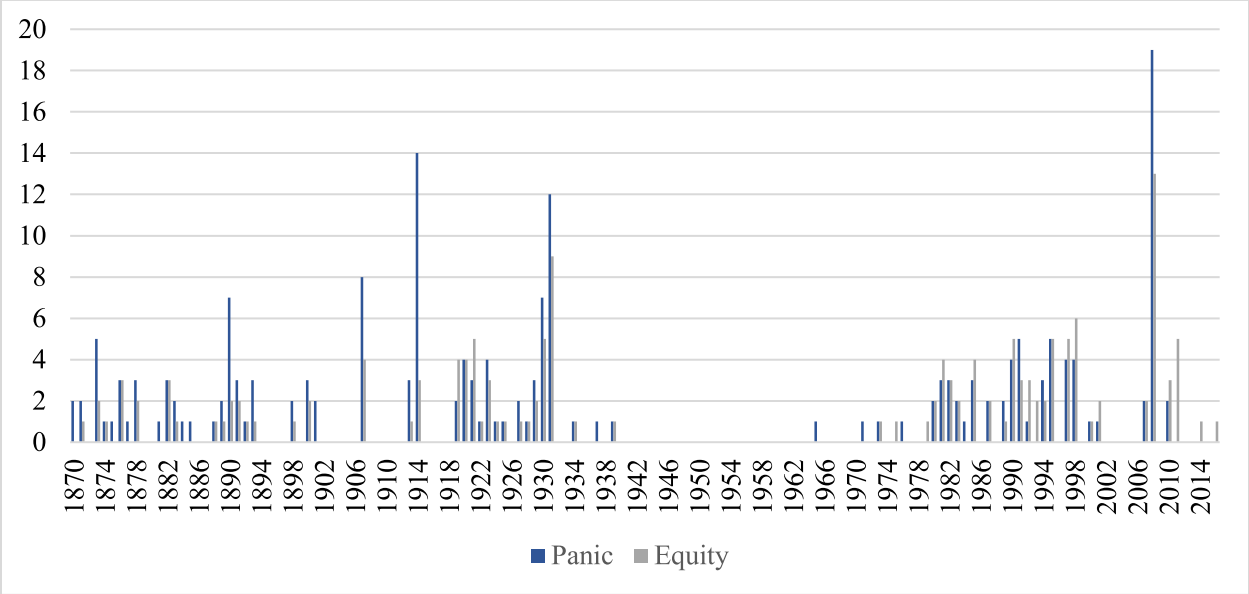

Figure 3. Percentage of countries that experience the start of a panic of equity financial crisis each year.

This figure uses the dataset of Baron, Verner, and Xiong (2021). The dataset covers 46 countries from 1870 to 2016. A financial crisis is defined as an event with a banking panic or a drop in bank stock prices of 30% or more together with widespread bank failures.

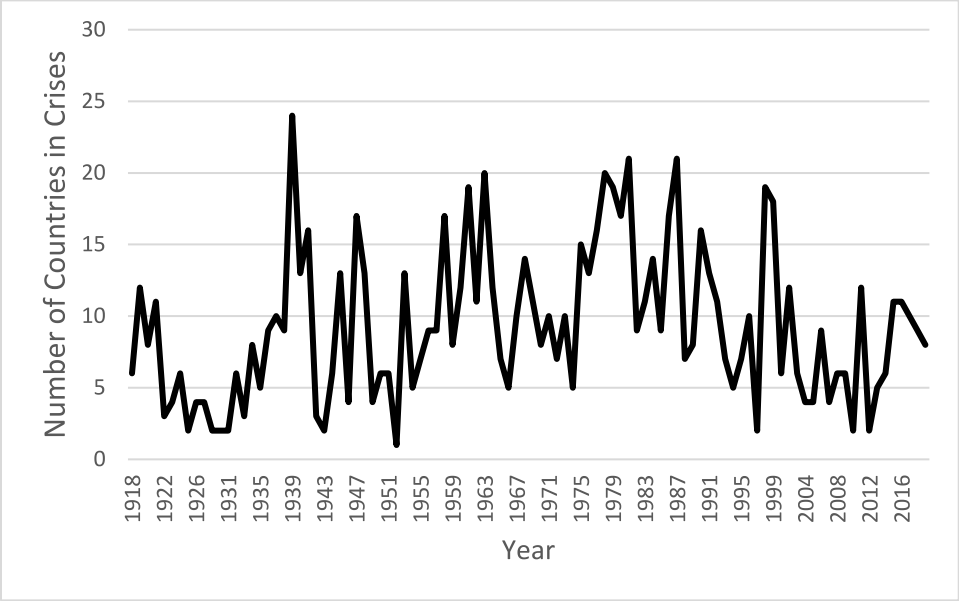

Figure 4. Number of countries in a political crisis.

This figure uses the data from the International Crisis Behavior project and plots the number of countries in a political crisis each year. The data is available at https://sites.duke.edu/icbdata/ .