

Financial stability

Cloud outsourcing and financial stability risks

Contact: alexander.harris@esma.europa.eu¹⁴⁵

Summary

The growing use of cloud service providers (CSPs) by financial institutions can provide benefits to individual firms and the financial system. However, high concentration in CSPs could create financial stability risks if an outage in a CSP affects many of its clients, increasing the likelihood of simultaneous outages. Analysis using a stylised model calibrated with operational risk data suggests that CSPs need to be significantly more resilient than firms to improve the safety of the financial system. In financial settings where only longer (multi period) outages cause systemic costs, the results suggest that CSPs can best address systemic risks by strongly reducing incident resolution times, rather than incident frequency. In the model, using a back-up CSP successfully mitigates the systemic risk caused by CSPs. Backup requirements may need to be mandated however, as the systemic risk is an externality to individual firms. Finally, there is a clear need for detailed data on outages by financial institutions and CSPs.

Introduction

The use of cloud services by financial institutions has risen in recent years, as firms are increasingly outsourcing parts of their IT infrastructure. Cloud computing is an innovation that allows for the use of an online network ('the cloud') of hosting processors to increase the scale and flexibility of computing capacity (FSB, 2019).

While cloud computing is still a topic of research, it has become key to the digital economy. The use of cloud has significantly increased in the last few years (RA.1), a trend which has been further accelerated by the COVID-19 pandemic, as firms have had to set up remote working facilities.

There are many benefits associated to using cloud computing in the financial system. Cloud technology can help firms reduce the costs of developing and maintaining IT systems, as

financial services firms seldom have the scale and capacity to set up such infrastructures.

RA.1
Percentage of EU firms purchasing cloud services
Firms increasingly purchasing cloud services



Note: Percentage of businesses purchasing cloud computing services by year in 22 EU countries, %. Countries included: AT, BE, CZ, DE, DK, EE, ES, FI, FR, GR, HU, IE, IT, LV, LT, LU, NL, PL, PT, SI, SK, SE. Firms across the economy with at least 10 employees were surveyed.
Sources: OECD, ESMA


¹⁴⁵ This article was written by Carolina Asensio, Antoine Bouveret and Alexander Harris. It summarises a more detailed analysis and discussion by Asensio, Bouveret and Harris (2021, forthcoming).

Likewise, CSPs can also increase the resilience of financial institutions as they invest heavily in security and spread their infrastructures across geographical areas.

Cloud computing can also help firms expedite and scale up processes, increase flexibility and operational efficiency, and enhance their ability to identify business opportunities and revenue streams. Another key benefit is risk mitigation through enhanced information security and disaster recovery plans, given that the cloud can provide efficient solutions to mitigate traditional technology risks, such as capacity, redundancy, and resiliency concerns. Equally, cloud migration plays a huge role in enabling the use of other innovative technologies such as AI, big data and DLT.

But while migrating to the cloud provides a range of benefits to firms, it can also raise challenges at firm level in terms of governance, data protection and information security. Operational risks are also relevant, as they result from inadequacies or failures of internal processes, people, and systems, or from external events, and they may impact financial institutions in different ways. For instance, data losses could happen due to failures, deletion or disasters that occur at CSPs, or when CSPs outsource some of their functions to third parties, or 'fourth parties'. Cyber risk is also important to consider, as massive amounts of data are stored in cloud ecosystems. 'Vendor lock-in' is also relevant when financial institutions rely strongly on the services of one CSP.

RA.2
Global market share of cloud infrastructure services
High concentration in CSP market



Note: Global market share of cloud infrastructure services in Q2 2020, by vendor
Sources: Synergy Research Group

In addition, the cloud can bring risks at the level of the wider financial system. Given that a limited amount of CSPs can meet the high standards of resiliency requirements that financial institutions

demand, there is high concentration in the provision of cloud services within the financial sector (RA.2). In this context, it is plausible that a sufficiently large number of financial institutions become dependent on a small number of CSPs, meaning that operational incidents may become more correlated. Concentration risk in this context is thus a form of systemic risk.

A model of concentration risk

We introduce a risk model to investigate the conditions under which outsourcing to the cloud by financial sector firms may generate systemic operational risk as in Asensio, Bouveret and Harris (2021).

Existing literature

The increasing use of CSPs has been accompanied by emerging literature on the risks and potential impact of CSP outages.

A series of studies estimate the costs related to an outage of cloud providers. Using scenario analysis, Lloyd's estimates global losses ranging from USD 4 bn to USD 53 bn for an outage duration of between 0.5 and 3 days (Lloyd's, 2017), and losses for the largest US firms (corporates and financials) at around USD 10 bn for an outage of the top three CSPs lasting between 3 and 6 days (Lloyd's, 2018).

Using a Value-at-Risk approach, Naldi (2017) provides a measure of potential losses for CSPs, based on outage data and estimated loss per minute. The author models outage frequency using a Poisson distribution and outage duration using a Generalised Pareto Distribution, frequently used to model fat tails in operational risk (Bouveret, 2019). Our model builds on this approach, distinguishing between outage frequency and duration. For tractability, and to prevent time-consistency (i.e. time-overlapping outages for a single area of a firm's operations), we do so in a two-state Markov chain framework. This allows us to analyse alternative technology-based approaches to mitigating systemic risk: preventing outages versus quickly resolving them.

A related strand of the literature examines the impact of using CSPs on the cost of cyber events for individual firms. Using a large dataset of cyber losses, Aldasoro et al. (2020) find that a higher dependence on CSPs, measured by investment in cloud services at country-level, is associated

with lower costs. However, the authors note that this result might not apply to more extreme events since they only have small losses in their database. Harmon, Vytelingum and Babaie-Harmon (2020) put forward an agent-based model with banks and CSPs in a settlement context. CSPs can face outages, the duration of which is assumed to follow an exponential distribution. When a CSP suffers an outage, banks using the CSP cannot proceed with settlement, creating credit risk. The authors estimate the impact on other banks in the network, using contagion measures based on market-based data for banks (Demirer et al., 2018).

Main features of the model

The model considers a set of financial sector firms in three main scenarios:

1. A setting where no cloud outsourcing is available (the 'no-cloud scenario');
2. A setting where each financial sector firm outsources the time-critical IT service to one of several CSPs (the 'cloud scenario'); and
3. A setting where each financial sector firm outsources the time-critical IT service to a primary CSP and to a secondary provider (the 'multi-cloud scenario').

The risk model does not explicitly consider the firms' decision on whether to outsource to the cloud. Instead, it focuses on the risk implications of the different scenarios. However, the model can readily be understood in a strategic context. Firms will have an incentive to move operations to the cloud – other things being equal and neglecting frictional costs – if cloud outsourcing prevents incidents or improves their resolution speed.¹⁴⁶

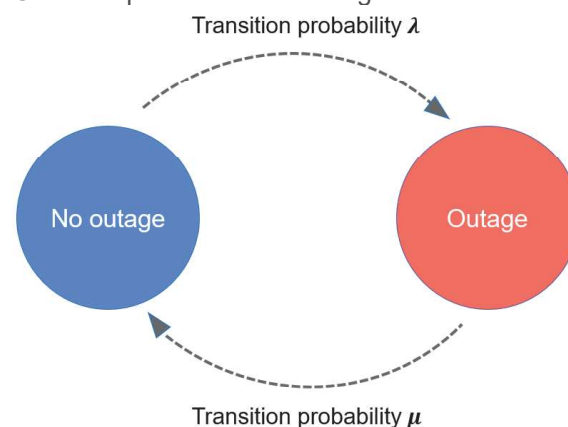
The model considers a set of firms over discrete time periods. In any time period, each firm is in one of two states: *outage* or *no outage*. A firm in an outage state in one period will resolve the outage (i.e. transition to the no outage state) in the next with a constant probability. Conversely,

a firm in a no outage state in one period will experience an outage in the next period (i.e. transition to the outage state) with another constant probability. Importantly, outages are assumed to arise independently across firms.¹⁴⁷

This arrangement is known as a Markov chain. Regardless of the system's initial configuration, it has long-run *steady state* properties that we can study. For example, given the transition probabilities we can calculate the average amount of time a firm spends in an outage, the average amount of time that two or more firms are in simultaneous outage, and the frequency with which a firm suffers a multi-period outage of a given duration.

In scenario 1, where firms do not outsource to the cloud, the per-period probability of suffering a new outage is denoted λ , known as the *incident rate*. The per-period probability that an outage is resolved is denoted μ and known as the *repair rate* (RA.3).

RA.3
Markov chain diagram for a firm in scenario 1
Constant probabilities of outage and resolution



Note: Markov chain diagram for a single firm in the no-cloud baseline scenario, in which possible states of the firm are represented by coloured circles.

Given these transition probabilities, the average time a firm spends in outage, τ , can be calculated as follows.

¹⁴⁶ Asensio, Bouveret and Harris (forthcoming) examine these incentives formally. A finding is that even if firms find it optimal to migrate to the cloud (scenario 2), they may not find it individually optimal to use a back-up cloud provider (scenario 3). This can happen even if the system would be more efficient if all firms were to back-up. In short, there is a potential externality that may warrant policy intervention.

¹⁴⁷ Independence can to some extent be justified by interpreting the model as a means to study the difference

in systemic risk between scenarios 1 and 2, abstracting away from those risk drivers that are common to both settings. For instance, to the extent the two scenarios face a common risk of a multi-firm malicious attack – which can be perpetrated directly against the firms or via the cloud – we can regard the effect as 'cancelling out' between the scenarios. However, the independence assumption clearly reduces baseline systemic risk in scenario 1, which therefore overstates the extent to which CSPs create additional systemic risk via concentration.

$$\tau = \frac{\lambda}{\lambda + \mu} \quad (1)$$

In scenario 2, firms outsource to the cloud. For a cloud provider, the per-period probabilities of new outages and of resolving existing outages are denoted λ' and μ' respectively. The average time in outage is denoted τ' and calculated analogously to equation (1). In scenario 2, the firms are assigned to a small number of CSPs, which each have an equal market share. If a CSP suffers an outage, we assume that all its client firms will suffer an outage at the same time.¹⁴⁸ In scenario 2, λ' and μ' therefore also represent the transition probabilities for any given firm.

As noted in the introduction the services offered by CSPs may bring a range of benefits as specialist technology providers to client firms, including enhanced operational resilience. This can be represented in the model via the following equation.

$$\tau \geq \tau' \quad (2)$$

Inequality (2) says that in the model, CSPs (and their client firms) have lower average outage time than firms in the no-cloud scenario. A key finding of the illustrative results of the model is that despite assuming this improved resilience for individual firms in the cloud scenario compared to the no-cloud baseline, the former may nonetheless create systemic operational risk. This is due to the assumption that outages in the cloud scenario are correlated, unlike in the no-cloud baseline where they are realised independently. Inequality (2) is consistent with an equilibrium framework in which all firms find it optimal to outsource to the cloud. It is also in line with the calibration data presented below, where we consider an illustrative application of the model to securities markets.

Applications

The simple, stylised nature of the model makes it versatile. It can be applied to any setting in which costs of simultaneous outages among several firms are greater than if the outages were separate. This is likely to be the case especially where a financial system relies on transactions between a relatively small number of

counterparties, such as in the banking system. Two possible applications to securities are as follows.

Clearing Members of Central Counterparties

Within financial market infrastructures, the clearing members that allow Central Counterparties (CCPs) to function constitute a possible real-world application of the model.

If clearing members outsource core services, and one or more CSPs suffer an outage, the impact on the financial system could be substantial. First, the failure of some clearing members to post collateral would lead to the liquidation of their positions according to the default management rules used by CCPs, entailing potential losses due to fire sales and the consumption of some of the resources in the default fund. In addition, outages affecting clearing members could prevent some of their clients from clearing transactions with them. This, in turn, could result in additional costs – either in the form of frictional costs incurred by clients switching to other clearing members (where possible) or, worse, the cancellation of transactions where clearing cannot be executed. In its 2020 stress test, ESMA estimated that the failure of the two largest counterparties to a CCP could lead to losses of around EUR 1 bn each for the two largest EU CCPs (ESMA, 2020).

CCPs might not have visibility to assess the concentration risks related to cloud outsourcing by the clearing members.

Primary dealers and market makers

The model could also be applied elsewhere in financial markets. For example, in sovereign bond markets, primary dealers play an important role not only at the issuance stage, but also by providing market making services in secondary markets. While each country has different rules, primary dealers are usually required to support the liquidity in sovereign markets (AFME, 2020). If a set of primary dealers were unable to operate due to CSP outage, secondary market liquidity would be significantly reduced.¹⁴⁹ Similar effects could also occur in equity markets, although the

¹⁴⁸ This assumption is a simplification and does not reflect the fact that some outages may be local, rather than global.

¹⁴⁹ In a different context, Bouveret et al. (2021) document how liquidity deteriorated on the Italian sovereign bond market on May 29, 2018, when primary dealers retrenched from quoting bonds on the MTS interdealer platform.

fragmentation of trading across venues and the diversity of market makers might mitigate the impact of an outage affecting a few institutions.

Example calibration

An example calibration using public data suggests that cloud outsourcing (scenario 2) may introduce systemic risk into securities markets compared with the no-cloud baseline (scenario 1). In particular, we consider the first of the applications described above, namely clearing members of CCPs. We set the transition probabilities in scenario 1 for clearing members using available public data¹⁵⁰, and likewise for the transition probabilities for the clearing members (via CSPs) in scenario 2. For this example calibration, we set the duration of each period at one hour.

To investigate systemic risk, we established the following condition:

- For a systemic event to occur, at least 3 clearing members must be simultaneously unable to operate¹⁵¹.

The intuitive assumption is that if large clearing members or a multitude of smaller ones are disrupted, then the CCP will be unable to operate in an orderly manner since several counterparties would be unable to post and receive margins.

This requirement is stricter than the one used for CCP stress tests, where CCPs should be able to withstand their two largest CMs defaulting simultaneously. However, in our model and application we only focus on the number of firms suffering an outage irrespective of their size. Setting a 3-firm minimum requirement for a systemic event is intended to counterbalance this effect.¹⁵²

Regarding the duration of the outages, the Principles for Financial Market Infrastructures (FMIs) put forward by CPMI and IOSCO explicitly specify that FMIs should have a business continuity plan that ensures that critical IT

systems are able to resume two hours after a disruptive event (CPMI-IOSCO, 2012).

RA.4

Illustrative example calibration of cloud outage model
Parameter values for clearing members and CSPs based on public data

Parameter	Interpretation	Value
n	Number of firms	20
n'	Number of CSPs	5
S	Minimum number of firms in simultaneous outage for systemic event	4
λ	Hourly probability of new outage ('incident rate') in no-cloud baseline	0.18%
μ	Per-period probability that an outage is resolved ('repair rate') in no-cloud baseline	78%
λ'	Per-period probability of new outage ('incident rate') in cloud scenario	0.056%
μ'	Per-period probability that an outage is resolved ('repair rate') in cloud scenario	24%

Note: λ , μ estimated as exponential decay parameters using CCP outage data as a proxy for clearing member outages. CCP outage from 10 CCPs for 2016-2020. λ' , μ' estimated as exponential decay parameters using data on outages and average duration of outages reported by Google Cloud for 2016-2020, taking averages across 16 different service areas. Observations that reported zero outages have been excluded from the analysis.

Sources: 10 CCPs (CME, DTCC, Eurex, ICC_CDS, ICE NGX, ICEU, ICUS_F&O, JSCC OTC-JGB, LCH.Clearnet.Ltd, LCH.Clearnet.SA). CSP parameter estimates: Google Cloud.

The longer the duration of the outage, the higher the probability that the event will be systemic. Any event that prevents or impairs end-of-day settlements could then be considered systemic (Brauchle, Göbel and Seiler, 2020).

We consider 3 different minimum-time conditions for systemic events to occur:

- A 1-hour condition, i.e. whenever 3 firms are in a simultaneous outage, a systemic event occurs.

¹⁵⁰ Data on operational risks for clearing members are not available. Instead, we use the quarterly quantitative disclosures by CCPs which provide information about the number of outages over the last 12 months and the total duration of the outages. For CSPs, we use publicly available data from one CSP.

¹⁵¹ We assume that authorities and CCPs do not react to the outages. However, it is likely that if such event were to occur, they would use back-up procedures (including manual transfer or margins) to mitigate risks for clearing services.

¹⁵² An extension of the model where systemic events are defined based on size could be analysed in future work.

- A 2-hour condition, in line with the CPMI-IOSCO target.
- An 8-hour condition. This reflects the fact that clearing is on a T+1 basis, and 8 hours is the approximate length of a trading day.

Comparing the different results gives insight into the role played by the recovery rate parameter μ' in mitigating systemic risk.

Illustrative results

Given the parameter values and the definition of a systemic event in the present application, we can investigate under what conditions scenario 2 introduces systemic risk compared with scenario 1, and whether these conditions are likely to be met in practice. To do this, we first define the *odds ratio*, R , as follows.

$$R = \frac{\text{Prob} [\text{Systemic event in scenario 2}]}{\text{Prob} [\text{Systemic event in scenario 1}]} \quad (4)$$

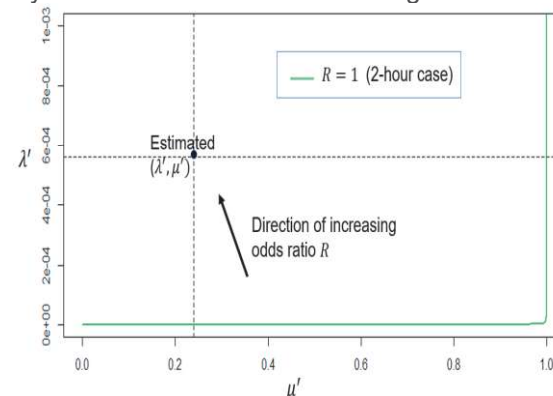
The odds ratio describes how many times more likely a simultaneous outage constituting a systemic event is in scenario 2 than in scenario 1. Intuitively, it describes how much more likely such an outage is made by concentration risk due to cloud outsourcing. If $R > 1$, then systemic risk is higher in the presence of cloud outsourcing, according to our stylised model, given the assumptions made and the calibration.

Using the parameter values for λ , λ' , μ and μ' yields the solutions for $R = 1$ (RA.5). The green line gives $R = 1$ under the specification that a systemic event requires the same 3 firms to have a simultaneous outage for at least 2 hours. The purple line gives $R = 1$ on the assumption that a systemic event simply requires the same 3 firms to be in a simultaneous outage.

The purpose of the analysis is not to provide accurate point estimates of the relative risk of systemic events between the two scenarios, given the limitations in the data discussed above and the stylized features of the model such as independence of outages across firms (Assumption 1) and the specification that an outage affecting 3 firms is the threshold for a systemic event. However, the results provide a useful framework for further analysis.

RA.5

Estimated incident and repair rate for cloud outsourcing compared with solutions for $R = 1$
Systemic risk arises in outsourcing scenario



Note: The lines plot values of cloud incident rate λ' and cloud repair rate μ' , expressed as per-hour quantities, for which systemic events have the same probability in the no-cloud baseline and the cloud scenario, given the parameter estimates for λ and μ based on CCP outage data. A systemic event occurs whenever the same 3 firms are out simultaneously for at least 2 hours. The y-axis is truncated at $\lambda' = 0.1\%$ for clarity.

The precise parameter values of CSP outage probability λ' and recovery probability μ' that we infer from the available data (using the assumptions discussed above) are approximate estimates only. Nonetheless, as order-of-magnitude estimates they appear to be *plausible*, in that they are close to the target values adopted by the CSP in question. Given that these plausible values of (λ', μ') lie far above the risk-equalization ($R = 1$) lines, we conclude that $R > 1$ in the present application. In other words, given the available data, our model suggests that outsourcing of core services by clearing members could create a new source of systemic risk, through simultaneous operational outages.

Consequently, as financial sector firms outsource to the cloud for core functions, policymakers should investigate the possibility of additional systemic risk arising. They can do this by:

- seeking and collecting more comprehensive data on outages by clearing members, or by other firms for whom simultaneous outages may have systemic effects; and
- investigating the extent to which the modelling assumptions hold true in practice and adjusting the modelling accordingly

The results (RA.3) indicate that in the most time-critical applications – where two hours of simultaneous outage represents a systemic

event – then there is a non-linear trade-off between the cloud incident rate and cloud repair rate in equalising risk with the no-cloud baseline.¹⁵³

So far, the analysis has only considered the 1 2-hour minimum time threshold for a systemic event. However, it could be argued that the systemic effects of an outage are less time-critical than that. For instance, we could instead assume that CCP outages only have systemic effects after 8 trading hours, given the T+1 clearing cycle. Using an 8-hour minimum makes the probability of a systemic event in the no-cloud baseline vanishingly small in our model for the parameter estimates based on CCP outage data. The implied probability of λ' for $R = 1$ would accordingly be vanishingly small – in effect requiring CSPs to prevent outages with perfect reliability.

In summary, where systemic events occur only after extended periods of simultaneous outages among firms, our modelling suggests that CSPs would need *perfect* service availability so as not to introduce additional systemic risk compared to the no-cloud baseline. Achieving equality of systemic risk with the no-cloud baseline (the $R = 1$ line in RA.5 and RA.6) is therefore effectively unattainable for CSPs in the case of an 8-hour minimum for systemic events. This finding illustrates certain limitations with the modelling, however:

- Policymakers may wish to tolerate more than the level of vanishingly small risk implied by the no-cloud baseline, given other benefits of the cloud computing paradigm.
- The no-cloud baseline risk is based on simplifying assumptions, as set out above.
- The CCP outage data may not provide a true guide to firm-level outage duration. One issue is that the data only report only total outage duration per firm per quarter, rather than the length of each outage. This makes it hard to test the goodness-of-fit of the geometric decay implied by our modelling (as opposed to a fat-tailed distribution). In particular, the data do not identify the number of day-long outages among CCPs.


One way to address these limitations is to consider the values of λ' and μ' that are required to achieve a less extreme mitigation of systemic risk, while retaining the 8-hour minimum for

systemic events. This can be done by plotting the $R = 1$ line while specifying that the repair rate in the no-cloud baseline is now equal to that implied by the CSP data (RA.6). In other words, we now set $\mu = 24\%$, rather than $\mu = 78\%$. The hourly probability of systemic risk in the no-cloud baseline is now around 1 in 10,000, or roughly one systemic event every 5 years.

With this more modest target for systemic risk, our model indicates that CCPs still have a greater risk of a simultaneous outage of one hour, but a greater risk of a simultaneous outage of 8 hours, i.e. a systemic event. Finally, the scenario where a systemic event is defined simply as occurring after 1-hour is included for comparison.

RA.6

Estimated incident and repair rate for cloud outsourcing compared with solutions for $R = 1$
Systemic risk arises in outsourcing scenario



Note: The lines plot values of cloud incident rate λ' and cloud repair rate μ' , expressed as per-hour quantities, for which systemic events have the same probability in the no-cloud baseline and the cloud scenario, given the parameter estimate for λ based on CCP outage data but a lower estimate of $\mu = 24\%$ (equal to that inferred from CSP data). Systemic event occurs whenever the same 3 firms are out simultaneously for at least 8 hours. The y-axis is truncated at $\lambda = 0.1\%$ for clarity.

The results in RA.4 suggest that starting with the estimates of λ' and μ' from the CSP data, systemic risk will be most effectively addressed by improving the cloud repair rate μ' . Doubling μ' will enable the systemic risk target to be met, while halving the incident rate λ' will not. If μ' is increased to nearly 50%, then a far higher outage frequency λ' can be tolerated without introducing systemic risk. Put simply, if cloud outages are almost always repaired in a matter of minutes, then even if they are relatively frequent, they will

¹⁵³ If systemic events cover outages lasting at least one hour, then the relationship is linear.

not introduce systemic risks that only emerge after several hours.

Mitigating risk through back-up: multi-cloud outsourcing

A simple extension of the analysis examines a scenario in which firms have access to a backup cloud service – either from a different provider, or from the same provider such that the back-up version of a given service operates fully independently of the primary version (known as a ‘multi-cloud’ approach). Our focus is on a multi-cloud approach for a given core service, to address risk arising from concentration at system level, in contrast to a multi-cloud approach across services to address risk to the operations of a single firm arising from concentration within the firm (ESMA, 2020b).¹⁵⁴

This risk mitigation strategy is already offered to some extent by some CSPs by constructing separate groups of cloud computing resources designed to be largely independent of each other, often known as ‘zones’. Zones may be connected to each other within a geographical region. Services can be provided at a regional level, meaning that even if one zone suffers an outage, the services are likely to remain in operation. For example, Google Cloud (2021) aims for each zone to achieve 99.9 % availability (i.e. $\tau' = 0.1\%$) but aims for each region to achieve 99.99 % availability (i.e. $\tau' = 0.01\%$).

To extend the analysis to a multi-cloud scenario, we suppose that each of the 20 clearing members in the application now uses a multi-cloud model – specifically, using a back-up service from a different provider to seamlessly enable them to carry out their functions if their primary CSP suffers an outage. As set out below, a key feature of this new scenario is that a systemic event (again triggered when 3 firms suffer simultaneous outage) now requires 2 providers to suffer a simultaneous outage, rather than one.

For simplicity, as in the general n -firm case we assume that providers’ clients are shared equally with the other firms. This implies that just as in the primary market, the 4 CSPs have equal market shares in the market for back-up services.

If just one CSP suffers an outage, then its client firms are instantly able to switch to the back-up

service, and their operations are interrupted. If two CSPs suffer a simultaneous outage, then a third of the 5 client firms of each provider suffer an outage (since each backs up one third of the market for the other firms), making a total of $\frac{10}{3}$ firms. Since the threshold for a systemic outage is $S = 3$, a systemic event now requires simultaneous outage by two CSPs.

Assuming a 2-hour minimum for systemic outage, the odds ratio of scenario 2 (cloud outsourcing without back-up) compared with the no-cloud baseline is $R \sim 10^3$. In other words, systemic risk is around a thousand times higher in the case with cloud outsourcing.

In contrast, the odds ratio of scenario 2 (cloud outsourcing with back-up) is $R \sim 1$, i.e. risk is reduced to around the level of the no-cloud baseline.

In summary, if firms back up their cloud services, the odds ratio decreases by several orders of magnitude. A multi-cloud model is a successful mitigant in the stylised model, based on the parameter calibration examined. However, our model only takes account of the efficacy of risk mitigants, neglecting the costs of improving resilience and security. A relevant policy consideration would be whether the risk reduction outweighs the associated costs.

An important caveat to the finding that back-up is a successful mitigant is that CSP outages are (like firm outages) assumed to be independent. Introducing positive correlation between CSP outages (stemming for example from shared vulnerabilities) would weaken the effectiveness of a multi-cloud policy. Nonetheless, discussion with market participants suggests that CSPs are likely to have different cybersecurity strategies and measures, which limits the scope for common vulnerabilities to malicious actions. Additionally, the scope for common vulnerabilities to natural disasters is limited by geography, in a similar manner to the crucial assumption made in the model of independence of firm-level outages in the no-cloud baseline.

Conclusion

The growing use of CSPs by financial institutions can provide benefits to individual firms and the

¹⁵⁴ ESMA (2020b) includes guidelines for firms to assess concentration risk both at firm level and at sectoral level,

and for competent authorities to monitor such risks once they are identified.

financial system. However, the high degree of concentration within CSPs might create financial stability risks if CSPs were to suffer an outage that affected their clients, as the likelihood of simultaneous outages might increase.

We discuss several options that can be pursued to mitigate this risk. First, if CSPs are substantially more resilient than individual firms, systemic risk could decline as the additional resilience gained by using CSPs more than compensates for concentration risk. Finally, multi-cloud solutions, where firms use one CSP and then another as backup – or alternatively, the successful provision of cloud services via independent groups of resources by the same provider – may significantly reduce systemic risk. This will only happen, however, if the different CSPs or groups of resources have limited shared vulnerabilities. It is also important to bear in mind that mitigation options are likely to involve costs, and so the optimal solution may be to tolerate a certain level of risk.

Our work also shows the need for detailed data on outages by financial institutions and CSPs. Having consistent data reported by firms and CSPs would allow for better calibration of the model and improve the assessment of trade-offs between different uses of CSPs by firms.

Given the ubiquity of CSPs and continuing migration to use of their services – a trend accelerated by the COVID-19 pandemic – it is crucial for policymakers and market participants to assess the benefits and risks of outsourcing to CSPs. An important example in the EU is the proposed Digital Operational Resilience Act, which envisages a mandate for the European Supervisory Authorities, working with other authorities, to oversee third party providers of critical financial services to address related systemic risks (European Commission, 2020).

References

- Aldasoro, I., Gambacorta, L., Paolo Giudici, P. and Leach, T., "[The Drivers of Cyber Risk](#)", BIS Working Papers 865, May 20, 2020.
- Asensio, C., Bouveret, A. and Harris, A., "Financial stability risks from cloud outsourcing", ESMA Working Paper, *forthcoming*.
- Association for Financial Market in Europe, "[European primary dealers handbook](#)", 2020.
- Bouveret, A., "[Estimation of Losses Due to Cyber Risk for Financial Institutions](#)" Journal of Operational Risk 14, no. 2 (May 16, 2019): 1–20.
- Bouveret, A., Haferkorn, M., Marseglia, G. and Panzarino, O., "Flash crashes on sovereign bond markets – EU evidence", ESMA Working Paper, *forthcoming*.
- Brauchle, P., Göbel, M. and Seiler, J., "Cyber Mapping the Financial System." *Cyber Policy Initiative Working Paper Series 6* (2020): 28.
- CPMI-IOSCO. "[Principles for Financial Market Infrastructures](#)", April 16, 2012.
- Demirer, M., Diebold, F., Liu L. and Yilmaz K., "[Estimating Global Bank Network Connectedness](#)", Journal of Applied Econometrics 33, no. 1 (2018): 1–15.
- ESMA (2020a). "[3rd EU-Wide CCP Stress Test](#)" Report. Paris, France, 2020.
- ESMA (2020b), "[Final Report: Guidelines on Outsourcing to Cloud Service Providers](#)".
- European Commission, "[Proposal for a Regulation of the European Parliament and of the Council on digital operational resilience for the financial sector and amending Regulations \(EC\) No 1060/2009, \(EU\) No 648/2012, \(EU\) No 600/2014 and \(EU\) No 909/2014](#)", COM/2020/595 final, 2020.
- FSB, "[Third-Party Dependencies in Cloud Services - Considerations on Financial Stability Implications](#)", December 9, 2019.
- Harmon, R., Vytelingum P., and Babaie-Harmon J., "[Cloud Concentration Risk: A Framework Agent Based Model For Systemic Risk Analysis](#)", June 22, 2020.
- Lloyd's. "[Cloud down Impacts on the US Economy](#)", Emerging Risk Reports, 2018.
- "[Counting the Cost Cyber Exposure Decoded](#)", Emerging Risk Reports, 2017.
- Naldi, M., "[Evaluation of Customer's Losses and Value-at-Risk under Cloud Outages](#)", 2017.