Portfolio Compression: Positive and Negative Effects on Systemic Risk

Steffen Schuldenzucker Sven Seuken Stefano Battiston
University of Zurich

First version: November 25, 2016
This version: March 7, 2018

We show how conservative portfolio compression, i.e., netting cycles in a financial network, can *increase* systemic risk even though the exposures of all banks to each other *decrease*. We argue that this is because cycles enable risk sharing between the involved banks and can thus dampen the effect of a shock on the rest of the financial system.

We first provide an example where under certain shock values, portfolio compression decreases total equity in the system and thus increases the overall loss in value due to default costs compared to the system when no compression is performed. We next argue that in our example, banks that are part of the cycle to be compressed always profit from compression and thus have an incentive to promote it. Thus, portfolio compression yields an instance of the "tragedy of the commons" problem where individually rational decisions lead to low efficiency. We then show via simulations that our example is robust to changes in the shock size and the strength of the cycle eliminated. We further show that depending on the shock size, no, partial, or complete compression of cycles may be most efficient, giving rise to a trade-off: portfolio compression can make a financial system more resilient to small shocks, but more vulnerable to large shocks.

1 Model

We assume the (Rogers and Veraart, 2013) model. There are banks $N = \{1, \ldots, n\}$ and each bank $i \in N$ has external assets $e_i \geq 0$. Any two banks i, j have fixed liabilities $l_{i,j} \geq 0$ between them. The liabilities could be due to debt contracts, i.e., the banks would know them in advance, or they could have arisen from derivatives on the same external event. The model thus includes, e.g., networks of credit default swaps (CDSs) written on the same external entity (like a country), but does not include networks where banks write CDSs on each other.[1] Given is further a *default cost parameter* $\alpha \in [0, 1]$. We assume that defaulting banks can only pay a share of α of their assets to their creditors while a share of $1 - \alpha$ is lost.[2]

Payments, or equivalently the values of contracts, are given as a collection of real numbers $p_{i,j} \in [0, l_{i,j}]$ for $i, j \in N$. Given a vector of payments p, we can compute the *assets* of a bank i as $a_i(p) := e_i + \sum_j p_{j,i}$. We can further compute the *total liabilities* of i as $l_i := \sum_j l_{i,j}$. p is called *clearing* or a *solution* if it satisfies the following equation for

[1] In (Schuldenzucker et al., 2016), we present an extension of the (Rogers and Veraart, 2013) model that *does* include CDSs on banks.

[2] The original (Rogers and Veraart, 2013) model allowed two different parameters α and β for default costs on external and interbank assets, respectively. We use a simplified model to ease the exposition; all results still hold in the original model.

all $i, j \in N$:

$$p_{i,j} = l_{i,j} \cdot \begin{cases} 1 & \text{if } a_i(p) \geq l_i \\ \frac{\alpha \cdot a_i(p)}{l_i} & \text{if } a_i(p) < l_i \end{cases}$$

That is, banks with sufficient assets to pay their liabilities in full do so while banks without sufficient assets pay their liabilities discounted by their *recovery rate*, i.e., the ratio of their assets after default costs and their liabilities. This means that the total payments of defaulting banks equal their assets after default costs and this amount is distributed to creditors in proportion to the respective liabilities.

Note that p occurs on both sides of the equation, making computing clearing payments non-trivial. We know from Rogers and Veraart (2013) that there is always a solution p that simultaneously maximizes all payments and, equivalently, the assets of all banks. While there could be further solutions, we only consider the maximal solution in this paper.

Given p, we define the *positive equity* of a bank i by $E_i^+ := \max(0, a_i - l_i)$. Note that if i is in default, then $E_i^+ = 0$. If the external assets represent the value that "enters" the financial system from outside, then equities represent the value that "leaves" the financial system after clearing.

We define two measures of efficiency of the conversion from external assets to equities: first, the *total equity*

$$E_*^+ := \sum_i E_i^+$$

and second, the *systemic default cost indicator*

$$\bar{\alpha} := \frac{\sum_i E_i^+}{\sum_i e_i}.$$

The two measures are closely related. E_*^+ is useful for comparing different network structures when the external assets are the same while $\bar{\alpha}$ can be used to compare both different network structures and different external asset vectors. $\bar{\alpha}$ describes the share of the total external assets that are realized as equity, so that $1 - \bar{\alpha}$ is the share of external assets lost due to default costs. The value is always in $[0, 1]$. If no bank defaults, then $\bar{\alpha} = 1$, i.e., nothing is lost. For $\alpha = 1$ (no default costs) Eisenberg and Noe (2001) showed that still $\bar{\alpha} = 1$ even if there are defaults. In acyclic networks, $\bar{\alpha}$ is bounded from below by α^κ where κ is the length of the longest path of defaulting banks that have positive assets. In cyclic networks, $\bar{\alpha}$ can be arbitrarily low.

2 Portfolio Compression Can Decrease Total Equity

Consider the two financial networks in Figure 1. These networks could have arisen in a two-period model in the following way:

- *Period 1:*
  - i) Banks form contracts with each other. Banks hold risky external assets.
  - ii) In case b), portfolio compression is performed and the cycle A–B–C is eliminated. In case a), nothing further happens.
- *Period 2:* Banks are hit by an exogenous shock. For bank A, this shock is negative and reduces the value of its external assets to 0.

We consider clearing payments after the shock and we compare the cases a) and b) with respect to our efficiency measures E_*^+ and $\bar{\alpha}$.

Figure 1 Financial system (top) and its variant where the cycle A–B–C is eliminated (bottom). Arrows represent liabilities ($l_{i,j}$ values) and numbers in boxes represent external assets (e_i values). Assume default costs of 50%, so $\alpha = 50\%$.

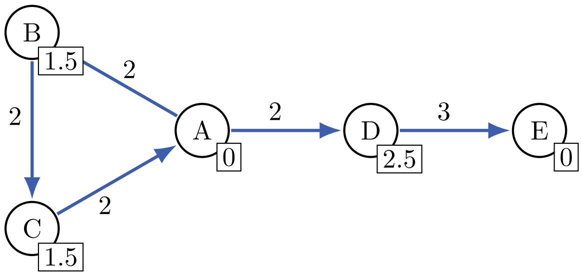

(a) Network without compression

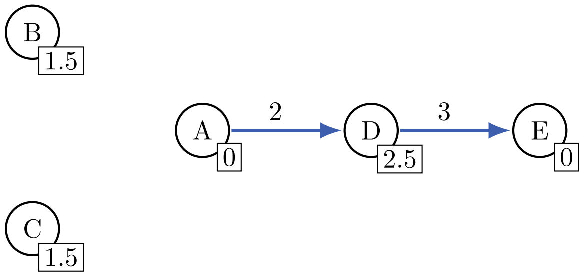

(b) Network with compression

Case a): Uncompressed Network For the (uncompressed) case a), we use a variant of Eisenberg and Noe's (2001) *fictitious default algorithm* to compute clearing payments. We need to resolve the cycle A–B–C; then we can compute clearing payments for the banks D and E.

- Round 1: Even if $p = l$, $a_A = 2 < 4 = l_A$, so A defaults and $a_A = 2$, so $p_{A,B} = l_{A,B} \cdot \frac{\alpha a_A}{l_A} = 2 \cdot \frac{0.5 \cdot 2}{4} = 0.5$.
- Round 2: $a_B = e_B + p_{A,B} = 1.5 + 0.5 = 2 = l_B$, so B does not default. Neither does C.
- Round 3: The cycle A–B–C is resolved and since $l_{A,D} = l_{A,B}$ also $p_{A,D} = p_{A,B} = 0.5$ and $a_D = e_D + p_{A,D} = 2.5 + 0.5 = 3 = l_D$. Thus, D does not default and $p_{D,E} = 3$.

For equities, we have $E_A^+ = 0$, $E_B^+ = 0$, $E_C^+ = 1.5$, $E_D^+ = 0$, and $E_E^+ = 3$. For banks B and D, assets exactly equal liabilities, so they have equity 0, but do not default. Now the total equity is

$$E_*^+ = 4.5$$

and the systemic default cost indicator is

$$\bar{\alpha} = \frac{4.5}{5.5} \approx 81.9\%.$$

Case b): Compressed Network In the (compressed) case b), we can compute clearing payments easily: $a_A = 0$, so $p_{A,D} = 0$ and $a_D = 2.5 < 3 = l_D$, so D defaults and $p_{D,E} = \alpha a_D = 0.5 \cdot 2.5 = 1.25$. We have $E_A^+ = E_D^+ = 0$, $E_B^+ = E_C^+ = 1.5$ and $E_E^+ = 1.25$, so the total equity is

$$E_*^+ = 4.25$$

and the systemic default cost indicator is

$$\bar{\alpha} = \frac{4.25}{5.5} \approx 77.3\%.$$

Interpretation We see that less value is lost to default cost in the uncompressed case a). The difference is small, but visible.[3] Intuitively, the reason why the uncompressed network is more efficient is that A's loss is shared between B and D, so D does not have to bear it alone. This prevents D from defaulting and thus prevents the external assets of D from falling victim to default costs. B on the other hand is highly enough capitalized not to default, either. On the other hand, 50% of the payment from C to A is lost to default costs at A while this payment does not exist in the compressed case. This is why the effect is not more pronounced.

At a higher level, we understand that cycles enable risk sharing: the relatively highly capitalized bank B supports the (post-shock) poorly capitalized bank A, which dampens the effect of the shock at A on the rest of the system. Note that the result in our example is qualitatively independent of the external assets of C because B does not default, so C is always paid in full and can thus always pay A in full.

Portfolio compression is related to the introduction of a seniority structure.[4] To see this, consider a third variant c) where the cycle A–B–C is not removed, but all claims in the cycle are made senior while the other claims become junior. Now assume that during clearing, a bank that can pay its senior liabilities in full is never considered in default for the purpose of paying these senior liabilities even if it might later default upon its junior liabilities. Then this case c) leads to the same equities as the compressed case b): A would pay its senior liability to B in full out of the (senior) payment it receives from C. Then A would not have anything left to pay its junior liability to D, so D would receive nothing and default. Note that this equivalence heavily depends on our assumptions to the clearing mechanism: if banks that cannot pay their junior liabilities are already considered in default while paying their senior liabilities, then after subtracting default costs, A would only pay 1 to B and nothing to D, leading to $E^+_* = 3.25$, a much worse result than both the compressed and non-compressed case.

3 Network Formation and Incentives

We now consider step 1.ii) in the description how the network may arise. We assume that banks can somehow influence whether or not compression is performed and we aim to understand for which banks compression is desirable and for which it is not. We further assume that the shock in step 2 is not fixed any more to reduce the external assets of A to 0, but is random. Banks do not know the shock in advance, but they may have some knowledge about the distribution of shocks and thus of the distribution of the final external assets. This knowledge could be due to the fact that banks typically know how risky their own external investments are and they might also have some information about the external investments of others. We assume that banks aim to maximize their expected equity; it may or may not be be a further goal of banks or the regulator to maximize the expected recovery rate of defaulting banks.

We now examine the incentives of banks to promote compression of the cycle. We first take an ex-post perspective to understand how banks should act if they know the vector of shocks upfront. We then consider an ex-ante perspective where banks only have some probabilistic information.

In our example, A was exposed to a strong negative shock and the cycle dampens the effect of the shock on its creditor D outside the cycle as well as D's creditor E while it

[3] A generalized variant of our example where the difference in $\bar{\alpha}$ is as large as $1 - \alpha$ can be found in Appendix A.

[4] We thank Helmut Elsinger for bringing this up.

exposes A's creditor B inside the cycle to the shock. So B prefers compression while D and E prefer no compression. For C, it does not make a difference whether or not the cycle is compressed since the "cycle part" of the shock is fully absorbed by B. A has equity $\max(0, e_A - 2)$ independently of whether or not the cycle is compressed while its recovery rate is higher in the non-compressed case for shock vectors where A defaults. Thus, A is either indifferent to compression or prefers no compression depending on whether banks are only interested in their equity or also in their recovery rate.

If instead of A, B or C are exposed to a strong negative shock, then the effect is fully contained in the compressed case (in fact, $\bar{\alpha} = 1$) while in the uncompressed case, it may negatively affect some or all of the banks, depending on the shock vector. Thus, all banks prefer compression over no compression.

Now consider an ex-ante perspective where shocks are random and assume that banks fall into one of two categories: risky banks hold risky external assets, making them susceptible to strong negative shocks, while safe banks hold safe external assets. Assume further that it is common knowledge to which of the two categories each bank belongs. We conclude from the above discussion for our example network:

- Safe banks that are part of a cycle prefer to compress the cycle.
- Risky banks that are part of a cycle prefer to compress the cycle if they are pure expected-equity maximizers. They may prefer to keep the cycle if they also take into account their own recovery rate in case of default and this default is rather likely.
- Outside creditors of safe banks prefer to compress the cycle.
- Outside creditors of risky banks prefer to keep the cycle.

Now assume that in step 1.ii) in the description how the network may arise, cycles are compressed iff the banks in the cycle agree to compress them and assume that it is common knowledge that B and C are safe and A is risky. If banks are pure expected-equity maximizers, then all banks A, B, C agree to compress the cycle. Yet, efficiency, as measured by the expected total equity or the expected systemic default cost indicator, is lower than in the non-compressed case.

4 Simulation: Varying Shock Size and Cycle Strength

An immediate question in the context of our example from Figure 1 is to which extent the result from Section 2 depends on the choice of the numeric values. In particular, one may wonder how efficiency of the compressed and the non-compressed case depends on the external assets of A (0 in our example) and the strength of the cycle A–B–C (2 in our example). We conducted simulations to answer these questions.

Figure 2 shows the total equity E_*^+ dependent on the external assets e_A of A where we consider five different values for the notional γ of the cycle A–B–C.[5] For $e_A = 2$ and all values of γ, A does not default. Thus, no other bank defaults, $E_*^+ = \sum_i e_i = 7.5$, and $\bar{\alpha} = 1$. For $e_A < 2$ and all values of γ, A defaults, which gives rise to a discontinuous loss in total equity due to default costs. This is the reason why we see a "jump" just left of $e_A = 2.0$ for all curves. For $\gamma = 2$, no further defaults occur as e_A decreases, so E_*^+ forms a straight line. For $\gamma \in \{1.5, 1, 0\}$, bank D defaults at some point as e_A is decreases, which gives rise to a second discontinuous loss in total equity. This happens sooner (i.e., for higher values of e_A) for smaller values of γ. The higher value of $\gamma = 2.5$ leads to the worst efficiency for all values of e_A. This is because the assets of A now contain the

[5] We could have plotted the systemic default cost indicator $\bar{\alpha}$ instead of E_*^+; the results are qualitatively the same.

Figure 2 Total equity E_*^+ (y axis) dependent on the external assets e_A of A (x axis) in the network from Figure 1 where the notional of the cycle A–B–C is replaced by γ. $\gamma = 2$ corresponds to Figure 1a while $\gamma = 0$ corresponds to Figure 1b. These two values of γ are marked by squares, the other values by circles.

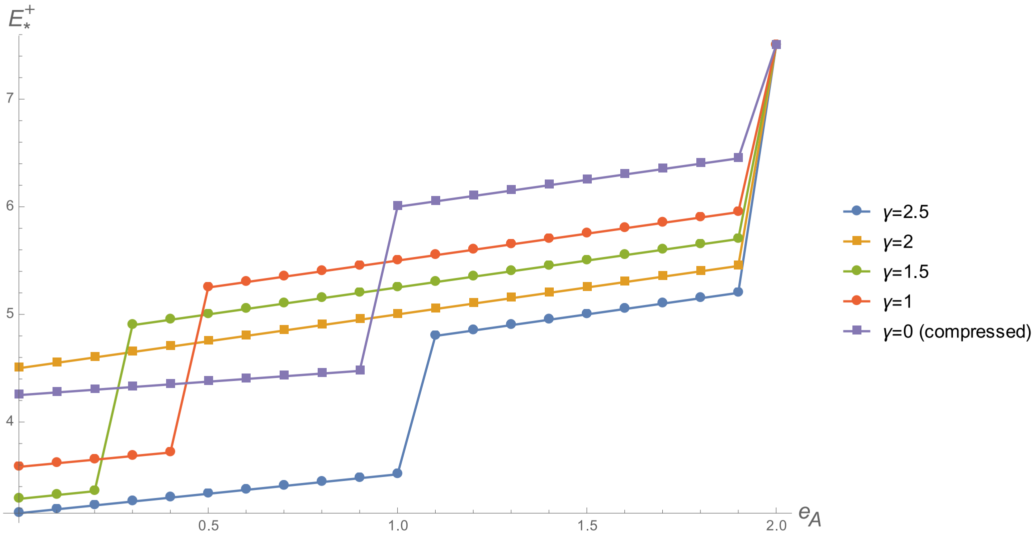

relatively large value $p_{C,A} = 2.5$ and are affected by default costs at A. The resulting loss is born to the larger part by B until it defaults at $e_A \leq 1$, while D never defaults.

We can also compare different values of γ for fixed values of e_A. If e_A is high, then total equity monotonically increases as γ is decreased while if e_A is low, then total equity is maximal for the relatively high value $\gamma = 2$. Note however that the effect is not monotonic if e_A is low: here, γ values from worst to best are: 2.5, 1.5, 1, 0, 2, so that a cycle of strength 2 should not be compressed, but cycles of strength 1, 1.5, or 2.5 (*both* higher and lower values!) should be compressed. We witness the worst possible impact of compression in this network at $e_A = 0.9$ and $\gamma = 1$: here, compression decreases total equity by 1 and $\bar{\alpha}$ by $\frac{1}{6.4} \approx 15.6\%$.

Our simulation result gives another indication for our earlier interpretation: if external assets are high, then portfolio compression is most efficient because it insulates banks from losses at other banks (in this case B is insulated from a loss at A). However, if external assets are low, then portfolio compression may not be optimal any more because it also prevents risk sharing. From an ex-ante perspective, this constitutes at trade-off: portfolio compression can make the system more resilient to small negative shocks, but less resilient to large negative shocks. If we follow the line of the topmost points from left to right, we see that in fact, different cycle strengths γ can be optimal depending on the value of e_A. This seems related to Elliott et al.'s (2014) discussion of diversification and integration and Acemoglu et al.'s (2015) work on the optimal level of interconnectedness dependent on the shock size.

References

D. Acemoglu, A. Ozdaglar, and A. Tahbaz-Salehi. Systemic risk and stability in financial networks. *American Economic Review*, 105(2):564–608, Feb 2015.

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

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

L. Rogers and L. A. Veraart. Failure and rescue in an interbank network. *Management Science*, 59(4):882–898, 2013.

S. Schuldenzucker, S. Seuken, and S. Battiston. Clearing payments in financial networks with credit default swaps [extended abstract]. In *Proceedings of the 17th ACM Conference on Economics and Computation*, EC '16, Maastricht, The Netherlands, 2016. ACM. Current Working Paper: `http://www.ifi.uzh.ch/ce/publications/Clearing_CDSs.pdf`.

Figure 3 Generalization of the system in Figure 1. Let α, γ, and δ be arbitrary such that $0 < \alpha < 1$, $\gamma > 0$, and $\delta \geq \frac{1}{2}\alpha\gamma$. Set $e_B = e_C = \gamma - \frac{1}{2}\alpha\gamma$ and $e_D = \delta - \frac{1}{2}\alpha\gamma$.

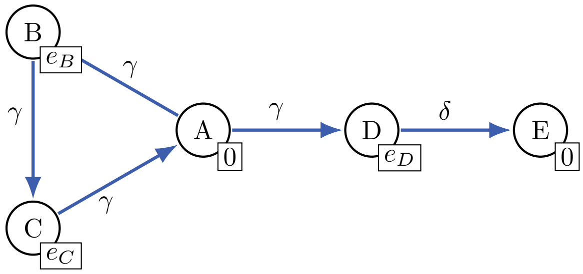

(a) Network without compression

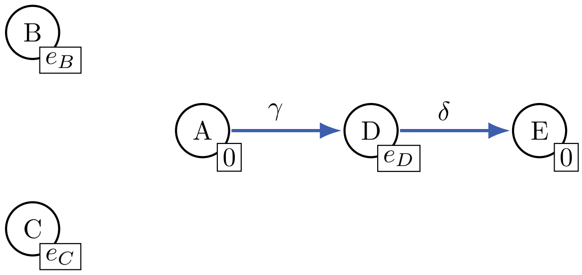

(b) Network with compression

APPENDIX

A Generalized Version of the Example from Figure 1

We present a class of financial systems where $\bar{\alpha}$ can be arbitrarily close to 1 in the case without compression, but close to α in the case with compression.

Consider Figure 3. This financial system is a generalization of Figure 1 where we kept the cycle strength γ and the liability δ from D to E variable and chose the external assets dependent on these values. Let α be fixed, but arbitrary. We first consider fixed, but arbitrary values of γ and δ. We then consider the limit case in these two parameters.

For any choice of α, γ, δ we compute

$$\sum_i e_i = 2e_B + e_D = \delta + 2(1 - \frac{3}{4}\alpha)\gamma.$$

We now proceed just like in Section 2 to compute E_*^+ and $\bar{\alpha}$ in cases a) and b).

Case a): Uncompressed Network

- Round 1: A defaults and $a_A = \gamma$, so $p_{A,B} = l_{A,B} \cdot \frac{\alpha a_A}{l_A} = \gamma \cdot \frac{\alpha\gamma}{2\gamma} = \frac{1}{2}\alpha\gamma$.
- Round 2: $a_B = e_B + p_{A,B} = \gamma - \frac{1}{2}\alpha\gamma + \frac{1}{2}\alpha\gamma = \gamma = l_B$, so B does not default. Neither does C.
- Round 3: The cycle A–B–C is resolved and $p_{A,D} = p_{A,B} = \frac{1}{2}\alpha\gamma$. We have $a_D = e_D + p_{A,D} = \delta - \frac{1}{2}\alpha\gamma + \frac{1}{2}\alpha\gamma = \delta = l_D$, so D does not default and $p_{D,E} = \delta$.

For equities, we have $E_A^+ = 0$, $E_B^+ = 0$, $E_C^+ = \gamma - \frac{1}{2}\alpha\gamma$, $E_D^+ = 0$, and $E_E^+ = \delta$. For banks B and D, assets exactly equal liabilities, so they have equity 0, but do not default.

Now the total equity is

$$E_*^+ = \delta + (1 - \frac{1}{2}\alpha)\gamma$$

and the systemic default cost indicator is

$$\bar{\alpha} = \frac{\delta + (1 - \frac{1}{2}\alpha)\gamma}{\delta + 2(1 - \frac{3}{4}\alpha)\gamma}.$$

Case b): Compressed Network $a_A = 0$, so $p_{A,D} = 0$ and $a_D = e_D = \delta - \frac{1}{2}\alpha\gamma < \delta = l_D$, so D defaults and $p_{D,E} = \alpha a_D$. We have $E_A^+ = E_D^+ = 0$, $E_B^+ = E_C^+ = e_B$ and $E_E^+ = \alpha a_D$, so the total equity is

$$E_*^+ = 2e_B + a_D = \alpha\delta + 2(1 - \frac{1}{2}\alpha - \frac{1}{4}\alpha^2)\gamma$$

and the systemic default cost indicator is

$$\bar{\alpha} = \frac{\alpha\delta + 2(1 - \frac{1}{2}\alpha - \frac{1}{4}\alpha^2)\gamma}{\delta + 2(1 - \frac{3}{4}\alpha)\gamma}.$$

Comparison and Limit Case Which of the two networks is more efficient depends on the choice of γ and δ relative to each other: if δ is slightly smaller[6] than γ, then the compressed case is more efficient while if δ is much larger than γ, then the non-compressed case is more efficient. In particular, for any fixed value of δ and for $\gamma \to 0$ we have $\bar{\alpha} \to 1$ in the non-compressed case a) and $\bar{\alpha} \to \alpha$ in the compressed case b). This discrepancy can be explained intuitively as follows: in the limit, almost all the value in the network is concentrated at the external assets of D. In the compressed case, D defaults, and these external assets are affected by default costs. In the non-compressed case, the small payment from A prevents this, so only part of the (vanishingly small) external assets of B is lost.

Our example demonstrates the greatest possible difference in $\bar{\alpha}$ between the two cases for this network topology: since we chose $e_A = 0$, the longest path of non-defaulting banks with positive assets in the compressed case b) consists of bank D only; thus, $\bar{\alpha} \geq \alpha$ for any choice of notionals (cf. Section 1). We cannot achieve a value $\bar{\alpha} = \alpha$ because that would require $e_B = e_C = 0$, in which case the compressed and non-compressed case would be essentially the same. Likewise, we cannot achieve $\bar{\alpha} = 1$ in the non-compressed case a) because this would mean that no bank defaults, but then also no bank defaults in the compressed case. The effect *can* be made more pronounced by changing the network topology: first, we could append to E a chain of banks with 0 external assets owing δ to each other. It is easy to see that this would preserve $\bar{\alpha} \to 1$ in the non-compressed case, but yield $\bar{\alpha} \to \alpha^\kappa$ where κ is the length of the chain. Alternatively, we could replace bank D by a cycle of intermediation to achieve $\bar{\alpha} \to 0$ in the compressed case.

[6] The threshold, where both values of E_*^+ and $\bar{\alpha}$ are equal, is at $\delta = \frac{1 - \frac{1}{4}\alpha^2}{1-\alpha}\gamma > \frac{1}{2}\alpha\gamma$.