Contagious disruptions and complexity traps in economic development

Charles D. Brummitt[a,b], Kenan Huremovic[c,d], Paolo Pin[e], Matthew H. Bonds[b,f], and Fernando Vega-Redondo[e]

[a]Center for the Management of Systemic Risk, Columbia University, New York, NY, 10027, USA
[b]Department of Global Health and Social Medicine, Harvard Medical School, Boston, MA 02115, USA
[c]Aix-Marseille School of Economics, Aix-Marseille University, 2 Rue de la Charité, 13002 Marseille, France
[d]Centre National de la Recherche Scientifique (CNRS), 3, rue Michel-Ange, 75794 Paris, France
[e]Department of Decision Sciences and Innocenzo Gasparini Institute for Economic Research, Università Bocconi, Via Roentgen 1, Milano 20136, Italy
[f]Earth System Science, Stanford University, Stanford, CA 94305, USA
[*]Corresponding author: charles_brummitt@hms.harvard.edu

September 12, 2016

Abstract

Poor economies not only produce less; they typically produce things that involve fewer inputs and fewer intermediate steps. Nevertheless, the supply chains of poor countries face more frequent disruptions—delivery failures, faulty parts, delays, power outages, theft, government failures—that systematically thwart the production process. To understand these effects on economic development, we model an evolving input–output network in which disruptions spread contagiously among optimizing agents. The key finding is that a poverty trap can emerge: agents adapt to frequent disruptions by producing simpler, less valuable goods, yet disruptions persist. Growing out of poverty requires that agents invest in buffers to disruptions. These buffers rise and then fall as the economy produces more complex goods, a prediction consistent with global patterns of input inventories. Large jumps in economic complexity can backfire, so the model supports policies that gradually increase technological complexity. This advice contrasts with "big push" policies for overcoming poverty traps.

Producing valuable goods and services is a complex, intricate process. One obtains inputs from a multitude of suppliers who must honour their contracts and deliver those inputs without them breaking, spoiling, or being stolen. These inputs must be stored safely and manipulated in interdependent stages, using labour from workers who may fall ill or shirk their duties, together with complex equipment and vast infrastructure that may malfunction. These complex interdependencies underlie specialization and trade that are the foundation of economic growth and material progress [1, 2].

Yet this progress, and the disruptions that thwart it, are unevenly distributed around the world. In low-income countries, disruptions can be frequent, long-lasting, and severe. They include power outages [3, 4], worker absenteeism [5], failed deliveries of products, water shortages, customs delays, damage from natural disasters, and epidemic diseases (Fig. 1). Poor countries also tend to produce simpler goods, especially primary resources like timber, mining, and subsistence agriculture [6, 7, 8].

In middle- and high-income countries, by contrast, inputs tend to be more reliable and goods produced tend to be more complex. Yet rich economies are not immune to disruptions: competition drives firms to build lean supply chains with buffers so small that disruptions can cascade around the globe, causing large aggregate losses [13, 14].

Might the mechanisms causing globalised supply chains to become fragile also be preventing poor economies from becoming more complex and global?

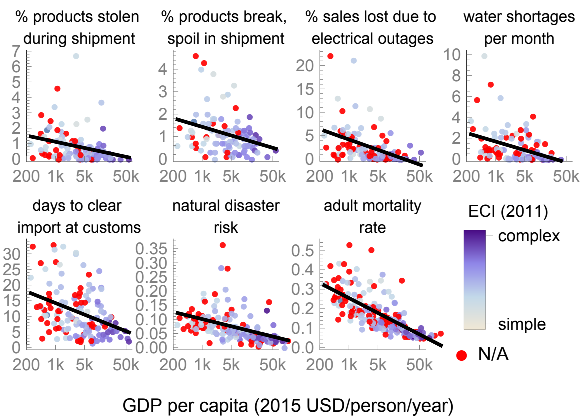

Figure 1: Disruptions to the production process tend to be more frequent in poorer, less complex economies. The color of each dot indicates the country's Economic Complexity Index (ECI) [9]; if ECI is missing, then the dot is red. Black lines are least-squares fits, with per-capita incomes [10] on a logarithmic scale. Natural disaster risk combines exposure and ability to cope [11]. Adult mortality rate is the chance that a 15-year old dies before age 60 [12]. The first five panels' data are from [4].

This question stretches the limits of our understanding of economic growth and complexity. Input–output linkages among firms—wherein one firm's output is another firm's input—are known to have large, nonlinear effects on economies. In theoretical models, these linkages propagate changes in productivity [15, 16, 17, 18], disruptions [19], and bankruptcies [20, 21, 22, 23, 24, 25]. Empirical research on industrialised economies finds that supply-chain disruptions often lead to lower stock prices and sales growth [26, 27, 28, 29]. Linkages and disruptions affect development by causing firms to use less capital [19], to fail to industrialise [15], to misallocate inputs [16, 17], or to form shorter supply chains [24]. However, in these models, disruptions are exogenous, and firms interact once in a static network, thereby precluding dynamic feedback of the kind that can generate complex outcomes such as poverty traps and periodic cycles.

Modelling dynamic production networks remains an unsolved problem because these networks involve heterogeneous input–output patterns and elasticities of myriad pairs of inputs [30]. Recent models consider firms that endogenously form input–output linkages [31, 32, 33]; others consider firms deciding how to source their inputs in a risky supply chain with one [34, 35] or more [36, 37, 38] tiers. Missing is an understanding of how fast dynamics in economic networks, such as disruptions in supply chains, affect their long-run evolution and their growth in complexity.

We aim to fill this theoretical gap by introducing a simple model that captures complex dynamics of disruptions spreading in an evolving input–output network. The main result is that poverty can emerge and reinforce itself: facing an unreliable environment of potential inputs, agents choose simple production processes that require few inputs, but disruptions remain frequent. Escaping this trap requires investing in buffers against disruption, such as arranging for extra suppliers or storing inventories of inputs. We find empirical support for the prediction that these buffers grow and then shrink as economies develop. When they shrink too much, disruptions can spike in number, as occurs in lean supply chains today. This mechanism also imperils developing economies: jumping abruptly to a more complex technology can backfire by causing greater dysfunction, suggesting that "big push" policies [39] may benefit from technological gradualism. We suggest that this alternative perspective—focused on contagion in evolving supply chains—may shed light on why some poor economies are not catching up with the more advanced ones.

Methods

We consider a large population of agents who represent entrepreneurs or firms producing goods and services that require inputs from other agents. The model framework is meant to correspond to a variety of situations that broadly represent the process of coordinating inputs and outputs for economic production: launching a business requires intermediate goods from suppliers; coordinating stakeholder meetings requires a quorum of attendees; repairing equipment requires parts supplied by others; and so on.

Balls-and-urn model of production and contagious dysfunction

At each time t, all agents exist in one of two states. A fraction $F(t)$ of agents are *functional*: they recently succeeded in producing and can provide inputs to others upon request. The remaining fraction $1 - F(t)$ are *dysfunctional*: they recently failed to produce and cannot provide inputs to others.

Agents become functional and dysfunctional as they succeed and fail, respectively, in producing goods or accomplishing tasks. Each agent attempts to produce a good requiring τ many inputs. Attempts at producing a good occur randomly at a constant rate (as a Markov process). We do not track types of inputs nor economic sectors. This simplification allows us to abstract from which pairs of inputs are substitutes by using a simple threshold rule: An agent attempts to obtain inputs from m agents in the population, and she succeeds in producing if and only if at least τ of those m many inputs are successfully produced and delivered to her (see Fig. 2).[1] We call m the number of *attempted inputs* (or the *in-degree* when viewing these interactions as an input–output network).

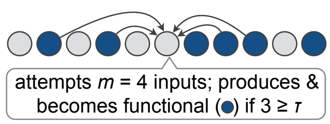

Figure 2: Illustration of the model. Agents of the model are people or firms who are either functional or dysfunctional at any moment in time. *Functional* means that has enough inputs needed to produce or to accomplish a task, and that other agents can rely on this agent for inputs. At constant rates, a random agent attempts to produce a good (or to have a meeting with other people, etc.) using m inputs drawn randomly from the population, and she succeeds if at least τ of them are functional.

These simplifications allow us to describe an evolving input–output network, together with disruptions spreading on it, using a single differential equation for the expected fraction of functional agents $F(t)$:

$$\mathrm{d}F/\mathrm{d}t := \mathbb{P}\left[\mathrm{Binomial}(m, F(t)) \geq \tau\right] - F(t) \qquad (1)$$

for $t \geq 0$ and integers $m \geq 0$ and $\tau > 0$. (If $\tau = 0$, then there is little to model, so we let $\mathrm{d}F/\mathrm{d}t := 0$.) This framework is a balls-and-urn model [43, 44] taken to an infinite-population, continuous-time limit, so that transition probabilities become deterministic rates of change in the mean-field master equation (1) [45]. For simplicity, the input–output network is random and "annealed": in each production attempt, inputs are chosen uniformly at random with replacement from the population.[2] All agents use the same value of (m, τ) (i.e., they play "symmetric strategies"). Therefore, the chance of successful production is the probability $\mathbb{P}$ that a binomial random variable with parameters m and $F(t)$ is $\geq \tau$, where τ is the number of critical inputs needed. This threshold rule resembles the essential inputs and critical subtasks in the "O-Ring theory" of Kremer [19], but here people can have buffers against failures: the number of attempted inputs (m) can exceed the number of inputs needed (τ). This threshold rule also appears in models of social contagion and collective behavior [46, 47, 48], but here we have an annealed network, bidirectional changes in state, and decision making, described later.

Functional agents—encouraged by their recent successful production—may attempt to produce more frequently than dysfunctional agents. To allow for this flexibility, we generalize equation (1) by assuming that functional agents attempt to produce at a rate L, whereas dysfunctional agents attempt to produce at rate 1. Then

$$\begin{aligned} \mathrm{d}F/\mathrm{d}t &= [1 - F(t)]\, P - LF(t)(1 - P) && \text{(2a)} \\ &= P\,[1 + F(t)(L - 1)] - F(t)L && \text{(2b)} \end{aligned}$$

where P, as before, is the probability that an agent successfully produces, $\mathbb{P}\left[\mathrm{Binomial}(m, F(t)) \geq \tau\right]$. The first term in equation (2a) is the rate $1 - F(t)$ at which dysfunctional agents attempt to produce; each attempt succeeds with probability P; if the attempt succeeds, then $F(t)$ rises; otherwise $F(t)$ stays the same, and vice versa for the second term. Equation (2) is derived in Sec. SI-1, and it recovers equation (1) with $L := 1$. The initial amount of dysfunction $1 - F(0)$ is exogenous; after that, disruptions are entirely endogenous, spreading from supplier to customer. Driving this contagion is the assumption that an agent delivers an input upon request if and only if she successfully produced in her most recent

[1] Some inputs are critical: without them, production halts or fails. For example, the March 11, 2011 earthquake near Japan closed the Hitachi factory that produces most of the world's airflow sensors, a critical input for cars; as a result, automobile factories on the other side of the globe had to curb production or close [14]. Drip irrigation has failed in Sub-Saharan Africa due to disruptions in water infrastructure and scarce knowledge for repair [40]; adulterated fertiliser sold in Ugandan markets yields negative average returns [41]; Internet-connected kiosks in India fell into disuse because of unreliable electricity and insufficient service from operators [42].

[2] This annealed network captures the idea that people do different tasks that require different inputs: an engineer fixes a machine on Monday and leads a meeting on Wednesday; an entrepreneur tries one business idea this year and another idea the next year, requiring different inputs for each idea.

attempt to produce. For example, a Ugandan farmer who discovers that her seeds were inauthentic [41]; an Ethiopian farmer whose drip irrigation system fails because of upstream failures [40]; or an automobile manufacturer who failed to produce due to missing parts [13, 14] all may subsequently fail to deliver output promised to a customer.

We have abstracted from considerations about market equilibrium and price formation: the price of every good and the quantity purchased are both normalised to one.[3] Only τ goods are used in production, even if more than τ of m suppliers are functional, because unused inputs are assumed to be perfect substitutes for used ones.

Deciding on complexity τ and on buffers against disruption $m - \tau$

The threshold τ loosely captures the complexity of the good or service being produced: more complex goods require more inputs [50, 51]. To capture the incentives to create high-value products, we present a simple, reduced-form model in which agents derive utility from successfully producing goods that require more inputs. We assume that when an agent successfully produces she gets some utility that rises with τ with decreasing returns, which for simplicity we take to be τ^β where $\beta \in (0, 1)$. We also assume that each attempted input costs $\alpha > 0$. This parameter α represents the marginal cost of finding suppliers, maintaining multiple suppliers for the same input,[4] incentivising suppliers to have multiple manufacturing sites,[5] or maintaining inventory of inputs.[6]

For simplicity, we assume that each agent knows the current likelihood $F(t)$ that a uniformly-random supplier would successfully produce and deliver an input upon request. Based on that reliability $F(t)$, agents revise their strategy of how complex a product to produce ($\tau \in \{0, 1, 2, \dots\}$) and how many inputs $m \in \{0, 1, 2, \dots\}$ to attempt to procure in order to produce that good. For instance, if suppliers are unreliable [i.e., $F(t)$ is small], then agents arrange for redundant inputs (i.e., $m - \tau > 0$) provided that they can afford it. Agents must commit to a certain technology and production technique for a certain amount of time T, so we assume that every T amount of time all agents simultaneously update their strategy to the "best response", the maximiser (m^*, τ^*) of the utility function

$$\mathcal{U}\left[m, \tau, F(t); \alpha, \beta\right] := P\left[m, \tau, F(t)\right] \tau^\beta - \alpha m. \quad (3)$$

Thus, agents' strategies at time t are

$$(m^*, \tau^*) = \underset{m,\tau \geq 0}{\arg\max}\ \mathcal{U}\left[m, \tau, F(kT); \alpha, \beta\right] \quad (4)$$

for $t \in [kT, (k+1)T)$ where $k \in \{0, 1, 2, \dots\}$. Together, equations (2)–(4) and the initial $F(0)$ define the model.

Results

Figure 3 illustrates the three phases of an economy in this model: trapped, emerging, and rich. To understand the figure, suppose that at time $t = 0$ agents successfully produce and deliver an input only 50% of the time [i.e., $F(0) = 50\%$]. Then, from equation (3), agents choose the strategy $(m^*, \tau^*) = (3, 1)$, meaning that agents produce a good requiring $\tau^* = 1$ input, but they arrange for $m^* - \tau^* = 2$ extra suppliers because disruptions are so common [$1 - F(0) = 50\%$]. Using this strategy in an economy with reliability $F(0) = 0.5$ causes disruptions to become less frequent ($\mathrm{d}F/\mathrm{d}t > 0$), indicated by the green curve marked "3, 1" in Fig. 3.

Figure 3 corresponds to an economy in which agents best respond arbitrarily quickly based on the reliability $F(t)$ of their fellow agents; that is, the best-response timescale T is arbitrarily close to 0. This $T \to 0$ limit is more analytically tractable because $\mathrm{d}F/\mathrm{d}t$ changes discontinuously wherever the best response (m^*, τ^*) changes as a function of $F(t)$. We relax this assumption later when we discuss cycles.

Next we explain the economy's three main phases and a pitfall in reaching the "industrialised" phase (high τ and F).

[3] In models of production networks such as [49], the price of a good is a function of its position in the input–output network. In our model, the network is an annealed m-regular graph, so all agents have a symmetric position in the network, which motivates our assumption that each good has the same price.

[4] Maintaining multiple suppliers for an input can be costly when that requires changing the product, working with suppliers to develop alternatives, or overcoming quality issues with alternative inputs [52].

[5] If a firm has a certain crucial input with just one supplier (a "strategic component"), it may "provid[e] incentives to [those] suppliers to have multiple manufacturing sites in different regions" [53].

[6] Inventory costs can be high: "[A]s product life cycle shortens and as product variety increases, the inventory holding and obsolescence costs of these additional safety stock inventories could be exorbitant" [54].

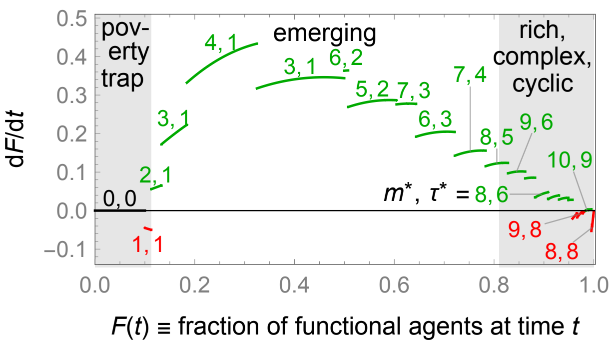

Figure 3: Representative phase portrait, showing the three phases of a model economy. Here, $(\alpha, \beta, L) := (0.1, 0.4, 1.5)$. The black, green, and red curves are the ODE (2) with labels indicating the best response (m^*, τ^*) and colours denoting the sign of $\mathrm{d}F/\mathrm{d}t$.

Poverty trap with simple technology and frequent disruptions

In an economy with frequent disruptions [$F(t)$ near zero], agents choose to withdraw from the economy by not relying on any inputs from others ($m^* = \tau^* = 0$). This strategy resembles subsistence agriculture, hunting, and pastoralism. Such an economy is in steady state: $\mathrm{d}F/\mathrm{d}t = 0$ [from equation (2)] and no agent wants to deviate from the strategy $(0, 0)$.

This steady state also has a basin of attraction: because we have assumed that functional agents attempt to produce $L := 1.5$ times more frequently than dysfunctional agents do, the strategy $(m^*, \tau^*) = (1, 1)$ causes dysfunction to rise (i.e., $\mathrm{d}F/\mathrm{d}t < 0$), indicated by the red line segment labelled "1, 1".[7] This basin of attraction is marked "poverty trap" in Fig. 3.

Emerging economies' buffers to disruption rise and then fall

If an economy is sufficiently reliable then it begins to develop. For instance, in Fig. 3 if $F(t) > 2^{-(\beta+1)}\left(1 - \sqrt{1 - \alpha 2^{\beta+2}}\right) \approx 11\%$ then agents choose to produce goods that require some inputs ($\tau^* > 0$). Provided that $F(t)$ is not too close to one (a case described later), the agents also arrange for some extra inputs ($m^* > \tau^*$) in anticipation that some inputs will not be functional. This strategy results in the economy becoming more functional over time ($\mathrm{d}F/\mathrm{d}t > 0$) and producing ever more complex goods [τ^* rises with $F(t)$]. As this economy develops, two features rise and then fall over time: the speed of development $\mathrm{d}F/\mathrm{d}t$ and the buffer against disruptions $m^* - \tau^*$. (Later we examine this inverted-U pattern empirically.) This inverted-U reflects the idea that in a very unreliable economy firms require costly buffers against disruption to produce even simple, low-value goods (complexity $\tau^* = 1$, for example), and the economy barely manages to produce simple goods using the small amounts of redundancy afforded by the low earnings (redundancy $m^* - \tau^* = 2$, for example). When the economy is more reliable [higher $F(t)$], more complex tasks become feasible with large buffers against disruption, such as complexity $\tau^* = 3$ with buffer $m^* - \tau^* = 4$. Finally, as the economy becomes maximally reliable [$F(t)$ approaches one], agents economise on their costly buffer against disruptions ($m^* - \tau^*$), which leads to new vulnerabilities.

Rich yet cyclic

When the economy becomes very reliable [large $F(t)$], agents produce very complex goods requiring many inputs (large τ). Yet this high reliability also induces agents to economise on their buffers to disruptions. In fact, when $F(t)$ is close to 1, they eliminate their buffer ($m^* = \tau^*$). Then disruptions spread like a virus to which no one is immune: $F(t)$ falls, indicated by the red curves in the bottom-right corner of Fig. 3, where $\mathrm{d}F/\mathrm{d}t < 0$. Falling $F(t)$ means that more and more agents are unable to produce, and the drop in output resembles a recession. Such downturns occur generically in rich, highly functional economies (see Theorem 1 of Sec. SI-3 of the Supplementary Information).

If agents commit to a strategy for a positive amount of time $T > 0$, then the economy's reliability $F(t)$ falls until either (i) the economy enters the poverty trap (which occurs only for very large T) or (ii) agents best respond in a way that causes $F(t)$ to begin to rise. $F(t)$ can rise because agents produce simpler, lower-value goods (smaller τ) or because they increase the buffer against disruption (larger m). After $F(t)$ rises for a while, agents best respond again, and because their economy is quite reliable they choose to produce very complex goods or to decrease their buffer against disruption, and the

[7] If $L = 1$, then $\mathrm{d}F/\mathrm{d}t \equiv P[1, 1, F] - F = F - F = 0$ for all $F \in [0, 1]$, so if the best response (m^*, τ^*) is $(1, 1)$ then the system is in a steady state. Assuming $L \neq 1$ breaks this symmetry, so that $\mathrm{d}F/\mathrm{d}t \not\equiv 0$ for $(m, \tau) = (1, 1)$.

process repeats, resembling business cycles.[8] This mechanism—of expanding and contracting production based on mutual reliability of inputs—appears to be unexplored in the literature on business cycles [55, 56]. It resembles the pressure that firms face today to build leaner supply chains, to invest in smaller buffers against disruptions, and to produce ever more complex goods, which occasionally result in cascading disruptions [13, 14].

Overshooting complexity can backfire

The core mechanism that causes downturns in the rich economy also makes it difficult for emerging economies to become rich and complex themselves. Specifically, if an economy tries to "prematurely jump" to a more complex technology, then it can slide backward and become more dysfunctional. To make this idea precise, suppose that agents do not use the best response (m^*, τ^*) but instead attempt a more complex strategy that requires s more inputs: $(m^* + s, \tau^* + s)$. The buffer against disruptions remains the same; it is still $m^* - \tau^*$. What is different is that agents try to produce goods that require more inputs, or they try to produce the same good as before but using technology that depends on more inputs, such as drip irrigation instead of traditional irrigation [40].

Figure 4 shows that this strategy (m^*+2, τ^*+2) often results in dysfunction rising over time ($\mathrm{d}F/\mathrm{d}t < 0$), indicated by the red curves. In these intervals with $\mathrm{d}F/\mathrm{d}t < 0$, agents are "overshooting" in complexity: they are attempting a production process more complex than what the surrounding system can support. This overshooting echoes failures to adopt complex technologies in developing countries because the technologies depend on myriad inputs prone to disruption, such as drip irrigation systems [40] and Internet kiosks [41].

Emerging economies are especially vulnerable to overshooting in complexity: notice in Fig. 4 that $\mathrm{d}F/\mathrm{d}t < 0$ for many intermediate values of $F(t)$. As a result, the poverty trap in Fig. 4 is dramatically larger than when agents use the best response (compare with Fig. 3). For example, an economy with $F(t)$ near 50% can fall into the poverty trap if it overshoots in complexity for a sufficiently long amount of time.

[8] If an economy could best respond instantaneously (i.e., $T \to 0$), then it would settle upon a value of $F(t)$ at which $\mathrm{d}F/\mathrm{d}t$ changes sign from positive to negative. More plausibly, agents must commit to a certain technology, production technique, or strategy for a while (i.e., $T > 0$), causing cycles.

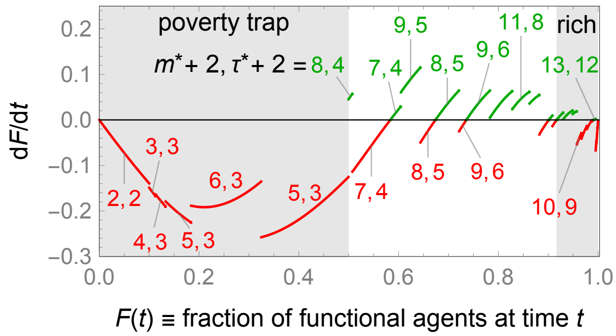

Figure 4: Jumping to a more complex technology can backfire by causing dysfunction to rise, especially for emerging economies. Here, agents use a strategy that requires two more inputs than they would have chosen given the reliability $F(t)$ of their potential inputs: they use the strategy $(m^* + 2, \tau^* + 2)$, where (m^*, τ^*) is the best response. The parameters are the same as in Fig. 3.

By contrast, a rich economy can typically accommodate a jump in complexity without causing dysfunction to rise: in Fig. 4 there are many large values of $F(t)$ with $\mathrm{d}F/\mathrm{d}t > 0$.

Comparing Figs. 3 and 4, we see the benefit of gradual growth in technological complexity. This prescription is at odds with the classic idea of a "big push" of simultaneously industrialising many sectors of an economy [39]: A big push overcomes coordination problems, but it can add fragility by introducing complex technologies that depend on unreliable inputs. This prescription for slow, gradual reform mirrors the suggestions given by a model of trust and social capital [57].

Phase diagram

To demonstrate that the phenomena in Fig. 3 are rather generic, Fig. 5 shows the sign of $\mathrm{d}F/\mathrm{d}t$ and the best response (m^*, τ^*) for many values of the parameter α, the marginal cost of each attempted input. A poverty trap occurs for $F(t) \in [0, \alpha]$; the boundary $F(t) = \alpha$ is the indifference curve between $(m, \tau) = (1, 1)$ and $(0, 0)$. If the cost to arrange for an input is too high ($\alpha > 1/4$ in Fig. 5), then the only long-run outcome is poverty. Otherwise, there exists a good outcome in the long run in which the economy is complex and highly functional, and buffers against disruptions $m^* - \tau^*$ tend to rise and then fall as the economy approaches this rich state.

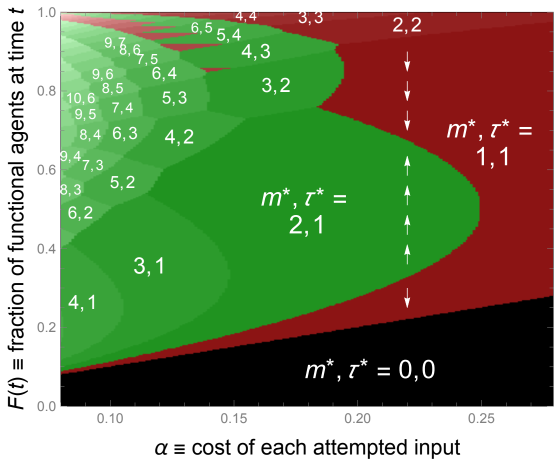

Figure 5: Phase diagram of the sign of $\mathrm{d}F/\mathrm{d}t$ ($\mathrm{d}F/\mathrm{d}t > 0$ is green; < 0 is red; $= 0$ is black) and the best response (m^*, τ^*) (white text) as a function of α and $F(t)$. The parameter α is assumed to change slowly, if at all. Vertical white arrows illustrate the dynamics $\mathrm{d}F/\mathrm{d}t$. Here, as in Fig. 3, $\beta := 0.4$, and agents attempt to produce $L := 1.5$ more frequently, so the strategy $(m^*, \tau^*) = (1, 1)$ results in $\mathrm{d}F/\mathrm{d}t < 0$. Poor economies get stuck in the region labelled "0, 0". Rich, complex economies settle upon limit cycles near where green lies below red at the top of the diagram.

However, there are pitfalls in reaching this rich state. One pitfall is the "overshooting" described above. Another is to decrease the cost α of each attempted input. The marginal cost α is exogenous but could change if, for example, communications technology makes it easier to arrange alternative suppliers. Decreasing α can trigger an escape from the poverty trap if it puts the economy in the green region in Fig. 5. But it can also make the economy more dysfunctional: if α is decreased into the red region, where agents choose $m^* = \tau^* = 1$, then dysfunction rises (provided that functional agents attempt to produce at a faster rate, i.e., $L > 1$). The intuition is that decreasing α incentivises people to attempt more complex production that uses more inputs (higher τ^*), which can lead to more failure than success, resulting in more frequent dysfunction in the new steady state [lower $F(t)$]. If policymakers sense this feedback, then they may avoid actions that decrease α, keeping the economy stuck in the trap.

Input inventories rise and then fall as economies become more complex

There is scant data—especially in developing countries—on supply-chain disruptions and on responses to them. Relevant data from the World Bank's Enterprise Survey are the number of days of inventory that firms keep of their main (i.e., highest-value) input [4]. Stockpiling inputs is one costly way to mitigate the risk of disruptions in one's supply chain [34], so it loosely corresponds to our model's buffer against disruption $m^* - \tau^*$. Macroeconomic research has focused on inventories of finished goods, but inventories of inputs have drawn increasing attention [58, 59], and some models of input inventories also consider intermediate goods and supply chains [60, 61].

Recall from Fig. 3 that our model predicts that buffers against disruptions to inputs $(m^* - \tau^*)$ tend to rise and then fall as economies develop. To test it qualitatively, we plot in Figure 6(a) the input inventory of firms averaged at the country level, for 95 countries for which we have an estimate of the complexity of the economy [9]. This Economic Complexity Index is calculated from the bipartite network of countries and of the products that they export [8]. We find that input inventory has a statistically significant inverted-U relationship with the complexity C of an economy's production [a least-squares fit of inventory with $\gamma_0 + \gamma_1 C + \gamma_2 C^2$ has $\hat{\gamma}_2 = -3.14$, with p-value $= 0.02$; see Fig. 6(a)]. This relationship qualitatively matches the inverted-U exhibited by the model [Fig. 6(b)].

Discussion

Poverty traps have long been used to explain disparities of incomes across countries and to justify a "big push," a coordinated investment in many sectors to unleash growth [39, 62, 63, 64]. Yet many big pushes have failed [65]. Our model suggests a reason why. Disruptions in supply chains, broadly defined, can spread contagiously. This systemic fragility can cause complex technologies to fail. Even if all firms coordinate their industrialisation (as suggested by big push theories [39, 63]), if the firms jump too far in technological complexity with insufficient buffers against disruptions, then the economy can slide backward, becoming poorer and less reliable. This work suggests that development policies be gradual; as in other complex systems [66], going slower may re-

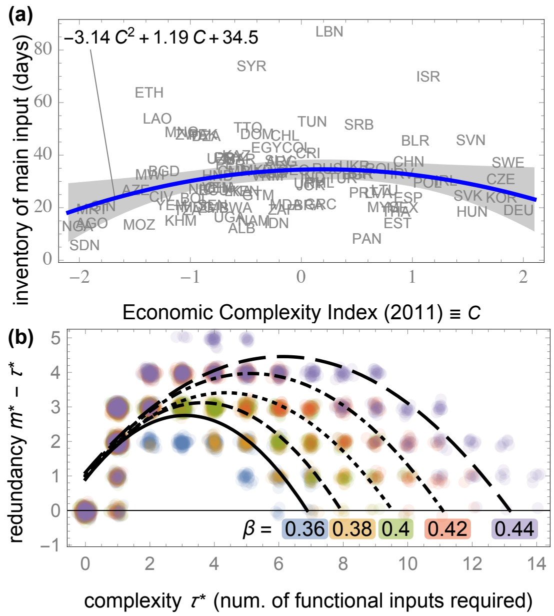

Figure 6: Qualitative match between (a) empirical data on input inventories and (b) the model's prediction that buffers to supply-chain disruptions rise and then fall as economies develop. (a) Input inventories of firms, averaged at the country level [4], have an inverted-U relationship with the complexity of the economy [9] (p-value 0.022 for the C^2 coefficient; $R^2 = 0.063$; 95 countries; 95% mean prediction band shown in gray). (b) Redundancy versus complexity for $\alpha := 0.1$ and five values of β, each a different colour. The curves show least-squares fits to $\delta_0 + \delta_1\tau^* + \delta_2\tau^{*2}$. The data is dispersed by $\mathcal{N}(\mathbf{0}, 0.008 \times \mathbf{1})$ to indicate density.

sult in collectively going faster. Unreliability affects economic performance in a multifaceted way, involving risk [67], network contagion, technology adoption [40], and psychology [68]. Understanding their interplay can elucidate the causes of persistent poverty.

Acknowledgments C.D.B. and M.H.B. acknowledge funding from the James S. McDonnell Foundation for the Postdoctoral Award and the Scholar Award (respectively) in Complex Systems.

Author Contributions All authors designed research, performed research, and wrote the paper; C.B. and K.H. analyzed data.

Competing Interests The authors declare that they have no competing financial interests.

Data and code availability All the empirical data and code used to produce the results in this paper are available at Github (DOI 10.5281/zenodo.60583).

References

[1] Schumpeter, J. A. *The Theory of Economic Development: An Inquiry Into Profits, Capital, Credit, Interest, and the Business Cycle*, vol. 55 of *Economics Third World studies* (Transaction Publishers, 1934).

[2] Romer, P. M. Growth based on increasing returns due to specialization. *The American Economic Review* **77**, 56–62 (1987).

[3] Allcott, H., Collard-Wexler, A. & O'Connell, S. D. How Do Electricity Shortages Affect Industry? Evidence from India. NBER Working Papers 19977, National Bureau of Economic Research, Inc (2014). URL `https://ideas.repec.org/p/nbr/nberwo/19977.html`.

[4] The World Bank. Enterprise Survey (2015). URL `http://www.enterprisesurveys.org/data/`. Accessed October 15, 2015.

[5] Chaudhury, N., Hammer, J., Kremer, M., Muralidharan, K. & Rogers, F. H. Missing in action: teacher and health worker absence in developing countries. *The Journal of Economic Perspectives* **20**, 91–116 (2006). URL `https://www.aeaweb.org/articles.php?doi=10.1257/089533006776526058`.

[6] Caselli, F. Accounting for cross-country income differences. In Aghion, P. & Durlauf, S. N. (eds.) *Handbook of Economic Growth*, vol. 1, 679–741 (Elsevier, 2005).

[7] Hidalgo, C. A., Klinger, B., Barabási, A. L. & Hausmann, R. The Product Space Conditions the Development of Nations. *Science* **317**, 482–487 (2007).

[8] Hidalgo, C. A. & Hausmann, R. The building blocks of economic complexity. *Proceedings of the National Academy of Sciences* **106**, 10570–10575 (2009).

[9] Atlas of Economic Complexity. Economic Complexity Index (2011). URL http://atlas.cid.harvard.edu/rankings/country/download/. Accessed November 8, 2015.

[10] World Bank. GDP per capita (in current US$) for most recent year available (2016). URL http://data.worldbank.org/indicator/NY.GDP.PCAP.CD. Downloaded February 15, 2016.

[11] Hilft, B. E. Worldriskreport. Tech. Rep., United Nations University Institute for Environment and Human Security (2014). URL http://i.unu.edu/media/ehs.unu.edu/news/4070/11895.pdf. Data "WorldRiskIndex" from pages 64–66.

[12] World Health Organization. Adult mortality rate (2013). URL http://apps.who.int/gho/data/node.main.11?lang=en. Accessed February 15, 2016.

[13] Powell, B. The global supply chain: So very fragile. *Fortune* (2011). URL http://fortune.com/2011/12/12/the-global-supply-chain-so-very-fragile/.

[14] Punter, A. Supply Chain Failures. Tech. Rep., Association of Insurance and Risk Managers in Industry and Commerce (Airmic), London, UK (2013). URL http://www.airmic.com/sites/default/files/supply_chain_failures_2013_FINAL_web.pdf.

[15] Ciccone, A. Input Chains and Industrialization. *The Review of Economic Studies* **69**, 565–587 (2002). URL http://www.jstor.org/stable/1556710.

[16] Jones, C. I. Intermediate Goods and Weak Links in the Theory of Economic Development. *American Economic Journal: Macroeconomics* **3**, 1–28 (2011).

[17] Jones, C. I. Misallocation, economic growth, and input-output economics. In *10th World Congress of the Econometric Society*, 47 (2011).

[18] Acemoglu, D., Carvalho, V. M., Ozdaglar, A. & Tahbaz-Salehi, A. The Network Origins of Aggregate Fluctuations. *Econometrica: Journal of the Econometric Society* **80**, 1977–2016 (2012).

[19] Kremer, M. The O-Ring Theory of Economic Development. *The Quarterly Journal of Economics* **108**, 551–575 (1993).

[20] Battiston, S., Delli Gatti, D., Gallegati, M., Greenwald, B. & Stiglitz, J. E. Credit chains and bankruptcy propagation in production networks. *Journal of Economic Dynamics and Control* **31**, 2061–2084 (2007).

[21] Weisbuch, G. & Battiston, S. From production networks to geographical economics. *Journal of Economic Behavior & Organization* **64**, 448–469 (2007).

[22] Mizgier, K. J., Wagner, S. M. & Holyst, J. A. Modeling defaults of companies in multi-stage supply chain networks. *Intern. Journal of Production Economics* **135**, 14–23 (2012).

[23] Henriet, F., Hallegatte, S. & Tabourier, L. Firm-network characteristics and economic robustness to natural disasters. *Journal of Economic Dynamics and Control* **36**, 150–167 (2012).

[24] Levine, D. K. Production chains. *Review of Economic Dynamics* **15**, 271–282 (2012).

[25] Contreras, M. G. A. & Fagiolo, G. Propagation of economic shocks in input-output networks: A cross-country analysis. *Physical Review E* **90**, 062812 (2014).

[26] Hendricks, K. B. & Singhal, V. R. The effect of supply chain glitches on shareholder wealth. *Journal of Operations Management* **21**, 501–522 (2003).

[27] Hendricks, K. B. & Singhal, V. R. An empirical analysis of the effect of supply chain disruptions on long-run stock price performance and equity risk of the firm. *Production and Operations Management* **14**, 35–52 (2005).

[28] Barrot, J.-N. & Sauvagnat, J. Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks (2014). URL http://www.ssrn.com/abstract=2427421. Working paper.

[29] Wang, Y. I., Li, J. & Anupindi, R. Risky Suppliers or Risky Supply Chains? An Empirical Analysis of Sub-Tier Supply Network Structure on Firm Performance in the High-Tech Sector. *SSRN Electronic Journal* (2015). URL http://www.ssrn.com/abstract=2705654.

[30] Carvalho, V. M. From Micro to Macro via Production Networks. *The Journal of Economic Perspectives* **28**, 23–48 (2014).

[31] Duan, W.-Q. Modelling the Evolution of National Economies Based on Input–Output Networks. *Computational Economics* **39**, 145–155 (2011).

[32] Oberfield, E. Business Networks, Production Chains, and Productivity: A Theory of Input-Output Architecture (2012). URL http://ssrn.com/abstract=1967148. FRB of Chicago Working Paper No. 2011-12.

[33] Carvalho, V. M. & Voigtländer, N. Input Diffusion and the Evolution of Production Networks (2014). URL http://www.nber.org/papers/w20025. NBER Working Paper No. 20025.

[34] Tomlin, B. On the Value of Mitigation and Contingency Strategies for Managing Supply Chain Disruption Risks. *Management Science* **52**, 639–657 (2006).

[35] Aydin, G., Babich, V., Beil, D. & Yang, Z. Decentralized Supply Risk Management. In *The Handbook of Integrated Risk Management in Global Supply Chains*, 387–424 (John Wiley & Sons, Inc., Hoboken, NJ, USA, 2011). URL http://doi.wiley.com/10.1002/9781118115800.ch14.

[36] Bimpikis, K., Fearing, D. & Tahbaz Salehi, A. Multi-sourcing and miscoordination in supply chain networks (2014). URL http://web.stanford.edu/~kostasb/publications/multisourcing_miscoordination.pdf. Working paper.

[37] Ang, E., Iancu, D. A. & Swinney, R. Disruption Risk and Optimal Sourcing in Multitier Supply Networks (2014). URL http://www.ideasforleaders.com/ideas/supply-chain-networks-and-csr-sourcing. Working paper.

[38] Bakshi, N. & Mohan, S. Mitigating disruption cascades in supply networks (2015). URL http://faculty.london.edu/nbakshi/Nw.pdf. Working paper, 14919920635029043942.

[39] Murphy, K. M., Shleifer, A. & Vishny, R. W. Industrialization and the Big Push. *Journal of Political Economy* **97**, 1003 (1989).

[40] Garb, Y. & Friedlander, L. From transfer to translation: Using systemic understandings of technology to understand drip irrigation uptake. *Agricultural Systems* **128**, 13–24 (2014).

[41] Bold, T., Kaizzi, K., Svensson, J. & Yanagizawa-Drott, D. Low quality, low returns, low adoption: Evidence from the market for fertilizer and hybrid seed in uganda. CEPR Discussion Papers 10743, C.E.P.R. Discussion Papers (2015). URL http://EconPapers.repec.org/RePEc:cpr:ceprdp:10743.

[42] Sanjay, A. K. & Gupta, V. Gyandoot: Trying to improve government services for rural citizens in India. Tech. Rep. 11, eGovernment for Development (2003). URL http://www.egov4dev.org/transparency/case/gyandoot.shtml.

[43] Johnson, N. & Kotz, S. *Urn models and their application: an approach to modern discrete probability theory*. Wiley Series in Probability and Statistics: Applied Probability and Statistics Section Series (Wiley, 1977). URL http://books.google.com/books?id=ZBfvAAAAMAAJ.

[44] Woodbury, M. A. On a probability distribution. *Ann. Math. Statist.* **20**, 311–313 (1949). URL http://dx.doi.org/10.1214/aoms/1177730043.

[45] Toral, R. & Colet, P. *Stochastic Numerical Methods*, chap. 8, 235–260 (Wiley-VCH Verlag GmbH & Co. KGaA, 2014). URL http://dx.doi.org/10.1002/9783527683147.ch8.

[46] Granovetter, M. Threshold models of collective behavior. *American Journal of Sociology* **83**, 1420–1443 (1978). URL http://www.jstor.org/stable/10.2307/2778111.

[47] Watts, D. J. A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences of the United States of America* **99**, 5766–5771 (2002).

[48] Jackson, M. O. & Yariv, L. Diffusion of behavior and equilibrium properties in network games. *The American Economic Review* **97**, 92–98 (2007). URL http://www.jstor.org/stable/30034427.

[49] Huremović, K. & Vega-Redondo, F. Production networks (2014). Manuscript in preparation.

[50] Ethier, W. J. National and international returns to scale in the modern theory of international trade. *The American Economic Review* **72**, 389–405 (1982). URL http://www.jstor.org/stable/1831539.

[51] Romer, P. M. Endogenous technological change. *Journal of Political Economy* **98**, S71–S102 (1990). URL http://www.jstor.org/stable/2937632.

[52] Proviti. Managing supply chain disruption risk (2011). URL http://www.protiviti.com/ja-JP/Downloads/RiskOversight_vol19_E.pdf.

[53] Simchi-Levi, D., Schmidt, W. & Wei, Y. From superstorms to factory fires: Managing unpredictable supply-chain disruptions. *Harvard Business Review* (2014). URL https://hbr.org/2014/01/from-superstorms-to-factory-fires-managing-unpredictable-supply-chain-disruptions.

[54] Tang, C. S. Robust strategies for mitigating supply chain disruptions. *International Journal of Logistics: Research and Applications* **9**, 33–45 (2006).

[55] Gabisch, G. & Lorenz, H.-W. *Business cycle theory* (Springer-Verlag, Berlin, Heidelberg, 1989), 2 edn. URL http://link.springer.com/10.1007/978-3-642-74715-1.

[56] Rebelo, S. Real business cycle models: Past, present and future. *Scandinavian Journal of Economics* **107**, 217–238 (2005). URL http://dx.doi.org/10.1111/j.1467-9442.2005.00405.x.

[57] Francois, P. & Zabojnik, J. Trust, Social Captial, and Economic Development. *Journal of the European Economic Association* **3**, 51–94 (2005).

[58] Humphreys, B. R., Maccini, L. J. & Schuh, S. Input and output inventories. *Journal of Monetary Economics* **47**, 347–375 (2001).

[59] Iacoviello, M., Schiantarelli, F. & Schuh, S. Input and output inventories in general equilibrium. *International Economic Review* **52**, 1179–1213 (2011).

[60] Khan, A. & Thomas, J. K. Inventories and the Business Cycle: An Equilibrium Analysis of (S, s) Policies. *The American Economic Review* **97**, 1165–1188 (2007).

[61] Wen, Y. Input and Output Inventory Dynamics. *American Economic Journal: Macroeconomics* **3**, 181–212 (2011).

[62] McArthur, J. *et al.* Ending Africa's Poverty Trap. *Brookings Papers on Economic Activity* **1**, 117–240 (2004). URL http://www.brookings.edu/about/projects/bpea/papers/2004/africa-poverty-trap-sachs.

[63] Azariadis, C. & Stachurski, J. Poverty Traps. In Aghion, P. & Durlauf, S. (eds.) *Handbook of Economic Growth*, 295–384 (Elsevier, 2005). URL http://hdl.handle.net/11343/34380.

[64] Barrett, C. B., Garg, T. & McBride, L. Well-being dynamics and poverty traps. *Annual Review of Resource Economics* **8** (2016). URL http://www.annualreviews.org/doi/abs/10.1146/annurev-resource-100815-095235. http://www.annualreviews.org/doi/pdf/10.1146/annurev-resource-100815-095235.

[65] Easterly, W. Reliving the 1950s: the big push, poverty traps, and takeoffs in economic development. *Journal of Economic Growth* **11**, 289–318 (2006).

[66] Gershenson, C. & Helbing, D. When slower is faster. *Complexity* **21**, 9–15 (2015). URL http://dx.doi.org/10.1002/cplx.21736.

[67] Morduch, J. Poverty and Vulnerability. *The American Economic Review* **84**, 221–225 (1994). URL http://www.jstor.org/stable/2117833.

[68] Haushofer, J. & Fehr, E. On the psychology of poverty. *Science* **344**, 862–867 (2014). URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2014Sci...344..862H&link_type=EJOURNAL.

Supplementary Information

SI-1 Derivation of the ODE (2)

Suppose that we have a large number of agents $N \in \mathbb{N}_0$, and let $n(t)$ be the number of agents who are functional at time t, so that

$$f(t) \equiv \frac{n(t)}{N}$$

is the fraction of agents who are functional at time t. $f(t)$ is a stochastic process. Here we derive a mean-field approximation for the master equation of $f(t)$, which is an ODE for the expected fraction of functional agents at time t, $F(t) \equiv \mathbb{E}\, f(t)$.

SI-1 Single agent

Attempts at producing are events that occur randomly according to a Markov process. Functional agents attempt to produce at rate L, while dysfunctional agents attempt to produce at rate 1.

First we focus attention on a single agent. Let $\mathrm{d}t$ be a small, positive amount of time. Let $\omega(f \to d)\mathrm{d}t$ be the chance that an agent that is functional at time t is dysfunctional at time $t+\mathrm{d}t$. Similarly, let $\omega(d \to f)\mathrm{d}t$ be the probability that a dysfunctional agent at time t is functional at time $t+\mathrm{d}t$.

There are many ways in which a functional agent at time t could become dysfunctional at time $t+\mathrm{d}t$. In that short amount of time, this agent could do the following transitions:

- $\mathtt{F} \to \mathtt{D}$ (the agent was chosen once to attempt to produce and failed to produce),
- $\mathtt{F} \to \mathtt{D} \to \mathtt{D}$ (the agent was chosen twice to produce and failed in both attempts),
- $\mathtt{F} \to \mathtt{F} \to \mathtt{D}$ (the agent was chosen twice to produce, succeeded in the first attempt, and failed in the second attempt),
- $\mathtt{F} \to \mathtt{D} \to \mathtt{F} \to \mathtt{D}$ (the agent was chosen three times to produce and succeeded only in the second of those attempts),
- ...and so on.

The chance of the first event is the chance that this agent is chosen to attempt to produce only once in the amount of time $\mathrm{d}t$, which occurs with probability $L\mathrm{d}t$, times the chance that the agent fails to produce in that attempt, which is $1-P$. The variable P is shorthand for the chance of successfully producing:

$$P\left[m, \tau, n(t)/N\right] \equiv \mathbb{P}\left[\mathrm{Binomial}(m, n(t)/N) \geq \tau\right].$$

The probabilities of all the other events are $O(\mathrm{d}t^2)$ because they require getting chosen at least twice to attempt to produce. Therefore, to first approximation in this small, positive amount of time $\mathrm{d}t$, the chance that a functional agent at time t is dysfunctional at time $t+\mathrm{d}t$ is

$$\omega(\mathtt{F} \to \mathtt{D})\mathrm{d}t = L\mathrm{d}t \times (1-P) + O(\mathrm{d}t^2). \tag{SI-5a}$$

Similarly, the chance $\omega(\mathtt{D} \to \mathtt{F})\mathrm{d}t$ that a dysfunctional agent at time t is functional at time $t+\mathrm{d}t$ is the chance $1 \times \mathrm{d}t$ that this agent is chosen to produce once in the time interval $(t, t+\mathrm{d}t)$ times the probability P of successfully producing in that one attempt, plus higher-order terms $O(\mathrm{d}t^2)$. Thus,

$$\omega(\mathtt{D} \to \mathtt{F})\mathrm{d}t = \mathrm{d}t \times P + O(\mathrm{d}t^2). \tag{SI-5b}$$

In the limit $\mathrm{d}t \to 0$, the events involving multiple jumps in the time interval $(t, t+\mathrm{d}t)$ occur with vanishing probability, so we approximate the system using only the terms linear in $\mathrm{d}t$ in (SI-5).

SI-2 Global rates 1

Now we consider a system of many agents. Again, let $\mathrm{d}t$ be a small, positive amount of time. Let $\Omega(n \to n+1)\mathrm{d}t$ be the chance that a system with n functional agents at time t has $n+1$ functional agents at time $t+\mathrm{d}t$. To first-order in $\mathrm{d}t$, we have

$$\Omega(n \to n+1)\mathrm{d}t = (N-n)\omega(\mathtt{D} \to \mathtt{F})\mathrm{d}t + O(\mathrm{d}t^2) \tag{SI-6a}$$

because each of the $N-n$ dysfunctional agents at time t becomes functional at time $t+\mathrm{d}t$ with probability $\omega(\mathtt{D} \to \mathtt{F})\mathrm{d}t$. Here we are neglecting the probabilities of events that are of order $O(\mathrm{d}t^2)$, such as the event in which two agents changing $\mathtt{D} \to \mathtt{F}$ and one changes from $\mathtt{F} \to \mathtt{D}$ during the time interval $(t, t+\mathrm{d}t)$. Similarly, the chance that a system with n functional agents at time t has $n-1$ functional agents at time $t+\mathrm{d}t$ is

$$\Omega(n \to n-1)\mathrm{d}t = n\omega(\mathtt{F} \to \mathtt{D})\mathrm{d}t + O(\mathrm{d}t^2). \tag{SI-6b}$$

2

SI-3 Mean-field approximation of the master equation 3

From [45, Eq. 8.95], the expected number $\mathbb{E}\,n(t)$ of functional agents at time t changes over time according to

$$\frac{\mathrm{d}\,\mathbb{E}\,n(t)}{\mathrm{d}t} = -\sum_{\ell} \ell\,\mathbb{E}[\Omega(n(t) \to n(t)-\ell)] \tag{SI-7}$$

where $\Omega[n(t) \to n(t)-\ell]$ is the instantaneous rate at which the system jumps from $n(t)$ functional agents to $n(t)-\ell$ functional agents. In our case, by combining (SI-5) and (SI-6) in (SI-7) and dividing both sides of the equation by N, we have

$$\frac{\mathrm{d}\,\mathbb{E}\,f(t)}{\mathrm{d}t} = -L\,\mathbb{E}\left[f(t)(1-P)\right] + \mathbb{E}\left[(1-f(t))P\right] + O(\mathrm{d}t^2).$$

The *mean-field approximation* [45, page 255] is that fluctuations of $f(t) \equiv n(t)/N$ can be ignored because the number of nodes N is large, so $\mathbb{E}\left[f(t)^k\right] \approx \left[\mathbb{E}\,f(t)\right]^k$. Thus, because P is a polynomial in f, we approximate $\mathbb{E}\,P \approx \mathbb{P}\left[\mathrm{Binomial}(m, F(t)) \geq \tau\right]$, which (for simplicity) is how we defined the chance of success P in the paper. Neglecting higher-order terms $O(\mathrm{d}t^2)$ gives (2a):

$$\frac{\mathrm{d}F(t)}{\mathrm{d}t} = -FL(1-P) + (1-F)P.$$

SI-2 Strategies that could be a best response 4

Recall that the agents' decision problem is to maximise the utility (3)

$$\mathcal{U}\left[m, \tau, F(t); \alpha, \beta\right] = \mathbb{P}\left[\mathrm{Binomial}(m, F) \geq \tau\right]\tau^{\beta} - \alpha m$$

over all pairs of non-negative integers $(m, \tau) \in \mathbb{N}^2$, where $\mathbb{N} \equiv \{0, 1, 2, 3, \ldots\}$. The set of strategies (m, τ) 5
that could be a best response turns out to be a finite set, which enables numerical simulations. 6

Lemma 1 *A best response (m^*, τ^*) must belong to the set*

$$\{(m, \tau) \in \mathbb{N}^2 : m = 0 \text{ or } 0 < \tau \leq m < \tau^{\beta}/\alpha\}. \tag{SI-8}$$

7

PROOF **1** *First observe that if $m = 0$ then the utility is zero. If $0 < m < \tau$, then the utility is negative, so a best response cannot have $0 < m < \tau$. Hence a best response must have $m = 0$ or $m \geq \tau > 0$.*

In the latter case (with $m \geq \tau > 0$), the utility must exceed zero (the utility obtained with $m = 0$), so

$$\mathbb{P}\left[\text{Binomial}(m, F) \geq \tau\right] \tau^{\beta} - \alpha m > 0,$$

or, after rearranging and dividing by $\alpha > 0$,

$$m < \mathbb{P}\left[\text{Binomial}(m, F) \geq \tau\right] \frac{\tau^{\beta}}{\alpha} \leq \frac{\tau^{\beta}}{\alpha}. \tag{SI-9}$$

■

Corollary 1 *A best response (m^*, τ^*) with $m^* > 0$ must satisfy $\tau^* \leq \alpha^{-1/(1-\beta)}$ and $m^* < \alpha^{-1/(1-\beta)}$.*

PROOF **2** *From Lemma 1, we know that if (m, τ) is a best response with $m > 0$ then $\tau \leq m < \tau^{\beta}/\alpha$. Now we equate these lower and upper bounds on m. The equation $\tau = \tau^{\beta}/\alpha$ has a two solutions: $\tau = 0$ and $\tau = \alpha^{-1/(1-\beta)}$. The latter provides provides an upper bound on τ^* in the best response. From Lemma 1 we know that a best response must have $m^* < \tau^{*\beta}/\alpha$, so $m^* < \tau^{*\beta}/\alpha \leq \alpha^{-1/(1-\beta)}$.* ■

Numerical calculations of a best response can be sped up slightly using the following two observations:

1. If $m = \tau > 0$ is a best response, then $\mathcal{U}\left[m, m, F(t); \alpha, \beta\right] = F(t)^m m^{\beta} - \alpha m > \mathcal{U}\left[0, 0, F(t); \alpha, \beta\right] = 0$, which can be rewritten as

$$m < \frac{(\beta - 1)}{\log\left[F(t)\right]} W\left(\frac{\left(\frac{1}{\alpha}\right)^{\frac{1}{1-\beta}} \log\left[F(t)\right]}{\beta - 1}\right), \tag{SI-10}$$

   where W is the product logarithm (i.e., the Lambert W function). This upper bound on the diagonal $m = \tau$ can be tighter than the one in inequality (SI-9).

2. If $F(t) = 1$, then every strategy with $m \geq \tau$ will certainly succeed [i.e., $\mathbb{P}\left[\text{Binomial}(m, F(t)) \geq \tau\right] = 1$], so the best response must have $m = \tau$, and the first-order condition for the utility $m^{\beta} - \alpha m$ shows that the best-response $m = \tau$ is either the floor or the ceiling of $(\beta/\alpha)^{1/1-\beta}$ (whichever one results in more utility).

If practice, if the best response is $(0, \tau)$, where $\tau \in \mathbb{N}$, then we take the best response to be $(0, 0)$ because it is not economically meaningful to have $\tau > m = 0$ (it would mean that an agent requires some inputs to produce but does not attempt to acquire any inputs).

In simulations, we compute the set of strategies that could be best responses as follows:

$$\begin{cases} \{(0,0)\} & \text{if } F = 0 \\ \{(\lfloor\gamma\rfloor, \lfloor\gamma\rfloor), (\lceil\gamma\rceil, \lceil\gamma\rceil)\} & \text{if } F = 1 \\ \{(m, \tau) \in \mathbb{N}^2 : 0 \leq m \leq \alpha^{-1/(1-\beta)}, \text{ and either} & \\ m = \tau \text{ satisfies (SI-10) or } 0 < \tau < m < \tau^{\beta}/\alpha & \text{if } 0 < F < 1 \end{cases} \tag{SI-11}$$

where

$$\gamma \equiv \left(\frac{\beta}{\alpha}\right)^{\frac{1}{1-\beta}}.$$

Given $(\alpha, \beta, F(t))$, we compute the finite set of strategies that could be a best response [given by expression (SI-11)] and among those strategies we select the strategy with the highest utility [(3)].

SI-3 Downturns in rich, highly functional economies

Here we show that rich, highly functional economies undergo occasional spikes in dysfunction, akin to economic recessions. In this section we assume that functional agents attempt to produce at a rate that is greater than or equal to the rate at which dysfunctional agents attempt to produce (i.e., we assume that $L \geq 1$).

Lemma 2 (A completely functional economy is in steady state) *For $L \geq 1$, there is a steady state at $F(t) = 1$, at which agents best respond by choosing $m^* = \tau^*$. Furthermore, $m^* = \tau^*$ equals the floor or ceiling of $(\beta/\alpha)^{1/(1-\beta)}$ (whichever yields more utility).*

Proof **3** *In a completely functional economy [with $F(t) = 1$], the ODE for $F(t)$ [from (2)] is*

$$\mathrm{d}F/\mathrm{d}t = [1 - F(t)]P - LF(t)(1 - P) = -L(1 - P) \tag{SI-12}$$

for $\tau > 0$ (and $\mathrm{d}F/\mathrm{d}t := 0$ for $\tau = 0$). Because $F = 1$, agents certainly succeed in producing (i.e., $P = 1$) as long as $m \geq \tau > 0$. If $P = 1$, then choosing m to be greater than τ presents only costs and no benefits, so agents choose $m^* = \tau^*$. By computing the value of τ that maximises the utility function $\tau^\beta - \alpha\tau$, we see that the best response at $F(t) = 1$ has $m^* = \tau^*$ equal to either the floor or ceiling of $(\beta/\alpha)^{1/(1-\beta)}$ (whichever yields more utility).

In summary, if $m^* = \tau^* = 0$ then $\mathrm{d}F/\mathrm{d}t = 0$ [by definition; see (2)], and if $m^* = \tau^* > 0$ then $P = 1$ and therefore $\mathrm{d}F/\mathrm{d}t = 0$ from (SI-12). ∎

Next we show in Lemma 3 that the steady state at $F(t) = 1$ is not stable to perturbations in $F(t)$.

Lemma 3 (The best response has no redundancy at and just below $F(t) = 1$) *Assume that $L \geq 1$, $\alpha > 0$, and $\beta \in (0, 1)$. There exists $\bar{F} \in (0, 1)$ such that the best response (m^*, τ^*) satisfies $m^* = \tau^*$ for $F(t) \in (\bar{F}, 1]$.*

Proof **4** *We will prove the claim using the continuity of the utility function $\mathcal{U}[m, \tau, F(t); \alpha, \beta]$ [defined in (3)] as a function of $F(t)$. $\mathcal{U}$ is continuous in $F(t)$ because it is a polynomial in $F(t)$ of degree m.*

From Lemma 2, we know that for $F(t) = 1$ the best response is $m^* = \tau^*$ and that τ^* equals the floor or the ceiling of $(\beta/\alpha)^{1/(1-\beta)}$ (whichever yields more utility). We denote this value of τ^* by γ.

From Lemma 1, we know that the set of strategies that could be a best response is finite and depends on α and on β but not on F. Denote using Γ the set of strategies that could be a best response, which is given in (1).

Let ϵ_k denote the difference between the first-best utility and kth-best utility, where the k-th best utility is defined to be the kth-highest utility among the set of strategies Γ that could be best responses.

Let $\epsilon := 1/2 \times \min\{\epsilon_k : 2 \leq k \leq |\Gamma|\}$. We know that $\epsilon > 0$ because it is a minimum of finitely many positive numbers (because $|\Gamma| < \infty$). Because $\mathcal{U}$ is continuous in $F(t)$, there exists $\delta \in (0, 1)$ such that if $F(t) > 1 - \delta$ then

$$|\mathcal{U}[m, \tau, F(t); \alpha, \beta] - \mathcal{U}[m, \tau, 1; \alpha, \beta]| < \epsilon \tag{SI-13}$$

for all $(m, \tau) \in \Gamma$.

Now fix $F_0 \in (1-\delta, 1)$ and let $(m_0, \tau_0) \in \Gamma$ be one of the possible best responses different from (γ, γ) [which is the best response at $F(t) = 1$]. Let k denote the utility-rank of this strategy (m_0, τ_0) at $F(t) = 1$, meaning that the strategy (m_0, τ_0) has the k-th highest utility at $F(t) = 1$. Then

$$\begin{aligned}
\mathcal{U}[\gamma, \gamma, F_0; \alpha, \beta] &> \mathcal{U}[\gamma, \gamma, 1; \alpha, \beta] - \epsilon && \text{by continuity of } \mathcal{U}[\ldots, F(t), \ldots] && \text{(SI-14a)}\\
&\geq \mathcal{U}[\gamma, \gamma, 1; \alpha, \beta] - \epsilon_k/2 && \text{by definition of } \epsilon && \text{(SI-14b)}\\
&= \mathcal{U}[m_0, \tau_0, 1; \alpha, \beta] + \epsilon_k/2 && \text{by definition of } \epsilon_k && \text{(SI-14c)}\\
&\geq \mathcal{U}[m_0, \tau_0, 1; \alpha, \beta] + \epsilon && \text{by definition of } \epsilon && \text{(SI-14d)}\\
&> \mathcal{U}[m_0, \tau_0, F_0; \alpha, \beta] && \text{by continuity of } \mathcal{U}[\ldots, F(t), \ldots], && \text{(SI-14e)}
\end{aligned}$$

so (γ,γ) *is the best response for all* $F \in (1-\delta, 1]$*. This argument completes the proof with* $\bar{F} := 1-\delta \in (0,1)$*.*
$\blacksquare$

Finally, we arrive at our main result of this appendix: if the marginal cost of each attempted input (α) is small enough, then a sufficiently functional economy undergoes spikes in disruptions, akin to recessions.

Theorem 1 (Highly functional economies undergo spikes in disruption) *If* $\alpha < 2^\beta - 1$*,* $\beta \in (0,1)$*, and* $L \geq 1$*, then there exists* $\widetilde{F} \in (0,1)$ *such that* $\mathrm{d}F/\mathrm{dt} < 0$ *for* $F \in (\widetilde{F}, 1)$*.*

Proof **5** *From Lemma 3 we know that there exists* $\bar{F} \in (0,1)$ *such that the best response* (m^*, τ^*) *satisfies* $m^* = \tau^*$ *for* $F \in (\bar{F}, 1]$*, and from Lemma 2 we know that this* $m^* = \tau^*$ *equals the floor or ceiling of* $(\beta/\alpha)^{1/(1-\beta)}$ *(whichever yields more utility), which we denote by* γ*. Note that* γ *is nonincreasing in* α*.*

Next we equate utilities from pairs of strategies:

- *An agent is indifferent between the strategies* $(m,\tau) = (1,1)$ *and* $(0,0)$ *when* $F(t) - \alpha = 0$*. This observation, together with the observation that* γ *is nonincreasing in* α*, implies that for* $F(t) = 1$ *and* $\alpha \geq 1$ *the best response is* $(0,0)$*.*
- *An agent is indifferent between the strategies* $(m,\tau) = (1,1)$ *and* $(2,2)$ *when* $F^2 2^\beta - 2\alpha = F(t) - \alpha$*, which is a quadratic polynomial in* $F(t)$*. The positive root is* $F_+ := 2^{-(\beta+1)}\left(1 - \sqrt{1 - \alpha 2^{\beta+2}}\right)$*. This positive root* F_+ *is equal to* 1 *when* $\alpha = 2^\beta - 1$*, which is smaller than* 1 *for any* $\beta \in (0,1)$*. Therefore, for* $\alpha \in (2^\beta - 1, 1)$ *and* $F(t) = 1$ *the best response is* $(m^*, \tau^*) = (1,1)$*. Also, for* $\alpha \in (0, 2^\beta - 1)$ *and* $F(t) = 1$ *the best response* (m^*, τ^*) *satisfies* $m^* = \tau^* \geq 2$*.*

Finally, note that if $m = \tau \geq 2$*, then the chance of successfully producing is* $P\left[m, \tau, F(t)\right] = F(t)^\tau$*, which is less than one for* $F(t) \in (0,1)$ *because* $\tau \geq 2$*. Therefore, for* $L = 1$ *the differential equation* $\mathrm{d}F/\mathrm{dt} = F(t)^\tau - F(t)$ *is* < 0 *for all* $0 < F(t) < 1$*.*

It remains only to show the claim for $L > 1$*. If* $m = \tau \geq 2$*, then*

$$\mathrm{d}F/\mathrm{dt} = 0 \text{ at } F(t) = 1, \tag{SI-15}$$

and the slope of the ODE at $F(t) = 1$ *is positive:*

$$\frac{\partial}{\partial F}\frac{\mathrm{d}F}{\mathrm{dt}}\bigg|_{F=1} = mL - 1 \geq 1 > 0. \tag{SI-16}$$

Combining (SI-15) *and* (SI-16) *with the fact that* $\mathrm{d}F/\mathrm{dt}$ *is a polynomial in* $F(t)$*, we conclude that there exists* $\widehat{F} \in (0,1)$ *such that* $\mathrm{d}F/\mathrm{dt} < 0$ *for* $F(t) \in (\widehat{F}, 1)$*, and the claim holds for* $\widetilde{F} := \max\{\widehat{F}, \bar{F}\}$*.* $\blacksquare$