

Quantification of the resilience of primary care networks by stress testing the health care system

Donald Ruggiero Lo Sardo^{a,b}, Stefan Thurner^{a,b,c,d,1}, Johannes Sorger^b, Georg Duftscheid^e, Gottfried Endel^f, and Peter Klimek^{a,b}

^aSection for Science of Complex Systems, Center for Medical Statistics, Informatics and Intelligent Systems (CeMSIS), Medical University of Vienna, A-1090 Vienna, Austria; ^bComplexity Science Hub Vienna, A-1080 Vienna, Austria; ^cInternational Institute for Applied Systems Analysis (IIASA), A-2361 Laxenburg, Austria; ^dSanta Fe Institute, Santa Fe, NM 85701; ^eSection for Medical Information Management, CeMSIS, Medical University of Vienna, A-1090 Vienna, Austria; and ^fMain Association of Austrian Social Security Institutions, A-1030 Vienna, Austria

Edited by Timothy George Buchman, Emory University School of Medicine, Atlanta, GA, and accepted by Editorial Board Member Simon A. Levin October 1, 2019 (received for review March 27, 2019)

There are practically no quantitative tools for understanding how much stress a health care system can absorb before it loses its ability to provide care. We propose to measure the resilience of health care systems with respect to changes in the density of primary care providers. We develop a computational model on a 1-to-1 scale for a countrywide primary care sector based on patient-sharing networks. Nodes represent all primary care providers in a country; links indicate patient flows between them. The removal of providers could cause a cascade of patient displacements, as patients have to find alternative providers. The model is calibrated with nationwide data from Austria that includes almost all primary care contacts over 2 y. We assign 2 properties to every provider: the “CareRank” measures the average number of displacements caused by a provider’s removal (systemic risk) as well as the fraction of patients a provider can absorb when others default (systemic benefit). Below a critical number of providers, large-scale cascades of patient displacements occur, and no more providers can be found in a given region. We quantify regional resilience as the maximum fraction of providers that can be removed before cascading events prevent coverage for all patients within a district. We find considerable regional heterogeneity in the critical transition point from resilient to nonresilient behavior. We demonstrate that health care resilience cannot be quantified by physician density alone but must take into account how networked systems respond and restructure in response to shocks. The approach can identify systemically relevant providers.

coevolving networks | dynamics of collapse | robustness | quality of care | patient-sharing network

For the last 50 y, health-related expenditures in almost all western countries have been growing faster than national incomes (gross domestic product) (1). This has raised concerns about the sustainability of health care systems all across the Organization for Economic Cooperation and Development (2). In several developed countries, health care demand will further increase, because the population is aging and the prevalence of chronic disorders is increasing (3). The situation is exacerbated by impending retirement waves (4). Is there a point beyond which these pressures will severely impair the quality of care? If so, how close are we to it? To answer these questions, a quantitative understanding of the resilience of health care systems is required. Resilience quantifies the rate of recovery and the extent to which a system is able to recover from disruptive events (5). In health care systems, such events include sudden increases in patient numbers or reductions in the number of health care providers within a specific region. Resilience captures how fast and the extent to which it is possible to deliver adequate health services to the entire population in the wake of such a shock.

Health care is undergoing a digital revolution (6) driven by the increasing availability of observational health care data (7). As countries adopt systems of shared electronic health records,

such data become available at national scales (8). These systems enable analysts to answer questions, such as “Who did what, when, for whom, where and at what costs?” for practically all medical services in a given country. For instance, in Austria, it has been shown that such data can be used to identify genetic, environmental, and epigenetic disease risks (9, 10) or to investigate how individual health care providers coordinate with each other in the treatment of patients (11–14). Health care providers are embedded in multiple formal and informal relationships, because they share information or patients. These relationships can be formalized in so-called patient-sharing networks, which consist of nodes (providers) connected by links if they share the same patients (15, 16). Such networks show large structural variations that can be related to differences in the cost and quality of care (17–19).

Here, we show how to quantify the resilience of health care systems with respect to changes in local densities of health care providers by means of dynamical simulations of structural changes in the patient-sharing networks. Further, we show how this method can be used to benchmark and stress test 121 Austrian regions (political districts) in terms of their resilience. The central idea of our approach is as follows.

Significance

We shock a full-scale simulation model of a national health care system by locally removing health care providers. We measure resilience of the system in terms of how fast and to what extent it can recover its ability to deliver adequate health services to the population. The model is based on actual regional primary care networks in Austria, where all patients and physicians are represented as anonymized avatars that are calibrated with nationwide data. After removal of a critical fraction of physicians, networks generically undergo a transition from resilient to nonresilient behavior, where it is impossible to maintain coverage for all patients. These “stress tests” allow us to quantify regional health care resilience and identify systemically risky health care providers.

Author contributions: D.R.L.S., S.T., and P.K. designed research; D.R.L.S. performed research; J.S. and G.D. contributed analytic tools; D.R.L.S., S.T., J.S., G.D., G.E., and P.K. analyzed data and reviewed the manuscript; and D.R.L.S., S.T., J.S., and P.K. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission. T.G.B. is a guest editor invited by the Editorial Board.

This open access article is distributed under Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 (CC BY-NC-ND).

¹To whom correspondence may be addressed. Email: stefan.thurner@muv.ac.at.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1904826116/-DCSupplemental.

First published November 11, 2019.




Fig. 1. Schematic representation of patient displacement dynamics. (A) Doctors are represented as nodes (size represents the number of patients treated per year). They are linked if they share patients in the patient-sharing network, A (black arrows). The color represents their current capacity; green means that they have capacity, and red means that they can no longer accept new patients. (B) Doctor a retires at time step 1; his/her patients are distributed to other doctors according to the weights of the links from a to b and from a to c (yellow arrows). This, in turn, changes the capacity of the other doctors. (C) As c has reached its capacity limit (red), he/she must send patients to other doctors (blue arrows from c to b and d). This creates a cascade of patient displacements of size 2. D – F show the same steps as in A – C in a simulation of a realistic environment. Doctors are localized (due to data protection) at random locations within a district, and a real patient-sharing network is used. (E) A doctor is removed, and his/her patients are shared (yellow). (F) Those doctors who reach their capacity send excess patients to others in a second round (blue). At this point, all patients are cared for, and the model dynamic terminates.

Consider 4 physicians a , b , c , and d who share patients with each other: say 20% of a 's patients have also visited b , and 10% have seen c (Fig. 1). Links between physicians may arise for a multitude of reasons (e.g., because b is a 's holiday locus, a is on maternity or sick leave, etc.). How doctors share patients is given by a network, A , in which doctors are nodes connected by patient-sharing relations. Assume that, at time $t = 1$, doctor a is closed for business. As 20% of a 's patients already have a treatment relationship with b , it is natural to assume that many of these patients will now seek treatment with b . The removal of node a induces a displacement flow of patients along the link from a to b but also from a to c . By receiving new patients from a , both b and c will get closer to their maximum capacity. In Fig. 1B, this is shown by the change in node color. Green means high spare capacity; red means that the capacity limit is reached. In the example, c now has exceeded its capacity (received more patients from a than can be treated within reasonable time). Doctor c must, therefore, in the next time step send the excess number of newly inherited patients to yet other doctors (along the links in the patient-sharing network), here to b and d (Fig. 1C). Nodes b and d get closer to their limits but are still capable of absorbing more patients. The removal of doctor a leads to a cascade of patient displacements of size 2. In other cases, where doctors are closer to their limits, cascades can become large and eventually span a large region of the patient-sharing network.

A highly resilient health care system should be able to redistribute a 's patients with a minimal number of patient displacements in a short interval of time. The initial shock (a 's removal) is then quickly absorbed, and the system becomes fully functional again soon afterward (all patients find a new doctor). A nonresilient system, however, is characterized by cascades of

patient displacements that push multiple doctors beyond their capacities. If a substantial number of patients do not find a new doctor, the health care system will essentially lose its ability to deliver adequate care. We can identify 2 related indicators to distinguish resilient from nonresilient behavior. The higher the resilience of a health care system, 1) the lower the number of displacements that the removal of doctors causes, and 2) the lower the number of patients unable to find a new doctor. The systemically beneficial doctors (i.e., those who contribute most to regional resilience) are those who take over the largest shares of patients.

Cascading processes are examples of dynamical phenomena that take place on networks (20). To model such processes, a localized perturbation is considered by shocking or removing a single node. This initial event spreads via the links of the perturbed node to other nodes, which might trigger another step in a cascade as those nodes propagate the shock to their neighbors and so on. Such processes can be formulated by means of recursive centrality measures (e.g., the PageRank algorithm) (21) or models that consider load distribution on networks (22). In concrete applications, these network measures often require modifications that reflect specific properties of the system under consideration, such as the propagation of shocks between financial institutions (23), failures in power grids (24), or cascading failures in interconnected infrastructures (25).

Here, we develop a data-driven computational framework to estimate the impact of doctor removals through cascading processes of displacements on patient-sharing networks of practically all physicians in Austria. We construct patient-sharing networks $A(\delta)$ of primary care providers (PCPs) for 121 districts δ from an extensive dataset containing about 97% of all outpatient contacts over 2 y in Austria (9–12) (SI Appendix, Fig. S1 and Text S1). We formulate a dynamical model that simulates the removal of one or several providers and computes the size and duration of the resulting cascades of patient displacements in the following way.

Every PCP i is a node in the patient-sharing network with weighted directed links from i to j . The link strength, A_{ij} , corresponds to the number of patients of i who occasionally also visit PCP j . In every quarter of a year q , every PCP i treats p_i^q unique patients. The average number of unique patients who a doctor sees in a quarter is $\mu_i = \sum_{q=1}^T p_i^q / T$, where T is the total timespan of the data. Every PCP i is further characterized by a fixed capacity c_i , which is estimated from historical data. In the simplest case, we assume $c_i = (1 + C)\mu_i$, where C is a free model parameter.

The model dynamic takes place on a timescale, t , that is shorter than a quarter, say days. Initially, each patient is assigned to the PCP who he/she most frequently consulted in the past. A PCP is in 1 of 3 internal states: available, fully booked, or unavailable (removed) (SI Appendix, Text S2). Assume that a PCP i is removed from the network of district δ at time t . Those μ_i patients who usually visited PCP i now transfer to j with probability $P_{ij} = A_{ij}(\delta) / \sum_k (A_{ik}(\delta))$. We allow for the possibility that not all patient displacements follow the links of the patient-sharing network. With probability Q , patients select a random doctor in the same district with a uniform probability (SI Appendix, Text S3). To every PCP i , we assign the average number of displacements, D_i , that i 's patients must undergo before finding a new and available PCP. Ranking PCPs according to their value of D_i (from high to low) identifies physicians with the largest contributions to systemic risk; we call this rank the CareRank of a PCP. For each PCP i , we also measure average systemic benefit, B_i , which is the fraction of displaced patients who end up at i averaged over removals of all other providers in the district. The displacements, D_i , and benefits, B_i , are proxies of the systemic risk

and benefit contributions of every doctor i ; a definition is in [SI Appendix, Text S4](#).

We use this model to quantify the resilience of individual regions in which multiple PCPs are removed. The set of doctors removed at time $t=0$ is denoted by S . The size of this initial shock f is the fraction of PCPs who become unavailable at $t=0$, $f=|S|/N(\delta)$, where $N(\delta)$ is the number of doctors in district δ . Following this shock, we count the number of patients in district δ unable to find an available PCP, $L_S(f, \delta)$ ([SI Appendix, Text S4](#)). We refer to $L_S(f, \delta)$ as the number of “lost patients.” For each district, δ , we are interested in the smallest shock size, $f_c(\delta)$, for which $L_S(f \geq f_c, \delta) > 0$ holds in a certain fraction of model runs. This means that there will be patients no longer able to find primary health care services within a given district. At this critical shock size, the district has reached its “resilience point.” The parameter, $f_c(\delta)$, serves as our resilience indicator. We show that, surprisingly, the critical doctor removal density $f_c(\delta)$ is practically uncorrelated with regional physician densities, a conventional indicator to assess health care coverage (26). To explore how the resilience indicators depend on properties of the PCPs and the networks that they are embedded in, we use 2 different types of linear regression model ([SI Appendix, Texts S5 and S6](#)).

We consider 4 alternative model variants. First, we assume that doctor capacity, c_i , can be estimated from the historically observed fluctuations in a doctor’s patient numbers (i.e., c_i is proportional to the variance of p_i^q). Second, c_i is implemented as a multistep function to take differences in staffing into account (that is, physicians hire additional staff, which increases their capacity by a constant factor). Third, the next variant is equal to the main variant except that patients seeking a new PCP do not contact the same doctor twice during their search (they perform a self-avoiding random walk on the patient-sharing network). Fourth, the dynamics of the main model variant is studied on the countrywide patient-sharing network without being broken down into districts. [SI Appendix, Text S7](#) has a description of these model variants.

Results

The model dynamic is illustrated in Fig. 1. Initially (Fig. 1D), all doctors operate well below their capacity (green). At time $t=1$, PCP a becomes unavailable (Fig. 1E). His/her former patients seek a new doctor on the patient-sharing network (yellow in Fig. 1E). PCP c now is fully booked. At $t=2$, patients can no longer be accommodated by c and move from PCP c to nodes b and d (Fig. 1F). After all patients find a new

PCP, the dynamic stops. A web-based interactive visualization of a simplified version of the model on a real regional primary care network is available online (https://csh.ac.at/vis/med_public/pcn-resilience) ([SI Appendix, Text S8](#)). Structural properties of these patient-sharing networks have been reported previously (11–14).

We now study the validity of 2 central model assumptions. First, we test whether patients who lose their PCP are indeed displaced along links in the patient-sharing network. We identify as removed doctors those who had at least 100 quarterly patients on average in the first year but no patients in the second half of the second year; 28,795 patients with at least 2 different physicians were displaced this way. Of those, 84% most frequently consulted a PCP in 2007 who they had already seen in 2006. In these cases, the removal of a doctor did indeed lead to displacements along the patient-sharing network. Second, we inquire to what extent the nationwide patient-sharing network can be decomposed into individual districts; 77% of links between doctors from the same district are nonzero compared with 3.6% of links between doctors of different districts being nonzero ([SI Appendix, Fig. S1 and Text S1](#)). How these interdistrict links influence the model results is investigated in the model variant that uses the countrywide patient-sharing network.

Systemic Relevance of PCPs. We next determine the average number of patient displacements, D_i , caused by the removal of doctor i , $S=\{i\}$. As the model is not deterministic, we performed 500 model runs. The median of μ_i , the average quarterly patient number per PCP, is 945 (corresponding to about 10 patients per day) (Fig. 2A). Fig. 2B shows the relation between μ_i and average displacements D_i . The most systemically relevant PCPs cause almost 3 displacements per patient on average, while many cause slightly more displacements than the theoretical minimum of 1. Doctors with displacements close to this minimum tend to have low patient numbers within the range from 20 to 500. The majority of physicians have patient numbers between 500 and 2,500, for which we observe displacements that vary between 1 and 3. These 2 “modes” of the bivariate distribution of physicians in the $\mu_i - D_i$ plane result in a weak linear correlation (Pearson’s $R=0.52$, p value of $p < 10^{-4}$). To explore how differences in D_i relate to other network or demographic properties, we perform univariate and multivariate regression analyses ([SI Appendix, Fig. S2, Table S1, and Text S6](#)). Overall, high-impact doctors tend to have high numbers of patients, low numbers of links with high weights, low numbers of closed triadic relationships (low clustering), and low




Fig. 2. Systemic risk profile of Austrian health care providers. (A) The distribution of average quarterly patient numbers μ_i of doctors has a median of 945 patients. (B) Displacements, D_i , for every doctor in Austria tend to correlate with patient numbers μ_i of doctors (Pearson’s $R=0.52$, $p < 10^{-4}$). The color encodes the number of doctors with a given (μ_i, D_i) pair. (C) Systemic risk contributions of doctors, D_i , show only little correlation with their systemic benefit (Pearson’s $R=0.42$, $p < 10^{-4}$), B_i . The 4 quadrants indicate regions where D_i and B_i are above or below their population medians, respectively.




Fig. 3. Number of patients, $L_S(f, \delta)$, who cannot be cared for as a function of the fraction of unavailable PCPs, f , in district δ . $L_S(f, \delta)$ is shown for 2 different scenarios where doctors are removed in a different order: sequence A (green) and sequence B (blue). Labels display indexes of the removed PCPs at each step (i.e., the green sequence first removes PCP 2 followed by 13, 10, and so forth, whereas the blue sequence first removes PCP 4 followed by 7, 3, etc.). The shaded area envelops all observed values of $L_S(f, \delta)$ (100% CI). Sequence B gives a scenario where 44% (8PCPs) have to be removed before losing patients, whereas 22% (4PCPs) are sufficient to put the district in a condition where it cannot care for all patients for sequence A. The red arrow marks the position of the critical fraction f_c , which is the smallest f such that $L_S(f, \delta) > 0$ holds for each observed sequence.

centrality in the network. In Fig. 2C, we show that the systemic risk contributions of doctors, D_i , can substantially deviate from their systemic benefit, B_i (Pearson's $R = 0.42$, p value of $p < 10^{-4}$). PCPs in the upper left region of Fig. 2C combine high systemic risk contributions with low benefit, whereas the lower right region shows physicians with high benefit and relatively small risk contributions.

Resilience of Districts. After removal of a single doctor, patients typically find a new doctor in district δ , $L_S(f, \delta) = 0$; no patients are “lost.” In the situation where a larger fraction f of PCPs is removed, this can change. We now ask at which critical fraction f_c we find the onset of lost patients, $L_S(f_c, \delta) > 0$. f_c indicates the location of a regime shift in the model behavior (in the Introduction).

Fig. 3 shows the number of lost patients, $L_S(f, \delta)$, as a function of the shock size f for the district of Reutte in Tyrol. Doctors are removed sequentially. We show 2 different sequences (green and blue in Fig. 3). The smallest value of f for which $L_S(f, \delta)$ becomes nonzero depends on the sequence order. A critical f_c can be defined using the sequence that leads to the largest (upper bound, red arrow in Fig. 3) or smallest (lower bound) f_c (SI Appendix, Text S4). We perform 500 model runs (50 different choices of specific sequences S for 10 model realizations) in which $|S| = fN(\delta)$ doctors have been removed initially. In Fig. 4, the upper bound for the resilience indicator $f_c(\delta)$ for each district is encoded in the district color from green (most resilient) to red (least resilient). For most districts, the transition occurs after about 30% of the doctors are removed (SI Appendix, Fig. S3). However, there are also districts where the transition occurs for substantially smaller (about 20%) or larger (about 40%) values of f . SI Appendix, Fig. S3 shows that the width of this transition varies substantially across districts. Note that $f_c(\delta)$ depends on the choice of the capacity parameter C and is, therefore, not in itself informative unless reasonable choices are made. However, the relative ranking of individual districts by their $f_c(\delta)$ for regional comparisons can be carried out for any C .

In Fig. 5, we compare the lower bounds of the resilience scores, f_c , with the de facto standard indicator for health coverage (i.e., physician density; number of PCPs per thousand population). SI Appendix, Fig. S4 shows a similar comparison using the upper bound of f_c . In both comparisons, districts with similar resilience scores, f_c , can have physician densities that vary by up to 1 order of magnitude. The regression analysis additionally shows a negative correlation of the resilience scores with the district-averaged clustering coefficient [$CC(\delta)$ Pearson's $R = -0.48$, $p < 10^{-4}$] and a positive correlation with district-averaged closeness centrality, [$CL(\delta)$ Pearson's $R = 0.38$, $p = 0.003$] (SI Appendix, Fig. S4). Both of these correlations are confounded by the demographic properties of the districts (SI Appendix, Table S3).

Robustness. We obtain qualitatively similar results in all 4 different model variants and for a wide range of choices in the model parameters (SI Appendix, Figs. S5 and S6). Results for the patient displacements, D_i , present no qualitative change with respect to the standard variant for values of C in the range from 0.01 to 0.1. For even larger values of C , cascade sizes approach 1 for many patient displacements, whereas for smaller values, the cascades might easily span the entire system, even for small shocks. We study 2 alternative definitions of the doctor capacity, namely by inferring c_i from the observed variance of patient numbers, p_i^q , and by assuming a multistep function of capacity to take different levels of staffing into account. The model was also evaluated on the countrywide patient-sharing network (patients can cross districts) and by assuming that patients perform a self-avoiding walk on the network. Due to computational costs, particularly for the latter 2 variants, the results of these variants are compared for the doctor displacements, D_i . Overall, we find substantial correlations between all model variants, in many cases with correlation coefficients close to 1 (SI Appendix, Fig. S5). Considering pairwise comparisons between the main model and the other variants (SI Appendix, Fig. S6), we find the lowest agreement with the variance definition of doctor capacity (for very low values of C) and with the multistep variant, where we observe correlations around 50%. All other correlation coefficients fall in the range between 70 and 95%. Finally, we confirmed that the relations between our doctor- and district-level systemic risk measures, D_i and f_c , show similar correlations with other demographic and network properties as in the main model (SI Appendix, Tables S2 and S4).

Discussion

The primary aim of this paper was to quantify the resilience of regional primary care networks on a fully data-driven basis. We were able to quantify resilience on 2 scales. First, we determined the systemic relevance of individual doctors by estimating the




Fig. 4. Map of Austria that shows the upper bound of the resilience indicator, $f_c(\delta)$, for all districts. Districts colored in green (red) have a particularly high (low) resilience: that is, critical removal fractions of $f_c(\delta)$.




Fig. 5. Resilience vs. (primary care) provider density. Every circle (size is proportional to the district population) is a district with its provider density (number of PCPs per thousand population) on the x axis and the lower bound of resilience indicator, $f_c(\delta)$, on the y axis. While there is some correlation between them (Pearson's $R = 0.38$, $p < 10^{-4}$; Spearman's $R = 0.37$, $p = 0.0002$), physician density can vary by up to 1 order of magnitude for districts of similar resilience.

number of patient displacements, D_i , caused by her/his removal. Second, we were able to estimate the critical fraction, f_c , of PCPs who can be removed before the regional health care service breaks down. We developed a full-scale simulation model for how regional patient flow networks reorganize after the removal of one or several doctors. By full scale, we mean that the actual data of each Austrian patient and each PCP are represented as a fully anonymized individual avatar in the model. Avatars are used to infer patient-sharing networks and capacities of doctors. The decisions and behavior of these avatars (that is, how they choose their doctors based on the patient-sharing networks in which they are embedded) have been formulated and calibrated on a large-scale database of observational health care data. The model has 2 relevant free parameters: the shortcut probability, Q , that captures whether patients choose new doctors through the network, or if they base their choice on other factors, and the capacity parameter, C , that quantifies the willingness of doctors to accept additional patients. The introduction of Q avoids dynamical traps in the model and has only a marginal impact on the results. We showed that our main results are robust with respect to the remaining free parameter, the capacity parameter, C . Therefore, it is unlikely that our results are idiosyncratic features of particular choices of the used free parameters, but rather, they reflect genuine structural and dynamical properties of primary care networks.

Regions show considerable differences in their resilience scores. Districts with similar resilience can have very different physician densities (a difference of up to 1 order of magnitude). Physician density assumes that doctors and patients circulate freely in their regions and meet with a probability that is independent of their actual position. This view neglects the actual structure of the underlying patient-sharing networks. Our results show that resilience of the primary care sector is largely a network effect that is determined by the onset of cascading processes of patient displacements triggered by the removal of one or several nodes. After we take the demographic characteristics of the districts into account (e.g., their population, number of doctors, population density, etc.), we find no significant correlations between the regional resilience scores, $f_c(\delta)$, and a number of conventional network measures (connectivity, clustering, centrality, and so on). Our measure of resilience,

therefore, quantifies a genuine network capacity effect of how efficiently the network distributes cascades of patient displacements. Consequently, policy makers should exercise caution when using physician density indicators to estimate the impact of changes in the density of care providers on health service accessibility or coverage, since ignoring the network structure of the health care system might severely under- or overestimate these impacts.

Alongside the systemic risk contributions of individual doctors, we show how to quantify their systemic benefit, B_i , in terms of how many patients they typically absorb in a patient displacement cascade. We find a large number of PCPs who combine relatively high systemic risk with rather low systemic benefits or low risk with high benefit. This is to be expected, since the first is determined by the flow of patients from the PCP, while the second is determined by the flow to the PCP. As the network is not symmetric, these need not be the same. Our results suggest that the health care system could be made more resilient by protecting doctors with high benefit, B_i (or prioritizing their immediate replacement after they leave). Our study also highlights the necessity for an investigation on how specific structural changes (e.g., increased number of multiprofessional primary care centers) impact resilience.

Several limitations originate from the quality of the underlying dataset; a thorough discussion is found in refs. 11 and 12. For resilience assessment, a relevant limitation is that the data allow us to reliably estimate only the quarterly patient numbers of a doctor; we do not observe the maximal number on a given weekday, which would serve as a better proxy for the capacity. Capacity might depend on several nonobservable characteristics in the data, such as working hours, age, sex, number of assistants, infrastructure, or characteristics of the patient set (27). Capacity might also show seasonal variations and be lower during holiday periods. However, while these factors may certainly be helpful in better determining the capacity of a single doctor, they can be expected to have little impact on the overall systemic properties of the networks, such as the existence of a critical fraction of removed doctors on which nonresilient behavior sets in. The factors mentioned could certainly shift the position of $f_c(\delta)$ but would not necessarily impact the comparison of individual districts.

There is a tradeoff between increasing resilience and decreasing overcapacity. One should be careful about interpreting the proposed resilience scores literally as excess capacity. Currently, we overestimate regional resilience by providing upper bounds for how much capacity is needed to avoid disruptions. For instance, if we choose $f_c(\delta)$ such that removals of this size will lead to patients being lost with certainty, there is a chance that patients may already be lost for smaller shocks. We also assumed that patients have near-unlimited patience in searching for new doctors, when in reality, they may just abstain from consulting any further doctors after 1 or 2 unsuccessful attempts. Factors like these may decrease resilience and thereby, overcapacity. A diagnosis of overcapacity would also require an overrepresentation of providers with low financial income.

Our approach can be easily modified to include scenarios other than removals of doctors. There could be surges in patient numbers due to an influx of refugees (28) or an epidemic (29). The method can be transferred to other settings as long as the construction of a patient-sharing network is feasible (i.e., there is a negligible number of isolated nodes or groups of nodes). The structure of patient-sharing networks has already been studied in the United States (15–19), Canada (30), Italy (31), and Australia (32). To transfer our model to other settings, one would, therefore, need to 1) identify suitable data, 2) identify the relevant health sector (e.g., primary care), and 3) confirm that the networks are connected (no substantial isolated components). Most model parameters are estimated from historical data and

therefore, take the heterogeneity of providers and health care delivery models explicitly into account. For instance, in Austria, each federal state has its own social security institution (as do certain occupation groups), which could confound the results. In the regression analysis, we showed that our regional resilience indicators are not driven by such state-level effects, while adjustments might be necessary to compare doctor-level results across different federal states. Finally, it should be noted that the underlying dataset is more than 10 y old and therefore, cannot be expected to adequately represent the current situation in Austria.

Our results clearly show that the resilience of health care systems cannot be described by trivial summary statistics, such as physician density. Resilience must be understood and measured

as the property of how networked systems absorb and restructure themselves in response to shocks (5). We show how resilience can be quantified and used to aid decisions on optimal allocations and how investments for the increase of regional PCP densities would be most beneficial. We can estimate the systemic relevance of individual providers and therefore, identify which providers it would be particularly important to replace immediately on their retirement.

ACKNOWLEDGMENTS. We thank Vito D. P. Servedio for help with the online visualization. We acknowledge support from the European Commission, Horizon 2020 SmartResilience Grant 700621, Wiener Wissenschafts-, Forschungs- und Technologiefonds (WWTF) Project MA16-045, and Österreichische Forschungsförderungsgesellschaft (FFG) Project 857136. We thank an anonymous referee for the idea of studying the multistep model variant.

1. OECD, Health spending (indicator). OECD. <https://doi.org/10.1787/777a9575-en>. Accessed 9 October 2018.
2. F. Pammolli, M. Riccaboni, L. Magazzini, The sustainability of European health care systems: Beyond income and aging. *Eur. J. Health Econ.* **13**, 623–634 (2012).
3. H. P. J. R. Beard, D. E. Bloom, Towards a comprehensive public health response to population ageing. *Lancet* **385**, 658–661 (2015).
4. M. P. Silver, A. Hamilton, A. Biswass, N. I. Warrick, A systematic review of physician retirement planning. *Hum. Resour. Health* **14**, 67 (2016).
5. D. Woods, Four concepts for resilience and the implications for the future of resilience engineering. *Reliab. Eng. Syst. Saf.* **141**, 5–9 (2015).
6. E. G. Martin, N. Helbig, N. R. Shah, Liberating data to transform health care. *J. Am. Med. Assoc.* **311**, 2481 (2014).
7. W. Raghupathi, V. Raghupathi, Big data analytics in healthcare: Promise and potential. *Health Inf. Sci. Syst.* **2**, 3 (2014).
8. World Health Organization, *From Innovation to Implementation, eHealth in the WHO European Region* (WHO Regional Office for Europe, 2016).
9. S. Thurner et al., Quantification of excess-risk for diabetes when born in times of hunger, in an entire population of a nation, across a century. *Proc. Natl. Acad. Sci. U.S.A.* **110**, 4703–4707 (2013).
10. P. Klimek, S. Aichberger, S. Thurner, Disentangling genetic and environmental risk factors for individual diseases from multiplex comorbidity networks. *Sci. Rep.* **6**, 39658 (2016).
11. S. Sauter, L. Neuhofer, G. Endel, P. Klimek, G. Duftschmid, “Analyzing healthcare provider centric networks through secondary use of health claims data” in *IEEE-EMBS International Conference on Biomedical and Health Informatics (BHI)* (IEEE, Valencia, Spain, 2014), pp. 522–525.
12. C. Rinner et al., Improving the informational continuity of care in diabetes mellitus treatment with a nationwide shared EHR system: Estimates from Austrian claims data. *Int. J. Med. Inform.* **92**, 44–53 (2016).
13. A. Geroldinger et al., Mortality and continuity of care—definitions matter! A cohort study in diabetics. *PLoS One* **13**, e0191386 (2018).
14. G. Duftschmid et al., Patient-sharing relations in the treatment of diabetes and their implications for health information exchange: Claims-based analysis. *JMIR Med Inform* **7**, e12172 (2019).
15. H. H. Pham, A. S. O'Malley, P. B. Bach, C. Salontz-Martinez, D. Schrag, Primary care physicians' links to other physicians through medicare patients: The scope of care coordination. *Ann. Intern. Med.* **150**, 236–242 (2009).
16. B. Y. Lee et al., Social network analysis of patient-sharing among hospitals in Orange County, California. *Am. J. Public Health* **101**, 707–713 (2011).
17. B. E. Landon et al., Variation in patient-sharing networks of physicians across the US. *J. Am. Med. Assoc.* **308**, 265–273 (2012).
18. M. L. Barnett et al., Physician patient-sharing networks and the cost and intensity of care in US hospitals. *Med Care* **50**, 152–160 (2012).
19. B. E. Landon et al., Using administrative data to identify naturally occurring networks of physicians. *Med Care* **51**, 715–721 (2013).
20. S. Thurner, R. Hanel, P. Klimek, *Introduction to the Theory of Complex Systems* (Oxford University Press, 2018).
21. M. E. J. Newman, *Networks. An Introduction* (Oxford University Press, 2010).
22. J. Lorenz, S. Battiston, F. Schweitzer, Systemic risk in a unifying framework for cascading processes on networks. *Eur. Phys. J. B* **71**, 441 (2009).
23. S. Battiston, M. Puliga, R. Kaushik, P. Tasca, G. Caldarelli, DebtRank: Too central to fail? Financial networks, the FED and systemic risk. *Sci. Rep.* **2**, 541 (2012).
24. Y. Yang, T. Nishikawa, A. E. Motter, Vulnerability and cosusceptibility determine the size of network cascades. *Phys. Rev. Lett.* **118**, 048301 (2017).
25. S. Buldyrev, R. Parshani, G. Paul, H. E. Stanley, S. Havlin, Catastrophic cascade of failures in interdependent networks. *Nature* **464**, 1025–1028 (2010).
26. World Health Organization, *World Health Statistics 2018: Monitoring Health for the SDGs, Sustainable Development Goals* (World Health Organization, Geneva, Switzerland, 2018).
27. S. M. Petterson et al., Projecting US primary care physician workforce needs: 2010–2025. *Ann. Fam. Med.* **10**, 503–509 (2012).
28. W. Ammar et al., Health system resilience: Lebanon and the Syrian refugee crisis. *J. Glob. Health* **6**, 020704 (2016).
29. R. M. C. Evans, Social networks, migration, and care in Tanzania. *J. Child. Poverty* **11**, 111–129 (2007).
30. T. A. Stukel et al., Multispecialty physician networks in Ontario. *Open Med.* **7**, e40–e55 (2013).
31. D. Mascia, F. Angeli, F. Di Vincenzo, Effect of hospital referral networks on patient readmissions. *Soc. Sci. Med.* **132**, 113–121 (2015).
32. S. Uddin, Exploring the impact of different multi-level measures of physician communities in patient-centric care networks on healthcare outcomes: A multi-level regression approach. *Sci. Rep.* **6**, 20222 (2016).