

## Adoption of a High-Impact Innovation in a Homogeneous Population

Curtis H. Weiss,<sup>1,\*</sup> Julia Poncela-Casasnovas,<sup>2,3</sup> Joshua I. Glaser,<sup>2</sup> Adam R. Pah,<sup>2</sup> Stephen D. Persell,<sup>4</sup> David W. Baker,<sup>4</sup> Richard G. Wunderink,<sup>1</sup> and Luís A. Nunes Amaral<sup>2,3,5,6,†</sup>

<sup>1</sup>*Division of Pulmonary and Critical Care Medicine, Feinberg School of Medicine, Northwestern University, Chicago, Illinois 60611, USA*

<sup>2</sup>*Department of Chemical and Biological Engineering, Northwestern University, 2145 Sheridan Road, Tech E136, Evanston, Illinois 60208, USA*

<sup>3</sup>*Howard Hughes Medical Institute, Northwestern University, Evanston, Illinois 60208, USA*

<sup>4</sup>*Division of General Internal Medicine, Feinberg School of Medicine, Northwestern University, Chicago, Illinois 60611, USA*

<sup>5</sup>*Department of Physics and Astronomy, Northwestern University, Evanston, Illinois 60208, USA*

<sup>6</sup>*Northwestern Institute on Complex Systems, Northwestern University, Evanston, Illinois 60208, USA*

(Received 26 February 2014; revised manuscript received 4 June 2014; published 15 October 2014)

Adoption of innovations, whether new ideas, technologies, or products, is crucially important to knowledge societies. The landmark studies of adoption dealt with innovations having great societal impact (such as antibiotics or hybrid crops) but where determining the utility of the innovation was straightforward (such as fewer side effects or greater yield). Recent large-scale studies of adoption were conducted within heterogeneous populations and focused on products with little societal impact. Here, we focus on a case with great practical significance: adoption by small groups of highly trained individuals of innovations with large societal impact but for which it is impractical to determine the true utility of the innovation. Specifically, we study experimentally the adoption by critical care physicians of a diagnostic assay that complements current protocols for the diagnosis of life-threatening bacterial infections and for which a physician cannot estimate the true accuracy of the assay based on personal experience. We show through computational modeling of the experiment that infection-spreading models—which have been formalized as generalized contagion processes—are not consistent with the experimental data, while a model inspired by opinion models is able to reproduce the empirical data. Our modeling approach enables us to investigate the efficacy of different intervention schemes on the rate and robustness of innovation adoption in the real world. While our study is focused on critical care physicians, our findings have implications for other settings in education, research, and business, where small groups of highly qualified peers make decisions about the adoption of innovations whose utility is difficult if not impossible to gauge.

DOI: [10.1103/PhysRevX.4.041008](https://doi.org/10.1103/PhysRevX.4.041008)

Subject Areas: Complex Systems, Medical Physics

### I. INTRODUCTION

Social scientists have long investigated the “forces” driving the successful adoption of new ideas, technologies, and products [1–4]. The availability of social media has

increased researchers’ ability to study the social aspects of the adoption process on a scale and depth unimaginable even a decade ago [5–7]. However, many studies have restricted their attention to the investigation of the large-scale adoption among heterogeneous populations of products, such as online apps, with low societal impact [5]. In contrast, the landmark studies of adoption investigated the adoption of innovations with high societal value (such as new drugs [2,3] or new hybrid crops [8]) but within populations of individuals with heterogeneous personal, social, and professional characteristics. Indeed, one of the driving forces of these latter studies was the investigation of the impact of the subjects’ differing personal and professional demographics on the process of adoption [2,9–11].

Importantly, numerous situations with great societal impact exist where the adoption process takes place within small groups of highly trained professionals. A school’s teachers determine whether a new teaching method is adopted or not [12]. A hospital’s physicians determine whether a new diagnostic method is adopted or not. In these

\*Corresponding author.

Division of Pulmonary and Critical Care Medicine, Feinberg School of Medicine, Northwestern University, 676 N. Saint Clair, Suite 1400, Chicago, IL 60611, USA.  
curtisweiss@northwestern.edu

†Corresponding author.

Department of Chemical and Biological Engineering, Northwestern University, 2145 Sheridan Road, Tech E136, Evanston, IL 60208, USA.  
amaral@northwestern.edu

*Published by the American Physical Society under the terms of the Creative Commons Attribution 3.0 License. Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI.*

and many other cases, there is clear societal interest in harnessing social networks to accelerate the adoption of the best practices and the obsolescence of poor practices [13].

Here, we investigate the adoption of a diagnostic method within a homogeneous group of critical care physicians who are assigned to teams according to a predetermined schedule and that are responsible for the treatment of specific patients. Because of its nature, our study limits the influence of confounders (such as the role of marketing or traditional forms of adoption promotion) and avoids the long time scales necessary in a larger, multicenter study. In addition to the homogeneity of the study population, our study has two additional features that clearly distinguish it from the landmark studies of adoption. First, knowledge of the innovation is spread almost entirely by peer-to-peer interactions. Second, physicians cannot possibly determine the efficacy of the innovation based on their personal experience and must trust the information provided by peers or by the research literature.

We collect experimental data on adoption over a 244-day period and model the data to identify the mechanisms driving adoption. We then use this understanding to design an intervention scheme that could be used to accelerate the adoption of the best practices.

## II. EXPERIMENTAL PROTOCOL

We examine the adoption of the serum procalcitonin (PCT) assay. PCT is the prohormone of calcitonin and is a marker for the presence and severity of bacterial infection and sepsis in critically ill patients [14]. The PCT assay renders a result in a considerably shorter amount of time (2–3 h) than a microbial culture (approximately 48 h), the current standard of care [15,16]. PCT-based management can safely avoid extraneous antibiotic use in patients with community-acquired pneumonia and guide antibiotic management in patients with sepsis, even with negative microbial cultures [17–19]. Despite the evidence from randomized controlled trials that demonstrate that PCT can be used to avoid extraneous antibiotic utilization without worsening patient outcomes [20], PCT assay use has not been incorporated into clinical practice guidelines as a strong recommendation, suggesting that there remains significant controversy regarding its use [16,21]. In addition, the fact that adopters do not have to abandon the current standard practice makes adoption of the PCT assay a particularly useful real-world experimental condition.

We study adoption among a population of 36 pulmonary and critical care physicians who provide high-intensity critical care coverage in the medical intensive care unit (ICU) at Northwestern Memorial Hospital [Fig. 1(a)]. At any given time, some of these physicians are on duty as part of independent multidisciplinary teams assigned to treat different groups of patients. The department creates the shift schedule prior to the study commencing. Each team conducts daily rounds during which team members

exchange information about each patient and make treatment decisions.

We have arranged with the hospital not to publicize the availability of the PCT assay prior to or during our experiment, thereby ensuring that physicians across the hospital are not aware of test availability, except for two study co-authors (C. H. W. and R. G. W.) who are considered adopters at the start of the experiment. On the first day of the experiment, we expose one additional ICU attending physician and one fellow to information regarding clinical indications for the PCT assay and inform them of its availability. This exchange occurs in one-on-one meetings lasting approximately 10 min.

### A. Practical and ethical constraints on protocol design

The gold standard for biomedical experimentation is a randomized double-blind trial with control groups. Because of the real-world nature of our experimental setting, we are severely limited in our ability to randomize certain features of the experimental setup or to establish traditional control groups. In the following, we detail interventions that could, in principle, be implemented but, in practice, were impractical.

#### 1. Randomization of the innovation

The effect of the degree of efficacy of an innovation on adoptions is very important. While it would be useful to investigate the degree of efficacy of the innovation on physician adoption, for ethical reasons, we could not assign some physicians to use the actual PCT assay and others to use a “placebo assay.” Life-or-death decisions are made daily by physicians based on the results of tests such as the PCT that suggest the presence of a bacterial infection.

#### 2. Randomization of communication patterns

We choose to initiate the experiment on a random date soon after IRB approval is obtained. In this manner, the selection of the participants to be exposed to information about the PCT assay is random. Other aspects of the information propagation process, however, cannot be randomized or perturbed. For example, we could not randomly suppress a communication pathway, as this suppression would affect patient care [22]. Randomization of the shift schedule would also be extremely difficult to implement. Clinical schedules are often the result of the particular clinical needs of an ICU or department as well as physician preferences [23]. Finally, we could randomly invite participants to meetings after the intervention is employed in which we intervene by allowing them to interact outside the work-shift schedule. However, these meetings would unblind the participants to the study and add a degree of external influence that we do not believe is necessary.

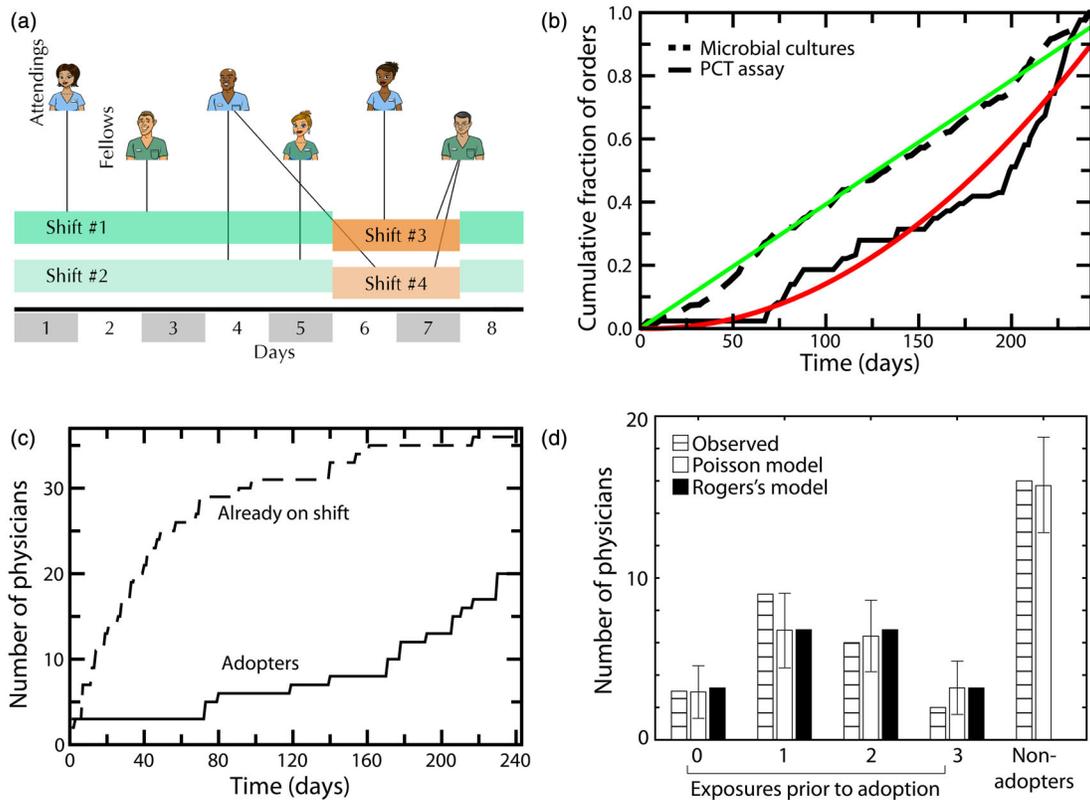


FIG. 1. Experimental conditions of innovation adoption in the medical ICU at Northwestern Memorial Hospital. (a) Schematic of attending and fellow physician schedules. On weekdays (shifts 1 and 2), each of two ICU teams include one attending physician (blue), one fellow (green), and several residents (not pictured). On weekends (shifts 3 and 4), each team has one attending physician, but a single fellow is assigned to both. A third team is established on 1/1/2012 (not pictured). (b) Cumulative fraction of microbial cultures (dashed black line) and PCT assays (solid black line) ordered during the experiment. Cumulative microbial culture orders increase linearly ( $y = 0.004x$ , correlation coefficient = 0.997, solid green line), as expected for the gold standard for diagnosing bacterial infection. A quadratic increase in cumulative PCT orders ( $y = -2 \times 10^{-4} + 1.6 \times 10^{-5}x^2$ , correlation coefficient = 0.979, solid red line) supports the hypothesis that adoption of the PCT assay increases over time. (c) Rate of adoption of the PCT assay among the physicians in the study. Two physicians are adopters prior to the start of the experiment. Adoption of the PCT assay increases over time (solid line) but is slower than the increase in the number of distinct physicians who have already been on shift (dashed line). The data suggest that adoption requires exposure over multiple shifts before adoption occurs. (d) Distribution of adopters according to number of exposures by a unique physician prior to adoption. An exposure is defined as being on shift with a unique adopter. The empirical distribution (striped bars) is compared with the expected distribution if the data were generated by a Poisson process (white bars) and with the expected distribution according to Rogers's categorization of adopters (black bars). In Rogers's categorization, we classify innovators and early adopters, early majority, late majority, and laggards as physicians who adopted after zero, one, two, and three prior exposures, respectively. (Rogers's framework does not consider nonadopters.)

### 3. Control groups

The difficulty of maintaining blinding and of obtaining approval for conducting network experiments that may affect patient care forces us to focus on a single unit at a single hospital. Fortunately, we can consider all the other units in the hospital as control groups, since they are not seeded with knowledge of the innovation. While obtaining results for units at other hospitals might be helpful to demonstrate the generalizability of our results, experimenting at two institutions would be impractical, as it would require the two institutions to be simultaneously ready to adopt the same innovation and yet that both agree not to publicize this innovation.

While our protocol is constrained by a number of practical and ethical issues, it nonetheless enables us to control the experimental conditions to a large degree. Specifically, our study focuses on a homogeneous population of physicians trained in pulmonary and critical care medicine who are blinded to the study (except the two co-authors).

### B. Data collection

We collect data on ICU physician schedules and all PCT and microbial culture orders until the end of the academic year (244 days). For control comparisons, we also collect PCT and microbial culture orders from all other inpatient

units in the hospital. By inference from the empirical data, we define adopters as attending or fellow physicians who are consistently associated with PCT orders during the observation period. (For the results of the analysis using an alternative definition of the adopter, see the Supplemental Material [24]). We assume that physicians who become adopters are unable to revert to nonadopter status. The reason for this decision is that most physicians are not on duty for long enough to enable us to determine with statistical confidence if a physician has reverted.

### III. EXPERIMENTAL RESULTS

First, we compare the time dependence of the number of PCT assay orders during the observation period with the number of microbial culture orders [Fig. 1(b)]. As expected for a widely adopted standard practice, we find an approximately constant rate of orders for microbial cultures. As expected for a test whose adoption is ongoing during the observation period, we find an accelerating rate of orders for PCT assays.

Confirming ongoing adoption, we find that the cumulative number of adopters among the study cohort of 36 physicians displays a concave time dependence [Fig. 1(c)]. This nonlinear pattern is at least partially due to the structure of the network of contacts imposed by the work-shift schedule. Since communication between physicians about the PCT assay is assumed to occur only during a shared work shift, a plateau can occur whenever adopters are absent from the system. Indeed, a plateau is observed in both number of PCT orders and number of adopters during the first several months of the study. This plateau is likely the result adopters only working on a few shifts during this period. At the end of the experiment, 20 physicians have met our definition of being an adopter (56% of the cohort).

Rogers’s seminal work on the diffusion of innovations classifies individuals into five archetypes: innovators, early adopters, early majority, late majority, and laggards [1]. These archetypes can be characterized by the typical number of exposures to the innovation prior to adoption and thus can be understood in the context of a Poisson process occurring in an unstructured population. We calculate the number of exposures (defined as prior shared work shifts with a unique adopter) needed for a physician to become an adopter [Fig. 1(d)]. We find that the distribution of the number of exposures needed by a physician to become an adopter is consistent with Rogers’s typology [1] ( $p = 0.74$ , chi-square test) and with a Poisson process with parameter  $\lambda = 2.5$  ( $p = 0.76$ , chi-square test). These results are consistent with past findings of adoption in unstructured populations. However, adoption in our experiment is occurring due to contacts that are imposed by the shift schedule. This constraint implies that our data hold the potential to enable us to select among mechanistic models responsible for the observed dynamics.

## IV. MODELING RESULTS

To attempt to understand the mechanism underlying adoption in our study, we analyze models that can explain the dynamics of the adoption process. A basic model assumes that the dynamics of adoption [Fig. 1(c)] follow a homogenous Poisson process. However, the time dynamics of adoption are not consistent with a Poisson process ( $p < 0.01$ , Monte Carlo hypothesis testing on coefficients of quadratic fit) (Supplemental Material Fig. 1 [24]).

A reason for the failure of the Poisson model is the fact that diffusion of the innovation in our experiment is occurring due to contacts that are imposed by the shift schedule. While diffusion of innovations within social networks is still mostly modeled as a contagion process [25–27], we hypothesize that opinion dynamics and decision-making mechanisms [28–32] would provide a more plausible explanation of the adoption process in this context—small teams of qualified individuals making important decisions. We therefore investigate both contagion and opinion models and compare their ability to describe and predict the outcomes of our experiment.

Since attending physicians and fellows make the most important decisions regarding medical intensive care unit (MICU) patient care, we assume they would be the most relevant for modeling the adoption process. (Other health professionals, e.g., resident physicians, medical students, and nurses, are part of MICU multidisciplinary teams.) If our assumption is correct, we would expect to observe only minimal leakage of information to other patient-care areas of the hospital—either due to resident physicians rotating in those other areas or due to diffusion of information outside the work-shift network. Indeed, we find that the rate of PCT assay use among attending physicians outside the MICU ( $29/376 = 7.7\%$ ) is an order of magnitude lower than the rate for MICU attending physicians ( $16/22 = 73\%$ ). We believe this result validates our decision to model the propagation of PCT assay adoption using only attending physicians and fellows through the work-shift schedule.

### A. Contagion models

Dodds and Watts introduced a set of models [25,26], denoted as generalized contagion models, that generalize and interpolate between the two standard classes of contagion models: independent interaction [27] and threshold [25]. In contagion models of adoption, adopters are infected and infectious, and nonadopters can be modeled as immune, unexposed, or exposed. If an adopter  $j$  comes into contact with a nonadopter  $i$ , then with probability  $P_{\text{inf}}$ , agent  $i$  receives a positive “dose”  $\delta$ , drawn randomly from some distribution  $f(\delta)$  of dose sizes; otherwise,  $\delta = 0$ . We denote the adoption state of a given individual  $i$  at a given time as  $B_i(t)$ , while the adoption threshold is denoted by  $B_{\text{adopt}}^i$ . A nonadopter starts as an unexposed individual with a value  $B_i(t = 0) = 0$  and maintains a memory of the doses

TABLE I. Summary of performance for all models and for different lengths of the training segment. Recall that the final number of adopters in the experiment is 20. BIC refers to the Bayesian information criterion. Probability  $<0.05$  indicates that the model is not compatible with the experimental data. Lower BIC values indicate that the model provides a better description of the experimental data for the same value of  $T_{\text{training}}$ . The persuasion model performs significantly better than the contagion models in all the metrics, and also for any training periods equal to or longer than 100 days.

Model	$T_{\text{training}}$ (days)	Estimated parameter values	Average (Standard Deviation) final number of adopters	Probability final number of adopters = 20	Probability final number of adopters [17–23]	Probability sum square residuals < 100	BIC score
Persuasion	75	$B_{\text{adopt}} = 0.5$ $R = 0$ $P = 0.1$ $M = 0$	24.9 (1.5)	$<10^{-3}$	0.165	$<10^{-6}$	18.42
	100	$B_{\text{adopt}} = 0.5$ $R = 0$ $P = 0.1$ $M = 0$	21.8 (1.5)	0.255	0.866	$<10^{-6}$	2.73
	125	$B_{\text{adopt}} = 0.5$ $R = 0$ $P = 0.1$ $M = 0$	21.8 (1.5)	0.277	0.998	0.283	2.57
	150	$B_{\text{adopt}} = 0.5$ $R = 0$ $P = 0.1$ $M = 0$	19.0 (1.2)	0.199	1.0	$2.1 \times 10^{-3}$	3.23
	175	$B_{\text{adopt}} = 0.5$ $R = 0$ $P = 0.1$ $M = 0$	19.0 (1.2)	0.222	0.999	0.48	3.01
Contagion multiple dose	75	$P_{\text{inf}} = 0.8$ $P_{\text{imm}} = 0$ dose = 0.2	24.6 (2.4)	0.02	0.366	$5.8 \times 10^{-5}$	7.82
	100	$P_{\text{inf}} = 0.8$ $P_{\text{imm}} = 0$ dose = 0.2	21.8 (1.1)	0.085	0.97	$6.9 \times 10^{-4}$	4.93
	125	$P_{\text{inf}} = 0.7$ $P_{\text{imm}} = 0$ dose = 0.2	17.4 (1.3)	0.043	0.75	0.010	6.29
	150	$P_{\text{inf}} = 0.7$ $P_{\text{imm}} = 0$ dose = 0.2	16.2 (1.8)	0.021	0.484	$<10^{-6}$	7.73
	175	$P_{\text{inf}} = 0.7$ $P_{\text{imm}} = 0$ dose = 0.2	16.6 (1.7)	0.036	0.556	0.081	6.65
Contagion single dose	75	$P_{\text{inf}} = 0.1$ $P_{\text{imm}} = 0$	21.0 (3.3)	0.11	0.68	$2 \times 10^{-4}$	4.41
	100	$P_{\text{inf}} = 0.1$ $P_{\text{imm}} = 0$	18.8 (3.0)	0.124	0.746	$7.3 \times 10^{-4}$	4.17
	125	$P_{\text{inf}} = 0.1$ $P_{\text{imm}} = 0$	15.6 (2.3)	0.028	0.335	$5.1 \times 10^{-3}$	7.15
	150	$P_{\text{inf}} = 0.1$ $P_{\text{imm}} = 0$	14.7 (2.2)	0.016	0.202	$3.6 \times 10^{-4}$	8.27
	175	$P_{\text{inf}} = 0.1$ $P_{\text{imm}} = 0$	16.0 (2.2)	0.033	0.4	$2 \times 10^{-4}$	6.82

(Table continued)

TABLE I. (*Continued*)

Model	$T_{\text{training}}$ (days)	Estimated parameter values	Average (Standard Deviation) final number of adopters	Probability final number of adopters = 20	Probability final number of adopters [17–23]	Probability sum square residuals < 100	BIC score
Contagion multiple dose (threshold from uniform distribution)	75	$P_{\text{inf}} = 1.0$ $P_{\text{imm}} = 0.1$ dose = 0.95	32.3 (2.1)	$< 1 \times 10^{-3}$	0.0	$< 10^{-6}$	18.42
	100	$P_{\text{inf}} = 0.5$ $P_{\text{imm}} = 0.1$ dose = 0.3	23.4 (2.8)	0.058	0.499	$8.4 \times 10^{-5}$	5.69
	125	$P_{\text{inf}} = 0.6$ $P_{\text{imm}} = 0.2$ dose = 0.2	16.9 (2.3)	0.064	0.575	0.015	5.50
	150	$P_{\text{inf}} = 0.7$ $P_{\text{imm}} = 0.6$ dose = 0.7	13.8 (2.28)	0.006	0.123	$7.3 \times 10^{-5}$	10.23
	175	$P_{\text{inf}} = 0.5$ $P_{\text{imm}} = 0$ dose = 0.15	16.0 (2.1)	0.037	0.385	0.05	6.59

received over some time period  $T$ , thus recording her current belief  $B_i$  in the innovation. (Note that we assume here that  $T$  is much longer than the observation period.) Susceptible individuals become adopters if at some time  $B_i > B_{\text{adopt}}^i$ , where the adoption threshold  $B_{\text{adopt}}^i$  of agent  $i$  is drawn randomly from a distribution  $g(B_{\text{adopt}})$ . We present here results for the case where  $B_{\text{adopt}} = 1$  for all individuals. (See Table I for results with  $B_{\text{adopt}}$  drawn from a uniform distribution.) If  $\delta > B_{\text{adopt}}^i$ , then this model reverts to a standard susceptible-infected model [Fig. 2(a)] [25–27].

We select two distinct cases within the generalized contagion model: independent interaction (single dose needed for adoption) and threshold (multiple doses needed for adoption). We generate 1000 independent realizations of each case using the physician interactions imposed by the ICU work schedule, and the empirical adoption and exposure initial conditions—the two initial adopters are set to  $B_i(t=0) = 1$ , whereas  $B_i(t=0) = 0$  for the rest of the population. We then allow the dynamics to take place among the physicians working together, in chronological order and as dictated by the shift schedule.

In the single-dose model, if an infected and a susceptible physician are on duty together, then the susceptible physician will become infected with probability  $P_{\text{inf}}$  or remain susceptible with probability  $1 - P_{\text{inf}}$ . A physician that becomes infected does not return to the noninfected state for the rest of the simulation. A physician has a probability  $P_{\text{imm}}$  of being immune. Immune physicians cannot change state and do not affect the states of other physicians. The single-dose model is characterized by two parameters  $P_{\text{inf}}$  and  $P_{\text{imm}}$ .

In the multiple-dose model, physicians remember exposure doses received over some period of time. A susceptible physician on duty with an infected physician receives an

exposure dose  $\delta$  with probability  $P_{\text{inf}}$ . If the cumulative value of all doses “remembered” by a susceptible physician exceeds the exposure threshold  $B_{\text{adopt}}$ , then the physician becomes infected. The multiple-dose model is thus characterized by two additional parameters  $\delta$  and  $B_{\text{adopt}}$ ; however, we can set  $B_{\text{adopt}} = 1$  without loss of generality.

For both contagion models, we obtain the corresponding simulated time evolution for the number of adopters. We then identify the parameter values that yield the best fit to the experimental data (the set of parameter values that renders the minimum sum of square distances between experimental and simulated curves) and obtain the mean number of adopters as a function of time [Fig. 2(b) and Supplemental Material Figs. 2 and 3].

## B. Persuasion model

Next, we consider a set of models inspired by the opinion models of Deffuant *et al.* and of Hegselmann and Krause [33,34]. In opinion models such as those, belief takes continuous values, and therefore, they can be formalized in a manner similar to the generalized contagion models. Again, we assume that an agent  $i$  has belief  $B_i$ , lying between 0 (no belief in the suitability of the innovation) and 1 (true believer in the suitability of the innovation) [35].

As in contagion models, above a threshold value  $B_{\text{adopt}}^i$ , where  $0 < B_{\text{adopt}}^i < 1$ , the agent believes enough in the suitability of the innovation to start using it, i.e., becomes an adopter [Fig. 2(a)]. Opinion models differ primarily from contagion models in that influence flows bidirectionally between agents and in proportion to the difference in their beliefs. Specifically, interactions between an adopter and a nonadopter result in movement of the beliefs of both individuals toward the initial mean belief of the two agents—the adopter receives a negative dose, whereas the

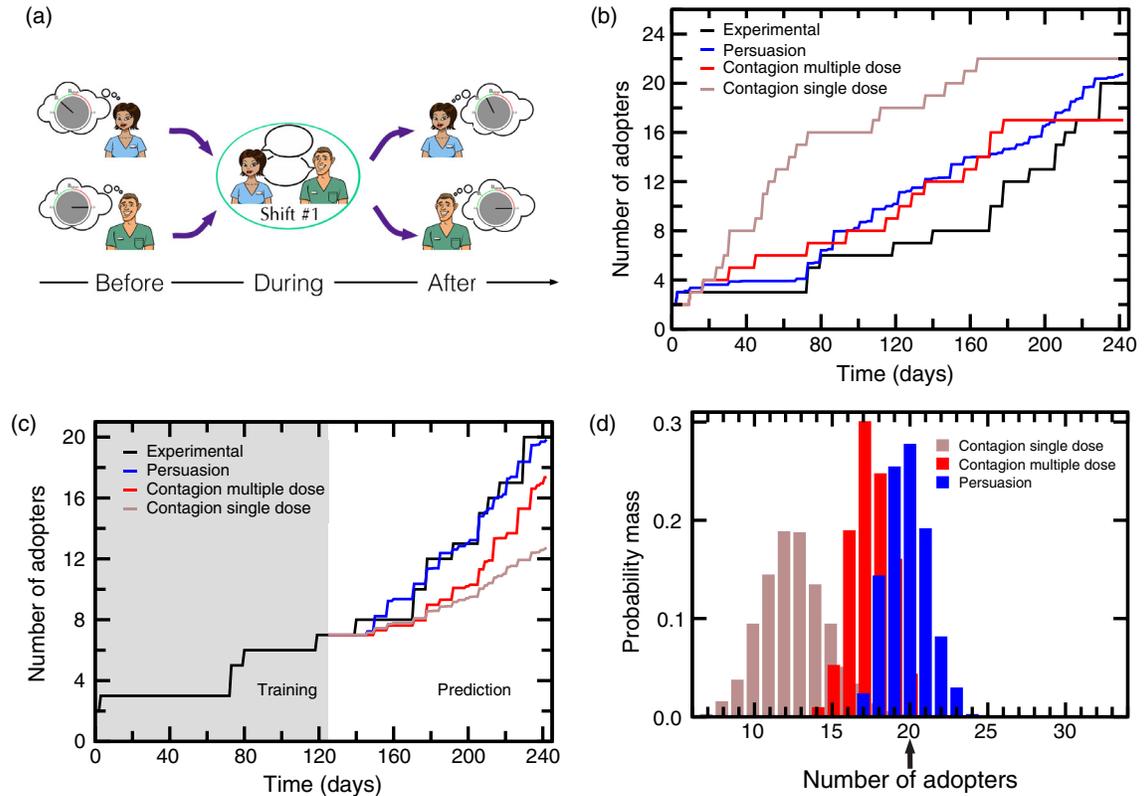


FIG. 2. Computational modeling of the adoption experiment. (a) Schematic of the adoption interaction dynamics between two physicians in the study. Each physician enters each interaction with a belief level  $B_i$  in the innovation with range  $[0, 1]$  represented as a dial. If  $B_i \geq B_{\text{adopt}}$ , the physician is an adopter (red dial). We assume an adopter can expose other physicians to knowledge about the innovation (the PCT assay) only during a shared work shift. As a result of the interaction, physicians may change their opinion. (b) Mean number of adopters over time for the different adoption models and their corresponding best-fit parameters (contagion with multiple doses  $P_{\text{inf}} = 1$ ,  $P_{\text{imm}} = 0$ , and dose = 0.2; contagion with a single dose  $P_{\text{inf}} = 1$  and  $P_{\text{imm}} = 0$ ; and persuasion  $B_{\text{adopt}} = 0.5$ ,  $P = 0.1$ ,  $M = 0$ , and  $R = 0$ ) determined for 1000 simulations compared to experimental data (black line). (c) To test the prediction power of the models, we fit the contagion (with single and multiple doses) and persuasion models to the data for a training period (gray area;  $T_{\text{training}} = 125$  days is illustrated). We determine the best-fit set of parameter values (contagion with a single dose  $P_{\text{inf}} = 0.1$  and  $P_{\text{imm}} = 0$ ; contagion with multiple doses  $P_{\text{inf}} = 0.1$ ,  $P_{\text{imm}} = 0$ , and dose = 0.4; and persuasion  $B_{\text{adopt}} = 0.5$  (fixed),  $R = 0$  (fixed),  $P = 0.1$ , and  $M = 0$ ). We then use these optimal parameter values and the empirical conditions at the end of  $T_{\text{training}}$  as new initial conditions to simulate the models for the remainder of the experiment ( $244 - T_{\text{training}}$ ), using 1000 independent realizations of each model. (d) Distribution of the final number of adopters predicted by the three models considered. We indicate the experimentally observed final number of adopters with the arrow. These data show that the persuasion model is the only one consistent with the experimental data. See the main text and Table I for more details on the model comparisons and performance.

nonadopter receives a positive dose. We further generalize traditional opinion models by assuming that (i) interactions between two adopters lead to both receiving a reinforcing positive dose, (ii) adopters can resist a decrease in their beliefs, (iii) nonadopters have different degrees of initial openness to adopting innovations  $B_i(t=0)$ , and (iv) the difference in beliefs of interacting individuals is not bounded (that is, there is no bounded confidence, as in Ref. [33]). We denote our generalization of opinion models as a “persuasion” model.

The time evolution of the belief of an agent behaving according to the persuasion model is controlled by four parameters, each with a value in the range  $[0, 1]$ . The agent’s persuadability  $P$  quantifies the extent to which the

belief of an agent may change (in any direction) after interaction with an agent in a different state; mutual encouragement  $M$  quantifies the extent to which an adopter increases her belief after interaction with another adopter; and resistance  $R$  quantifies the extent to which an adopter resists a decrease in her belief and the adoption threshold  $B_{\text{adopt}}$ . Additionally, the initial conditions for nonadopters are randomly set to a value  $B_i$  in the range  $[0, B_{\text{adopt}}]$ . We allow the dynamics to take place among the physicians working together, in chronological order and as dictated by the network. When two physicians with beliefs  $B_1$  and  $B_2$  interact, their beliefs are updated. If both physicians are nonadopters, no persuasion occurs, so the beliefs remain unchanged:

$$B_1(t+1) = B_1(t),$$

$$B_2(t+1) = B_2(t).$$

If both physicians are adopters, then their beliefs are increased due to mutual encouragement:

$$B_1(t+1) = B_1(t) + M,$$

$$B_2(t+1) = B_2(t) + M,$$

Finally, if physician 1 is an adopter but physician 2 is a nonadopter, then their beliefs are updated to

$$B_1(t+1) = B_1(t) - RP[B_1(t) - B_2(t)],$$

$$B_2(t+1) = B_2(t) + P[B_1(t) - B_2(t)].$$

For simplicity, we set  $B_{\text{adopt}}$  to the same value for all physicians in the cohort and estimate its value from the data. Even though our model allows for any value of  $R$ , based on the fact that adopters do not revert to being nonadopters in our data,  $R$  is set to 0, leaving three parameters to be fitted. By systematically investigating parameter space, we find that any one of the remaining three free parameters can be fixed without any loss in the model's ability to reproduce the experimental data (Supplemental Material Fig. 4 [24]). This feature of the persuasion model mirrors the presumed characteristics of the physicians in our population: While we choose the same values for the parameters for all physicians, we are modeling a number of different individuals with their own, differing, personal beliefs and "persuadabilities." Had we found that the model only works in a small region of parameter space, we would be highly suspicious, given the fact that physicians are assumed to have differing individual beliefs. We set  $B_{\text{adopt}} = 0.5$  and  $R = 0$  for all the simulations presented here. (See Supplemental Material Figs. 4 and 5 for examples [24] of other best fits and landscapes when fixing different pairs of parameters in the persuasion model.)

For each pair of values of  $P$  and  $M$ , we generate 1000 independent realizations of the persuasion model using the empirical shift network of physicians and the empirical adoption and exposure initial conditions [the two initial adopters are set to  $B_i(t=0) = 1$ , with  $B_i(t=0) = 0$  for the rest of the population] and obtain the corresponding simulated time evolution for the number of adopters. From that parameter exploration, we identify the set of parameter values that yield the best fit to the experimental data and obtain the mean number of adopters as a function of time [Fig. 2(b) and Supplemental Material Fig. 4 [24]]. We observe that the best fit is considerably better for persuasion than for any of the contagion models. These features are at least partially due to the impossibility of the contagion models to replicate the initial plateau, where there are no new adopters.

### C. Model selection

Next, we investigate the ability of the contagion and persuasion models to be predictive of the time evolution of the number of adopters. Specifically, we fit each of these cases to the experimental data for a training period with duration  $T_{\text{training}}$  and then determine to what extent the models (using best-fit parameters from the training period) predict the evolution of adoption for the remaining  $244 - T_{\text{training}}$  days of the experiment. The experimentally observed adoption status for all physicians on day  $T_{\text{training}}$  is used as the initial condition for the prediction simulations [Fig. 2(c); 1000 independent realizations].

Since we cannot *a priori* know what is the optimal duration of the training period, and in order to test the robustness of the models' predictions, we repeat all calculations for five distinct values:  $T_{\text{training}} = 75, 100, 125, 150, \text{ and } 175$  days. (See a complete summary of the results in Table I.) Clearly, the more data available to estimate model parameters, the more accurate the estimates of the parameters should be. Indeed, one finds that the parameter estimates rapidly converge as  $T_{\text{training}}$  increases for all models. Moreover, the longer the duration of the training period  $T_{\text{training}}$  (or, equivalently, the shorter the duration of the prediction period  $244 - T_{\text{training}}$ ), the more accurate one expects the predictions of a model to be. We find this effect to be the case for the persuasion model but not for either of the contagion models, for which performance does not improve with increasing  $T_{\text{training}}$  (Table I).

For concreteness, we focus here on the case  $T_{\text{training}} = 125$  days when discussing model performance [Figs. 2(c) and 2(d)]. The probability that the models will return the experimentally observed final number of adopters is 2.8% for the contagion model with a single dose, 4.3% for the contagion model with multiple doses, and 27.7% for the persuasion model (Table I). The probability that the models will predict a final number of adopters at least as large as that experimentally observed (that is, the  $p$  value for the models) is 0.043 for the contagion models with single and multiple doses and 0.58 for the persuasion model. These  $p$  values show that we can reject the contagion models at the 95% confidence level but not the persuasion model.

## V. MODELING OF INTERVENTIONS

In order to establish the practical impact of our findings, we next explore how the persuasion model can be used to design interventions that lead to faster and more robust adoption. Knowledge translation interventions where all physicians are exposed to an innovation in a general, impersonal manner (e.g., reading a journal or guideline, attending a conference) are known to have low efficacy [36]. However, the resources, cost, and time required to intervene on all physicians individually is impractical. Interventions that intermittently target physicians, such as encouraging or prompting them to change their behavior,

are incrementally effective models of knowledge translation [37–44], but they have not been developed and deployed based on a mechanistic understanding of adoption as a social process.

Our experimental data suggest that the physician contact network imposed by the shift schedule gives rise to long time periods in the experiment during which no adoption can occur because few or no adopters are on duty [Fig. 1(c)]. It thus follows that by designing interventions that target additional physicians whenever it is known that no adopters will be on shift for some period of time, one could potentially speed up the spread of adoption and make the outcome more robust. Specifically, in our proposed intervention strategy, we audit the experimental data to identify all five-day periods during which no adopters are present (“adoption pause”). We then select one team member at random and simulate prompting that physician with information about the PCT assay.

We study the impact of this intervention strategy by performing 1000 simulations of the persuasion model on the experimental shift schedule with the best-fit parameters and a prompting intervention impact  $\Delta_I$ . We observe that for  $\Delta_I/B_{\text{adopt}} = 0.8$ , there is an acceleration in the rate of adoption throughout the simulations, leading to a significantly higher mean number of final adopters compared to the persuasion model without this intervention ( $22.8 \pm 3.0$  vs  $20.8 \pm 2.3$  final adopters,  $p < 0.01$ , student’s  $t$  test) [Figs. 3(a) and 3(b)]. Additionally, we find that as  $\Delta_I$  increases, a concomitant increase in the number of adopters occurs during the experiment, again leading to an increased number of final adopters [Fig. 3(c)]. We determine that the average number of interventions per simulation needed to avoid a five-day adoption pause is 0.5. Regarding the specific moment when the interventions occur, we observe that they tend to occur around day 20 [Fig. 3(d)]. The implication of these results is that a small number of

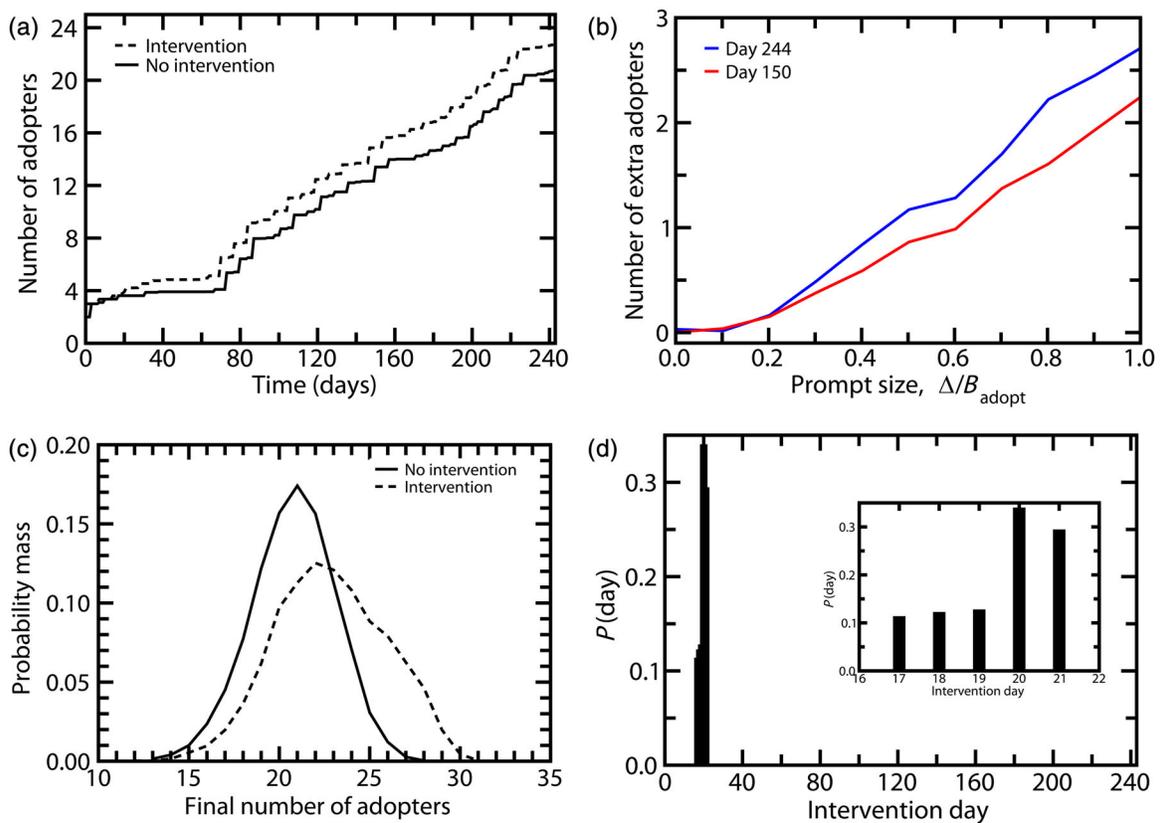


FIG. 3. Using the persuasion model to predict the impact of a prompting intervention (see the main text) on the dynamics of innovation adoption for the persuasion model with best-fit parameters and the experimental work-shift schedule and initial conditions. The intervention (dashed line) accelerates adoption during the simulation, leading to a significantly higher mean final number of adopters compared to the model without intervention (solid line) ( $22.8 [3.0]$  vs  $20.8 [2.3]$  final adopters,  $p < 0.01$ , student’s  $t$  test). (b) Number of additional adopters created by the prompting intervention on the persuasion model as a function of prompt magnitude  $\Delta/B_{\text{adopt}}$  on two different days during the experiment. (c) Distribution of the final number of adopters in the persuasion model with (dashed line) and without (solid line) intervention ( $p < 0.01$ , Kolmogorov-Smirnov test). (d) Distribution of intervention days obtained for the 1000 realizations of the persuasion model with prompting intervention. Inset: Exploded view of the intervention day distribution. For all four panels, the length of the adoption pause is five days. For (a), (b), and (d), the prompting intervention impact is  $\Delta_I = 0.8B_{\text{adopt}}$ . See Supplemental Material Fig. 8 for details [24] of other cases.

interventions at an early point in the process may be all that is necessary in order to increase the adoption of a worthwhile innovation.

We also explore the impact of varying the adoption pause length (e.g., 1, 7, and 10 days). We find that a shorter adoption pause (therefore requiring more interventions) leads to a faster acceleration of the adoption process and more adopters. On the other hand, allowing adoption pauses  $\geq 7$  days does not lead to a significant increase in the number of adopters during or at the end of the process, regardless of the size of the intervention prompt. These results suggest a strategy for determining the optimal magnitude of the prompt or the length of time needed to achieve adoption criticality. (See also Supplemental Material Figs. 6–8 [24].) The intervention we explore here is arguably the simplest imaginable. Other, more sophisticated interventions, such as seeding more physicians, modifying the work schedule to maximize the exposure of nonadopters to adopters, or using knowledge about individual physicians' intrinsic persuadability, would likely lead to a more robust intervention.

## VI. DISCUSSION

Our results are remarkable in that they enable us to select from among a number of mechanistic models for describing the adoption of innovations. Our study clearly rejects the notion that adoption, within a homogenous group of highly trained professionals, propagates as a standard contagion process. Rejection of contagion-type models for adoption of an innovation with a low burden for potential adopters is particularly significant. Contagion models would seem most suited for this simplified situation, where individuals are not asked to abandon their current practice or place a value on the innovation prior to deciding on its adoption. Given the failure of contagion models in our experiment, they are even less likely to explain adoption in more complex contexts, where the risk of adoption is not negligible. In such a complex system, individuals must ponder the innovation more intensely, either because adoption demands the total replacement of standard practice, because the cost associated with adoption is high, or because the agent must personally be able to assess the suitability of the innovation. The persuasion model provides a satisfactory explanation of the adoption of a low-burden and high-impact innovation and is a plausible model for innovations involving greater burdens to adoption.

## ACKNOWLEDGMENTS

We thank Jacob Iasha Sznajder (Northwestern University) for his assistance with study design. We acknowledge support from National Institutes of Health Grants No. K23HL118139 and No. T32HL0761139 (C. H. W.) and Grants No. 5UL1RR025741 and No. 8UL1TR000150 (to Northwestern University Clinical and Translational

Sciences Institute for the Enterprise Data Warehouse), from the Parker B. Francis Fellowship Program (C. H. W.), from the Northwestern Pre-doctoral Biotechnology Training Grant (A. R. P.), and from the Chicago Biomedical Consortium with support from the Searle Funds at the Chicago Community Trust (A. R. P.). The authors declare no conflicts of interest.

C. H. W. and J. P.-C. contributed equally to this work.

- 
- [1] E. M. Rogers, *Diffusion of Innovations*, 5th ed. (Free Press, New York, 2003).
  - [2] J. Coleman, E. Katz, and H. Menzel, *The Diffusion of an Innovation Among Physicians*, *Sociometry* **20**, 253 (1957).
  - [3] F. J. Garjon, A. Azparren, I. Vergara, B. Azaola, and J. R. Loayssa, *Adoption of New Drugs by Physicians: A Survival Analysis*, *BMC Health Serv. Res.* **12**, 56 (2012).
  - [4] L. N. Joppa, G. McNerny, R. Harper, L. Salido, K. Takeda, K. O'Hara, D. Gavaghan, and S. Emmott, *Computational Science: Troubling Trends in Scientific Software Use*, *Science* **340**, 814 (2013).
  - [5] S. Aral and D. Walker, *Identifying Influential and Susceptible Members of Social Networks*, *Science* **337**, 337 (2012).
  - [6] D. Centola, *The Spread of Behavior in an Online Social Network Experiment*, *Science* **329**, 1194 (2010).
  - [7] J. Kleinberg, *The Convergence of Social and Technological Networks*, *Commun. ACM* **51**, 66 (2008).
  - [8] B. Ryan and N. Gross, *The Diffusion of Hybrid Seed Corn in Two Iowa Communities*, *Rural Sociol.* **8**, 15 (1943).
  - [9] E. Katz, *The Two-Step Flow of Communication: An Up-to-Date Report on an Hypothesis*, *Publ. Opin. Q.* **21**, 61 (1957).
  - [10] C. Van den Bulte and G. L. Lilien, *Medical Innovation Revisited: Social Contagion Versus Marketing Effort*, *Am. J. Sociology* **106**, 1409 (2001).
  - [11] R. S. Burt, *Social Contagion and Innovation: Cohesion Versus Structural Equivalence*, *Am. J. Sociology* **92**, 1287 (1987).
  - [12] L. Deslauriers, E. Schelew, and C. Wieman, *Improved Learning in a Large-Enrollment Physics Class*, *Science* **332**, 862 (2011).
  - [13] In medicine, because traditional adoption promotion methods such as lectures, continuing medical education, audit and feedback, and facilitation by opinion leaders are modestly and inconsistently effective at accelerating adoption [36,45–47], health-care innovations take as long as 17 years to be widely adopted in clinical practice [48].
  - [14] K. Reinhart and M. Meisner, *Biomarkers in the Critically Ill Patient: Procalcitonin*, *Critical Care Clinics* **27**, 253 (2011).
  - [15] N. P. O'Grady *et al.*, *Guidelines for Evaluation of New Fever in Critically Ill Adult Patients: 2008 Update from the American College of Critical Care Medicine and the Infectious Diseases Society of America*, *Critical care medicine* **36**, 1330 (2008).
  - [16] R. P. Dellinger *et al.*, *Surviving Sepsis Campaign: International Guidelines for Management of Severe Sepsis and Septic Shock: 2012*, *Critical care medicine* **41**, 580 (2013).

- [17] P. Schuetz *et al.*, *Effect of Procalcitonin-Based Guidelines Vs Standard Guidelines on Antibiotic Use in Lower Respiratory Tract Infections: The ProHOSP Randomized Controlled Trial*, *JAMA, J. Am. Med. Assoc.* **302**, 1059 (2009).
- [18] L. Bouadma *et al.*, *Use of Procalcitonin to Reduce Patients' Exposure to Antibiotics in Intensive Care Units (PRORATA Trial): A Multicentre Randomised Controlled Trial*, *Lancet* **375**, 463 (2010).
- [19] V. Nobre, S. Harbarth, J. D. Graf, P. Rohner, and J. Pugin, *Use of Procalcitonin to Shorten Antibiotic Treatment Duration in Septic Patients: A Randomized Trial*, *Am. J. Respir. Crit. Care Med.* **177**, 498 (2008).
- [20] N. J. Soni, D. J. Samson, J. L. Galaydick, V. Vats, D. L. Pitrak, and N. Aronson, *Procalcitonin-Guided Antibiotic Therapy. Comparative Effectiveness Review No. 78, AHRQ Report No. 12(13)-EHC124-EF, 2012* [[www.effectivehealthcare.ahrq.gov/reports/final.cfm](http://www.effectivehealthcare.ahrq.gov/reports/final.cfm)].
- [21] B. M. Tang, G. D. Eslick, J. C. Craig, and A. S. McLean, *Accuracy of Procalcitonin for Sepsis Diagnosis in Critically Ill Patients: Systematic Review and Meta-analysis*, *Lancet Infect. Dis.* **7**, 210 (2007).
- [22] The communication pathways in our study are the daily rounds conducted by physicians to transmit information and make medical decisions for their critically ill patients. Suppressing or constraining this communication would prevent physicians from providing adequate medical care.
- [23] Indeed, the schedule-creation process is already fairly random, not to mention that it is virtually impossible to prevent random schedule changes among individual physicians, which commonly occur during a year-long clinical schedule.
- [24] See Supplemental at <http://link.aps.org/supplemental/10.1103/PhysRevX.4.041008> for analyses based on an alternative definition of adoption.
- [25] P. S. Dodds and D. J. Watts, *Universal Behavior in a Generalized Model of Contagion*, *Phys. Rev. Lett.* **92**, 218701 (2004).
- [26] P. S. Dodds and D. J. Watts, *A Generalized Model of Social and Biological Contagion*, *J. Theor. Biol.* **232**, 587 (2005).
- [27] R. Pastor-Satorras and A. Vespignani, *Epidemic Spreading in Scale-Free Networks*, *Phys. Rev. Lett.* **86**, 3200 (2001).
- [28] C. Castellano, S. Fortunato, and V. Loreto, *Statistical Physics of Social Dynamics*, *Rev. Mod. Phys.* **81**, 591 (2009).
- [29] M. Stone, *The Opinion Pool*, *Ann. Math. Stat.* **32**, 1339 (1961).
- [30] S. Chatterjee and E. Seneta, *Towards Consensus: Some Convergence Theorems on Repeated Averaging*, *J. Appl. Probab.* **14**, 89 (1977).
- [31] J. E. Cohen, J. Hajnal, and C. M. Newman, *Approaching Consensus Can Be Delicate When Positions Harden*, *Stoch. Proc. Appl.* **22**, 315 (1986).
- [32] C. Castellano, M. Marsili, and A. Vespignani, *Nonequilibrium Phase Transition in a Model for Social Influence*, *Phys. Rev. Lett.* **85**, 3536 (2000).
- [33] G. Deffuant, F. Amblard, and G. Weisbuch, *Mixing Beliefs Among Interacting Agents*, *Adv. Compl. Syst.* **03**, 87 (2000).
- [34] R. Hegselmann and U. Krause, *Opinion Dynamics and Bounded Confidence Models, Analysis, and Simulation*, *J. Art. Soc. Social Sim.* **5**, 3 (2002).
- [35] The belief axis would need to have a more complex arrangement if some agents would be actively against the proposed innovation.
- [36] C. D. Naylor, *Better Care and Better Outcomes: The Continuing Challenge*, *JAMA, J. Am. Med. Assoc.* **279**, 1392 (1998).
- [37] D. R. Gifford, R. G. Holloway, M. R. Frankel, C. L. Albright, R. Meyerson, R. C. Griggs, and B. G. Vickrey, *Improving Adherence to Dementia Guidelines through Education and Opinion Leaders: A Randomized, Controlled Trial*, *Ann. Intern. Med.* **131**, 237 (1999).
- [38] S. B. Soumerai *et al.*, *Effect of Local Medical Opinion Leaders on Quality of Care for Acute Myocardial Infarction: A Randomized Controlled Trial*, *JAMA, J. Am. Med. Assoc.* **279**, 1358 (1998).
- [39] J. A. Kelly, J. S. St Lawrence, Y. E. Diaz, L. Y. Stevenson, A. C. Hauth, T. L. Brasfield, S. C. Kalichman, J. E. Smith, and M. E. Andrew, *HIV Risk Behavior Reduction Following Intervention with Key Opinion Leaders of Population: An Experimental Analysis*, *American Journal of public health : JPH / American Public Health Association* **81**, 168 (1991).
- [40] J. R. Stanford, L. Swaney-Berghoff, and K. Recht, *Cardiac Surgical Outcomes Improvement Led by a Physician Champion Working with a Nurse Clinical Coordinator*, *Am. J. Med. Qual.* **27**, 5 (2011).
- [41] D. C. Scales, K. Dainty, B. Hales, R. Pinto, R. A. Fowler, N. K. Adhikari, and M. Zwarenstein, *An Innovative Telemedicine Knowledge Translation Program to Improve Quality of Care in Intensive Care Units: Protocol for a Cluster Randomized Pragmatic Trial*, *Implement. Sci.* **4**, 5 (2009).
- [42] D. C. Scales, K. Dainty, B. Hales, R. Pinto, R. A. Fowler, N. K. Adhikari, and M. Zwarenstein, *A Multifaceted Intervention for Quality Improvement in a Network of Intensive Care Units: A Cluster Randomized Trial*, *JAMA, J. Am. Med. Assoc.* **305**, 363 (2011).
- [43] O. Berwanger *et al.*, *Effect of a Multifaceted Intervention on Use of Evidence-Based Therapies in Patients with Acute Coronary Syndromes in Brazil: The BRIDGE-ACS Randomized Trial*, *JAMA, J. Am. Med. Assoc.* **307**, 2041 (2012).
- [44] C. H. Weiss *et al.*, *Prompting Physicians to Address a Daily Checklist and Process of Care and Clinical Outcomes: A Single-Site Study*, *Am. J. Respir. Crit. Care Med.* **184**, 680 (2011).
- [45] R. Grol and J. Grimshaw, *From Best Evidence to Best Practice: Effective Implementation of Change in Patients' Care*, *Lancet* **362**, 1225 (2003).
- [46] N. Ivers, G. Jamtvedt, S. Flottorp, J. M. Young, J. Odgaard-Jensen, S. D. French, M. A. O'Brien, M. Johansen, J. Grimshaw, and A. D. Oxman, *Audit and Feedback: Effects on Professional Practice and Healthcare Outcomes*, *Cochrane Database Syst. Rev.* **6**, CD000259 (2012).
- [47] G. Flodgren, E. Parmelli, G. Doumit, M. Gattellari, M. A. O'Brien, J. Grimshaw, and M. P. Eccles, *Local Opinion Leaders: Effects on Professional Practice and Health Care Outcomes*, *Cochrane Database Syst. Rev.* **8**, CD000125 (2011).
- [48] E. A. Balas and S. A. Boren, *Yearbook of Medical Informatics*, edited by J. Bemmell and A. T. McCray (Schattauer, Stuttgart, Germany, 2000), p. 65–70.