

The power of a good idea: Quantitative modeling of the spread of ideas from epidemiological models

Luis M.A. Bettencourt^a, Ariel Cintrón-Arias^{b,c,*}, David I. Kaiser^d,
Carlos Castillo-Chávez^{b,c}

^aCCS-3, Computer and Computational Sciences, Los Alamos National Laboratory, Los Alamos, NM 87545, USA

^bCenter for Applied Mathematics, Cornell University, 657 Rhodes Hall, Ithaca, NY 14853, USA

^cDepartment of Mathematics and Statistics, Arizona State University, P.O. Box 871804, Tempe, AZ-85287-1804, USA

^dCenter for Theoretical Physics and Department of Physics, Massachusetts Institute of Technology, 7
7 Massachusetts Ave., Cambridge, MA 02139, USA

Received 5 May 2005; received in revised form 18 August 2005

Available online 18 October 2005

Abstract

The population dynamics underlying the diffusion of ideas hold many qualitative similarities to those involved in the spread of infections. In spite of much suggestive evidence this analogy is hardly ever quantified in useful ways. The standard benefit of modeling epidemics is the ability to estimate quantitatively population average parameters, such as interpersonal contact rates, incubation times, duration of infectious periods, etc. In most cases such quantities generalize naturally to the spread of ideas and provide a simple means of quantifying sociological and behavioral patterns. Here we apply several paradigmatic models of epidemics to empirical data on the advent and spread of Feynman diagrams through the theoretical physics communities of the USA, Japan, and the USSR in the period immediately after World War II. This test case has the advantage of having been studied historically in great detail, which allows validation of our results. We estimate the effectiveness of adoption of the idea in the three communities and find values for parameters reflecting both intentional social organization and long lifetimes for the idea. These features are probably general characteristics of the spread of ideas, but not of common epidemics.

© 2005 Elsevier B.V. All rights reserved.

Keywords: Scientific idea-diffusion; Epidemiological models; Transition parameter estimation; Rumor models

1. Introduction

Dynamical population models are used to predict average behavior, generate hypotheses or explore mechanisms across many fields of science including ecology [1–3], epidemiology [4–7] and immunology [8], to name but a few. Traditionally, epidemiological models focus on the dynamics of “traits” transmitted between individuals, communities, or regions (within specific temporal or spatial scales). Traits may include (i) a communicable disease such as measles [4] or HIV [9]; (ii) a cultural characteristic such as a religious belief, a

*Corresponding author.

E-mail address: ariel@cam.cornell.edu (A. Cintrón-Arias).

fad [10–13], an innovation [14], or fanatic behavior [15]; (iii) an addiction such as drug use [16] or a disorder [17]; or (iv) information spread through, e.g., rumors [18,19], e-mail messages [20], weblogs [21], or peer-to-peer computer networks [22].

The earliest and by now the most thoroughly studied population models are those used to map disease progression through a human population [23–25]. These models typically divide a population into classes that reflect the epidemiological status of individuals (e.g. susceptible, exposed, infected, etc.), who in turn transit between classes via mutual contact at given average rates. In this way the models can capture average disease progression by tracking the mean number of people who are infected, who are prone to catch the disease, and who have recovered over time. In addition, these models can be used to identify the role of specific population characteristics such as age, variable infectivity, and variable infectious periods [24]. The division of epidemiological classes according to such characteristics gives rise to more complex models with so-called heterogeneous mixing.

In this paper we apply models similar to those used in epidemiology to the spread of ideas. By the term “idea” we refer generally to any concept that can be transmitted from person to person [26–29]. It may refer to a technology, which may require effort and apprenticeship to be learned, but it may also be a more fickle piece of information such as a colloquialism or a piece of news. What is important is that it is possible to tell if someone has adopted the idea, understands and remembers it, and is capable of and/or active in spreading it to others.

Pioneering contributions to the modeling of social contagion processes, based on epidemiological models, date back to as early as 1953 [10,18,30–35]. Nearly a decade later, population models were applied to the spread of scientific ideas [36,37]. Around the same time, a stochastic model for the spread of rumors was proposed and analyzed [19]. In this model, a closed population is divided into three “social” states: ignorant, spreaders, and stiflers. Transitions from the ignorant state to spreaders may result from contacts between the two classes, whereas encounters between individuals who already know the rumor may lead to its cessation. Various recent extensions of this model include a general class of Markov processes for generating time-dependent evolution [38], and studies of the effects of social landscapes on the spread, either through Monte Carlo simulations over small-world [39] and scale-free [40] networks, or by derivation of mean-field equations for a population with heterogeneous ignorant and spreader classes [41]. Other interesting mathematical models, that attempt to capture the capacity for a population or idea to persuade others, or imitate, have also been developed to generate predictions regarding product adoption [42] or public opinion trends [43,44]. Despite this revival in the modeling of information spread, few of these models have been directly validated by empirical data [36,37,42,44].

Beyond obvious qualitative parallels there are also important differences between the spread of ideas and diseases. The spread of an idea, unlike a disease, is usually an intentional act on the part of the transmitter and/or the adopter. Some ideas that take time to mature, such as those involving apprenticeship or study, require active effort to acquire. There is also no simple automatic mechanism—such as an immune system—by means of which an idea may be cleared from an infected individual. Most importantly, it is usually advantageous to acquire new ideas, whereas this is manifestly not so for diseases. This leads people to adopt different, often opposite, behaviors when interested in learning an idea compared to what they may do during an epidemic outbreak. Thus we should expect important qualitative and quantitative differences between ideas and diseases when using epidemiological models in a sociological context. We explore some of these points below in greater detail, in the context of specific models and data.

In spite of these differences, quantifying how ideas spread is very desirable as a means of testing sociological hypotheses. For example, we can apply dynamical population models to the spread of an idea to validate statements about how effectively it is transmitted, the size of the susceptible population, the speed of its spread, as well as its persistence. Estimating the population numbers and rates is useful in constraining explanatory frameworks. It is also useful for studying how cultural environments may affect adoption, as happens when the same idea is presented to communities in different nations, or conversely when different ideas are presented to the same community.

In this paper, we apply several general models, inspired by epidemiology and informed by our knowledge of the sociology of the spread dynamics, to the diffusion of a specific scientific idea in three different communities. Our test case is the spread of Feynman diagrams, since the late 1940s the principal

computational tool of theoretical high-energy physics, and later also used extensively in other areas of many-body theory such as atomic physics and condensed-matter theory. The primary reason to choose this example is that we have detailed historical information about the network of contacts, person by person, by means of which the diagrams spread during the first 6 years after their introduction [45–47].

This example of the spread of an idea may not transcend automatically to other cases of idea diffusion. Feynman diagrams are primarily a tool for complex calculation. As such their study and assimilation require a period of apprenticeship and familiarization. Transmission of the technique almost invariably proceeded, in the early years, through personal contact, from informal teacher to student and among peer groups of users. In later years the idea became familiar and available in accessible forms so that (in principle) it could more easily have been learned from books and lecture notes. Thus, although our example will clearly not cover every class of ideas it will point, we believe, to features of epidemic models that apply to idea diffusion. It will also reveal features of these models that require modification, thereby producing more realistic candidate models that we expect will prove useful beyond our present analysis.

In Section 2 we give some historical background on the spread of Feynman diagrams in the United States, Japan, and the Soviet Union. We discuss our data sources and the organization of the datasets. Section 3 presents several classes of models of epidemiology (or directly inspired by them), some of their mathematical properties, and the circumstances under which we expect them to apply to the spread of ideas. We apply each model to the historical data in Section 4, and discuss the estimated values for the model parameters in the light of our independent knowledge of how the diagrams spread. Finally, in Section 5 we present our conclusions and give some outlook on the general population modeling of the spread of ideas. Appendix A contains details about our parameter estimation procedure.

2. Data sources, time series reconstruction, and state determination

Feynman diagrams occupy a central role in modern theoretical physics. Realistic models of high-energy physics, as well as in condensed-matter, atomic, and nuclear physics cannot be solved exactly to generate predictions that can be confronted with experiments. In special circumstances, however, such as when interactions are weak, series expansions in a small parameter permit very good systematic approximations.

In models of particle physics, such as the relativistic quantum theory of electromagnetism—quantum electrodynamics—most terms of this series beyond zeroth order (tree level) are formally infinite. The procedure of removing unphysical infinities to generate predictions is called “renormalization.” It is vital for renormalization to work that commensurate terms be grouped together. This is a relatively simple procedure for the lowest orders in the expansion series but becomes absolutely confounding at higher orders, in which many terms contribute and infinities must cancel precisely between them. For example, in quantum electrodynamics, second-order calculations (involving the first non-trivial corrections within the perturbative expansion) typically involve 10 or so distinct terms to be delimited, calculated, and added together, while eighth-order calculations involve nearly 1000 such terms. Both challenges to making calculations in quantum electrodynamics—the presence of infinities and the accounting difficulties of perturbative calculations—were well known to physicists during the 1930s, and the problems remained unsolved after World War II. Throughout 1947 and 1948 several approaches to rendering quantum electrodynamics well-defined were being attempted in the USA and Japan, but it remained unclear if any renormalization program could succeed systematically [48].

It was then that Freeman Dyson, following up on an idea by Richard Feynman, was able to show how a diagrammatic representation of particle interactions could be used to organize the series expansion. Using the diagrams, Dyson further demonstrated that the infinities could be systematically identified and cancelled to any perturbative order. This conceptual breakthrough unified Feynman’s approach (then at Cornell University) with that of Julian Schwinger (at Harvard University) and Sin-Itiro Tomonaga (at Tokyo Education University). For their contributions Feynman, Schwinger, and Tomonaga were awarded the Nobel Prize in 1965 [48]. Feynman diagrams opened the floodgates for computation (and prediction) in quantum electrodynamics and beyond, creating enormous research opportunities for a new generation of theoretical physicists. Tests of quantum electrodynamics and later quantum field theories of the weak and strong nuclear interactions continue today in multibillion-dollar particle accelerators at CERN and Fermilab, as well as at

smaller installations. These quantum field theories taken together constitute the “standard model of particle physics,” which summarizes our most fundamental (and most exact) understanding of matter and radiation to date. Almost all quantitative predictions of the standard model, on which modern particle physics and cosmology are based, are computed using series of Feynman diagrams.

Because of their extraordinary importance in enabling a good part of modern theoretical physics, the advent of quantum electrodynamics and of Feynman diagrams in particular has been very well documented. Our data was collected in large part for a new book by one of the authors [45]. For the United States and Britain one of us (Kaiser) reconstructed the network of contacts—author by author—for the spread of the diagrams during the first 6 years after their introduction, between 1949 and 1954. For this he relied upon unpublished correspondence, preprints, lecture notes, and publications from the period, along with more recent interviews and published recollections. With the aid of two colleagues, he used similar materials to study how the diagrams spread to young physicists in Japan and the Soviet Union. Although less information is readily available about these communities of physicists, a reasonably complete picture of contacts and spread can also be inferred [45,46].

Data for the number of authors adopting Feynman diagrams were collected for the first 6 years in the USA and Japan, from the beginning of 1949 to the end of 1954. For the Soviet Union, where the diagrams were introduced later and where the spread was initially slower, we assembled data for the first 8 years, from the beginning of 1952 to the end of 1959. We identified adopters of the idea (or members of the “infected” class) based on published uses of (or discussion of) Feynman diagrams in the main physics research journals of each country: *Physical Review* in the USA, *Progress in Theoretical Physics* in Japan, and *Zhurnal Eksperimental'noi i Teoreticheskoi Fiziki* (*Journal of Experimental and Theoretical Physics*) in the USSR. The data were identified by manual page-by-page counts. We found this to be necessary because no citation search or even keyword search would suffice. Often in the early years authors would cite the Feynman and/or Dyson papers without making any use of the actual diagrammatic techniques, and, conversely, by the early 1950s many would use the diagrams without necessarily citing the Feynman or Dyson papers. Given the quasi-exponential nature of the adoption process these identification methods become impractical for longer times. This is the principal reason why we have not extended the study to later years. Additionally, we have detailed historic accounts of the spreading process spanning these initial periods only [45,46]. Such knowledge will allow us to build models below that reflect fundamental social dynamics, different qualitatively from those underlying standard models of epidemics.

The identification of adopters with published authors can clearly lead to underestimation. Nevertheless the procedure is dictated by practical considerations and has become the standard. For example Tabah, reviewing an extensive literature on historical scientific trends grappling with this issue, suggests that the most practical measuring unit of a (scientific) idea is the published article [37]. And Goffman, using mean-field epidemiological models [36,49–51], chose the publication as the transmission “unit”, when he validated models for the diffusion of scientific ideas with aggregate longitudinal data about mast cell research [36], and the growth of symbolic logic [51].

Similarly the identification of national communities with specific journal publication is imperfect, although we find almost no cross-national publications, apart from a few British authors who were in active contact with developments in the USA and published in the *Physical Review*. As such they are counted as part of the diagram-using community in the USA. With these choices the evolution of cumulative numbers of Feynman-diagram authors is shown in Fig. 1.

We see that none of the data sets shows saturation in the growth of the adoption of Feynman diagrams. There are good reasons for this, spanning the initial period covered by the data, shown in Fig. 1, and beyond. The physics graduate student enrollment grew rapidly in the US after World War II, faster than any other field, and was growing especially quickly during the late 1940s when Feynman diagrams were introduced [52]. This growth persisted until the late 1960s, with an average doubling time of 6.24 years. Among all subfields of physics, particle and nuclear physics, where the diagrams first spread, grew the fastest. The numbers of new physics graduate students in the USSR also increased exponentially, at a rate comparable to that in the US during the postwar period, but quantitative estimates are more uncertain. In Japan, we know that membership in the new Elementary Particle Theory Group (which consisted largely of interested grad students and postdocs) grew rapidly during this period [45]. Moreover during the mid and late 1950s, the range of




Fig. 1. The time evolution of the cumulative number of authors using Feynman diagrams in the USA, Japan, and the USSR. The method was first discovered in the USA and quickly spread both there and in Japan. Adoption was particularly fast in Japan where researchers had already developed similar methods. At the same time, new institutions were developed throughout Japan after World War II that helped the nation's physicists share information from the international scientific community that might otherwise have been difficult to access. Adoption in the USSR occurred later because of scientific isolation from physicists in the West with the onset of the Cold War, and proceeded more slowly because of institutional resistance. For details of these institutional and pedagogical factors, see Refs. [45,46].

applications to which physicists applied Feynman diagrams widened considerably. No longer the province for nuclear and high-energy physicists alone, many people began to apply the diagrams for problems in solid-state physics and beyond. This led to another surge in growth of diagram adopters, as new cohorts encountered the diagrams across a growing number of subfields of physics. Compounding this growth, a new generation of textbooks appeared that featured the diagrammatic techniques prominently, ensuring even wider adoption within graduate students' curricula.

Analogies to other population states commonly used in epidemiological models are natural but must be properly qualified. The identification of susceptibles is usually problematic both for diseases and ideas. For simplicity, one may consider the entire population that is not infected (or recovered), but if the spreading process requires such features as direct contact with those already infected this may turn out to be a gross overestimate. With the benefit of hindsight we can see what fraction of the population actually became infected, but such estimates can clearly underestimate the class of susceptibles.

Finally, it is interesting to discuss the recovered state. For some communicable diseases such a state does not exist; as it happens in HIV and tuberculosis, for which infected individuals remain latent for extensively long periods. On the other hand, there are infectious diseases in which an individual acquires immunity right after recovery and will not get re-infected. This is not true with ideas, a case in which culture is manifestly different from biology. An idea can recur again and again, whenever it becomes useful, once it becomes part of an individual's repertoire. In many cases (and this is clear in our data for several authors), an individual might publish in areas where Feynman diagrams are used, only to later leave the area for good or to return to it later. For very prolific authors, publication in several areas simultaneously occurs frequently.

With these caveats in mind we proceed in the next sections to apply epidemic models to our data. Model parameters will be estimated on the basis of how well they fit the evolution of adopters. Furthermore, the results of these estimates will be subjected to broad bounds imposed by the solutions' plausibility, given our knowledge of the historical facts.

3. Population models: drawing parallels between epidemics and idea diffusion

Below we shall concentrate on the classical, simplest epidemiological models, based on “homogeneous mixing” in which state variables are only functions of time. In a review of epidemiological models, Hethcote [24] introduced their compartmental characterization (e.g. SIR, SIS, SEIR, etc.) within a global analysis of the field. Such survey also discusses how more complex models can be used to assess the impact of population structure (age, risk, gender, etc.), epidemiological variability (age of infection, variable infectivity, distributed incubation periods, etc.), and scale (spatial, temporal, etc.) on disease dynamics and control. Although we have knowledge of some population characteristics (e.g. academic level, institutional location) in our data set we feel it may not be large enough to make such distinctions in a way that will lead to useful quantitative discrimination.

As such we explore below a large class of “mean-field” models, illustrated by Fig. 2. At the onset of the spread of the idea most of the population will be in the susceptible class (S), with a few individuals in the incubator class (E)—having been in contact with the idea—and a small number of adopters (I) manifesting it. These are the principal classes in the models below. In addition, inspired by the approaches of Daley and Kendall [19], we also explore models in which there may be competing and mutually exclusive ideas (e.g. where susceptibles are turned off from the idea and become skeptics or idea stiflers, represented by the class Z). Furthermore, individuals may recover or become immune (R), and not manifest the idea again. Different models combine subsets of these states and admit different couplings between them.

The total population is denoted by $N(t)$, where $N = S + E + I + Z + R$ (see Table 1). In the epidemic models used in this study, the demographic dynamics are modeled by $dN/dt = \mathcal{B}(N) - \mu N$, where $\mathcal{B}(N)$ is referred to as the recruitment function. In our case, this denotes the arrival rate of new individuals susceptible to the idea, such as new graduate students starting in the field as well as other scientists who find the idea relevant for their research. The parameter $\mu > 0$ denotes the rate at which physicists stop using Feynman diagrams (the exit terms in Fig. 2). Thus, the maximum value that $1/\mu$ can take is the average lifespan of the idea within a generation of researchers in the relevant community.

Whenever $\mathcal{B}(N) > 0$ and $\mu > 0$, then the system in Fig. 2 is said to have *vital dynamics*. If $\mathcal{B}(N) \equiv \Lambda > 0$, then $N(t)$ varies over time and approaches a stable fixed point, Λ/μ , as $t \rightarrow \infty$, in other words, the community




Fig. 2. The basic scheme of population dynamics models for the spread of ideas, inspired by similar models in epidemics. An individual can be recruited into the susceptible (S) class, then be exposed (E) to the idea, incubate it, and eventually manifest it, becoming a member of the adopter or infected class (I). An individual might instead move into a competing infective class (e.g., skeptics, Z). It is possible that part of the population may eventually recover (R), meaning that it will not manifest the idea again. Individuals can also exit any class, thus reducing the total population.

Table 1
Nomenclature for the state variables of the several population models used to describe the spread of ideas

Variable	Definition
<i>S</i>	Susceptible
<i>E</i>	Idea incubators
<i>I</i>	Idea adopters
<i>Z</i>	Skeptics
<i>R</i>	Recovered
<i>N</i>	Total population: $N = S + E + I + Z + R$

approaches its “carrying” capacity. In order to illustrate generic model features we discuss below a few partial implementations of this general scheme, including explicit parameterizations. At the end of this section we emphasize the role of the basic reproductive number, R_0 , as a measure of effectiveness of adoption.

3.1. Models without incubation: SIR model

The classical epidemic model consists of three states: susceptibles (*S*), adopters (or infected, *I*), and recovered (*R*). In this SIR model, susceptible individuals transit directly to the adopter class through contact with other adopters, without any delay period or incubation. The recovered state consists of those individuals who no longer manifest the idea. This state allows for the decay of adopters by recovery and thus leads to a regulation of the idea spread. The model is defined by the following system of ordinary differential equations (where overdots denote derivatives with respect to time):

$$\begin{cases} \dot{S} = \Lambda - \beta S \frac{I}{N} - \mu S, \\ \dot{I} = \beta S \frac{I}{N} - (\gamma + \mu)I, \\ \dot{R} = \gamma I - \mu R, \\ \dot{N} = \Lambda - \mu N, \end{cases} \tag{1}$$

where $1/(\gamma + \mu)$ is the average time spent manifesting the idea as an adopter (γ denotes the recovery rate from infection). The term $\beta SI/N$ is usually referred to as the standard incidence. The parameter β is the per capita idea adoption rate. It can in turn be thought of as the product between the mean contact rate per capita and the probability of adoption per contact.

As noted above, although recovery is a natural concept in epidemiology (since organisms naturally may become immune after exposure and/or infection), there is no strict parallel when discussing ideas. Loose analogies are possible, e.g. once one loses interest in an idea it is usually harder to have an individual express it, whereas novelty may make it more attractive. Nevertheless there is no systematic cognitive process, analogous to the immune system, that actively clears out ideas. As such many ideas are remembered for life.

Many ideas may be short-lived, say from years to days, compared to the lifetime of the individual. In this case, we may consider a single outbreak by setting $\Lambda = \mu = 0$. The sign of the right hand side of the second equation in system (1) then determines the spread of the idea and depends on the initial fraction of susceptibles, $S(t_0)/N$. If the initial state of the population can be such that $S(t_0)/N < \gamma/\beta$, then the number of infectives can only decrease. This is the basis of immunization campaigns, whereby members of the susceptible class are turned into members of the immune class, and hence become part of $R(t_0)$. Thus knowledge of the infection rate, β , and of the lifetime of the infection, $1/\gamma$, results in the recommendation for the fraction of immune (recovered) necessary for an epidemic not to develop, namely $R(t_0)/N > 1 - (\gamma/\beta)$. For a very infectious disease or idea (large β) or one with a slow recovery rate (small γ) almost all of the population must be immune in order to halt the spread.

Due to the less clear definition of immunity to an idea, the concept of what may constitute immunization is also ill-defined. Clearly, the novelty of an idea and a perception of its potential are often its most attractive features. Changing this perception through education (e.g., about the consequences of a certain behavior, ideology, or practice) may lead to an increase of skepticism and consequently greater “immunity” upon exposure. Moreover, we should keep in mind that this concept of immunization, just as in standard epidemics but for different reasons, is usually only valid for the lifetime of an individual. Although some biological immunity can be passed e.g. from mother to infant, it is usually the case that young individuals are more susceptible to new diseases and ideas alike. In the Feynman diagram case this is borne out historically: over 80% of the early adopters of the diagrams in each country were either graduate students or postdocs when they first began using the diagrams; older physicists simply did not re-tool [45].

The asymptotic late-time dynamics of model (1) are well known, and will form the basis for the analyses of more complex models discussed below. Suppose that $\Lambda > 0$ and $\mu > 0$. For long times, and regardless of the distribution of infectives and susceptibles, recruitment and exits will balance each other so that $\lim_{t \rightarrow \infty} N(t) = N^* = \Lambda/\mu$. There are up to two different non-negative steady states (fixed points), known in epidemiology as the *disease-free equilibrium* with $S^* = N^* = \Lambda/\mu$, $I^* = R^* = 0$, and the *endemic state* (whenever $\beta/(\gamma + \mu) > 1$) with

$$S^* = \frac{\gamma + \mu}{\beta} N^*, \quad I^* = \left[\frac{\mu}{\gamma + \mu} - \frac{\mu}{\beta} \right] N^*, \quad R^* = \frac{\gamma}{\mu} I^* = \left[\frac{\gamma}{\gamma + \mu} - \frac{\gamma}{\beta} \right] N^*. \quad (2)$$

The eigenvalues around the disease-free state equilibrium are $(-\mu, -\mu, \beta - (\gamma + \mu))$. Thus it is stable provided that $\beta < \gamma + \mu$, i.e., if the decay rate (due to exit and recovery) is larger than the idea adoption rate. The instability of the disease-free state corresponds to stability of the endemic state. The eigenvalues of the linearized system around the endemic equilibrium are

$$-\mu, \quad -\frac{\beta\mu \pm A}{2(\gamma + \mu)}, \quad (3)$$

where $A = \sqrt{\mu(\beta^2\mu - 4\beta(\gamma + \mu)^2 + 4(\gamma + \mu)^3)}$. All eigenvalues are negative provided that $\beta > \gamma + \mu$, guaranteeing the local stability of the endemic state.

As a result a transcritical bifurcation (where the two equilibria exchange stability) takes place at $R_0 \equiv \beta/(\gamma + \mu) = 1$. In the mathematical epidemiology literature the dimensionless quantity R_0 is known as the basic reproductive number. R_0 has an intuitive and useful interpretation as the average number of secondary cases produced by a “typical” infected individual during his/her entire life as infectious, when introduced in a population of susceptibles (assumed to be at a demographic steady state). We will discuss the role of R_0 further in Section 3.3.

3.2. Competition and incubation: SIZ and SEIZ models

In the spread of ideas, but almost never in standard epidemics, the exposure of individuals to an idea almost invariably leads to both enthusiasts and skeptics. In the case of Feynman diagrams, skeptics did indeed emerge. Julian Schwinger, for example, who developed a non-diagrammatic method of renormalization, quipped years later that Feynman diagrams had “brought computation to the masses”—hardly a good thing, as far as Schwinger was concerned. Although his graduate students at Harvard did learn something about the diagrams, they made little use of them in their dissertations and early articles. Oppenheimer, too, was initially skeptical, and effectively blocked Dyson’s recruitment efforts at the Institute for Advanced Study in Princeton for several weeks, before Hans Bethe interceded directly on Dyson’s behalf. In Moscow, meanwhile, the influential Lev Landau made his distaste for Feynman diagrams clear during the early 1950s, blocking any discussion of them in his famous seminar (even chastising one young graduate student who had expressed interest in the diagrams that it would be “immoral” to chase such “fashions” as Feynman diagrams!) [45]. Thus inclusion of skeptics alongside enthusiasts is quite important. This can be modeled by considering two competing and mutually exclusive infected states, say I and Z . The simplest such model (SIZ)

is given by

$$\begin{cases} \dot{S} = \Lambda - \beta S \frac{I}{N} - bS \frac{Z}{N} - \mu S, \\ \dot{I} = \beta S \frac{I}{N} - \mu I, \\ \dot{Z} = bS \frac{Z}{N} - \mu Z, \end{cases} \tag{4}$$

where b and β denote the per capita rates of idea rejection and adoption by susceptibles, respectively.

The interesting new feature about this type of model is that it can support up to three fixed points (see Table 2). The first is the usual disease-free state $S = N^* = \Lambda/\mu$, $I = 0$, $Z = 0$ (extinction of both adopters and skeptics), and two endemic states, one for each strand I , Z :

$$S = \frac{\mu}{\beta} N^*, \quad I = \left(1 - \frac{\mu}{\beta}\right) N^*, \quad Z = 0 \text{ (extinction of skeptics),} \tag{5}$$

or

$$S = \frac{\mu}{b} N^*, \quad Z = \left(1 - \frac{\mu}{b}\right) N^*, \quad I = 0 \text{ (extinction of adopters).} \tag{6}$$

Observe that model (4) does not support the steady state co-existence of adopters and skeptics. For the disease-free state the eigenvalues are $(b - \mu, \beta - \mu, -\mu)$. Thus for stability one needs both $b < \mu$ and $\beta < \mu$. This means that there are two R_0 's, $R_0^I = \beta/\mu$, and $R_0^Z = b/\mu$.

Under these circumstances, which of the endemic states becomes stable? To investigate this question we inspect the eigenvalues around the I endemic state. This gives

$$-\mu, \quad \left(\frac{b}{\beta} - 1\right)\mu, \quad -\beta + \mu. \tag{7}$$

Similarly, we obtain the eigenvalues for the endemic Z state by replacing b with β and vice versa in (7). This result implies that only one of the two endemic states can be stable, depending on the relative magnitude of the contact rates b and β . We note, however, that because there is no contact term between the I and Z , the way one class ends up dominating relies on long-time changes in the population through cycles of recruitment and exit. This time scale can be very long, diverging in the limit where $b \rightarrow \beta$. For $\beta > b$ it will take on average $\beta/(\beta - b)$ generations until the disappearance of skeptics.

The model generalizes immediately to an arbitrary number, n_Z , of alternative endemic states, Z_i (in which we include the usual I), with associated contact rates b_i . There will then be $n_Z + 1$ fixed points, one disease-free and n_Z endemic corresponding to each strand. As in the SIZ model above only the state with the largest b_i will be locally stable. The stability of the fixed point associated with Z_i for decay in favor of an alternative state Z_j

Table 2
Parameter definitions used in the several population models of Section 3

Parameter	Definition
Λ	Recruitment rate
$1/\mu$	Average lifetime of the idea
$1/\varepsilon$	Average idea incubation time
$1/\gamma$	Average recovery time
β	Per-capita S - I contact rate
ρ	Per-capita E - I contact rate
b	Per-capita S - Z contact rate
l	$S \rightarrow Z$ transition probability given contact with skeptics
$1 - l$	$S \rightarrow E$ transition probability given contact with skeptics
p	$S \rightarrow I$ transition probability given contact with adopters
$1 - p$	$S \rightarrow E$ transition probability given contact with adopters

is characterized by an eigenvalue $[(b_i/b_j) - 1]\mu$. The disease-free equilibrium will be locally stable if and only if all $R_0^i = b_i/b_j < 1$, $\forall_{i=1}^m$.

As above, consider the case in which recovery can take place in the SIZ model, and proceeds with rate γ_I from the I class, and with rate γ_Z from the Z class. $R_0^{I,Z}$ change by the simple modification $\mu \rightarrow \mu + \gamma_{I,Z}$. In the absence of vital dynamics, it then becomes a necessary and sufficient condition for the growth of the strand I , Z that $S(t_0)/N > \gamma_{I,Z}/\beta$, respectively. What is interesting now is that the reduction of the susceptibles can be achieved by having a suitably large fraction of the population in the complementary infective strand(s). For example, I will not grow if $[Z(t_0) + R(t_0)]/N > 1 - (\gamma_I/\beta)$. This observation quantifies the fact that in a population with a large fraction of skeptics an idea will not take hold. In this sense, complementary strands effectively act like recovery states. This may be the most natural explanation for why old ideas seldom re-surface, in spite of being preserved for very long times in the population and various archives.

One important drawback of SIR and SIZ models is that once exposed to an infected person, a susceptible individual transits immediately to the infected class. This feature is often unrealistic, especially for ideas that require long periods of apprenticeship, which is common in scientific research and is a significant feature of the Feynman-diagram user data discussed below. The simplest way of incorporating some delay in an SIZ model is to introduce a new class of incubators (or exposed), denoted by E , between the susceptible and adopter states. Upon contact with an adopter, a susceptible individual transits with a given probability to the E class. This class has a given characteristic lifetime, $1/\varepsilon$, before the individual manifests the idea and transits to the I class. That is, $1/\varepsilon$ is the average incubation (or maturation) time of the idea [53]. It is expected to be a function of personal effort on the part of the adopter as well as environment (adverse or supportive). There may also be population losses due to vital dynamics, which we will continue to assume occur on a timescale $1/\mu$. In this sense not all of the exposed population will become infected.

This extension leads to an SEIZ model. In addition, this model can be enriched with extra processes to generate a better description of the data. Below we present a version of the SEIZ model in which skeptics recruit from the susceptible pool with rate b , but their action may result either in turning the individual into another skeptic (with probability l), or it may have the unintended effect of sending that person into the incubator class (with probability $1 - l$). We also introduce a probability, p , that a susceptible individual will become immediately infected with the idea upon contact. Conversely, with probability $1 - p$ that person will transit to the incubator class instead, from which the individual may later become an adopter. Furthermore, the transition of individuals from the incubator class to the adopter class can be promoted by contact, with rate ρ . With these choices the model is given by

$$\begin{cases} \dot{S} = \Lambda - \beta S \frac{I}{N} - b S \frac{Z}{N} - \mu S, \\ \dot{E} = (1 - p)\beta S \frac{I}{N} + (1 - l)b S \frac{Z}{N} - \rho E \frac{I}{N} - \varepsilon E - \mu E, \\ \dot{I} = p\beta S \frac{I}{N} + \rho E \frac{I}{N} + \varepsilon E - \mu I, \\ \dot{Z} = lb S \frac{Z}{N} - \mu Z. \end{cases} \quad (8)$$

As expected, the system has a disease-free state with $S^* = N$, $E^* = I^* = Z^* = 0$. Analysis of the local stability of this fixed point (utilizing next generation operator [54,55]) reveals that the basic reproductive numbers are given by

$$R_0^{I,Z} = \left(\frac{\beta(\varepsilon + p\mu)}{\mu(\varepsilon + \mu)}, \frac{lb}{\mu} \right). \quad (9)$$

As in the SIZ model the first number, R_0^I , is the one of interest, as it corresponds to an eigenvector with a component of adopters. The second value, R_0^Z , corresponds to the exclusive growth of a population of skeptics, without acceptance of the idea.

Table 3
Basic reproductive number R_0^I for the SIR, SEI, and SEIZ models discussed in Section 3

Model	SIR	SEI	SEIZ
R_0^I	$\frac{\beta}{\gamma + \mu}$	$\frac{\beta\varepsilon}{\mu(\mu + \varepsilon)}$	$\frac{\beta(\varepsilon + p\mu)}{\mu(\varepsilon + \mu)}$

3.3. Speed of idea propagation and effectiveness of adoption

From the discussion of the models above we can define several important intuitive quantities that characterize the spread of ideas. For example, a simple measure of the speed of propagation of the idea is the number of new adopters per unit time. This is simply given by \dot{I} .

For simple models, such as the ones discussed above, in which there is only one growing eigenvalue λ^+ for each infective strand, the initial velocity of the spread is simply

$$v_{\text{ini}} \equiv \dot{I}(t_0) \simeq \lambda^+ I(t_0). \tag{10}$$

The quantity v_{ini} gives a measure of how fast the idea will initially spread but not of its overall adoption effectiveness. In order to determine the latter we must consider the number of new adoptions that a spreader of the idea can lead to during his/her lifetime. Since there is no a priori good reason to suspect that ideas are short-lived, the effectiveness of an idea may result from slow spread over long periods of time and thus may not be well characterized by v_{ini} .

The number of secondary adoptions induced by a typical idea spreader in a population of susceptibles over that person’s lifetime as an adopter, t_{idea} , is called the *basic reproductive number*, R_0 , in ecology and epidemiology (see Refs. [3,5,24]). As such R_0 is the invasion criterion for adopters in a population of susceptibles, or analogously the average branching ratio (the number of offspring) of the typical adopter over his/her lifetime in this state. If $R_0 > 1$ then the idea will spread. The greater R_0 , the more effective the idea adoption will be.

In practice R_0 can be computed in simple models through the linearization of $\dot{I}(t)$ around the disease-free equilibrium. These expressions are summarized in Table 3. For the computation of R_0 in models with heterogeneous populations other methods are necessary [5,54,55]. In Section 4 we will estimate the statistical distributions for R_0 (see Fig. 4) subject to fitting the data for the spread of Feynman diagrams in the USA, Japan, and the USSR. The mean of this distribution provides a measure of the effectiveness of the adoption of Feynman diagrams in the three countries.

4. Results and discussion

We now analyze the results of estimating parameters by matching the data on the spread of Feynman diagrams for three distinct countries to several population models discussed above. These models allow us to discuss the effects of the recovered class, of latency, and of competitive idea strands. They also explore several classes of transition mechanisms, both by progression and by contact between population classes.

Table 4 summarizes the results. To gauge the applicability of each model to each data set we used the simplest measure of goodness of fit, by computing the absolute value of the deviation between model prediction and data. Average deviations per data point are shown in Table 4. Details of our ensemble estimation procedure are given in Appendix A.

Here we note simply that parameter estimation must, by practical necessity, be confined to given numerical ranges, with upper and lower bounds dictated by general empirical considerations. Our choices of estimation intervals are shown in Table 5. This procedure is familiar from epidemiology, where knowledge about such quantities as the length of incubation and infectious periods is often used to restrict various model parameters to plausible values (Refs. [8,56] also employ assumptions of this nature in immunology).

Table 4

The smallest (absolute value) average deviation per data point between the best fit parameters of each model and data on the number of Feynman diagram adopters for the USA, Japan, and the USSR

Model	USA	Japan	USSR
SIR	2.816	1.788	1.487
SEI	1.963	1.638	1.437
SEIZ	1.467	1.568	1.437

Table 5

Parameters used in the SIR, SEI and SEIZ population models, their allowed ranges in our estimation procedure and units

Parameter	Baseline range	Unit
$S(t_0)$	[0, 500]	people
$E(t_0)$	[0, 100]	people
$I(t_0)$	[0, 20]	people
$R(t_0)$	[0, 10]	people
$Z(t_0)$	[0, 100]	people
ε	[0.2, 6]	1/year
β	[0, 12]	1/year
b	[0, 12]	1/year
l	[0, 1]	1
γ	[0, 12]	1/year
λ	[0, 50]	people/year
μ	[0.025, 12]	1/year
p	[0, 1]	1
ρ	[0, 12]	1/year

4.1. Results for models without incubation: SIR

We start by presenting our results for simple models without incubation. Parameter estimates are given in Table 6 for the USA, Japan, and the USSR, while the model solutions are compared to the data in Fig. 3.

The estimates for the initial population paint a picture of a considerably larger scientific community susceptible to learn Feynman diagrams in the USA than in the other two countries. In Japan, $S(t_0)$ appears more than three times smaller than in the USA, while in the Soviet Union our estimates indicate a very small number of susceptibles around 1952. Nevertheless both the USA and the USSR show strong levels of recruitment (slightly higher λ in the USA), as compared to Japan.

This makes sense given each community's rates of growth during this time period. In the postwar United States, the rate at which new Ph.D.s in physics were granted grew by nearly twice that of any other field between 1945 and 1951, quickly exceeding (by a factor of three) the prewar rate at which new physicists had been trained. Meanwhile, building on the wartime Manhattan Project pattern, the federal government pumped money into physics at more than ten times the prewar levels. Most singled out for support during the early postwar period was high-energy physics, precisely that branch of the discipline in which Feynman diagrams were first developed and from which the earliest adopters came [52]. These factors led to a large population of susceptibles when Feynman and Dyson first introduced Feynman diagrams.

Japan, on the other hand, had a strong tradition of high-energy physics before the war, but massive shortages of funding and basic supplies during the early postwar years hampered the growth of that country's physics community (lower λ). Although absolute numbers of new physicists in Japan did not grow at anything like the pace in the United States after World War II, several institutional changes were introduced in Japan right around the time that Feynman diagrams were invented, greatly facilitating the diagrams' spread

Table 6

Parameter estimation (SIR model) for the spread of Feynman diagrams in the USA, Japan, and the USSR. The three columns show our best-fit estimate, the mean computed over an ensemble of parameter set realizations, and corresponding standard deviation (Std)

Parameter	Best-fit	Mean	Std
USA			
$S(t_0)$	114.092	96.463	76.726
$I(t_0)$	11.948	10.982	0.542
$R(t_0)$	0.830	0.550	0.432
β	0.534	0.663	0.052
γ	8.542×10^{-3}	0.049	0.034
A	40.417	42.864	5.130
μ	0.036	0.058	0.023
R_0	12.029	6.752	2.008
Japan			
$S(t_0)$	33.901	24.534	3.537
$I(t_0)$	4.018	3.799	0.348
$R(t_0)$	1.925	0.864	0.714
β	1.990	2.255	0.131
γ	8.668×10^{-3}	0.054	0.034
A	12.466	20.759	2.646
μ	0.031	0.087	0.037
R_0	49.582	16.922	4.308
USSR			
$S(t_0)$	1.347	1.156	1.088
$I(t_0)$	1.935	1.583	0.218
$R(t_0)$	9.742	4.928	2.415
β	1.251	1.258	0.045
γ	0.030	0.092	0.062
A	32.822	32.031	6.894
μ	0.188	0.134	0.063
R_0	5.739	6.053	1.963

throughout Japan. This fact is reflected in the highest adoption rate, β , for Japan, compared to the other two countries. This in turn leads both to the fastest speed of adoption and the highest value of R_0 .

Contacts between Japanese and American physicists began again in 1948 (while Japan was still under US occupation), including visits by several Japanese theoretical physicists to the Institute for Advanced Study in Princeton, New Jersey, where Freeman Dyson was honing the new diagrammatic techniques. A new organization in Japan, known as the Elementary Particle Theory Group, was also founded in 1948, and began to publish its own informal newsletter and preprint organ, *Soryushi-ron Kenkyu*, which helped to spread news of the new diagrammatic techniques. And finally, the Japanese university system quickly expanded tenfold, beginning in 1949, allowing young physicists to establish new groups and visit new institutions throughout the country, putting the new techniques into rapid circulation [45,46].

The Soviet Union was the only country in the world after World War II in which the growth in the numbers of new physicists and in government spending on physics was comparable with the United States. This may explain why our estimates of the recruitment rates A are so high and commensurate for the two nations. But the onset of the Cold War in the late 1940s effectively ended all informal communication between physicists in the USA and USSR just months before Feynman and Dyson introduced Feynman diagrams.

These geopolitical constraints severely limited the exchange of information for several years and explain why Feynman diagrams took hold in the Soviet Union only later and at a slower initial pace (smallest v_{ini}). Only with the “Atoms for Peace” initiatives, starting in 1955, did physicists from both countries begin to meet informally for extended visits. And only after these lengthy face-to-face “exposures” did Soviet physicists begin to adopt Feynman diagrams at a comparable rate to those in the USA and Japan [45,46]. Over time the effectiveness of adoption, R_0 , was nevertheless comparable between the USSR and the USA.




Fig. 3. The best fit trajectory (see Table 6) for the rise of Feynman diagram adoption obtained for the SIR model vs. the data for the USA, Japan, and the USSR.




Fig. 4. The probability distribution of the basic reproductive number R_0 estimated from the SIR model for the USA data on the spread of Feynman diagrams. R_0 measures the effectiveness of the idea adoption. Its value estimates the average number of adopters induced by a typical spreader in a population of susceptibles.

Finally we notice that both the exit and recovery rates, μ and γ , are small in every case, their sum being comparable to a career lifetime (5–25 years). The fact that γ is estimated to be smaller than μ is a consequence of our imposed lower bound on the exit rate and the fact that the data only constrains their sum. Although this estimate cannot be made with good confidence for data which only covers the first 6 years, it is an indication that ideas are not naturally forgotten.

Incidentally, we do know in some cases that the time to change research subject was much shorter for a few prominent authors. Richard Feynman was working almost exclusively on his theory of superfluidity by 1953 (although some of his students continued to use the diagrams under his supervision), while Freeman Dyson was persuaded to change research direction, to condensed-matter theory, at a meeting with Enrico Fermi also in 1953. (See Dyson’s testimony in Ref. [57]; see also, Ref. [45].)

The long exit and recovery times, combined with finite, plausible values of the contact rate β , lead in turn to large values of R_0 . The fact that an infected individual, when introduced in a population of susceptibles, can lead to many adopters (here 6–50) is associated *not* with high adoption rate for the idea, β , but rather with a long time (many years) over which the idea can be transmitted, $1/(\gamma + \mu)$. This is a feature that we will see repeated in more complex models and that is manifestly different between biological infection and the spread of ideas.

4.2. Results for models with incubation: SEI

We now analyze the effects of including latency in the models. In the simplest SEI model, susceptibles transit to an intermediate class of incubators (E) upon contact with adopters, in which they remain for a characteristic “incubation” time $1/\varepsilon$, after which they manifest the idea. Note that due to exit processes the average time spent in the incubator class is actually $1/(\varepsilon + \mu)$, and that some individuals exit the population and never manifest the idea. In practice μ will be estimated to be small and the time spent in the incubators class is indeed essentially the incubation time. The simplest SEI model is a subset of the SEIZ model of Eq. (8), and is given by

$$\begin{cases} \dot{S} = A - \beta S \frac{I}{N} - \mu S, \\ \dot{E} = \beta S \frac{I}{N} - \varepsilon E - \mu E, \\ \dot{I} = \varepsilon E - \mu I. \end{cases} \tag{11}$$

Results of the parameter estimates are presented in Table 7.

The most important qualitative difference, relative to models without latency, is that the model can now better fit data at early times for the USA and Japan (see Fig. 5). This accounts for the bulk of the improvements in Table 4. In both cases this is made possible in the SEI model because starting with a number of individuals in the incubator class allows a two-stage growth process for the adopters. Initially, the incubators are depleted, allowing for a growth of adopters with a negative second derivative. This is the main feature of SEI solutions, accounting for their better fit of the data relative to the SIR model. The two-stage process is a property of the growth curve for adopters in the USA from the initial time until the early 1950s, and to a lesser extent for Japan over the same period, after a slightly later start. The characteristic time at which the curve changes concavity can be computed from the initial growth as

$$t^* \simeq \frac{1}{\varepsilon} \ln \left[\frac{E(t_0)}{E(t_0) + I(t_0)} \left(1 + \frac{\varepsilon + \mu N(t_0)}{\beta S(t_0)} \right) \right]. \tag{12}$$

This time is longest for the USA, on the order of 10 months, shortest for Japan at 2.3 months, and about 5 months for the USSR, reflecting the relative values of the parameters ε and β estimated for the three countries. Beyond this point in time the effect of the incubator class is relatively negligible. For Japan and the USSR, where ε is largest, the class becomes essentially non-dynamical beyond the initial transient, with $E \rightarrow \beta SI/(N\varepsilon)$, and as a consequence the solutions look much like those of the SIR model.

Table 7

Parameter estimations for the SEI model for the adoption of Feynman diagrams in the USA, Japan, and the USSR

Parameter	Best-fit	Mean	Std
USA			
$S(t_0)$	478.515	398.691	61.990
$E(t_0)$	60.989	44.686	4.728
$I(t_0)$	0.020	0.160	0.135
ε	0.257	0.391	0.055
β	1.041	0.951	0.086
μ	0.025	0.040	0.012
λ	45.385	40.052	6.467
R_0	37.711	23.172	5.798
Japan			
$S(t_0)$	30.248	31.037	2.190
$E(t_0)$	11.569	12.022	1.400
$I(t_0)$	0.153	0.165	0.129
ε	2.361	2.009	0.279
β	5.956	4.417	0.787
μ	0.039	0.044	0.013
λ	12.067	12.578	1.093
R_0	150.136	105.372	35.223
USSR			
$S(t_0)$	3.074	0.810	0.722
$E(t_0)$	3.344	3.462	0.647
$I(t_0)$	0.682	0.738	0.266
ε	1.713	1.613	0.476
β	3.715	3.589	0.753
μ	0.067	0.075	0.035
λ	17.819	19.372	3.668
R_0	53.257	55.892	28.788

In practice, the incubation periods estimated for the three countries are quite different. For the USA (see Fig. 6), the best fit solutions prefer to start in 1949 with a relatively large number of incubators and an incubation time of order 3–4 years. In both Japan and the Soviet Union the initial population included fewer incubators but had a considerably shorter incubation time, of the order of 5–6 months in Japan and 7–8 months in the USSR. These incubation period estimates for Japan and the Soviet Union are unexpectedly short, since most of the papers were authored by graduate students who took on average a few years of training (“incubation”) before publishing. The small values for ε thus reveal some limits of the simple SEI model: in particular, simple progression to adoption (parameterized by ε) does not capture the dynamics adequately, since (as we know historically) the role of multiple contacts was important. We return to this issue below.

Beyond the role played by the incubator class, we observed the same relative hierarchy of several important quantities among the different national communities. Japan had the largest effectiveness of adoption, R_0 , whereas both the USA and the USSR displayed smaller and statistically commensurate values for R_0 . (Data were only collected for the USA and Japan for the period 1949–1954, because the steep rate of growth made longer collection times infeasible. The slow rise of diagram adoption in the USSR, on the other hand, encouraged us to collect data for a longer period, 1949–1959, making direct comparisons between late-time behavior in the USA and the USSR difficult.) In every case the large values of R_0 are essentially due to a long lifetime of the idea, $1/\mu$, of 13–40 years. The recruitment rates, λ , similarly to the SIR estimates, are highest for the USA, followed by the USSR, reflecting these national efforts to increase the numbers of new physicists.

In spite of all these qualitative similarities one should also keep in mind that the numerical values for each of these parameters are generally different between the SIR and SEI models, and not always statistically




Fig. 5. Best fit trajectories corresponding to the parameter estimations for the SEI model, for the USA, Japan, and the USSR. The SEI model fits the data better at early times, especially for the USA, when compared to the SIR model.

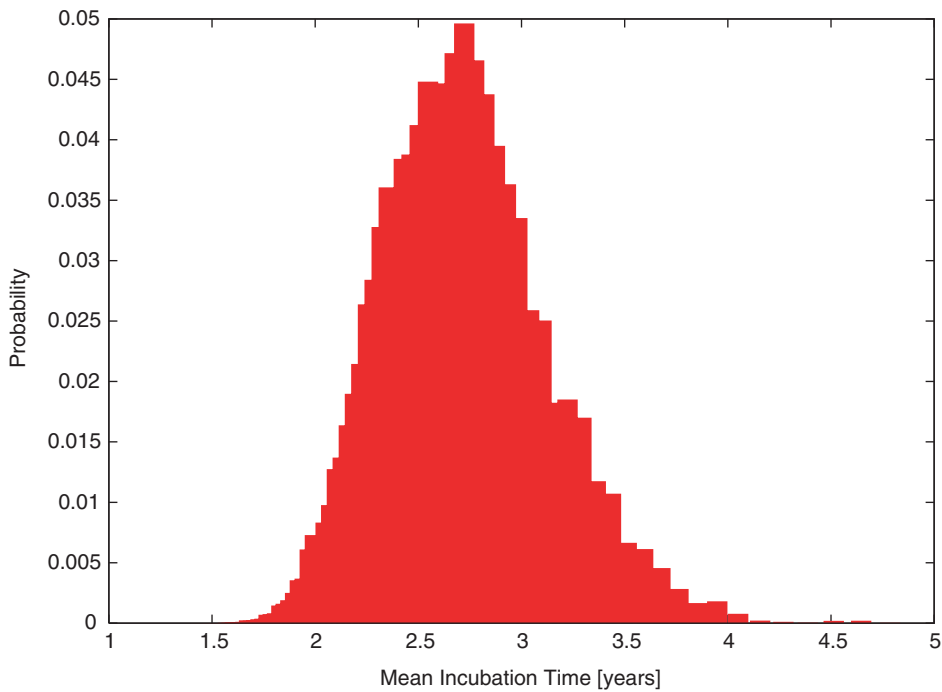


Fig. 6. Histogram of the mean incubation time $1/\varepsilon$ for Feynman diagrams in the USA, estimated by fitting the SEI model to data, see Table 7.

compatible. Thus preference of one model over another can be determined via consideration of the goodness of fit (Table 4), but should also take into consideration qualitative knowledge of the processes at play.

4.3. Results for models with incubation and competition: SEIZ

Finally, we consider the most complex model of our set, which includes an additional class Z much like that of adopters, but which competes with I for susceptibles. Results of the parameter estimation procedure are given in Table 8 and in Fig. 7.

Table 8

Parameter estimations for the SEIZ model and data for the spread of Feynman diagrams for the USA, Japan, and the USSR. We restricted the estimation procedure to the regime where $R_0^I > R_0^Z$, see Eq. (9)

Parameter	Best-fit	Mean	Std
USA			
$S(t_0)$	98.973	108.662	5.852
$E(t_0)$	24.515	24.984	0.447
$I(t_0)$	5.916×10^{-5}	0.031	0.027
$Z(t_0)$	0.114	0.160	0.119
ε	0.202	0.210	0.009
β	0.488	0.496	0.012
b	0.164	0.156	0.117
l	0.311	0.252	0.171
μ	0.025	0.032	0.006
p	0.570	0.566	0.052
ρ	11.893	11.549	0.330
A	49.527	47.860	1.555
R_0^I	18.412	14.975	2.227
Japan			
$S(t_0)$	24.806	24.798	1.356
$E(t_0)$	16.123	15.292	0.781
$I(t_0)$	1.35×10^{-3}	0.092	0.076
$Z(t_0)$	0.333	0.517	0.452
ε	0.995	0.976	0.077
β	2.365	2.341	0.115
b	0.077	0.378	0.351
l	0.365	0.406	0.227
μ	0.031	0.036	0.009
p	0.007	0.068	0.051
ρ	3.897	4.008	0.461
A	11.553	12.033	0.634
R_0^I	74.821	65.245	13.808
USSR			
$S(t_0)$	1.064	0.957	0.609
$E(t_0)$	4.129	2.660	0.481
$I(t_0)$	0.954	0.980	0.151
$Z(t_0)$	1.176	1.162	0.522
ε	0.230	0.482	0.145
β	1.818	1.731	0.102
b	0.0112	0.267	0.187
l	0.730	0.649	0.247
μ	0.075	0.070	0.023
p	0.097	0.104	0.071
ρ	3.340	3.341	0.506
A	18.134	18.288	1.785
R_0^I	18.806	25.055	10.614




Fig. 7. The best fit solutions of the SEIZ model (see Table 8) vs. the data for the USA, Japan, and the USSR.

It is clear both from Table 4 and from Fig. 7 that the SEIZ model gives the best fits to data, particularly in the case of the USA.

We observed that good solutions (with very similar smallest deviation per point) are possible with either idea strand I or Z having the largest R_0 . In this sense our initial 6 years of data cannot determine which idea strand, adopters or skeptics, will eventually win out over many generation times. In the parameter estimates presented in Table 8 we have restricted the solutions to have $R_0^I > R_0^Z$, thus limiting the search space to the historically sanctioned eventual domination of Feynman diagrams over other techniques. This does not preclude the skeptics from growing initially in a population of susceptibles, and we find in fact that degenerate solutions with and without a growing number of skeptics are possible.

A novelty of the SEIZ model relative to the SEI is that the progression to adoption can result from multiple contacts, both while susceptible (parameterized by β) and while incubating (parameterized by ρ). For every country the fact that p is small and ρ sizable makes adoption favored and faster through contact with adopters while incubating, relative to simple progression as in the SEI model. This may indeed be the case in reality since the learning of Feynman diagrams in the early years was characterized by extensive interpersonal contacts at several stages of physicists' apprenticeship. We know of only one case in all three countries in which a few physicists learned about the diagrams sufficiently well from articles or textbooks alone. Practically every adopter in all three countries is known to have interacted repeatedly with other adopters before using the diagrams in their research [45].

We also observe that for the SEIZ model the relative magnitudes of the recruitment rates for the USA, USSR, and Japan follow the trends observed in simpler models, while the same is approximately true also for the effectiveness of adoption, R_0 . The estimated probability distribution function for R_0^I for Japan in the SEIZ model is shown in Fig. 8. As with the previous model, the large values of R_0^I estimated in the SEIZ model are mainly due to the very long lifetime of the idea.

Among the models discussed above we are, therefore, inclined to prefer the SEIZ model. Not only does it best fit the empirical data, but it also includes effects that we know to have been important, such as latency (apprenticeship), adoption through multiple contacts, and institutional and intellectual resistance. Estimated




Fig. 8. The probability distribution of the basic reproductive number R_0' estimated from the SEIZ model for the Japan data on the spread of Feynman diagrams. The Japanese effectiveness of adoption is consistently greater than that for the other two countries, regardless of the specific model considered.

parameters, both in their orders of magnitude and (more important) in their relative sizes, reflect properties of the idea's spread in each national community that match qualitative expectations based on our empirical knowledge of the process.

5. Conclusions

In this paper, we applied several population models, inspired by epidemiology, to the spread of a scientific idea, Feynman diagrams, in three different communities undergoing very different social transformations during the middle years of the 20th century. There is always a tradeoff between the use of models that include more detail (heterogeneous populations) and highly aggregate simple models with a manageable number of parameters. Although a model built under very simplistic assumptions is expected to have deep limitations, the use of simple epidemic-type models has had tremendous success in the recent past, partly due to their ability to use existing data to make predictions (treatment for HIV [8]) or recommendations (control measures for SARS [56]). This is the thinking behind the model choices made above, which form in our opinion the simplest synthesis of population standard modeling with sociological information at our disposal. Moreover, given the relative sparsity of quantitative data on social dynamical processes at present such models may well prove to be the most useful starting points for modeling.

We have found that suitably adapted epidemic models do a good job of fitting the empirical data, provided we allow their parameters to be very different from those normally estimated for standard epidemics. In this rough sense, the spread of Feynman diagrams appears analogous to a very slowly spreading disease, with characteristic progression times of years instead of days or weeks. The spread of the diagrams also shows an enormous effectiveness of adoption due primarily to the very long lifetime of the idea, rather than to abnormally high contact rates.

The models give a quantification of parameters that are characteristic both of the idea and of the mixing population in which it spreads. This allows a more precise discussion of the sociological reasons why the idea evolved differently in distinct national communities. The initial velocity of spread of Feynman diagrams was

fastest in Japan, followed by the USA, and slowest in the USSR, probably as a result of geopolitical constraints that severely limited access to the idea and its practitioners. The effectiveness of the adoption, encapsulated by R_0 , was consistently largest for Japan, most likely reflecting the high level of organization of its scientific community in the difficult times that followed the end of World War II. To our knowledge, this is the first time that basic reproductive number distributions have been estimated for the spread of an idea. The USA and the USSR also show high recruitment rates, following the two countries' massive investment in nuclear and high-energy physics during the early Cold War. In this study, we have done what seems to be yet uncommon in epidemiology, namely the estimation not only of model parameters and their variability, but also of the effective population sizes of the communities involved.

In the process of constructing epidemiological-type models and estimating their parameters for the spread of Feynman diagrams, we had to confront several conceptual issues concerning why the spread of ideas is or is not analogous to that of a disease. One interesting aspect of the spread of ideas is the inadequacy (or irrelevance) of the recovered state. In fact, many ideas may never be forgotten at all, as that would be in the worst interest of the adopter. As a result, our parameter estimates consistently find very long recovery times, $1/\gamma$. The same holds for the exit rates, $1/\mu$.

In spite of these slow rates of exit and recovery, individuals commonly have to acquire many ideas, and these may in some cases be mutually exclusive, or at least may adversely affect the adoption of others. We introduced a new class of simple models with multiple Z classes representing these strands. It is a curious, and we believe important fact that the recruitment of individuals from a class of susceptibles to other ideas has the same mathematical effect as vaccination against disease. In this sense “immunity” to an idea may be obtained either by education about its possible implications (perhaps analogous to actual immunization), or by distraction with other, more easily acquired concepts embodied by the Z classes.

We must emphasize that the behavior of individuals when exposed to ideas may be very different, indeed opposite, to what they may do during an epidemic outbreak. First, people intentionally seek ways to extend the infectious period of an idea, usually by recording it and storing it in various documents. In this sense, the lifetime of an idea can largely transcend that of individuals. Second, short of vaccination the most effective strategy to stop a disease epidemic is through isolation, which reduces the contact rate. Ideas, unlike diseases, are usually beneficial and thus people's behavior tends to maximize effective contacts. This pattern can be captured through the mapping of the social network of contacts that underlie the spread of the idea, which we analyze elsewhere [47]. There we show that the communities where Feynman diagrams spread the fastest had created intentional social and behavioral structures that ensured very efficient communication of scientific knowledge.

We finish by remarking that the SEIZ model, which included both skeptic and incubator classes, as well as acceleration to adoption from incubation (parameterized by ρ), captures most adequately the role of such classes in the transmission process, since it yields the best fits (smallest average deviations in Table 4). Nevertheless the modeling of the spread of ideas discussed above is but a simple caricature of the complex social dynamical processes involved. Our hope is that this work may bring a new and hopefully useful quantitative perspective into the study of the diffusion of ideas, by the simplest means possible.

Acknowledgements

We thank Gerardo Chowell, Ed MacKerrow, Miriam Nuño, Steve Tenenbaum and Alun Lloyd, for discussions and comments. A. Cintrón-Arias acknowledges financial support from Mathematical and Theoretical Biology Institute and Center for Nonlinear Studies at Los Alamos National Laboratory. Collaboration was greatly facilitated through visits by several of the authors to the Statistical and Applied Mathematical Sciences Institute (SAMSI), Research Triangle Park, NC, which is funded by NSF under grant DMS-011209. The authors thank the hospitality of Santa Fe Institute, where portions of this work were undertaken.

Appendix A. Ensemble parameter estimation procedure

Here we give a short description of our parameter estimation procedure and uncertainty quantification.

The problem of generating estimates for model parameters describing the spread of ideas is the absence of clear quantitative expectations, both concerning which model should apply best and what the quantitative value of its parameters should be. As such we devised a novel search method capable of both finding the best fit to the data possible given a choice of model, but also of producing an ensemble of solutions that are compatible with the data within a certain admissible error.

As a starting point we take the fact that simple population models cannot be expected to give perfect descriptions of the data, resulting in a minimum level of discrepancy. We chose to parameterize this discrepancy by a collective measure of the average absolute value of the deviation between the best model prediction and each data point. This measure allows us to discuss and compare how good models are at describing a specific data set. Our results are given in Table 4.

Second, we expect in general that data contains errors, e.g. early underestimation, false positives, accounting errors. Thus a given level of uncertainty in the data will translate into parameter statistical distributions that are compatible with those allowable deviations. This is a stochastic optimization problem (see, e.g., Ref. [58] for a general discussion). Based on this idea we perform an estimation of the joint parameter distribution of model parameters, conditional on a set of allowable deviations at each datum. To be specific we can write that the unknown exact data point $I^E(t_i)$, measured at time $t = t_i$, can be written in terms of the observed datum $I^O(t_i)$ and an error $\xi(t_i)$ as

$$I^E(t_i) = I^O(t_i) + \xi(t_i). \quad (13)$$

The error $\xi(t_i)$ is only known statistically. In order to proceed we must specify a model for ξ . Here we assumed a simple Gaussian distribution such that

$$P[\xi(t_i)] = P[I^E(t_i) - I^O(t_i)] = \mathcal{N}e^{-\xi^2(t_i)/2\sigma^2(t_i)}, \quad (14)$$

where \mathcal{N} is the normalization factor and $\sigma(t_i)$ parameterizes the expected error at time $t = t_i$.

This expectation for the errors can be translated into a commensurate fitness function (analogous to a Hamiltonian in statistical physics) that can in turn be minimized in order to produce parameter estimates through a search procedure. For each model realization (in terms of a set of parameters $\mathcal{S} = (S(t = t_0), E(t = t_0), \dots, \beta, \gamma, \dots)$) we take this function to be

$$H(\mathcal{S}) = \sum_i \frac{[I^M(t_i) - I^O(t_i)]^2}{2\sigma^2(t_i)}, \quad (15)$$

which is an implicit function of \mathcal{S} . If the model could generate exact results we could then make the natural association $I^E(t_i) \rightarrow I^M(t_i)$. This is usually not the case, since a residual minimal deviation always persists. To account for this, we normalize this function to zero by taking $H'(\mathcal{S}) = H - H_0$, i.e., by subtracting the minimal value of H , obtained for the best parameter set.

Given this choice of H' we can produce, in analogy with standard procedures in statistical physics, a joint probability distribution for model parameters given by

$$P(\mathcal{S}) \sim e^{-H'}. \quad (16)$$

This choice guarantees that all statistical moments are finite. This joint probability distribution can then be used to compute any moment of any set of parameters, including single parameter distribution functions, and cross-parameter correlations such as covariances. In Section 4, we show results for the single parameter averages and their standard deviations. We also show some single parameter probability distribution functions.

In general the estimation of this probability distribution can be obtained by randomly generating many model parameter sets and weighing them according to Eq. (16). The procedure is slightly complicated because we are dealing with an inverse problem in which, given a trial set of parameters, comparison with the data is performed only after the non-linear model dynamical equations have been solved. Fortunately, for models that consist of small numbers of ordinary differential equations the computational effort is not prohibitive.

In practice, we used an ensemble of trial solutions, from which we select a number of best strings, according to a standard Monte Carlo procedure, weighted by Eq. (16), to generate the next generation of the ensemble.

In order to do this we introduce a mutation implemented in terms of random Gaussian noise around the best parameter sets. This yields an effective minimization method, capable of exploring large regions of parameter space. It also creates as a byproduct an ensemble of good strings with small deviations to the data. For small enough deviations from the best string we can sample parameter space in an unbiased manner. It is this ensemble, and its best string, that is then used to estimate Eq. (16). Results given in Section 4 involve ensembles with several million realizations and a choice of σ , common to all points, corresponding to 10% deviation between the best parameter estimate and other ensemble members.

References

- [1] R.M. May (Ed.) in: *Theoretical Ecology: Principles and Applications*, second ed., Sinauer, Sunderland, 1981.
- [2] P. Yodzis, *Introduction to Theoretical Ecology*, Harper & Row, New York, 1989.
- [3] H. Thieme, *Mathematics in Population Biology*, Princeton University Press, Princeton, 2003.
- [4] R. Anderson, R. May, *Infectious Diseases of Humans: Dynamics and Control*, Oxford University Press, Oxford, 1991.
- [5] O. Diekmann, J.A.P. Heesterbeek, *Mathematical Epidemiology of Infectious Diseases: Model Building, Analysis and Interpretation*, Wiley, New York, 2000.
- [6] L. Allen, *An Introduction to Stochastic Processes with Applications to Biology*, Pearson Education–Prentice-Hall, New Jersey, 2003.
- [7] F. Brauer, C. Castillo-Chávez, *Mathematical Models in Population Biology and Epidemiology*, Springer, New York, 2001.
- [8] A.S. Perelson, et al., *Science* 271 (5255) (1996) 1582–1586.
- [9] W. Huang, K. Cooke, C. Castillo-Chávez, *SIAM J. Appl. Math.* 52 (1990) 835–854.
- [10] D. Watts, *Proc. Natl. Acad. Sci. USA* 99 (2002) 5766–5771.
- [11] S. Bikhchandani, D. Hirshleifer, I. Welch, *J. Polit. Econ.* 100 (1992) 992–1026.
- [12] L.M.A. Bettencourt, (<http://xxx.lanl.gov/abs/cond-mat/0212267>).
- [13] D. Strang, M. Macy, *Am. J. Sociol.* 107 (2001) 147–182.
- [14] E. Rogers, *Diffusion of Innovations*, Free Press, New York, 1995.
- [15] C. Castillo-Chávez, B. Song, in: H.T. Banks, C. Castillo-Chávez (Eds.), *Bioterrorism: Mathematical Modeling Applications in Homeland Security*, SIAM Frontiers in Applied Mathematics, vol. 28, SIAM, Philadelphia, 2003, p. 155.
- [16] F. Sánchez-Peña, C. Castillo-Chávez, in preparation.
- [17] B. González, et al., *J. Math. Psychol.* 47 (2003) 515–526.
- [18] A. Rapoport, *Bull. Math. Biophys.* 15 (1953) 523–533.
- [19] D.J. Daley, D.G. Kendall, *J. Inst. Math. Appl.* 1 (1965) 42–55.
- [20] L. Adamic, B. Huberman, in: E. Ben-Naim, H. Frauenfelder, Z. Toroczkai (Eds.), *Complex Networks*, Lecture Notes in Physics, vol. 650, Springer, Berlin, 2004, p. 371.
- [21] E. Adar, Z. Li, L.A. Adamic, R. Lukose, in: *Workshop on the weblogging ecosystem*, 13th International World Wide Web Conference, New York, 2004; L.A. Adamic, E. Adar, *Soc. Networks* 25(3) (2003) 211–230.
- [22] D. Kempe, J. Kleinberg, in: *Proceedings 43rd Symposium on Foundations of Computer Science*, IEEE Computer Society, Los Alamitos, 2002, pp. 471–480.
- [23] W. Kermack, A. McKendrick, *Proc. R. Soc. London Ser. A* 115 (772) (1927) 700–721.
- [24] H. Hethcote, *SIAM Rev.* 42 (2000) 599–653.
- [25] Y. Moreno, R. Pastor-Satorras, A. Vespignani, *Eur. Phys. J. B* 26 (2002) 521–529.
- [26] R.L. Rosnow, *Am. Psychol.* 46 (1991) 484–495.
- [27] P. Bordia, N. DiFonzo, *Asian J. Soc. Psychol.* 5 (2002) 49–61.
- [28] J.J. Brown, P.H. Reingen, *J. Consum. Res.* 14 (1987) 350–362.
- [29] D. Kempe, J. Kleinberg, E. Tardos, in: *Proceedings of 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ACM, New York, 2003.
- [30] B. Allen, *J. Math. Sociol.* 8 (1982) 265–281.
- [31] L.L. Cavalli-Sforza, M.W. Feldman, *Cultural Transmission and Evolution: A Quantitative Approach*, Princeton University Press, Princeton, NJ, 1981.
- [32] D.J. Bartholomew, *Stochastic Models for Social Processes*, Wiley, New York, 1982.
- [33] G.R. Funkenhouser, M.E. McCombs, *J. Math. Sociol.* 2 (1972) 121–130.
- [34] R.K. Karmeshu, R.K. Pathria, *J. Math. Sociol.* 7 (1980) 59–71.
- [35] P.S. Dodds, D.J. Watts, *J. Theor. Biol.* 232 (2005) 587–604.
- [36] W. Goffman, *Nature* 212 (5061) (1966) 449–452.
- [37] For a recent review see A.N. Tabah, in: M.E. Williams (Ed.), *Annual Review of Information Science and Technology (ASIS)*, vol. 34, Information Today, Medford, 1999, p. 249.
- [38] R.E. Dickinson, C.E.M. Pearce, *Math. Comput. Model.* 38 (2003) 1157–1167.
- [39] D.H. Zanette, *Phys. Rev. E* 65 (2002) 041908.
- [40] Y. Moreno, M. Nekovee, A.F. Pacheco, *Phys. Rev. E* 69 (2004) 066130.
- [41] K. Thompson et al., *Mathematical and Theoretical Biology Institute Technical Report*, 2003, unpublished.

- [42] F.M. Bass, *Manage. Sci.* 15 (1969) 215–277.
- [43] D.P. Fan, *J. Math. Sociol.* 11 (1985) 1–23.
- [44] D.P. Fan, R.D. Cook, *J. Math. Sociol.* 27 (2003) 29–51.
- [45] D.I. Kaiser, *Drawing Theories Apart: The Dispersion of Feynman Diagrams in Postwar Physics*, University of Chicago Press, Chicago, 2005.
- [46] D.I. Kaiser, K. Ito, K. Hall, *Soc. Stud. Sci.* 34 (6) (2004) 879–922.
- [47] L.M.A. Bettencourt, D.I. Kaiser, in preparation.
- [48] S.S. Schweber, *QED and the Men Who Made It: Dyson, Feynman, Schwinger, and Tomonaga*, Princeton University Press, Princeton, 1994.
- [49] W. Goffman, V.A. Newill, *Nature* 204 (1964) 225–228.
- [50] W. Goffman, V.A. Newill, *Proc. R. Soc. London A* 298 (1967) 316–334.
- [51] W. Goffman, *J. Assoc. Comput. Mach.* 18 (1971) 173–185.
- [52] D.I. Kaiser, *Hist. Stud. Phys. Bio. Sci.* 33 (1) (2002) 131–159.
- [53] It is customary to model the movements out of the class E into the next class I by a term like εE . This corresponds to having exponentially distributed waiting times in the E class. In other words, the simple progression rate εE corresponds to $\mathcal{P}(\tau) = \exp(-\varepsilon\tau)$ as the fraction that is still in the incubator class τ units after entering this class, and to $1/\varepsilon$ as the mean waiting time.
- [54] C. Castillo-Chávez, Z. Feng, W. Huang, in: C. Castillo-Chávez, et al. (Eds.), *Mathematical Approaches for Emerging and Reemerging Infectious Diseases*, The IMA Volumes in Mathematics and its Applications, vol. 125, Springer, New York, 2002, p. 229.
- [55] P. van den Driessche, J. Watmough, *Math. Biosci.* 180 (2002) 29–48.
- [56] G. Chowell, P.W. Fenimore, M.A. Castillo-Garsow, C. Castillo-Chávez, *J. Theor. Biol.* 224 (2003) 1–8.
- [57] F. Dyson, *Nature* 427 (2004) 297.
- [58] J.C. Spall, *Introduction to Stochastic Search and Optimization*, Wiley-Interscience, Hoboken, New Jersey, 2003.