

Modeling bursts and heavy tails in human dynamics

Alexei Vázquez,^{1,2} João Gama Oliveira,^{2,3} Zoltán Dezső,² Kwang-Il Goh,^{1,2} Imre Kondor,⁴ and Albert-László Barabási^{1,2}

¹*Center for Cancer System Biology, Dana Farber Cancer Institute, Harvard Medical School, 44 Binney St, Boston, Massachusetts 02115, USA*

²*Department of Physics and Center for Complex Networks Research, University of Notre Dame, Notre Dame, Indiana 46556, USA*

³*Departamento de Física, Universidade de Aveiro, Campus Universitário de Santiago, 3810-193 Aveiro, Portugal*

⁴*Collegium Budapest, Szentháromság u. 2, H-1014 Budapest, Hungary*

(Received 18 October 2005; published 24 March 2006)

The dynamics of many social, technological and economic phenomena are driven by individual human actions, turning the quantitative understanding of human behavior into a central question of modern science. Current models of human dynamics, used from risk assessment to communications, assume that human actions are randomly distributed in time and thus well approximated by Poisson processes. Here we provide direct evidence that for five human activity patterns, such as email and letter based communications, web browsing, library visits and stock trading, the timing of individual human actions follow non-Poisson statistics, characterized by bursts of rapidly occurring events separated by long periods of inactivity. We show that the bursty nature of human behavior is a consequence of a decision based queuing process: when individuals execute tasks based on some perceived priority, the timing of the tasks will be heavy tailed, most tasks being rapidly executed, while a few experiencing very long waiting times. In contrast, priority blind execution is well approximated by uniform interevent statistics. We discuss two queuing models that capture human activity. The first model assumes that there are no limitations on the number of tasks an individual can handle at any time, predicting that the waiting time of the individual tasks follow a heavy tailed distribution $P(\tau_w) \sim \tau_w^{-\alpha}$ with $\alpha = 3/2$. The second model imposes limitations on the queue length, resulting in a heavy tailed waiting time distribution characterized by $\alpha = 1$. We provide empirical evidence supporting the relevance of these two models to human activity patterns, showing that while emails, web browsing and library visitation display $\alpha = 1$, the surface mail based communication belongs to the $\alpha = 3/2$ universality class. Finally, we discuss possible extension of the proposed queuing models and outline some future challenges in exploring the statistical mechanics of human dynamics.

DOI: [10.1103/PhysRevE.73.036127](https://doi.org/10.1103/PhysRevE.73.036127)

PACS number(s): 89.75.Da, 02.50.-r

I. INTRODUCTION

Humans participate on a daily basis in a large number of distinct activities, from electronic communication, such as sending emails or browsing the web, to initiating financial transactions or engaging in entertainment and sports. Given the number of factors that determine the timing of each action, ranging from work and sleep patterns to resource availability, it appears impossible to seek regularities in the apparently random human activity patterns, apart from the obvious daily and seasonal periodicities. Therefore, in contrast with the accurate predictive tools common in physical sciences, forecasting human and social patterns remains a difficult and often elusive goal. Yet, the need to understand the timing of human actions is increasingly important. Indeed, uncovering the laws governing human dynamics in a quantitative manner is of major scientific interest, requiring us to address the factors that determine the timing of human actions. But these questions are application driven as well: most human actions have a strong impact on resource allocation, from phone line availability and bandwidth allocation in the case of Internet or Web use, all the way to the design of physical space for retail or service oriented institutions. Despite these fundamental and practical driving forces, our understanding of the timing of human initiated actions is rather limited at present.

To be sure, the interest in addressing the timing of events in human dynamics is not new: it has a long history in the mathematical literature, leading to the development of some of the key concepts in probability theory [1], and has re-emerged at the beginning of the 20th century when the design problems surrounding the phone system required a quantitative understanding of the call patterns of individuals. But most current models of human activity assume that human actions are performed at constant rate, meaning that a user has a fixed probability to engage in a specific action within a given time interval. These models approximate the timing of human actions with a Poisson process, in which the time interval between two consecutive actions by the same individual, called the waiting or interevent time, follows an exponential distribution [2]. Poisson processes are at the heart of the celebrated Erlang formula [3], predicting the number of phone lines required in an institution, and they represent the basic approximation in the design of most currently used Internet protocols and routers [4]. Yet, the availability of large datasets recording selected human activity patterns increasingly question the validity of the Poisson approximation. Indeed, an increasing number of recent measurements indicate that the timing of many human actions systematically deviate from the Poisson prediction, the waiting or interevent [73] times being better approximated by a heavy tailed or Pareto distribution [5–8]. The difference between a Poisson and a heavy tailed behavior is striking: the

exponential decay of a Poisson distribution forces the consecutive events to follow each other at relatively regular time intervals and forbids very long waiting times. In contrast, the slowly decaying heavy tailed processes allow for very long periods of inactivity that separate bursts of intensive activity.

We have recently proposed that the bursty nature of human dynamics is a consequence of a queuing process driven by human decision making [5]: whenever an individual is presented with multiple tasks and chooses among them based on some perceived priority parameter, the waiting time of the various tasks will be Pareto distributed. In contrast, first-come-first-serve and random task execution, common in most service oriented or computer driven environments, lead to a uniform Poisson-type dynamics. Yet, this work has generated just as many questions as it resolved. What are the different classes of processes that are relevant for human dynamics? What determines the scaling exponents? Do we have discrete universality classes (and if so how many) as in critical phenomena [9], or the exponents characterizing the heavy tails can take up arbitrary values, as it is the case in network theory [10–13]? Is human dynamics always heavy tailed?

In this paper we aim to address some of these questions by studying the different universality classes that can appear as a result of the queuing of human activities. We first review, in Sec. II, the frequently used Poisson approximation, which predicts an exponential distribution of interevent times. In Sec. III we present evidence that the interevent time probability density function $P(\tau)$ of many human activities is characterized by the power law tail

$$P(\tau) \sim \tau^{-\alpha}. \quad (1)$$

In Sec. IV we discuss the general characteristics of queuing models that govern how humans time their various activities. In Secs. V and VI we study two classes of queuing models designed to capture human activity patterns and document the existence of two distinct universality classes, one characterized by $\alpha=3/2$ (Sec. V) and the other by $\alpha=1$ (Sec. VI). In Sec. VII we discuss the relationship between interevent and waiting times. Finally, in Sec. VIII we discuss the applicability of these models to explain the empirical data, as well as outline future challenges in modeling human dynamics.

II. POISSON PROCESSES

Consider an activity performed with some regularity, such as sending emails or browsing the web. We can keep track of this activity by recording the timing of each event, for example, the time each email is sent by an individual. The time between two consecutive events we call *the interevent time* for the monitored activity and will be denoted by τ . Given that the interevent time can be explicitly measured for selected activities, it serves as a test of our ability to understand and model human dynamics: proper models should be able to capture its statistical properties.

The most primitive model of human activity would assume that human actions are fundamentally periodic, with a period determined by the daily sleep patterns. Yet, while some periodicity is certainly present, the timing of most hu-

man actions are highly stochastic. Indeed, periodic models are hopeless in capturing the time we check out a book from the library, beyond telling us that it should be within the library's operation hours. The first and still most widely used stochastic model of human activity assumes that the tasks are executed independently from each other at a constant rate λ , so that the time resolved activity of an individual is well approximated by a Poisson process [2]. In this case the probability density function of the recorded interevent times has the exponential form

$$P(\tau) = \lambda e^{-\lambda\tau}. \quad (2)$$

In practice this means that the predicted activity pattern, while stochastic, will display some regularity in time, events following each other on average at $\tau \approx \langle \tau \rangle = 1/\lambda$ intervals. Indeed, given that for a Poisson process $\sigma = \sqrt{\langle \tau^2 \rangle - \langle \tau \rangle^2} = \langle \tau \rangle$ is finite, very long interevent times are exponentially rare. This is illustrated in Fig. 1(a), where we show a sequence of events generated by a Poisson process, appearing uniformly distributed in time (but not periodic).

The Poisson process was originally introduced by Poisson in his major work applying probability concepts to the administration of justice [14]. Today it is widely used to quantify the consequences of human actions, such as modeling traffic flow patterns or accident frequencies [2], and is commercially used in call center staffing [15], inventory control [16], or to estimate the number of congestion caused blocked calls in mobile communications [4]. It has been established as a basic model of human activity patterns at a time when data collection capabilities on human behavior were rather limited. In the past few years, however, thanks to detailed computer based data collection methods, there is increasing evidence that the Poisson approximation fails to capture the timing of many human actions.

III. EMPIRICAL RESULTS

Evidence that non-Poisson activity patterns characterize human activity has first emerged in computer communications, where the timing of many human driven events is automatically recorded. For example, measurements identifying the distribution of the time differences between consecutive instant messages sent by individuals during on-line chats [17] have found evidence of heavy tailed statistics. Professional tasks, such as the timing of job submissions on a supercomputer [18], directory listings and file transfers (FTP requests) initiated by individual users [19] were also reported to display non-Poisson features. Similar patterns emerge in economic transactions [20,21], in the number of hourly trades in a given security [22] or the time interval distribution between individual trades in currency futures [23]. Finally, heavy tailed distributions characterize entertainment related events, such as the time intervals between consecutive online games played by users [24]. Note, however, that while these datasets provide clear evidence for non-Poisson human activity patterns, most of them do not resolve individual human behavior, but capture only the aggregated behavior of a large number of users. For example, the dataset recording the timing of the job submissions looks




FIG. 1. The difference between the activity patterns predicted by a Poisson process (top) and the heavy tailed distributions observed in human dynamics (bottom). (a) Succession of events predicted by a Poisson process, which assumes that in any moment events take place with probability q . The horizontal axis denotes time, each vertical line corresponding to an individual event. Note that the interevent times are comparable to each other, long delays being virtually absent. (b) The absence of long delays is visible on the plot showing the delay times τ for 1000 consecutive events, the size of each vertical line corresponding to the gaps seen in (a). (c) The probability of finding exactly n events within a fixed time interval is $\mathcal{P}(n; q) = e^{-qt}(qt)^n/n!$, which predicts that for a Poisson process the interevent time distribution follows $P(\tau) = qe^{-q\tau}$, shown on a log-linear plot in (c) for the events displayed in (a) and (b). (d) The succession of events for a heavy tailed distribution. (e) The waiting time τ of 1000 consecutive events, where the mean event time was chosen to coincide with the mean event time of the Poisson process shown in (a)–(c). Note the large spikes in the plot, corresponding to very long delay times. (b) and (e) have the same vertical scale, allowing to compare the regularity of a Poisson process with the bursty nature of the heavy tailed process. (f) Delay time distribution $P(\tau) \approx \tau^{-2}$ for the heavy tailed process shown in (d) and (e), appearing as a straight line with slope -2 on a log-log plot. The signal shown in (d)–(f) was generated using $\gamma=1$ in the stochastic priority list model discussed in Appendix C.

at the timing of *all jobs* submitted to a computer, by any user. Thus for these measurements the interevent time does not characterize a *single* user but rather a *population* of users. Given the extensive evidence that the activity distribution of the individuals in a population is heavy tailed, these measurements have difficulty capturing the origin of the observed heavy tailed patterns. For example, while most people send only a few emails per day, a few send a very large number on a daily basis [25,26].

If the activity pattern of a large number of users is simultaneously captured, it is not clear where the observed heavy tails come from: are they rooted in the activity of a single individual, or rather in the heavy tailed distribution of user activities? Therefore, datasets that capture the long term activity pattern of a *single* individual are of particular value. To our best knowledge a few papers have taken this approach, capturing the timing of printing jobs submitted by users [27], the email activity patterns of individual email users [5,25] and the browsing pattern of users visiting a major web portal [7] and the correspondence patterns of famous scientists [6]. These measurements offer direct evidence that the heavy

tailed activity patterns emerge at the level of a *single* individual, and are not a consequence of the heterogeneous distribution of user activity. Despite this evidence, a number of questions remain unresolved: Is there a single scaling exponent characterizing all users, or rather each user has its own exponent? What is the range of these exponents? Next we aim to address these questions through the study of six datasets, each capturing individual human activity patterns of different nature. First we describe the datasets and the collection methods, followed by a quantitative characterization of the observed human activity patterns.

Web browsing: Automatically assigned cookies allow us to reconstruct the browsing history of approximately 250 000 unique visitors of the largest Hungarian news and entertainment portal (<http://www.origo.hu>), which provides online news and magazines, community pages, software downloads, free email and search engine, capturing 40% of all internal Web traffic in Hungary [7,28]. The portal receives 6 500 000 HTML hits on a typical work day. We used the log files of the portal to collect the visitation pattern of each visitor between 11/08/02 and 12/08/02, recording with second resolu-

tion the timing of each download by each visitor [7]. The interevent time, τ , was defined as the time interval between consecutive page downloads (clicks) by the same visitor.

Email activity patterns: This dataset contains the email exchange between individuals in a university environment, capturing the sender, recipient, and the time of each email sent during a 3 and 6 month period by 3188 [25] and 9665 [26] users, respectively. We focused here on the data collected by Eckmann [25], which records 129 135 emails with second resolution. The interevent time corresponds to the time between two consecutive emails sent by the same user.

Library loans: The data contains the time with minute resolution at which books or periodicals were checked out from the library by the faculty at University of Notre Dame during a 3 year period. The number of unique individuals in the dataset is 2247, together participating in a total of 48 409 transactions. The interevent time corresponds to the time difference between consecutive books or periodicals checked out by the same patron.

Trade transactions: A dataset recording all transactions (buy/sell) initiated by a stock broker at a Central European bank between June 1999 and May 2003 helps us quantify the professional activity of a single individual, giving a glimpse on the human activity patterns driving economic phenomena. In a typical day the first transactions start at 7:00 am and end at 7:00 pm and the average number of transactions initiated by the dealer in one day is around 10, resulting in a total of 54 374 transactions. The interevent time represents the time between two consecutive transactions by the broker. The gap between the last transaction at the end of one day and the first transaction at the beginning of the next trading day was ignored.

The correspondence patterns of Einstein, Darwin, and Freud: We start from a record containing the sender, recipient, and the date of each letter [29–31] sent or received by the three scientists during their lifetime. The databases used in our study were provided by the Darwin Correspondence Project (<http://www.lib.cam.ac.uk/Departments/Darwin/>), the Einstein Papers Project (<http://www.einstein.caltech.edu/>) and the Freud Museum of London (<http://www.freud.org.uk>). Each dataset contains the information about each sent/received letter in the following format: SENDER, RECIPIENT, DATE, where either the sender or the recipient is Einstein, Darwin or Freud. The Darwin dataset contained a record of a total of 7591 letters sent and 6530 letters received by Darwin (a total of 14 121 letters). Similarly, the Einstein database contained 14 512 letters sent and 16 289 letters received (total of 30 801). For Freud we have 3183 (2675) sent (received) letters. Note that 1541 letters in the Darwin database and 1861 letters in the Einstein database were not dated or were assigned only potential time intervals spanning days or months. We discarded these letters from the dataset. Furthermore, the dataset is naturally incomplete, as not all letters written or received by these scientists were preserved. Yet, assuming that letters are lost at a uniform rate, they should not affect our main findings. For these three datasets we do not focus on the interevent times, but rather the *response* or *waiting times* τ_w . The waiting time, τ_w , represents the time interval between the date of a letter received from a given person, and the date of the next letter from Darwin, Einstein

or Freud to him or her, i.e., the time the letter waited on their desk before a response is being sent. To analyze Einstein, Darwin, and Freud's response time we have followed the following procedure: if individual A sent a letter to Einstein on DATE1, we search for the next letter from Einstein to individual A, sent on DATE2, the response time representing the time difference $\tau = \text{DATE2} - \text{DATE1}$, expressed in days. If there are multiple letters from Einstein to the recipient, we always consider the first letter as the response, and discard the later ones. Missing letters could increase the response time, the magnitude of this effect depending on the overall frequency of communication between the respective correspondence partners. Yet, if the response time follows a distribution with an exponential tail, then randomly distributed missing letters would not generate a power law waiting time. Thus the observed power law cannot be attributed to data incompleteness.

In the following we will break our discussion into three sections, each focusing on a specific class of behavior observed in the studied individual activity patterns.

A. The $\alpha=1$ universality class: Web browsing, email, and library datasets

In Figs. 2(a)–2(c) we show the interevent time distribution between consecutive events for a single individual for Web browsing, email, and library visitation. For these datasets we find that for intermediate times the interevent time distribution has a power law tail

$$P(\tau) \sim \tau^{-\alpha} \quad (3)$$

with exponent $\alpha \approx 1$ independent of the nature of the activity. Given that for these activity patterns we collected data for thousands of users, we need to calculate the distribution of the exponent α determined separately for each user whose activity level exceeds a certain threshold [i.e., avoiding users that have too few events to allow a meaningful determination of $P(\tau)$]. As Figs. 2(e)–2(g) shows, we find that the distribution of the exponents is peaked around $\alpha=1$. Note that the scaling in (3) cannot hold for infinitely long times because the resulting probability density cannot be normalized. Therefore a natural cutoff should emerge for large τ .

The scattering around $\alpha=1$ in the measured exponents could have two different origins. First, it is possible that each user is characterized by a different scaling exponent α . Second, each user could have the same exponent $\alpha=1$, but given that the available dataset captures only a finite time interval from one to several months, with at best a few thousand events in this interval, there are uncertainties in our ability to determine numerically α . To demonstrate that such data incompleteness could indeed explain the observed scattering, in Figs. 2(h) and 2(i) we show the result of a numerical experiment, in which we generated 10 000 time series, corresponding to 10 000 independent users, the interevent time of the events for each user being taken from the same distribution $P(\tau) \sim \tau^{-1}$. The total length in time of each time series was chosen to be 1 000 000. We then used the automatic fitting algorithm employed earlier to measure the exponents in Figs. 2(e)–2(g) to determine numerically the exponent α

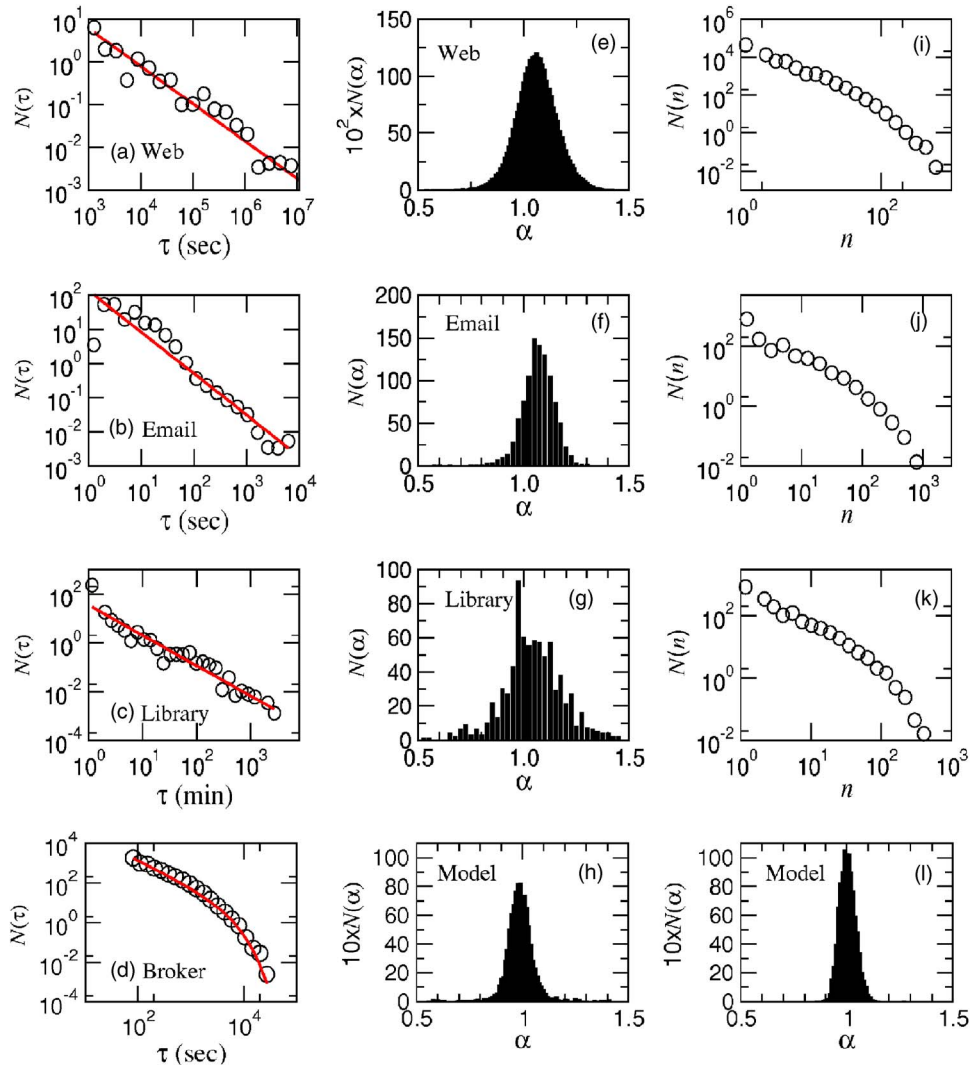


FIG. 2. (Color online) (a) The interevent time distribution between (a) two consecutive visits of a web portal by a single user; (b) two consecutive emails sent out by a user and (c) two consecutive library loans made by a single individual. For (a)–(c) we show as a straight line the $\alpha=1$ scaling. (d) The interevent time distribution between two consecutive transactions made by a stock broker. The distribution follows a power law with the exponential cutoff $P(\tau) \sim \tau^{-1.3} \exp(-\tau/\tau_0)$. (e)–(g) The distribution of the exponents (α) characterizing the interevent time distribution of users browsing the web portal (e), individual loans from the library (f) and the emails sent by different individuals (g). The exponent α was determined only for users whose total activity levels exceeded certain thresholds, the values used being 15 web visits (e), 15 emails (f) and 10 books (g). (h) and (l) We numerically generate for 10 000 individuals interevent time distributions following a power law with exponent $\alpha=1$. The distribution of the measured exponents follows a normal distribution similar to the distribution observed in (e)–(g). If we double the time window of the simulation (h) the deviation around the average becomes much smaller (l). (i)–(k) The distribution of the number of events in the studied systems: number of HTML hits for each user (i), the number of books checked out by each user (j) and the number of emails sent by different individuals (k), indicating that the overall activity patterns of individuals is also heavy tailed.

for each user. In principle for each user we should observe the same exponent $\alpha=1$, given that the datasets were generated in an identical fashion. In practice, however, due to the finite length of the data, each numerically determined exponent is slightly different, resulting in the histogram shown in Fig. 2(h). As the figure shows, even in this well controlled situation we observe a scattering in the measured exponents, obtaining a distribution similar to the one seen in Figs. 2(e)–2(g). The longer the time series, the sharper the distribution is [Fig. 2(l)], given that the exponent α can be determined more accurately.

The distributions obtained for the three studied datasets are not as well controlled as the one used in our simulation: while the length of the observation period is the same for each user, the activity level of the users differs widely. Indeed, as we show in Figs. 2(i)–2(k), the activity distribution of the different users, representing the number of events recorded for each user, also spans several orders of magnitude, following a fat tailed distribution. Thus the degree of scattering of the measured exponent α is expected to be more significant than seen in Figs. 2(h) and 2(l), since we can determine the exponent accurately only for very active users, for

which we have a significant number of datapoints. Therefore, the obtained results are consistent with the hypothesis that each user is characterized by a scaling exponent in the vicinity of $\alpha=1$, the difference in the numerically measured exponent values being likely rooted in the finite number of events we record for each user in the datasets. This conclusion will be eventually corroborated by our modeling efforts, that indicate that the exponents characterizing human behavior take up discrete values, one of which, providing the empirically observed $\alpha=1$.

As we will see in the following sections, an important measure of the human activity patterns is the *waiting (or response) time*, τ_w , representing the amount of time a task waits on an individual's priority list before being executed. For the email dataset we define the waiting time as the difference between the time user A receives an email from user B, and the time A sends an email to user B. In looking at this quantity we should be aware of the fact that not all emails A sends to B are direct responses to emails received from B, thus there are some false positives in the data that could be filtered out only by reading the text of each email (which is not possible in the available datasets). We have measured the empirically obtained waiting time distribution in the email dataset, finding that the distribution of the response times indeed follows a power law with exponent $\alpha=1$ [Fig. 3(a)].

B. The $\alpha=3/2$ universality class: The correspondence of Einstein, Darwin, and Freud

In the case of the mail correspondence patterns we will focus on the response time of the authors, partly because we will see later that this has the most importance from the modeling perspective. As shown in Fig. 4, the probability that a letter will be replied in τ_w days is well approximated by a power law (3) with $\alpha=3/2$, the scaling spanning four orders of magnitude, from days to years. Note that this exponent is significantly different from $\alpha=1$ observed in the earlier datasets, and we will show later that modeling efforts indeed establish $\alpha=3/2$ as a scaling exponent characterizing human dynamics.

The dataset allows us to determine the interevent times as well, representing the time interval between two consecutive letters sent by Einstein, Darwin or Freud to any recipient. We find that the interevent time distribution is also heavy tailed, albeit the quality of scaling is not as impressive as we observe for the response time distribution, and the scaling has a different origin than the one observed in the response time distribution. This is due to the fact that we know only the day when the letter is written, providing a day resolution, in contrast with the email, which is known with second resolution. Given that both Einstein and Darwin wrote at least one letter most days, this means that long interevent times are rarely observed. Furthermore, owing to the long observational period (over 70 years), the overall activity pattern of the two scientists has changed significantly, going from a few letters per year to as many as 400–800 letters/year during the later, more famous phase of their professional life. Thus the interevent time, while it appears to follow a power law distribution, it is by no means stationary. In fact, we believe that this

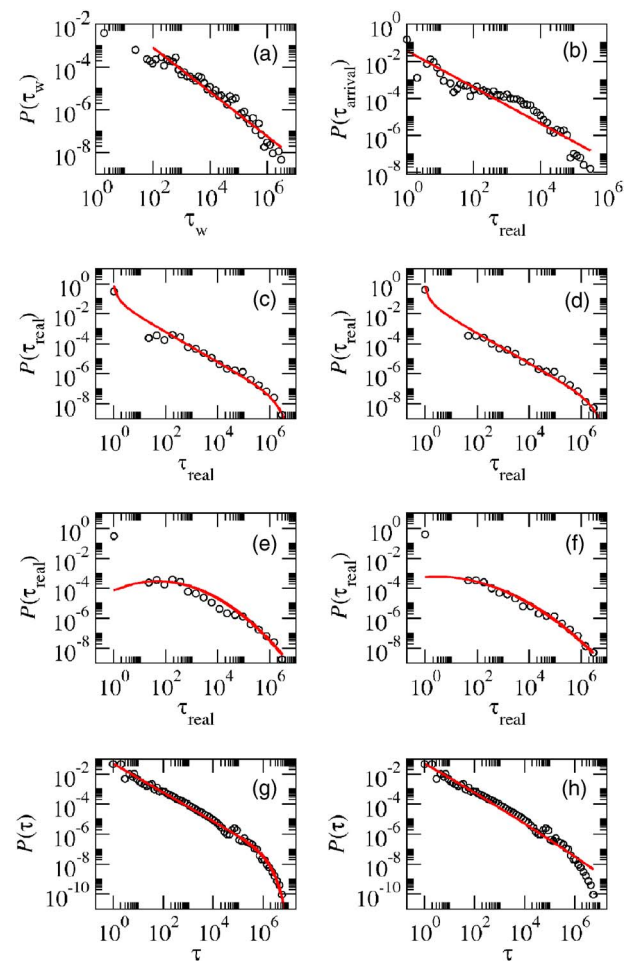


FIG. 3. (Color online) Distribution of the response and arrival time intervals of the email user shown in Fig. 2(b). (a) Given two email users A and B, the response times of user A to B are the time intervals between A receiving an email from B and A sending an email to B. The response time distribution of user A is then computed taking into account the response times to all users he/she communicates with. The continuous line is a power law fit with exponent $\alpha=1.0$. (b) Given a user A, the interarrival times are the time intervals between the two consecutive arrivals of an email to user A, independently of the sender. The arrival time distribution of user A is obtained taking into account all the interarrival times for that user. The continuous line is a power law fit with exponent 0.98. (c)–(f) The real waiting time distribution of two email users (left and right), where τ_{real} represents the time between the time the user first sees an email and the time she sends a reply to it. The black symbol shown in the upper left-hand corner represents the messages that were replied to right after the user has noticed it. (c) and (d) The line is a fit to (8) resulting in $p=0.999\,999\pm 0.000\,005$ for both users. (e) and (f) The line is a fit to a log-normal distribution.

increase and nonstationarity is responsible for the power law interevent times. Indeed, when we look at only a 10 year slice of their activity pattern, an exponential distribution is observed for the interevent times. For these reasons the interevent time distribution cannot be used as a proper measure of the mail dynamics, forcing us to focus on the response time distribution only.




FIG. 4. (Color online) Distribution of the response times for the letters replied to by Einstein, Darwin, and Freud, as indicated on each plot. Note that the distributions are well approximated with a power law tail with exponent $\alpha=3/2$. While for Darwin and Einstein the datasets provide very good statistics (the power law regime spanning 4 orders of magnitude), the plot corresponding to Freud’s responses is not so impressive, yet still being well approximated by the power law distribution. Note that while in most cases the identified reply is indeed a response to a received letter, there are exceptions as well: many of the very delayed replies represent the renewal of a long lost relationship.

C. The stock broker activity pattern

For the stock broker we again focus on the interevent time distribution, finding that the best fit follows $P(\tau) \sim \tau^{-\alpha} \exp(-\tau/\tau_0)$ with $\alpha=1.3$ and $\tau_0=76$ min [see Fig. 2(d)]. This value is between $\alpha=1$ observed for the users in the first three other datasets and $\alpha=3/2$ observed for the correspondence patterns. Yet, given the scattering of the measured exponents, it is difficult to determine if this represents a standard statistical deviation from $\alpha=1$ or $\alpha=3/2$, the two values expected by the modeling efforts (see Secs. V and VI), or it stands as evidence for a new universality class. At this point we believe that the former case is valid, something that can be decided only once data for more users will become available. The exponential cutoff is not inconsistent with the modelling efforts either: as shown in Ref. [8], such cutoffs are expected to accompany all human activity patterns with $\alpha < 2$.

D. Qualitative differences between heavy tailed and Poisson activity patterns

The heavy tailed nature of the observed interevent time distribution has clear visual signatures. Indeed, it implies that an individual’s activity pattern has a bursty character: short time intervals with intensive activity (bursts) are separated by long periods of no activity [Figs. 1(d)–1(f)]. Therefore, in contrast with the relatively uniform activity pattern predicted by the Poisson process, for a heavy tailed process very dense successions of events (bursts) are separated by very long gaps, predicted by the slowly decaying tail of the power law distribution. This bursty activity pattern agrees with our experience of an individual’s normal email usage pattern: during a single session we typically send several emails in quick succession, followed by long periods of no email activity, when we focus on other activities.

IV. CAPTURING HUMAN DYNAMICS: QUEUING MODELS

The empirical evidence discussed in the preceding section raises several important questions: Why does the Poisson

process fail to capture the temporal features of human activity patterns in human dynamics? To address these questions we need to inspect closely the processes that contribute to the timing of the events in which an individual participates.

Most of the time humans face simultaneously several work, entertainment, and family related responsibilities. Indeed, at any moment an individual could choose to participate in one of several tasks, ranging from shopping to sending emails, making phone calls, attending meetings or talks, going to a theater, getting tickets for a sports event, and so on. To keep track of the various responsibilities ahead of them, individuals maintain a *to do* or *priority* list, recording the upcoming tasks. While this list is occasionally written or electronically recorded, in many cases it is simply kept in memory. A priority list is a dynamic entity, since tasks are removed from it after they are either executed or become irrelevant and new tasks are added continuously. The tasks on the list compete with each other for the individual’s time and attention. Therefore, task management by humans is best described as a queuing process [32,33], where the queue represents the tasks on the priority list, the server is the individual which executes them and maintains the list, and some selection protocol governs the order in which the tasks are executed. To define the relevant queuing model we must clarify some key features of the underlying queuing process, ranging from the arrival and service processes to the nature of the task selection protocol, and the restrictions on the queue length [32].

Server: The server refers to the individual (or agent) that maintains the queue and executes the tasks. In queuing theory we can have one or several servers in parallel (like checkout counters in a supermarket). Human dynamics is a *single server* process, capturing the fact that an individual is solely responsible for executing the tasks on his/her priority list.

Task arrival pattern: The arrival process specifies the statistics of the arrival of new tasks to the queue. In queuing theory it is often assumed that the arrival is a Poisson process, meaning that new tasks arrive at a constant rate λ to the queue, randomly and independently from each other. We will use this approximation for human queues as well. If the ar-

rival process is not captured by a Poisson distribution, it can be modeled as a renewal process with a general distribution of interarrival times [32]. For example, our measurements indicate that the arrival time of emails follows a heavy tailed distribution, thus a detailed modeling of email based queues must take this into account. We must also keep in mind that the arrival rate of the tasks to the list is filtered by the individual, who decides which tasks to accept and place on the priority list and which to reject. In principle the rejection of a task is also a decision process that can be modeled as a high priority short lived task.

Service process: The service process specifies the time it takes for a single task to be executed, such as the time necessary to write an email, explore a web page or read a book. In queuing theory the service process is often modeled as a Poisson process, which means that the distribution of the time devoted to the individual tasks has the exponential form (2). However, in some applications the service time may follow some general distribution. For example, the size distribution of files transmitted by email is known to be fat tailed [34,35], suggesting that the time necessary to review (read) them could also follow a fat tailed distribution. In queuing theory it is often assumed that the service time is independent of the task arrival process or the number of tasks on the priority list. While we adopt this assumption here as well, we must also keep in mind that the service time can decrease if too many tasks are in the queue, as humans may devote less time to individual tasks when they have many things to do.

Selection protocol or queue discipline: The selection protocol specifies the manner in which the tasks in the queue are selected for execution. Most human initiated events require an individual to weight and prioritize different activities. For example, at the end of each activity an individual needs to decide what to do next: send an email, do some shopping or place a phone call, allocating time and resources for the chosen activity. Normally individuals assign to each task a priority parameter, which allows them to compare the urgency of the different tasks on the list. The time a task waits before it is executed depends on the method the individual uses to choose the task to be executed next. In this respect three selection protocols are particularly relevant for human dynamics.

(i) The simplest is the first-in-first-out (FIFO) protocol, executing the tasks in the order they were added to the list. This is common in service oriented processes, like the first-come-first-serve execution of orders in a restaurant or getting help from directory assistance and consumer support.

(ii) The second possibility is to execute the tasks in a random order, irrespective of their priority or time spent on the list. This is common, for example, in educational settings, when students are called on randomly, and in some packet routing protocols.

(iii) In most human initiated activities task selection is not random, but the individual tends to execute always the highest priority item on his/her list. The resulting execution dynamics is quite different from (i) and (ii): high priority tasks will be executed soon after their addition to the list, while low priority items will have to wait until all higher priority tasks are cleared. In the following we show that this selection mechanism, practiced by humans on a daily basis, is the

likely source of the fat tails observed in human initiated processes.

Queue length or system capacity: In most queuing models the queue has an infinite capacity and the queue length can change dynamically, depending on the arrival and the execution rate of the individual tasks. In some queuing processes there is a physical limitation on the queue length. For example, the buffers of Internet routers have finite capacity, so that packets arriving while the buffer is full are systematically dropped. In human activity one could argue that, given the possibility to maintain the priority list in a written or electronic form, the length of the list has no limitations. Yet, if confronted with too many responsibilities, humans will drop some tasks and not accept others. Furthermore, while keeping track of a long priority list is not a problem for an electronic organizer, it is well established that the immediate memory of humans has finite capacity of about seven tasks [36,37]. In other words, the number of priorities we can easily remember, and therefore the length of our priority list, is bounded. These considerations force us to inspect closely the difference between finite and unbounded priority lists, and the potential consequences of the queue length on the waiting time distribution.

In this paper we follow the hypothesis that the empirically observed heavy tailed distributions originate in the queuing process of the tasks maintained by humans, and seek appropriate models to explain and quantify this phenomenon. Particularly valuable are queuing models that do not contain power law distributions as inputs, and yet generate a heavy tailed output. In the following we will focus on priority queues, reflecting the fact that humans most likely choose the tasks based on their priority for execution.

In the empirical datasets discussed in Sec. III we focused on both the interevent time and the waiting (response) time distribution of the tasks in which humans participate. In the following two sections we focus on the *waiting time* of a task on the priority list rather than the interevent times. In this context the waiting time, τ_w , represents the time difference between the arrival of a task to the priority list and its execution, thus it is the sum of the time a task waits on the list and the time devoted to executing it. In Sec. VII we will return to the relationship between the empirically observed interevent times and the waiting times predicted by the discussed models.

V. MODELS WITH VARIABLE QUEUE LENGTH: $\alpha=3/2$ UNIVERSALITY CLASS

Our first goal is to explore the behavior of priority queues in which there are no restrictions on the queue length. Therefore, in these models an individual's priority list could contain arbitrary number of tasks. As we will show below, such models offer a good approximation to the surface mail correspondence patterns, such as that observed in the case of Einstein, Darwin, and Freud (see Sec. III B). Therefore, we will construct the models with direct reference to the the datasets discussed in Sec. III. We assume that letters arrive at rate λ following a Poisson process with exponential arrival time distribution. Replacing letters with tasks, however, pro-

vides us a more general model, in principle applicable to any human activity. The responses are written at rate μ , reflecting the overall time a person devotes to his correspondence. Each letter is assigned a discrete priority parameter $x = 1, 2, \dots, r$ upon arrival, such that always the highest priority unanswered letter (task) will be always chosen for a reply. The lowest priority task will have to wait the longest before execution, and therefore it dominates the waiting time probability density for large waiting times. This model was introduced in 1954 by Cobham [38] to describe some manufacturing processes. Most of the analytical work in queuing theory has concentrated on the waiting time of the lowest priority task, finding that the waiting time distribution follows [39]

$$P(\tau_w) \sim A \tau_w^{-3/2} \exp\left(-\frac{\tau_w}{\tau_0}\right), \quad (4)$$

where A and τ_0 are functions of the model parameters, the characteristic waiting time τ_0 being given by

$$\tau_0 = \frac{1}{\mu(1 - \sqrt{\rho})^2}, \quad (5)$$

where $\rho = \lambda/\mu$ is the traffic intensity. Therefore, the waiting time distribution is characterized by a power law decay with exponent $\alpha=3/2$, combined with an exponential cutoff.

The model can be extended to the case where the priorities are not discrete, but take up continuous values $0 \leq x < \infty$ from an arbitrary $\eta(x)$ distribution. The Laplace transform of the waiting time distribution for this case has been calculated in Ref. [32], but the resulting equation is difficult to invert, forcing us to study the model numerically (Fig. 5). The natural control parameter is $\rho = \lambda/\mu$, allowing us to distinguish three qualitatively different regimes.

Subcritical regime, $\rho < 1$: Given that the arrival rate of the tasks is smaller than the execution rate, the queue will be often empty. This significantly limits the waiting time, most tasks being executed soon after their arrival. The simulations indicate that the waiting time distribution exhibits an asymptotic scaling behavior consistent with (4) (Fig. 5). While in the $\rho \rightarrow 0$ limit we observe mainly the exponential decay, as ρ approaches 1 a power law regime with exponent $\alpha=3/2$ emerges, combined with the exponential cutoff.

Critical regime, $\rho = 1$: When the arrival and the response rate of the letters are equal, according to (4) and (5) we should observe a power law waiting time distribution with $\alpha=3/2$ (Fig. 5). In this case it is easy to show that the queue length performs a one-dimensional random walk bounded at $l=0$. These fluctuations in the queue length will limit the waiting time distribution, as the tasks will wait at most as long as it takes for the queue length to return to $l=0$. Therefore, the waiting time distribution will have as upper bound the return time distribution of a one-dimensional random walk. It is known, however, that the return time distribution of a random walker follows $P(t) \sim t^{-3/2}$ [40,41], which is the origin of the $3/2$ exponent in Eq. (3). This argument indicates that (4) is related to the fluctuations in the length of the priority list.

Supercritical regime, $\rho > 1$: Given that in this regime the




FIG. 5. (Color online) Waiting time distribution for tasks in the queuing model discussed in Sec. V with continuous priorities. The numerical simulations were performed as follows: At each step we generate an arrival τ_a and service time τ_s from an exponential distribution with rate λ and μ , respectively. If $\tau_a < \tau_s$ or there are no tasks in the queue then we add a new task to the queue, with a priority $x \in [0, 1]$ from uniform distribution, and update the time $t \rightarrow t + \tau_a$. Otherwise, we remove from the queue the task with the largest priority and update the time $t \rightarrow t + \tau_s$. The waiting time distribution is plotted for three $\rho = \lambda/\mu$ values: $\rho = 0.9$ (circles), $\rho = 0.99$ (squares), and $\rho = 0.999$ (diamonds). The data has been rescaled to emphasize the scaling behavior $P(\tau_w) = \tau_w^{-3/2} f(\tau_w/\tau_0)$, where $\tau_0 \sim (1 - \sqrt{\rho})^{-2}$. In the inset we plot the distribution of waiting times for $\rho = 1.1$, after collecting up to 10^4 (plus) and 10^5 (diamonds) *executed* tasks, showing that the distribution of waiting times has a power law tail even for $\rho > 1$ (supercritical regime). Note, however, that in this regime a high fraction of tasks are never executed, staying forever on the priority list whose length increases linearly with time, a fact that is manifested by a shift to the right of the cutoff of the waiting time distribution.

arrival rate exceeds the response rate, the average queue length grows linearly as $\langle l(t) \rangle = (\lambda - \mu)t$. Therefore, a $1 - 1/\rho$ fraction of the letters is never responded to, waiting indefinitely in the queue. Given that the measured response rates are 0.32 (Darwin), 0.24 (Einstein), and 0.31 (Freud) [6], this regime captures best the studied correspondence patterns. We can measure the waiting time for each letter that is responded to. In Fig. 5 we show the waiting time probability density obtained from numerical simulations, indicating that it follows a power law with exponent $\alpha=3/2$. Thus the supercritical regime follows the same scaling behavior as the critical regime, but only for the letters that are responded to. The rest of the letters wait indefinitely in the list ($\tau_w = \infty$).

While the discussed model can indeed generate power law waiting time distributions, a critical comparison with the empirical datasets reveals some notable deficiencies. First, a power law distribution emerges only in the critical ($\rho = 1$) and the supercritical ($\rho > 1$) regimes. The critical regime requires a careful tuning of the human execution rate, so that the execution and the arrival rates are exactly the same. In contrast, for $\rho > 1$ no tuning is necessary, but the number of tasks on the list increases linearly with time, thus many tasks are never executed. This limit is probably the most realistic for human dynamics: we often take on tasks that we never execute, and technically stay on our priority list forever. As

we discussed above, this is documentedly the case for Einstein, Darwin, and Freud, who answer only a fraction of their letters. However, we must not overlook the second important feature of the discussed model: the only exponent it can predict is $\alpha=3/2$, rooted in the fluctuations of the queue length. While this agrees with the correspondence patterns of Einstein, Darwin, and Freud, it is significantly higher than the values observed for web browsing, email communications or library visits, which we found α to be scattered around 1.

VI. MODELS WITH FIXED QUEUE LENGTH: $\alpha=1$ UNIVERSALITY CLASS

To understand the limitations of the model discussed in the preceding section we must remember that when the arrival and execution rates are equal ($\rho=1$) the length of the priority list follows a random walk, and can thus occasionally take up very large values. The situation is even worse for $\rho>1$, when the queue length increases linearly with time. Therefore, according to the model an individual must have the capacity to keep track of hundreds or thousands of tasks at the same time. This may be appropriate for surface mail, where the letters pile on our desk until replied to. In contrast, there is extensive evidence from the psychology literature that the number of tasks humans can easily keep in their short term memory is bounded [36], therefore it is unrealistic that we will remember hundreds or thousands of tasks at any given time. This forces us to inspect a model in which the length of the priority list remains unchanged [5], a new task being added only when an old task is removed from the list (executed).

We assume that an individual maintains a priority list with L tasks, each task being assigned a priority parameter $x_i, i=1, \dots, L$, chosen from an $\eta(x)$ distribution. At each time step with probability p the individual selects the highest priority task and executes it, removing it from the list. At that moment a new task is added to the list, its priority x_i is again chosen from $\eta(x)$, thus the length L of the list remains unchanged. With probability $1-p$ the individual executes a randomly selected task, independent of its priority, being replaced again with a new task. The $p \rightarrow 1$ limit of the model describes the deterministic highest-priority-first protocol, when always the highest priority task is chosen for execution, while $p \rightarrow 0$ corresponds to the random choice protocol, introduced to mimic the fact that humans occasionally select some low priority items for execution, before all higher priority items are executed. If the model time is discrete, each task execution corresponds to one unit of time. Implicit in this assumption is the approximation that the service time distribution follows a delta function, i.e., each task takes one unit time to execute.

To understand the dynamics of the model we first study it via numerical simulations with priorities chosen from a uniform distribution $x_i \in [0, 1]$. The simulations show that in the $p \rightarrow 1$ limit the probability that a task spends τ_w time on the list has a power law tail with exponent $\alpha=1$ [Fig. 6]. In the $p \rightarrow 0$ limit $P(\tau_w)$ follows an exponential distribution [Fig. 6(a)], as expected for the random selection protocol. As the typical length of the priority list differs from individual to




FIG. 6. (Color online) Waiting time probability distribution function for the model discussed in Sec. VI for $L=2$ and a uniform new task priority distribution function, $\eta(x)=1$, in $0 \leq x \leq 1$, as obtained from (6) (lines) and numerical simulations (symbols), for $p=0.9$ (squares), $p=0.99$ (diamonds), and $p=0.999$ (triangles). The inset shows the fraction of tasks with waiting time $\tau_w=1$, as obtained from (6) (lines) and numerical simulations (symbols).

individual, it is important for the tail of $P(\tau_w)$ to be independent of L . Numerical simulations indicate that this is indeed the case: changes in L do not affect the scaling of $P(\tau_w)$ [5]. The fact that the scaling holds for $L=2$ as well indicates that it is not necessary to have a long priority list: even if an individual balances only two tasks at the same time, a bursty heavy tailed interevent dynamics will emerge. Next we focus on the $L=2$ case, for which the model can be solved exactly, providing important insights into its scaling behavior that can be generalized for arbitrary L values as well.

A. Exact solution for $L=2$

For $L=2$ the waiting time distribution can be determined exactly [8], obtaining

$$P(\tau_w) = \begin{cases} 1 - \frac{1-p^2}{4p} \ln \frac{1+p}{1-p}, & \tau_w = 1, \\ \frac{1-p^2}{4p(\tau_w-1)} \left[\left(\frac{1+p}{2} \right)^{\tau_w-1} - \left(\frac{1-p}{2} \right)^{\tau_w-1} \right], & \tau_w > 1 \end{cases} \quad (6)$$

independent of $\eta(x)$ from which the task priorities are selected. In the limit $p \rightarrow 0$ from (6) follows that

$$\lim_{p \rightarrow 0} P(\tau_w) = \left(\frac{1}{2} \right)^{-\tau_w}, \quad (7)$$

i.e., $P(\tau_w)$ decays exponentially, in agreement with the numerical results [Fig. 6(a)]. This limit corresponds to the random selection protocol, where a task is selected with probability $1/2$ in each step. In the $p \rightarrow 1$ limit we obtain

$$\lim_{p \rightarrow 1} P(\tau_w) = \begin{cases} 1 + O\left(\frac{1-p}{2} \ln(1-p)\right), & \tau_w = 1, \\ O\left(\frac{1-p}{2}\right) \frac{1}{\tau_w - 1}, & 1 < \tau_w \ll \tau_0, \end{cases} \quad (8)$$

where

$$\tau_0 = \left(\ln \frac{2}{1+p}\right)^{-1}. \quad (9)$$

In this case almost all tasks have a waiting time $\tau_w = 1$, being executed as soon as they were added to the priority list. The waiting time of tasks that are not selected in the first step follows a power law distribution, decaying with $\alpha = 1$. This behavior is illustrated in Fig. 6(a) by a direct plot of $P(\tau_w)$ in (6) for a uniform distribution $\eta(x)$ in $0 \leq x \leq 1$. For $p < 1$ the $P(\tau_w)$ distribution has an exponential cutoff, which can be derived from (6) after taking the $\tau_w \rightarrow \infty$ limit with p fixed, resulting in

$$P(\tau_w) \sim \frac{1-p^2}{4} \frac{1}{\tau_w} \exp\left(-\frac{\tau_w}{\tau_0}\right). \quad (10)$$

When $p \rightarrow 1$ we obtain that $\tau_0 \rightarrow \infty$ and, therefore, the exponential cutoff is shifted to higher τ_w values, while the power law behavior $P(\tau_w) \sim 1/\tau_w$ becomes more prominent. The $P(\tau_w)$ curve systematically shifts, however, to lower values for $\tau_w > 1$, indicating that the power law applies to a vanishing task fraction [see Fig. 6(a) and (10)]. In turn, $P(1) \rightarrow 1$ when $p \rightarrow 1$, corroborated by the direct plot of $P(1)$ as a function of p [see inset of Fig. 6(a)].

B. Numerical results for $L > 2$

Based on the results discussed above, the overall behavior of the model with a uniform priority distribution can be summarized as follows. For $p = 1$, corresponding to the case when *always* the highest priority task is removed, the model does not have a stationary state. Indeed, each time the highest priority task is executed, there is a task with smaller priority x_m left on the list. With probability $1 - x_m$ the newly added task will have a priority x'_m larger than x_m , and will be executed immediately. With probability x_m , however, the new task will have a smaller priority, in which case the older task will be executed, and the new task will become the resident one, with a smaller priority $x'_m < x_m$. For a long period all new tasks will be executed right away, until another task arrives with probability x''_m that again pushes the nonexecuted priority to a smaller value $x''_m < x'_m$. Thus with time the priority of the lowest priority task will converge to zero, $x_m(t) \rightarrow 0$, and thus with a probability converging to one the new task will be immediately executed. This convergence of x_m to zero implies that for $p = 1$ the model does not have a stationary state. A stationary state develops, however, for any $p < 1$, as in this case there is always a finite chance that the lowest priority tasks will also be executed, thus the value of x_m will be reset, and will converge to some $x_m(p) > 0$. This

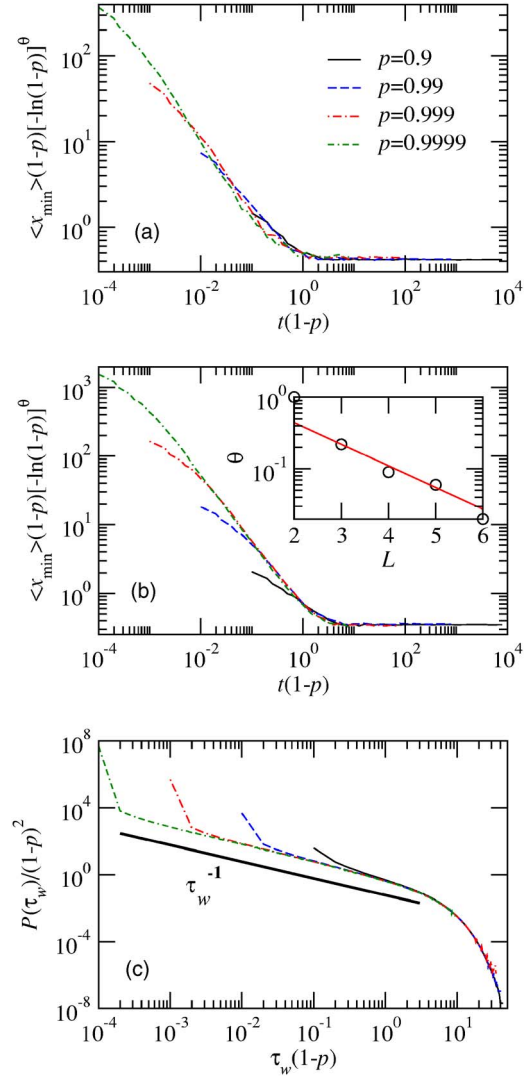


FIG. 7. (Color online) Rescaled plot of the average priority of the lowest task priority in the list for $L=2$ (a) and $L=3$ (b) and different values of p (see legend). The inset in (b) shows the exponent θ_L for different L (points), indicating that $\theta_L = \theta_3/2^{L-3}$ for $L > 2$ (continuous line). (c) Rescaled plot of the waiting time distribution for $L=3$. Similar plots are obtained for larger values of L (data not shown).

qualitative description applies for arbitrary $L > 2$ values.

To quantify this qualitative picture we studied numerically the $L > 2$ case assuming that $\eta(x)$ is uniformly distributed in the $0 \leq x \leq 1$ interval. To investigate how fast the system approaches the stationary state we compute the average priority of the lowest priority task in the queue, $\langle x_{\min}(t) \rangle$ [see Figs. 7(a) and 7(b)] since it represents a lower bound for the average of any other priorities on the list. We find that for any L values $\langle x_{\min}(t) \rangle$ decreases exponentially up to a time scale t_0 , when it reaches a stationary value $\langle x_{\min}(\infty) \rangle$. The numerical simulations indicate that

$$t_0 \sim \frac{1}{1-p}, \quad (11)$$

$$\langle x_{\min}(\infty) \rangle \sim (1-p)[- \ln(1-p)]^{\theta_L}. \quad (12)$$

For $L=2$ we can calculate $\langle x_{\min}(\infty) \rangle$ exactly [8], obtaining

$$\langle x_{\min}(\infty) \rangle = \frac{1-p}{2p} \left(\frac{1+p}{2p} \ln \frac{1+p}{1-p} - 1 \right) \approx \frac{1-p}{2} [- \ln(1-p)], \quad (13)$$

and therefore $\theta_2=1$. For $L>2$ we determined θ_L from the best data collapse, obtaining the values shown in the inset of Fig. 7(b), indicating that

$$\theta_L = \frac{\theta_3}{2^{L-3}},$$

where $\theta_3=0.22$ is the value of θ_L for $L=3$. These results support our qualitative discussion, indicating that for all $L \geq 2$ and $0 \leq p < 1$ values the system reaches a stationary state.

Finally we measured the waiting time distribution after the system has reached the stationary state. The results for $L=3$ are shown in Fig. 7(c), and similar results were obtained for other $L>2$ values. The data collapse of the numerically obtained $P(\tau)$ indicates that

$$P(\tau) \sim (1-p)^2 \frac{1}{\tau} \exp\left(-\frac{\tau}{\tau_0}\right), \quad (14)$$

when $L>2$ and $\tau \gg 1$, where

$$\tau_0 \sim \frac{1}{1-p} \quad (15)$$

in the $p \rightarrow 1$ limit. The simulations indicate that the model's behavior for $L>2$ is qualitatively similar to the behavior derived exactly for $L=2$, but different scaling parameters characterize the scaling functions. For any $L \geq 2$ and $1 \ll \tau \ll \tau_0$, however, the waiting times scale as $P(\tau_w) \sim \tau_w^{-1}$, i.e., we have $\alpha=1$.

C. Comparison with the empirical data

As we showed above, for truly deterministic execution ($p=1$) the model does not have a stationary state. The problem was cured by introducing a random task execution ($p < 1$), which leads to stationarity. In this case, however, a p dependent fraction of tasks are executed immediately, and only the rest of the long lived tasks follow a power law. As p converges to one, the fraction of tasks executed immediately diverges, developing a significant gap between the power law regime, and the tasks displaying $\tau=1$ waiting time. Is this behavior realistic, or represents an artifact of the model? A first comparison with the empirical data would suggest that this is an artifact, as measurements shown in Figs. 2 and 3(a) do not provide evidence of a large number of tasks that are immediately executed. However, we should keep in mind that we are measuring the intervent times, and not the waiting times. When the waiting time can be directly measured, like in the email or mail based correspondence, there is some ambiguity to the real waiting time. Indeed, in the email data, for example, the waiting time was measured as the time dif-

ference between the arrival of an email and the response sent to it. While this offers a reasonable approximation, from the user's priority queue's perspective this is not the real waiting time. Indeed, consider the situation when an email arrives at 9:00 am, and the recipient does not check her email until 11:56 am, at which point she replies to the email at 11:58 am. From the perspective of her priority list the waiting time was 2 minutes, as she replied as soon as she saw the email. In our dataset, however, the waiting time is 2 hours and 58 minutes. Thus the way we measured the waiting times cannot identify the true waiting time of a task on a user's priority list. The email dataset allows us, however, to get a much better approximation of the real waiting times than we did before. Indeed, for an email e_1 received by user A we record the time t_1 it arrives, and then the time t_2 of the first email sent by user A to any other user *after* the arrival of the selected email. We start measuring the waiting time for email e_1 at t_2 . Thus if user A replies to e_1 at time t_3 , we consider that the email's waiting time $\tau_{\text{real}}=t_3-t_2$, instead of t_3-t_1 considered in Fig. 3(a). The results, shown in Figs. 3(c) and 3(d) for two users, displays the same power law scaling with $\alpha=1$ as we have seen in Fig. 3(a), but in addition there is a prominent peak at $\tau_{\text{real}}=1$, corresponding to emails responded to immediately. Note that the peak's magnitude is orders of magnitude larger than the probabilities displayed by the large waiting times. Most important, we find that the exact solution (8) offers a perfect fit to the whole dataset, capturing the large peak at $\tau_{\text{real}}=1$, the $\alpha=1$ scaling regime and the exponential cutoff for high τ_{real} . Given the many factors that contribute to human activity patterns and the simplicity of the model, the quality of the fit is exceptional. These results suggest that what we could have easily considered a model artifact in fact captures a common feature of email communications. Indeed, a high fraction of the emails we receive are responded to immediately, right after our first chance to read them, as predicted by the priority model discussed in this section.

Identifying fat tailed distributions has always been an issue of much discussion in the mathematical, engineering and physics literature. Indeed, one issue of major importance, occurring in many different fields, is what is the best fit to an apparently fat tailed distribution. The answer is complicated by the fact that for many real finite datasets often a power law, log-normal, and stretched exponential distributions offer comparable fits [43]. Properly distinguishing them requires very extensive datasets, that often push the boundaries of the data collection capabilities. Given the limited datasets a log-normal distribution appears to also offer a reasonable fit of $P(\tau)$ [44]. The reason is that both the log-normal and the power law distributions have the same leading behavior τ^{-1} , differing only in the functional form of the exponential correction: for a log-normal distribution the correction has the form $e^{-[\ln(\tau/\tau_0)]^2/2\sigma^2}$ while the exact solution for $L=2$ predicts $e^{-\tau/\tau_0}$ (7). Strong support for the power law nature of the distribution comes from considering the aggregated distribution of all interevent times observed in the email system (i.e., collecting all interevent times for all users). Note, that the aggregated distribution collapses together users of rather different activity pattern. Yet, as Fig. 3(g) shows, this distribution follows a power law for over five orders of magnitude, followed by an exponential cutoff at $\tau_0=12$ days. As Fig.

3(h) shows, a log-normal distribution fails to provide a reasonable fit to this dataset. In summary, given the relatively small amounts of data, it is hard to definitively distinguish the relative validity of the two distributions at this point for single users. Yet, the aggregated data, which is sufficient in size, clearly rules out the log-normal, and offers strong support for the power law with exponential cutoff.

At this stage, however, the most important drawback of this fitting exercise is the lack of theoretical understanding that a log-normal distribution could emerge in a reasonable model of human dynamics. Indeed, the power law fit shown in Figs. 3(c) and 3(d) is not a mere fit, but represents the prediction of a simple model of human dynamics, and captures all major features of the empirical dataset, from the small τ peak to the power law regime and the exponential cutoff.

Therefore, in the absence of a mechanism indicating that a log-normal distribution could also emerge in a model with relevance to the discussed human communication patterns, the observation that a log-normal distribution offers an equally good fit for some users is a mere exercise in statistics, one that has little hope to be conclusive until considerably longer records of email patterns become available, and does not provide much valuable insight in human dynamics.

VII. RELATIONSHIP BETWEEN WAITING TIMES AND INTEREVENT TIMES

As we discussed above, the empirical measurements provide either the interevent time distribution $P(\tau)$ (Secs. III A and III C) or the waiting time distribution $P(\tau_w)$ (Sec. III B) of the measured human activity patterns. In contrast the model predicts only the waiting time τ_w of a task on an individual's priority list. What is the relationship between the observed interevent times and the predicted waiting times? The basic thesis of our paper is that the waiting times the various tasks experience on an individual's priority list is responsible for the heavy tailed distributions seen in the interevent times as well. The purpose of this section is to discuss the relationship between the two quantities.

The model predictions, that the waiting time distribution of the tasks follows a power law, is directly supported by one dataset in each universality class: the email data and the correspondence data. As discussed in Sec. III, we have measured the response time distribution in both datasets, finding that the distribution of the response times indeed follows a power law with exponent $\alpha=1$ (email) and $\alpha=3/2$ (correspondence mail) as predicted by the models. For the other three datasets, however, such as web browsing, library visits, and stock purchases, we cannot determine the waiting time of the individual events, as we do not know when a given task is added to the individual's priority list.

To explore the broader relationship between the response times and the interevent times we must remind ourselves that while during the measurements we monitor a specific task (like email), the models assume the knowledge of *all* tasks that an individual is involved in. Thus the empirical measurements offer only a selected subset of an individual's activity pattern. To see the relationship between τ and τ_w next we discuss two different approaches.

Queuing of different task categories: The first approach acknowledges the fact that tasks are grouped in different categories of priorities: we often do not keep in mind specific emails to be answered, but rather remember that we need to check our email and answer whatever needs attention. Similarly, we may remember a few things that we need to shop for, but our priority list would often contain only one item: go to the supermarket. When we monitor different human activity patterns, we see the repetitive execution of these categories, like going to the library, or doing emails, or browsing the web. Given this, one possible generalization of the discussed models would assume that the tasks we monitor correspond to specific activity categories, and when we are done with one of them, we do not remove it from the list, but we just add it back with some changed priority [45]. That is, checking our email does not mean that we deleted email activity from our priority list, but only that next has some different priority. If we monitor only one kind of activity, then a proper model would be the following: we have L tasks, each assigned a given priority. After a task is executed, it will be reinserted in the queue with a new priority chosen from the same distribution $\eta(x)$. If we now monitor the time at which the different tasks exit the list, we will find that the interevent times for the *monitored* tasks correspond exactly to the waiting time of that task on the list. Note that this conceptual model would work even if the tasks are not immediately reinserted, but after some delay τ_d . Indeed in this case the interevent time will be $\tau = \tau_w + \tau_d$, and as long as the distribution from which τ_d is selected from is bounded, the tail of the interevent time distribution will be dominated by the waiting time statistics.

Interaction between individuals: The timing of specific emails also depends on the interaction between the individuals that are involved in an email based communication. Indeed, if user A receives an email from user B, she will put the email into her priority list, and answer when she gets to it. Thus the timing of the response depends on two parameters: the receipt time of the email, and the waiting time on the priority list. Consider two email users, A and B, that are involved in an email based conversation. We assume that A sends an email to B as a response to an email B sent to A, and vice versa. Thus, the interevent time between two consecutive emails sent by user A to user B is given by $\tau = \tau_w^A + \tau_w^B$, where τ_w^A is the waiting time the email experienced on user A's queue, and τ_w^B is the waiting time of the response of user B to A's email. If both users prioritize their tasks, then they both display the same waiting time distribution, i.e., $P(\tau_w^A) \sim (\tau_w^A)^{-\alpha}$ and $P(\tau_w^B) \sim (\tau_w^B)^{-\alpha}$. In this case the interevent time distribution $P(\tau)$, which is observed empirically if we study only the activity pattern of user A, follows also $P(\tau) \sim \tau^{-\alpha}$. Thus the fact that users communicate with each other turns the waiting time (or response) into an observable interevent times.

In summary, the discussed mechanisms indicate that the waiting time distribution of the tasks could in fact drive the interevent time distribution, and that the waiting time and the interevent time distributions should decay with the same scaling exponent. In reality, of course, the interplay between the two quantities can be more complex than discussed here,

and perhaps even better mapping between them could be found for selected activities. But the discussed mechanisms indicate that if the waiting time distribution is heavy tailed, we would expect that the interevent time distribution would be also affected by it.

VIII. DISCUSSION

Universality classes: As summarized in the introduction, the main goal of the present paper was to discuss the potential origin of the heavy tailed distributed interevent times observed in human dynamics. To start we provided evidence that in five distinct processes, each capturing a different human activity, the interevent time distribution for individual users follow a power law. Our fundamental hypothesis is that the observed interevent time distributions are rooted in the mechanisms that humans use to decide when to execute the tasks on their priority list. To support this hypothesis we studied a family of queuing models, that assume that each task executed by an individual waits some time on the individual's priority list and we showed that queuing can indeed generate power law waiting time distributions. We find that a model that allows the queue length to fluctuate leads to $\alpha = 3/2$, while a model for which the queue length is fixed displays $\alpha = 1$. These results indicate that human dynamics is described by at least two universality classes, characterized by empirically distinguishable exponents. Note that while we have classified the models based on the limitations on the queue length, we cannot exclude the existence of models with fixed queue length that scale with $\alpha = 3/2$, or models with fluctuating length that display scaling with $\alpha = 1$, or some other exponents.

In comparing these results with the empirical data, we find that email and phone communication, web surfing and library visitation belong to the $\alpha = 1$ universality class. The correspondence patterns of Einstein, Darwin, and Freud offer convincing evidence for the relevance of the $\alpha = 3/2$ exponent, and the related universality class, for human dynamics. In contrast the fourth process, capturing a stock broker's activity, shows $\alpha = 1.3$. Given, however, that we have data only for a single user, this value is in principle consistent with the scattering of the exponents from user to user, thus we cannot take it as evidence for a new universality class. One issue still remains without a satisfactory answer: why does email and surface mail (Einstein, Darwin, and Freud datasets) belong to different universality classes? We can comprehend why should the mail correspondence belong to the $3/2$ class: letters likely pile on the correspondent's desk until they are answered, the desk serving as an external memory, thus we do not require to remember them all. But the same argument could be used to explain the scaling of email communications as well, given that unanswered emails will stay in our mailbox until we delete them (which is one kind of task execution). Therefore one could argue that email based communication should also belong to the $3/2$ universality class, in contrast with the empirical evidence, that clearly shows $\alpha = 1$ [5,25].

Some difficulty in comparing the empirical data with the model predictions is rooted in the fact that the models predict

the waiting times, while for many real systems only the interevent times can be measured. It is encouraging, however, that for the email and the surface mail based communication we were able to determine directly the waiting times as well, and the exponents agreed with those determined from the interevent times. In addition we argued that in a series of processes the waiting time distribution determines the interevent time distribution as well (see Sec. VII). This argument closes the loop of the paper's logic, establishing the relevance of the discussed queuing models to the datasets for which only interevent times could be measured. We do not feel, however, that this argument is complete, and probably future work will strengthen this link. In this respect two directions are particularly promising. First, designing queuing models that can directly predict the observed interevent times as well would be a major advance. Second, establishing a more general link between the waiting and interevent times could also be of significant value.

The results discussed in this paper leave a number of issues unresolved. In the following we will discuss some of these, outlining how answers to them could further our understanding of the statistical mechanics of human driven processes.

Tuning the universality class: As we discussed above, the discussed models provide evidence for two distinct universality classes in human dynamics, with distinguishable exponents. The question is, are there other universality classes, characterized by exponents different from 1 and $3/2$? If other universality classes do exist, it would be valuable not only to find empirical support for them, but also to identify classes of models that are capable of predicting the new exponents.

In searching for new exponents we need to explore several different directions. First, if one inserts some power law process into the queuing model, that could tune the obtained waiting time distribution, and the scaling exponents. There are different ways to achieve this. One method, discussed in Ref. [5], is based on the hypothesis that while we always attempt to select the highest priority task, circumstances or resource availability may not allow us to achieve this. For example, our highest priority may be to get cash from the bank, but we cannot execute this task when the bank is closed, moving on to some lower priority task. One way to account for this is to use a probabilistic selection protocol, assuming that the probability to choose a task with priority x for execution in a unit time is $\Pi(x) \sim x^\gamma$, where γ is a parameter that interpolates between the random choice limit (ii) ($\gamma = 0$, $p = 0$) and the deterministic case, when always the highest priority item is chosen for execution (iii) ($\gamma = \infty$, $p = 1$). As shown in Appendix C, the exponent α will depend on γ as $\alpha = 1 + 1/\gamma$. At this moment we do not have evidence that such preferential selection process acts in human dynamics. However, detailed datasets and proper measurement tools might help us decide this by measuring the function $\Pi(x)$ directly, capturing the selection protocol. Such measurements were possible for complex networks, where a similar function drives the preferential attachment process [10,46–51].

In this paper we focused on models with Poisson inputs, meaning that both the arrival time and the execution time are

bounded. In some situations, however, the input distributions can be themselves heavy tailed. This could have two origins: (i) Heavy tailed arrival time distribution. As we show in Fig. 3(b), there is direct evidence for this in the email communication datasets, where we find that the interevent time of arriving emails can be roughly approximated with a power law with exponent $\alpha_{in}=1$. (ii) The execution time could also be heavy tailed, describing the situation when most tasks are executed very rapidly, while a few tasks require a very long execution time. For example, the file sizes transmitted by email are known to follow a heavy tailed distribution [34,35]. Therefore, if we read every line of an email, in principle the execution time should also be heavy tailed (i.e., the time we actually take to work on the response, including reading the original email). Note, however, that measurements failed to establish a correlation between email size and the response time [5]. It is not particularly surprising that both (i) and (ii) would significantly impact the waiting time distribution, generating a heavy tailed distribution for the waiting times even when the behavior of the model otherwise would be exponential, or change the exponent α . Some aspects of this problem were addressed recently by Blanchard and Hongler [52]. However, to illustrate the impact of the heavy tailed inputs in Appendix B we study the model of Sec. V with a heavy tailed service time distribution $h(s) \sim s^{-\beta}$ with $0 < \beta < 1$.

Finally, could the power law distributed arrival and execution times serve as the proper explanation for the observed heavy tailed interevent time distribution in human dynamics? Note that in a number of systems we observe heavy tailed distributed events without evidence for power law inputs. For example, the timing of the library visits or stock purchases by brokers does not appear to be driven by any known power law inputs, and they have negligible execution time compared with the average observed interevent times. Similarly, the beginning of online games or instant messages is not driven by file sizes either, but only by the time availability for playing a game or sending a message, which is mostly a priority driven issue. Therefore, while it is important to understand the impact of power law inputs on the scaling properties of various models, attempts to explain the waiting times solely based on the heavy tailed inputs only delegate the problem to an earlier cause (the origin of the power law input).

Potential model extensions: Guided by the desire of constructing the simplest models that capture the essence of task execution, we have neglected many processes that are obviously present in human dynamics. For example, we assumed that the priority of the tasks is assigned at the moment the task was added to an individual's priority list, and remains unchanged for the rest of the queuing process. In reality the priorities themselves can change in time. For example, many tasks have deadlines [52], and one could assume that a task's priority diverges as the deadline approaches. Even in the absence of a clear deadline some priorities may increase in time [52], others may decrease. Sometimes external factors change suddenly a task's priority—for example, the priority of watering the lawn suddenly diminishes if it starts raining. The possibility of dropping tasks, either by not allowing them on the queue, or by simply deleting them from the

queue, could also affect the waiting time distributions. Tasks could be dropped if they were not executed for a considerable time interval, and thus become irrelevant, or when the individual is very busy, or some may be simply forgotten. Obviously, the precise impact on the waiting time distribution will depend on the implementation of the task dropping conditions. It is important to understand if any or all of these processes could change the universality class of the waiting time distribution.

Model limitations: The studied datasets do not capture all tasks an individual is involved in, but only the timing of selected activities, like sending emails or borrowing books from the library. Yet, we must consider the fact that between any two recorded events individuals participate in many other nonrecorded activities. For example, if we find that an individual clicks on a new document every few seconds, likely he/she is fully concentrating on web browsing. However, when we notice a break of hours or days between two consecutive clicks, it is clear that in the meantime the individual was involved in a myriad of other activities that were not visible to us. The queuing models discussed here were designed to take into consideration all human activities, as we assume that the priority list of an individual contains all tasks the person is involved in. Currently an understanding of the interplay between the *recorded* and the *invisible* activities is still lacking.

Seasonality and fluctuations: Human activities are affected by circadian and seasonal patterns, such as the 24 hour succession of sleeping and working hours as well as the 7 day working week induced periodicity. This raises an important question: could these seasonal patterns be responsible in some way for the observed heavy tailed distributions [42]? Recently Hidalgo offered a general argument, indicating that fluctuations in the execution rate, and in particular periodic changes, can lead to a power law interevent time distribution decaying with exponent $\alpha=2$ [42]. This not only offers evidence of a novel universality class, but opens new and promising avenues for research addressing the role of fluctuations in human dynamics.

Task optimization: The order in which we execute different tasks is often driven by optimization: we try to minimize the total time, or some cost functions. This is particularly relevant if the execution times depend on the order in which the tasks are executed. For example, often executing a certain task might be faster if we execute some other preparatory tasks before, and not in the inverse order. In principle optimization could be incorporated into the studied models by assuming that they determine the priority of the tasks. Optimization raises several important questions for future work: How should we model optimization driven queuing processes? Can they also lead to power laws, and if so, will they result in new universality classes?

Correlations: So far we have focused on the origin of the various distributions observed in human dynamics. Distributions offer little information, however, about potential correlations present in the observed time series. Such correlations were documented in Ref. [27], observing that the correlation function of the interevent time series for printing job arrivals decayed as a power law. Are such temporal correlations present in other systems as well? What is their origin? Can

the queuing models predict such correlations? Answers to these questions could not only help better understand human dynamics, but could also aid in distinguishing the various models from each other.

Network effects: In searching for the explanation for the observed heavy tailed human activity patterns we limited our study to the properties of single queues. In reality none of our actions are performed independently—most of our daily activity is embedded in a web of actions of other individuals [53,54]. Indeed, the timing of an email sent to user A may depend on the time we receive an email from user B. An important future goal is to understand how the various human activities and their timing is affected by the fact that the individuals are embedded in a network environment.

Nonhuman activity patterns: Heavy tailed interevent time distributions do not occur only in human activity, but emerge in many natural and technological systems. For example, earthquakes records reveal heavy tailed interevent times between consecutive seismic activities [55–57]; measurements indicate that the fishing patterns of seabirds also display heavy tailed statistics [58]; plasticity patterns [59], avalanches in lungs [60] and volatility return intervals in financial markets [61,62] show similar power law interevent times. While a series of models have been proposed to capture some of these processes individually, there is also a possibility that some of these modeling frameworks can be reduced to various queuing processes. Some of the studied queuing models show a close relationship to several models designed to capture self-organized criticality [63–68]. Could the mechanisms be similar at some fundamental level? Even if such higher degree of universality is absent, understanding the mechanisms and queuing processes that drive human dynamics could help us better understand other natural phenomena as well, from the timing of chemical reactions in a cell to the temporal records of major economic events [53] or the timing of events in manufacturing processes and supply chains [69–71] or panic [72].

ACKNOWLEDGMENTS

The authors wish to thank Diana Kormos Buchwald and Tilman Sauer for providing us the dataset capturing the Einstein correspondence, and for helping us understand many aspects of Einstein’s life and correspondence patterns. Similarly, the authors wish to thank Alison M. Pearn from the Darwin Correspondence Project for providing the record of Darwin’s communications and Thurston Miller for providing the library visit dataset. Three of the authors (A.-L.B., A.V., and Z.D.) are supported by NSF ITR 0426737, NSF ACT/SGER 0441089 awards and by an award from the James S. McDonnell Foundation. One of the authors (J.G.O.) is supported by FCT (Portugal) Grant No. SFRH/BD/14168/2003. One of the authors (I.K.) thanks B. Jankó and the Institute for Theoretical Sciences at University of Notre Dame for their hospitality during this collaboration, as well as the National Office of Research and Technology in Hungary for support.

APPENDIX A: TRANSITIONS BETWEEN THE TWO UNIVERSALITY CLASSES

A basic difference between the models discussed in Sec. V and Sec. VI is the capacity of the queue. Our results indi-

cate that the model without limitation on the queue length displays $\alpha=3/2$, rooted in the fluctuations of the queue length. In contrast, the model with fixed queue length (Sec. VI) has $\alpha=1$, rooted in the queuing of the low priority tasks on the priority list. If indeed the limitation in the queue length plays an important role, we should be able to develop a model that can display a transition from the $\alpha=3/2$ to the $\alpha=1$ universality class as we limit the fluctuations in the queue length. In this section we study such a model, interpolating between the two observed scaling regimes. We start from the model discussed in Sec. V, and impose on it a maximum queue length L . This can be achieved by altering the arrival rate of the tasks: when there are L tasks in the queue no new tasks will be accepted until at least one of the tasks is executed. Mathematically this implies that the arrival rate depends on the queue length as

$$\lambda_\ell = \begin{cases} \lambda, & 0 \leq \ell < L, \\ 0, & \ell = L. \end{cases} \quad (\text{A1})$$

In the stationary state the queue length distribution $P(\ell)$ satisfies the balance equation

$$\lambda_{\ell-1}P(\ell-1) + \mu_{\ell+1}P(\ell+1) = (\lambda_\ell + \mu_\ell)P(\ell), \quad (\text{A2})$$

where

$$\mu_\ell = \begin{cases} 0, & \ell = 0, \\ \mu, & 0 < \ell \leq L. \end{cases} \quad (\text{A3})$$

From (A2) we obtain the queue length distribution as

$$P(\ell) = \frac{1-\rho}{1-\rho^{L+1}}\rho^\ell, \quad (\text{A4})$$

suggesting the existence of three scaling regions.

Subcritical regime, $\rho \ll 1$: If the arrival rate of the tasks is much smaller than the execution rate, the fact that the queue length has an upper bound has little significance, since ℓ will rarely reach its upper bound L , but will fluctuate in the vicinity of $\ell=0$. This regime can be reached either for $\rho \ll 1$ and L fixed or for $\rho < 1$ and $L \gg 1$. Therefore, in this case the waiting time distribution is well approximated by that of the model with an unlimited queue length, displaying the scaling predicted by Eq. (4), i.e., either exponential, or a power law with $\alpha=3/2$, coupled with an exponential cutoff [see Fig. 8(a)].

Critical regime: For $\rho=1$ we observe an interesting interplay between the queue length and L . Normally in this critical regime $\ell(t)$ should follow a random walk with the return time probability density scaling with exponent $3/2$. However, the limitation imposed on the queue length limits the power law waiting time distribution predicted by Eq. (4), introducing a cutoff [see Fig. 8(b)]. Indeed having the number of tasks in the queue limited allows each task to be executed in a finite time.

Supercritical regime: When $\rho \gg 1$ from (A4) follows that




FIG. 8. (Color online) Waiting time distribution for tasks in the queuing model discussed in Appendix A, with a maximum queue length L . The waiting time distribution is plotted for three L values: $L=10$ (circles), $L=100$ (squares), and $L=1000$ (diamonds). The data has been rescaled to emphasize the scaling behavior $P(\tau_w) = \tau_w^{-3/2} f(\tau_w/\tau_0)$, where $\tau_0 \sim L^2$. In the inset we plot the waiting time for $\rho=10^6$, showing the crossover to the model discussed in Sec. VI in the limit $\rho \rightarrow \infty$ and L fixed.

$$\mathcal{L}_\ell = \begin{cases} O(\rho^{-1}), & 0 \leq \ell < L, \\ 1 - O(\rho^{-1}), & \ell = L, \end{cases} \quad (\text{A5})$$

i.e., with probability almost once the queue is filled. Thus, in the supercritical regime $\rho \gg 1$ new tasks are added to the queue immediately after a task is executed. If we take the number of executed tasks as a new reference time then this model corresponds to the one discussed in Sec. VI, displaying $\alpha=1$ [5], as supported by the numerical simulations [see Fig. 8(a)].

APPENDIX B: HEAVY TAILED INPUT DISTRIBUTIONS

In this Appendix we study the model discussed in Sec. V with a heavy tailed service time distribution $h(s) \sim s^{-\beta}$ with $0 < \beta < 1$. In this case it has been shown that [39]

$$P(\tau_w) \sim \tau_w^{-\beta}. \quad (\text{B1})$$

This result is a consequence of the generalized limit theorem for heavy tailed distributions [1]. Let us focus on a selected task and assume that m tasks need to be executed before it. Therefore, the selected task's waiting time is given by

$$\tau_w = \sum_{l=1}^m s_l, \quad (\text{B2})$$

where s_l is the service time of the l th task executed before the given task. Equation (B2) represents the sum of m independent and identically distributed random variables, with pdf $h(s) \sim s^{-\beta}$, which is known to follow a pdf with the same heavy tail, and thus resulting in (B1). Hence, in this case the heavy tail in the waiting time distribution is a consequence of the heavy tails in the service time distribution.

APPENDIX C: PREFERENTIAL SELECTION PROTOCOL

As we discussed in Sec. VIII, one possible modification of the priority model introduced and studied in Sec. VI involves the assumption that we do not always choose the highest priority task for execution, but rather the tasks are chosen stochastically, in increasing function of their priority. That is, the probability to choose a task with priority x for execution in a unit time is $\Pi(x) \sim x^\gamma$, where γ is a predefined parameter of the model. This parameter allows us to interpolate between the random choice limit (ii) ($\gamma=0$, $p=0$) and the deterministic case, when always the highest priority item is chosen for execution (iii) ($\gamma=\infty$, $p=1$). Note that this parametrization captures the scaling of the model discussed in Sec. VI only in the $p \rightarrow 0$ and $p \rightarrow 1$ limits, but not for intermediate p values. That is, the two limits of this model map into extreme limits of the model introduced in Sec. VI, but the intermediate p and γ values do not map into each other.

The probability that a task with priority x waits a time interval t before execution is $f(x, t) = [1 - \Pi(x)]^{t-1} \Pi(x)$. The average waiting time of a task with priority x is obtained by averaging over t weighted with $f(x, t)$, providing

$$t_w(x) = \sum_{t=1}^{\infty} t f(x, t) = \frac{1}{\Pi(x)} \sim \frac{1}{x^\gamma}, \quad (\text{C1})$$

i.e., the higher an item's priority, the shorter is the average time it waits before execution. To calculate $P(\tau)$ we use the fact that the priorities are chosen from the $\eta(x)$ distribution, i.e., $\eta(x)dx = P(\tau)d\tau$, which gives

$$w(t_w) \sim \frac{\eta(t_w^{-1/\gamma})}{t_w^{1+1/\gamma}}, \quad (\text{C2})$$

providing the relationship $\alpha=1+1/\gamma$ between α and γ , and indicating that with changing γ we can continuously tune α as well. In the $\gamma \rightarrow \infty$ limit, which converges to the strictly priority based deterministic choice ($p=1$) in the model, Eq. (C2) predicts $w(t_w) \sim t_w^{-1}$, in agreement with the numerical results [Fig. 3(a)], as well as the empirical data on the email interarrival times [Fig. 2(a)]. In the $\gamma=0$ ($p=0$) limit $t_w(x)$ is independent of x , thus $w(t_w)$ converges to an exponential distribution, as shown in Fig. 3(b).

The apparent dependence of $w(t_w)$ on the $\eta(x)$ distribution from which the agent chooses the priorities may appear to represent a potential problem, as assigning priorities is a subjective process, each individual being characterized by its own $\eta(x)$ distribution. According to Eq. (C2), however, in the $\gamma \rightarrow \infty$ limit $w(t_w)$ is independent of $\eta(x)$. Indeed, in the deterministic limit the uniform $\eta(x)$ can be transformed into an arbitrary $\eta'(x)$ with a parameter change, without altering the order in which the tasks are executed [32]. This insensitivity of the tail to $\eta(x)$ explains why, despite the diversity of human actions, encompassing both professional and personal priorities, most decision driven processes develop a heavy tail.

- [1] W. Feller, *An Introduction to Probability Theory and its Applications*, 2nd ed. (Wiley, New York, 1971), Vol. 2.
- [2] F. A. Haight, *Handbook of the Poisson Distribution* (Wiley, New York, 1967).
- [3] A. K. Erlang, *Nyt. Tidsskrift for Matematik B* **20** (1909); *Elektroteknikeren* **13** (1917).
- [4] H. R. Anderson, *Fixed Broadband Wireless System Design* (Wiley, New York, 2003).
- [5] A.-L. Barabási, *Nature (London)* **207**, 435 (2005).
- [6] J. G. Oliveira and A.-L. Barabási, *Nature (London)* **437**, 1251 (2005).
- [7] Z. Dezső, E. Almaas, A. Lukács, B. Rácz, I. Szakadát, and A.-L. Barabási, e-print physics/0505087.
- [8] A. Vazquez, *Phys. Rev. Lett.* **95**, 248701 (2005).
- [9] H. E. Stanley, *Introduction to Phase Transitions and Critical Phenomena* (Oxford University Press, Oxford, 1987).
- [10] R. Albert and A.-L. Barabási, *Rev. Mod. Phys.* **74**, 47 (2002).
- [11] S. N. Drogovtsev and J. F. F. Mendes, *Evolution of Networks: From Biological Nets to the Internet and the WWW* (Oxford University Press, Oxford, 2003).
- [12] R. Pastor-Satorras and A. Vespignani, *Evolution and Structure of the Internet: A Statistical Physics Approach* (Cambridge University Press, Cambridge, 2004).
- [13] M. E. J. Newman, A.-L. Barabási, and D. J. Watts, *The Structure and Dynamics of Complex Networks* (Princeton University Press, Princeton, NJ, 2006).
- [14] S.-D. Poisson, *Recherches sur la Probabilité des Jugements en Matière Criminelle et en Matière Civile, Précédées des Règles générales du Calcul des Probabilités* (Bachelier, Paris, 1837).
- [15] P. Reynolds, *Call Center Staffing* (The Call Center School Press, Lebanon, TN, 2003).
- [16] J. H. Greene, *Production and Inventory Control Handbook*, 3rd ed. (McGraw-Hill, New York, 1997).
- [17] C. Dewes, A. Wichmann, and A. Feldman, *Proceedings of the 2003 ACM SIGCOMM Conference on Internet Measurement (IMC-03)* (ACM, New York, 2003).
- [18] S. D. Kleban and S. H. Clearwater, *Proceedings of SC'03, 2003, Phoenix, AZ, USA*.
- [19] V. Paxson and S. Floyd, *IEEE/ACM Trans. Netw.* **3**, 226 (1995).
- [20] F. Mainardi, M. Raberto, R. Gorenflo, and E. Scalas, *Physica A* **287**, 468 (2000).
- [21] M. Raberto, E. Scalas, and F. Mainardi, e-print cond-mat/0203596.
- [22] V. Plerou, P. Gopikrishnan, A. N. Amaral, X. Gabaix, and H. E. Stanley, *Phys. Rev. E* **62**, R3023 (2000).
- [23] J. Masoliver, M. Montero, and G. H. Weiss, *Phys. Rev. E* **67**, 021112 (2003).
- [24] T. Henderson and S. Nhatti, *Proceedings of the ACM Multimedia 2001, Ottawa, Canada, 2001*, pp. 212–220.
- [25] J.-P. Eckmann, E. Moses, and D. Sergi, *Proc. Natl. Acad. Sci. U.S.A.* **101**, 14333 (2004); A. Johansen, *Physica A* **338**, 286 (2004).
- [26] H. Ebel, L.-I. Mielsch, and S. Bornholdt, *Phys. Rev. E* **66**, R35103 (2002).
- [27] U. Harder and M. Paczuski, *Physica A* **361**, 329 (2006).
- [28] B. Rácz and A. Lukács, *Data Compression Conference* (IEEE Computer Society Press, 2004).
- [29] *The Collected Papers of Albert Einstein* (Princeton University Press, Princeton, 1993–2004), Vols. 1, 5, 8, 9.
- [30] *The Correspondence of Charles Darwin* (Cambridge University Press, Cambridge, 1984–2004), Vols. 1–14.
- [31] Freud Museum, London, <http://www.freud.org.uk/>
- [32] J. W. Cohen, *The Single Server Queue* (North-Holland, Amsterdam, 1969).
- [33] D. Gross and C. M. Harris, *Fundamentals of Queueing Theory* (Wiley, New York, 1998).
- [34] M. Crovella and A. Bestavros, *IEEE/ACM Trans. Netw.* **5**, 835 (1997).
- [35] M. Mitzenmacher, *Internet Mathematics* **1**, 226 (2004).
- [36] G. A. Miller, *Psychol. Rev.* **63**, 8197 (1956).
- [37] A. Baddeley, *Psychol. Bull.* **101**, 353 (1994); **101**, 668 (1994).
- [38] A. Cobham, *J. Oper. Res. Soc. Am.* **2**, 70 (1954).
- [39] J. Abate, *Queueing Syst.* **25**, 173 (1997).
- [40] S. Redner, *A Guide to First-Passage Processes* (Cambridge University Press, New York, 2001).
- [41] *Diffusion and Reactions in Fractals and Disordered Systems*, edited by D. Ben-Avraham and S. Havlin (Cambridge University Press, Cambridge, 2000).
- [42] C. Hidalgo, *Physica A* e-print cond-mat/0512278 (to be published); Y. Ikegaya and N. Matsuki (unpublished).
- [43] M. Mitzenmacher, *Internet Mathematics* **1**, 226 (2004).
- [44] D. B. Stouffer, R. Dean Malmgren, and L. A. N. Amaral, e-print physics/0510216.
- [45] A similar idea was proposed independently by R. V. R. Pandya, e-print physics/0512086.
- [46] H. Jeong, Z. Néda, and A.-L. Barabási, *Europhys. Lett.* **61**, 567 (2003).
- [47] A.-L. Barabási and R. Albert, *Science* **286**, 509 (1999).
- [48] A.-L. Barabási, H. Jeong, R. Ravasz, Z. Néda, T. Vicsek, and A. Schubert, *Physica A* **311**, 590 (2002).
- [49] R. Pastor-Satorras, A. Vázquez, and A. Vespignani, *Phys. Rev. Lett.* **87**, 258701 (2001).
- [50] A. Vázquez, R. Pastor-Satorras, and A. Vespignani, *Phys. Rev. E* **65**, 066130 (2002).
- [51] M. E. J. Newman, *Phys. Rev. E* **64**, 025102(R) (2001).
- [52] Ph. Blanchard and M. O. Hongler (unpublished).
- [53] K. Yamasaki, L. Muchnik, S. Havlin, A. Bunde, and H. E. Stanley, *Proc. Natl. Acad. Sci. U.S.A.* **102**, 9424 (2005).
- [54] P. Holme, *Phys. Rev. E* **71**, 046119 (2005).
- [55] F. J. Omori, *Coll. Sci. Imper. Univ. Tokyo* **7**, 111 (1894).
- [56] B.-F. Apostol (unpublished).
- [57] A. Corral, *Phys. Rev. Lett.* **92**, 108501 (2004).
- [58] G. M. Viswanathan *et al.*, *Nature (London)* **381**, 413 (1996).
- [59] S. Zapperi, A. Vespignani, and H. E. Stanley, *Nature (London)* **388**, 658 (1997).
- [60] B. Suki, A.-L. Barabási, Z. Hantos, F. Petak, and H. E. Stanley, *Nature (London)* **368**, 615 (1994).
- [61] K. Yamakasi, L. Muchnik, S. Havlin, A. Bunde, and H. E. Stanley, *Proc. Natl. Acad. Sci. U.S.A.* **102**, 9424 (2005).
- [62] A. Bunde, J. F. Eichner, J. W. Kantelhardt, and S. Havlin, *Phys. Rev. Lett.* **94**, 048701 (2005).
- [63] P. Bak and K. Sneppen, *Phys. Rev. Lett.* **71**, 4083 (1993).
- [64] H. J. Jensen, *Self-Organised Criticality* (Cambridge University Press, Cambridge, 1998).
- [65] P. Bak, C. Tang, and K. Wiesenfeld, *Phys. Rev. Lett.* **59**, 381 (1987).
- [66] M. Paczuski, S. Maslov, and P. Bak, *Phys. Rev. E* **53**, 414 (1996).

- [67] H. Flyvbjerg, K. Sneppen, and P. Bak, *Phys. Rev. Lett.* **71**, 4087 (1993).
- [68] J. de Boer, B. Derrida, H. Flyvbjerg, A. D. Jackson, and T. Wettig, *Phys. Rev. Lett.* **73**, 906 (1994).
- [69] B. A. Huberman and D. Helbing, *Europhys. Lett.* **47**, 196 (1999).
- [70] D. Helbing, S. Lammer, T. Seidel, P. Seba, and T. Platkowski, *Phys. Rev. E* **70**, 066116 (2004).
- [71] D. Helbing, S. Lammer, U. Witt, and T. Brenner, *Phys. Rev. E* **70**, 056118 (2004).
- [72] D. Helbing, I. Farkas, and T. Vicsek, *Nature (London)* **407**, 487 (2000).
- [73] A. Vazquez, e-print physics/0603064.