process is highly effective, most recruits would not reach the intended food sources without the use of odour and visual cues in the final stages of their flight. We hope that together with earlier studies, particularly those of Gould[10], Srinivasan *et al.*[13] and Esch *et al.*[14], our results will also be accepted as a vindication of the von Frisch hypothesis. □

Methods

Harmonic radar

Most of our knowledge of insect flight behaviour at high altitudes has been derived from the use of radar modified for entomological observations[29]. More recently, it has become possible to apply this technique to low-flying insects by tagging them with tiny harmonic transponders, weighing only a few milligrams (6–20 mg, depending on the degree of mechanical robustness required). The transponders return signals to the radar at twice the original transmitted frequency, and because these signals can be distinguished from returns from ground features and other unwanted targets, the position of tagged insects can be determined while they are in flight[15–17]. In our experiments we caught bees as they exited their hive, attached transponders, and released them from the hive (or from remote release points), and recorded their subsequent flight trajectories. Nineteen of the 23 recruits (83%) released from the hive flew off satisfactorily, and produced good radar tracks, all to the east. Of the remaining four, three also went east, but were not detected by the radar enough times to produce a satisfactory track, probably because they were flying very low. Only one failed to leave the vicinity of the hive. Similar results were achieved for the remote releases, except during a period when the tubes in which the recruits were being transported became accidentally contaminated with sucrose solution, with the result that these bees did not fly away from the release point.

Experimental arena

The flight observations were made over a carefully selected[30], large area of mown pastureland, approximately 1 × 1.5 km, where the terrain was unusually flat and free from obstacles that would have obscured the radar's field of view[28]. The radar was positioned on the southern edge of the arena, so that it overlooked an observation hive and a feeding station 200 m to the east of the hive. Three release points were set 200–250 m in the sector to the southwest of the hive. There were very few natural sources of pollen and nectar present during our study period (late July/early August 2000).

Description of the wind field

Wind speed and direction were recorded at 10-s intervals at a height of 2.7 m by anemometers and wind vanes placed at the corners of a 500 m × 600 m rectangle centred on the hive. We also set up a mast near to the centre of the rectangle, holding anemometers at heights of 0.65, 1.3, 2.7 and 8.2 m, and a wind vane at 2.7 m. The clocks of the recording data loggers were synchronized each morning with a master clock at the radar to ensure that wind data were recoverable for the duration of each individual recorded flight. Using interpolation methods described elsewhere[16,27] the data collected by these instruments were then combined to describe the mean wind field within which each of the flights recorded by the radar took place.

The observation hive

We used a two-frame colony, equipped with a transparent side panel that faced directly into a small, low tent attached to the hive. From within this darkened enclosure we could observe the dance behaviour of the bees, and their entry into and exit from the hive. Most of the bees in the small colony were marked with numbered tags. The entrance to the hive was in the form of clear plastic tube, so that observers stationed outside the tent could also observe entry and exit, and capture selected bees for tagging with transponders.

The flight experiments

We began our study by establishing in our experimental bees (European species *Apis mellifera carnica*) a route memory of the position of a feeding station relative to the hive. The feeder was placed directly to the east of the hive, and supplied with 0.2–1 M sucrose solution, and training to a distance of 200 m was accomplished over two days. No artificial odour cues were used in the procedure. From the start, observers at the feeder recorded the identification number of every marked bee that arrived there, and the hive was never opened unless an observer was at the feeder. Once foraging flights between hive and feeder were well established, observers in the tent watched the waggle dances. Whenever a numbered bee was seen to follow a dance and then move directly towards the exit, an observer outside the tent was alerted to catch the bee as it attempted to leave. If its identification number indicated that it had never previously visited the feeder, the bee was confirmed as a recruit, and a transponder attached. The bee was then either released directly from the hive exit, or taken in an opaque tube to one of three release points 200–250 m from the hive, and allowed to fly from there. Bees fitted with transponders could be detected while in flight within a 190° arc of radius 900 m, centred on the radar; their positions were shown once every 3 s on the screen of a desktop personal computer, and their coordinates recorded[17].

Received 28 October 2004; accepted 8 March 2005; doi:10.1038/nature03526.

1. von Frisch, K. *The Dance Language and Orientation of Bees* (Harvard Univ. Press, Cambridge, Massachusetts, 1967).
2. Dyer, F. C. The biology of the dance language. *Annu. Rev. Entomol.* **47,** 917–949 (2002).
3. Gould, J. L. Honey bee recruitment: the dance-language controversy. *Science* **189,** 685–693 (1975).
4. Sherman, G. & Visscher, P. K. Honeybee colonies achieve fitness through dancing. *Nature* **419,** 920–922 (2002).
5. Wenner, A. M. & Johnson, D. L. Honeybees: do they use direction and distance information provided by their dancers? *Science* **158,** 1076–1077 (1967).
6. Wells, P. H. & Wenner, A. M. Do honey bees have a language? *Nature* **241,** 171–175 (1973).
7. Wenner, A. M. & Wells, P. H. *Anatomy of a Controversy: the Question of a "Language" Among Bees* (Columbia Univ. Press, New York, 1990).
8. Wenner, A. M., Meade, D. E. & Friesen, L. J. Recruitment, search behavior, and flight ranges of honey bees. *Am. Zool.* **31,** 768–782 (1991).
9. Wenner, A. M. The elusive honey bee dance "language" hypothesis. *J. Insect Behav.* **15,** 859–878 (2002).
10. Gould, J. L. The dance-language controversy. *Q. Rev. Biol.* **51,** 211–244 (1976).
11. Dyer, F. C. & Smith, B. H. A review of "Anatomy of a Controversy: the Question of a 'Language' Among Bees" by A.M. Wenner & P.H. Wells. *Anim. Behav.* **47,** 1242–1244 (1994).
12. Seeley, T. D. Bee warned. *Nature* **349,** 114 (1991).
13. Srinivasan, M. V., Zhang, S. W., Altwein, M. & Tautz, J. Honeybee navigation: nature and calibration of the "odometer". *Science* **287,** 851–853 (2000).
14. Esch, H. E., Zhang, S. W., Srinivasan, M. V. & Tautz, J. Honeybee dances communicate distances measured by optic flow. *Nature* **411,** 581–583 (2001).
15. Riley, J. R. *et al.* Tracking bees with harmonic radar. *Nature* **379,** 29–30 (1996).
16. Riley, J. R. & Osborne, J. L. in *Insect Movement: Mechanisms and Consequences* (eds Woiwod, I. P., Reynolds, D. R. & Thomas, C. D.) 129–157 (CAB International, Wallingford, UK, 2001).
17. Riley, J. R. & Smith, A. D. Design considerations for an harmonic radar to investigate the flight of insects at low altitude. *Comput. Electron. Agric.* **35,** 151–169 (2002).
18. Towne, W. F. & Gould, J. L. The spatial precision of the honey bee's dance communication. *J. Insect Behav.* **1,** 129–155 (1988).
19. Esch, H. & Bastian, J. A. How do newly recruited honeybees approach a food site? *Z. Vergl. Physiol.* **68,** 175–181 (1970).
20. Mautz, D. Der Kommunikationseffekt der Schwänzeltänze bei *Apis mellifica carnica* (Pollm.). *Z. Vergl. Physiol.* **72,** 197–220 (1971).
21. Seeley, T. D. Division of labor between scouts and recruits in honeybee foraging. *Behav. Ecol. Sociobiol.* **12,** 253–259 (1983).
22. Tautz, J. & Sandeman, D. C. Recruitment of honeybees to non-scented food sources. *J. Comp. Physiol. A* **189,** 293–300 (2003).
23. Menzel, R. *et al.* Honeybees navigate according to a map-like spatial memory. *Proc. Natl Acad. Sci. USA* **102**(8), 3040–3045 (2005).
24. Baker, T. C. & Haynes, K. F. Pheromone-mediated optomotor anemotaxis and altitude control exhibited by male oriental fruit moths in the field. *Physiol. Entomol.* **21,** 20–31 (1996).
25. Cardé, R. T. & Mafra-Neto, A. in *Insect Pheromone Research: New Directions* (eds Cardé, R. T. & Minks, A. K.) 275–290 (Chapman & Hall, New York, 1997).
26. Riley, J. R. *et al.* Harmonic radar as a means of tracking the pheromone-finding and pheromone-following flight of male moths. *J. Insect Behav.* **11,** 287–296 (1998).
27. Riley, J. R. *et al.* Compensation for wind drift by bumble-bees. *Nature* **400,** 126 (1999).
28. Riley, J. R. *et al.* The automatic pilot of honeybees. *Proc. R. Soc. Lond. B* **270,** 2421–2424 (2003).
29. Reynolds, D. R. *et al.* in *Methods in Ecological and Agricultural Entomology* (eds Dent, D. R. & Walton, M. P.) 111–145 (CAB International, Wallingford, UK, 1997).
30. Chittka, L. & Geiger, K. Honeybee long-distance orientation in a controlled environment. *Ethology* **99,** 117–126 (1995).

Acknowledgements Our field studies depended on the contributions of E. Schüttler, S. Watzl, G. Bundrock, J. Stindt, S. Berger, S. Hülse, S. Brunke and T. Plümpe, and we are also grateful to A. Edwards for assistance with the reduction of the radar data and manuscript preparation. Data analysis and writing up by J.R.R. was supported by the Leverhulme Trust Emeritus Fellowship programme. The research was carried out with joint funding from the Deutsche Forschungsgemeinschaft and the Biotechnology and Biological Sciences Research Council of the United Kingdom (BBSRC). Rothamsted Research receives grant aided assistance from the BBSRC.

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

Correspondence and requests for materials should be addressed to J.R.R. (joe@radarent.freeserve.co.uk).

The origin of bursts and heavy tails in human dynamics

Albert-László Barabási

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

The dynamics of many social, technological and economic phenomena are driven by individual human actions, turning the quantitative understanding of human behaviour 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[1–3]. In contrast, there is increasing evidence that the timing of many human activities, ranging from communication to entertainment and work patterns, follow non-Poisson statistics, characterized by bursts of rapidly occurring events separated by long periods of inactivity[4–8]. Here I show that the bursty nature of human behaviour is a consequence of a decision-based queuing process[9,10]: when individuals execute tasks based on some perceived priority, the timing of the tasks will be heavy tailed, with most tasks being rapidly executed, whereas a few experience very long waiting times. In contrast, random or priority blind execution is well approximated by uniform inter-event statistics. These finding have important implications, ranging from resource management to service allocation, in both communications and retail.

Humans participate on a daily basis in a large number of distinct activities, ranging from electronic communication (such as sending e-mails or making telephone calls) to browsing the Internet, 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 seems impossible to seek regularities in human dynamics, 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.

Current models of human activity are based on Poisson processes, and assume that in a dt time interval an individual (agent) engages in a specific action with probability qdt, where q is the overall frequency of the monitored activity. This model predicts that the time interval between two consecutive actions by the same individual, called the waiting or inter-event time, follows an exponential distribution (Fig. 1a–c)[1]. Poisson processes are widely used to quantify the consequences of human actions, such as modelling traffic flow patterns or accident frequencies[1], and are commercially used in call centre staffing[2], inventory control[3], or to estimate the number of congestion-caused blocked calls in calls in mobile communication[4]. Yet, an increasing number of recent measurements indicate that the timing of many human actions systematically deviates from the Poisson prediction, the waiting or inter-event times being better approximated by a heavy tailed or Pareto distribution (Fig. 1d–f). The differences between Poisson and heavy-tailed behaviour are striking: a Poisson distribution decreases exponentially, forcing the consecutive events to follow each other at relatively regular time intervals and forbidding 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 (Fig. 1).

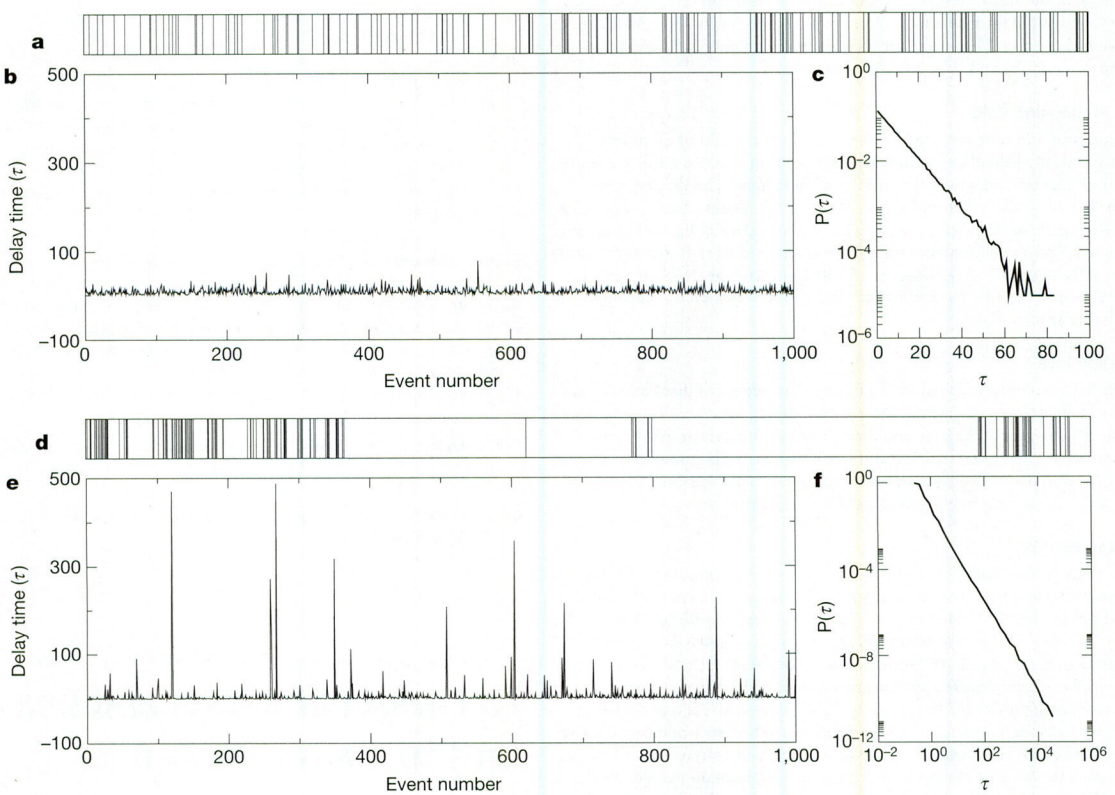

Figure 1 The difference between the activity patterns predicted by a Poisson process and the heavy-tailed distributions observed in human dynamics. **a**, Succession of events predicted by a Poisson process, which assumes that in any moment an event takes place with probability q. The horizontal axis denotes time, each vertical line corresponding to an individual event. Note that the inter-event 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 1,000 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 $P(n; q) = e^{-qt}(qt)^n/n!$, which predicts that for a Poisson process the inter-event time distribution follows $P(\tau) = qe^{-q\tau}$, shown on a log-linear plot in **c** for the events displayed in **a**, **b**. **d**, The succession of events for a heavy-tailed distribution. **e**, The waiting time τ of 1,000 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 the comparison of the regularity of a Poisson process with the intermittent nature of the heavy-tailed process. **f**, Delay time distribution $P(\tau) \simeq \tau^{-2}$ for the heavy-tailed process shown in **d**, **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 the Supplementary Information.

To provide direct evidence for non-Poisson activity patterns in individual human behaviour, I study the communication between several thousand e-mail users based on a data set capturing the sender, recipient, time and size of each e-mail[11–12]. As Fig. 2a shows, the distribution of the time differences between consecutive e-mails sent by a selected user is best approximated with $P(\tau) \approx \tau^{-\alpha}$, where $\alpha \simeq 1$, indicating that an individual's e-mail pattern has a bursty non-Poisson character: during a single session a user sends several e-mails in quick succession, followed by long periods of no e-mail activity. This behaviour is not limited to e-mail communications. Measurements capturing the distribution of the time differences between consecutive instant messages sent by individuals during online discussions[5] show a similar pattern. Professional tasks, such as the timing of job submissions on a supercomputer[6], directory listing and file transfers (FTP request) initiated by individual users[7], or the timing of printing jobs submitted by users[13] were also reported to display non-Poisson features. Similar patterns emerge in economic transactions, describing the time interval distributions between individual trades in currency futures[8]. Finally, heavy-tailed distributions characterize entertainment-related events, such as the time intervals between consecutive online games played by the same user[14].

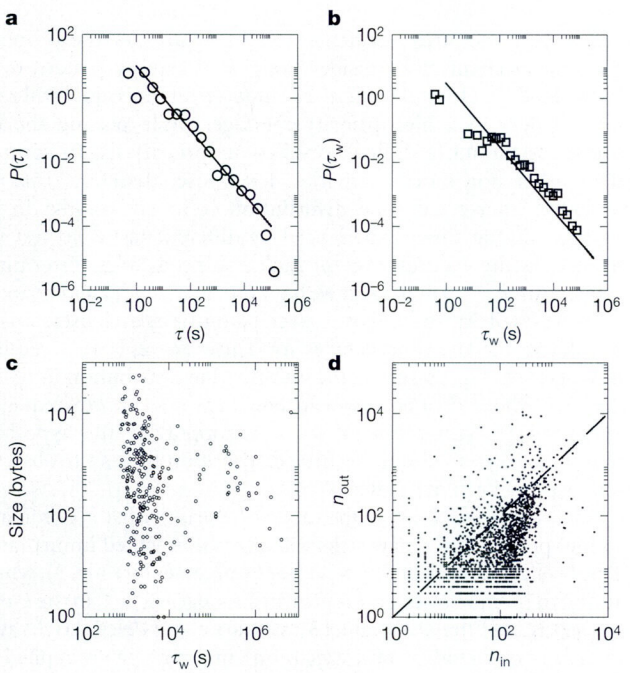

Figure 2 Heavy-tailed activity patterns in e-mail communications. **a**, The distribution of the time intervals between consecutive e-mails sent by a single user over a three-month time interval, indicating that $P(\tau) \approx \tau^{-1}$ (the solid line in the log–log plot has slope -1). Although the exponent differs slightly from user to user, it is typically centred around $\alpha = 1$. **b**, The distribution of the time taken by the user to reply to a received message. To determine τ_w we recorded the time the user received an e-mail from a specific user, and the time it sent a response to that user, the difference between the two providing τ_w. For consistency the figure shows the data for the user whose inter-event time distribution is shown in **a**. The solid line in the log–log plot has slope -1. **c**, A scatter plot showing the waiting time τ_w and the size for each e-mail responded to by the user discussed in **a**, **b**, indicating that the file size and response time do not correlate. **d**, Scatter plot showing the number of e-mails received and sent by 3,188 users during a three-month interval. Each point corresponds to a different user, indicating that there are significant differences between the number of received and responded e-mails. The dashed line corresponds to $n_{in} = n_{out}$, capturing the case when the classical queuing models also predict a power law waiting time distribution (see Supplementary Information), albeit with exponent $\alpha = 3/2$.

The fact that a wide range of human activity patterns follow non-Poisson statistics suggests that the observed bursty character reflects some fundamental and potentially generic feature of human dynamics. Yet, the mechanism responsible for these marked non-random features remains unknown. Here, I show that the bursty nature of human dynamics is a consequence of a queuing process driven by human decision making: 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 follow a Pareto distribution. In contrast, first-come-first-serve and random task execution, common in most service-oriented or computer-driven environments, lead to uniform Poisson-like dynamics.

Most human-initiated events require an individual to assess and prioritize different activities. Indeed, at the end of each activity an individual needs to decide what to do next—for example send an e-mail, do some shopping, or make a telephone call—allocating time and resources for the chosen activity. Consider an agent operating with a priority list of L tasks. After a task is executed, it is removed from the list, offering the opportunity to add another task. The agent assigns to each task a priority parameter x, which allows it to compare the urgency of the different tasks on the list. The question is, how long will a given task have to wait before it is executed. The answer depends on the method the agent uses to choose the task to be executed next. In this respect three selection protocols[10] are particularly relevant for human dynamics.

(i) The simplest selection rule is the first-in-first-out protocol, executing the tasks in the order that they were added to the list. This is common in service-oriented process, such as the first-come-first-serve execution of orders in a restaurant or getting help from directory assistance and consumer support. The time period an item stays on the list before execution is determined by the cumulative time required to perform all tasks added to the list before it. If the time necessary to perform the individual tasks are chosen from a bounded distribution (that is, the second moment of the distribution is finite), then the waiting time distribution will develop an exponential tail, indicating that most tasks experience approximately the same waiting time.

(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 in Internet communications. The waiting time distribution of the individual tasks (that is, the time between two calls on the same student) in this case is also exponential.

(iii) In most human-initiated activities task selection is not random, but the individual executes the highest-priority item on its list. The resulting execution dynamics is quite different from the first (i) and second (ii) selection protocols: high-priority tasks will be executed soon after their addition to the list, whereas low-priority items will have to wait until all higher-priority tasks are cleared, forcing them to stay on the list for considerable time intervals. Below, I show that this selection mechanism, practiced by humans on a daily basis, is the probable source of the fat tails observed in human-initiated processes.

I assume that an individual has a priority list with L tasks, each task being assigned a priority parameter x_i, where $i = 1, \ldots, L$, chosen from a $\rho(x)$ distribution. At each time step the agent selects the highest-priority task from the list and executes it, removing it from the list. At that moment a new task is added to the list, its priority x_i being again chosen from $\rho(x)$. This simple model ignores the possibility that the agent occasionally selects a low-priority item for execution before all higher-priority items are done—common, for example, for tasks with deadlines. This can be incorporated by assuming that the agent executes the highest-priority item with probability p, and with probability $1 - p$ executes a randomly

selected task, independent of its priority. Thus the $p \to 1$ limit of the model describes the deterministic protocol (iii), when always the highest-priority task is chosen for execution whereas $p \to 0$ corresponds to the random choice protocol (ii) discussed above.

To establish that this priority list model can account for the observed fat-tailed inter-event time distribution, we first studied its dynamics numerically with priorities chosen from a uniform distribution $x_i \in [0, 1]$. Computer simulations show that in the $p \to 1$ limit the probability that a task spends τ time on the list has a power law tail with exponent $\alpha = 1$ (Fig. 3a) in agreement with the exponent obtained for e-mail communications (Fig. 2a). In the $p \to 0$ limit $P(\tau)$ follows an exponential distribution (Fig. 3b), as expected for the case (ii). As the typical length of the priority list differs from individual to individual, it is particularly important for the tail of $P(\tau)$ to be independent of L. Numerical simulations indicate that this is indeed the case: changes in L do not affect the scaling of $P(\tau)$. The fact that the scaling holds for $L = 2$ indicates that it is not necessary to have a long priority list: as long as individuals balance at least two tasks, a bursty, heavy-tailed inter-event dynamics will emerge.

To determine the tail of $P(\tau)$ analytically I consider a stochastic version of the model in which the probability to choose a task with priority x for execution in a unit of time is $\Pi(x) \approx x^{\gamma}$, where γ is a parameter that 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 parameterization captures the scaling of the model only in the $p \to 0$ and $p \to 1$ limits, but not for intermediate p values, thus it is chosen only for mathematical convenience. The probability that a task with priority x is executed at time t 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

$$\tau(x) = \sum_{t=1}^{\infty} t f(x,t) = \frac{1}{\Pi(x)} \approx \frac{1}{x^{\gamma}} \qquad (1)$$

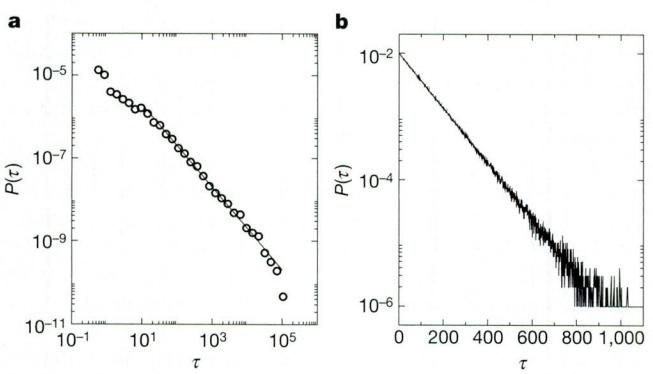

Figure 3 The waiting time distribution predicted by the investigated queuing model. The priorities were chosen from a uniform distribution $x_i \in [0,1]$, and I monitored a priority list of length $L = 100$ over $T = 10^6$ time steps. **a**, Log–log plot of the tail of probability $P(\tau)$ that a task spends τ time on the list obtained for $p = 0.99999$, corresponding to the deterministic limit of the model. The continuous line of the log–log plot corresponds to the scaling predicted by equation (2), having slope -1, in agreement with the numerical results and the analytical predictions. The data were log-binned, to reduce the uneven statistical fluctuations common in heavy-tailed distributions, a procedure that does not alter the slope of the tail. For the full curve, including the $\tau = 1$ peak, see Supplementary Fig. 3. **b**, Linear-log plot of the $P(\tau)$ distribution for $p = 0.00001$, corresponding to the random choice limit of the model. The fact that the curve follows a straight line on a linear-log plot indicates that $P(\tau)$ decays exponentially.

that is, the higher an item's priority, the shorter the average time it waits before execution. To calculate $P(\tau)$ I use the fact that the priorities are chosen from the $\rho(x)$ distribution; that is, $\rho(x)\mathrm{d}x = P(\tau)\mathrm{d}\tau$, which gives

$$P(\tau) \approx \frac{\rho(\tau^{-1/\gamma})}{\tau^{1+1/\gamma}} \qquad (2)$$

In the $\gamma \to \infty$ limit, which converges to the strictly priority-based deterministic choice ($p = 1$) in the model, equation (2) predicts $P(\tau) \approx \tau^{-1}$, in agreement with the numerical results (Fig. 3a), as well as the empirical data on the e-mail inter-arrival times (Fig. 2a). In the $\gamma = 0$ ($p = 0$) limit, $\tau(x)$ is independent of x thus $P(\tau)$ converges to an exponential distribution, as shown in Fig. 3b (see Supplementary Information).

The apparent dependence of $P(\tau)$ on the $\rho(x)$ 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 $\rho(x)$ distribution. According to equation (2), however, in the $\gamma \to \infty$ limit, $P(\tau)$ is independent of $\rho(x)$. Indeed, in the deterministic limit the uniform $\rho(x)$ can be transformed into an arbitrary $\rho'(x)$ with a parameter change, without altering the order in which the tasks are executed[11]. This insensitivity of the tail to $\rho(x)$ explains why, despite the diversity of human actions encompassing both professional and personal priorities, most decision-driven processes develop a heavy tail.

To obtain empirical evidence for the validity of the proposed queuing mechanism I consider the e-mail activity pattern of an individual[11,12]. Once in front of a computer, an individual will reply immediately to a high-priority message, while placing the less urgent or the more difficult ones on its priority list to compete with other non-e-mail activities. I propose, therefore, that the observed inter-event time distribution is in fact rooted in the uneven waiting times experienced by different tasks. To test this hypothesis the waiting time for each task needs to be determined directly. In the e-mail data set we have the time, sender and recipient of each e-mail transmitted over several months by each user, thus we can determine the time it takes for a user to reply to a received message[11]. As Fig. 2b shows, the waiting time distribution $P(\tau_w)$ for the user whose $P(\tau)$ is shown in Fig. 2a is best approximated by $P(\tau_w) \approx \tau_w^{-\alpha_w}$ with exponent $\alpha_w = 1$, supporting the hypothesis that the heavy-tailed waiting time distribution drives the observed bursty e-mail activity patterns.

As in the $p \to 1$ limit of the model the priority list is dominated by low-priority tasks, new tasks will often be executed immediately. This results in a peak at $P(\tau = 1)$ (see Supplementary Fig. 3), which, although in some cases may represent a model artefact, in the e-mail context is not unrealistic: most e-mails are either deleted right away (which is one kind of task execution) or immediately replied to. Only the more difficult or time consuming tasks will queue on the priority list. The e-mail data set does not allow us to resolve this peak, however, because a message which the user deletes or replies to right away will appear to have some waiting time, given the delay between the arrival of the message and the time the user has a chance to check her e-mail.

Although I have illustrated the queuing process for e-mails, in general the model is better suited to capture the competition between different kinds of activities an individual is engaged in; that is, the switching between various work, entertainment and communication events. Indeed, most data sets displaying heavy-tailed inter-event times in a specific activity reflect the outcome of the competition between tasks of different nature. For example, the starting of an online gaming session often implies that all higher-priority work-, family-, and entertainment-related activities have been already executed.

Detailed models of human activity require us to consider the impact of a number of additional mechanisms on the queuing

process. First, in the priority list model I have assumed that the time necessary to execute a task (service time) is the same for all tasks. The size distribution of e-mails is heavy tailed[15,16], however, thus the waiting time distribution could be driven entirely by the time it takes to read an e-mail (that is, the message size). Yet, as Fig. 2c shows, there is no correlation between the size of the e-mail received by a user and the time the user takes to reply to it. Although detailed analysis should also consider the role of attachments, Fig. 2c suggests that the priority of a response is more important than the message size. Furthermore, the priorities assigned to tasks are often driven by optimization processes, as agents aim to maximize profits or minimize the overall time spent on some activity.

A natural extension of the model is to assume that tasks arrive at a rate λ and are executed at a rate μ, allowing the length of the priority list L to change in time. In this case the model maps into Cobham's priority queue model[9], which has a power-law-distributed waiting time with $\alpha = 3/2$ when $\lambda = \mu$ (see Supplementary Information). Thus, to account for the power law waiting times the model requires an additional mechanism that guarantees $\lambda = \mu$ (which, as Fig. 3d indicates, is not satisfied for most e-mail users). In contrast, in the proposed priority list model I have assumed that for humans the length of the priority list remains relatively unchanged (L is constant). To understand the orgin of this assumption we must realize that for $\lambda = \mu$ the length of the priority list fluctuates widely and can occasionally grow very long. Although keeping track of a long priority list is not a problem for a computer, it is well established that the immediate memory of humans has finite capacity[17]. In other words, the number of priorities that we can easily remember, and therefore the length of the priority list, is bounded, motivating the choice of a finite L.

Although other generalizations are possible and often required, the main finding is that the observed fat-tailed activity distributions can be explained by a simple hypothesis: humans execute their tasks based on some perceived priority, setting up queues that generate very uneven waiting time distributions for different tasks. In this respect the proposed priority list model represents only a minimal framework that allows us to demonstrate the potential origin of the heavy-tailed activity patterns, and offers room for further extensions to capture more complex human behaviour. As the exponent of the tail could depend on the details of the prioritizing process, future work may allow the empirical data to discriminate between different queuing hypotheses. A mapping into punctuated equilibrium models (see Supplementary Information and refs 18, 19) with the mathematical framework of queuing theory could help the systematic classification of the various temporal patterns generated by human behaviour.

There is overwhelming evidence that Internet traffic is characterized by heavy-tailed statistics[20], rooted in the Pareto size distribution of the transmitted files[15,16]. As I have shown (Fig. 2c), a user's e-mail activity does not correlate with e-mail size. Similarly, the timing of online games[14] or sending an instant message[5] cannot be driven by file size either. This suggests that Internet traffic is in fact driven by two separate processes: the heavy-tailed size distribution of the files sent by the users, and the human decision-driven timing of various Internet-mediated activities individuals engage in. In some environments this second mechanism, the origin of which is addressed in this paper, can be just as important as the much investigated first one. Given the differences in routing performance under Poisson and Pareto arrival time distributions[20–22], a queuing–based model of human-driven arrival times could also contribute to a better understanding of Internet traffic.

Uncovering the mechanisms governing the timing of various human activities has significant scientific and commercial potential. First, models of human behaviour are indispensable for large-scale models of social organization, ranging from detailed urban models[23,24] to modelling the spread of epidemics and viruses, the development of panic[25] or capturing financial market behaviour[26]. Understanding the origin of the non-Poisson nature of human dynamics could fundamentally alter the dynamical conclusions these models offer. Second, models of human behaviour are crucial for better resource allocation and pricing plans for telephone companies, to improve inventory and service allocation in both online and "high street" retail, and potentially to understand the bursts of ideas and memes emerging in communication and publication patterns[27]. Finally, heavy tails have been observed in the foraging patterns of birds as well[28], raising the intriguing possibility that animals also use some evolutionarily encoded priority-based queuing mechanisms to decide between competing tasks, such as caring for offspring, gathering food, or fighting off predators. □

Received 7 November 2004; accepted 15 February 2005; doi:10.1038/nature03459.

1. Haight, F. A. *Handbook of the Poisson Distribution* (Wiley, New York, 1967).
2. Reynolds, P. *Call Center Staffing* (The Call Center School Press, Lebanon, Tennessee, 2003).
3. Greene, J. H. *Production and Inventory Control Handbook* 3rd edn (McGraw-Hill, New York, 1997).
4. Anderson, H. R. *Fixed Broadband Wireless System Design* (Wiley, New York, 2003).
5. Dewes, C., Wichmann, A. & Feldman, A. in *Proc. 2003 ACM SIGCOMM Conf. Internet Measurement (IMC-03)* (ACM Press, New York, 2003).
6. Kleban, S. D. & Clearwater, S. H. Hierarchical Dynamics, Interarrival Times and Performance. Proc. SC2003 ⟨http://www.sc-conference.org/sc2003/paperpdfs/pap222.pdf⟩ (2003).
7. Paxson, V. & Floyd, S. Wide-area traffic: The failure of Poisson modeling. *IEEE/ACM Trans. Netw.* **3,** 226 (1996).
8. Masoliver, J., Montero, M. & Weiss, G. H. Continuous-time random-walk model for financial distributions. *Phys. Rev. E* **67,** 021112 (2003).
9. Cobham, A. Priority assignment in waiting line problems. *J. Oper. Res. Sec. Am.* **2,** 70–76 (1954).
10. Cohen, J. W. *The Single Server Queue* (North Holland, Amsterdam, 1969).
11. Eckmann, J.-P., Moses, E. & Sergi, D. Entropy of dialogues creates coherent structure in e-mail traffic. *Proc. Natl Acd. Sci. USA* **101,** 14333–14337 (2004).
12. Ebel, H., Mielsch, L. I. & Bornholdt, S. Scale-free topology of e-mail network. *Phys. Rev. E* **66,** R35103 (2002).
13. Harder, U. & Paczuski, M. Correlated dynamics in human printing behavior. Preprint at ⟨http://xxx.lanl.gov/abs/cs.PF/0412027⟩ (2004).
14. Henderson, T. & Nhatti, S. Modelling user behavior in networked games. *Proc. 9th ACM Int. Conf. on Multimetia* 212–220 (ACM Press, New York, 2001).
15. Crovella, M. & Bestravros, A. Self-similarity in World Wide Web traffic: evidence and possible causes. *IEEE/ACM Trans. Netw.* **5,** 835–846 (1997).
16. Mitzenmacher, M. A brief history of generative models for power law and lognormal distributions. *Internet Math.* **1,** 226–251 (2004).
17. Miller, G. A. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychol. Rev.* **63,** 8197 (1956).
18. Bak, P. & Sneppen, K. Punctuated equilibrium and criticality in a simple model of evolution. *Phys. Rev. Lett.* **71,** 4083–4086 (1993).
19. Jensen, H. J. *Self Organised Criticality* (Cambridge Univ. Press, Cambridge, 1998).
20. Park, K. & Willinger, W. *Self-Similar Networks Traffic and Performance Evaluation* (Wiley, New York, 2000).
21. Leighton, F. T., Maggs, B. M. & Rao, S. B. Packet routing and job-shop scheduling in O(congestion + dilation) steps. *Combinatorica* **14,** 167–186 (1994).
22. Harris, C. M., Brill, P. H. & Fischer, M. J. Internet-type queues with power-tailed interarrival times and computational methods for their analysis. *INFORMS J. Comp.* **12,** 261–271 (2000).
23. Eubank, H. *et al.* Controlling epidemics in realistic urban social networks. *Nature* **429,** 180–184 (2004).
24. Manrubia, S. C., Zanette, D. & Sole, R. V. Transient dynamics and scaling phenomena in urban growth. *Fractals* **7,** 1–8 (1999).
25. Helbing, D., Farkas, I. & Vicsek, T. Simulating dynamic features of escape panic. *Nature* **407,** 487–490 (2000).
26. Caldarelli, G., Marsili, M. & Zhang, Y.-C. A prototype model of stock exchange. *Europhys. Lett.* **40,** 479–484 (1997).
27. Kleinberg, J. in *Proc. 8th ACM SIGKDD Intl Conf. Knowledge Discov. Data Mining*, 91–101 (2002).
28. Viswanathan, G. M. *et al.* Optimizing the success of random searches. *Nature* **401,** 911–914 (1999).

Supplementary Information accompanies the paper on **www.nature.com/nature**.

Acknowledgements I have benefited from discussions with A. Vazquez on the mathematical aspects of the model. I also thank L. A. N. Amaral, Z. Dezsö, P. Ivanov, J. Kelley, J. Kertész, A. Motter, M. Paczuski, K. Sneppen, T. Vicsek, W. Whitt and E. Zambrano for useful discussions and comments on the manuscript; J.-P. Eckmann for providing the e-mail database; and S. Aleva for assisting me with manuscript preparation. This research was supported by NSF grants.

Competing interests statement The author declares that he has no competing financial interests.

Correspondence and requests for materials should be addressed to A.-L. B. (alb@nd.edu).