Notes and Comment

1/f noise in human cognition: Is it ubiquitous, and what does it mean?

SIMON FARRELL

University of Bristol, Bristol, England

ERIC-JAN WAGENMAKERS

University of Amsterdam, Amsterdam, The Netherlands

and

ROGER RATCLIFF

Ohio State University, Columbus, Ohio

Researchers in psychology are paying increasing attention to temporal correlations in performance on cognitive tasks. Recently, Thornton and Gilden (2005) introduced a spectral method for analyzing psychological time series; in particular, this method is tailored to distinguish transient serial correlations from the persistent correlations characterized by 1/f noise. Thornton and Gilden applied their method to word-naming data to support the claimed ubiquity of 1/f noise in psychological time series. We argue that a previously presented method for distinguishing transient and persistent correlations (e.g., Wagenmakers, Farrell, & Ratcliff, 2004) compares favorably with the new method presented by Thornton and Gilden. We apply Thornton and Gilden's method to time series from a range of cognitive tasks and show that 1/f noise is not a ubiquitous property of psychological time series. Finally, we assess the theoretical developments in this area and argue that the development of well-specified models of the principles or mechanisms of human cognition giving rise to 1/f noise is long overdue.

In conducting psychological research, it is essential to have models of the processes or principles theorized to underlie observed behavior and to have experimental and statistical methods sufficient to determine the theoretical adequacy of these models. This has become increasingly evident in recent work in which the dynamics of human behavior have been examined, particularly in investigations of the presence of 1/f noise in human cognition (Gilden, 2001; Gilden & Wilson, 1995; Van Orden, Holden, & Turvey, 2003, 2005; Wagenmakers, Farrell, & Ratcliff, 2004, 2005). 1/f noise is a particular type of stochastic process

that possesses the unique characteristic of self-similarity: 1/f noise looks the same (both visually and statistically) at time scales varying over orders of magnitude (e.g., Beran, 1994). Temporal self-similarity also implies long-range dependence: Observations separated by a large number of intervening observations in a 1/f time series tend to be correlated, and these correlations decrease slowly with increasing separation (i.e., the correlations are persistent). Finally, the presence of 1/f noise is taken as a signature of complexity in time series (Thornton & Gilden, 2005), and its presence has been used to argue for the abandonment of standard scientific methods employed in psychology in favor of a focus on emergent properties (Van Orden et al., 2003, 2005). A full presentation of the properties and implications of 1/f noise can be found in Wagenmakers et al. (2004) and Gilden (2001).

In a recent article, Thornton and Gilden (2005) introduced a new method for detecting 1/f noise in psychological time series. Two central claims are made in Thornton and Gilden's presentation of their spectral classifier. One claim is that 1/f processes are distinguishable from alternative processes that do not possess the characteristics of 1/f noise but may mimic the statistics of such processes (Wagenmakers et al., 2004), given that appropriate statistical methods are used. A second, more contentious claim is that applying their spectral classifier reveals that 1/f noise generally provides a better description of psychological time series than do alternative models.¹

Although we applaud the adoption of a more rigorous method for detecting 1/f noise in psychological time series (in line with previous suggestions: Wagenmakers et al., 2004), we feel that there are several important statistical and theoretical issues arising from Thornton and Gilden's (2005) presentation that must be addressed. We will examine the relationship between Thornton and Gilden's method and model selection methods previously presented for detecting 1/f noise (Wagenmakers et al., 2004, 2005) and will argue that the new method presented by Thornton and Gilden improves detection little, in comparison with that presented by Wagenmakers et al. (2004). We then will discuss Thornton and Gilden's claim of the ubiquity of 1/f noise in psychological time series and will show, by application of their method to data that we have previously collected (Wagenmakers et al., 2004), that 1/f noise is not a general property of human behavior. Finally, we will address an important issue neglected in Thornton and Gilden's article, that of the lack of theoretical development in research on 1/f noise in psychology.

Methods for Detecting 1/f Noise

Thornton and Gilden (2005) focus primarily on the presentation of a method for distinguishing 1/f noise from stochastic processes that are not true 1/f noise but may

The second author was supported by a VENI grant from the Dutch Organization for Scientific Research (NWO). The third author was supported by NIMH Grant R37-MH44640. We thank Michael Lee, David Gilden, Thomas Thornton, and an anonymous reviewer for their insightful reviews of a previous version of the manuscript. We are also grateful for Tom Thornton's assistance in clarifying the specifics of the spectral classification method. Correspondence concerning this article should be addressed to S. Farrell, Department of Experimental Psychology, University of Bristol, 12a Priory Road, Clifton, Bristol BS8 1TU, England (e-mail: simon.farrell@bristol.ac.uk).

mimic the statistical properties of $1/f$ noise—that is, distinguishing persistent serial correlations from transient correlations. The method presented by Thornton and Gilden is a spectral classifier, in which the likelihood of a time series (or set of time series) is estimated by comparing the power spectrum (the frequency domain representation) of a time series with a library of spectra derived from two candidate models of serial correlations. One model considered by Thornton and Gilden is $1/f$ noise, treated by Thornton and Gilden as fBmW: fractional Brownian motion (fBm)² with added Gaussian white noise (white noise possessing no systematic serial correlations). The added white noise is sometimes interpreted as independent variability in motor processes (e.g., Gilden, 1997). The other model they consider is an autoregressive moving average model, the ARMA(1, 1) model, in which the value of a series at time t depends only on the state of the system at time $t - 1$; that is,

$$X_t = \phi X_{t-1} + \varepsilon_t + \theta \varepsilon_{t-1}. \quad (1)$$

For a full description of ARMA models, see Brockwell and Davis (1996). The purpose in comparing these two models is that although possessing only transient correlations, the temporal statistics of an ARMA(1, 1) time series can resemble those of $1/f$ noise (Wagenmakers et al., 2004). Thornton and Gilden (2005) present some simulation results showing that their method is well suited to discriminating series generated from these two specific models.

It is clear that the method presented by Thornton and Gilden (2005) is more rigorous than a method that has been commonly employed. Previous investigations (Gilden, 1997, 2001; Gilden & Wilson, 1995; Van Orden et al., 2003) of serial correlations in psychology have fit only a single model, a fractal model, and have not considered alternative models of the fluctuations in psychological series. Wagenmakers et al. (2004) showed that this approach is inappropriate, since short-range stochastic processes not possessing the characteristics of $1/f$ noise could be misidentified as $1/f$ noise, using standard procedures such as spectrum fitting. Accordingly, Wagenmakers et al. (2004, 2005) argued that $1/f$ noise cannot be considered in isolation but must be accompanied by examination of alternative models, such as the ARMA model, that can give rise to temporal patterns similar to those of $1/f$ noise but that do not possess long-range dependence. It is promising to see recognition of this important point in the method presented by Thornton and Gilden.

Despite the implication in Thornton and Gilden's (2005) presentation, however, it is not clear that their spectral classification method is also superior to a method previously suggested for distinguishing persistent and transient correlations. Wagenmakers et al. (2004) presented a method in which short-range processes, represented by the ARMA model, are compared with an extension of the ARMA model that incorporates long-range dependencies (see also Beran, Bhansali, & Ocker, 1998). This ARFIMA model (fractionally integrated ARMA model; see, e.g.,

Beran, 1994) incorporates an additional parameter that scales the extent of long-range dependence. Wagenmakers et al. (2004) advocated an approach in which ARMA and ARFIMA models were competitively tested in a model selection framework. Specifically, Wagenmakers et al. (2004) advocated determination of the maximum likelihood of a time series under the ARMA and the ARFIMA models and then selection of one of these models, using an information metric such as Akaike's information criterion (AIC; Akaike, 1974).

The existence of two methods for the estimation of $1/f$ noise demands a comparison of the approaches. Table 1 lists the similarities and differences between the spectral classifier introduced by Thornton and Gilden (2005) and the ARFIMA model selection approach of Wagenmakers et al. (2004). Looking first at the similarities, it is apparent that the two methods are, for the most part, very similar in their approach and application. Both methods, in contrast to those employed previously in psychology (e.g., Gilden, 2001; Gilden & Wilson, 1995), involve the comparison of a long-range dependence model with an alternative model, such as the ARMA(1, 1) model, that displays only short-range dependence. This model selection approach is one of the main strengths of the two procedures, since it ensures that short-range dependence processes will not be misidentified as long-range dependence (see Thornton & Gilden, 2005; Wagenmakers et al., 2004). The ARFIMA approach can be extended to allow estimation of a long-range component, such as ARFIMA(0, d , 0) contaminated by an independent white noise source (e.g., Crato & Ray, 2002; Hsu & Breidt, 2003). The ARFIMA(0, d , 0) plus white noise time series model is very similar to the fBmW model estimated by Thornton and Gilden's spectral classifier.

A related point is that despite objections leveled at the ARFIMA modeling framework by Thornton and Gilden (2005), neither method requires the nesting of the short-range dependence model in a more general long-range dependence model or assumption of the ARMA model as a null hypothesis. Thornton and Gilden devoted several pages to criticism based on the supposed necessity of nesting in the ARMA/ARFIMA framework, claiming that "the sole utility [of the ARFIMA model] arises from its nesting relationship to the ARMA" (p. 29). Wagenmakers et al. (2004) did use an approach in which the ARMA(1, 1) model was competitively compared with the ARFIMA(1, d , 1) model, which can be treated as a nested comparison. Our choice of nested models in Wagenmakers et al. (2004) was motivated by the belief that psychological series are unlikely to be "pure" and that a long-range dependence process would likely be contaminated by short-range dependencies. However, Wagenmakers et al. (2004) did not extend the nesting to statistical comparison: The two models were compared using a general model selection metric, the AIC (Akaike, 1974), rather than a likelihood ratio test naturally suggested by the nested framework. More important, as has been shown in Wagenmakers et al. (2005; see also Beran et al., 1998), a broad range of non-nested ARMA and ARFIMA models can be compared in

Table 1
Similarities and Differences Between the ARFIMA Approach Recommended by Wagenmakers, Farrell, and Ratcliff (2004) and the Spectral Classifier Proposed by Thornton and Gilden (2005)

Similarities	Differences
LRD model compared with SRD alternative [e.g., ARMA(1, 1)]	Fractional model in T&G is LRD + white noise; in ARFIMA, fractional model is LRD with or without SRD
Neither method requires nesting of models	Analysis in time domain (ARFIMA) versus frequency domain
Availability of maximum likelihood and Bayesian methods	Standard software available for ARFIMA (Ox, R, S+); software not freely available for T&G
	ARFIMA method generalizes easily to different series length and alternative models incorporating SRD
	Properties of ARFIMA well known and widely applied; statistical properties of T&G not well explored

Note—LRD, long-range dependence; SRD, short-range dependence.

the ARMA/ARFIMA framework; indeed, Thornton and Gilden acknowledged in their note 4 that the ARMA/ARFIMA framework does not require nesting. One of the important strengths of the ARMA/ARFIMA approach is that a range of ARMA and ARFIMA models can be compared for an observed series. Torre, Delignières, and Lemoine (in press) have recently shown that the generalized model selection approach adopted by Wagenmakers et al. (2005) can reliably estimate fractal noise, with few false responses to ARMA series.

Finally, both methods can be applied in maximum likelihood and Bayesian frameworks. Thornton and Gilden (2005) have demonstrated the use of maximum likelihood and Bayesian estimation in the spectral classifier; others have demonstrated the use of exact maximum likelihood (Hauser, 1999; Sowell, 1992), approximate maximum likelihood (Haslett & Raftery, 1989), prequential (Wagenmakers, Grünwald, & Steyvers, 2006), and Bayesian (Hsu & Breidt, 2003; Pai & Ravishanker, 1998) ARFIMA modeling.

Table 1 also lists some differences between the two approaches. Some of these are inessential differences that could easily be modified in either approach. For example, the ARFIMA model is estimated in the time domain (using the autocovariance function), whereas the spectral classifier requires estimation in the frequency domain. However, the spectral classifier could easily be adapted for analysis in the time domain, and methods exist for estimating ARFIMA models in the frequency domain (e.g., Fox & Taquq, 1986). Other differences between the models are less trivial and generally favor the ARFIMA modeling framework on pragmatic grounds. ARFIMA methods are available in popular statistics and numerical programs, such as Ox (Doornik, 2001), R (Maechler, 2005), and S-Plus, whereas the spectral classifier is not freely available. Accordingly, we have made code for the procedure available on the Web (seis.bristol.ac.uk/~pssaf), as well as details of simulations comparing the spectral classifier with the ARFIMA method. The ARFIMA packages also easily extend to different lengths of time series and higher order models [e.g., ARFIMA (2, d , 2)], whereas the spectral classifier requires generation of a new covariance library for each new model or series length. Finally, given its popularity, the ARFIMA model's properties are well

known (e.g., Beran, 1994; Haslett & Raftery, 1989; Sowell, 1992), whereas those of the spectral classifier have yet to be rigorously explored.

Is $1/f$ Noise Ubiquitous in Psychology?

Thornton and Gilden (2005) advance a second claim that if one has a method that is able to detect $1/f$ noise, application of such a method to psychological time series will reveal $1/f$ noise to be a general property of human behavior. As an example, Thornton and Gilden apply their spectral classifier to the word-naming data of Van Orden et al. (2003). These data are of particular interest, given that Wagenmakers et al. (2005) analyzed the same data, using a range of ARMA and ARFIMA models, and found inconsistent evidence for the presence of $1/f$ noise. In contrast, Thornton and Gilden found that application of their spectral classifier revealed the presence of $1/f$ noise in the majority of series collected by Van Orden et al. (2003).

The claim that $1/f$ noise is "the best explanation for the fluctuations that characterize psychological time series" (Thornton & Gilden, 2005, p. 430) is a strong, general claim that is easily falsified by examining other sets of data. As an example, Table 2 gives the log-likelihood ($\ln L$) differences resulting from application of Thornton and Gilden's spectral classifier to the data collected by Wagenmakers et al. (2004). Wagenmakers et al. (2004) ran 6 participants on three different types of tasks involving responses to numbers (see Wagenmakers et al., 2004, for methodological details). Wagenmakers et al.'s (2004) participants completed a simple reaction time (RT) task (press a single key as soon as a number appeared), a choice RT task (classifying numbers as odd or even), and a time estimation task (press a key 1 sec after onset of a stimulus). Wagenmakers et al. (2004) also manipulated response-stimulus interval (RSI: short vs. long). Table 2 shows that in none of the tasks did Thornton and Gilden's method classify more than half the series as fBmW. The bottom two rows of the table show that aggregating the results across participants by summing $\ln L$ differences reveals convincing evidence for $1/f$ noise in only a single task: time estimation with a long RSI. For all the other series, the summed $\ln L$ differences are less than 0 (i.e., evidence for the ARMA model), or the odds are only very slightly in favor of fBmW. Indeed, in three of these con-

Table 2
Log-Likelihood Differences ($\ln L_{fBmW} - \ln L_{ARMA(1,1)}$)
From Application of the Spectral Classifier to Data From
Wagenmakers, Farrell, and Ratcliff (2004)

Participant	SS	SL	CS	CL	ES	EL
1	1.35	2.52	-1.75	-1.86	-7.82	-2.96
2	-1.00	0.56	2.07	-2.00	-1.11	7.47
3	0.55	-1.57	-0.30	2.56	-0.88	4.41
4	0.93	-2.43	-0.12	0.07	-1.93	10.70
5	-0.68	0.05	1.27	-1.63	2.47	-4.20
6	-2.90	-0.03	-0.83	0.26	-0.05	-1.99
Sum	-1.75	-0.90	0.35	-2.59	-9.32	13.45

Note—The first letter in each condition label designates the task (S, simple RT; C, choice RT; E, time estimation); the second letter designates the RSI condition (S, short; L, long). An $\ln L$ difference greater than 0 indicates that the fBmW model is the best-fitting model. Likelihood ratios can be obtained from summed $\ln L$ differences by $\exp(\Sigma \ln L)$.

ditions (simple RT with a short RSI, choice RT with a long RSI, and time estimation with a short RSI) the data are at least 18 times more likely under the ARMA model than under the fBmW model. Notably, these results are in line with those from the ARFIMA method presented in Wagenmakers et al. (2004; see their Table 1), who found systematic evidence for $1/f$ noise only in the estimation task with a long RSI.

One possible objection to the analysis above is that the conclusions appear to differ from those reached by Wagenmakers et al. (2004) in examining the same data set. First, it is important to note that these apparent differences do not arise due to differences between the spectral classifier and the ARFIMA model. Five of the 17 series classified as $1/f$ noise by the ARFIMA method (Wagenmakers et al., 2004) are classified as ARMA by the spectral classifier, and 4 of the 16 series classified as $1/f$ by the spectral classifier are identified as ARMA by the ARFIMA method; the two methods exhibit considerable overlap in their classification of series. What differs from the analysis of Wagenmakers et al. (2004) is the use of a group analysis. Wagenmakers et al. (2004) were interested only in determining whether $1/f$ noise could be witnessed in psychological series; their conclusion was that in some tasks, for some participants, $1/f$ noise might be observed. Here, in addressing the stronger claim of the ubiquity of $1/f$ noise, we have combined information from several participants by aggregating log-likelihoods to draw more general conclusions. This aggregation reveals a lack of compelling evidence for $1/f$ noise as a general feature of cognitive performance, although it is possible that in some cases, individuals may show behavior consistent with $1/f$ noise.

This example shows that $1/f$ noise is by no means ubiquitous in psychology, even when classification is carried out using the method presented by Thornton and Gilden (2005). The finding that the evidence for the ARMA model was overwhelming in three of the tasks supports the argument that in many cases, the ARMA model is a more appropriate statistical model of the serial correlations in human performance and that the presence of these

transient correlations should be considered in developing theories of the dynamics of human performance.

The Provenance of $1/f$ Noise

One issue that remains unaddressed by Thornton and Gilden (2005), despite the promise of the title of their article, is the provenance of the serial correlations in psychological time series. It must be asked, given that we have several methods for distinguishing $1/f$ noise from potential short-range dependent lures, what types of experiments should we now run in order to advance psychological theory? If it does turn out that $1/f$ correlations are generally observed in cognitive psychology, what does this tell us about human cognition?

Thornton and Gilden (2005) reject the ARMA and ARFIMA models as being “theoretically beholden to autoregression” (p. 418). Although the ARMA/ARFIMA framework is intended as a statistical framework, we would welcome any framework in which the models of serial correlations under analysis are “theoretically beholden,” regardless of the creditor. However, we wonder whether the fBmW model favored by Thornton and Gilden is such a theoretical framework. It is clear that Thornton and Gilden think of the underlying processes as “fractal,” but it is not clear that calling those processes fractal goes any further than describing the statistical properties of $1/f$ noise. Thornton and Gilden discussed several processes that generate $1/f$ -like noise (pp. 411–412), but the charge of a lack of psychological realism can be leveled at these as convincingly as at autoregressive models (indeed, the random-random walk model discussed on their p. 411 is an autoregressive model similar to the ARFIMA model). We puzzle over the implementation of a psychological theory of word naming in which responses arise from “the natural fluctuations emanating from metastable systems that have converged to a phase transition between order and chaos” (Thornton & Gilden, 2005, p. 412).

A comparison can be made between research in $1/f$ noise and research on the power law of practice. The power law of practice is the finding that the time taken to do a task decreases as a power function with amount of practice (see Ritter & Schooler, 2001, for an overview).³ Just as is claimed for $1/f$ noise, this power law of practice has been found to be ubiquitous in psychology, applying to a diverse range of tasks, including cigar rolling (A. Newell & Rosenbloom, 1981) and book writing (Ohlsson, 1992). However, observations of a power law of practice have been accompanied by a clear theoretical development, with several well-formulated explanations being offered for this form of speedup, including the accumulation of instances (Logan, 1988), chunking (K. M. Newell, 1990), and a power function of learning reflecting the statistics of the environment (Anderson & Lebiere, 1998). Notably, these theories have gone beyond the reduction in mean task completion time to account for other learning phenomena, such as the power-shaped drop in variability with practice (e.g., Anderson & Lebiere, 1998; Logan, 1988). We note that the demonstration of $1/f$ in psychological time series

has not been accompanied by such a theoretical development and wonder how such progress can be made.

It is clear from their closing sentence that Thornton and Gilden (2005) are aware of the need for the development of well-specified and testable theories of temporal correlations in human performance. However, looking back, we note that research on $1/f$ noise in cognitive psychology has been ongoing for 10 years now (the first publication on this topic being that of Gilden & Wilson, 1995) in the absence of theoretical development. Despite the advances reflected in the work of Wagenmakers et al. (2004) and Thornton and Gilden, the accounts offered for $1/f$ noise (such as self-organized criticality [SOC]) have as much currency now as they did 10 years ago (Gilden, 2001; Thornton & Gilden, 2005; Van Orden et al., 2003, 2005). Although explanations such as SOC, random-random walks, and iterative maps incorporating bifurcations may be explanations for $1/f$ noise (Thornton & Gilden, 2005; Van Orden et al., 2003, 2005), these theories, which originate outside psychology, need computational formulation in terms of psychological processes. Meanwhile, we question the usefulness of continued investigation of $1/f$ noise in human behavior in the absence of well-specified psychological theories. The prospect of $1/f$ noise being a universal feature of human cognition is certainly an exciting one, but this excitement must be accompanied by theoretical and empirical development.

REFERENCES

- AKAIKE, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, **19**, 716-723.
- ANDERSON, J. R., & LEBIERE, C. (1998). *The atomic components of thought*. Mahwah, NJ: Erlbaum.
- BERAN, J. (1994). *Statistics for long-memory processes*. New York: Chapman & Hall.
- BERAN, J., BHANSALI, R. J., & OCKER, D. (1998). On unified model selection for stationary and non-stationary short- and long-memory autoregressive processes. *Biometrika*, **85**, 921-934.
- BROCKWELL, P. J., & DAVIS, R. A. (1996). *Introduction to time series and forecasting*. New York: Springer.
- CRATO, N., & RAY, B. K. (2002). Semi-parametric smoothing estimators for long-memory processes with added noise. *Journal of Statistical Planning & Inference*, **105**, 283-297.
- DOORNIK, J. A. (Ed.) (2001). *Ox: An object-oriented matrix language*. London: Timberlake Consultants.
- EKE, A., HERMAN, P., BASSINGTHWAIGHTE, J. B., RAYMOND, G. M., PERCIVAL, D. B., CANNON, M., ET AL. (2000). Physiological time series: Distinguishing fractal noises from motions. *European Journal of Physiology*, **439**, 403-415.
- FOX, R., & TAQQU, M. S. (1986). Large-sample properties of parameter estimates for strongly dependent stationary Gaussian time series. *Annals of Statistics*, **14**, 517-532.
- GILDEN, D. L. (1997). Fluctuations in the time required for elementary decisions. *Psychological Science*, **8**, 296-301.
- GILDEN, D. L. (2001). Cognitive emissions of $1/f$ noise. *Psychological Review*, **108**, 33-56.
- GILDEN, D. L., & WILSON, S. G. (1995). Streaks in skilled performance. *Psychonomic Bulletin & Review*, **2**, 260-265.
- HASLETT, J., & RAFTERY, A. E. (1989). Space-time modelling with long-memory dependence: Assessing Ireland's wind power resource. *Journal of the Royal Statistical Society*, **38**, 1-50.
- HAUSER, M. A. (1999). Maximum likelihood estimators for ARMA and ARFIMA models: A Monte Carlo study. *Journal of Statistical Planning & Inference*, **80**, 229-255.
- HEATHCOTE, A., BROWN, S., & MEWHORT, D. J. K. (2000). The power law repealed: The case for an exponential law of practice. *Psychonomic Bulletin & Review*, **7**, 185-207.
- HSU, N. J., & BREIDT, F. J. (2003). Bayesian analysis of fractionally integrated ARMA with additive noise. *Journal of Forecasting*, **22**, 491-514.
- LOGAN, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, **95**, 492-527.
- MAECHLER, M. (2005). *fracdiff: Fractionally differenced ARIMA(p, d, q) models*. Vienna: R Foundation for Statistical Computing. Available at www.R-project.org.
- NEWELL, A., & ROSENBLUM, P. S. (1981). Mechanisms of skill acquisition and the law of practice. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition* (pp. 1-51). Hillsdale, NJ: Erlbaum.
- NEWELL, K. M. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- OHLSSON, S. (1992). The learning curve for writing books: Evidence from Asimov. *Psychological Science*, **3**, 380-382.
- PAI, J. S., & RAVISHANKER, N. (1998). Bayesian analysis of autoregressive fractionally integrated moving-average processes. *Journal of Time Series Analysis*, **19**, 99-112.
- RITTER, F. E., & SCHOOLER, L. J. (2001). The learning curve. In W. Kintsch, N. Smelser, & P. Baltes (Eds.), *International encyclopedia of the social and behavioral sciences* (pp. 8602-8605). Oxford: Pergamon.
- SOWELL, F. (1992). Maximum likelihood estimation of stationary univariate fractionally integrated time series models. *Journal of Econometrics*, **53**, 165-188.
- THORNTON, T. L., & GILDEN, D. L. (2005). Provenance of correlations in psychological data. *Psychonomic Bulletin & Review*, **12**, 409-441.
- TORRE, K., DELIGNIERES, D., & LEMOINE, L. (in press). Detection of long-range dependence and estimation of fractal exponents through ARFIMA modelling. *British Journal of Mathematical & Statistical Psychology*.
- VAN ORDEN, G. C., HOLDEN, J. G., & TURVEY, M. T. (2003). Self-organization of cognitive performance. *Journal of Experimental Psychology: General*, **132**, 331-350.
- VAN ORDEN, G. C., HOLDEN, J. G., & TURVEY, M. T. (2005). Human cognition and $1/f$ scaling. *Journal of Experimental Psychology: General*, **134**, 117-123.
- WAGENMAKERS, E.-J., FARRELL, S., & RATCLIFF, R. (2004). Estimation and interpretation of $1/f^\alpha$ noise in human cognition. *Psychonomic Bulletin & Review*, **11**, 579-615.
- WAGENMAKERS, E.-J., FARRELL, S., & RATCLIFF, R. (2005). Human cognition and a pile of sand: A discussion on serial correlations and self-organized criticality. *Journal of Experimental Psychology: General*, **134**, 108-116.
- WAGENMAKERS, E.-J., GRÜNWALD, P., & STEYVERS, M. (2006). Accumulative prediction error and the selection of time series models. *Journal of Mathematical Psychology*, **50**, 149-166.

NOTES

1. Throughout their article, Thornton and Gilden (2005) variously refer to psychological (e.g., p. 430) and psychophysical (e.g., p. 429) time series. Here, we discuss the claim applying to psychological time series.
2. We note that the noise considered by Thornton and Gilden (2005) is not technically fractional Brownian motion and is better treated as fractional Gaussian noise (fGn); fBm is obtained by integrating fGn, has a spectral slope steeper than -1 , and is nonstationary (Eke et al., 2000).
3. Recent evidence suggests that the power function is actually a combination of exponentials (Heathcote, Brown, & Mewhort, 2000).

(Manuscript received September 6, 2005;
revision accepted for publication January 3, 2006.)