

Ambiguity and nonidentifiability in the statistical analysis of neural codes

Asohan Amarasingham^{a,b,1,2}, Stuart Geman^{c,2}, and Matthew T. Harrison^{c,1}

^aDepartment of Mathematics, The City College of New York, The City University of New York, New York, NY 10031; ^bDepartments of Biology and Psychology and CUNY Neuroscience Collaborative, The Graduate Center, The City University of New York, New York, NY 10016; and ^cDivision of Applied Mathematics, Brown University, Providence, RI 02912

Contributed by Stuart Geman, April 3, 2015 (sent for review January 2, 2014; reviewed by Moshe Abeles)

Many experimental studies of neural coding rely on a statistical interpretation of the theoretical notion of the rate at which a neuron fires spikes. For example, neuroscientists often ask, “Does a population of neurons exhibit more synchronous spiking than one would expect from the covariability of their instantaneous firing rates?” For another example, “How much of a neuron’s observed spiking variability is caused by the variability of its instantaneous firing rate, and how much is caused by spike timing variability?” However, a neuron’s theoretical firing rate is not necessarily well-defined. Consequently, neuroscientific questions involving the theoretical firing rate do not have a meaning in isolation but can only be interpreted in light of additional statistical modeling choices. Ignoring this ambiguity can lead to inconsistent reasoning or wayward conclusions. We illustrate these issues with examples drawn from the neural-coding literature.

temporal coding | spike timing | spike count variability | trial-to-trial variability | doubly stochastic

Among the most important open questions in neurophysiology are those regarding the nature of the code that neurons use to transmit information. Experimental approaches to such questions are challenging, because the spike outputs of a neuronal subpopulation, as typically recorded in behaving animals, are influenced by a vast array of factors. Such factors span all levels of description, from the microscopic (e.g., ion fluctuations and states of presynaptic neurons) to the macroscopic (e.g., sensation and attention), but only a small fraction of these is measured, or even understood. As a consequence, it is not clear to what degree variations in unknown and uncontrolled variables alternately reveal or confound the underlying signals that observed spikes are presumed to encode. Another consequence, very much related, is that these uncertainties also disturb our intuitive comfort with common models of statistical repeatability in neurophysiological signal analysis. In this context, there is an increasingly popular strain of thought in the neural-coding literature that “doubly stochastic point processes” (1–8) provide a way to think about and model fundamental questions about the relationship between sources of “trial-to-trial variability” (4, 9) and the observed variability of spike responses.

Imagine an experiment consisting of repeated presentations of a sensory stimulus. In a typical probabilistic model of spike responses, the probability that a particular neuron emits (fires) a spike at one time is described by a theoretical firing rate function (a function of time relative to stimulus onset). [Note that theoretical firing rates are distinct from the basic “observed” or “empirical” firing rate, which is a report of how many spikes occur in a window of a specified time length. Here, we discuss the theoretical firing rate. Observed firing rates are used to infer theoretical firing rates or their properties (*SI Appendix, section S1*)]. For example, in a generic firing rate model, spikes are treated as completely random, beyond the structure induced by the time-varying firing rate function. An analogy can be made with coin tossing: every neuron is associated with a sequence of weighted coins; each coin is associated with a moment in time, representing the absence or presence of a spike. Trials consist of

tossing the coins. In a doubly stochastic spiking model, there is another level of complexity in that the firing rate function is itself random and varies from trial to trial (illustrated in *SI Appendix, Fig. S1*).

Thus, in such doubly stochastic models, there are “two layers of variability” (1). For each trial and neuron, a firing rate function (analogously, a sequence of coins) is randomly chosen. The variation in the firing rate function across trials is the first layer of variability. The firing rate function, in turn, generates the observed spike train (analogously, the coins are tossed), contributing the second layer of variability. Importantly, there is a mental affinity—never quite pinned down, perhaps, but nevertheless influential—between rate-coding and firing rate functions. In this view, the firing rate function is the “rate” of the “rate code.” For example, the affinity is sometimes exploited in experimental attempts to contrast rate codes from temporal codes, in which precise temporal structures or dependencies in spike patterns play a role in the neural code.

It is then natural to ask if can we use statistics to relate this probabilistic picture of two-layered variability to real, experimentally observed spike trains. Answers might then be related to questions about neural coding. In the words of Churchland and Abbott, “Experimentalists often attempt to segregate response variability into firing rate variability and spiking noise. It is generally assumed that the former can influence behavior or perception, where the latter effectively acts as measurement noise” (1).

This image is seductive (generative, easy to picture, and easy to simulate) but also potentially misleading, especially when we turn to inference. The root of the problem is that the definition of firing rate functions in the image that we have sketched—essentially, that a firing rate function specifies the probability of

Significance

Among the most important open questions in neurophysiology are those regarding the nature of the code that neurons use to transmit information. Experimental approaches to such questions are challenging because the spike outputs of a neuronal subpopulation are influenced by a vast array of factors, ranging from microscopic to macroscopic scales, but only a small fraction of these is measured. Inevitably, there is variability from trial to trial in the recorded data. We show that a prominent conceptual approach to modeling spike-train variability can be ill-posed, confusing the interpretation of results bearing on neural codes. We argue for more careful definitions and more explicit statements of physiological assumptions.

Author contributions: A.A., S.G., and M.T.H. conceived the ideas; A.A. and M.T.H. performed research; and A.A. and M.T.H. wrote the paper.

Reviewers included: M.A., Bar-Ilan University.

The authors declare no conflict of interest.

Freely available online through the PNAS open access option.

¹A.A. and M.T.H. contributed equally to this work.

²To whom correspondence may be addressed. Email: amarasingham@ccny.cuny.edu or stuart_geman@brown.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1506400112/-DCSupplemental.

a spike—is underconstrained, and therefore, firing rate functions are usually not unique. The six columns in Fig. 1 provide an illustration in which six strikingly different firing rate processes give rise to spike train processes which are, in fact, identical (i.e., statistically indistinguishable). This example is not singular. In fact, as exemplified by the sixth random firing rate process in Fig. 1, every (discrete time) spiking process can be perfectly described by a doubly stochastic model that attributes all variability to firing rate variability. Additional modeling assumptions would be needed to resolve these ambiguities. For example, we may require firing rates to be slowly varying in time, prohibit the firing rates from being too high at any instant, or invoke other kinds of statistical constraints on the firing rate process. Modeling assumptions of some kind are necessary to make reasonable use of the doubly stochastic concept.

Consider typical neurophysiological settings in which a theoretical firing rate is invoked. For example, do a pair of neurons exhibit more synchronous firing than one would expect from their (theoretical) firing rates? For another example, how much of a neuron's observed spiking variability can be attributed to the variability of the (theoretical) firing rates, and how much can be attributed to spike timing variability? Theoretical firing rates specify the probability of a spike in a neuron. However, the probability of an event is not meaningful unless conditions of observation are specified (10–13). Rényi himself found it necessary to emphasize this: “Every probability is in reality a conditional probability” (10). This observation might appear elementary—indeed, Rényi added: “[t]his evident fact is somewhat obscured by the practice of omitting the explicit statement of the conditions if it is clear under what conditions the probability of an

event is considered” (10)—but in the neural-coding literature for various reasons such conditions may not be specified. This state of affairs causes at least two kinds of serious problems. In the first kind, different investigators both use the term firing rate, but in reality, they have different conceptions of firing rate in mind and are not discussing comparable quantities. In the second kind, an investigator assumes that the firing rate in a doubly stochastic model is unique or well-defined independent of a precise account of conditioning and then, using apparently ordinary statistical reasoning, draws conclusions about neural coding. This oversight can demonstrably lead to mathematical inconsistencies. A recurring problem stems from this basic fact that firing rate functions in doubly stochastic models may not be unique and, in any case, should be defined with more care.

Our aim here is to illustrate these ambiguities in the context of typical and recent neurophysiological questions. We choose these particular examples entirely for the sake of concreteness; the general issues are widespread in the literature of neural coding, if often unarticulated (14, 15). The manuscript is organized as follows. Our main point is developed in Fig. 1, where the essential ambiguity of inference is related to the statistical concept of identifiability. In *Trial-to-Trial Variability and Doubly Stochastic Decompositions*, we show how this ambiguity is present in, and dooms, prototypical model-free attempts to interpret the contribution of trial-to-trial firing rate variability to spike count and spike timing variability. In *Implicit Modeling Assumptions and Solutions*, we point out that closely related issues concerning modeling assumptions are pervasive in neurophysiology, and we provide a brief overview of potential solutions in this context. We summarize our conclusions in *Summary and Discussion*. Technical details are

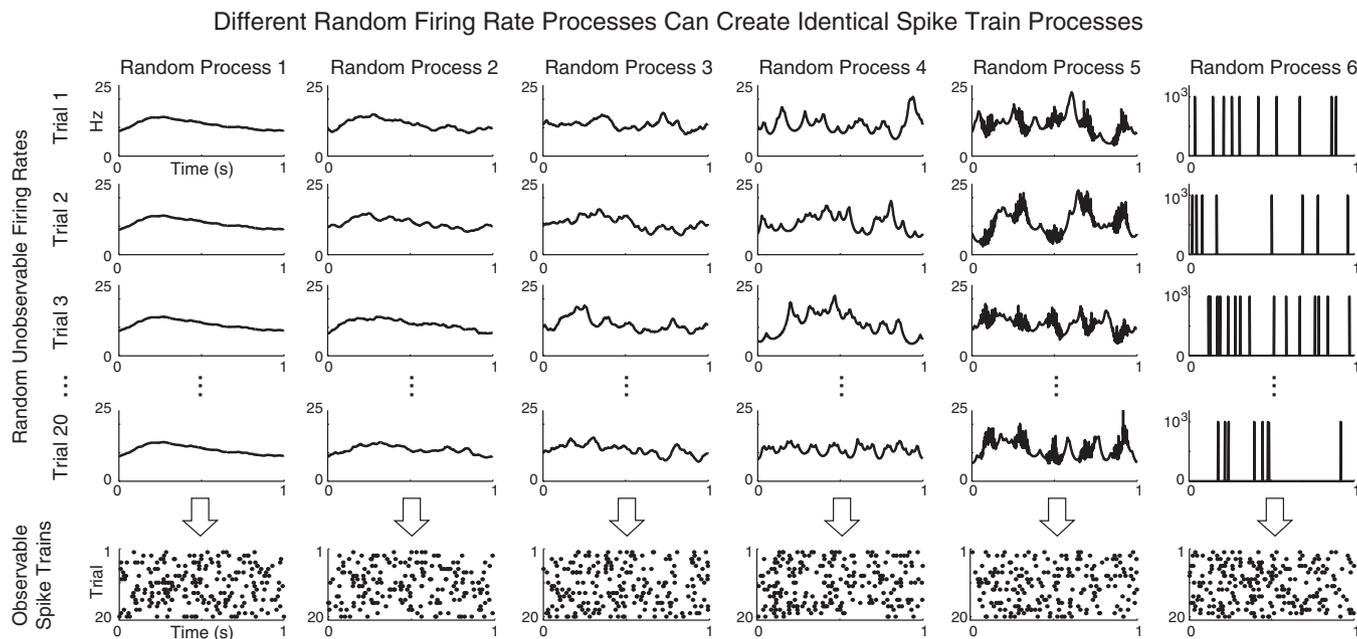


Fig. 1. Random firing rates are not identifiable. Each of six columns shows a different way to generate the same probabilistic spike train process using different versions of a hidden, random firing rate. In each case, a firing rate function is selected randomly on each of 20 trials. The selected functions for trials 1, 2, 3, and 20 are shown. (Here, the firing rates are plotted in units of spikes per second.) Then, the spike times for that trial are generated according to the selected firing rate function. The corresponding spike times (rasters) for all 20 trials are shown at the bottom of each column. Although each of the six processes has a different distribution over firing rate functions, they all generate the identical distribution over observable spike trains. Thus, the spike trains observed in different columns are statistically indistinguishable; any apparent differences are illusory (entirely caused by sampling variability). As a consequence, the evident distinctions in the firing rate processes across the six examples are not identifiable from spike data: there is no correct or true version of the theoretical firing rate. In essence, the different cases merely reflect modeling choices affecting what proportion of total variability we wish to attribute to the hidden firing rate process. Random process 1 attributes none of the variability to the firing rates. Random process 6 attributes all of the variability to the firing rates. There are infinitely many alternatives in between these two extremes, which are illustrated by processes 2–5. Processes 4–6 satisfy the additional constraint that spikes in different time bins are conditionally independent, given knowledge of the firing rate function (as in *SI Appendix, Fig. S1*). Even this strong modeling assumption, which is also reminiscent of rate coding, is still too weak to ensure that firing rate functions are uniquely specified. Processes 1–3 have history dependence in the spiking process in the sense that spikes in different time bins are not conditionally independent, given the firing rate function. Complete details for the simulations are provided in *SI Appendix, section S10*.

provided in *SI Appendix*. Our chief conclusion is that the assumptions justifying statistical models should be prominently described and scrutinized when applied to the quantification of neural-coding questions, particularly in regards to trial-to-trial variability and related phenomena.

Trial-to-Trial Variability and Doubly Stochastic Decompositions

Here, we will first develop definitions and notation to facilitate a precise discussion of theoretical firing rates. The key property that a theoretical firing rate clearly must have is that the expected number of spikes per unit time is equal to the theoretical firing rate when the theoretical firing rate is known. Mathematically, we can write this as

$$E[\text{rate of spiking}|\text{firing rate}] = \text{firing rate},$$

where E denotes expected value. (We will now typically use firing rate in place of theoretical firing rate.) More formally, if the term N denotes the observed data (such as the number of spikes per time bin) and λ denotes the corresponding (theoretical) firing rate, then we must have

$$E[N|\lambda] = \lambda. \quad [1]$$

This equation is read: “the expected value of N , given knowledge of λ , is itself λ .” Eq. 1 simply states that λ is a firing rate for N .

For example, if we know that the firing rate in a 1-s interval is uniformly 10 Hz, so that $\lambda = 10$, then the expected number of spikes in that interval given that knowledge, $E[N|\lambda]$, should be 10. The actual number of spikes may differ from 10 on a particular trial, but if we repeat the experiment enough, always with a 10-Hz firing rate, then the average number of spikes in the interval, averaged across the repetitions, will approach 10.

Allowing N and λ to further depend on time does not change this basic relationship. We simply have that $E[N(t)|\lambda(t)] = \lambda(t)$, where t indexes time.

Complexities involving trial-to-trial variability, as described in the Introduction, often motivate modeling λ as random, in which case the pair (N, λ) becomes a doubly stochastic model: both N and λ are random variables (variables for which values are specified by the outcome of a random experiment). λ itself is not observed. It is sometimes treated, for example, as the signal of interest (“the quantity that the neuron is supposed to represent”) (3).

When λ is random, then it also has an expected value $E[\lambda]$ and variance $\text{Var}[\lambda]$ [the latter quantity denotes the variability in the (random) firing rate]. Furthermore, when λ is random, quantities like $E[N|\lambda]$ and $\text{Var}[N|\lambda]$ become random as well, because they inherit the randomness in λ . Because $\text{Var}[N|\lambda]$ is a random variable, it also has an expected value, $E[\text{Var}[N|\lambda]]$, which quantifies the average conditional variability in spiking, given the firing rate. *SI Appendix, section S2* reviews the notation regarding expectation $E[\cdot]$, variance $\text{Var}[\cdot]$, conditional expectation $E[\cdot|\cdot]$, and conditional variance $\text{Var}[\cdot|\cdot]$.

Ambiguity of the Firing Rate. There are many random variables that satisfy Eq. 1. This fact immediately dooms model-free attempts to estimate firing rates from experimental data. For example, given any random variable R (or roughly, any piece of information), the random variable λ defined by

$$\lambda = E[N|R] \quad [2]$$

will satisfy Eq. 1. (In fact, the following definition is equivalent to Eq. 1: λ is a firing rate if $\lambda = E[N|R]$ for some R .) It is, thus, clear that, by considering different random variables R , a spectrum of different theoretical firing rates can be obtained in general. For example, in one traditional experimental formulation, R is simply taken to specify time since the onset of an experimental trial. Alternatively, R can draw from an essentially limitless combination

of measured or unmeasured information sources. These sources range from relatively coarse, such as attributes of the experimental stimulus or the spiking activity of a simultaneously measured population of neurons (16–18), to finely-detailed, such as extracellular ion concentrations or somatic voltage. Consequently, there are many firing rates (Fig. 1). Furthermore, strikingly distinct firing rates can give rise to the same observed data. Thus, the observed data cannot be used to distinguish between candidate firing rate models. In statistical parlance, we say that such models are non-identifiable (a formal definition is in *SI Appendix, section S3*): there is no correct or true version of the theoretical firing rate.

Essentially, the way to address this ambiguity is to introduce constraints on λ . However, not all constraints will work. For example, consider the spike time series $N(t)$ (a sequence of spike counts or a binary spike train) as a time series with the associated firing rate $\lambda(t)$ [satisfying $E[N(t)|\lambda(t)] = \lambda(t)$]. Then, $\lambda(t)$ is ambiguous in the sense discussed above (Fig. 1). An additional constraint, which is reminiscent of certain descriptions of rate coding, treats $N(t)$ as conditionally independent across time (8) given the time series $\lambda(t)$, but this constraint is also too weak. Processes 4–6 in Fig. 1 illustrate this point. Another way to look at this is as follows: consider a discrete spike train of length T bins and discretize probability into $M > 2$ levels. A distribution on $\lambda(1), \lambda(2), \dots, \lambda(T)$ is specified by $M^T - 1$ parameters. Moreover, the distribution on $\lambda(1), \lambda(2), \dots, \lambda(T)$ (continuously) determines the distribution on $N(1), N(2), \dots, N(T)$, which is characterized by $2^T - 1$ parameters. This observation alone implies non-identifiability and suggests its pervasiveness.

We treat more subtle characteristics of such ambiguities in detail below.

Spike Count Variability. Consider the following interesting statistics problem (related, in varying degrees, to refs. 3, 7, and 19–21). Fix a recording interval (of, for example, 100 ms, with some start time specified relative to a stimulus onset), according to which we record multiple observations of spike counts, N_1, N_2, \dots, N_m , from a neuron corresponding to m trials of a repeated experiment. For each of the m trials, there is a separate firing rate (i.e., λ_i is the firing rate on trial i). Trial-to-trial variability is here modeled by the fact that $\lambda_1, \lambda_2, \dots, \lambda_m$ are not necessarily identical. A typical neural-coding intuition is that of rate coding: λ_i represents an “intensity command” or “the quantity that the neuron is supposed to represent” in trial i (3). In practice, one only observes N_1, N_2, \dots, N_m , and it is variable. How much of the observed variability is because of variability in the “intensity command” (3) (firing rate)?

To formalize this question, suppose that N and λ are the spike count and firing rate, respectively, in a randomly chosen or future trial. The variability in N can be decomposed as follows (2, 3, 7, 19):

$$\text{Var}[N] = \text{Var}[E[N|\lambda]] + E[\text{Var}[N|\lambda]] \quad [3]$$

$$= \underbrace{\text{Var}[\lambda]}_{\text{firing rate variability}} + E[\underbrace{\text{Var}[N|\lambda]}_{\text{spiking variability}}], \quad [4]$$

where Eq. 3 is a standard decomposition valid for any pair of finite random variables, and Eq. 4 follows from the definition that λ is a firing rate for N (Eq. 1). It is natural to ask how we can use observations of N to distinguish the two variability terms $\text{Var}[\lambda]$ (“firing rate variability”) and $E[\text{Var}[N|\lambda]]$ (“point process variability”) (3, 7, 19).

To consider this last question, suppose now that we return to our observations of N_1, \dots, N_m in m trials and use these to estimate $\text{Var}[N]$ in the usual manner. In practice, some approximations are made to also estimate $E[\text{Var}[N|\lambda]]$ from N_1, \dots, N_m . This estimate can then be used (3) to infer $\text{Var}[\lambda]$ using Eq. 4. Even without examining those approximations directly, it is clarifying to take a broader perspective on the statistical model underlying this decomposition. The model is fairly minimal: N_1, \dots, N_m and

$\lambda_1, \dots, \lambda_m$ are modeled as random variables, which have a joint probability distribution. In principle, what can we learn about $\text{Var}[\lambda]$ and $E[\text{Var}[N|\lambda]]$ from the data N_1, \dots, N_m ? Put aside the question of sampling variability: given complete knowledge about the probability distribution on N , what could one learn about the probability distribution generating $\lambda_1, \dots, \lambda_m$? Essentially nothing. Consider the following two elementary cases of generative models for the λ_i and the spike counts. In both, c is a constant.

Case I. $\lambda_1, \lambda_2, \dots, \lambda_m$ are independent and identically distributed (iid) Poisson random variables with parameter c . $N_i = \lambda_i$. Thus, N_1, N_2, \dots, N_m are iid Poisson(c). Under this generative model, $\text{Var}[\lambda] = c$, and $E[\text{Var}[N|\lambda]] = 0$.

Case II. $\lambda_1 = \lambda_2 = \dots = \lambda_m = c$. (The λ_i values are a constant.) Each N_i is an independent Poisson random variable with parameter λ_i . Thus, as in case I, N_1, N_2, \dots, N_m are iid Poisson(c). Therefore, under this generative model, $\text{Var}[\lambda] = 0$ and $E[\text{Var}[N|\lambda]] = c$.

In both cases I and II, N_1, \dots, N_m has exactly the same spiking distribution. Thus, the underlying variance decomposition is nonidentifiable: no use of the data (i.e., N_1, \dots, N_m) could possibly distinguish $\text{Var}[\lambda]$ from $E[\text{Var}[N|\lambda]]$, regardless of how much data we collect. [Mixing cases I and II produces intermediate models that are equally indistinguishable. In fact, our use of a Poisson example is just a pedagogical choice; the non-identifiability phenomenon is completely general (SI Appendix, section S4).] The problem could only be worsened if we introduce sampling variability, such as when the (probability) distribution of N is inferred from the data.

Additional constraints are needed. For example, refs. 3 and 4 invoke another modeling assumption that $\text{Var}[N|\lambda] = \phi\lambda$, where ϕ is an unknown constant parameter that does not vary with λ or time. Ref. 3 includes an extensive discussion of this assumption as well as its motivation from the model of rate-modulated renewal processes, in which spike trains are the outcomes of a renewal process that has been rescaled in time (SI Appendix, section S5). However, this strong assumption is not strong enough to ensure identifiability, which the counterexamples show, because under this model, ϕ is itself not identifiable (despite the fact that it is crucial to conclusions drawn about $\text{Var}[\lambda]$) (3). [Note, for example, that ϕ is zero and one, respectively, in cases I and II above (additional discussion of this nonidentifiability of ϕ and the non-identifiability of the constancy of $\text{Var}[N|\lambda]/\lambda$ is in SI Appendix, section S6).] Indeed, the rate-modulated renewal process equipped with a random rate function can generate any discrete time spike process, even when ϕ is restricted to zero (SI Appendix, section S5). This observation alone implies the nonidentifiability of ϕ and theoretical firing rates in general. The supplemental information in ref. 3 further describes how different assumptions about the numerical value of ϕ can influence scientific results.

However, if ϕ is known and this assumption holds a priori, then $\text{Var}[\lambda]$ is identifiable, and it is straightforward to estimate from the spike counts (a mathematical explanation is in SI Appendix, section S7). In terms of scientific considerations, the modeling assumption and a priori knowledge of ϕ are the key assumptions to be debated.

Indeed, if ϕ is not known or $\text{Var}[N|\lambda]/\lambda$ is not a constant, $\text{Var}[\lambda]$ and $E[\text{Var}[N|\lambda]]$ are entangled, and their distinction does not even quite make sense.

Spike Train Variability and Cross-Correlogram Statistics. Another prominent case is synchrony analysis and other analyses based on the cross-correlogram. Is it possible to disambiguate the relative contributions of firing rate variability and spike timing variability to the observed synchrony between a pair of neurons? This question arises naturally from experimental attempts to distinguish rate-coding from more complex temporal-coding hypotheses.

Consider the following generative prescription for a doubly stochastic spike train model. First, in each trial, a firing rate

function is randomly chosen. As above, the randomness of the firing rate function is a mechanism for modeling complexities of trial-to-trial variability. Second, spikes are randomly generated, with probabilities governed, in an appropriate sense, by the associated firing rate function. Then, we can ask (2, 3, 7) whether the variability of the firing rate functions can be separated from the variability in spike timing, given the firing rate. If so, the goal is to quantify the contribution of each component of variability to the observed synchrony (or time-lagged synchrony).

Suppose that we model time as discrete so that we are dealing with (otherwise arbitrary) binary time series. Let N and λ refer to the observed spiking and firing rate, respectively, of one neuron in a (e.g., random) time bin, and let M and γ refer to the observed spiking and firing rate of a second neuron, possibly lagged by a fixed amount τ . Note that N and M are now binary random variables signifying the presence or absence of a spike in a time bin, in neurons one and two, respectively. [It is also possible to include time directly in the equations so that these quantities are related to an experimental clock, but our conclusion—non-identifiability—will be the same (SI Appendix, section S8).] We can study the covariability between N and M using the covariance $\text{Cov}[N, M]$, which measures the covariability in firing for two neurons (2, 3) (SI Appendix, section S2).

Firing rates λ and γ are hidden (“latent”) random variables that have the property that knowledge of the firing rate function specifies the likelihood of a spike:

$$\begin{aligned} E[N|\lambda, \gamma] &= E[N|\lambda] = \lambda \\ E[M|\lambda, \gamma] &= E[M|\gamma] = \gamma. \end{aligned} \tag{5}$$

Verbally, (i) given knowledge of a neuron’s firing rate function, the knowledge of other neurons’ firing rate functions is irrelevant to the (conditional) probability of a spike, and (ii) the (conditional) probability of a spike in neuron 1 or 2 at a randomly selected moment, given knowledge of λ or γ , respectively, is λ or γ , respectively.

A covariance decomposition can then be applied:

$$\text{Cov}[N, M] = \text{Cov}[E[N|\lambda, \gamma], E[M|\lambda, \gamma]] + E[\text{Cov}[N, M|\lambda, \gamma]] \tag{6}$$

$$= \underbrace{\text{Cov}[\lambda, \gamma]}_{\text{firing rate covariability}} + E[\underbrace{\text{Cov}[N, M|\lambda, \gamma]}_{\text{spiking covariability (spike coordination)}}], \tag{7}$$

where Eq. 6 is the analogous covariance decomposition valid for any pair of finite random variables, and Eq. 7 follows from definitions in Eq. 5. The reasoning, now based on covariance, is the same as in the previous section (compare Eq. 7 with Eq. 4) For example, Staude et al. (2) conclude: “The decomposition achieved [i.e., Eq. 7] shows that spike coordination and rate covariation are mathematically distinguishable components of the cross-correlation function.”

In other words, the setup is just as in the previous section, except that, now, N and M indicate the event of a spike in a time bin for each neuron, and λ and γ are the associated (random) probabilities of a spike in the bin. By the same logic, this decomposition suggests the strategy (2) of quantifying the “spike coordination” component of variability by subtracting an estimate of “firing rate covariability” from an estimate of $\text{Cov}[N, M]$.

However, once again and for the same reasons, the fundamental difficulty is that (λ, γ) is not well-defined in terms of the random variables N and M . In particular, for any pair of (discrete-time) random spike trains, there is always an interpretation of firing rates available that is (i) completely consistent with the definitions in Eq. 5, but (ii) under which, there is no spike coordination in the sense of Eq. 7. More generally, any population of (discrete time) random spike trains can be represented as a doubly stochastic model, in which spikes are conditionally independent, given their firing rate functions. (This fact is the implication of the example of

some of the problems that arise when insufficient attention is devoted to the question. These problems include inconsistent statistical reasoning, an incomplete understanding of quantitative conclusions regarding trial-to-trial variability, and even nonidentifiability.

Most likely, the source of these problems is that we are collectively misled by intuitions that are overgeneralized from simple generative models. For example, consider again the example of defining spike coordination. Firing rate functions have a natural definition for inhomogeneous Poisson processes. In that setup, it is natural to define spike coordination in terms of violations of independence beyond the firing rates. However, intuitions about trial-to-trial variability reasonably challenge Poisson-like assumptions (6), leading one to further parameterize the firing rate function with trial-varying parameters. Then, one follows the same approach of associating departures from independence with spike coordination. However, the analysis is delicate, because the probabilistic structure that we evoke to model trial-to-trial variability will have significant implications for what we mean by spike coordination. We cannot avoid such modeling assumptions, although they may be implicit (see *Implicit Modeling Assumptions*). This fact is most exposed by the extreme observation that, if we simply generalize the Poisson-like assumptions to accommodate arbitrary trial-to-trial variability, then spike coordination and trial-to-trial variability, in such a model, are conflated.

Ambiguous language might also bear some of this blame: many things are called a firing rate in the neurophysiology literature (23). For context, recall the spike count measurements discussed in *Trial-to-Trial Variability and Doubly Stochastic Decompositions, Ambiguity of the Firing Rate*, where (λ_i, N_i) forms a doubly stochastic model producing spike counts N_1, N_2, \dots, N_m across m trials. Consider that (i) the firing rate of the process is $E[N_i]$ vs. (ii) the firing rate is $(1/m)\sum_{i=1}^m N_i$. Or, accommodating trial-to-trial variability, consider that the firing rate is defined individually on each trial: (i) the firing rate on trial i is N_i vs. (ii) the firing rate

on trial i is $E[N_i|\lambda_i]$ vs. (iii) the firing rate on trial i is $E[N_i]$. All four definitions of firing rate are well-defined in any of the generative models discussed in *Trial-to-Trial Variability and Doubly Stochastic Decompositions, Ambiguity of the Firing Rate* and have relationships to the different ways in which the term firing rate is used in the literature. Some correspond to random variables, and some correspond to parameters. However, they are completely different. Perhaps different notions of firing rate are vaguely or inadvertently interchanged, or simply misinterpreted.

What is more, perhaps, even the relationship between firing rate functions and firing rate codes, although arguably a principal justification for firing rate analysis in the literature, is a cloudy one (5, 39–41). (The neural-coding schemes proposed in refs. 42–44, for instance, are counterexamples to a traditional view of the relationship between a conventional firing rate code and noise in neural computation.) We suspect that such analyses are often motivated by buried but arguably underjustified assumptions regarding exactly how neurons are representing information to downstream targets.

In terms of neural coding, another way to look at this is as follows. There is no a priori separation between a signal and noise. Absent a model, the difference is semantic. It follows that variability decompositions cannot solely be based on mathematical, as opposed to neuroscientific, distinctions.

ACKNOWLEDGMENTS. We thank B. Anderson, D. Butts, G. Buzsáki, P. Cloud, K. Diba, S. Fujisawa, A. Groszmark, J. Levitt, R. Kass, S. Kelly, W. Levy, G. Maimon, M. Paradiso, A. Renart, J. de la Rocha, E. Schomburg, L. Sjulson, W. Truccolo, and B. Watson, as well as M. Abeles and an anonymous referee for helpful comments on earlier versions of this manuscript. We thank Chip Levy for encouraging our collaboration. Research reported in this publication was based, in part, on work supported by NIH Grant R01MH102840 and NIH Institutional Development Award P20GM103645, a City College of New York City SEED Grant, National Science Foundation Grants DMS-1007593 and DMS-1309004, Office of Naval Research Contract ONR N000141010933, and Defense Advanced Research Projects Agency Contract FA8650-11-1-7151.

- Churchland MM, Abbott LF (2012) Two layers of neural variability. *Nat Neurosci* 15(11):1472–1474.
- Stauder B, Rotter S, Grün S (2008) Can spike coordination be differentiated from rate covariation? *Neural Comput* 20(8):1973–1999.
- Churchland AK, et al. (2011) Variance as a signature of neural computations during decision making. *Neuron* 69(4):818–831.
- Churchland MM, et al. (2010) Stimulus onset quenches neural variability: A widespread cortical phenomenon. *Nat Neurosci* 13(3):369–378.
- Amarasingham A, Chen T-L, Geman S, Harrison MT, Sheinberg DL (2006) Spike count reliability and the Poisson hypothesis. *J Neurosci* 26(3):801–809.
- Brody CD (1999) Correlations without synchrony. *Neural Comput* 11(7):1537–1551.
- Goris RLT, Movshon JA, Simoncelli EP (2014) Partitioning neuronal variability. *Nat Neurosci* 17(6):858–865.
- Grün S (2009) Data-driven significance estimation for precise spike correlation. *J Neurophysiol* 101(3):1126–1140.
- Faisal AA, Selen LP, Wolpert DM (2008) Noise in the nervous system. *Nat Rev Neurosci* 9(4):292–303.
- Rényi A (1970) *Foundations of Probability* (Holden-Day Inc., San Francisco).
- Keynes JM (1921) *A Treatise on Probability* (Macmillan and Co., London).
- Jeffreys H (1939) *Theory of Probability* (Oxford Univ Press, London).
- Jaynes ET (2003) *Probability Theory: The Logic of Science* (Cambridge Univ Press, Cambridge, UK).
- Cohen MR, Kohn A (2011) Measuring and interpreting neuronal correlations. *Nat Neurosci* 14(7):811–819.
- Nienborg H, Cohen MR, Cumming BG (2012) Decision-related activity in sensory neurons: Correlations among neurons and with behavior. *Annu Rev Neurosci* 35:463–483.
- Renart A, et al. (2010) The asynchronous state in cortical circuits. *Science* 327:5965587–590.
- Okun M, et al. (2012) Population rate dynamics and multineuron firing patterns in sensory cortex. *J Neurosci* 32(48):17108–17119.
- Ecker AS, et al. (2014) State dependence of noise correlations in macaque primary visual cortex. *Neuron* 82(1):235–248.
- Ecker AS, Tolias AS (2014) Is there signal in the noise? *Nat Neurosci* 17(6):750–751.
- Shadlen MN, Newsome WT (1998) The variable discharge of cortical neurons: Implications for connectivity, computation, and information coding. *J Neurosci* 18(10):3870–3896.
- Marcos E, et al. (2013) Neural variability in premotor cortex is modulated by trial history and predicts behavioral performance. *Neuron* 78(2):249–255.
- Ventura V, Cai C, Kass RE (2005) Trial-to-trial variability and its effect on time-varying dependency between two neurons. *J Neurophysiol* 94(4):2928–2939.
- Harrison MT, Amarasingham A, Kass RE (2013) Statistical identification of synchronous spiking. *Spike Timing: Mechanisms and Function*, Frontiers in Neuroscience, eds DiLorenzo PM, Victor JD (CRC, Boca Raton, FL), pp 77–120.
- Dayan P, Abbott LF (2001) *Theoretical Neuroscience* (MIT Press, Cambridge, MA).
- Rieke F, Warland D, van Steveninck R, Bialek W (1997) *Spikes: Exploring the Neural Code* (MIT Press, Cambridge, MA).
- Amarasingham A, Harrison MT, Hatsopoulos NG, Geman S (2012) Conditional modeling and the jitter method of spike resampling. *J Neurophysiol* 107(2):517–531.
- Cox DR (1955) Some statistical methods connected with series of events. *J R Stat Soc Series B* 17(2):129–164.
- Feller W (1943) On a general class of “contagious” distributions. *Ann Math Stat* 14(4):389–400.
- Maimon G, Assad JA (2009) Beyond Poisson: Increased spike-time regularity across primate parietal cortex. *Neuron* 62(3):426–440.
- Date A, Bienenstock E, Geman S (1998) *On the Temporal Resolution of Neural Activity. Technical Report* (Brown University Division of Applied Mathematics, Providence, RI).
- Harrison MT, Geman S (2009) A rate and history-preserving resampling algorithm for neural spike trains. *Neural Comput* 21(5):1244–1258.
- Harrison MT, Amarasingham A, Truccolo W (2015) Spatiotemporal conditional inference and hypothesis tests for neural ensemble spiking precision. *Neural Comput* 27(1):104–150.
- Smith MA, Kohn A (2008) Spatial and temporal scales of neuronal correlation in primary visual cortex. *J Neurosci* 28(48):12591–12603.
- Shmiel T, et al. (2006) Temporally precise cortical firing patterns are associated with distinct action segments. *J Neurophysiol* 96(5):2645–2652.
- Fujisawa S, Amarasingham A, Harrison MT, Buzsáki G (2008) Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nat Neurosci* 11(7):823–833.
- Hirabayashi T, Takeuchi D, Tamura K, Miyashita Y (2013) Functional microcircuit recruited during retrieval of object association memory in monkey perirhinal cortex. *Neuron* 77(1):192–203.
- Diba K, Amarasingham A, Mizuseki K, Buzsáki G (2014) Millisecond timescale synchrony among hippocampal neurons. *J Neurosci* 34(45):14984–14994.
- Truccolo W, et al. (2014) Neuronal ensemble synchrony during human focal seizures. *J Neurosci* 34(30):9927–9944.
- Barlow HB (1972) Single units and sensation: A neuron doctrine for perceptual psychology? *Perception* 1(4):371–394.
- Amarasingham A (2004) Statistical methods for the assessment of temporal structure in the activity of the nervous system. PhD thesis (Brown University, Providence, RI).
- Masquelier T (2013) Neural variability, or lack thereof. *Front Comput Neurosci* 7:7.
- Geman S (2006) Invariance and selectivity in the ventral visual pathway. *J Physiol Paris* 100(4):212–224.
- Abeles M (1991) *Corticonics: Neuronal Circuits of the Cerebral Cortex* (Cambridge Univ Press, Cambridge, UK).
- von der Malsburg C (1981) *The Correlation Theory of Brain Function. Technical Report*. Available at cogprints.org/1380/1/vdM_correlation.pdf. Accessed January 3, 2014.