

How You Named Your Child: Understanding the Relationship Between Individual Decision Making and Collective Outcomes

Todd M. Gureckis,^a Robert L. Goldstone^b

^a*Department of Psychology, New York University*

^b*Psychological and Brain Sciences, Indiana University*

Received 7 July 2008; received in revised form 12 August 2009; accepted 16 August 2009

Abstract

We examine the interdependence between individual and group behavior surrounding a somewhat arbitrary, real-world decision: selecting a name for one's child. Using a historical database of the names given to children over the last century in the United States, we find that naming choices are influenced by both the frequency of a name in the general population, and by its "momentum" in the recent past in the sense that names which are growing in popularity are preferentially chosen. This bias toward rising names is a recent phenomena: In the early part of the 20th century, increasing popularity of a name from one time period to the next correlated with a decrease in future popularity. However, more recently this trend has reversed. We evaluate a number of formal models that detail how individual decision-making strategies, played out in a large population of interacting agents, can explain these empirical observations. We argue that cognitive capacities for change detection, the encoding of frequency in memory, and biases toward novel or incongruous stimuli may interact with the behavior of other decision makers to determine the distribution and dynamics of cultural tokens such as names.

Keywords: Collective and individual choice behavior; Imitation; Social sampling; Decision making

1. Introduction

Psychologists, economists, profiteers, and politicians alike have long been interested in understanding how people make decisions. However, despite many advances in our understanding, we are still unable to explain many of the choices people make in their daily lives.

Correspondence should be sent to Todd M. Gureckis, Department of Psychology, New York University, 6 Washington Place, New York, NY 10003. E-mail: todd.gureckis@nyu.edu

How do people decide what type of music they like? Why do people prefer one political candidate to others? Although there are many potential answers to these kinds of questions, one thing that these decisions have in common is that even though they are (ostensibly) individual expressions of taste or preference, they are fundamentally linked to the behavior and decisions of others. There is little use in supporting a candidate that no one else does, just as you are unlikely to hear a local band on the radio until a great number of other people have heard their music as well. In these and many other naturally occurring contexts, the only way to meaningfully understand individual choice is to take seriously the interaction between those individuals and the groups in which they are embedded (Gureckis & Goldstone, 2006).

In this paper, we attempt to illuminate the relationship between individual and group behavior in a simple, real-world decision-making task: choosing a name for one's child. In addition to being a topic of fascination for many expecting parents, baby names provide a unique opportunity for studying the intersection of individual and group behavior. First, naming is an important real-life decision to which parents devote much time and energy. Second, given names are discrete tokens for which extensive historical records exist. This not only allows direct measurement of the actual choices that a large number of parents make but also an estimate of the social context in which those decisions were made (by considering the popularity of the name in the years leading up to any individual choice). Third, at least to a reasonable approximation, different names have similar intrinsic value (i.e., there is nothing particular to a name like *Joshua*, a common boy's name in 2007, compared with *Damarion*, a relatively uncommon boy's name in 2007, that would favor one over the other), making patterns of convergence and coordination in choice behavior all the more interesting (Ford, Mirua, & Masters, 1984; Fryer & Levitt, 2004; Hahn & Bentley, 2003). Finally, names may be unique in that they are not subject to the external marketing and advertising forces that complicate the analysis of collective behavior in other domains such as the Internet, the stock market, or fashion trends (Lieberman, 2000).

We will ultimately argue that the perceived value of a name is determined not by some intrinsic property of the name itself, but is rather an emergent property of the behavior of other parents who are themselves making naming decisions. In developing this argument, we present a number of novel analyses of naming behavior in the United States that give new insights into the changing dynamics and distribution of these cultural tokens. Most importantly, we show that, contrary to the predictions of existing formal models of cultural evolution (Bentley, Hahn, & Shennan, 2004; Hahn & Bentley, 2003; Xu et al., 2008), parents in the United States are increasingly sensitive to the *change* in frequency of a name in recent times, such that names that are gaining in popularity are seen as more desirable than those that have fallen in popularity in the recent past. This bias then becomes a self-fulfilling prophecy as names that are falling continue to fall, while names on the rise reach new heights of popularity, in turn influencing a new generation of decision makers. Through a number of formal analyses, we demonstrate how such dynamics might arise from an interaction between the cognitive decision strategies enacted by individual parents and the social environment in which the decision takes place (i.e., the naming choices of other parents). Our analysis shows that decision makers can be subtly influenced by the statistical patterns

in their environment (in this case, the frequency of names and changes in those frequencies) and suggest ways in which cognitive processes originating within the individual may contribute to, and even reinforce, the emergent dynamics of the group.

2. Aggregate naming patterns in the United States and what they reveal about individual decisions

We begin our analysis by considering some of the key aggregate statistical properties of naming behavior in the United States and what they reveal about individual name choice. One striking, but particularly revealing aspect of name choice is that there are large disparities in the prevalence of different names. This pattern can be most clearly illustrated as a frequency distribution where one counts the number of names that appear at a given frequency in the general population. For example, we can compare the number of names that appear with frequency one per million babies to the number that appear with frequency 10,000 per million babies. In our first analysis, we computed the frequency distribution of names in the U.S. population for each of the last 127 years based on records published by the United States Social Security Administration (SSA).¹ Fig. 1 displays the cumulative proportion of baby names in the top 1,000 that occur at a given frequency (normalized for population size) on a log–log scale for a number of selected years and for both male and female names.

The shape of this distribution reveals a considerable degree of convergence and coordination in the choice preferences of individual parents. For example, a large number of names

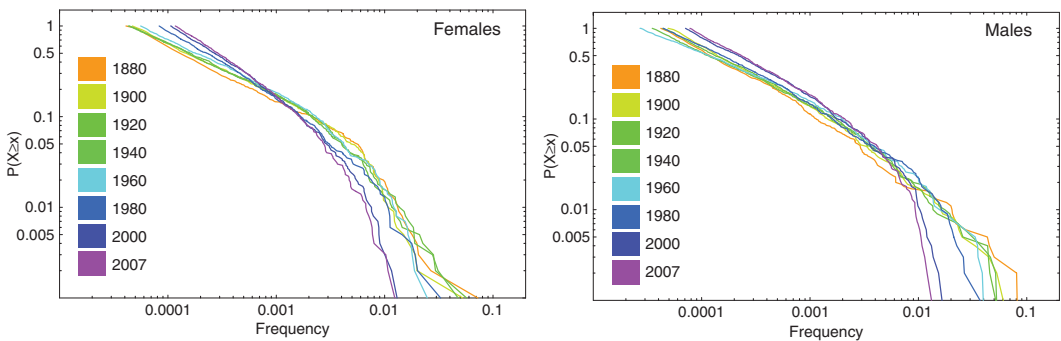


Fig. 1. A log–log plot of the cumulative frequency distribution for names (see Appendix for more details). Plotted is a selected subset of years since 1880 for both females (left) and males (right). Each point in a curve measures the probability that a name appears with frequency X or greater (i.e., $P(x \geq X)$). The leftmost points are all 1.0 as all names appeared with at least the lowest measurable frequency or greater, while the rightmost points reflect only the most popular names (e.g., those that appear with frequencies as high as 5% of the total population in any given year). The overall distribution follows an approximate power-law relationship, as evidenced by the roughly linear relationship between the percentages of names in the top 1,000 at each level of frequency. Note that like Hahn and Bentley (2003), we do not claim that these empirical distributions are necessarily best fit by a power-law each year. Rather, the approximately linear relationship between the log of frequency and the log of occurrence provides a standard against which systematic deviations in the distribution can be compared.

are relatively infrequent (given to only a couple of hundred babies per year), while a much smaller set of names is given to a large number of individuals. For example, in 1880 approximately 8.2% of the registered male babies born were named *Robert* (a raw frequency of 9,655), while there were approximate 609 names that appeared with raw frequency less than or equal to 20 that year. In fact, there were more registered boys named *Robert* in 1880 than all of these 609 uncommon names put together (a combined tally of 5,981)! The approximately linear relationship between the cumulative proportion of names at a given prevalence on a log–log scale suggests that the distribution conforms to an approximate power-law relationship (Barabási & Albert, 1999).²

Across our entire sample of 127 years, the general shape of this power-law-like distribution is somewhat stable (see also, Bentley et al., 2004; Hahn & Bentley, 2003) with an equivalent power-law exponent (α) between 1.75 and 2.0. However, the closer analysis of the cumulative distribution in Fig. 1 reveals that, particularly over the last 50 years, there have been systematic changes in the slope of the distribution. To help visualize this, Fig. 2 (left) plots the best-fit power-law exponent over the entire sample for data aggregated on both a per-year and per-decade basis. Particularly over the last 70 years, the frequency distribution of female names is matched with a consistently steeper slope than for males. For example, the yearly best-fit male exponent was on average lower than the female exponent, $t(126) = 5.5$, $p \ll .001$. This finding reflects differences in the cultural practice of naming females versus males, where female naming is associated with a more diverse choice set and less favoritism for the most popular names. Consistent with this, male names were generally associated with a slightly better overall linear fit (average r^2 over all years was .985, max $r^2 = .988$ in 1880, min $r^2 = .957$ in 2006) than for female names (mean $r^2 = .975$, max $r^2 = .977$ in 1880, min $r^2 = .961$ in 2006). Finally, note that except for a

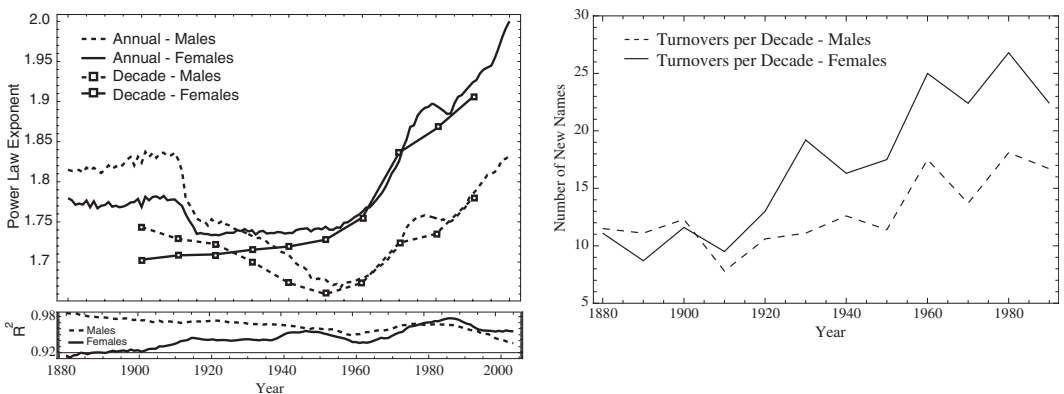


Fig. 2. (Left) The slope of the best-fit t line to the cumulative frequency distribution transformed to the equivalent power-law exponent (α) for each year since 1880. The plot shows the trends for both the annual lists and the top 1,000 per decade. The results show a sharp increase in the power-law exponent starting in the 1950s. The panel underneath shows the r^2 value for the linear fit for each year and for both male and female names. (Right) The number of new names introduced in the top 1,000 list between successive decades.

small period in the late 1970s and early 1980s, the best-fitting power-law exponent has been steadily increasing for both male and female names at roughly the same rate.

As the best-fit line becomes steeper, it reflects a decrease in the relative market share of the most popular names and increasing relative popularity of low and moderately popular names. In fact, one reason for the increase in the best-fitting slope is increasing departures from the canonical power-law distribution in recent years, especially for the most popular names (a point often not acknowledged in previous analyses of naming distributions). For example, while the most popular boy's name in 1880, *Robert*, accounted for 8.2% of the male babies counted that year, in 2007 the most popular boy's name, *Jacob*, accounts for a meager 1.1%. In Fig. 1, this tendency is captured by the fact that the rightmost tail of the distribution is increasingly deflected downwards as the most popular names lose market share. Indeed, in the SSA data, changes in the best-fitting power-law slope were accompanied by changes in the quality of the linear fit.³ Consistent with these general trends, Fig. 2 (right panel) shows the number of names that were replaced on the top 1,000 lists between successive decades. These rates of turnover qualitatively match the changes in the best-fitting power-law slope with more names being replaced in successive top 1,000 lists in recent decades compared with the middle part of the last century, and an overall higher turnover rate for female names.

2.1. Imitation, innovation, copying, and mutation

While the frequency distribution of names at the level of an entire culture is interesting in and of itself, what does it reveal about the strategies that individuals use to select names? Can the historical shifts in naming patterns give insights into the sources of information that individuals use in making these decisions? With respect to the near power-law phenomena in naming, a number of theoretical models have been proposed in order to account for how power-law distributions of elements in a system may form (Mitzenmacher, 2003; Newman, 2005). The most popular class (often referred to as preferential-attachment models) formalizes the intuitive rich-get-richer adage by adding new elements (links, tokens, words) to the system in a biased way such that already popular elements gain even more connections or references by virtue of their popularity, making them yet more popular (Barabási & Albert, 1999).

Hahn and Bentley (2003) developed a related account of how cultural elements (such as baby names) might become power-law distributed on the basis of random copying. In their model (borrowed from work in population genetics), names are considered value-neutral elements (much like junk DNA), which are copied from one generation to the next based on frequency-dependent sampling along with random mutation (Bentley et al., 2004; Hahn & Bentley, 2003; Kumar et al., 2000; Xu et al., 2008). By this account, parents in a given generation choose a name at random by copying the name that some previous parent gave to their child. Given that names are chosen at random and with replacement, the probability of any particular parent selecting a particular name is proportional to the frequency of that name in the previous generation. More popular names are more likely to acquire new adherents, offering them the opportunity for continued

market share in the following generation. Repeating this sampling process in a fixed population often leads to global convergence on a single token, as random growth of some names results in an absolute loss in frequency of others.

However, when this process is augmented by allowing some small percentage of parents to invent a new name at random in a manner akin to genetic mutation, the resulting frequency distribution of names in the final population will be approximately power-law distributed with a slope positively related to the amount of mutation in the system (Cavalli-Sforza & Feldman, 1981; Kimura & Crow, 1964). In this way, certain discrete tokens (which have no intrinsic value) can become extremely common through repeated (but random) sampling while others fail to catch on. In addition, the random-drift model provides a process-level account of the observed deviations in the tails of the power-law shown in Fig. 1. As more novel names (i.e., mutations) are introduced to the system, extremely popular names are likely to lose market share. As a result, there is what appears to be a bias against the most popular names. For example, consider the extreme case of a population of 100 people: if 99 were named *Charles* and one was named *Taran*, then in the next generation, a single individual choosing outside this set of two names will mostly likely mean a loss of adherents for *Charles*. Similarly, losses for less common names like *Taran* simply mean replacement of that name with another, causing little impact on the overall frequency distribution. Thus, the random-drift model explains the recent changes in the shape of the power-law in terms of increased novelty seeking or “mutation,” an observation consistent with many sociological studies of naming behavior in Western societies (Evans, 1997; Fryer & Levitt, 2004; Lieberman, 2000).

2.2. Memory, the mere exposure effect, and the effect of the social environment on individual choice

Previous applications of the random-drift model used it primarily as a null model of cultural change and loosely tied the generative principals of the model to understand the frequency distribution of first names (Hahn & Bentley, 2003) and other cultural artifacts (Bentley et al., 2004; Herzog, Bentley, & Hahn, 2004). While simplistic, the model suggests a tight relationship between individual choices and the structure of an individual’s social environment. For example, in the standard model, the probability of a parent deciding to select name i in generation t , denoted p_i^t , is approximately given by:

$$p_i^t \approx \frac{n_i^{t-1}}{\sum_{j=1}^L n_j^{t-1}} \quad (1)$$

where n_i^{t-1} is the number of agents named i in the previous generation ($t-1$) and L is the total number of distinct name tokens in the previous generation.⁴ At an individual level, one might think of this decision rule as following from well-established effects of fluency on judgment and memory (Hasher & Zacks, 1984; Maddox & Estes, 1997; Malmberg, Steyvers, Stephens, & Shiffrin, 2002). In particular, the relative “value” of a name (and

thus its probability of being chosen) is approximated by the relative frequency of that token in the overall culture. Our later analyses build upon this idea in order to quantify the sources of information that influence name choice.

There are many individual-level psychological mechanisms consistent with such a biased choice process. For example, the well-known mere exposure effect demonstrates how people rate objects that they have seen before (and more frequently) as more pleasant or appealing (Bornstein & Agostino, 1992; Zajonc, 1968). By this account, names that are more common in the social environment of a parent may result in stronger traces in memory, such that at the time the decision is actually made, these names are more perceptually fluent or accessible. Interestingly, parents need not be explicitly aware of the bias that the name environment has on their choices given that the affective enhancement of stimuli due to familiarity can operate independent of recognition of those stimuli (Kunst-Wilson & Zajonc, 1980; Whittlesea & Williams, 1998). Similarly, the recognition heuristic (by which choices which are recognized are preferred; Goldstein & Gigerenzer, 1999) can provide a link between individuals memory and experience, and the aggregate outcomes. For example, Todd and Heuvelink (2007) present simulations showing how social agents following the recognition heuristic (i.e., preferring choices they recognize) can drive convergence of preferences in a population in the absence of external signals of utility (Boyd & Richerson, 1985; Todd & Kirby, 2001).

The second component of the random-drift model, namely, the mutation operator, has a similar grounding in basic cognitive phenomena at the level of the individual. For example, besides obvious desire for individuation by selecting interesting or uncommon names, at a cognitive level, memory biases toward novel or incongruous stimuli (Jenkins & Postman, 1948; von Restorff, 1933; Sakamoto & Love, 2006) might support the differentiate coding of unusual or unique names which facilitate their growth as cultural motifs.

3. Do names “drift” or “march” over time?

While the random-drift model and related variants provide one explanation for the frequency distribution of name tokens and can explain increasing deviations from the power-law relationship as innovation or novelty-seeking increases, the theory ultimately assumes that popular names arrive at their success through an accumulation of random samples that are themselves independent from one generation to the next. For example, in the standard random-drift model, if in one generation 100 individuals are named *Thomas*, then in the following generation, 120 individuals might be given this name, instead of the frequency-matched equivalent of 100. However, on the next iteration, an increase or decrease in frequency is equally likely. In this model, the single best estimate of a name’s popularity in the future is simply its current popularity or frequency in the culture. Due to the fact that decision making depends only on the last generation, change from one period has no impact on change in the next (this is also true for the Xu et al., 2008 model mentioned earlier, although in this model, generations are replaced one person at a time).

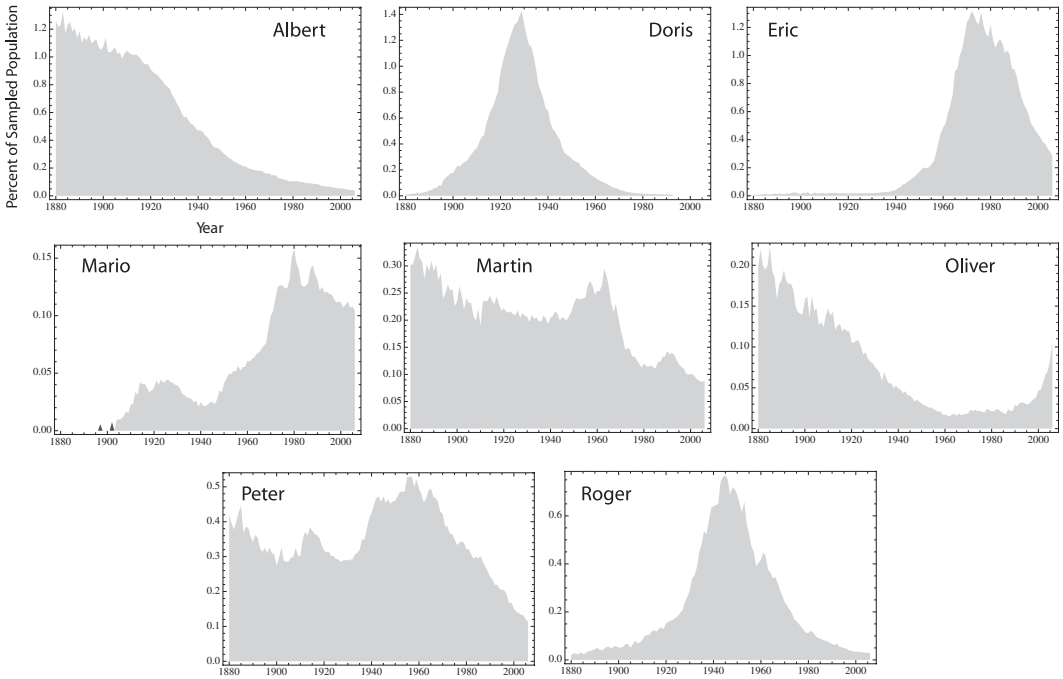


Fig. 3. The normalized frequency of the names of the Nobel Prize winners for 2007 with names that appear consistently in the SSA data: Eric Maskin (Economics), Roger Myerson (Economics), Doris Lessing (Literature), Mario Capecchi (Medicine), Martin Evans (Medicine), Oliver Smithies (Medicine), Albert Gore (Peace), Albert Fert (Physics), and Peter Grünberg (Physics). Each graph plots the number of registered individuals given each name divided by the total number of individuals registered that year for each year from 1880 to 2007.

However, a second striking feature of the popularity of many names in the United States is that they follow surprisingly consistent trajectories over time. For example, Fig. 3 shows the normalized popularity of the names of the 2007 Nobel Prize winners over the recorded lifetime of those names. Note the remarkable stability in the direction of change from year to year that these names exhibit, with increasing popularity in one year being strongly associated with increasing popularity the next year, and vice versa for decreasing popularity. Fig. 4 measures this phenomena in the aggregate (i.e., for all names in our sample) by computing the probability that a particular name goes up in overall incidence (normalized for changes in population size) at time t given that it rose in frequency the preceding year, $t-1$, as a function of the rank of the name in the top 1,000 list at time $t-1$. Also shown are the patterns for $P(\text{up}_t | \text{down}_{t-1})$, $P(\text{down}_t | \text{down}_{t-1})$, and $P(\text{down}_t | \text{up}_{t-1})$. In the earlier time period (1880–1905), names tended to fluctuate in overall frequency from one year to the next. A name that increases its relative frequency one year was more likely to decrease rather than increase its frequency in the following year. Similarly, decreases in frequency are more likely to be followed by increases than further decreases. However, in the more recent data (1981–2007), names move in consistent ways such that a change in popularity one year is predictive of the same direction of change in the following year. In essence, names appear

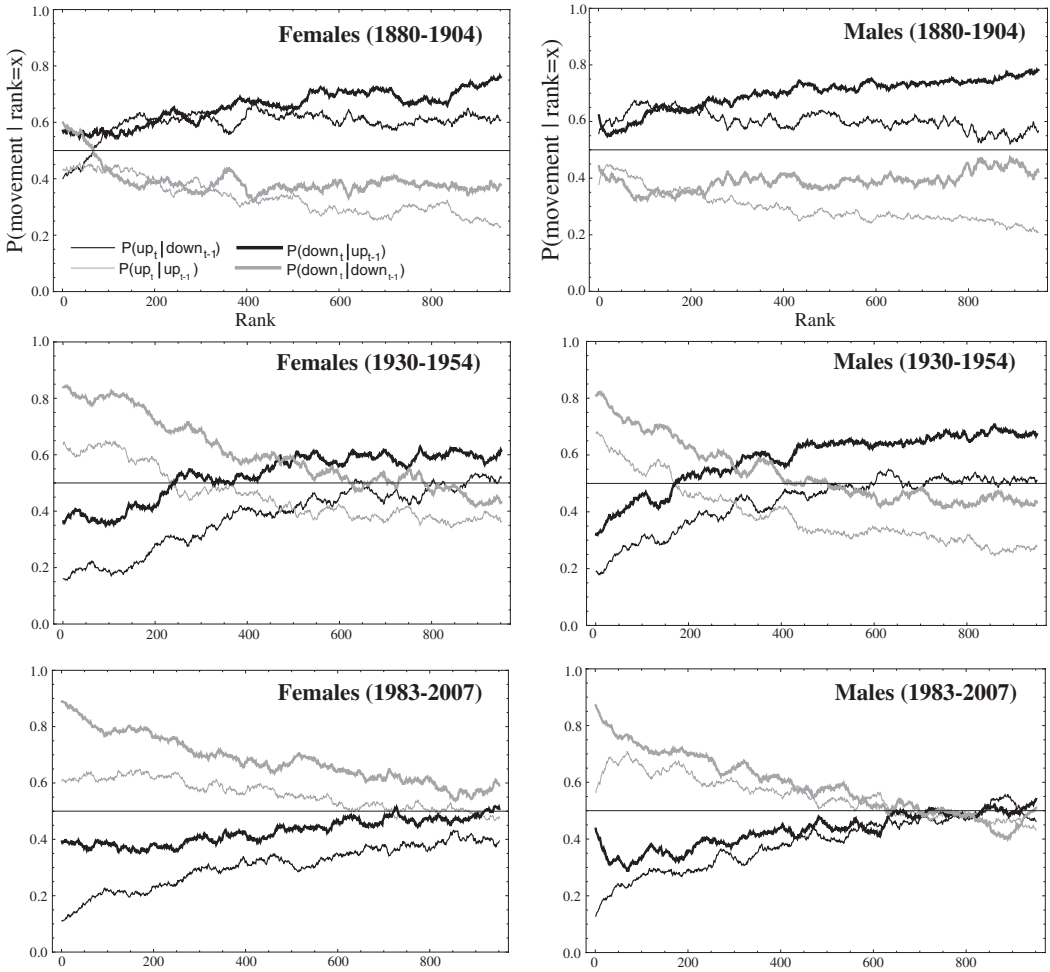


Fig. 4. The conditional probability that a name moves either up or down in overall frequency (normalized for population size), given that it moved either up or down on the previous year, as a function of the rank of the name on the top 1,000 list. The top row plots these values averaged across a sliding window of 25 years from 1880 to 1904. The middle row shows 1930–1954, and the bottom row shows the pattern in recent times (1982–2007). The results show a growing tendency for names to move in a consistent direction from year to year, particularly for the top 200–300 names.

to carry with them a “momentum” that tends to push changes in popularity in the same direction year after year.

Rather than an abrupt shift, the middle row of Fig. 4 (1930–1954) shows how this effect has steadily emerged over the last 100 years. In addition, Fig. 4 shows that names that are already more common (i.e., higher ranked) appear to be more strongly influenced by year-to-year momentum than are relatively uncommon (i.e., lower ranked) names. As a further test of this observation, we attempted to predict which names would rise in popularity and which would fall in popularity in a given year using only the direction of change in

frequency from the preceding year. Starting in the 1880s, using consistent change as a cue for future change is a poor estimator (this strategy is only correct approximately 30% of the time because, in fact, changes in popularity from one year to the next are negatively correlated). However, by the year 2006 our ability to predict year-to-year changes has steadily risen to around 60% for both male and female names (see Fig. 5). A similar analysis collected on naming data from France over a similar period reveals a similar, although less extreme trend.

3.1. The “ratchet effect” and choice momentum

What causes this path-dependent momentum in choice behavior? In order to answer this question, it is first useful to explain what are *not* likely to be the causes. First, note that consistent year-after-year change is relatively common in fashion or product markets. For example, the average width of women’s skirts (measured as the diameter of the hem) has gone through a number of smooth oscillations from relatively short to relatively long and back again (Kroeber, 1919; Lieberman, 2000; Richardson & Kroeber, 1940). Lieberman (2000) describes this in terms of a “ratchet effect” where change in a style is always made with reference to the previous style and the direction the current style took with respect to past styles. In other words, the dynamics of fashion change are incremental: skirt length does not jump from 1880s ankle-length skirts to 1980s miniskirts overnight. Instead, small incremental changes are made in successive generations with respect to established norms.

Sustained directionality for change can also arise as a result of competition between imitative and reactive groups. For example, continuing with the women’s skirt example, higher class and affluent women might lengthen their skirts to show off their ability to afford expensive fabric. The desire of lower class women to imitate the higher class leads to widespread

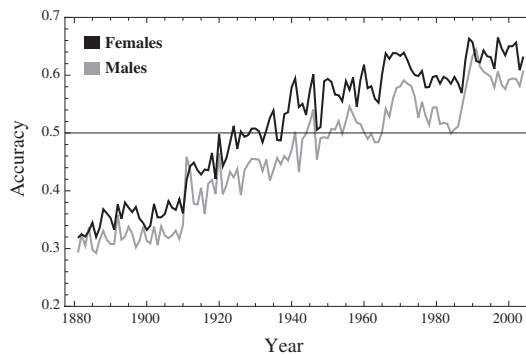


Fig. 5. The accuracy of year-to-year prediction of which direction (increase or decrease in frequency) a name will take using movement in the past year as a predictor. Direction of change one year was a good predictor of future change in the 1880s because an increase in one year was more likely to be followed by a decrease than an increase, and vice versa. However, since the 1950s, the direction of change one year has steadily become more positively related to future popularity. Names present on the top 1,000 list in year x but not present in year $x + 1$ were scored as going down in popularity. Names not present in year x but were present in year $x + 1$ were scored as going up in popularity.

imitation of this fashion trend. However, as this once elite fashion is commoditized there is a renewed pressure for differentiation among the higher class group, and thus the richer women make their skirts even longer. This desire of “trend-setters” to stand out from the norm and for the general public to quickly assimilate these novel fashions can lead to sustained directional change (Lieberman, 2000). A similar competitive principle drives the sustained oscillations in predator–prey systems (Lotka, 1925; Volterra, 1926). However, it is unclear whether this principle is directly applicable in the case of naming. For example, the ratchet effect for skirts requires change along a continuous cultural variable (like skirt length), which can be either “more” or “less” than before. However, name choice is a one-off decision made by a continually changing set of individuals. There is no way for an individual to name a child more *Michael* than last year’s *Michael*. Instead, continued growth in the popularity of a name reflects the desire of more independent individuals to select the name (in other words, the perceived value of the name is modulated by its recent change).

The empirical observation of “momentum” for names is also puzzling from the perspective of contemporary models of cultural evolution in which generation-to-generation change in the frequency of value-neutral tokens is essentially random. One hypothesis (which we evaluate in the next section) is that in the early part of the last century, names may have held longer lasting cultural value (e.g., it was far more common to name a child after a relative and to preserve names within families). Assuming that the long-term value of a name is a slowly changing function of the aggregate number of people given that name, then a random increase one year to a higher frequency will tend to be accompanied by a decrease the following year by virtue of regression toward the mean. Thus, name frequency would (on average) be negatively correlated from one year to the next. However, in recent times, names have become more fad-like, fueling consistent patterns of growth or decline over successive years. Like momentum traders in the stock market, parents appear to be increasingly influenced by recent changes in the relative market share of particular names and to integrate these temporal changes into their estimate of current desirability of a name.

One recent study (Berger & Mens, 2009) provides partial support for this idea. In their study, expecting parents (recruited through naming-related websites) were asked to rate a restricted set of names from the 2006 SSA list on a number of measures, including how likely they would be to give a particular name to their child and how popular they believed each name to be. Overall, adoption velocity (measured as the average rate of change in the name over the period 2001–2006) was positively related to ratings of perceived popularity even after controlling for the actual popularity of the name. Thus, names that had more rapid increases were viewed by contemporary parents as more popular. Furthermore, names that were rated as more popular were also rated as more likely to be actually selected, suggesting the positive effect that change in recent time may have on subjective judgments.

4. Predicting future naming behavior using random-drift principles

In the following section, we develop a modeling approach that allows us to quantify the sources of information that enter into individual naming decisions, and how these

information sources have changed over time. Our primary goal in these analyses was to predict the entire popularity distribution of names for each year of the SSA data set given assumptions about how the social environment leading up to that year influenced the decisions of individual parents. By making a number of simple assumptions about the decision strategies used by individual parents, and the sources of information that enter into these decisions, we are able to infer changes in the “parameters” of naming behavior over the last century and to model the effect that these changes have had on the overall distribution of names. In addition, our modeling analyses allowed us to test the hypothesis developed in the last section that present-day parents are increasingly influenced by the “momentum” of a name in the recent past.

4.1. Using the model to predict the distribution of names

Recall that in the random-drift model, naming choices are made according to the relative frequency of a name in the previous generation (Eq. 1). Thus, we assume that the probability that an individual parent would choose name i at time t should be a function of the relative prevalence of name i among previously named babies and the degree of novelty-seeking/innovation in the culture (i.e., mutation rate in the random-drift model). Building on this basic framework, we included a number of additional assumptions that we felt enriched the ecological validity of the model. For example, in the Hahn et al. model, decision makers sample a name with replacement from the immediately preceding generation (owing to its foundation as a model of *genetic* drift processes). While this assumption is necessary to model genetic transmission, in a cultural context this assumption is somewhat unrealistic. Regardless of how one treats time in the random-drift model (i.e., does one generation equal 1 year or 1 day?), in real life, names can be selected not just from the immediately preceding time period but from any number of previous time periods. Thus, in our extension of the model, we assumed that expecting parents’ decisions reflect the overall probability of encountering an individual with a given name, which is biased over time. Formally, for each year, a recency-weighted estimate of the popularity of each name was computed via the following temporal difference (TD) equation (Sutton & Barto, 1998):

$$e_i^t = (1 - \alpha_g)e_i^{t-1} + \alpha_gp_i^{t-1} \quad (2)$$

where e_i^t is a long-running estimate of the value/frequency of name i at time t , p_i^t is the probability of encountering name i among those born in year t as given by the SSA data set, and α_g is a parameter that controls the degree to which the current estimate e_i^t depends on the most recent naming information, p_i^{t-1} . Note that this equation is equivalent to a simple exponential decay over time, the rate of which is controlled by the parameter α_g with larger values reflecting a stronger weighting of recent information. Thus, Eq. 2 holds that the estimated value of a name is an exponentially weighted average of the probabilities of encountering that name, such that one is more likely to hear more recent name tokens, and less likely to come across infrequent or particularly old names.

When attempting to predict the naming distribution for year t , Eq. 2 was iteratively applied for all years up to $t-1$ using a single value of α_g . For example, when trying to predict the relative name frequencies for 1950, the value of e_i^t for each name was initialized to zero, then Eq. 2 was used to update the values of the e_i^t for each successive year starting at 1880 and ending at 1949. This process was repeated for each year, thus predicting 1951 means starting over at 1880 with a new setting of the α_g and stopping at 1950. Using this procedure, we found the setting of α_g that best predicted the entire name distribution for each overlapping 5-year interval of the SSA data set (year-to-year fits find a similar result but are more influenced by idiosyncratic noise from year to year). We interpret changes in the best-fit value of α_g from one period to the next as measuring generational differences in the types of information that parents rely on in making their naming decisions.

The resulting estimates of the “value” or prevalence of each token are assumed to then bias individual decisions. In particular, the final values of e_i^t were converted into a predicted choice probability according to:

$$\hat{p}_i^t = \frac{e_i^t}{\sum_j^N e_j^t} (1 - \mu_g) \quad (3)$$

where \hat{p}_i^t reflects the *predicted* probability of a parent choosing name i in year t (as opposed to p_i^t , which is the empirical probabilities) and N is the total number of non-zero name tokens. In addition, we assumed that there is some probability that an individual invents a novel name. This tendency was captured by a single parameter, μ_g , and was implemented by subtracting a small probability, $\mu_g e_i^t / \sum_j^N e_j^t$, from each name.⁵ Thus, each existing (predictable) name loses some of its market share in order to accommodate innovation or novelty seeking (cf. Xu et al., 2008). Like the “cultural memory” parameter, μ_g was assumed to be shared among all members of a generation but could vary from one generation to the next.

Our model instantiates the basic principles of the random-drift model in a fairly direct way (names are chosen in proportion to their popularity in the recent past, while a small percentage of individuals choose novel names). In addition, this analysis allowed us to assess the predictive utility of the model’s central principals (relative to its already established ability to generate power-law-shaped distributions; Hahn & Bentley, 2003). Most importantly, the fitting of the model allowed estimation of period-to-period changes in the “memory” (α_g) and novelty-seeking (μ_g) parameters in the population by comparing the best-fit value of these parameters for each period (parameters were estimated by maximizing log-likelihood of the observed name distribution: The full details of the fitting procedure are described in the Appendix). In order to verify that the model provides an adequate account of the data (over and above some less interesting alternatives), we compared the fit quality for the random-drift account against a number of baseline models (see Appendix). These analyses confirmed a superior fit for each time period for this model.

Fig. 6 shows the results of these fits. The left panel shows the changes in the best-fit innovation rate parameter (μ_g) and the right panel shows the best-fit memory parameter (α_g) for each decade and for both male and female names. Overall, the model recovers our intuition that in the period following 1950, there has been an increase in the probability of parents




Fig. 6. The mean best-fit parameters from the predictive model mean smoothed using overlapping 10-year windows. The left panel displays the best-fit setting of the innovation parameter, μ_g . The right panel shows the best-fit setting of the cultural memory parameter, α_g .

choosing a name that goes against the current name distribution (i.e., increasing μ_g , faster for female than for male names). In addition, it appears that the more recent name lists are better fit by considering a smaller window of prior history relative to the early part of the century (i.e., a larger α_g). For example, the best-fit value of α_g steadily increases until around 1960. Interestingly, in the last 20 years, the best-fit value of α_g has trended downward, suggesting that present-day parents may be integrating over a longer window of recent history than were parents in the middle part of the last century (perhaps reflecting the ability in recent time to search for names through online resources). In addition, early sex differences in the α_g parameter for males and females appear to be dissipating. Importantly, the fact that the α_g value remains below 1.0 means that one does a better job predicting each year's list using a recency-weighted estimate of past naming behavior than using last year's list alone. This challenges the assumption in the standard random-drift model that names are copied from only a previous time step and confirms the intuition that name choice is best described as a sampling process that aggregates over time. The negative year-to-year correlations for names in the early part of the century are thus described as a consequence of the longer cultural memory over which naming choices integrated (leading to regression to the mean). Overall, the basic principals in the random-drift model appear to provide both a predictive (demonstrated here) and generative account (Hahn & Bentley, 2003) of the distribution of names in the culture by positing a process of frequency-dependent sampling and random mutation.

4.2. The MILEY model: Measuring cultural changes in memory, choice momentum, and innovation

In our second set of model-based analyses, we extended the predictive random-drift model above to include a bias that favored names that have increased in popularity in recent time and away from names that have fallen. The new model, named MILEY (Momentum Influences Liking Each Year), derives its name from the fastest growing girl's name in 2007

(which was not present in 2006, but debuted at no. 278 in 2007). As before, our goal was to fit the entire distribution of names each year using past choice data and to recover parameters that reveal the changes in individual decision strategies over time.

For each year, a recency-weighted estimate of the popularity of each name was computed using Eq. 2. In MILEY, a second equation is used to estimate the more recent popularity of each name:

$$f_i^t = (1 - \gamma_g)e_i^{t-1} + \gamma_g p_i^{t-1} \tag{4}$$

where $\gamma_g > \alpha_g$. Thus, both f_i^t and e_i^t estimate the popularity of a name in the recent past. However, given the parameter constraint, f_i^t is an estimate of the more recent popularity of the name, while e_i^t tracks the longer term popularity. Parents are assumed to compare the recent popularity of a name, f_i^t , with the long-running average, e_i^t , in order to detect the ‘‘momentum’’ associated with the name:

$$m_i^t = f_i^t - e_i^t \tag{5}$$

Names that, in the recent past, have gathered more adherents relative to the long-running average will have a positive momentum score. In contrast, names that very recently have gathered fewer adherents than would be expected given the long-running popularity will have a negative momentum score. Thus, the momentum term m_i^t indexes the degree of surprise or deviation that an observer would have about a recent popularity of a name relative to its long-term popularity. In other words, names that people detect are outpacing their long-term average popularity are assumed to be positively biased, while names that are underperforming relative to the average are negatively biased. Note that this prediction is also somewhat consistent with Berger and Mens (2009) in that names that grow slowly are expected to have less momentum associated with them because their long- and short-term estimates are always similar (i.e., m_i^t is closer to zero). In contrast, very fad-like names that in one year strongly outpace their average long-term growth are predicted to continue to rise more quickly. Estimates of long-term popularity provided by Eq. 2, and the estimates of the direction of recent change provided by Eq. 5 were combined to generate a final choice probability:

$$\hat{p}_i^t = \frac{\max(0, e_i^t + \beta_g e_i^t m_i^t)}{\sum_j^N \max(0, e_j^t + \beta_g e_j^t m_j^t)} (1 - \mu_g) \tag{6}$$

where \hat{p}_i^t once again reflects the probability of a person choosing name i in year t . Parameter β_g controls the influence that momentum has on the current estimate of the value of a name. The momentum term in this equation ($\beta_g e_i^t m_i^t$) multiplicatively combines the current estimate of the value of the name and the estimate of the change in time capturing the finding in Fig. 4 that momentum appears stronger for more common names. Thus, $\beta_g e_i^t m_i^t$ is positive if the short-term popularity of a name is higher than average, and it is negative if the short-

term popularity of a name is lower than average. The combined sum, $e_i^t + \beta_g e_i^t m_i^t$, was constrained to be positive; thus, relatively unpopular names for whom the contribution of momentum would make the estimated prevalence negative were simply predicted to disappear from next year's list. As before, we assumed that some percentage of the parents choose a novel name according to parameter u_g and that this tendency simply reduces the probability of existing names by a small factor. The predictive random-drift model is thus a special case of the more general MILEY account (where $\beta_g = 0$). As in our fits with the simpler random-drift model, the best-fit values for μ_g , α_g , γ_g , and β_g were found by maximizing the log-likelihood of the actual distribution reported by the SSA in overlapping 5-year windows.

Fig. 7 shows the results of these fits. The top left panel shows the changes in the best-fit innovation rate parameter (μ_g), the top right panel shows the best-fit memory parameter (α_g), the bottom left shows the best-fit "recent" memory parameter (γ_g), and the bottom right shows changes in the weighting of momentum (β_g) for each year and for both male and female names. For almost every period the inclusion of the momentum term in MILEY provided an improved fit when compared with the random-drift model described in the previous section (see the Appendix for more details on the model comparison). Consistent with our previous fits, MILEY captures the fact that in the period following 1950, there has been

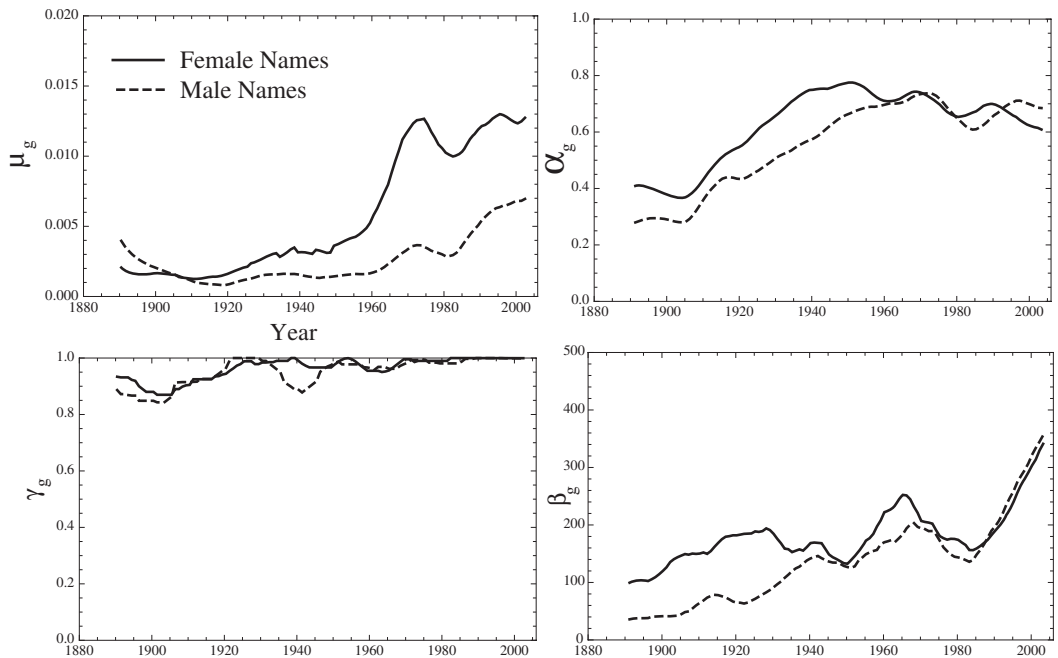


Fig. 7. The mean best-fit parameters from the predictive model mean smoothed using overlapping 10-year windows. The top left panel displays the best-fit setting of the innovation/novelty-seeking parameter, μ_g . The top right panel shows the cultural memory parameter, α_g . The lower left panel shows the memory parameter for recent changes, γ_g , and the bottom right panel shows the best-fit value of the momentum bias term, β_g .

steady increases in the probability of parents choosing a name that goes against the current name distribution (increasing μ_g).

Most importantly, MILEY captures changes in the way that recent changes in name prevalence influence choice. This is particularly clear in the panel showing the best-fit value of β_g , which shows a gradual increase in the weight given to the momentum term in Eq. 6 over the entire data set (generally the model adjusted γ_g so that the more recent name popularity was influenced only the previous year). Overall, the recovered period-to-period changes in the model parameters are broadly consistent with the idea that recent naming decisions more heavily weight both recent name frequency information *and* recent changes in popularity. In this sense, the model provides additional insights into the data patterns reported above. Our simple model assumes that agents are influenced by the distribution of names in the past and are biased toward names whose recent popularity outstrips its long-running popularity. The shift from anti-correlated year-to-year changes to positively correlated changes thus reflects the combined forces of parents basing decisions on both recent popularity, and recent deviations from the norm (i.e., “momentum”). Also significant is that the incorporation of momentum in the MILEY model significantly improved the fit over the extended random-drift account for each year in our data set (see Appendix).

5. Discussion

In this paper, we undertook a historical analysis of the relationship between the statistical structure of an individual’s social environment and his or her choice preferences. In this regard, our work is aligned with the psychological tradition emphasizing the structure of the environment in shaping behavior (Anderson & Schooler, 1991; Gigerenzer, Todd, & the ABC Research Group, 2006). The decision-making domain we examined—naming—provides an excellent opportunity for studying the relationship between individual and group behavior. Compared with a standard psychological experiment, the scope of the Social Security data and its number of “participants” is clearly impressive. By virtue of the extensive historical records on naming, we are able to characterize with reasonable accuracy the probability distribution of names that each individual would have experienced in his or her life prior to having a child, and of the subsequent influence that this distribution had on his or her decisions. Furthermore, our sample set included as many as 90% of the individuals registered with the SSA in a given year (i.e., our data were very close to the complete population distribution of name choices in a given year).

One contribution of the present paper is empirical: We report a number of novel analyses showing how change-in-time is an important factor influencing the desirability of a name. This finding is interesting both for its implications for individual parents (names that are rising are even more likely to be selected than names that are falling; thus, if you want your child to have a unique name the latter may be a better choice!), and for the implications this has for models of cultural evolution. The emphasis on explaining the power-law distribution of names (which we show has also been systematically changing over time) may have led many researchers away from the interesting dynamics that underlie the aggregate pattern of

name choice. Like the stock market, cycles of boom and bust appear to arise out of the interactions of a large set of agents who are continually influencing one another.

Overall, our results and model-based analyses support the idea that individual naming choices are subtly shaped by the social environment of the expecting parents. In particular, we showed how we could improve our prediction of the names list each year by including estimates of long-term popularity along with deviations of the short-term popularity of name compared with its long-term popularity. Unlike an arbitrary data-fitting procedure such as logistic regression, each of the proposed mechanisms in our model has a grounding in well-known psychological processes. By constraining our model predictions in this way, we were able to recover changes in the decision-making variables that seem to vary over time within the culture. In particular, we found evidence of increased innovation and an increasing bias toward names that are growing in popularity.

5.1. The use of archival data to inform cognitive science research

As mentioned earlier, the ecological validity of our investigation is quite unique relative to a typical laboratory experiment. However, the use of archival data in cognitive science research presents a number of interesting challenges. For example, data such as the Social Security Administration name database are sampled from the general population using a variety of methods. In addition, these sampling methods probably change over time. For example, the name distribution as of today (when most babies get a social security card, and records are kept electronically) is probably much more accurate than in 1880. In addition, there may be subtle changes in the sampling procedures that under-represent particular social groups (for example, rural names compared with urban names). Thus, one challenge is assessing how these sampling biases influence the conclusions that can be made from any data set. One approach, which we considered, is to randomly down-sample the survey data in different ways to measure the robustness of our effects assuming a less representative sample. For example, in one simulation, we considered what would happen to our results if each year we excluded 50% of the recorded babies born in any given year. We found that the increasing momentum effect remained robust to this distortion of the data. However, there are many potential biases that might arise. The ideal solution may be to consider data from multiple countries where the nature of the sampling biases is likely to be somewhat different. For example, we examined a data set of French names that shows a similar, although less extreme momentum bias, suggesting that at least some aspects of our results may generalize to other cultures (see also Berger & Mens, 2009).

5.2. When does it stop?

In the MILEY model, we found that including a momentum term in our model improved our ability to explain the name distribution. However, one obvious question is what causes names to stop rising? If rising names are preferred, which in turn causes them to rise, then a momentum bias might quickly lead to convergence on a single token. As yet we do not have a complete explanation for why the situational factors contribute to peaks and reversals in

the actual name data. However, one possibility is that names can only sustain a particular level of growth over a certain time before momentum starts to decay. This might be modeled by an additional factor in the model that causes the bias toward momentum to dampen for names that have been growing too quickly or for too long. However, note that the continued desire of individuals to choose novel or new names (captured by the mutation rate parameter in the model) ensures that no single token will ever overtake the entire population. Thus, continued competition between novel and existing names may help to prevent the runaway growth of any single token and ultimately lead most names to eventually drop in popularity.

5.3. How social environments constrain choice

Overall, our analysis suggests how, in many situations, it may be more productive to extend the analysis of the traditional judgment and decision-making literature beyond the individual to consider the way that individual decisions and social value interacts (Gureckis & Goldstone, 2006). For example, while certain stimuli in the environment have inherent value (i.e., food, sex, and drugs), others only derive their value from the social system in which they are embedded. While this is particularly true for domains that one might think of as fashions (names, clothing, music, etc.), where the very definition of value may be social, it is likely that similar influences make their mark in other domains as well. For example, the value of a particular computer system is tied to the ubiquity and interoperability of the device. One advantage that Microsoft enjoys is that as more people use their systems, more software and hardware are developed that depend on their systems, making the value of these systems still greater. More subtly, social influence is known to increase in domains where individual uncertainty is higher. Thus, when people have less direct information on the quality of choice options, it is more likely that their choices become influenced by the actions of others (Boyd & Richerson, 1985; Laland, 2004). Similarly, as the number of possible options grows, learning through direct experience is less effective, and as a result, social information may become a more useful proxy signal for choice value (cf. Garcia-Retamero, Takezawa, & Gigerenzer, 2006). Naming behavior and fashions may be just a special case of the shift to social information in the face of this uncertainty. Furthermore, in these cases, changes in frequency may become an even more powerful clue concerning “future” social value.

5.4. Reflections of cognition in culture

In their classic paper “Reflections of the Environment in Memory,” Anderson and Schooler (1991) suggest that human cognition has adapted to mirror the information structure of the environment. We would like to take this one step further and argue that the structure and dynamics of our culture and communities may reflect aspects of our cognitive abilities in much the same way. Our results show that decision makers are influenced by the statistical patterns in their environment (in this case, the frequency of names and changes in those frequencies), but that this influence may reflect an aspect of our individual cognition.

Our account suggests that people perceive these sources of information and integrate them into their subjective judgments of value. If this is true, then cognitive processes within the individual in fact may contribute to, and reinforce, the dynamics of the group. Relative to the “null” evolutionary model of Hahn and Bentley (2003), our observations suggest that uniquely cognitive capacities like change detection (Brown & Steyvers, 2005; Silka, 1989), frequency effects on memory (Hasher & Zacks, 1984), and biases toward novelty or incongruous stimuli (Ranganath & Rainer, 2003; von Restorff, 1933) all lead the name market to drift in particular ways. Our formal modeling allowed simultaneous consideration of the constraints from multiple levels of social and behavioral organization. On the one hand, our results show that carefully studying the social environment in which choices are made can reveal the constraints that guide such choices. On the other hand, we find that the structure of that environment is built from and in fact is determined by the choices, preferences, and behaviors of individuals.

Notes

1. Data for our analyses were derived from public records provided by the U.S. Social Security Administration (SSA) reporting the frequency of the top 1,000 first names given to newborn babies for every year since 1880 for both male and female names (<http://www.ssa.gov/OACT/babynames/>). In some analyses, we also considered a separate list from the same source showing the top 1,000 names for each decade. Prior to 1937, the accuracy of the data may be slightly less representative because prior to this time, registration of newborns with the SSA was optional. Overall, the top 1,000 names account for between 70% and 90% of the full distribution of registered babies in a given year.
2. In a power-law frequency distribution, a small number of highly popular elements coexist with a large number of low-popularity elements with the total distribution following a fat-tailed, power-law relationship. The form of this relationship typically relates the magnitude of some quantity (such as the number of citations a paper receives, x) to the probability of encountering objects with this magnitude such that $P(x) \propto x^{-\alpha}$.
3. For example, the average yearly $r^2 = .982$ for male names over the period 1880–1890, while from 1996 to 2006 this value was $r^2 = .944$, $t(18) = 18.4$, $p < .001$. Interestingly, the quality of the linear fit has actually been increasing for female names over the same period (from 1880 to 1890 the average $r^2 = .917$, but by 1996–2007 this value rose to $r^2 = .956$, $t(18) = -40.7$, $p < .001$). Note, however, that the best linear fit (i.e., most faithful power-law relationship) was for male names during the late 1880s (see Fig. 2, lower right panel).
4. Xu, Reali, and Griffiths (2008) extend this basic model to relax the assumption about discrete “generations.” In their model, each agent makes a name decision and replaces the oldest member of the population. The convergence of the process to a power-law distribution is ensured by a small correction to the choice probabilities.

5. Note that in the reported model fits, all parents in a given period are assumed to share the same estimate of e'_i and μ_g for computational efficiency. However, one extension of the model might be to model each agent as “learning” the value of names through limited, individual experience. As a result, we are not necessarily committed to interpreting the μ_g parameter as reflecting the bias of individuals toward novelty (as in the Hahn model) or the average desire across all individuals. In either case, the effect would be a flatter (i.e., uniform) name distribution as μ_g increases.

Acknowledgments

This work was supported by NIH-NIMH training grant TM32 MH019879-12 to T.M. Gureckis and Department of Education, Institute of Education Sciences grant R305H050116 and National Science Foundation grant 0527920 to R.L. Goldstone. The authors thank Luís Bettencourt, Jonah Berger, Jason Gold, Thomas Hills, Julia Hollifield, Winter Mason, Tim Pleskac, Michael Roberts, and Peter Todd for helpful conversations in the development of this project.

References

- Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science*, 2(6), 396–408.
- Barabási, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286(15), 509–512.
- Bentley, R., Hahn, M. W., & Shennan, S. (2004). Random drift and culture change. *Proceedings of the Royal Society of London B*, 271, 1443–1450.
- Berger, J., & Mens, G. (2009). How adoption speed affects the abandonment of cultural tastes. *Proceedings of the National Academy of Sciences*, 106(20), 8146–8150.
- Bornstein, R. F., & Agostino, P. R. (1992). Stimulus recognition and the mere exposure effect. *Journal of Personality and Social Psychology*, 63, 545–552.
- Boyd, R., & Richerson, P. (1985). *Culture and the evolutionary process*. Chicago: University of Chicago Press.
- Brown, S., & Steyvers, M. (2005). The dynamics of experimentally induced criterion shifts. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31(4), 587–599.
- Cavalli-Sforza, L., & Feldman, M. (1981). *Cultural transmission and evolution*. Princeton, NJ: Princeton University Press.
- Evans, C. (1997). *Unusual and most popular baby names*. New York: Consumer Guide.
- Ford, M., Mirua, I., & Masters, J. (1984). Effect of social stimulus value on academic achievement and social competence: A reconsideration of children’s first-name characteristics. *Journal of Educational Psychology*, 76(6), 1149–1158.
- Fryer, R., & Levitt, S. (2004). The causes and consequences of distinctively black names. *The Quarterly Journal of Economics*, 119(3), 767–805.
- García-Retamero, R., Takezawa, M., & Gigerenzer, G. (2006). How to learn good cue orders: When social learning benefits simple heuristics. In R. Sun & N. Miyake (Eds.), *Proceedings of the 28th annual meeting of the cognitive science society* (pp. 1352–1357). Mahwah, NJ: Erlbaum.
- Gigerenzer, G., Todd, P., & the ABC Research Group. (2006). *Simple heuristics that make us smart*. New York: Oxford University Press.

- Goldstein, D., & Gigerenzer, G. (1999). The recognition heuristic: How ignorance makes us smart. In G. Gigerenzer, P. Todd, & the ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 37–58). Oxford, England: Oxford University Press.
- Gureckis, T., & Goldstone, R. (2006). Thinking in groups. *Pragmatics & Cognition*, 14, 293–311.
- Hahn, M. W., & Bentley, R. (2003). Drift as a mechanism for cultural change: An example from baby names. *Proceedings of the Royal Society of London B (Suppl.)*, 270, 120–123.
- Hasher, L., & Zacks, R. (1984). Automatic processing of fundamental information: The case of frequency of occurrence. *American Psychologist*, 39(12), 1372–1388.
- Herzog, H., Bentley, R., & Hahn, M. W. (2004). Random drift and large shifts in popularity of dog breeds. *Proceedings of the Royal Society of London B, Biology Letters*, 271, S353–S356.
- Jenkins, W., & Postman, L. (1948). Isolation and “spread of effect” in serial learning. *American Journal of Psychology*, 61, 214–221.
- Kimura, M., & Crow, J. (1964). The number of alleles that can be maintained in a finite population. *Genetics*, 49, 725–738.
- Kroeber, A. (1919). One the principal of order in civilization as exemplified by changes in fashion. *American Anthropologist*, 21(3), 235–263.
- Kumar, R., Raghavan, P., Rajagopalan, S., Sivakumar, D., Tomkins, A., & Upfal, E. (2000). Stochastic models for the web graph. In *IEEE symposium on foundations of computer science (FOCS) symposium on foundations of computer science (FOCS)* (pp. 57–65).
- Kunst-Wilson, W., & Zajonc, R. B. (1980). Affective discrimination of stimuli that cannot be recognized. *Science*, 207, 557–558.
- Laland, K. (2004). Social learning strategies. *Learning and Behavior*, 31(1), 4–14.
- Liebertson, S. (2000). *A matter of taste: How names, fashions, and culture changes*. New Haven, CT: Yale University Press.
- Lotka, A. (1925). *Elements of physical biology*. Baltimore, MD: Williams & Wilkins Co.
- Maddox, W., & Estes, W. (1997). Direct and indirect stimulus-frequency effect in recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23(3), 539–559.
- Malmberg, K., Steyvers, M., Stephens, J., & Shiffrin, R. (2002). Feature frequency effects in recognition memory. *Memory & Cognition*, 30(4), 607–613.
- Mitzenmacher, M. (2003). A brief history of generative models for power law and lognormal distributions a brief history of generative models for power law and lognormal distributions. *Internet Mathematics*, 1, 226–251.
- Newman, M. (2005). Power laws, Pareto distributions, and Zipf’s law. *Contemporary Physics*, 46(5), 323–351.
- Ranganath, C., & Rainer, G. (2003). Neural mechanisms for detecting and remembering novel events. *Nature Reviews Neuroscience*, 4(3), 193–202.
- von Restorff, H. (1933). Analyse von vorgangen in spurenfeld. i. uber die wirkung von bereichsbildungen im spurenfeld [Analysis of processes in the memory trace. I. On the effect of group formations on the memory trace]. *Psychologische Forschung*, 18, 299–342.
- Richardson, J., & Kroeber, A. (1940). Three centuries of women’s dress fashions: A quantitative analysis. *Anthropological Records*, 5(2), 111–153.
- Sakamoto, Y., & Love, B. (2006). Vancouver, Toronto, Montréal, Austin: Enhanced oddball memory through differentiation, not isolation. *Psychonomic Bulletin and Review*, 13, 474–479.
- Silka, L. (1989). *Intuitive judgements of change*. New York: Springer-Verlag.
- Sutton, R., & Barto, A. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- Todd, P., & Heuvelink, A. (2007). Shaping social environments with simple recognition heuristics. In P. Carruthers, S. Laurence, S. Stich (Eds.), *The innate mind: Culture and cognition* (pp. 165–181). Oxford, England: Oxford University Press.
- Todd, P., & Kirby, S. (2001). I like what i know: How recognition-based decisions can structure the environment. In J. Kelemen & P. Sosik (Eds.), *Advances in artificial life: 6th European conference proceedings*. Berlin: Springer-Verlag.

- Volterra, V. (1926). Variazioni e fluttuazioni del numero d'individui in specie animali conviventi. *Memorie della R. Accademia Nazionale dei Lincei*, 2(VI), 31–113.
- Whittlesea, B., & Williams, L. (1998). Why do strangers feel familiar, but friends don't? The unexpected basis of feelings of familiarity. *Acta Psychologica*, 98(2), 141–166.
- Xu, J., Reali, F., & Griffiths, T. L. (2008). A formal analysis of cultural evolution by replacement. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the 30th annual conference of the cognitive science society* (pp. 1435–1400). Austin, TX: Cognitive Science Society.
- Zajonc, R. (1968). Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology*, 9(2), 1–27.

Appendix

Estimating power-law distributions

Fig. 1 plots the estimated cumulative distribution function (cdf) for names. The cdf has a number of advantages over other methods for estimating power-law behavior, including the fact that the function is continuous, well defined for all values of frequency, and does not require exponential binning of the data. As a result, the cdf can be more sensitive to changes in the tail of the distribution. When plotted in log–log scale, the slope of the best-fit line to the cdf can be used to estimate the exponent of the related power-law function. In particular, the cdf takes the following form:

$$P(x) = C \int_x^{\infty} x'^{-\alpha} dx' = \frac{C}{\alpha - 1} x^{-(\alpha-1)} = \frac{C}{\beta} x^{-\beta}$$

As a result, the slope of the best-fit line in our cdf plots, β , is related to the standard power-law exponent α , via the formula $\alpha = \beta + 1$ (Newman, 2005).

Modeling fitting procedure

We assumed that individual name choices were independent and mutually exclusive, and could thus be modeled using a multinomial distribution. The goal of our fitting was to find parameters that maximized the log likelihood of the actual name distribution given the predictions of the model (i.e., $\log(L(\theta_t|D_t))$, where θ_t is a vector of probabilities given by Eqs. 3 and 6 along with the mutation probability μ_g ; here, D_t is the actual frequency counts for each name in year t). We were concerned with predicting the popularity of names that appeared in the to-be-predicted lists based on past choice behavior rather than attempting to guess which new completely novel names would “appear” in each year’s list (which is highly unpredictable and, for the most part, arbitrary). Thus, when computing our likelihood, new names that did not appear on any previous list but did appear on the to-be-predicted list were grouped into a single name token representing “new names” and given a cumulative frequency n_{new} , which was the sum of the individual frequency counts (generally this captured around 30–50 names out of approximately 1,000 names). Thus, our model was

not penalized for not correctly predicting the sudden appearance of a particular name token but had to predict the relative frequency with which these “sudden appearances” happened each year. For multinomial random variables, the log-likelihood was thus,

$$\log(L(\theta_t|D)) = n_{\text{new}} \log(\mu_g) + \sum_{i=1}^k n_i \log(\hat{p}_i^t) \quad (7)$$

where n_{new} is the number of new or novel names in year t , n_i is the number of people given name i , \hat{p}_i^t is the probability (as given by the model) of name i , and k is the total number of non-zero names. Parameter searches were found for each model that maximized the value of Eq. 7 using the Nelder–Mead simplex algorithm.

Model comparison

In order to verify that the two-parameter random-drift model provides an adequate account of the data (over and above some less interesting alternatives), we compared the fit quality for the random-drift account against a simpler baseline model. The first baseline model we considered simply used the choice probabilities from the previous year (e.g., predicting the name frequency from the year 2000 using the frequencies from 1999). This is equivalent to setting $\alpha_g = 1.0$ in the random-drift model and fitting the value of μ_g . As these models are nested, we can compute a likelihood ratio test of the models. In particular, for each year we computed $G^2 = 2(L_{\text{model}} - L_{\text{baseline}})$, where L_{model} is the total log-likelihood given in Eq. 7 for the random model and L_{baseline} is the same value for the baseline model. The variable G^2 is distributed according to a chi-squared distribution with degrees of freedom equal to the difference in the number of parameters between the two models (in this case one parameter differed). Thus, for each year we computed G^2 for the random-drift model and this restricted alternative and found that the random-drift model provided a significantly better fit even against a Bonferroni-corrected α , ($p < .0001$). Thus, the inclusion of the α_g parameter in the random-drift model significantly improved the fit of the model compared with simply using the previous year’s list. We also tested to see if restricting the windows over which names were integrated in a more discrete way (e.g., by using only that previous 25 years to predict each year) would improve the fit of the model. However, we found in all cases, that the exponentially weighted estimate provided by the full model provided a significantly better G^2 fit for each year.

A similar analysis was conducted on the MILEY model to verify that this model provided a significantly better fit over the random-drift model. Once again, G^2 was computed as the random-drift model and is a nested case of the more general MILEY (where $\beta_g = 0.0$). The difference in the number of parameters between the random-drift model and MILEY is two (2.0). For each year of the data set, we found that the MILEY model provided a significantly better fit ($p < .0001$). Thus, despite the additional complexity in the MILEY model, the inclusion of momentum bias provides a more accurate account of the empirical naming distributions.