From Karen to Katie: Using Baby Names to Understand Cultural Evolution

Jonah Berger
The University of Pennsylvania

Eric T. Bradlow
The University of Pennsylvania

Alex Braunstein
Chomp.com

Yao Zhang
The University of Pennsylvania

Contact Information: University of Pennsylvania, 3730 Walnut St. Philadelphia, PA 19104; jberger@wharton.upenn.edu. Thanks to Keelan Evanini for help parsing the names, Uri Simonsohn for suggesting the hurricane analysis, and Ann Kronrod, Christophe Van den Bulte, and Duncan Watts for helpful comments on the manuscript.

Abstract

How do psychological processes shape cultural evolution? We use over 100 years of first names data to investigate how the popularity of other, similar variants shapes cultural success. Names can be broken into phonetic parts, or phonemes, and we examine how a name's popularity is influenced by the popularity of that name's component phonemes in *other* names in the past year. Building on mere exposure research, we show that names are more likely to become popular when similar names have been popular recently. These effects are non-linear, however, indicating that over-popularity hurts adoption, and vary based on phoneme position. Further evidence for the causal impact of similarity on cultural success comes from a natural experiment using hurricane names. An exogenous shock to phoneme frequency, due to prominent hurricanes containing those sounds, boosts the popularity of phoneme-sharing names. Taken together, our results suggest how inter-item similarity shapes popularity and cultural evolution.

Is it possible to predict which summer movies will be hits? Which political ideologies will catch on? Whether “Madison” or “Margaret” will be a more popular baby name next year?

Academics, popular press writers, and practitioners alike have long been interested in trying to predict cultural evolution, or which cultural tastes and practices will become popular next (Farrell, 1998; Gladwell, 2006; Simonton, 1980). Some ideas, styles, norms, and social movements catch on while others languish. But while some companies have suggested they can predict hit songs (Elberse, Eliashberg, & Villaneuva, 2006), there is no empirical data that actually supports that claim. Further, even domain experts have great difficulty forecasting what will become popular (Bielby & Bielby, 1994; Hirsch, 1972). In fact, cultural evolution often seems random (Hahn & Bentley, 2003; Pinker, 2007; Salganik, Dodds, & Watts, 2006), with one style, song, or idea catching on due to seemingly happenstance events and the fact that people tend to imitate one another.

But might it be possible to predict how culture evolves? Might there actually be some regularity in the way culture changes over time (Lieberman et al., 2007)? Most diffusion research has assumed the successes of different cultural items are independent of one another (Bass, 1969; Rogers, 1995). In such models, the popularity of one technological innovation or style is neither facilitated nor inhibited by the popularity of other innovations or styles.

In contrast, we suggest that inter-item similarity shapes cultural evolution. Just as culture impacts psychological processes (Markus & Kitiyama, 1991), the converse is also true: psychological processes shape the norms, tastes, and choices that make up culture (Berger & Heath, 2008; Berger & Milkman, 2012; Heath, Bell, & Sternberg, 2001; Kashima, 2008; Schaller & Crandall, 2004). In particular, research finds that exposure to a stimulus increases preference not only for that stimulus, but also for related stimuli that share common features (e.g., nonidentical but similar looking shapes, Gordon & Holyoak, 1983; Landwehr, Labroo, and Hermann, 2011; Monahan, Murphy, & Zajonc, 2000). Taken to the collective level, this suggests that the popularity of similar items should impact cultural success. Songs, for example, may be more likely to become popular if their chord progressions are similar to recent hits (Simonton, 1980).

To study this phenomenon quantitatively, we examine the popularity of first names from 1882 to 2006. Names are composed of phonemes, or perceptually distinct units of sound (e.g., the first name "Karen" = K EH R AH N, The CMU Pronouncing Dictionary). In addition to being influenced by its own past popularity (i.e., same item effects), we examine how a name's popularity (i.e. number of babies given that name in a given year) is impacted by the popularity of that name's component phonemes in *other* names in the past year (i.e., inter-item effects). We investigate whether the popularity of the name "Karen," for example, is influenced by the recent popularity of other names that start with a hard "K" sound (e.g., Carl and Katy) or end with an "N" sound (e.g., "Darren" and "Warren"). We also examine whether these relationships vary based on whether the phonemes appear at the beginning (in this example, "K"), middle ("EH," R," or "AH") or end ("N") of the name.

We focus on names for a number of reasons. First, in many domains, producers determine the available options. Consequently, any patterns in inter-item similarity over time could merely be a result of production incentives. Movie studios or car manufacturers might make new products similar to old products because it allows them to retain the same personnel or production line. In contrast, name choice is driven by individuals and essentially unconstrained (i.e., parents can select any name they like), making it an ideal domain to study internal drivers of cultural evolution. Second, there is relatively little influence of commercial effort or advertising on name choice. Third, unlike technology domains there is little (if any) relative advantage or quality differences between items (Gureckis & Goldstone, 2009; Lieberson, 2000). The absence of these factors makes it easier to examine the relationship between inter-item similarity and cultural success.

Study 1: 280 Million Births

Study 1 uses a hierarchical Bayesian model to estimate the impact of inter-item similarity on name popularity. We acquired data from the U.S. Social Security Administration on the number of babies born with different names from 1882 to 2006 (over seven thousand names and 280 million births). Controlling for the number of babies born with a given name in the past year, the model estimates how the usage of a name in a given year is impacted by the usage of its first, last, and internal phonemes in *other* names the previous year.

Preliminary Results and Model Building

Figure 1 shows that the more phonemes a pair of names has in common, the more their popularities are correlated over time ($p < 0.01$). To provide a more rigorous examination, we use a formal statistical model which controls for a name's direct effect on itself, isolating how cross-item effects shape name popularity.

To reflect both the count nature of the data, and the large number of zeros due to names being introduced and dying off (being unused for a number of years), we decompose our probability model into two parts commonly known as a spike-at-zero model (Morrison & Schmittlein, 1988). As shown in Figure 2, Stage 1 models whether a given name (i) is potentially used in year (*t*). If yes, usage is determined by a hierarchical Poisson-regression model (Stage 2, described below). If no, then usage is zero.

We rely on four sets of variables to estimate what drives the time-varying propensity for potential name usage, p_{it}, and given potential usage, the average usage intensity (the Poisson mean, μ_{it}).

Set one consists of two control effects. A baseline name effect which accounts for the fact that some names are inherently more popular (across years, "Michael" may be more popular than "Max"). A time effect accounts for the fact that there are more babies born in some years than others. Thus, we include α_{pi} and α_{Ui} in the models for name *i*'s impact on propensity (*p*) and usage (*U*) respectively, and δ_{pt} and δ_{Ut} for year *t* propensity and usage effects.

Set two consists of further controls: the recent popularity of a given name *i* in the previous time period, *t*-1 (same item effect). We include this variable in both the propensity and usage model because we are interested in understanding the impact of phonetic similarity on popularity *after* controlling for a name's own previous popularity. These effects are denoted as $\beta_{p1}*N_{pi,t-1}$ and $\beta_{U1}*N_{pi,t-1}$ where N_{t-1} is the name's usage in the previous year.

Set three are our primary variables of interest: for a given name *i*, the usage of the first, middle and last phoneme in the previous year. In particular, we include lagged terms for propensity ($\beta_{p3}*F_{i,t-1}, \beta_{p4}*M_{i,t-1}, \beta_{p5}*L_{i,t-1}$) and usage ($\beta_{U3}*F_{i,t-1}, \beta_{U4}*M_{i,t-1}, \beta_{U5}*L_{i,t-1}$).[1]

Lastly, to test whether over-popularity kicks-in, we include both linear and quadratic effects in both the propensity and usage models. Thus, in addition to controls for name and time (described in variable set one), our explanatory variables are a 4 (same item, first phoneme, internal phoneme, last phoneme) x 2 (linear, quadratic) design.

---

[1] In the few cases that names did not have a middle phoneme, that variable was assigned a value of 0 in the analysis.

Mathematical Model and Estimation

To account for heterogeneity in effects and small sample inferences, we utilize a hierarchical Bayesian formulation as follows. For brevity, we describe only the model for Stage 1, noting an identical specification for Stage 2 provided in the supplementary materials. Inferences from the model were derived by obtaining samples from the posterior distribution of model parameters using Markov Chain Monte Carlo (MCMC) sampling.

Stage 1 model: Potential use of baby name i at time t

$Y_{it} \sim \text{Bernoulli}(p_{it})$ [1]

$\text{logit}(p_{it}) = \alpha_{pi} + \delta_{pt} + \beta_{p1} * N_{pi,t-1}$ [Controls]

$+ \beta_{p2} * F_{i,t-1} + \beta_{p3} * M_{i,t-1} + \beta_{p4} * L_{i,t-1}$ [Main Parameters of Interest]

$+ \beta_{p5} * N^2_{pi,t-1} + + \beta_{p6} * F^2_{i,t-1} + \beta_{p7} * M^2_{i,t-1} + \beta_{p8} * L^2_{i,t-1}$ [Main Parameters of Interest]

Results

Figure 3 illustrates the relationship between name usage in year *t* and the popularity of other names that include those phonemes in year *t*-1. For the name “Karen” in 2000, for example, the model predicts the popularity of that name based on the usage of names that begin with the K sound, end with the N sound, or have EH, R, or AH internally in 1999.

Results show that even after controlling for the past popularity of a name, names are more popular when their component phonemes were more popular in *other* names the previous year.

These inter-item effects are non-linear, however, and vary based on phoneme position (i.e., whether phonemes appear at the beginning or end of a name; see Figure 4 and Table S2).

Take first phoneme. Figure 4 shows the independent impact of first, middle, and last phoneme, fixing the other variables at their mean observed values (see also Table S2). Increased usage of a name's first phoneme in *other* names one year initially has a positive influence on that name's popularity the next year. A shift from a first phoneme being used in 100,000 births in other names the past year to 125,000, for example, is associated with a 36% increase in name usage. The impact eventually wanes around 183,000 births using a given sound, however, and turns negative such that further increases in usage decrease predicted name usage the following year. This suggests that over-popularity or tedium may kick in after some saturation point (Berger & Le Mens, 2009) and reduce the desirability of particular phonemes.

The effects are similar, but smaller, for internal or last phonemes ($p < 0.01$, see Table S2). This indicates that while names will be more popular when their last or internal phonemes have been used more in other names the previous year, this increase is even larger when their first phoneme has been used more the previous year. This is consistent with prior work on primacy effects (Whitney, 2001; Perea & Lupker, 2003), whereby earlier things are more salient, receive more attention, and have less lateral interference.[2]

Study 2: Hurricanes as a Natural Experiment

These findings suggest that inter-item similarity shapes name popularity over time but one might wonder whether these relationships are truly causal. To test this, we use a natural experiment. We examined how an exogenous increase in the frequency of hearing a phoneme impacts the popularity of other names that share that phoneme.

[2] Results are similar using a measure of phoneme popularity from the past 10 years rather than just the last one, or allowing for cross-position effects (Supplementary Materials, Figure S3)

We focus on hurricanes. When hurricanes cause more damage, their names are mentioned more frequently, leading their component phonemes to be heard more often. Importantly, however, hurricanes are automatically assigned names from pre-existing lists and are named long before their damages are known. Consequently, a hurricane's occurrence provides an exogenous shock to the frequency of hearing certain phonemes (there is no effect of prior name or phoneme popularity on hurricane size; see Supplementary Information). This allows us to examine the causal impact of phoneme frequency on the popularity of similar names. How the incidence of Hurricane Katrina, for example, affects the popularity of not just the name "Katrina," but all names that begin with "K" such as "Katie" and "Carl".

Data

Since 1953, the United States National Hurricane Center has maintained a list of pre-approved names for tropical storms and hurricanes. In the Atlantic Ocean, for example, there are six lists of 21 names (each name starts with a different letter from A to W, not including Q and U). One list is used every year, and they repeat such that each list is used every seventh year. Names are assigned to storms in alphabetical order, such that the first hurricane of the season starts with the letter A, the next starts with B, and so on down the list. When an unusually destructive hurricane occurs, its name is retired and replaced by an alternate name.

We collected the names of all hurricanes from 1950-2009 using http://www.wunderground.com, as well as the amount of damages they caused (adjusted for inflation).

Hurricane Model

We model the impact of hurricane damages on the popularity of names with shared phonemes. The model retains the spike-at-zero structure from Study 1 and adds hurricane effects at both stages. In either stage, two parameters are introduced to capture linear and quadratic effects respectively. For brevity, we include Stage 1 here and Stage 2 in the supplementary materials.

Stage 1 model: Potential use of baby name i at time t

$Y_{it} \sim \text{Bernoulli}(p_{it})$ [2]

$\text{logit}(p_{it}) = [\text{eqn 1 effects}]$

$+ \theta_{p2} * HD_{i,t-1} + \theta_{p4} * HD_{i,t-1}^2$ [Hurricane effect]

where HD_{it} is the sum of the damages caused by hurricanes in year *t* which shared phonemes with baby name *i*.

Results

The results provide causal evidence that increased phoneme frequency boosts the success of similar, phoneme-sharing names. The more attention a hurricane name receives (i.e., the more damage it caused), the more popular first names with those phonemes become (Figure 5). Following Hurricane Katrina, for example, names that begin with "K" saw an approximately 9% increase in usage. The effect is non-linear, however, and after a certain point additional increases in damages no longer increase the popularity of phoneme-sharing names. This suggests a potential tradeoff between (a) a positive boost from phonetic familiarity and (b) either over-popularity or the negative connotations associated with an extremely damaging hurricane.

General Discussion

Taken together, these findings suggest some regularity in how culture changes over time. Inter-item similarity shapes cultural evolution: names are more likely to be popular when similar variants have been popular recently. While we focused on phonetic similarity, comparable effects should also occur for other types of perceptual and potentially even conceptual similarity. Thus one would expect similar results for other cultural domains. Songs, technological innovations, and other cultural products may not only become popular based on their own characteristics, or whether they are better or worse than competing items, but also based on whether they sound like, look like, or share common features with other things that have been prevalent recently (Simonton, 1980).

These results underscore the importance of considering inter-item similarity when predicting diffusion and cultural success. While prior work has shown that patterns of random copying can predict the aggregate distribution of cultural popularity (Hahn & Bentley, 2003), such models are less useful in predicting the success of individual cultural items. Further, our results (Table S1) show that including cross-item effects greatly improves model fit and leads to more accurate predictions of name popularity. This indicates that incorporating the popularity of similar items improves predictions of cultural evolution above and beyond random copying, or just considering an item's past popularity by itself.

Moderately similar cultural variants may be particularly successful because they provide "optimal" innovation. Repeated exposure to a stimulus increases familiarity, thus improving affective responses, but tedium or satiation eventually kicks in (Berlyne, 1970; Jakobovits, 1966). Further, familiarity boosts preferences more when it is unexpected (Schwarz, 2004). Consequently, moderately discrepant stimuli may be the ideal blend of familiarity and novelty. Similar enough to evoke the warm glow of familiarity, but different enough to feel fresh and have the familiarity be unexpected (Flavell, Miller, & Miller, 2001). Further, because cultural parts (e.g., phonemes, chord progressions, or components of ideas) are combined with one another to form larger cultural units, individual parts may be able to exist in a larger percentage of the population before tedium ensues.

More generally, this work supports recent theorizing on the psychological foundations of culture (Kashima, 2008; Schaller & Crandall, 2004). When shared across individuals, psychological processes can shape the beliefs, norms, tastes and institutions that make up culture (Berger & Milkman, 2012; Heath, et al., 2001; Markus & Kitayama, 1991). In this instance, preferences for familiarity may underlie the link between inter-item similarity and cultural success. The findings also speak to the reciprocal influence between individual decision making and collective outcomes (Gureckis & Goldstone, 2009). Name popularity impacts individual choices, which in turn shape the collective patterns that influence the future choices of others, and in turn, cultural evolution.

References

Bass, F. (1969). A new product growth model for consumer durables. *Management Science*,*15* (5), 215–227.

Berger, J. & Heath, C. (2008). Who Drives Divergence? Identity-Signaling, Outgroup Dissimilarity, and the Abandonment of Cultural Tastes. *Journal of Personality and Social Psychology*, *95*(3), 593-607.

Berger, J. & Le Mens, G. (2009). How Adoption Speed Affects the Abandonment of Cultural Tastes. *Proceedings of the National Academy of Sciences*, *106*, 8146-8150.

Berger, J. & Milkman K. (2012). What makes Online Content Viral? *Journal of Marketing Research*. Forthcoming.

Berlyne, D. E. (1970). Novelty, complexity, and hedonic value. *Attention, Perception, & Psychophysics, 8*, 279-286.

Bielby, W. T. & Bielby, D. D. (1994). All hits are flukes: Institutionalized decision making and the rhetoric of network prime-time program development. *American Journal of Sociology, 99*, 1287-1313.

Elberse, A., Eliashberg, J. & Villaneuva, J. (2006). Polyphonic HMI: Mixing Music and Math. *Harvard Business School Case* #9-506-009.

Farrell, W. (1998). *How Hits Happen: Forecasting Predictability in a Chaotic Marketplace*. New York: Harper Business.

Flavell, J. H., Miller, P. H., & Miller, S. A. (2001). *Cognitive Development* (Prentice Hall, Saddle River).

Gladwell, M. (2006). The Formula. *The New Yorker*, October 16.

Gordon, P. C. & Holyoak, K. J. (1983). Implicit learning and generalization of the “mere exposure” effect. *Journal of Personality and Social Psychology, 45*, 492-500.

Gureckis, T. M. & Goldstone, R. L. (2009). How You Named Your Child: Understanding the Relationship Between Individual Decision Making and Collective Outcomes. *TopiCS in Cognitive Science, 1*, 651-674.

Hahn, M. W. & Bentley, R.A. (2003). Drift as a mechanism for cultural change: an example from baby names. *Proceedings of the Royal Society of London. 270*, S120-S123.

Heath, C., Bell, C., & Sternberg, E. (2001). Emotional selection in memes: The case of urban legends. *Journal of Personality and Social Psychology, 81*, 1028-1041.

Hirsh, P. (1972). Processing fads and fashions: An organization-set analysis of cultural industry systems. *American Journal of Sociology*, *77*, 639-659.

Jakobovits, L.A. (1966). Studies of fads: I. The 'Hit Parade.' *Psychological Reports*, *18*, 443–50.

Kashima, Y. (2008). A Social Psychology of Cultural Dynamics: Examining How Cultures Are Formed, Maintained, and Transformed. *Social and Personality Psychology Compass, 2*, 107-120.

Landwehr, J. R., Labroo, A., & Herrmann, A. (2011). Gut Liking for the Ordinary: Incorporating Design Fluency Improves Automobile Sales Forecasts. *Marketing Science*, *30*(3), 416-429.

Lieberman, E., Michel, J. B., Jackson, J., Tang, T., & Nowak, M. A. (2007). Quantifying the Evolutionary Dynamics of Language. *Nature*, *449*, 713–716.

Lieberson, S. (2000). *A Matter of Taste: How Names, Fashions, and Culture Change.* New Haven: Yale University Press.

Markus, H., & Kitayama, S. (1991). Culture and the self: Implications for cognition, emotion, and motivation. *Psychological Review*, *98*, 224-253.

Monahan, J. L., Murphy, S. T., & Zajonc, R. B. (2000). Subliminal mere exposure: Specific, general, and diffuse effects. *Psychological Science, 11(6)*, 462-466.

Morrison, D. G. & Schmittlein, D. C. (1988). Generalizing the NBD Model for Customer Purchases: What Are the Implications and Is It Worth the Effort? *Journal of Business & Economic Statistics, 6*, 145-159.

Perea, M. & Lupker, S. J. (2003). Does jugde activate COURT? Transposed-letter similarity effects in masked associative priming. *Memory & Cognition, 31*, 829-841.

Pinker, S. (2007). *The Stuff of Thought: Language as a Window into Human Nature*. New York: Viking.

Rogers, E. (1995). *Diffusion of Innovations*. New York: Free Press.

Salganik, M. J., Dodds, P. S., & Watts, D. J. (2006). Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market. *Science*, *311*, 854–56.

Schaller, M. & Crandall, C. (2004). *The Psychological Foundations of Culture.* Mahwah: Lawrence Erlbaum Associates.

Schwarz, N. (2004). Metacognitive Experiences in Consumer Judgment and Decision Making. *Journal of Consumer Psychology, 14*, 332-348.

Simonton, D. K. (1980). Thematic fame, melodic originality, and the musical zeitgeist: A biographical and transhistorical content analysis. *Journal of Personality and Social Psychology*, *38*, 972-983.

Speech at CMU (2011). *The CMU Pronouncing Dictionary* (Carnegie Mellon University, Pittsburgh); available at <http://www.speech.cs.cmu.edu/cgi-bin/cmudict>.

Whitney, C. (2001). How the brain encodes the order of letters in a printed word: The SERIOL model and selective literature review. *Psychonomic Bulletin & Review, 8*, 221-243.

Figure 1. More similar names have more correlated popularities over time. Average correlation between the popularity of pairs of names over time based on the number of phonemes they have in common. Error bars show 95% confidence intervals

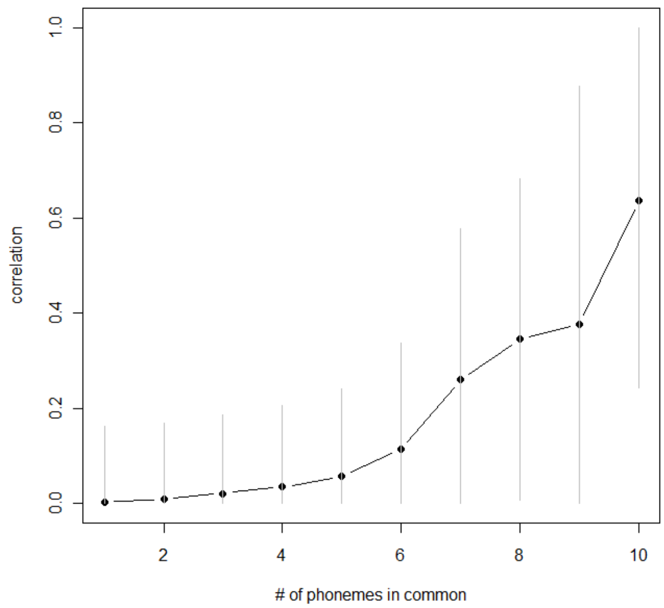

Figure 2: Two-stage model for name popularity

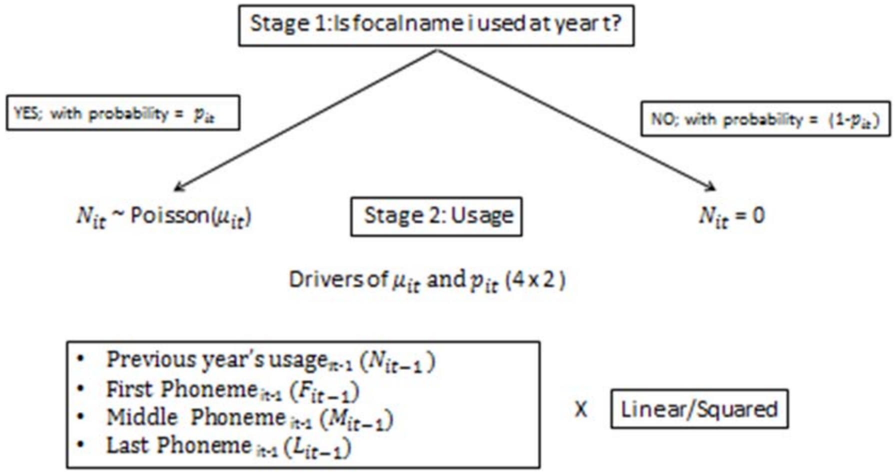

Figure 3. Names are more popular when their component phonemes were more popular in *other* names the previous year. Average number of babies born with a given name as a function of births the previous year with other names that include those phonemes. Includes control for a name's own past popularity. Each point represents one percentile of phoneme usage.

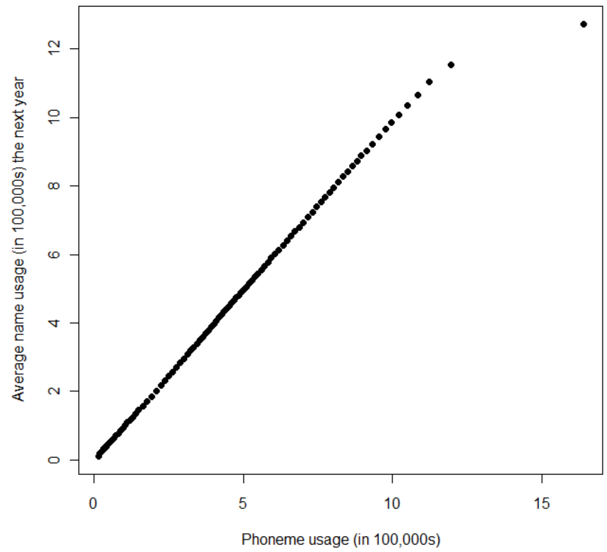

Figure 4: Phoneme popularity in other names has a non-linear impact on name popularity and is strongest for first phonemes. Controlling for a name's past popularity, predicted births of a given name as a function of the popularity of its first, internal (middle) and last phonemes in births the previous year. Predicted lines are shown for 97.5% of the distribution of data.

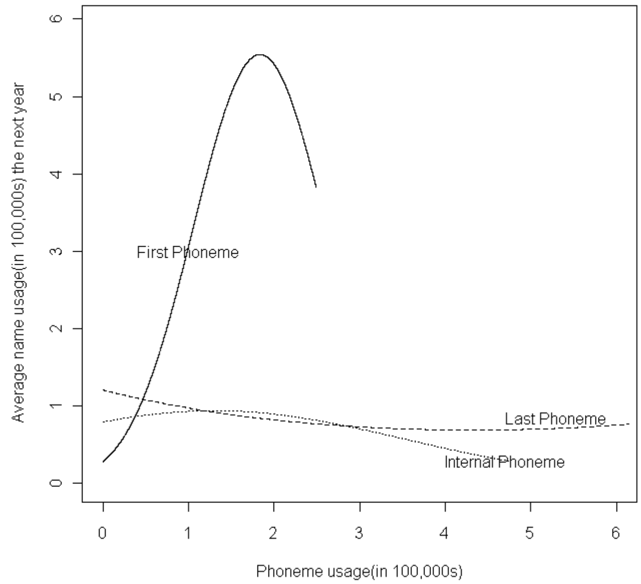

Figure 5: When a hurricane gets more attention, first names that share phonemes with the hurricane's name become more popular. Many hurricanes have zero damages but the mean value of hurricanes where damages occurred is $100 million.