

STATEMENT ABOUT PAST AND FUTURE RESEARCH: KEITH CHEN

I am a behavioral economist who studies individual choice behavior. While my work has ranged over a number of topics and subjects, I am drawn to questions where I believe that the prevailing explanation for an important human behavior is unnecessarily complex. I try to find simpler, more general drivers, located at a more basic level of explanation. I believe several of my projects shed new light on the way people make important economic decisions. I will briefly summarize the major areas of my research, and locate my work in the broader literature.

Early Work with Monkeys

My earliest, most numerous, and most cited papers involve experimental work with monkeys. I began this work at Harvard (working with tamarin monkeys), and have continued it at Yale, where I work with capuchin monkeys. When I started this work, though experimental methods with human subjects were already quite common for economists, work on animals was unusual.

A few economists had studied animal behavior before, most notably Kagel, Battalio & Green who studied a variety of economic decisions (e.g., consumer demand, labor supply, risk aversion, and intertemporal choice) in rats and pigeons. Similar to the point made by Gary Becker in his 1962 paper “Irrational Behavior and Economic Theory”, KB&G’s papers demonstrated experimentally that the central insights of price theory (such as downward sloping demand curves) work well even in very low-rationality environments. In a parallel literature, psychologists had studied monkeys to examine the evolutionary origins of cognitive systems such as visual processing, numeracy (monkeys can count to about 4), working memory, other-regarding preferences, and even theory of mind. My work on capuchin monkeys (with co-author Laurie Santos) links these two literatures by using a behavioral theory of decision making (prospect theory) to examine the decisions of closely-related primates.

Specifically, we present monkeys with decisions similar to those in which humans display cognitive biases, biases which have been implicated in systematic departures from optimal decision making. By looking at how closely-related primates make these same decisions, we examine the continuity or discontinuity of the psychological processes underlying the biases humans display. We argue that if we find that a particular behavioral bias is present not just in humans but in closely related primates, then it is likely that that bias is an innate behavior rather than a learned response to uniquely human experiences such as markets, institutions, or culture.

To do this, we began by introducing a fiat currency to a colony of capuchin monkeys. This allowed us to present them with a large set of potential trades and purchase opportunities. A monkey would first be introduced to small metal washers that we were using as coins. At first, monkeys were rewarded for merely playing with the coins. Later, we introduced them to “traders”, human experimenters who would offer a monkey a piece of food in one outstretched hand, but give the food only if the monkey put a coin in the trader’s other hand. Different traders (who wore different color uniforms) each traded for different types of food. Our monkeys

quickly began to recognize these people: an apple human, a pineapple human, and a grape human. Gradually, our monkeys displayed many signs that they understood the value of these coins: a monkey who found a coin would immediately rush to the trader who sells their favorite food. Our monkeys also quickly learned to recognize and respond to changes in price, buying more when food was on sale and less when it was expensive. Our monkeys also began to understand that other monkeys also valued these coins: trading coins with each other emerged spontaneously.

Overall, training took about six months, after which almost all of our monkeys were making basic budgeting and purchasing decisions. Once they seemed comfortable with prices and trading, we would introduce a monkey to even more complex trades involving risk. Several new traders acted similarly to the old ones except for one key difference: these traders would sometime add or subtract food before handing it over. For example, one trader would always show a monkey two pieces of apple, then if given a coin, would either take a piece away and only hand over one piece, or add an extra piece and hand over three (each with 50% probability).

With a bit of experience our monkeys seemed to understand gambles, and this new ability allowed us to present them with choices designed to be similar to those used to study decision-making biases in humans. In our first set of experiments, we looked at gambles in which humans typically display both reference-dependence (outcomes are evaluated as gains and losses, not as levels) and loss-aversion (losses motivate more than comparably sized gains). We found that capuchin monkeys display the same prospect-theoretic biases, both qualitatively and quantitatively; monkey and human estimates look identical. Specifically, capuchin monkeys displayed a coefficient of loss aversion (how much more losses feel bad than gains feel good) of around 2.5, almost exactly the mean of human estimates.

The paper we wrote about these first experiments was published in the *Journal of Political Economy* and has been well-cited: according to Google Scholar it is the JPE's 7th most cited paper from 2006. My co-authors and I have extended our earlier work to other components of prospect theory, allowing us to show that the theory as a whole appears to hold broadly in closely-related monkeys.

Our most important follow up paper shows that monkeys, like people, systematically switch from risk-averse to risk-seeking behavior, when identical gambles switch from being described as gains, to being described as losses. This flip in risk-preference is called the reflection effect, and is one of the three principle components of prospect theory. The experiment we run is simple to describe: in two different experiments, monkeys have to decide between two pieces of apple for sure (option A), or a 50-50 gamble between one and three pieces (option B). What differs between experiments is what a monkey starts with: monkeys begin with either one or three pieces of apple.

If they start with one piece of apple, option A (two for sure) is implemented by always giving the monkey a gain of one piece, while option B entails a 50% chance of gaining two pieces (and a 50% chance of no gain). Conversely, if the monkey starts off with three pieces, Option A is a sure loss of a piece, while Option B is a 50% chance of a loss of two pieces. What we show in our 2010 paper is that monkeys prefer the sure thing when their options are presented as gains, but prefer to gamble when their options are presented as losses. This mirrors the reflection effect in human behavior, an effect which is thought to be responsible for the disposition effect, one of the most robust anomalies in behavioral finance.

We also wrote a paper testing for the endowment effect, one of the most studied anomalies in behavioral economics. The endowment effect (an experiment replicated hundreds of times in humans), is the pattern that when randomly given either good A or B, the vast majority of people will refuse to trade A for B if given A, but also refuse to trade B for A if given B. That is, people seem to systematically value a good more after they own it, and are reluctant to trade, even if they did not initially receive their preferred object. We show that this also holds with monkeys, even after controlling for trading costs, uncertainty, and waiting times.

Collectively, my monkey research has been cited over 400 times, and has led many other teams of animal researchers around the world to focus on how behavioral biases do or do not manifest themselves in closely related species.

Early Empirical Work: Crime and Recidivism

In an influential paper on deterrence and prisons, Katz, Levitt and Shustrovich show that very unpleasant prison conditions (specifically, very high in-prison death rates), seem to discourage contemporaneous crime. KLS argued that would-be criminals consider just how bad prison is before committing crimes. This suggested that very harsh and unpleasant prison conditions, though bad for prisoners, might reduce crime rates at large.

My paper with Jesse Shapiro (*American Law and Economics Review*, 2007) asks whether this deterrence effect could be offset by the possibility that harsher prison conditions would make *current* prisoners more likely to re-commit crimes after release. Aggregate numbers suggest this is a large issue: in the United States over 600,000 prisoners are released each year, and over two-thirds will be re-arrested within three years. This alone accounts for a substantial share of crime. Yet, prior to our work, not only did criminologists not have good studies of a recidivism-effect, but they were widely split on whether the effect was negative or positive.

While we had thought that the most natural prediction was a hardening-effect (where tough prisons *increase* post-prison crime) there was a lot of sociological work that suggested harsher prison conditions *reduce* post-prison crime. One of the dominant theories in the sociological literature on crime was Specific-Deterrence Theory, which posited that criminals learn about the severity of prison and the attractiveness of crime from their own experiences: hence terrible prison conditions would scare inmates straight. This theory was backed by careful sociological

work. Inmates who were incarcerated in harsher prisons appeared to leave with a significantly worse view of prison, and self-report being much more determined to avoid returning. As economists, we wondered if this intention carried through to post-prison behavior.

Measuring this poses a hard inference problem, because prisons are not randomly assigned. The type of prison a person is sent to depends on the type of crime they have been convicted of, their past record, and many individual characteristics, including sex, age, and gang membership. There are clear identification issues in drawing conclusions from noting that people released from high-security prisons recommit crime at a much higher rate than those released from low-security. Many predictors of recidivism (most notably the severity of the crime for which they were convicted) are what put them in harsher prisons in the first place.

To overcome the identification problem, we looked for data that would allow us to explicitly control for the assignment process of people to different types of prisons. We found a small, federal-prison data repository which, for a little over one-thousand inmates released in 1987, had linked these inmates to a newly formed set of state records, allowing us to connect the federal data to state re-arrest records for each inmate's first three years post-release. Also, the federal prison system had very clear rules for determining prison assignment. Each inmate is given a "security-custody score" based on objective criteria such as age, criminal history and severity of current crime. This score is then used to assign inmate to prisons of different security-levels.

Our data allowed us to reconstruct, for each inmate in our sample, the security-custody score they were given upon entry to the federal system. Using this score, in a regression-discontinuity design, we estimated the *causal* impact of prison assignment on behavior, comparing inmates who found themselves just on either side of a cutoff sending them to different types of prisons. The regression-discontinuity (RD) design identifies off of the fact that the treatment (here, what type of prison you go to) changes *discontinuously* in your security-custody score, while other characteristics of the inmate are changing continuously.

We find very large effects in this paper: each additional prison security level causally increases the 3 year re-arrest likelihood by more than 10 percentage points. It also has a dramatic effect on the severity of crimes: harsh prison conditions make criminals more likely to recidivate into violent crimes, not profit-driven (pecuniary) crimes.

We published this paper in the ALER, following the Katz, Levitt & Shustorovich paper. Our paper won the American Law and Economics Association's best-paper prize (for the best paper of the last two years), and is the most cited paper in the ALER since 2007, the year it was published. Subsequent researchers have used our technique on richer data sets that allow for more precision, and in total these papers have now contributed a tremendous amount to our understanding of how incarceration affects many aspects of ex-inmates post prison lives and behavior.

Current Theory Work on Cognitive Dissonance: the Effect of Choice on Preferences

The central methodological assumption of micro-economic theories of behavior is revealed-preference: economic actors reveal a coherent set preferences with the choices make. Interestingly, one of the most important theoretical paradigms in social psychology flips this logic on its head. Developed by Leon Festinger in 1957, the theory of cognitive dissonance posits that when two or more cognitions are inconsistent with one another, an uncomfortable state of ‘dissonance’ is produced. People are motivated to resolve this by changing their cognitions, and because knowledge about one’s recent behavior is especially resistant to change, the study of dissonance reduction has typically focused on attitudes that shift in the direction of recent behavior. In other words, dissonance theory posits that behaviors shape preferences, rather than just reflect them.

Since 1957, there have been hundreds of studies empirically testing the specific theory that cognitive dissonance drives people, after making a simple choice between two options, to increase their liking of the option they choose, and devalue the option they reject. My work shows that all of the standard methods of testing this theory, however, are flawed, and yield false positive rates approaching 1. In later work, Jane Risen and I develop and test a research methodology that correctly tests for this form of “mere-choice induced” dissonance.

Dissonance theory has received support from the results of three experimental paradigms. Developed by one of Festinger’s earliest students (Brehm, 1956), experiments using the “free-choice paradigm” find that after making a difficult choice between two items, participants’ rating of their chosen alternative tend to rise, and the rating of rejected alternative tend to fall. This phenomena has come to be called ‘spreading of alternatives’, and was the earliest and most reliable demonstrations of dissonance reduction. Another set of experiments used the “induced-compliance paradigm”. Festinger and Carlsmith (1959) asked participants who had just finished a boring task to tell another student that the task was interesting. The researchers found that participants who were paid \$20 to lie did not change their attitudes because they could easily explain the inconsistency between their action and their attitude. However, participants who were paid only \$1 to lie, expressed a more positive attitude toward the boring task. Finally, using an “effort-justification paradigm”, Aronson and Mills (1959) and Gerard and Mathewson (1966) found that participants who had to go through a severe initiation to join (what turned out to be) a pretty dull group liked the group more than those who went through only a mild initiation.

Early dissonance theory researchers hypothesized that attitude change occurred because people were motivated to undo the unpleasant state that was caused by the inconsistency between their attitude and their recent behavior. Thus, the interpretation was that to reduce the experience of dissonance, participants come to like chosen objects more, form a more positive attitude toward boring task, and come to see dull groups as more interesting.

More recent work in social psychology, using children and monkeys as subjects (see Egan, Santos, and Bloom *Psychological Science*, 2007) argue that cognitive dissonance is actually a much more basic process: it occurred just as reliably in animals that don't display a well-developed sense of self (capuchin monkeys), and in amnesiacs who can't remember the choices which were the source of purported inconsistencies. These radical new findings suggested that dissonance processes didn't require conscious knowledge of inconsistent cognitions, or even really, cognitions. All of these new papers used the free-choice paradigm, which had become the most widespread experimental test of cognitive dissonance. I will use the set up in the ES&B paper to provide the intuition behind my critique of the empirical work in this area.

ES&B's main experiment begins with children rating a number of objects on a discrete five-point scale. After rating enough objects, three objects that were rated equally (say rated 4) are chosen for use in a second stage of the experiment. Call these three items A, B, and C. The experiment then asks children to choose between a randomly chosen two of these items, say A and B. Calling the object which the child chooses A, the child is then asked to choose between B (the initially rejected item), and C (a random third item that was also initially rated 4). The hypothesis here is that the act of choosing between A and B will *cause* the child to devalue the rejected item B. ES&B look for this by examining how children then choose between B and C. If a child is more likely to choose C than B in this choice, they are said to have engaged in dissonance reduction. I argued that this was to be expected, and that in fact, a *perfectly* rational person should be *expected* to choose good C 66% of the time. To see this, consider every possible (strict) way people could feel about the three goods A, B, and C, listed in the table below.

Table: All Possible Preference Orderings

	Case 1:	Case 2:	Case 3:	Case 4:	Case 5:	Case 6:
Best:	A	A	B	B	C	C
Middle:	B	C	A	C	A	B
Worst:	C	B	C	A	B	A
Cases where A is preferred to B are shaded. In these cases:						
B vs. C:	B	C			C	

Note that if every possible preference profile is equally likely (which must be true if A, B, and C were presented in random order), then after a subject has revealed themselves to prefer A over B (by choosing A over B), then we should in fact predict that they would choose C over B two-thirds of the time. This is just a consequence of liking A more than B telling us a bit both about a person feels about A, and on average, how they feel about B. In ES&B, children choose C 63% of the time and capuchins choose C 60% of the time. ES&B argue that since both of these numbers are significantly greater than 50%, both monkeys and young children experience dissonance. I argued that both of these results should be compared to 2/3rds, and that only

numbers significantly higher than 2/3rds would demonstrate cognitive dissonance (both number are in fact, statistically indistinguishable from 2/3rds).

Note the fundamental problem here: this experimental design assumed that three goods that were all initially lumped in the “4” bucket must be completely identical, and that people’s choices between them could be thought of as random (good A is no different than good B, even though it was chosen). But this is only true if asking a child to sort stickers into five discrete buckets is enough to completely exhaust how they feel about stickers (and false if one of the stickers is actually a 4.1, one 4.26, and one a 4.9). Indeed absent this strong assumption of equality, the belief that C and B should be equally chosen in the third round is mathematically equivalent to a well-known logical fallacy, popularly known as the three-door (or Monty-Hall) problem. Fundamentally, even though subjects are asked to make choices between randomly chosen options, *which good they choose* is not random, and ignoring that may introduce powerful biases into an experiment.

At the most fundamental level, this critique focuses attention on one of the most fundamental principles of microeconomics: choices reveal preferences.

Now while ES&B is relatively novel in using a three-goods version of the FCP, what I discovered was that all FCP paper suffered from the same fundamental problem, and that this could be easily quantified. I found that this methodological flaw could explain EVERY pattern in this literature, including the finding in children, amnesiacs, and the interactions researchers had been finding between dissonance with culture and self-affirmation.

Thinking that the best place for this work to have impact was in Psychology journals, I began collaborating with psychologist Jane Risen, and together, we did two things. First, I re-wrote the proof to be much more detailed, with separate sections for each sub-type of free-choice dissonance experiments, rather than tying the proof to the information structure common to every FCP experiment. Second, Risen and I ran a series of free-choice experiments which demonstrated that (at least in the logic of this literature), cognitive dissonance was able to make causality run backwards in time. That is, we ran experiments that showed (as predicted by my theorem), that if dissonance was measured the way the literature had been measuring it, then choices that a subject is going to make *tomorrow* (but does not know about) causes dissonance-driven preference change in how they feel *today*. Under certain conditions, these experimental methods also provide a control condition that can be used to isolate the bias we identify, and correctly measure choice-driven dissonance. This became our main paper on this topic, and is published in the Journal of Personality and Social Psychology, widely considered the best journal in Psychology.

Risen and I followed up on this early paper with several that show that this problem holds even more broadly than we had originally claimed, and offer experimental techniques that avoid the revealed preference problem. This work has remained controversial and some work continues

using the old techniques. By and large though, I think that this work is slowly having a positive impact on this important literature: several papers have now appeared that take these issues seriously. These new papers explore alternative methods of measuring choice-induced effects in ways that do not suffer from the problems I identified. Unfortunately, none of these techniques have yet produced results that (at least to my read) establish dissonance, and several initial reported successes (most notably the blind-choice studies) have failed to replicate. The control conditions that Jane and I suggested have also proven to be less useful than we had hoped: recent experiments have shown the identifying assumptions we discussed do not generally hold (for example, mere-exposure liking effects invalidate the approach we suggested). Overall though, many reasonable people in this literature now understand that most of what they thought they knew has not been shown, and many are now actively working with much more reliable techniques.

Current Work: Language and Behavior

Most recently, I've been working on a new research project looking at the connection between the structure of people's language and their decision making. So far, this research agenda has produced one working paper which has been accepted at the *American Economic Review*, and several ongoing research projects that I'm working on with psychologists and linguists. This work has already generated a considerable amount of excitement and attention among linguists (who are historically skeptical of work on the connection between language and thought or behavior). I've presented the paper not only at numerous economics and psychology departments, but also at the Yale, Penn, and Stanford Linguistics departments. Also, I've written an invited piece in an ongoing debate about the work on "Language Log", a forum on which many of the world's top linguists communicate with each other about both new research and news in linguistics.

The idea that language can impact the way people think and act has a rich history in linguistics, philosophy, and psychology and has been a topic of considerable controversy. The idea that language can influence thought has become known as the Sapir-Whorf hypothesis (SWH, Whorf 1956). This hypothesis has generated several interesting lines of research in cognitive linguistics and psychology, which have found robust effects across a number of cognitive domains.

For example, Russian makes an obligatory distinction between light blue (goluboy) and dark blue (siniy). Winawer et al. (2007) finds that Russian speakers do better than English speakers in distinguishing blues when the two colors span the goluboy /siniy border (but not when they do not), and these differences are eliminated when subjects must simultaneously perform a verbal (but not a spatial) distractor task. Further implicating language in this differential precision, Franklin et al. (2008) finds that this difference holds for adults, but not for pre-linguistic infants. What this (and several other examples in the psychology literature) teach us, is that obligatory linguistic categories can affect people's ability to categorize, encode, and recall subtle distinctions in the world around them.

As an economist, I became interested in whether these sorts of obligatory distinctions could affect decision making. In my first project on this topic, I focused on how time is encoded in a language, which connects to the large literature in behavioral economics on discounting. To my knowledge, this is the first work testing for a link between language structure and decision making.

English and German evolved from the same proto-language (English is a Germanic language). Linguistically, however, there are fundamental differences between the two languages on time perception. When you're talking about a future event in German, it is almost always permissible to speak in the present tense, as long as the context prevents confusing it with present time. The same is true for Chinese, Japanese, and all of the Nordic languages. What first struck me, was that all of these exemplars of languages with very weak grammatical distinctions between the present and the future are also countries which have extraordinarily high savings rates; rates difficult to explain by purely economic factors.

Now, of course this observations isn't really proof of anything: but it started this research agenda. The hypothesis this led me to is that languages which do not grammatically distinguish between present and future events (what linguists call weak-FTR languages) lead their speakers to take more future-oriented actions. Intuitively, fewer distinctions between present and future events puts future benefits more on par with present costs, making it easier to save.

I show how this prediction arises naturally when well-documented effects of language on cognition are merged with models of decision making over time. Then, I examine whether this hypothesis manifests itself in people's real-world economic decisions. To do this, I adopt the most detailed typological classification I could find in Linguistics, an examination of cross-linguistic future-time reference strategies produced by the EUROTYP project (a large European Science Foundation funded language typology involving hundreds of linguists from 1990-94). I find a strong correlation between their measure of how a language treats future-time reference (FTR), and the choices that speakers of those languages make when thinking about the future. Specifically, in large data sets that survey families across hundreds of countries, I find a strong and robust correlation between the obligatory marking of FTR in the language a family speaks, and a whole host of forward-looking behaviors, like saving, exercising, and not smoking. What's remarkable is that these correlations hold both across countries and within countries, even when comparing effectively identical families born and living in near each other.

The empirical work revolves around trying to eliminate sources of possible common causation. Obviously, savings behavior of current families can't have CAUSED their language structure. What we have to worry about is the possibility of something like correlated-cultural adoption: the German language spreading along with German values towards savings. More than this, we might need to worry about a whole host of other possible factors: might some languages have co-diffused (spread together) with attitudes towards education? With attitudes towards work? With a particular religion?

Most of my paper is spent checking for exactly these types of concerns. For example, if we thought languages co-diffused with proclivities towards education, one way to see that would be to compare only families with identical levels of education. If we thought the co-diffusion was mainly spatial, we might think to compare families born and living in the same country, or even city. Effectively, I try to control for confounding factors that could drive spurious correlations using conditional-logistic regression. What I do, in effect, is drop families around the world into one of 1.4 billion buckets, where two families fall into the same bucket if and only if they are identical in country of birth and residence, age, sex, income, family structure, number of children, and religion (broken up into 74 types). What I then look at is pairs of families who fall into the same bucket, but who report speaking different languages at home.

Now, at this level of detail, even very large economic data sets only allow me to look at around twenty-five-thousand families around the world, who live in nine countries with enough native linguistic diversity to form such pairs. Those countries are Belgium, Burkina Faso, the Democratic Republic of the Congo, Estonia, Ethiopia, Malaysia, Nigeria, Singapore, and Switzerland. What I find is that in every one of those countries, the categorization of languages I adopt from linguistic typologies (typologies built by hundreds of linguists) seems to have a large and consistent correlation with savings behavior. That is, the direction and the magnitude of these effects are statically identical in every country, despite spanning different regions of the world, and very different languages.

I also investigate a few other ways these effects might be spurious, and find no evidence of any alternative hypotheses. These effects can be measured not only in savings behavior, but in many different future-regarding behaviors studied by economists, smoking, exercise, obesity, and condom use. I find exactly the same effect in all three of those behaviors, in the predicted directions, comparing sets of nearly identical families. The OECD has collected savings data going back through the 70's, and when I look at these effects over time, I find stable effects in every decade from the 70s to the present. These effects are not driven by minority languages, nor are they affected by the inclusion of several controls for a cultural value on savings. In short, the data simply don't seem to suggest that confounding factors play a significant role in explaining the consistent set of correlations between the strength of future reference in language and future-regarding behavior.

I'm excited about this research agenda, and I am working with several linguists and psychologists to extend and deepen this initial work. Some of this work tries to look more carefully at exactly which future-referring strategies in languages seem to drive this effect. To do this I'm working with several linguists who are doing a more focused typological classification of languages, focusing on properties of language that behavioral models suggest might affect intertemporal choice. I've also begun a series of experiments that randomly manipulate language structure in targeted way and have already found intriguing evidence of small, short-lived, but robust effects on intertemporal choice.