

SESSION A: BEHAVIORAL ECONOMICS AND POLICY DECISION RULES.

MODERATOR: JESSE LEARY, Federal Trade Commission

PRESENTER: COLIN CAMERER, Cal Tech

DISCUSSANTS: JONATHAN KLICK, Florida State

PAUL PAUTLER, Federal Trade Commission

JACK CALFEE, American Enterprise Institute

Mr. CAMERER: Should I jump up and start?

Mr. LEARY: We're going to dive right into the next session so we can try to -- well, actually we're ahead of time, but we're going to keep moving because I'm sure we're going to have a lot to talk about in the next session, if I can find my notes.

The next session is going to start with a presentation by Colin Camerer on behavioral economics and consumer policy decision rules, and this is going to be I think an overview of how the various findings in behavioral economics might be helpful to developing and implementing consumer protection policy.

After about a 25 minute presentation by Professor Camerer, we're going to have three folks comment on it and offer additional discussion about the issue, and if things go smoothly, I'm hoping that we'll have about 15 minutes at the end to get to continue the discussion and involve and bring the audience in, and I'll also try to give Professor Camerer a couple minutes if

he has any specific issues that come up in the other presentations to respond to.

The folks that are going to be talking during this session are Professor Colin Camerer. He has an MBA in finance and a Ph.D. in decision theory from the University of Chicago. Before he joined Cal Tech, he taught at the Kellogg, Wharton and University of Chicago business schools. He studies behavioral and experimental economics and has published extensively in these areas.

After Professor Camerer, Jonathan Klick will be speaking. He is the Jeffrey A. Stoops Professor of Law at Florida State University. Professor Klick has a Ph.D. in economics as well as a law degree, and his research focuses on empirical analyses of how legal and regulatory changes affect individual behavior.

After Professor Klick will be Paul Pautler. Paul is the Deputy Director for Consumer Protection here within the Bureau of Economics here at the Federal Trade Commission. In that role he has overall responsibility for the economic input into consumer protection case work and policy making, as well as the various studies that we're doing here at the Commission.

He previously held several positions on the antitrust side including being assistant director for antitrust. He's published on antitrust economics and policy, health economics regulation,

and most recently his work has focused on mergers, health care and consumer protection.

Finally we'll have Jack Calfee. Dr. Calfee is a resident scholar at the American Enterprise Institute. His job after graduate school was actually as an economist here at the Federal Trade Commission, where he worked on consumer protection matters, especially on advertising, tobacco, marketing and the regulation of information more broadly.

He's since taught marketing and consumer behavior at the business schools at Maryland and Boston University. His more recent work at the AEI, his academic work and his opinion pieces cover a variety of topics including tort liability, advertising and information, FDA regulation and the pharmaceutical market more broadly.

I probably should have said at the beginning who I am. I'm Jesse Leary. I'm Assistant Director For Consumer Protection in the Bureau of Economics here at the Federal Trade Commission.

So I'll ask Professor Camerer to come up and give his talk. Do you know how to get it started or do you need help?

PROFESSOR CAMERER: I should be able to get it. Thanks. This is really a great deal, and I'm glad you all put this together. Some people who have been critical of neuroeconomics, which I won't say very much about today, but it's kind of a

common, natural extension of the behavioral economics to understanding brain process.

It was said in a paper that a behavioral economist is "one who enlists the services of the state to conceal facts to help people," and I have no interest in eliciting the services of the state, and apparently in fact the state is interested in eliciting in my services or at least is interested in hearing what people are thinking from a very academic view, very basic lab experiments.

We're interested in what we have to say about the emerging approach of the behavioral economics, but I'm really actually much more interested in what comes later today which is very concrete studies of things which is where we should be going, so I'm going to be talking about several things. I'll give you very quick prelude. I'll just give you a taste.

Many of you know at least some of the buzz words and some of the basic studies that have established things like loss aversion and optimism and status quo bias or default effects, and we'll hear more about them in context which is great, but I'll cover that kind of quickly.

I'll talk a little bit about some frontiers just to remind you that a lot of things which often people are sensibly critical about are emerging as the latest research in sort of a normal

science way. I'll say a little about private and public solutions. We're going to hear a lot more about that later.

I'll speculate a little bit about what are the domains of consumer protection where there might be big mistakes that the Commission should be thinking of and have some conclusions.

So here's roughly my position about paternalism and the influence of behavioral economics on such. I'm not an advocate for frequent changes in laws and constitutions, but laws and constitutions must go hand in hand with the progress of the human mind, and that progress might mean --

UNIDENTIFIED SPEAKER: Can you please use the microphone?

PROFESSOR CAMERER: Yes. I'm not an advocate for frequent changes in laws and constitutions, but laws and constitutions must go hand and hand with the progress of the human mind. That might mean that if consumers get better, smarter and smarter, we don't need to protect them as much. It might mean as firms get better and better at using things like viral marketing and internet tactics to fool consumers, we need to regulate that better, or it might be that progress in understanding the human mind by behavioral economists, psychologists and neuroscientists should be incorporated into policy.

Not only is this my point of view, in fact I got it from the Jefferson Memorial. This is Thomas Jefferson's quote. So a lot

of behavioral economics, as Nava Ashraf and George Lowenstein and I have noted and others, there was really a lot of interesting speculation about psychology in the roots of economic thought, in Adam Smith, Riccardo, Marshall, Keynes obviously and some of the behavioral economics is just revisiting those areas and kind of backing up to the fork in the road and taking a more psychological route.

So my definition is one adapted from Mullen, Nathan and Faylett, (phonetics) who have a very nice review paper on this topic. I define behavioral economics as using facts and methods from other social sciences to understand limits on rationality, will power and self interest and their implications, and I should add kind of post Eddie's talk that many of us, not all, but many of us, David Laibson is here, Matthew Rabin see this as really a kind of economics. It's not fundamentally opposed so the models often include concern for equilibrium, concern for heterogeneity and sorting, as they should.

So we're interested in applying standard tools like equilibrium analysis, also maybe developing tools like disequilibrium analysis in the form of learning or persistence of mistakes. We care about field data, and we try to do welfare analysis. Obviously that's often complicated if you have things like multiple cells over time, but it's no more complicated than

this stuff that people do in economic theory and mechanism design or any other difficult field.

Here's a short kind of a list of some of the kinds of limits on rationality that people have studied with sort of examples that are roughly relevant to consumer protection, so one is numerical mistakes, like underestimating the power of compound growth or there's some evidence I think from Raj Chetty at Berkeley that if you simply post the price of a consumer object including the taxes on the shelf, people buy less of it, so they're just not adding the tax on in their head because it's not concrete.

David I think may talk a little bit about shrouding and fees. I think that's telling us about limited attention. One thing economists haven't studied too directly, although it's often alluded to by people like Jacob Marshak long ago, what's really scarce in the brain is attention, and the Internet is not going to help that. It's just going to make things worse.

So models in which people have limited attention and attention is either attention to visible things and stronger than invisible things, or attention is drawn away rapidly by blinking signs is one way to think of behavioral economics. I've been interested in limits on strategic thinking. That is, some people don't realize that the choices of others may be revealing some

private information. You might call this gullibility or something like that.

Quite a few things people have studied fall under the rubric of motivated cognition; that is, the things we believe and the things we would like to believe, and we often get them confused so wishful thinking is an example; the various kinds of optimism, thinking that life will be great for me in the future or thinking it will be greater for me than for somebody else that I'm competing with and overconfidence. These are things that I think have been studied quite a bit, including some field data.

Probably the biggest initial win in sort of applied behavioral economics has been the recognition of the power of defaults to actually increase savings in Benartzi and Thaler's so-called SMarT plan and in some other interventions.

So here's a useful framework which some of you may not have seen, and I think it provides the best available language now, although it doesn't look like economic language. It's psychological to think about a unifying framework for judgment mistakes and also maybe too a language to think about what the FTC should be doing and has done.

So this is from a paper by Danny Kahneman, his Nobel lecture actually in 2003 in the AER, so it's very prominently published, and Danny, following many other psychologists in the last 15

years, suggests that there are two kinds of systems in the brain which are responsible for judgment. There's an intuitive system one, which is very associative of some of the properties that are given here. It's kind of perception like. It's often affective. It's very fast. It's something like Alex Torteroff's (phonetic) study where he shows people, two senatorial candidates, and he says: How competent is this person? And they make very reliable competent judgments in probably a fraction of a second, 500 milliseconds, and that's highly correlated with who gets elected to the Senate. That's system one.

It's not perfectly correlated, so system 2 is entering. System 2 is kind of a deliberate system which may recognize a kind of mistake or snappiness of judgment in system 1, and then produce some kind cognitive override.

So system 1 is in effect when people are rapidly doing their homework, watching TV and not checking, system 2 checks, so Shane Frederick has developed some beautiful simple examples of this which are essentially brain teaser questions such as there is an algae growing in a pond, and the algae doubles every day. The algae is going to cover the entire pond in 48 days. In how many days is half the pond covered?

It's covered in 48 days. The correct answer is 47, one day before it's half covered because it doubles, but many people

divide 48 by two and say 24 days, and remember, you're smarter than the average or at least more analytically trained and college and often graduate educated. You can read Shane's paper in the Journal of Economics Perspectives I think in '05 or '06, and he reports the scores on these types, three items questions scale for people from many different universities, and MIT is the best. He doesn't have Cal Tech in there. We're going to do it at Cal Tech. We'll see. MIT is the Cal Tech of the east as it's called by us.

I think 15 percent of the students only get zero or one out of three questions right, so they're mostly very good at either knowing the right answer through just incredible intuition. They may have intuition, which is deliberative or realizing that 24 is wrong and very quickly correcting it, but students at other Universities that are less selective on analytical skill often get two or three of these wrong, including the one I gave you, okay?

So think about that. When we say system 1 mistakes, what that is is snap judgments which are plausible but often incorrect and often parametrically very far from the correct answer. 24 is far from 47.

The reason I think this is useful is it kind of provides a way to think about -- again it's vague and I asked Danny a couple

years ago, "Did you mean this to be an invitation to a neuroscientist to look for these two systems in the brain?" He said, "Yeah, yeah, nobody has done that yet, and I've been thinking about how to do it, and now I know why nobody's done it. Probably they're not two super separable systems like there was with visual cortex and stuff like that."

But there are many experiments that have tried to essentially invoke system 1 and not let system 2 come into play by doing things very quickly, for example, or provoke system 2 in some other way, and I think this general idea that there are two systems of this sort is pretty well established.

The reason it's useful is I think in a sense, when economists are critical of psychological results, the review is that system 2 must be in play. If they're not, then people are somehow confused or they're not incentivized, and if you paid them more -- system 2 doesn't work for cheap, so if you paid more money, people would automatically work harder at a system 2 variety. Like system 2, get off the couch, there's a lot of money involved here.

It also gives us a way, I was looking at Pauline Ippolito's discussion later, and in the '70s I gather just from her slides and some email, that the FTC policy involved prohibition of claims that have the capacity to mislead "the ignorant,

unthinking and incredulous." That's system 1 roughly speaking, and the modern policy as to the reasonable avoidance, and so I think one way to frame the policy question is: Is system 2 necessary for reasonable avoidance or should it be required that reasonable avoidance could be system 1 type reasoning?

Again, I wish we could be clear about this in a formal way, and I had time to give you the body of evidence, and I think one thing behavioral economics would be good at is to formalize these concepts and tell us also how to go in there and look at these things experimentally and their implications.

Okay. Limits on will power. I'm just going to mention this to set up the day, but you'll hear a lot are more from the real experts, so the idea we have here is that it's hard to resist temptation. Some people like George Lowenstein and Dorjian Prelick (phonetics) think that the way consumer optimization is implemented in the mind is probably something like there's a pain of pains, like a Lagrangian multiplier, and when you have things like credit cards with no immediate pain, that might lead to overspending blah, blah, blah.

We've done some experiments on dynamic optimization. If you think about things like saving for retirement, it's actually an extremely complicated problem if you take into account possible policy changes, taxes, longevity and so forth, so it may be that

it's just very hard to do, and people adopt rules of thumb which are sometimes pretty good and sometimes not.

Then importantly, I think we are very interested in irreversibilities. Those are decisions where an initial mistake is hard to undo, like your first few times using crystal meth, which is highly addictive or even possibly things like cognitive dissonance will bolster irreversibility, so you buy a treadmill and don't use it for six months, but you're in denial about how often you're going to use it in the future.

Okay. Some frontiers of behavioral economics, so this is really a response to the question: How come you guys don't have formal theories of these things? It's not economics if you don't have math, and all you show is experiments with students. Both of the things are false, and they're getting false.

So the AER and the QJE, not the JPD, is filled with very normal theories of some of the judgment mistakes I talked about. I haven't given a conclusive list. This is just to show you a taste, so one example is from Joel Schrag and Matthew [Rabin] who are here on confirmation bias. Their idea is that if you think, hypothesis A is likely; then if you see evidence against it, perhaps it's ambiguous or perhaps it's rapidly presented or it's in a forest of information that tends to often get discounted. This is very precise, I used to call these mis-

bayesian, but now I think a better term is bonded bayesian. These are models that take a standard bayesian framework, and they just build in one kind of small mistake, and they do a lot of analysis.

So they're actually generalizations of the standard model. That also means if you believe in the standard model, you are really comitting to these models with particular parameters, so there's a clear empirical fight.

A couple other people have thought about over-confidence and what that means and wishful thinking. There's a lot of research coming on about field data, just a tremendous amount, and I'm not even including here some papers you'll see today.

So several people have studied overconfidence, amongst CEOs, among investors. We'll hear a little bit about cell phone plans. Michael Grubb from Stanford has a new paper on this, a dissertation. David will talk about savings and household finance where I think probably the best -- there's been the most rapid accumulation of very good quality data.

In the area of limited strategic thinking, we have a paper we're working on and showing we think tentatively that when movies aren't reviewed at all, usually it's a sign that they're bad and they are bad. The very best one in our sample is Snakes on a Plane. That's pretty good according to Metacritic, which is

a big collection of critics. I haven't seen it. If you think that consumers don't realize that the lack of a movie review on Friday means the movie is bad, they'll go to too many of these movies and be disappointed and there will be a box office premium and we have some evidence for that.

This is a case I'm very proud of because it's my own work, but also these models that we estimate were actually developed initially to explain laboratory data, and the idea is to take them and go to the field, and in doing so you have to kind of tinker with the models at tryouts and functional forms, but it also addresses the long standing question of how generalizeable things we learn in the lab are to important field settings.

We'll hear from Dean about -- I forgot to put Dean Karlan's name here, but he'll remind you about field experiments, and that area is coming on extremely rapidly, extremely rapidly and it will just snowball because what will happens is firms and regulatory agencies see these things as useful and powerful and will invite more people to do them, and the samples are often quite big, very carefully controlled, and they're often important things in the economy.

Let me remind you of some empirical issues we have to think about. One is heterogeneity, so an interesting question Alan Schwartz probably may talk about is really one of the basic

questions here which behavioral economics has kind of avoided largely is: What if some people are sophisticated and some are naive because that's the truth? Do the sophisticated protect the naive by kind of -- like if you go to Mowbray Park in Los Angeles where I live and you know nothing about Chinese food, just go to where there are the fewest white people and you will have good Chinese food because you cannot run a Chinese restaurant in Mowbray Park and attract Chinese people and have it be crummy. The Chinese are very sophisticated food consumers and it will be cheap too, or do firms actually sort them out, kind of cull the herd, and take advantage of naive people through shrouding and things like that?

Second issue is endogeneity of thinking. I kind of alluded to this before in the system 2, system 1 discussion, but a really crucial question of which I'm a little skeptical about the view, but quite open to seeing lots more data is: Will consumers gather more information when they think they might need to? Is there kind of a kick-in system that says, "Gosh, I don't know much about mortgages, I better read, I better talk to my friends, I better ask questions?"

One reason to be skeptical about the force of this or the universality is evidenced from David Dunning and others showing that in tests of overconfidence, a lot of the overconfidence

phenomena is driven by the fact in these psychology experiments that people who don't know very much don't realize how little they know, right?

As opposed to people knowing very little and thinking, I have to learn more, so if knowledge is correlated with low self knowledge, what psychologists call metacognition, then the endogeneity hypothesis may not work because again this I think is one of the premier empirical questions, and behavioral economics is nothing if not about letting the facts help us inform economists, so this is something people should be researching much more carefully.

Finally, another argument that comes up that Jonathan will talk about is moral hazard, so is it possible that if there's too much public regulation, that will cry out consumer sophistication and the incentive -- it will deter the incentive to build up sophistication?

One way to study that by the way might be to look at adjustments in most communist regimes. So I've heard stories from places like North Korea where refugees coming to South Korea, you think they would be absolutely thrilled with all the free choice, but they're actually kind of miserable because they're kind of used to being taken care of, and they often get ripped off and they feel ashamed because the government's been

taking care of them.

I'm going to have to skip this unfortunately.

Private and public solutions: Learning is one. I think the issue there is how rapid the learning is. It may be that people learn from financial mistakes, but it may be that it takes until age 53 to do so. There's some evidence of that.

As I mentioned before, a lot of the important decisions we might worry about are early decisions that are somewhat irreversible involving durable purchase, houses, choosing colleges and even private schools, weddings and fertility and things like that.

Then the empirical question is how well advice markets work, so whether informal advice works will depend on kind of the wisdom of the crowds type questions. Are the people who you seek advice from and who advise you pay attention to actually the best of it, or are they just the tallest or the most honest looking or the most attractive?

Professional advice obviously needs to be put into the equilibrium analysis, so in most of these complex decisions, you can buy advice from people, and the issue there is whether agency problems actually help or hurt. They probably don't hurt so much that you would want to say ban professional advice markets, but they may mitigate some of the advantages of being able to buy

good advice.

How well does competition work? Well, that will depend upon a bunch of things. It will depend on the structure and the nature of the product and again human psychology, so I just got a new Palm Pilot, and I was really struck by the fact that it doesn't tell you much about how to use all the features in the manual. It goes -- I got the lowest end one, so it goes on and on about how you have to organize your life, which is true. It's like, man, they customized this to me.

So some of the effect of competition will depend upon on how well products will be bundled with advice about say good decision-making and also how well negative ads work and other status quo bias, so if you've brought a product that's really not right for you and I'm the competitor and point that out to you, do you believe me and switch?

There's some other devices that should be part of an equilibrium analysis, and I'm just listing them here: Class action lawsuits. Those may be the best tool for when millions of people are harmed in a very small way by some kind of contractual breach or overcharging like Netflix and Blockbuster cases. Reputations, I don't have that much faith in actually, and media reporting, and I guess this would be important for consumer reporting, although again I think media reporting is usually

about how great some products are rather than which ones are bad.

The public solutions are kind of the obvious things. Taxation, restricting contract terms and things like cooling off. This is what my law review paper basically discusses. I think it would be a good idea to not allow auto-renewal to be a default contractually, that you must opt-in to auto-renew, like a magazine subscription.

Again the point about these defaults is very simple. You have to choose a default, so forcing people to opt-in to automatically renew imposes a tiny cost, tiny, and in the age of the Internet it could not be smaller, on a rational consumer who knows for sure they want to auto-renew.

Consumer education's obviously great. I don't know the extent to which that's the FTC's province or they have any influence, but I think people should teach high school economics and some behavioral economics, the best examples might be don't make this mistake.

Some possible big mistakes, I basically talked about these, small overcharges that add up. Those may be hard to regulate. Maybe class action lawsuits are the best policy tool, important decisions that are irreversible.

A lot of things also involved very powerful wishful thinking and emotion, so diets, cosmetics, testimonials are very powerful.

I'm kind of curious about viral marketing where you sell products to people, and you say, If you like this, tell your friends, and the claim among people that do viral marketing is that people won't tell their friends unless they really like the products, but who knows how that psychology actually works. It may just be high-jacking the desire of the viral marketers to appear fashionable and trendy and the trust the friends have in them.

Okay. Conclusion: So I think if I were to pick two things out I would say that were researchable, I would say it's the interaction of sophisticated and naive agents. People have been thinking about that recently; the speed of learning about important purchases. If the argument is that learning will save consumers from themselves, the speed in which it takes place is absolutely crucial. If it's on the order of days or minutes on the Internet is one thing, but if it's years. My wife likes to watch The Bachelor, and every year you get one observation on the TV series The Bachelor about whether the bachelor actually marries the woman, and so it's taken seven or eight years for consumers to realize that the answer is no, because you only get one observation per year, right?

So the speed of learning becomes -- you really need to be able to parameterize that. It's not something that lab experience might not be very good at, but field experiments

might.

The only thing I'll point to is eye tracking, so Eric Johnson and I did the earliest studies in economics where eye tracking just refers to either a mouse driven system where you put things behind boxes, and you have to move the mouse to open them up like the Jeopardy game, or the system we use has a little rack that sits on your head that has cameras that look at your eyes. That means you can also measure the person's pupil dilation, which is correlated with arousal and cognitive difficulty.

It seems to me that one could do studies about informed consent and whether people are actually reading the fine print. You can use this technology which is free if you use Mouse Lab or close to free at the margin if you buy an eye tracker to see whether people literally read the fine print, and studies like that could be quite useful.

So if the variable you're interested in in your economic theory or your regulation is whether people are processing information or paying attention to things, we can actually measure those things quite directly, and that's very commonly done in psychology, like in studying reading. It's only been done in economics by me and Eric and David Laibson and Vince Crawford, and it's not that hard to do. It's not something you

normally learn in graduate school, but thousands and thousands of scientists have done it, so my pitch here is that there could be lot of things we could measure rather than treating information as an unobservable. We have to infer from choices. We can really go into the other tools, and that the psychology will just help us generate a wider vista of scientific evidence.

Mr. LEARY: Thank you very much.

(Applause.)

Mr. LEARY: Professor Klick?

PROFESSOR KLICK: So in some ways, this isn't quite a discussion of Colin's paper. I mentioned to Colin before he got up here that really my article with Greg Mitchell is more about the folks in the law and policy communities that have jumped on this research to say: Well, because of all this, we need to throw out half of the laws that exist and we need to change them and we need to change all of contract law and so on and so forth.

That kind of claim or those kinds of claims don't exist in Colin's work or most of the economists that work on this stuff, although Colin does have a law review paper, and so we'll blame any sort of undue advocacy on Zackeroff for that one, but Greg and I, Greg Mitchell who is a Ph.D. psychologist and a lawyer at UVA wrote our paper, "Government Regulation of Irrationality: Moral and Cognitive Hazards" to kind of present some blow back on

the sort of behavioral econ movement, which has, as I've said, sort of tried to jump from this research to totally sort of reconceptualizing most of what we do in the law and presumably the policy community.

Now, it's not solely an effect of this new economics literature. The original work on why standard form contracts are bad, Kessler and things like that, goes back to the first half of the 20th Century, but clearly there are lots of folks who were just sort of waiting for this research to sort of write dozens and dozens, if not hundreds of law review articles saying how this should change the way that we view things.

The main point that Greg and I point to is that we need to be worried about the extent to which paternalism in general reduces the incentive to invest in decision-making processes.

In some sense paternalism limits the opportunities for learning and could in principle lower the stock of human capital in both sort of domain specific areas, but also domain general areas, so what do we mean by that?

Well, there's lots of evidence for learning of various extents, both in sort of the newer economics literature and the older developmental psychology literature. You're going to see one of the good papers from the economics literature today by Laibson, et al., and they have some other papers on this issue

that suggest there is some learning as time goes on.

Now, they also show that there's depreciation of this stock of knowledge, but as Colin seemed to suggest, this is an area that we really need to do a lot more research on. John List has a bunch of papers on while we find in lab experience people exhibiting lots of biases in trading environments, if we go in and look at professional traders, we don't see this necessarily or we don't see it to the same extent.

Vernon Smith has done a bunch of stuff where he shows if you sort of run the same experiments on the same group of people, over time they seem to exhibit learning. Now, Vernon's stuff doesn't look over long time frames, so there may very well be appreciation going on there too, but at least we have some evidence that people do or that the biases that we notice aren't necessarily immutable, and I think we need to sort of worry about -- worry about sort of the endogeneity before we start changing contract law and tort law and policy.

Some of the work that Greg had us focus on is James Burnes' work on self regulation model of decision-making where he stressed the value of feedback from successes and failures for improving decision-making, and there have been lots of folks in that literature who have done some tests, both with children and unfortunately fewer tests on adults, sort of looking at what

happens when we allow feedback and when we don't allow feedback from, quote, bad decisions.

Sort of the kind of paternalism that's been championed by a number of the law professors and policy folks, how can it affect these sort of moral and what we call cognitive hazards? Well, if you think of ex-ante paternalism, so the paternalism where you restrict choices upfront or where you incentivize the right choice, in principle this could restrict the opportunities to develop successful decision-making capabilities.

That is, if it mutes the positive and negative feedback from good and bad decisions, there's going to not be as much or potentially any learning going on from these good or bad decisions. How about ex post paternalism, which actually has existed in sort of legal theory for quite some time, so bailouts or somebody signs a standard form contract, do we enforce it or do we not?

Well, in principle, this is sort of the standard or this has the potential to generate the standard moral hazard problems that we often worry about. Does it reduce the motivation to invest in actually learning about or even considering what's actually in the contract?

I think we have some cites in the papers to cases where people are now citing work in behavioral economics and behavioral

law and economics in their briefs in court cases and they say, "Look, I shouldn't be held to this contract because as this evidence all shows I couldn't possibly have paid attention to the 57 terms in this contract."

Well, is that sort of somebody acting honestly or is that somebody sort of exhibiting the kind of strategic moral hazard that we might be worried about? That's an important question to answer, and also this ex post paternalism is also going to meet the negative feedback from bad choices even if we're not worried about the strategic type of behavior.

So what do we mean by the different domains of learning? Well, there's both a specific context; that is, if we have paternalism at time 1, is it the case that this is going to increase the need for paternalism in the same context in time 2; that is, if it were the case if we didn't have paternalism, if somebody does learn from their bad choices, presumably in a world where we don't have paternalism, when time 2 comes along, to the extent that they've learned, we no longer need to protect them.

If there's no learning, well then it necessarily would be the case that there's no improvement in period 2, but if there is learning, providing paternalism in period 1 means we're necessarily going to have to provide it in period 2 to get the same kind of outcome.

What's potentially even more important? Well, what if it's the case that decision strategies are transportable across domains, right? So paternalism in domain A might lead to a worse outcome in domain B. So what if it's the case that we act paternalistically to induce people to save early on in their lives for college or something like that, and then that leads to problems when the person needs to go then save for retirement.

Maybe it's the case if we can learn and learn from mistakes in sort of the earlier domain, we won't necessarily need as much intervention in the other domain as well. It may be that there are even external effects; that is, perhaps when we see others being protected from their bad choices, this may induce even, quote, good decision makers to cut back on their decision-making inputs, and that in principle is something that we may want to be worried about.

As usual, this all comes down to the elasticities, right? How sensitive are people to the feedback, right, both in the specific domain and the general domain. We need to worry about this. We need to get some parameter values for this. We need to also -- as the work of Laibson and others point out, we need to worry about: Is there depreciation of these effects over time?

Those are important things that I don't think for most contexts, even for any context in terms of robust results that we

have good sort of parameter values here. How much do individuals cut back on their own decision-making inputs when their outcomes are less dependent on their actual decisions, right?

Furthermore something that is interesting from the psychology experiments is: We would like to know how useful it is or how effective it is to have education, right? There are some experiments out there by psychologists that suggest people are bad bayesian updaters, but if we sort of give them some lessons early on in the experiment, they can improve over time.

Now, is it the case that the education in a particular domain where we want to act paternalistically, education is actually even more effective than the protection or even cheaper than the protection? That would be something that would be interesting to know for policy decisions and legal decisions. As Colin noted, just as there is heterogeneity in the biases themselves, there's likely to be heterogeneity in these relevant elasticities.

Now, here we have Kobayashi, the hot dog eating champion, and we might be worried about, well, do we need to induce him not to eat those 53 hot dogs, right? If we stop him from doing that, are we going to stop him from learning about the bad effects here, right? Is he never going to get that sort of human capital? Furthermore, are we going to stop him from sort of

reaching what's the right outcome for him. See Kobayashi eating 53 hot dogs, he wins a big metal. He wins a big trophy, right? So what are the kinds of things that we want to think about when we decide whether or not to stop Kobayashi from eating these hot dogs or from being paternalistic in general?

Well, in the paper, Greg and I sort of layout the kind of input considerations you would like to think about before making legal or regulatory changes. First, you want to look at the efficiency loss from current underinvestment in individual decision inputs; that is, how much do we lose from people's own investments and decision-making processes by protecting them?

In that particular instance we also want to look at the capitalized cost of future efficiency losses related to this current underinvestment. This is sort of the cognitive hazard point. We sort of view the moral hazard as the current decision-making process and the cognitive hazard being what's the effect on future decision-making processes through the loss of this decision-making capital?

Furthermore, we want to look at the cost of these private inputs, right? If it's the case that yes, there's moral hazard but people themselves investing in these inputs is quite expensive relative to the paternalistic alternative or maybe paternalism is fine even in the face of this moral hazard, but we

would sort of need some framework or some way to sort of compare these costs I think before we go and make large scale regulatory changes.

Furthermore, we need to look at what's the efficacy of education, how much does education cost in these various situations. We also, and these are -- number 4 and 5 are the ones that at least the folks in the legal and policy literature have focused on so far. We want to look at: What are the efficiency gains from limiting decision-making, right, so the asymmetric paternalism that Colin writes about says: If there are some people who gain a whole lot from the paternalism here and the people that lose a little bit, that's an interesting thing to know, and that's obviously quite important when making these decisions, so we need to look at those but we need to look at these other things, too.

Number 6: Are there welfare gains from discouraging non-productive exploitation of biases, right? Presumably if it's the case that people are biased and we sort of don't have any regulatory intervention there, it may well be the case that some firms set out to sort of exploit those biases so they invest resources in that exploitation that are largely a loss to society, so that's something we would like to worry about, too.

But sort of the flip side of that is there's going to be

rent seeking on the regulatory side as well. If we start making decisions on what food goes on the front of the buffet and what food goes on the end of buffet, if that sort of makes a real difference in terms of profits, it's going to be the case that people are going to engage in some sort of rent seeking to be the food that's on the front of the display, right?

So we would like to know that. We would like to worry about that, and that's an area in this whole literature I think that has so far gone unremarked upon, right? These regulatory decisions aren't just going to be made by the Colin Camerer's and the Cass Sunsteins of the worlds. They're going to be made by political folks so that necessarily is going to create some costs that may well, in certain contexts, make or entirely wash out the benefits from this kind of paternalism, and lastly of course as with anything we need to worry about just administrative costs in general.

I think at least mine and Greg's impression of what's gone on, again not in economics literature so much but in the legal and policy literature is no one has sort of worked with any kind of consistent framework on when we can decide, yes, we take this information on biases and things like that, and we need to change law and policy.

People so far at least in those literatures have simply

said, Look, this research suggests that contracts can't work, people can't really pay attention to contracts. Well, there are lots of concerns here, and I think people on that side of the literature need to be much more concerned about those.

Thanks.

(Applause.)

Mr. PAUTLER: I wanted to provide a few comments about asymmetric paternalism which is a paper that Colin referred to earlier. Because he didn't actually present it, I ought to give you a little bit of an idea of what's in the paper.

Asymmetric paternalism is essentially an approach of thinking about a cost benefit test in which you try to figure out whether the benefits of a government remedy to the cognitively impaired will be greater than the cost to those who are cognitively aware. You could change the words for those if you want. The cost-benefit trade-off also considers the implementation costs of the government providing the remedy and the profits that might be lost to the firms who previously used to make profits because of the cognitively impaired people.

By design, what asymmetric paternalism does is try to look at problems where the remedies are not terribly invasive. They have little impact on rational consumers, but would provide significant gains to the people who need protection.

The question Joe Mulholland asked me originally to think about was: Well, how does what the FTC does fit into the asymmetric paternalism framework? What we do fits into the framework in the sense that we worry about costs and benefits, but that of course doesn't get you much of anywhere, because everybody knows where the real action is, and deciding what kind of public policy you want depends on how you measure the costs and benefits, and what kinds of evidence you use to make that decision.

Using individual survey data probably isn't going to be very useful if the data that would come out of a market is very different. On the other hand, market-based data can also be misleading. If a large enough percentage of the population falls into the category that is cognitively impaired, then the market might produce evidence that doesn't really show you what the correct decisions would be by consumers who were in fact correctly making decisions based on the information set they should have had.

So it's hard to know whether you want to use survey information, whether you want to use market data. Certainly we've seen surveys where people will simply be asked, "Do you like privacy?" and of course the answer is "yes," but the important point is that it depends on what it costs you to obtain

more privacy.

I did want to just mention a few instances in which defaults really seem to matter, and that's one thing that the FTC has worried about a good bit. We worry about it in the area of privacy a lot. Many defaults that are chosen result in significant potential costs, mainly because we're not really sure what's going to happen in the future.

For instance, in setting a few examples of the case of privacy, when you choose a default, and we aren't the ones choosing these defaults normally. Congress is, or some other regulatory body. Almost all privacy policies have both hidden defaults and defaults that you can obviously see. A couple of examples, recent offers of credit, we now allow an 800 number opt-out so people who can manage to reduce their costs of getting junk mail.

In credit history use, which is the first item on the list, you can now have a new option, at least in some states, to freeze your credit report, so people can't use the information to make additional offers of credit. Now, what you've done when you do that is you've opted out of the "instantaneous availability of credit to strangers," which is the incredibly interesting thing about our credit system currently, that actually strangers can get credit almost instantaneously.

I think I want to move on to the one area of a default that Colin mentioned in his discussion. I think he called it "auto-renew." Around here we tend to call it negative option, they are the same thing.

In the case of negative option transactions, what happens is the normal default in the market has really changed immensely. In negative options, "silence is acceptance," and oftentimes there is no offer made to the consumer. The way this works in a lot of the cases the FTC sees is you get a free offer of something, and you get a freebie for a month, and then you are signed up for the indefinite future, and the question is, did you know it, and in an awful lot of our cases we it's pretty clear the firm doesn't tell you much of anything, and in fact in the worst cases, instead of you go from being free to paying, it's switching from a free trial offer to theft. They didn't tell you anything about the follow-on period. They got information to debit your account, provided you with something in period one, and then they sell you the product in the future, but you don't even know you bought it, and that happens more often than you might think.

We have brought a number of cases along those lines, but as you back up from the obvious case of the free trial to essentially theft, you move back toward free trials with not very

good disclosure, and how good does that disclosure need to be in order for us to decide that there's really no deception or unfairness going on in terms of negative option purchases?

With all the defaults we see, I think one of the things we have to worry about at the FTC - and I don't know that we're capable of figuring out all the costs - but we really want to worry about the costs that are imposed in the future when defaults are chosen by any policy maker. In the privacy area and in a number of other areas, those defaults can affect how markets evolve, and I think that's probably the biggest cost that we need to be aware of, and where we need the most help assessing our choice of defaults.

(Applause.)

Mr. CALFEE: Good morning. I do not have a PowerPoint, so we will not deal with that. I'm Jack Calfee. As Paul mentioned earlier, I'm an economist at the American Enterprise Institute here in Washington not very far away, but I spent a half dozen or so years at the FTC in the Bureau of Economics in the early 1980s, and it's always a delight to be back at this agency and talking to old friends and so and rediscovering that many of the issues that we mulled over in the early 1980s are still relevant, in some cases even still cutting edge.

I read the paper that Paul referred to by Colin, a very

interesting paper on basically trying to rescue the behavioral economics approach to regulation from a conservative perspective. That's a distortion of the paper, but it's a very interesting one which I strongly recommend to you, and that is the spirit of it is sort of to reconcile some of the kinds of regulatory initiatives that have been suggested from people who study this thing called behavioral economics and to point out that in many cases, those recommendations can be reconciled with a fairly free market approach to regulation.

I have my doubts about some of this in practice, which I wanted to go through very quickly. We are here at the FTC. The FTC is not the only agency that regulates the kinds of things that are relevant to behavioral economics or not the only agency that regulates information a lot. I'm thinking about agencies like the Food and Drug Administration, the Securities and Exchange Commission, especially the FDA, which I know the most about these days.

And what strikes me in looking at the research agenda that's been coming out of the economics profession in the last 10 or 15 years, what people are talking about today, is that there's been a great deal of research at the consumer level how people do things. There's been some research on the kinds of decisions that are made as a result of institutions. There's been

amazingly little research on regulation itself and what the effects of those regulations are and what are the results of implementing policies that are different from the way they were 10 or 15 or 20 years ago or the policies that I would suggest are actually more important, the ones that have not changed for the last 10 or 15 or 20 years.

And I think here there's one point to keep in mind at the very beginning, and that's a distinction between regulations that rely mainly upon rules or regulations per se as opposed to a case by case approach, and although some of the FTC's consumer protection rules are well known, most of what they do is done on a case by case enforcement of a prohibition on unfairness and especially deception.

I think there are a lot of advantages to that, but most of the discussion in the behavioral economics in the literature that I've been mulling over, mulling through in the last few days tends to focus on rules rather than a case by case enforcement, and I think it's important to bear in mind there is another approach, and there are actually I think reasons to believe that the other approach, that is the case by case, may avoid a lot of problems that arise in connection with rules.

Thinking a little bit about the FDA which is a very rules oriented place, and in fact they don't really bring many

enforcement actions because they're an unusual agency and they regulate the industry in a rather unusual way, and in most cases they don't have to bring an enforcement action against the pharmaceutical firm. The firm does what the FDA tells them to do, which mainly is to obey FDA rules, which are very restrictive, and so they have rules about advertising both to consumers and to physicians, and those rules prescribe in great detail the content of advertising, and they prohibit information that is not permitted as it were.

There's a lot of warning information, and there's a small but alarming literature about the over-warning effect, which in fact the FDA staff are very aware of and which they try to avoid when they can, but often they can't, and the most recent example is the black box warnings for antidepressants that are given to children and adolescents where some very strong warnings that came from the FDA at the behest of outsiders has greatly reduced prescribing to children and apparently there's no or little benefit from that, and it looks like there's probably a pretty substantial cost to what looks like a textbook case of over-warning.

And the flip side of this is when the information is so tightly regulated, it's really hard for competitive forces to come in to play, so there's a lot of important information that's

quite relevant to pharmaceutical, foods, et cetera which we don't hear much about because advertising is restricted, and you literally do not get much competition in terms of information about pharmaceuticals.

I wanted to say one thing about a topic that has been barely mentioned, if at all today, and which as far as I can tell was barely mentioned in Colin's paper, which is the use of information benchmarks, sort of like the EPA mileage rules, the FTC's tar and nicotine ratings for cigarettes, food labels, et cetera.

You would think that one thing the government could do at very little loss, it would be really an example of the kind of behavioral-based interventions that's might be very useful, would be to establish some kind of benchmark. There are a lot of different ways to measure something. We, the government, we're going to put out this particular way to do it because we think it works quite well, and if you use it, that will facilitate competition, et cetera, which in theory should be an example of what you could think of as what Colin calls asymmetrical paternalism, that is the kind of intervention that might prevent some problems with consumers who are not very well informed who have problems with self control without really inhibiting benefits to consumers who don't have those problems and just go

ahead and use the information as it appears in the market.

There's quite a few of these benchmarks out there, and my sense is that they don't really work very well, and the reason they don't work very well is first of all, instead of being an adjunct to competitive markets, they often are replacement for competition in terms of information.

I'm old enough to remember when there was a lot of advertising about fuel economy until all that advertising was replaced by the EPA mileage calculations. The same thing was true of tar and nicotine for cigarettes until that was replaced almost 40 years by the FTC's rating and other things, food labels, et cetera.

And what we find with these things is that once the benchmark is in place, once the alternative information has been more or less prohibited from the market, these benchmarks tend to get gained. There are parties that have vested interests in them. There is rent seeking. It becomes very difficult to change the benchmarks. It's very difficult to update them.

I don't know how often the EPA mileage calculations are updated, but it isn't very often, and the tar and nicotine ratings are almost never updated, and it gives them a way that would otherwise be useful competition. In fact I would say with the tar and nicotine ratings it's probably a very wide spread

consensus on two points: One being the ratings work very, very badly; the second being that they should not be changed or dispensed with.

So I guess what I'm suggesting in general is that since one of the topics of today's conference is research agendas and what we ought to see more research on, what I'm suggesting is we need a lot more research on how regulations actually work with special attention to situations in which interventions that are based upon some of the more interesting findings of behavioral economics, interventions that in practice inhibit competition or in some cases outright prohibit competition to see how these things actually work in practice.

And in many respects, this debate, things to worry about and so on are very reminiscent of the things we were talking about in the early 1980s that was not called behavioral economics, but it was very much the same kind of issues. We've had a lot of research since then so we know a lot more about how consumers do things in the marketplace, next to the marketplace and so on.

We still don't know nearly as much as I think we should about how the regulations actually work, and in particular regulations that are either motivated by or are consistent with or at least strongly informed by behavioral economics. So those are my comments. Thank you very much.

(Applause.)

Mr. LEARY: Thank you, John. Given the time, I think we are going to defer questions. Professor Camerer will be taking part in a roundtable this afternoon, and that should offer the opportunity for some additional questions, but if you have any particular responses you would like to make based on the three discussions that followed yours, I want to give you a couple minutes to follow-up.

Mr. CAMERER: Yeah, I'll just try to take maybe just one minute. I think these are -- this is exactly the right kind of research agenda, which is if mileage ratings are easily accessible and consumers come to trust them or they're gained or they're not updated correctly, certainly it's not ideal. I don't know if that means you shouldn't say widely publicized mileage ratings.

I think we're going to hear later today from Jim Lacko and Jan -- Janis it's listed here as Pappalardo, yeah, that's right, and I'm sorry I don't know who you are but I will soon -- about disclosure, and Alan Schwartz and I had this discussion maybe eight or ten years ago, and he was the first person for whom it resonated with me that if people make mistakes like they don't appreciate risks, warning them about risks is very tricky.

If those same people have limited attention, and I think the

main thing is they don't necessarily process the information we do, which is by adding up numbers and calculating things and computing, so disclosure is a very tricky art, and the way you think of it is one principle in behavioral economics is people may be uninformed, so you give them information.

Another principle is they may not process information like we do, so the form of the disclosure is really a matter of mostly now art and a little science, and it would be great if there was a little more science and less art.

I totally agree with Jack also about studying these regulations because it's only by looking at case studies and maybe cross country studies and cross domains that we are really going to learn about these things, and even though I'm a big fan of lab experiments, I think they've provided sort of the fuel and the basic insights that gave some of these phenomena a name, and examples: There's certain things that lab experiments aren't particularly well built for like long longitudinal studies for high stakes decisions, and it's only going to be case studies and field experiments that will get us further.

Mr. LEARY: Thank you very much, and we'll reconvene in ten minutes.

(Whereupon, a brief recess was taken.)