9 A Placebo Trilogy

Academic researchers are specialists: they develop expertise in a particular area of study. Therefore, we expect their work to exhibit some amount of repetitiveness. Economic theorists who develop an expertise in a subfield are likely to revisit the same class of models or economic environments.

In this chapter I address a different kind of recurrence in theorists' work; a more personal and "artistic" recurrence of themes or mottos that are in principle independent of modeling choices or substantive economic questions.[1] I am thinking of the analogue of things like the Coen brothers' recurring theme of "evil meets stupidity"—or even smaller-scale mannerisms, like the image of a fat old man yelling at the protagonist from behind a large desk, which appears in many of their films. These are motifs, themes, or schticks that, for whatever reason, the artist keeps returning to, but placed in such different contexts that this ceases to be a simple matter of repeating oneself.

Within the world of economic theory, a rare example of a theorist's attempt to introspect about "recurring motifs" in his work is Avinash Dixit's semi-humorous piece "My System of Work (Not!)":[2]

> As you can see, my approach to research is too opportunistic to have a constant direction. But taking stock of it for the purpose of writing this piece, I could see a recurrent if not dominant theme. Scale economies and sunk costs keep appearing in my papers with great regularity. Imperfect competition is the norm, and market equilibria are not socially optimal (but government interventions have more subtle effects than naive intuition would suggest, and may actually make matters worse).

If you think that the above quote holds no interest ("Who cares if you like to put sunk costs in your models? What are you, Fellini?"), Dixit offers the following broader interpretation of his professional mannerisms:

> And therein lies an irony. The left-wing critics of the late 1960s and 1970s, who influenced many youngsters when I started out, reserved their strongest criticism for the perfectly competitive equilibrium of the neoclassical system. Of course they did little by the way of offering a viable alternative. It has been the unexciting incremental work, to which I have contributed a little, that has built into a major shift in our understanding of how the economic system operates when the assumptions of neoclassical economics fail.

Now this is an observation with interesting political connotations. The "personal tics" of sunk costs and economies of scale suddenly become Dixit's way of expressing a "third way": a moderately progressive approach to economics, against the background of radical left-wing progressivism that surrounded him during his formative student years.

This chapter is an attempt at a more systematic introspection, using my own work for raw material. I hope this won't come across as too self-indulgent. I am perfectly aware that few would take intrinsic interest in recurring mottos in an economic theorist's oeuvre. I am not Fellini. Nevertheless, I think that whoever takes an interest in the culture of economic theory may find something to learn from such an introspective exercise. And I hope the examples themselves have entertainment value.

I refer to my recurring motto as the *placebo theme*. It is the idea that certain actions or products may have zero intrinsic value but generate stable demand nonetheless, due to some error of reasoning committed by economic agents. I have been fascinated by this idea for many years and I keep returning to it. I'd like to share three of these examples, which are based on a trio of papers (Spiegler 2006, 2013, 2016). Though different from each other in terms of economic substance and modeling technique, they all share the placebo theme. I conclude by speculating about the political and psychological significance of this fixation.

Consumers: The Dieter's Dilemma (2016)

We are all familiar with the motto "correlation doesn't imply causation." Economists use it to admonish or make fun of laypeople (or, even better, other scientists) for jumping into conclusions about causal effects from observed correlations. Yet, in the spirit of behavioral economics, shouldn't we also try to *model* how such people reason, and what the behavioral implications of confusing correlation with causation might be? Of course we should. So here's a story.

Imagine a population of identical consumers. Each consumer chooses whether to buy a food supplement, thinking this might have an effect on her long-term health. In reality, the food supplement has no such effect: the consumer has a 50% chance of being healthy, independently of whether she consumes the supplement. In the population of consumers, we will see two statistically independent variables: consumers' choices and their health outcome. However, if consumers realized their choice has no effect on their health, they would all refrain from buying the costly supplement, and so there would be no variation in the consumer-behavior variable.

Now add a third variable to our story, the blood level of some chemical—the like of cholesterol, vitamin D, or serotonin. In reality, the chemical blood level is determined by the other two variables: it is abnormal if and only if the consumer's underlying health is poor *and* she doesn't consume the supplement. Since the consumer doesn't care about the chemical blood level per se, and since it plays no role in the mapping from her action to the health outcome, the consumer should ignore this variable altogether.

However, suppose that our consumer believes that the three variables are causally related in a way that can be described by the following diagram:

$$s \rightarrow c \rightarrow h$$

In this diagram, s stands for the supplement consumption quantity, c stands for the chemical blood level and h stands for the health outcome. In what follows, consumption quantity can take only the values 0 and 1. Likewise, good and bad health are denoted $h = 1$ and $h = 0$. Normal and abnormal chemical levels are denoted $c = N$ and $c = A$.

This diagram is an example of a *directed acyclic graph* representing a causal model: s is perceived to be a direct cause of c, which is in turn perceived to be a direct cause of h, and there are no other causal transmissions. There is a wonderful literature at the intersection of artificial intelligence and statistics that applies directed acyclic graphs to probabilistic and causal inference. Judea Pearl's recent *The Book of Why* is a highly recommended entry point.[3] The model I am about to describe makes use of very basic concepts from this literature.

Our consumer forms a belief about the health implications of her consumption decision by *fitting her causal model* to objective observational data. This means measuring the empirical distribution of chemical blood levels conditional on supplement consumption, as well as the empirical distribution of health outcomes conditional on the chemical blood level, and putting them together in accordance with the consumer's causal model. Formally, the consumer's subjective probability of h conditional on s is

$$\Pr(c = N \mid s) \cdot \Pr(h = 1 \mid c = N) + \Pr(c = A \mid s) \cdot \Pr(h = 1 \mid c = A)$$

The consumer regards this conditional probability as a causal quantity, measuring the causal effect of her action on the probability of being in good health. As a result, her subjective estimate of the health effect of buying the supplement is

$$[\Pr(c = N \mid s = 1) - \Pr(c = N \mid s = 0)] \cdot [\Pr(h = 1 \mid c = N) - \Pr(h = 1 \mid c = A)]$$

All the terms in this formula can be measured from the objective empirical joint distribution over the variables s, c, and h. However, the formula delivers a *wrong* estimate of the causal effect of s on h, because it is the result of imposing a wrong causal model on objective data, and therefore a wrong causal interpretation of observed correlations.

The correct causal model that underlies the objective joint distribution over the three variables is represented by the following graph:

$$s \rightarrow c \leftarrow h$$

The consumer's subjective model inverts the true causal relation between chemical levels and health. In reality, the chemical level is a consequence of the consumer's health condition, yet in his own mind the direction of causality is flipped. In other words, the consumer is committing a *reverse causality* fallacy.

Let us now compute the terms in the formula that describes the consumer's estimated health effect of buying the supplement. Denote $\Pr(s = 1) = q$. This is the fraction of consumers in the population who buy the supplement.

- $\Pr(\boldsymbol{c} = \boldsymbol{N} \mid \boldsymbol{s} = 1) = 1$. Because by assumption, each consumer in the population ensures a normal chemical level if she consumes the supplement.
- $\Pr(\boldsymbol{c} = \boldsymbol{N} \mid \boldsymbol{s} = 0) = 0.5$. Because by assumption, if a consumer doesn't buy the supplement, her chemical level is determined by her underlying health condition. Since half the consumers in the population are healthy, it follows that half the consumers who do not buy the supplement have a normal chemical level.
- $\Pr(\boldsymbol{h} = 1 \mid \boldsymbol{c} = \mathrm{A}) = 0$. Because by assumption, the consumer must be unhealthy in order for her chemical to be abnormal.

- The fourth term, $\Pr(\boldsymbol{h}=1 \mid \boldsymbol{c}=\boldsymbol{N})$, requires a more elaborate calculation. The fraction of consumers with normal chemical blood level is

$$\Pr(s=1)+\Pr(s=0)\cdot\Pr(h=1)=q+(1-q)\cdot 0.5$$

The fraction of consumers who are healthy and have a normal chemical level is 0.5. Therefore, by Bayes' rule,

$$\Pr(h=1|c=N)=\frac{0.5}{q+(1-q)\cdot 0.5}=\frac{1}{1+q}$$

Plugging these four terms in the formula, we obtain the consumer's estimated effect of buying the supplement on the probability of being in good health:

$$[1-0.5]\cdot\left[\frac{1}{1+q}-0\right]=\frac{1}{2+2q}$$

Two things about this formula are noteworthy. First, it is strictly above zero, which is the objective health effect of buying the supplement. In other words, the consumer's reverse causality error leads her to assign value to something with no intrinsic value.

Second, the consumer's subjective estimate of the health effect of her action depends on the action frequencies in the *consumer population*. This dependence would never arise if the consumer had a correct model: by definition, the conditional objective probability $\Pr(y \mid x)$ is invariant to the marginal probability $\Pr(x)$. This invariance does not extend to the *subjective* conditional distribution, which is based on a wrong causal model.

Moreover, the formula is *decreasing* in q. The larger the fraction of consumers who buy the supplement, the lower its subjective value. For example, suppose that the cost of buying the supplement is 0.4. Then, as long as $q<0.25$, consumers will conclude that the benefit from buying the supplement exceeds its cost. Conversely, when $q>0.25$, consumers will conclude that the cost exceeds the benefit. This suggests a natural *equilibrium* characterization of subjectively optimal behavior: when $q=0.25$, all consumers are indifferent between the two actions because the estimated benefit of buying the supplement is equal to its cost. Therefore, the pattern in which 25% of the consumer population purchase the supplement is stable.

We have thus learned two lessons from this example. First, reverse causality can lead consumers to assign value to intrinsically worthless products. Second, subjective optimization under reverse causality needs to be analyzed as an equilibrium phenomenon, even though we are dealing with *individual* decision-making.

I am proud of this little parable, which I called *the dieter's dilemma*. With it, I managed to produce a coherent description of a decision-maker who "mistakes correlation for causation." In Spiegler (2016), I developed a broader modeling framework out of this little example, which put the language of directed acyclic graphs to productive use and enabled me to study in greater generality behavioral implications of causal misperceptions. Yet it all started with this "placebo" story about the poor consumer who throws good money at a no-good remedy.

Industries: The Market for Quacks (2006)

Our next story turns from the individual consumer, contemplating the purchase of an intrinsically worthless product, to an entire industry devoted to selling such products.[4]

Imagine a population of identical consumers having some underlying problem. Every consumer is willing to pay 1 for anything that will fix her problem. The consumers enter a market with n revenue-maximizing firms that propose solutions to their problem. Firms incur an arbitrarily small cost when selling their products. Each consumer must choose one of the $n+1$ available alternatives: the firms' products and the outside option ("doing nothing"). The probability that a consumer's problem is fixed is q, independently for each consumer, and—most importantly—independently of the action she takes. In particular, firms' products have no advantage over the outside option! For this reason, I refer to the firms as "quacks" and to this industry as a "market for quacks." If consumers understood how this market works, they would have no business entering it: in order to turn a profit, firms would have to charge a strictly positive price, and yet the outside option is free of charge and offers consumers the same quality. Therefore, the market for quacks would be inactive if all market agents were individually rational.

Every reader will think of her own real-life analogue of this market for quacks. For some, many kinds of alternative medicine would qualify. This is actually trickier in practice because of placebo effects: a patient's mere belief in the healing power of a potential remedy may trigger physiological processes that effectively endow it with actual healing powers. My example assumes away such placebo effects. A more indirect parallel (and an obviously more important one) is with the market for active money management. A strong form of the *efficient market hypothesis* maintains that since market prices of financial assets reveal all underlying information, active money managers have no advantage over a passive index fund. From this perspective, active money managers are "quacks" relative to the outside option, which is a relevant index fund.

Back to the parable. Assume the firms play a simultaneous-move game: each firm independently chooses a product price between 0 and 1. This is the script of the textbook model of Bertrand competition. The difference is in how consumers make their choices. In the standard model, they correctly perceive that all firms sell the same product and know how to evaluate this product. If the product had value relative to the outside option, they would choose to buy from the cheaper firm, as long as its price is below the product's value.

In contrast, in the current model, consumers rely on a *sampling procedure* to evaluate alternatives. Each consumer obtains an "anecdote" about each of the $n+1$ alternatives. An anecdote about an alternative is an independent random draw from the binary lottery that is associated with it: success in fixing a consumer's problem with probability q, and failure with probability $1-q$. Having gathered her $n+1$ anecdotes, the consumer chooses the best alternative in her sample. This means choosing the cheapest among all successful alternatives in the sample. If the sample contains no success story, the consumer opts out. Some tie-breaking rule is needed in case of ties in the sample. Although this is immaterial, for the present purposes it is simplest to assume that the consumer breaks ties between a firm and the outside option in favor of the firm.

The consumers' choice rule captures in a stylized manner two familiar psychological biases. First, the consumer behaves as if she believes that a small sample is representative of the underlying probability distribution. In particular, if the consumer hears a success story about one firm, she acts as if she believes that the firm will fix her problem with certainty. This type of exaggerated inference from small samples was studied by Tversky and Kahneman (1971) who gave it the ironic name "the law of small numbers." Second, anecdotes are by nature rich with contextual detail, and this makes them more memorable and affective than dry statistical data. This "vividness" of anecdotes makes them more likely to sway impressionable decision-makers, against their better judgment.

The consumers' choice model defines the payoff function in the firms' simultaneous-move game. For instance, when the firms' prices satisfy

$p_n > p_{n-1} > \ldots > p_1$, firm 1's market share is $q(1-q)$. The reason is that the firm's clientele consists of consumers who heard a success story about its own product and a bad anecdote about the outside option. Whether or not they heard a good anecdote about other products is irrelevant, because they are more expensive. The firm's revenue is $p_1 \cdot q(1-q)$. By the same logic, firm 2's revenue is $p_2 \cdot q(1-q)^2$. For any $k > 2$, firm k's revenue is $p_k \cdot q(1-q)^k$. Since the firms are conventional revenue-maximizing players who play a simultaneous-move game, using Nash equilibrium to describe a stable outcome of their strategic interaction is also conventional.

Looking at the expression for firms' market share, we can develop an alternative interpretation for consumer demand, which is entirely consistent with conventional rationality. While we assumed that firms sell an objectively homogeneous product (which happens to be worthless), suppose instead that they sell *differentiated* products, for which consumers have idiosyncratic tastes. For each consumer and each alternative, the consumer values the alternative at 1 with independent probability q and at 0 with probability $1-q$. In other words, q is the probability that the consumer likes the alternative. From this point of view, q is a parameter that determines the distribution of consumer valuations. These valuations reflect intrinsic, subjective tastes, rather than random estimates of commonly valued, objective quality.

This picture of consumer demand is formally equivalent to our sampling-based procedure. Therefore, Nash equilibrium analysis of firms' behavior will be the same as if they were facing rational consumers with differentiated demand parameterized by q. In this sense, there is nothing really "new" in the market-for-quacks model. What is unusual is the interpretation of consumer demand and the normative conclusions one draws from the market equilibrium analysis. Note that in many settings, the differentiated-taste reinterpretation will be patently absurd (can we imagine it being relevant when firms are money managers?) and therefore only useful as a purely formal analogy.

There is a unique Nash equilibrium in this game, which is therefore symmetric. Firms play a mixed strategy: a continuous probability distribution over the interval $[(1-q)^{n-1}, 1]$. The exact distribution does not matter. What *does* matter is two immediate conclusions we can draw from what we already know.

First, the market for quacks is active. Indeed, the industry profits it generates in equilibrium are equal to $nq(1-q)^n$. Why is that? Mixed-strategy Nash equilibrium requires that every price in the interval

$[(1-q)^{n-1}, 1]$ is optimal for a firm against the mixed strategies played by its rivals. In particular, the price $p=1$ must be optimal. It is easy to calculate the market share that this price generates. When a firm charges $p=1$, it is the most expensive alternative in the market. Therefore, only consumers who hear a good anecdote about this firm and a bad anecdote about each of the other n alternatives will pick the firm. This implies a market share—and therefore a revenue—of $q(1-q)^n$. This is the payoff that an individual firm earns in equilibrium, and there are n such firms in the market.

When we examine the expression for equilibrium industry profits, we see they are increasing in the number of firms n as long as it is below some critical level (which can be awfully large if q is small). How can greater competition lead to higher industry profits? Well, in the market for quacks, supply creates its own demand. The larger the number of firms, the higher the chances that a consumer will hear a good anecdote about one of them and decide to enter the market on the strength of this anecdote. This "aggregate demand" effect is countered by the standard competitive effect (the larger the number of firms, the lower the expected price they charge), but for low n, the former effect dominates.

The second interesting observation is that when q is close to zero—that is, when the consumers' underlying problem is nearly hopeless—equilibrium prices in this market for quacks will be very high, close to the "monopoly price" of 1. The reason is simple. When a firm gauges the competitive forces it faces, it cares about the number of good anecdotes in the consumer's sample. When q is small, this number will be small: with high probability, the firm will be effectively a monopolist if the consumer hears a good anecdote about it. Therefore, the firm faces very little effective competition, which drives its price upward. The differentiated-taste reinterpretation of the model helps clarifying this effect: a small q corresponds to a market for highly differentiated products, such that a consumer who likes one product is unlikely to find a substitute.

The market equilibrium when q is close to zero reminded me of the phenomenon of *guruism* (in the derogatory sense). Many consumers opt out entirely because they fail to hear a good story about any market alternative. But each firm attracts a small coterie of fans, who see value in the firm and no value in any other alternative. The firm can exploit this fan base and charge high prices from them, which the fans are happy to pay because they see no substitute.

To summarize the lessons from this parable, when consumers' value judgments are based on naïve extrapolation from anecdotes, equilibrium in the market for quacks turns firms into charlatans who command undeserved fees. When the consumers' problem is nearly hopeless, these charlatans become "gurus." Industry profits do not represent value added: they are a pure wealth transfer from consumers to firms. Moreover, this welfare loss for consumers can increase with the number of firms. Quite a market. Are some real-life industries essentially "markets for quacks"?

Voters: Placebo Reforms (2013)

For our final story, let us turn from consumers evaluating market alternatives to voters evaluating policymakers. Although I use the term "voters," there will be no explicit model of the electoral process. Rather, "voters" are in the background, exerting an accountability pressure on policymakers.

The model is dynamic. At every time period $t = 1, 2, 3, \ldots$, a distinct policymaker chooses an action that may affect the evolution of a variable of public interest—say, GDP. Let's use $x(t)$ to denote the value of GDP at time t.

Two of the available actions are salient, and the public recognizes them as *interventions*. Let's call them s and r. Taking such an action a at time t implies that at every subsequent period $t' > t$ until some future policymaker chooses an intervention,

$$x(t') = x(t'-1) + b_a + u_a(t') - u_a(t'-1)$$

We need to explain what b_a and u_a are. The term b_a is a deterministic trend parameter associated with the action a. The term $u_a(t')$ represents the period-t' realization of a random noise variable that takes two possible values, $-k_a$ and $+k_a$, with equal probability (independently across time periods). Prior to the first intervention, x evolves as if one of the actions s or r was taken at period 0. The subtraction of $u_a(t'-1)$ means that the stochastic process exhibits *mean reversion*: shocks are transient and do not affect the future evolution of x.

Suppose the trend parameters are $b_s = \varepsilon$ and $b_r = 1$, where $\varepsilon > 0$ is arbitrarily small. The noise parameters satisfy $0 < k_s < 1$ and $k_r > 3$. This means that the intervention s is a relatively *safe* action that induces nearly zero growth and carries small fluctuations around this trend, whereas the intervention r is a *risky* action that induces a larger growth rate as well as wider noise fluctuations.

But there is also a *third*, non-salient action, referred to as a *default* and denoted d. This action cannot affect the course of x: when a policymaker chooses d, x continues to evolve according to the most recent intervention. There is deliberate redundancy in this description: whether a policymaker chooses d or a salient action that happens to coincide with the most recent intervention, the evolution of x is the same. The only thing that can change the evolution of x is an intervention that *differs* from the latest one. An intervention that replicates the most recent one is pure theatre, a *placebo reform*.

What motivates individual policymakers? Assume all a policymaker cares about is the *credit* that voters give her for affecting a change in GDP. This assumption is based on the everyday observation that real-life policymakers seem obsessed with claiming credit for good developments and dodging blame for bad ones. Of course, this motivation is largely the result of basic short-term goals such as winning elections. However, policymakers care about posterity even if they never plan to run for office again: It's why they always come out right in those boring autobiographies they write in their retirement. At any rate, we'll take this motivation as given, without trying to derive it from more basic assumptions.

And here comes the key piece in this model. Our voters attribute credit or blame according to a *simple rule*. If a policymaker chooses the default action d, she gets zero credit. If, however, she chooses an intervention, her credit is the change in x from the moment she acts until the next time some policymaker chooses an intervention. (To close the model, we need to define what happens if no future policymaker ever intervenes. This need not bother us here.)

As with the consumers in the dieter's dilemma and the market for quacks, the voters' rule in this model captures an intuitive model of causal misattribution. I see it all around me. Daniel Kahneman's famous flight instructor story is one example.[5] Here's another one. Esther Duflo and Abhijit Banerjee (2011, chap. 3) argue that false inferences of this kind are partly responsible for major distortions in the demand for medical interventions, such as the overuse of antibiotics:

> Because most diseases that prompt visits to the doctor are self-limiting (i.e., they will disappear no matter what), there is a good chance that patients will feel better after a single shot of antibiotics. This naturally encourages spurious causal associations: Even if the antibiotics did nothing to cure the ailment, it is normal to attribute any improvement to them. By contrast, it is not natural to attribute causal force to inaction: If a person with the flu goes to the doctor,

and the doctor does nothing, and the patient then feels better, the patient will correctly infer that it was not the doctor who was responsible for the cure.

The voters in our story act like the patients in Duflo and Banerjee's example. The question is how their attribution rule affects our policymakers' behavior. The basic observation is that a policymaker has a stronger incentive to pick a salient action following a bad shock. The stochastic process exhibits mean reversion, as in the antibiotics example. Therefore, a negative shock artificially improves the credit that a policymaker gets if she intervenes when this shock hits. In contrast, a positive shock increases the policymaker's incentive to pick the default action, because a salient action is more likely to generate negative credit.

The strategic interaction between policymakers exacerbates this incentive. When a policymaker contemplates whether to intervene, she takes into account the *selective* subsequent interventions. Future policymakers' tendency to intervene after bad shocks pushes the expected credit that the present policymaker receives from an intervention *downward*. This strategic consideration gives a boost to the tendency to reserve interventions to bad shocks.

How does this play out? The policymakers are rational players in an infinite-horizon dynamic game, in which each player moves once. The conventional solution concept for such games is *subgame perfect equilibrium*: each policymaker's action is optimal (given the history at which she acts) against the strategies of all subsequent policymakers.

And this is what subgame perfect equilibrium looks like. Policymakers always intervene after a negative shock (whether it is $-k_r$ or $-k_s$) and opt for the default action after a positive shock (whether it is k_r or k_s). Moreover, when they intervene, they choose the *safe* action s.

What this result means is that all the interventions voters witness along the equilibrium path, except possibly the first one, are pure theater. They are placebo reforms that do not change the evolution of x; their only role is to take advantage of voters' intuitive attribution rule in pursuit of credit. On average, interventions take place every two periods. But neither of these interventions, except maybe the first one, has any real effect. At the same time, since these placebo reforms maintain the safe action, the long-run growth rate is nearly zero. This entails a significant loss of societal welfare relative to the alternative of playing r, which is risky in the short run but brings much higher returns in the long run.

Let us see why this is an equilibrium. Consider a policymaker who moves at some time period t following a shock $u(t)$. She takes it as given that whatever she does, the next policymakers will act only after negative shocks. Since positive and negative shocks are equally likely, the next negative shock will arrive after two periods in expectation. Therefore, if the policymaker chooses the safe intervention s, the expected credit she'll get is $2\varepsilon - k_s - u(t)$. Likewise, if she takes the risky action r, her expected credit will be $2 - k_r - u(t)$. Since $k_r - k_s > 2$, the policymaker will prefer s to r. And since ε is nearly zero, $2\varepsilon - k_s - u(t) > 0$ only when the shock $u(t)$ is negative. We have established that if the policymaker expects subsequent policymakers to intervene only after negative shocks, her optimal response is to do the same and adopt the safe action when that happens. This means that our guessed strategy is consistent with subgame perfect equilibrium. Showing no other equilibria exist is more intricate, and so I'll skip this part.

The assumption that policymakers are replaced every period is critical for this argument. Consider an alternative model in which policymakers move every T periods, where T is very large. Then, an individual policymaker faces much weaker competition for credit by subsequent policymakers. The difference between the two interventions' growth rates dwarfs the effect of adverse selection in future policymakers' timing of interventions. It is the long-run trend, not the short-term fluctuations, that dictates policymakers' choices, and they will always intervene and choose the risky action. Voters will still witness placebo reforms on the equilibrium path, but the selective intervention and risk aversion will disappear. This comparison resonates with the common intuition that short government terms inhibit major reforms. In the model, the reason is that short-term fights for credit lead to adversely selective interventions, which in turn lead policymakers to prefer low-risk, low-return measures.

In this last installment of the placebo trilogy, it is no longer consumers in a particular market but the general public that pays the price of using intuitive but ultimately wrong methods of causal attribution.

What Does It Mean?

Perhaps the most basic tenet of modern economics is its theory of value. According to this theory, the value of products and services inheres in peoples' willingness to pay for it—that is, in their preferences. These preferences are primitive and should be respected by the economic analyst: "*de gustibus non est disputandum*," as Stigler and Becker (1977) argued.

The placebo theme challenges this basic attitude. The economic value of objects or actions, as measured by how much people are willing to pay for them, can be a consequence of systematic attribution errors. If economic agents confuse correlation with causation, if they draw exaggerated inferences from anecdotal evidence, or if they fail to understand mean reversion, they may end up assigning value to intrinsically worthless things. And the broader implications? Transactions that are normally recorded as part of GDP are mere transfers. Sectors that seem to have great added value are nothing but a machine for transferring wealth from one group of people to another. Careers and reputations of business executives and policymakers live or die by stakeholders' poor methods for assigning credit and blame. The broad economic significance of the three placebo parables is that we need to take seriously the role of attribution errors in the formation of subjective value.

Analyze This

More personally, what is it about the placebo theme that I find so irresistible as a researcher? I can only speculate.

Part of it has to do with aesthetic or pedagogic appeal: when studying the economic consequences of inference errors, it is convenient to consider a case with a very clear correct-beliefs benchmark. Actions with zero objective value offer such a neat benchmark. In this way, anything "interesting" will be due to the novel model of inference errors. There is a downside to this methodology. An "uninteresting" benchmark is extreme and usually unrealistic, which makes the exercise less plausible as a realistic description of the economic system in question. Once again, we have a clash between "pure" and "applied" attitudes, which has been a running theme in this book. The "pure theory" mentality will welcome the crisp benchmark, while the "applied theory" mentality will aim at a more realistic one, even at the cost of muddying the task of figuring out which effects are due to the novel behavioral element.

Part of it could be politics, an aspect of a "leftist" mentality that doubts the value system underlying our economic and political systems. I do seem to be obsessed with the broader idea that some economic agents receive unmerited rewards. My very first research paper, written when I was still a master's student at Tel Aviv University, proposed that intermediaries who contribute nothing to overall welfare can use exclusive-dealership contracts to extract the entire surplus that other agents generate.[6] This theoretical effect had nothing to do with inference errors, yet it foreshadowed the "why do the wicked prosper" mindset of the placebo trilogy.

Or it might have to do with *professional* politics and with a childish rebellious streak that my wife insists (against my mild protest) that I possess. During my formative years, the economics profession had a "center-right" flavor, and so mocking it from the left was the appropriately juvenile prank to pull. With the profession's leftward shift in recent years, I can imagine the same childish impulse producing pieces that would attack it from the right—for example, making fun of some of my "applied micro" peers' apparent belief that they can reconcile dispassionate scientific objectivity in their papers with vigorous activism on their Twitter account.

Looking inward more deeply, beyond the aesthetic and political layers, is it possible that the placebo theme springs from professional insecurity, a fear of being called out as an overpaid know-nothing? Under this interpretation, I am the charlatan who commands undeserved fees. Or is it envy of former schoolmates who pursued careers as executives and financiers? Do I get a kick from showing that they earn their money peddling worthless propositions, if only in the fantasy world of my models? In other words, is my obsession with the placebo theme an expression of something I share with every male primate, especially in a time of soaring inequalities: status anxiety?

If this kind of inward probing seems overly self-indulgent, note that it can be directed outwardly just as well. There is a flipside to the final question of the previous paragraph. Those countless studies that go out of their way to show how the alpha males of our economy rightfully earn their money and status; could they, too, spring from status anxieties of a similar nature?

Whatever the answers to these questions may be, the placebo theme illustrates how economic theory can serve as an "artistic" medium that channels raw psychological and political motives. Does this diminish the value of economic theory? Or does it make it more interesting?

This is a section of doi:10.7551/mitpress/14884.001.0001

The Curious Culture of Economic Theory

By: Ran Spiegler

Citation:

The Curious Culture of Economic Theory
By: Ran Spiegler
DOI: 10.7551/mitpress/14884.001.0001
ISBN (electronic): 9780262379038
Publisher: The MIT Press
Published: 2024

The open access edition of this book was made possible by generous funding and support from MIT Press Direct to Open

The MIT Press

The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Palatino by Westchester Publishing Services.

The author's preference is to use binary pronouns in this work. The MIT Press recognizes all gender identities and recommends the use of gender-neutral pronouns.

Library of Congress Cataloging-in-Publication Data

Names: Spiegler, Ran, author.
Title: The curious culture of economic theory / Ran Spiegler.
Description: Cambridge, Massachusetts : The MIT Press, [2024] | Includes bibliographical references and index.
Identifiers: LCCN 2023028463 (print) | LCCN 2023028464 (ebook) | ISBN 9780262548229 (paperback) | ISBN 9780262379021 (epub) | ISBN 9780262379038 (pdf)
Subjects: LCSH: Economics. | Economics—Sociological aspects.
Classification: LCC HB171 .S7155 2024 (print) | LCC HB171 (ebook) | DDC 330.01—dc23/eng/20230713
LC record available at https://lccn.loc.gov/2023028463
LC ebook record available at https://lccn.loc.gov/2023028464