Eastern Economic Journal, **2008**, **34**, (115–125)

www.palgrave-journals.com/eej

Symposium article

Why the Con Hasn't Been Taken Out of Econometrics

Martin Zelder[a,b]
[a]Department of Economics, Northwestern University, Evanston, IL 60208, USA
E-mail: m-zelder@northwestern.edu
[b]Center for Health and the Social Sciences, University of Chicago, Chicago, IL 60637, USA

Economists often decry the perceived tendency towards selective reporting of empirical results (''specification search'') in scholarly work. Yet, economists have largely neglected to analyze the incentive structures underlying this phenomenon of econometric ''cons''. This paper endeavors to provide this analysis, posing a game-theoretic model of specification search. In this three-player game (author, journal, and profession), academic authors choose whether to report the ''true'' *t*-statistic associated with an empirical result, or whether to ''con'' by reporting a distorted *t*-statistic. Subsequently, both journal and profession must choose whether to bear the cost of ''scrutinizing'' the author's work (e.g., by reanalyzing his data). Multiple perfect Bayesian equilibria are found, including one where authors ''con'' and arc not detected, an equilibrium which may be Kaldor–Hicks efficient. Moreover, public and private mechanisms (existing or proposed) to curtail ''conning'' seem ineffectual.
Eastern Economic Journal (2008) **34,** 115–125. doi:10.1057/palgrave.eej.9050006

Keywords: econometrics; specification search; academic refereeing

JEL: C52; C72; A14

INTRODUCTION

Most people, including those who are not sure how to calculate a mean, mistrust statistical analysis. Even those among us presumably able to calculate means correctly have their doubts: ''Econometricians have found their Philosophers' Stone; it is called regression analysis and is used for transforming data into 'significant' results! Deception is easily practised from false recipes intended to simulate useful findings'' [Hendry 1980]. Although these concerns have evolved into systematic prescriptive analysis (e.g., Leamer [1978]), an irony remains: despite undoubtedly wide agreement among economists that many econometric estimates are suspect, few among us openly acknowledge our own shortcomings. Moreover, as Levy and Peart [2003] point out, economics/econometrics (unlike, e.g., dental hygiene) lacks even a basic formalized code of ethics to implement appropriate principles to guide estimation.

Perhaps one reason for our lack of clear principles that should and do govern econometric estimation is the limited body of economic analysis on the subject. Indeed, the vast majority of writing on econometric practice is didactic/prescriptive, and as such is extraordinarily useful.[1] Economic analysis of econometric practice, is, however, much rarer, with exceptions found in papers by De Long and Lang [1992], Feigenbaum and Levy [1993a, b; 1996], Fölster [1995], Engers and Gans [1998], and Freedman [2000], although none of these presents an explicit model of choice involving econometricians, journals, and the economics profession.

It is, therefore, the task of this paper to attempt to explore such a model, formalizing the insights from the analytical literature noted above while infusing

them with the spirit of the economic analysis of fraud and crime [including Becker 1968; Akerlof 1970; Darby and Karni 1973; Matsumura and Tucker 1992]. The orientation of this model is game-theoretic, with relevant choices being made in sequence by three players: an econometrician, a journal, and the rest of the economics profession. The next section defines "cons" in econometrics. The subsequent sections present the structure of the "con" game, examine the equilibria of the game, and assess the normative properties of these equilibria, respectively. The final section offers concluding thoughts.

DEFINING "CONS"

For the purposes of the analysis in this paper, outcomes in the market for econometric estimates will be depicted as arising from a sequential three-player game of incomplete information. The three players are the author (A), the journal (J), and the rest of the economics profession (P). In this game, moves are taken in sequence by A, then J, and finally, P.

Underlying A's choice problem is a distribution of t-statistics corresponding to a parameter of primary interest, β_1.[2] This distribution is defined by the full set of estimates associated with the specifications generated by A's private estimation process. Among these t-statistics is what can be referred to as the "true" t, t_T, which is the t-statistic associated with the specification that is best justified on a priori grounds. Operationally, this best a priori specification can be discerned by the other players, J and P, but only if additional costs are borne. Thus, t-statistics are credence goods as defined by Darby and Karni [1973], which "cannot be evaluated in normal use," with "normal use" in this context entailing reading and evaluating the article without actually reanalyzing the data.

Although one could conceive of A's choice among potential t-statistics as continuous (or at least multinomial), for simplicity, that choice is presented here as dichotomous. That is, the article that A submits to J is characterized by either the true t or a t different from the true t. The true t is private information known only by A unless J or P incur costs of scrutinizing the data; without incurring these costs, J and P view $t^*=a^*t_T$. A's choice of t^* reflects the extent to which he chooses to distort the true t, t_T. In other words, after A's analysis generates t_T, he chooses an optimal value, a^*, by which to inflate or deflate t_T in producing the publicly reported t value, t^*. The feasible values of a^* are constrained, however, to lie in the set implied by the *actual* set of t values estimable from A's data; that is, I assume that a^* cannot be chosen fancifully by A. The exclusion of fanciful explicitly frandulent choices of a^* is assumed in order to focus on the notion of "cons" as the product of specification searching as originally formulated by Leamer [1978; 1983].[3]

In this context, it is not obvious what constitutes a "con." As suggested above, cons might be considered to occur when t^* either overstates *or* understates t_T. The indeterminate direction of bias associated with a con can be thought of as arising from ideological, methodological, or commercial attachment to a prior belief about what t^* should be. Thus, for example, a researcher intent on propagating the view that lower marginal income tax rates are unrelated to real GDP growth (despite his own evidence that they are growth-enhancing) might choose $0<a^*\ll 1$ so as to report a t-value far below conventional significance levels. Another researcher, who wished to advance the view that tax-rate reduction actually reduced GDP growth (but was also confronted by his own evidence that it is growth-enhancing), might choose $a^*<-1$.

In the strictest sense, one might argue that selecting $a^* \neq 1$ signifies the occurrence of a "con." Alternatively, one might argue that a con is only committed if the selected distortion exceeds some designated threshold. The definition of the threshold depends on whether the con involves overstatement, understatement, or sign reversal. Thus, we might define cons as the following:

(i) $a^* > k$ (overstatement)
(ii) $0 < a^* < k$ (understatement)
(iii) $a^* < k < 0$ (sign reversal)

Choosing between the strictest definition ($a^* \neq 1$) and this set of threshold definitions is not the task of this paper, although it is a vital element in implementing any system to penalize cons, and hence merits some discussion in future work. In deferring this issue, I will assume that the threshold definitions are those relevant for analysis of *actual* current practice within the economics profession. Indeed, while these thresholds are not clearly defined, it is arguably realistic to depict the profession as concerned with whether thresholds of some sort are violated.

THE "CON" GAME: STRUCTURE AND SETUP

Turning to the specific setup of the "con" game (Figure 1), the author moves first, and chooses between the strategies "no con" (*NC*) and "con" (*C*), where "no con" entails adherence to the threshold and "con" involves violating the threshold. At the second stage of the game, journal *J* must decide whether to scrutinize (*S*) or not scrutinize (*NS*) *A*'s data and estimates in order to review the paper. The decision to scrutinize entails evaluation beyond typical assessment by referees, and hence

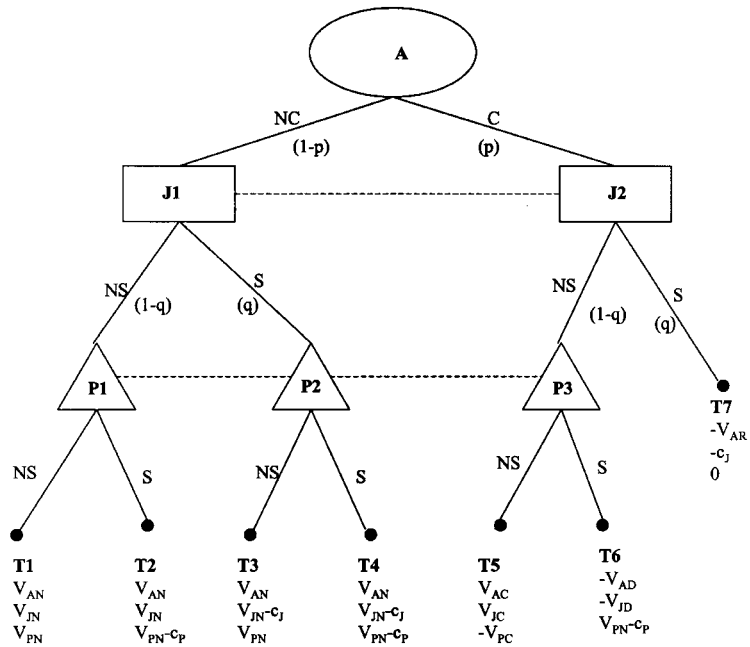

Figure 1. The "Con" Game

involves additional costs to the journal of c_J. These costs are constituted by such activities as recompiling A's data from its original source and reestimating A's regressions, or more generally, any extraordinary effort devoted to reviewing the paper that permits the journal to assess whether a con has occurred. Because of A's private information regarding whether he has conned or not, J's information set consists of multiple (specifically 2) nodes, indicating that A chooses to con (C) with probability p and not to con (NC) with probability $1-p$. J's disposition of the paper then depends on its decision to scrutinize or not, conditional on whether A has conned or not.

Specifically, if A chooses NC, the article will be published, regardless of whether the journal (J) scrutinizes the article prior to acceptance.[4] If, however, A chooses C, the article will not be published if J subjects it to scrutiny, but it will be published if J chooses NS. If J scrutinizes when A has chosen C, and J therefore chooses to reject the article, the game concludes after the second stage, and P (the rest of the profession) does not get to move.

If the article is accepted, however, the game proceeds to its third and final stage. Because P has imperfect information regarding the previous choices made by A (con or not) and J (scrutinize or not), its information set is not a singleton, consisting instead of the nodes $P1$, $P2$, and $P3$. P only observes that the article has been published, and knows A's probability of conning (p) and J's probability of scrutinizing (q) and must determine whether to scrutinize (S) or not (NS). If P scrutinizes the published article, it confirms the article's reported t if no con has occurred, but reveals the true t and the existence of the author's con when a con has occurred. The profession's move, and the corresponding confirmation or discrediting of the article's reported t, concludes play in the game.

Given the structure of play, the payoffs associated with each of the seven terminal nodes can be examined. If the game terminates at nodes T1, T2, T3, or T4, meaning that the author has not committed a con, then the value of the article for each player (gross of scrutiny costs) will be unaffected by whether either or both of J and P engage in scrutiny: any scrutiny will simply provide confirmation of A's claim. In this sense, the value of A's article (when no con is involved) is a collective good for J and P. Thus, at nodes T1–T4, the article's gross values for each of A, J, and N are, respectively, V_{AN}, V_{JN}, and V_{PN}. The payoffs at each of these nodes are then determined by subtracting scrutinizing costs for J and P where relevant. Hence, payoffs at these four terminal nodes are:

T1 (J does not scrutinize, P does not scrutinize):	(V_{AN}, V_{JN}, V_{PN})
T2 (J does not scrutinize, P scrutinizes):	$(V_{AN}, V_{JN}, V_{PN}-c_P)$
T3 (J scrutinizes, P does not scrutinize):	$(V_{AN}, V_{JN}-c_J, V_{PN})$
T4 (J scrutinizes, P scrutinizes):	$(V_{AN}, V_{JN}-c_J, V_{PN}-c_P)$

The payoffs at the other three terminal nodes, T5–T7, correspond to cases where A has chosen to commit a con. At node T5, where neither J nor P scrutinizes, the con remains unexposed, harming the profession by the amount $-V_{PC}$ (as a result, e.g., of a misallocation of future research effort that is induced by the con), but generating positive payoffs for author and journal of V_{AC} and V_{JC}, respectively. If, however, the profession chooses to scrutinize the article post-publication, the signs of the payoffs are reversed between T5 and T6, with reputational penalties suffered by author ($-V_{AD}$) and journal ($-V_{JD}$) and a potentially positive payoff to the

profession of $V_{PN}-c_P$. Finally, terminal node T7 results when an author whose article contains a con has that article scrutinized by the journal and consequently has the article rejected. In this circumstance, the author suffers a diminished reputation with the journal in question corresponding to the payoff $-V_{AR}$, while the journal incurs the cost of scrutinizing (c_J), and the rest of the profession receives neither benefit nor harm. Hence, the payoffs associated with nodes T5–T7 are:

T5 (J does not scrutinize, P does not scrutinize): $(V_{AC}, V_{JC}, -V_{PC})$
T6 (J does not scrutinize, P scrutinizes): $(-V_{AD}, -V_{JD}, V_{PN}-c_P)$
T7 (J scrutinizes, P does not move): $(-V_{AR}, -c_J, 0)$

Some assumptions are made regarding the payoffs in order to enhance realism and provide analytical content to the game. For the author, $V_{AC} > V_{AN}$; that is, A receives a higher payoff from getting away with a con than if he chooses not to con. Similarly, for the journal it is assumed that $V_{JC} > V_{JN}$; this reflects the private gain to a journal from publishing an unrefuted article that contains (say) an inflated t value relative to publishing the true t. These assumptions provide the private motivations for A to engage in and for J to facilitate cons.

THE "CON" GAME: EQUILIBRIA

Given these payoffs, the solution to the game can be derived by backward induction. In particular, equilibria must be perfect Bayesian equilibria, requiring that players' beliefs are sequentially rational and are determined by Bayes' rule and the players' equilibrium strategies [Gibbons 1992]. To determine the solution, consider first P's beliefs regarding which node has been reached conditional on the article having been published. These beliefs rely on underlying probabilities that each node has been reached that sum to 1. P computes these probabilities by applying Bayes' rule to the probabilities that condition its choice: p (the probability that A attempts a con) and q (the probability that J scrutinizes A's submitted paper). According to Bayes' rule, the probabilities that P is at nodes P_i (i=1, 2, 3) are given by

$$\begin{aligned} Pr(P_1|J_1 \cup J_2) &= \frac{Pr(P_1|J_1)Pr(J_1)}{Pr(P_1|J_1)Pr(J_1) + Pr(P_2|J_1)Pr(J_1) + Pr(P_3|J_2)Pr(J_2)} \\ &= \frac{(1-q)(1-p)}{(1-q)(1-p) + q(1-p) + (1-q)p} \\ &= \frac{(1-q)(1-p)}{1-qp} = \pi_1 \end{aligned}$$

$$\begin{aligned} Pr(P_2|J_1 \cup J_2) &= \frac{Pr(P_2|J_1)Pr(J_1)}{Pr(P_1|J_1)Pr(J_1) + Pr(P_2|J_1)Pr(J_1) + Pr(P_3|J_2)Pr(J_2)} \\ &= \frac{q(1-p)}{(1-q)(1-p) + q(1-p) + (1-q)p} \\ &= \frac{q(1-p)}{1-qp} = \pi_2 \end{aligned}$$

$$\begin{aligned} Pr(P_3|J_1 \cup J_2) &= \frac{Pr(P_3|J_2)Pr(J_2)}{Pr(P_1|J_1)Pr(J_1) + Pr(P_2|J_1)Pr(J_1) + Pr(P_3|J_2)Pr(J_2)} \\ &= \frac{(1-q)p}{(1-q)(1-p) + q(1-p) + (1-q)p} \\ &= \frac{(1-q)p}{1-qp} = \pi_3 \end{aligned}$$

Given these conditional probabilities, the expected payoffs for P associated with his strategy set (NS, S) can be formulated and compared. Specifically, P chooses not to scrutinize the article if:

(1) $$E(\text{Payoff } S)_P < E(\text{Payoff } NS)_P$$

$$\begin{aligned} \pi_1(V_{PN} - c_P) + \pi_2(V_{PN} - c_P) + \pi_3(V_{PN} - c_P) &< \pi_1 V_{PN} + \pi_2 V_{PN} + \pi_3(-V_{PC}) \\ V_{PN} - c_P &< (\pi_1 + \pi_2)V_{PN} - \pi_3 V_{PC} \\ \pi_3(V_{PN} + V_{PC}) &< c_P \\ \pi_3 &< \frac{c_P}{V_{PN} + V_{PC}} \end{aligned}$$

Intuitively, P chooses not to scrutinize if the expected benefit from doing so, $\pi_3(V_{PN} + V_{PC})$, is exceeded by the certain cost of scrutinizing, c_P.

Because it is ambiguous as to whether inequality (1) holds for P — neither strategy strictly dominates the other — J's analogous choice over (NS, S) must be evaluated for each of J's relevant possible beliefs regarding π_3. These are:

Belief 1A	$\pi_3 < \frac{c_P}{V_{PN}+V_{PC}}$	NS strictly dominates S for P
Belief 1B	$\pi_3 > \frac{c_P}{V_{PN}+V_{PC}}$	S strictly dominates NS for P

Because J's information is also imperfect — it does not know whether the article it receives is based on a con or not — it must then choose whether to scrutinize or not based on the expected payoffs associated with those two options, payoffs that depend on P's choice, and thus on J's beliefs regarding P's beliefs and P's choice.

Consider each belief in turn. Suppose that J holds Belief 1A, that is, that no scrutiny is the dominant strategy for P. Then, J will choose to scrutinize if:

(2) $$E(\text{Payoff } S)_J > E(\text{Payoff } NS)_J$$

$$\begin{aligned} (1-p)(V_{JN} - c_J) + p(-c_J) &> (1-p)V_{JN} + pV_{JC} \\ -c_J &> pV_{JC} \end{aligned}$$

Condition (2) cannot hold for any p (given non-negative values for c_J and V_{JC}). Hence, NS is a dominant strategy for J, given that NS is a dominant strategy for P. Intuitively, if J knows that the profession will never scrutinize, then J can never suffer reputational penalties from its own failure to scrutinize.

Given that NS dominates for J when it dominates for P, an equilibrium strategy for A can be discerned. Specifically, if A knows that failure to scrutinize

by P will induce failure to scrutinize by J, then the only terminal nodes relevant for A are T1 and T5. At T1, A does not con, and receives payoff V_{AN}, while at T5, A does con (and gets away with it), receiving payoff V_{AC}. By assumption, $V_{AN} < V_{AC}$. Hence, C is an equilibrium strategy for A, given dominant strategies of NS for both J and P. This reveals one perfect Bayesian equilibrium: (C, NS, NS).

Alternatively, if J holds Belief 1B, then J compares the expected payoffs from scrutinizing and not scrutinizing associated with terminal nodes T2, T4, T6, and T7. This means that J chooses not to scrutinize if:

$$E(\text{Payoff } S)_J < E(\text{Payoff } NS)_J \tag{3}$$

$$(1-p)(V_{JN} - c_J) + p(-c_J) < (1-p)V_{JN} + p(-V_{JD})$$

$$-c_J < -pV_{JD}$$

$$c_J > pV_{JD}$$

$$p < \frac{c_J}{V_{JD}}$$

According to (3), when J believes that P will scrutinize no matter what, then J will only choose not to scrutinize when its expected private reputational benefit from scrutinizing, pV_{JD}, is exceeded by its private cost of scrutinizing. Thus, neither S nor NS is a dominant strategy for J, given that S is a dominant strategy for P.

As a result, we must consider two cases with regard to J's beliefs regarding p:

Belief 2A	$p < \frac{c_J}{V_{JD}}$	NS strictly dominates S for J
Belief 2B	$p > \frac{c_J}{V_{JD}}$	S strictly dominates NS for J

Suppose J holds Belief 2A. Then, NS is its dominant strategy, given that P chooses S as its dominant strategy. Given these choices, A then compares the terminal nodes T2 (where he does not con) and T6 (where he does con). Given that it is dominant for P to scrutinize, A receives a higher payoff, V_{AN}, by refraining from conning, than from conning, $-V_{AD}$. Hence, another perfect Bayesian equilibrium is: (NC, NS, S).

Finally, suppose that J holds Belief 2B. In that case, S is J's dominant strategy, given that P also chooses S as its dominant strategy. The relevant payoffs for A, as a consequence, are those associated with not conning at node T4 (V_{AN}) and conning at node T7 ($-V_{AR}$). Therefore, A again chooses not to con, generating a third perfect Bayesian equilibrium: (NC, S, S).

WELFARE IMPLICATIONS

Given these three pure-strategy perfect Bayesian equilibria (enumerated below), it is important to consider their welfare properties:

Node	*Strategies*	*Payoffs*
T1	(NC, NS, NS)	(V_{AN}, V_{JN}, V_{PN})
T4	(NC, S, S)	$(V_{AN}, V_{JN} - c_J, V_{PN} - c_P)$
T5	(C, NS, NS)	$(V_{AC}, V_{JC}, -V_{PC})$

Using Pareto criteria, the equilibria at T1 and T4 can be unambiguously compared with one another, while T5 is not Pareto-comparable with T1 or T4. Both T1 and T4 are equilibria at which A chooses not to con, but T4 involves both J and P bearing costs of scrutinizing. Comparing T1 with T4, A is just as well off at T1, and J and P are both better off. Hence, T1 is preferred to T4 on Pareto grounds.

Given T1's Pareto-superiority to T4, we will set aside T4 in making the further comparison with T5. Unlike T1 (and T4), T5 involves A choosing to engage in a con. As noted previously, A and J are each better off at T5 than at T1. P, however, is worse off at T5. Because T1 and T5 are Pareto-non-comparable, an alternative welfare standard, such as Kaldor–Hicks efficiency, must be considered if these two equilibria are to be compared. Applying this standard (and assuming transferable utility), the gains to P at T1 (relative to T5) can be compared with the losses to A and J at T1 (relative to T5), providing a basis to determine the condition under which T5 is Kaldor–Hicks inefficient:

$$\text{Gains to } P > \text{Losses to } A + \text{Losses to } J \tag{4}$$

$$V_{PN} + V_{PC} > (V_{AC} - V_{AN}) + (V_{JC} - V_{JN})$$

The implication of this condition is that the equilibrium at T1, where A chooses not to con and J and P choose not to scrutinize, is Kaldor–Hicks efficient if the gain to the profession (represented by the value it attaches to a no-con result being published, V_{PN}, plus the loss the profession avoids, V_{PC}) exceeds the aggregated losses to A and J as determined by the difference for each between the value of a con and the value of a non-con article.

This result, that the "successful" con equilibrium, T5, may be Kaldor–Hicks inefficient, is perhaps reflective of many economists' concerns regarding published econometric research. It is thus worth considering the likelihood that this equilibrium occurs and is in fact inefficient. A necessary condition for T5 to occur is that P chooses not to scrutinize, or, from equation (1):

$$\pi_3 < \frac{c_P}{V_{PN} + V_{PC}} \tag{1}$$

This is more likely to be true the larger P's cost of scrutinizing, c_P, are. Engers and Gans [1998] suggest an explanation for why *referees*' costs of reviewing papers are considerable, and that is potentially applicable to the profession's cost of scrutinizing. They argue that the cost to referees associated with review includes not only their time costs but also an uninternilizable external cost. This cost arises because the act of refereeing itself, by raising the quality of the journal (a public good), raises the cost to other referees of choosing to referee, as they can free-ride on the efforts of others. Further, because payments to referees exacerbate this problem, they conclude that subsidies to refereeing may be efficient.

Also, (1) is more likely to hold the smaller the denominator, $V_{PN} + V_{PC}$. The denominator reflects P's private gains from scrutinizing, gains that are necessarily related to the profession's return to scrutinizing. In this context, Freedman's [2000] observations regarding the incentives characterizing journal referees and editors are probably pertinent. As Freedman notes, while journal marketing, printing, and distribution are increasingly governed by strict profit-maximization considerations, editorial operations are still best modeled as a non-profit enterprise. Consequently, one "would expect either Type I errors (publishing a paper that should not be

published) or the more serious Type II errors (not publishing a paper that should be published) to abound'' [Freedman 2000, p. 376], with successful cons exemplifying Type I errors. Analogous concerns undoubtedly characterize the incentives, and thus the payoffs, for P. As presented in stylized form here, P can be thought of as representing other academic economists besides A and other journals besides J. Given the deviation from strict for-profit incentives that typifies many economists' workplaces, and the vulnerability of other journals to Type II errors regarding refutations of cons, it is plausible to think that $V_{PN} + V_{PC}$ is relatively small, and that (1) is relatively likely to be satisfied. At the same time, as long as $V_{PN} + V_{PC}$ is not too small relative to A's and J's private gains from conning, efficiency condition (4) will still hold.

CONCLUSION: THE ECONOMETRICIAN'S NEW CLOTHES?

All of this suggests an irony. If it is likely that the equilibrium at which the econometrician cons, and the journal and the profession look aside, characterizes much of what is published,[5] then the scrutiny that we purport to focus on the real world might perhaps be better focused on ourselves.

How might we refocus? State intervention, as proposed by Engers and Gans [1998], is problematic in that any effective intervention would necessarily target the penalties and rewards associated with the review process, adding an additional layer of judgment by ''experts'' with incentives no better than those guiding current referees, editors, and economists in the general public [Darby and Karni 1973]. The state might impose large penalties administered rarely [Becker 1968], although given the inevitable difficulty of drawing the line between cons and non-cons, this might deter more of the latter than the former.

It might be thought that these views are too pessimistic. If, in fact, the reputational penalty imposed on journals for publishing cons, V_{JD}, is large, then they will be more likely to scrutinize (if they believe that P will also scrutinize). In the context of corporate crime, Karpoff and Lott [1993] find that reputational penalties are substantial. Within the academic world, however, Feigenbaum and Levy [1993a, b] suggest that such penalties are minimal, a phenomenon LaFollette [1992] attributes to fear of legal repercussions and reputational damage for *con exposers*, noting the chilling observation of a congressional fraud investigator: ''[F]or 'every one of the cases we looked at, there was hardly a single person you could point to that ever raised a question and survived'" [LaFollette 1992, p. 148, quoting from Kuznik, 1991].

Remaining hope might come from shifts in social norms surrounding academic publishing, an area recently explored by Ellison [2002]. Without attaching normative interpretation, Ellison depicts the growing length of economics papers and of the attendant review process as the product of a shift in emphasis away from one quality dimension (the main idea of the paper) and towards the other (improvements on, elaborations of, and robustness checks on the main idea, *as required by referees and editors*). While this latter form of quality resembles what has been referred to as ''scrutinizing'' in this paper, the results of this paper suggest that norms may not have evolved far enough in the direction of criteria proposed by Leamer and Leonard [1983] or McAleer et al. [1985].

The painful joke known seemingly to all is that economists assume a can opener in order to open a can of beans. This paper, and its important predecessors

(e.g., Leamer [1978]), suggests a new, and more sobering, version: do we, as econometricians, assume a validated methodology when opening cans of data? Figuring out the optimal penalty for cons might make us less prone to be the butt of jokes, and increase the confidence placed in econometrics — and econometricians.

Notes

1. Included here are the many texts available, as well as prescriptive writing more directly focused on specification search, such as work by Leamer [1978; 1983], Leamer and Leonard [1983], Hendry [1980], Lovell [1983], Feige [1975], McAleer et al. [1985], and Cooley and LeRoy [1986].
2. While it is undoubtedly more realistic to imagine that A has preferences over the t-statistics associated with a vector of β's, the analysis is simplified by focusing on a single parameter of interest without any apparent loss of qualitative insight.
3. Of course, it is possible that some t values reported in the literature are based on such fanciful distortions by researchers, but this explicitly fraudulent "con" is not the subject of consideration here.
4. Rather than assuming that all papers submitted by A's who choose NC and are scrutinized by journals are accepted, one could specify a game in which A's who choose NC and are scrutinized by journals have their papers accepted with probability $r<1$. This modification, while providing additional realism and complexity, does not appear to imply qualitatively different results for the analysis.
5. An anonymous referee raises the question of how common cons are. Given their nature — the non-reporting of results from a subset of regressions — empirical evidence of this is difficult to find. Indeed, LaFollette [1992], in her book on scientific misconduct, describes it as "an inherently unanswerable question." Nevertheless, commentary by eminent economists has for some time warned of the prevalence of cons. In beseeching the *Journal of Political Economy* to adopt novel refereeing protocols, Feige [1975] maintained that current policy was "contaminating our published literature with a proliferation of Type I errors", and that "data 'experimentation' in search of statistical significance is a widespread practice." Much more recently, in commenting on the adoption of the new policy by *Labour Economics* of encouraging replications, Hedges [1997] asserts that "Most empirical research involves some degree of specification searching, which may reduce interpretability of results." Moreover, when Dewald et al. [1986] simulated aspects of a specification search in 500 trials that attempted to replicate a specific result that had been published in the *Journal of Money, Credit and Banking*, they found that the result was quite fragile to the specification chosen.

References

Akerlof, George A. 1970. The Market for 'Lemons': Quality Uncertainty and the Market Mechanism. *Quarterly Journal of Economics*, 84: 488–500.

Becker, Gary S. 1968. Crime and Punishment: An Economic Approach. *Journal of Political Economy*, 76: 169–217.

Cooley, Thomas F., and Stephen F. LeRoy. 1986. What Will Take the Con Out of Econometrics? A Reply to McAleer, Pagan, and Volker. *American Economic Review*, 76: 504–507.

Darby, Michael R., and Edi Karni. 1973. Free Competition and the Optimal Amount of Fraud. *Journal of Law and Economics*, 16: 67–88.

De Long, J.Bradford, and Kevin Lang. 1992. Are All Economic Hypotheses False? *Journal of Political Economy*, 100: 1257–1272.

Dewald, Willliam G., Jerry G. Thursby, and Richard G. Anderson. 1986. Replication in Empirical Economics: The *Journal of Money, Credit and Banking* Project. *American Economic Review*, 76: 587–603.

Ellison, Glenn. 2002. Evolving Standards for Academic Publishing: A $q-r$ Theory. *Journal of Political Economy*, 110: 994–1034.

Engers, Maxim, and Joshua S. Gans. 1998. Why Referees Are Not Paid (Enough). *American Economic Review*, 88: 1341–1349.

Feige, Edgar L. 1975. The Consequences of Journal Editorial Policies and a Suggestion for Revision. *Journal of Political Economy*, 83: 1291–1296.

Feigenbaum, Susan, and David M. Levy. 1993a. The Market for (Ir)Reproducible Results in Econometrics. *Social Epistemology*, 7: 215–232.

———., 1993b. Response to the Commentaries. *Social Epistemology*, 7: 286–292.
———., 1996. The Technological Obsolescence of Scientific Fraud. *Rationality and Society*, 8: 261–276.
Fölster, S. 1995. The Perils of Peer Review in Economics and Other Sciences. *Journal of Evolutionary Economics*, 5: 43–57.
Freedman, Craig. 2000. Do Economic Journals Obey Economic Prescriptions? *Review of Industrial Organization*, 17: 371–384.
Gibbons, Robert. 1992. *Game Theory for Applied Economists*. Princeton: Princeton University Press.
Hedges, Larry V. 1997. The Promise of Replication in Labour Economics. *Labour Economics*, 4: 111–114.
Hendry, David F. 1980. Econometrics — Alchemy or Science? *Economica*, 47: 387–406.
Karpoff, Jonathan, and John R. Lott, Jr.. 1993. The Reputational Penalty Firms Bear from Committing Criminal Fraud. *Journal of Law and Economics*, 36: 756–802.
Kuznik, Frank. 1991. Fraud Busters. *Washington Post Magazine* (April 14): 22–26, 31–33.
LaFollette, Marcel C. 1992. *Stealing into Print: Fraud, Plagiarism, and Misconduct in Scientific Publishing*. Berkeley: University of California Press.
Leamer, Edward E. 1983. Let's Take the Con Out of Econometrics. *American Economic Review*, 73: 31–43.
———., 1978. *Specification Searches: Ad Hoc Inference with Nonexperimental Data*. New York: John Wiley & Sons.
———., and Herman Leonard. 1983. Reporting the Fragility of Regression Estimates. *Review of Economics and Statistics*, 65: 306–317.
Levy, David M., and Sandra J. Peart. 2003. Why Would Truth-Seekers Need a Code of Ethics?, Unpublished manuscript.
Lovell, Michael C. 1983. Data Mining. *Review of Economics and Statistics*, 65: 1–12.
Matsumura, Ella Mae, and Robert R. Tucker. 1992. Fraud Detection: A Theoretical Foundation. *Accounting Review*, 67: 753–782.
McAleer, Michael, Adrian R. Pagan, and Paul A. Volker. 1985. What Will Take the Con Out of Econometrics? *American Economic Review*, 75: 293–307.