

Marshall's Theory Should be Discarded

Steve Keen* and Carmen Costea[⊗]

Though it is no longer an active area of research by economists, the Marshallian theory of the firm is still central to introductory pedagogy in economics. It has withstood numerous criticisms over the years—of its internal consistency, its empirical relevance, its uni-dimensional description of the motives of firms, its "black box" treatment of the firm, and so on. In this article we put one further critique: it is, quite simply, mathematically false. When the errors in the theory are corrected, nothing of substance remains: Equating marginal revenue & marginal cost does not maximize profits, competition does not lead to price equaling marginal cost, and the welfare loss previously attributed to monopoly is due instead to profit maximizing behavior, independent of the number of firms in an industry.

1. A personal introduction by Steve Keen

When I wrote *Debunking Economics* seven years ago, my intention was simply to produce an accessible collation of the extant criticisms of neoclassical economics for a non-technical audience. Apart from my own critique of Marxian economics (Keen 1993a, 1993b), I had no intention of putting anything “new” into the book.

I was therefore somewhat surprised when, as I set out to document the standard criticisms of the neoclassical theory of the firm, I spotted something that I thought *was* new. It appeared that the “horizontal demand curve”, an essential aspect of the model of perfect competition, was logically incompatible with another essential aspect of the model, the downward-sloping market demand curve. When this logical incompatibility was acknowledged, the demand curve for the individual firm in Marshallian competition had the same slope as the market demand curve, and as a result, a “competitive” industry of atomistic profit-maximizers produced the same output as a comparable monopoly.

I outlined this argument verbally in Chapter 4, gave the chapter a suitably provocative title (“Size Does Matter”), went on with the remainder of the book, and delayed exploring the issue in mathematical detail until after the book was finished.

Once *Debunking Economics* was published, I had no choice but to do so. The book was well received by its target audience, but as I expected, neoclassical economists ignored it—except for that one chapter. They were, of course, convinced that I had made serious logical errors, and endeavored to tell me so—in email conversations, on discussion lists, during seminars, in a smattering of reviews, and in negative referee reports from neoclassical journals (*The Economic Journal*, *AER* and *Journal of Economic Education*).

As I addressed each objection, new ones were made, and as a result the critique went from something quite simple to something quite elaborate. This paper follows that developmental chronology, from the original insight—which, I discovered, was not

* University of Western Sydney; s.keen@uws.edu.au

[⊗] Academy of Economic Studies, Bucharest, cecoste@ase.ro.

new—through to original arguments, to a method not covered here (derived with the assistance of Russell Standish)¹ that enables Marshallian and Cournot theories of competition to be integrated. The critiques, in increasing order of analytic complexity, are:

- The demand curve for the individual competitive firm can't be horizontal under the Marshallian assumption of atomism;
- A competitive firm increases profit if it reduces output below where price equals marginal cost;
- Equating marginal cost and marginal revenue isn't profit-maximizing behavior in a multi-firm industry;
- Given comparable cost functions, the profit-maximizing output level for the individual firm results in a market level output that is independent of the number of firms;
- A simulated market of instrumentally rational profit-maximizing firms converges to the "Keen" level of output, rather than the neoclassical prediction;
- Price-taking behavior is irrational, and a degree of irrationality is needed to cause a competitive industry to converge to the output level at which price equals marginal cost;
- Finally, all of the above takes for granted that markets are in fact characterized by price-driven demand, homogenous products, and diminishing marginal productivity in a context of certainty. There is overwhelming empirical evidence that real markets are characterized by diversity-driven demand, heterogeneous products, and constant or rising marginal productivity in a context of uncertainty. Neoclassical theory is therefore not only wrong but also irrelevant. The continued teaching of Marshallian fantasies—and most research into "Industrial Organization"—are hindrances to the task of developing a theory of competition that has any relevance to what we witness in the real world.

2. Stigler 1957

The proposition that Keen had thought original in *Debunking Economics*—that, under conditions of "atomism", the slope of the demand curve facing the individual competitive firm was the same as the slope of the market demand curve—had in fact been made in 1957 by that arch defender of neoclassicism, George Stigler, and in a leading neoclassical journal: *The Journal of Political Economy* (Stigler 1957—see Figure 1).

THE JOURNAL OF POLITICAL ECONOMY

 Volume LXXV

FEBRUARY 1957

 Number 2

PERFECT COMPETITION, HISTORICALLY CONTEMPLATED

GEORGE J. STIGLER

³¹ Let one seller dispose of q_i , the other sellers each disposing of q . Then the seller's marginal revenue is

$$\frac{d(pq_i)}{dq_i} = p + q_i \frac{dp}{dQ} \frac{dQ}{dq_i},$$

where Q is total sales, and $dQ/dq_i = 1$. Letting $Q = nq_i = nq$, and writing E for

$$\frac{dQ}{dp} \frac{p}{Q},$$

we obtain the expression in the text.

Figure 1: Stigler 1957

Stigler's simple application of the chain rule showed that the underlying assumption of the Marshallian model—atomism, that firms in a competitive industry do not react strategically to the hypothetical actions of other firms—is incompatible with each firm facing a horizontal demand curve. In an n firm industry where the output of the i^{th} firm is q_i , this assumption, means that $\frac{\partial q_i}{\partial q_j} = 0 \forall i \neq j$. As a result, $\frac{dQ}{dq_i} = 1$, and hence

$$\frac{dP}{dq_i} = \frac{dP}{dQ};$$

$$\begin{aligned}
\frac{dP}{dq_i} &= \frac{dP}{dQ} \frac{dQ}{dq_i} \\
&= \frac{dP}{dQ} \left(\sum_{j=1}^n \frac{\partial q_j}{\partial q_i} \right) \\
&= \frac{dP}{dQ} \left(\frac{\partial q_i}{\partial q_i} + \sum_{j \neq i}^n \frac{\partial q_j}{\partial q_i} \right) \\
&= \frac{dP}{dQ} \left(1 + \sum_{j \neq i}^n 0 \right) \\
&= \frac{dP}{dQ}
\end{aligned} \tag{1.1}$$

It is thus impossible for the market demand function $P(Q)$ (where $Q = \sum_{i=1}^n q_i$) to have the dual properties that $P'(Q) < 0$ and $P'(q_i) = 0$ —and Stigler had shown this in 1957! Yet the claim that the market demand curve is negatively sloped, while the individual perfectly competitive firm faces a horizontal demand curve, has graced the opening chapters of every economics textbook published in the last half century.

3. Mendacity in education—another personal observation

One of my motivations for writing *Debunking Economics* was my belief that an education in economics was mendacious. I had in mind the failure to note the Cambridge Controversy arguments when discussing the concept of an aggregate production function (see Chapter 6 of *Debunking Economics*), or the avoidance of the Sonnenschein-Mantel-Debreu conditions when deriving a market demand curve from the aggregation of individual ones (Chapter 2).

When I discussed these issues with any of the minority of neoclassical economists who were themselves aware of those critiques, the even smaller minority who did not dismiss them outright would raise the pedagogic defense of difficulty. These topics are complex, and require an advanced knowledge, not only of economics, but of mathematics. Better to give new students a simple introduction—well behaved aggregate production functions, nice downward sloping market demand curves, and so on—and cover the nuances when they have more knowledge.

No such defense applies here: the only mathematical knowledge needed to comprehend that Marshallian atomism is incompatible with a horizontal demand curve for the firm is elementary calculus.

The responses I have received on this point from neoclassical economists to date have been disingenuous. At best, they have referred to Stigler's attempt to recast perfect competition as the limiting case as the number of firms in an industry increases (discussed in the next section).² At worst, they have claimed that the laws of mathematics do not apply to economics.³

The latter claim is of course nonsense for an approach to economics which, from its founding father to today’s leading exponents, exalted itself over its rivals because it *was* mathematical:

those economists who do not know any mathematics ... can never prevent the theory of the determination of prices under free competition from becoming a mathematical theory. Hence, they will always have to face the alternative either of steering clear of this discipline ... or of tackling the problems of pure economics without the necessary equipment, thus producing not only very bad pure economics but also very bad mathematics. (Walras 1900 [1954]: 47)

This raises the question of why neoclassical economists defend commencing an education in economics with such bad mathematics? We expect it is because the fantasy of perfect competition is essential to fulfilling the vision of rational self-interested behavior being compatible with welfare maximization. If one admits that the individual firm faces a downward-sloping demand curve, then the elimination of deadweight loss that is the hallmark of perfect competition can’t possibly be compatible with individual profit-maximization.

This is easily illustrated using another standard mathematical technique, the Taylor series expansion.⁴

4. Perfect competitors aren’t profit maximizers

Consider a competitive industry where all firms are producing at the “perfect competition” level where price equals marginal cost. In general, profit for the i^{th} firm is:

$$\pi_i(q_i) = P(Q) \cdot q_i - TC(q_i) \tag{1.2}$$

What happens to the i^{th} firm’s profits if it changes its output by a small amount δq_i ? Under the Marshallian condition of atomism, industry output also changes by the same amount. The change in profit $\delta\pi(\delta q_i)$ is thus

$$\pi_i(q_i + \delta q_i) - \pi_i(q_i) = (P(Q + \delta q_i) \cdot (q_i + \delta q_i) - TC(q_i + \delta q_i)) - (P(Q) \cdot q_i - TC(q_i)) \tag{1.3}$$

This can be approximated by applying the first order Taylor series expansion, and by making the substitution that, at this output level, price equals marginal cost: $P(Q) = TC'(q_i)$. The symbolic mathematics engine in Mathcad makes fast work of this approximation:⁵

$$\left[P(Q + \delta q_i) \cdot (q_i + \delta q_i) - TC(q_i + \delta q_i) \right] - (P(Q) \cdot q_i - TC(q_i)) \left| \begin{array}{l} \text{expand} \\ \text{series, } \delta q_i, 1 \\ \text{substitute } \frac{d}{dq_i} TC(q_i) = P(Q) \end{array} \right. \rightarrow q_i \cdot \delta q_i \cdot \frac{d}{dQ} P(Q)$$

Figure 2: Mathcad’s symbolic solution for change in a firm’s profit from perfect competition output level

Therefore $\delta\pi(\delta q_i) \approx q_i \cdot \delta q_i \cdot \frac{d}{dQ} P$. Since $\frac{d}{dQ} P < 0$, if $\delta q_i < 0$ —if, in words, the firm reduces its output—its profit will rise. Thus the output level at which price equals marginal cost is *not* a profit maximum for the individual competitive firm, and if such a firm is indeed a profit maximizer, it will reduce its output below this level.

Some neoclassical economists have thrown the “perfect knowledge” assumption at us at this point: perfectly informed consumers will instantly stop buying from the firm that has reduced its output and increased its price, and switch to those that are still setting price equal to marginal cost. But this argument is still based on the “horizontal demand curve” assumption, which itself is a furphy,⁶ and the market price in the model has already risen because of the change in output by one firm—there is no “cheaper supplier” to whom omniscient consumers can turn.

“Price equals marginal cost” is, therefore, not an equilibrium under the Marshallian assumption of atomism. As a result, the coincidence of collective welfare and the pursuit of individual profit is impossible: if neoclassical economists want to pull that particular rabbit out of a hat, they need another hat. Stigler attempted to provide one.

5. Stigler's limiting case

Stigler, of course, was not trying to bury perfect competition when he showed that $\frac{dP}{dq_i} = \frac{dP}{dQ}$: he was one of the pre-eminent defenders of the neoclassical model against empirically-oriented researchers like Eiteman and Means (see Freedman 1998). He therefore devised an alternative explanation of perfect competition, as the limiting case of competition as the number of firms in an industry increased. His analysis, shown in Figure 1, footnoted the derivation of the expression shown in Figure 3:

$$\begin{aligned} \text{Marginal revenue} &= \text{Price} \\ &+ \frac{\text{Price}}{\text{Number of sellers} \times \text{Market elasticity}} \end{aligned}$$

Figure 3: Stigler's expression for marginal revenue (Stigler 1957: 8)

Stigler then asserted that “this last term goes to zero as the number of sellers increases indefinitely” (Stigler 1957: 8). Marginal revenue for the i^{th} firm thus converges to market price. Perfect competition thus appeared to be saved, despite a downward-sloping firm's demand curve: profit-maximizers would set marginal cost equal to marginal revenue, and this would converge to price as more firms entered a market.

Stigler's convergence argument is technically correct, but in conflict with the proof shown above that “price equals marginal cost” is *not* a profit maximum for the individual firm. The resolution of this conflict led to Keen's first truly original contribution to this literature: *the proposition that equating marginal revenue and marginal cost maximizes profit is also a furphy.*

6. Equating MC and MR doesn't maximize profits

Generations of economists have been taught the simple mantra that “profit is maximized by equating marginal cost and marginal revenue”. The proof simply differentiates (1.2) with respect to q_i . However, the individual firm's profit is a function, not only of its own output, but of that of all other firms in the industry. This is true regardless of whether the firm reacts strategically to what other firms do, and regardless of whether it can control what other firms do. The objectively true profit maximum is therefore given by the zero of the *total* differential: the differential of the firm's profit with respect to total industry output.

We stress that this issue is independent of whether the individual firm can or cannot work out this maximum for itself, whether the firm does or does not interact with its competitors, and whether the firm does or does not control the variables that determine the profit maximum. Given a mathematically specified market inverse demand function that is a function of the aggregate quantity supplied to the market, and a mathematically specified total cost function for the individual firm that is a function of its output, the question “what is the level of the firm's output that maximizes its profit?” is completely independent of the question of “will the firm, in any given environment, or following any given behavioral rule, actually determine or achieve this level?”. That objective, profit-maximizing level is given by the zero of the *total* differential of profit:

$$\frac{d}{dQ} \pi(q_i) = \frac{d}{dQ} (P(Q)q_i - TC(q_i)) = 0 \quad (1.4)$$

This total derivative is the sum of n partial derivatives in an n -firm industry:

$$\frac{d}{dQ} \pi(q_i) = \sum_{j=1}^n \left\{ \left(\frac{\partial}{\partial q_j} \pi(q_i) \right) \cdot \frac{d}{dQ} q_j \right\} \quad (1.5)$$

In the Marshallian case, atomism lets us set $\frac{d}{dQ} q_j = 1 \forall j$ (the Cournot case is addressed in a subsequent paper). Expanding (1.5) yields

$$\frac{d}{dQ} \pi(q_i) = \sum_{j=1}^n \left(\frac{\partial}{\partial q_j} (P(Q)q_i - TC(q_i)) \right) \quad (1.6)$$

Continuing with the product rule, (1.6) can be expanded to:

$$\frac{d}{dQ} \pi(q_i) = \sum_{j=1}^n \left(P(Q) \frac{\partial}{\partial q_j} q_i + q_i \cdot \frac{\partial}{\partial q_j} P(Q) - \frac{\partial}{\partial q_j} TC(q_i) \right) \quad (1.7)$$

Under the Marshallian assumption of atomism, the first term in the summation in (1.7) is zero where $j \neq i$, and $P(Q)$ where $j = i$. The second term is equal to $q_i \cdot \frac{d}{dQ} P(Q) \forall j$; the third is zero where $j \neq i$, and equal to $\frac{d}{dq_i} TC(q_i)$ (or marginal cost $MC(q_i)$) where $j = i$. Equation (1.7) thus reduces to

$$\frac{d}{dQ}\pi(q_i) = P(Q) + n \cdot q_i \cdot \frac{d}{dQ}P(Q) - MC(q_i) \quad (1.8)$$

The true profit maximum—under the Marshallian condition of atomism—is thus given by equation (1.9):

$$\pi(q_i)_{\max} : MC(q_i) = P + n \cdot q_i \cdot \frac{dP}{dQ} \quad (1.9)$$

The error in the standard “Marshallian” formula is now obvious: it omits the number of firms in the industry (n) from the expression for the individual firm’s marginal revenue. With this error corrected, the correct profit-maximizing rule for a competitive firm is very similar to that for a monopoly: set marginal cost equal to *industry level* marginal revenue.⁷

7. Monopoly, competition, profit and hyper-rationality

Neoclassical economics assumes that, given revenue and cost functions, there is some output level that will maximize profits, and another that will maximize social welfare (by eliminating deadweight loss).⁸ The argument that the two coincide under perfect competition has been shown to be nonsense. So too is the argument that a single rational firm could work out the profit maximum, but a bunch of rational *non-interacting* firms couldn’t, as the calculus in the previous section shows.

Of course, an objection can be made to the above mathematical logic that solving equation (1.9) requires knowledge of the number of firms in the industry, which the individual competitive firm can’t be assumed to have.⁹ Here, we can turn Milton Friedman’s methodological defense of the theory of the firm against itself. Friedman, as is well known, argued that while firms didn’t in fact do calculus to work out their profit-maximizing output levels, we could model their behavior “as if” they did, because

unless the behavior of businessmen in some way or other approximated behavior consistent with the maximization of returns, it seems unlikely that they would remain in business for long. (Friedman 1953: 22)

We are not arguing that firms do the calculus to work out this profit-maximizing level either.¹⁰ Instead, we are simply showing that the calculus *can* be done, and the profit-maximizing level is not the one asserted by neoclassical economists. However, it is possible now—in a way that wasn’t possible in 1953—to actually carry out Friedman’s “billiard players” experiment. Citing him again:

Now, of course, businessmen do not actually and literally solve the system of simultaneous equations in terms of which the mathematical economist finds it convenient to express this hypothesis, any more than leaves or billiard players explicitly go through complicated mathematical calculations or falling bodies decide to create a vacuum. The billiard player, if asked how he decides where to hit the ball, may say that he “just figures it out” but then also rubs a rabbit’s foot just to make sure; and the businessman may well say that he prices at average cost, with of course some minor deviations when the market

makes it necessary. The one statement is about as helpful as the other, and neither is a relevant test of the associated hypothesis. (Friedman 1953: 22)

A “relevant test of the associated hypothesis” is to set up a virtual market that conforms to neoclassical assumptions—with a static downward sloping market demand curve, and given cost functions subject to diminishing marginal productivity, so that there is indeed a profit-maximizing level of output for each firm—and see what happens. Figure 4 shows a Mathcad program that implements this.¹¹

```

Firms := Seed(rand)
for i ∈ firmsmin, firmsmin + firmssteps .. firmsmax
    Q0 ←  $\begin{cases} \text{round}(\text{runif}(i, q_K(i), q_C(i))) & \text{if } i > 1 \\ q_C(i) & \text{otherwise} \end{cases}$ 
    P0 ←  $\begin{cases} P(\sum Q_0, a, b) & \text{if } i > 1 \\ P(q_C(i), a, b) & \text{otherwise} \end{cases}$ 
    dq ←  $\begin{cases} \text{round}\left(\text{rnorm}\left(i, 0, \frac{q_C(i)}{100}\right)\right) & \text{if } i > 1 \\ \frac{q_C(i)}{100} & \text{otherwise} \end{cases}$ 
    for j ∈ 0..runs - 1
        Qj+1 ← Qj + dq
        Pj+1 ←  $\begin{cases} P(\sum Q_{j+1}, a, b) & \text{if } i > 1 \\ P(Q_{j+1}, a, b) & \text{otherwise} \end{cases}$ 
        dq ←  $\frac{\text{sign}[(p_{j+1} \cdot Q_{j+1} - p_j \cdot Q_j) - (tc(Q_{j+1}, i) - tc(Q_j, i))]}{dq}$ 
        Fj, i-1 ← Qj
F

```

Figure 4: Simulation of instrumental profit maximizers

Working through the program line by line:

1. A random number generator is seeded
2. A *for* loop iterates from a minimum number to a maximum number of firms
3. If there is more than one firm in the industry, each firm is randomly allocated an initial output level. The amounts are uniformly distributed from a minimum of the Keen prediction for a profit-maximizing firm, q_K to a maximum of the neoclassical prediction q_C .
4. If there is only one firm in the industry, its output starts at the level predicted by the neoclassical model—which coincides with q_K .

5. An initial market price is set, based on the sum of initial outputs.
6. Line 6 sets the market price in the case of a monopoly.
7. Each firm is randomly allocated an amount by which it varies output. The distribution has a mean of zero and a standard deviation of 1% of the neoclassical prediction for a profit-maximizing firm's output (this is the last aspect of the program that involves probability).
8. Line 8 allocates a change amount of 1% of the predicted output for a monopoly.
9. A *for* loop iterates over a number of runs where each firm varies its output trying to increase its profit from the initial level.
10. Firstly each firm adds its change amount to its initial output. This is a vector operation: if there are 100 firms in the industry Q_0 is a vector with 100 initial output amounts, and dq is a vector with 100 (positive or negative) output changes.
11. A new market price is calculated on the basis of the new aggregate output level.
12. Line 12 again allows for a monopoly.
13. Each firm then calculates whether its profit has risen or fallen as a result of its change in output, and the collective impact of all the changes in output on the market price. If a firm finds that its profit has risen, it continues to change output in the same direction; if its profit has fall, it changes its output by the same amount but in the opposite direction.
14. Each step in the iteration is stored in a multi-dimensional array.¹²
15. The multidimensional array is returned by the program.

The program was run with identical cost functions for each firm, set up so that the market aggregate marginal cost curve was independent of the number of firms in the industry (we return to this issue in the Appendix). The number of firms was varied from 1 to 100. The eventual aggregate output at the end of 1000 iterations is shown in Figure 5, and the corresponding market price is shown in Figure 6, against the predictions of the Neoclassical and the Keen approach respectively.




Figure 5: Aggregate output

As is obvious, the number of firms in an industry had no impact on the eventual market output level or price: the Neoclassical prediction that price would converge to the level at which price equals marginal cost clearly was not fulfilled.




Figure 6: Market price

Some neoclassical referees thought that the results occurred because, though all firms were acting independently, they were all doing the same thing (reducing output from the initial level), and thus acting in a semi-collusive way.¹³ In fact, as Figure 7 and Figure 8 show, though the average outcome conformed to Keen's predictions, the individual firms all pursued very different strategies. The aggregate outcome, which contradicted the neoclassical prediction and confirmed Keen's, was the result of quite diverse individual firm behavior—despite all firms having identical cost functions.

Figure 7 shows the output levels of 3 randomly chosen firms from the 100 firm industry, the average for all firms, and the predictions of the Keen and neoclassical formulae. Firm 1 began near the neoclassical output level, rapidly reduced output towards the "Keen" level, but then reversed direction; Firm 2 began halfway between the neoclassical and Keen predictions, then reduced output below the Keen level and stayed there; Firm 3 began closer to the neoclassical level and meandered closer to the Keen level.




Figure 7: Firm outputs in 100 firm industry

The sole source of the volatility of each firm's behavior is the complex impact of interactions between firms, in the context of a very simply defined market—there is no random number generator causing this volatility. As Figure 8 shows, each firm made its changes in response to the impact of both its changes in output, and the collective changes in output, on its profit. Some firms made larger profits than others—notably the firms with the larger output made the larger profits. However, the average profit was much higher than predicted by the neoclassical model.

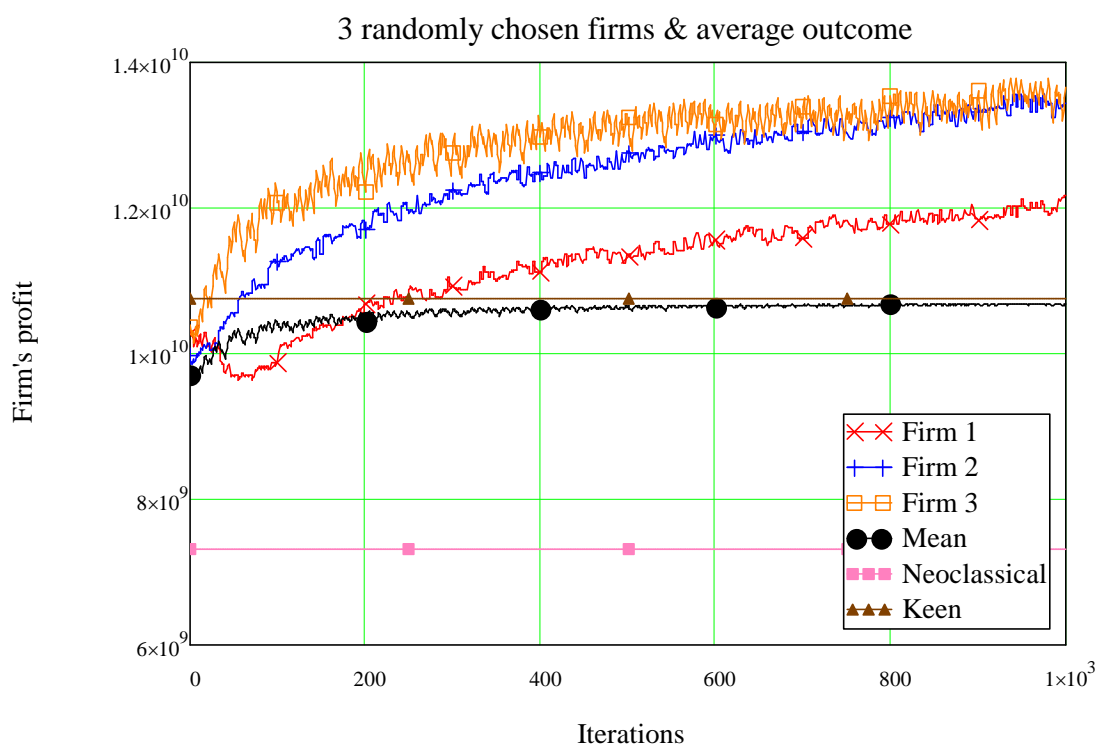


Figure 8: Firm profits in 100 firm industry

This model indicates that, in this game of competitive profit maximization, the virtual equivalents of Friedman's "billiard players" follow the laws of mathematics in their search for profits, as Friedman argued. However, these laws of mathematics are incompatible with the beliefs of neoclassical economists.

Since hyper-rational profit-maximizers cannot be relied upon to save neoclassical belief, there are only two avenues left: irrational behavior, and Cournot-Nash game theory.

8. Price-taking behavior is irrational

A regular neoclassical rejoinder to our analysis has been that we are "cheating" by not assuming rational, price-taking behavior. Our standard reply that the assumption of "price-taking" behavior is itself cheating, with regard to the laws of mathematics: as shown in Section 2, the assumption that $P'(q_i) = 0$ is incompatible with the assumption of a downward-sloping market demand curve ($P'(Q) < 0$). However, it is also easily shown that "price-taking behavior" is irrational.

The assumption of price-taking behavior appears regularly in neoclassical economics, from the level of Marshallian analysis through to the foundations of general equilibrium analysis (see for example Mas-Colell et al 1995: 314, 383). Neoclassical economists do not seem to realize that this is a classic "rabbit in the hat" assumption: if it is assumed,

then the “perfectly competitive” result of price equaling marginal cost follows, *regardless of how many firms there are in the industry*.

The essence of price-taking is the belief that a firm's change in its output doesn't affect market price: this amounts to setting $\frac{\partial}{\partial q_j} P(Q) = 0$ in equation (1.7). This results in the “profit-maximizing strategy” of setting price equal to marginal cost, independently of the number of firms—that is, once this assumption is made, *even a monopoly produces where price equals marginal cost*. This behavior is clearly irrational for a monopoly, and it is only the “fog of large numbers”—the confusion of infinitesimals with zero, as Keen noted in *Debunking Economics*—that led neoclassical economists to regard price-taking as rational behavior for competitive firms.

Figure 9 illustrates that price-taking behavior is *irrational*: an agent who behaves this way is necessarily making a logical error. If the market demand curve slopes downwards, then the *a priori* rational belief is that *any* increase in output by the firm will depress market price.

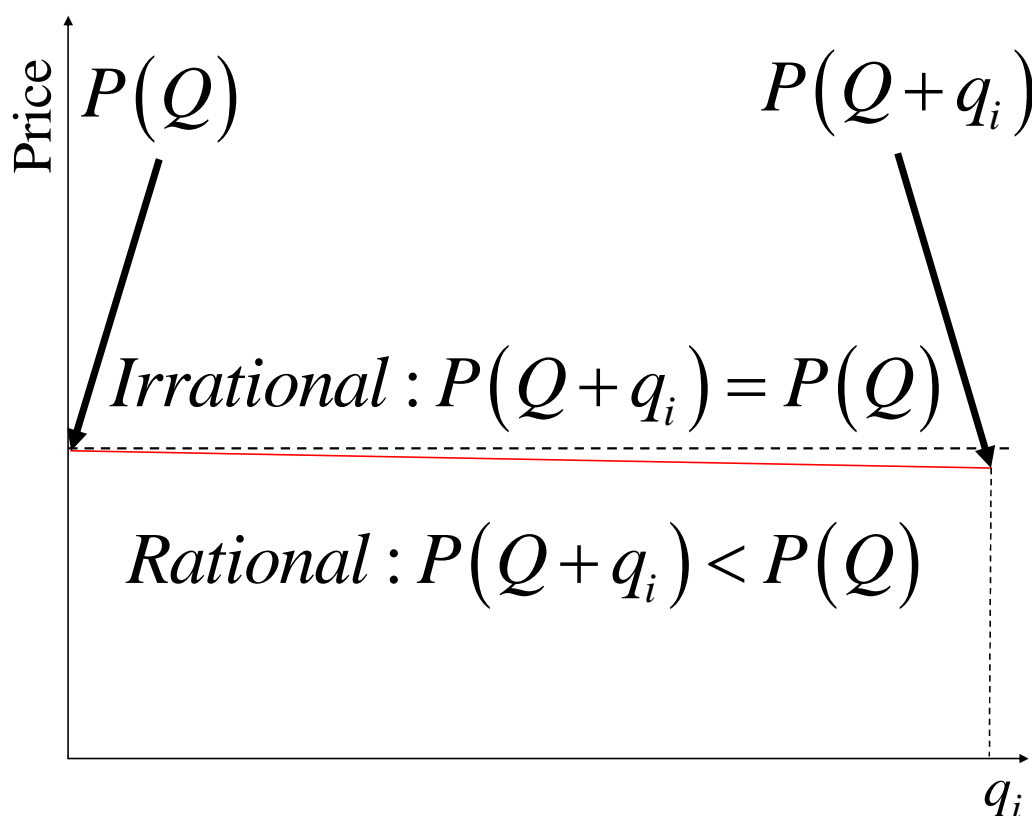


Figure 9: Irrationality of "price-taking" behavior

The desired neoclassical outcome of price equal to marginal cost is thus dependent on irrational behavior (in the context of Marshallian competition—Cournot competition is covered in a subsequent paper). We quantify the degree of irrationality needed with a variant on the program shown in Figure 4, in which a proportion of firms in an industry

of a specified size actually do behave irrationally: if a strategy has caused an increase in profit, a fraction of firms respond by *reversing* that strategy.

The program is shown in Figure 10. The outer loop (line 2) iterates a counter i from 0 to 50, with the value representing the fraction of firms who behave irrationally at each iteration. The only change to the inner loop is that the change in output by each firm is now reversed for $i\%$ of firms at each iteration.¹⁴

```

Firms := Seed(rand)
        for i ∈ 0..50
            Q0 ← round(runif(firms, qK(firms), qC(firms)))
            p0 ← P(∑ Q0, a, b)
            dq ← round(rnorm(firms, 0,  $\frac{q_C(\text{firms})}{100}$ ))
            for j ∈ 0..runs - 1
                Qj+1 ← Qj + dq
                pj+1 ← P(∑ Qj+1, a, b)
                dq ←  $\left[ \text{sign} \left[ \text{runif} \left( \text{firms}, \frac{-i}{100}, \frac{-i}{100} + 1 \right) \cdot \left[ (p_{j+1} \cdot Q_{j+1} - p_j \cdot Q_j) - (tc(Q_{j+1}, \text{firms}) - tc(Q_j, \text{firms})) \right] \right] \right] \cdot dq$ 
                Fj,i ← Qj
        F

```

Figure 10: Analyzing the impact of irrationality

Figure 11 shows the aggregate outcome for a 100 firm industry. With no irrationality, the industry produces the amount predicted by the Keen formula. Output then increases almost monotonically as the degree of irrationality rises—until, when 20 per cent of firms are behaving irrationally at each iteration, market output converges to near the neoclassical output level.

For a degree of irrationality between 20% and 45%, the neoclassical outcome continues to dominate the simulation results. Then as irrationality rises above this level, the market effectively follows a random walk—where, curiously, profits in general tend to be *higher* than what would apply if each firm equated marginal revenue and marginal cost.




Figure 11: Market output as a function of the degree of irrationality

Figure 12 shows the behavior of three randomly chosen firms, and the average behavior, at a 20% level of irrationality—i.e., when one firm in five reverses any strategy that benefited it on the previous iteration.

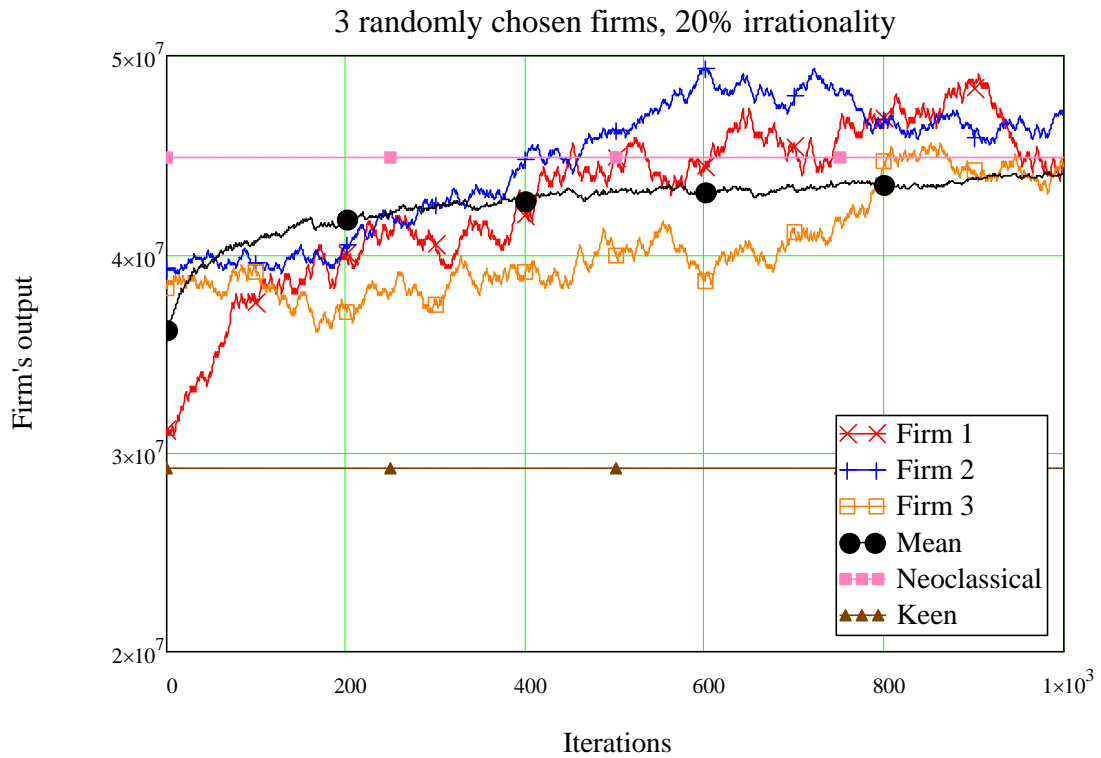


Figure 12: Sample outputs at 20% irrationality

Figure 13 shows the impact that a degree of irrationality of 20% has on firms' profits. Profit falls throughout the run, until by the end, it is almost (but not quite) as low as that caused by equating marginal revenue and marginal cost.




Figure 13: Impact of 20% irrationality on firms' profits

Ironically, higher profits apply if firms simply follow a random walk than if they apply the neoclassical formula (see Figure 14).

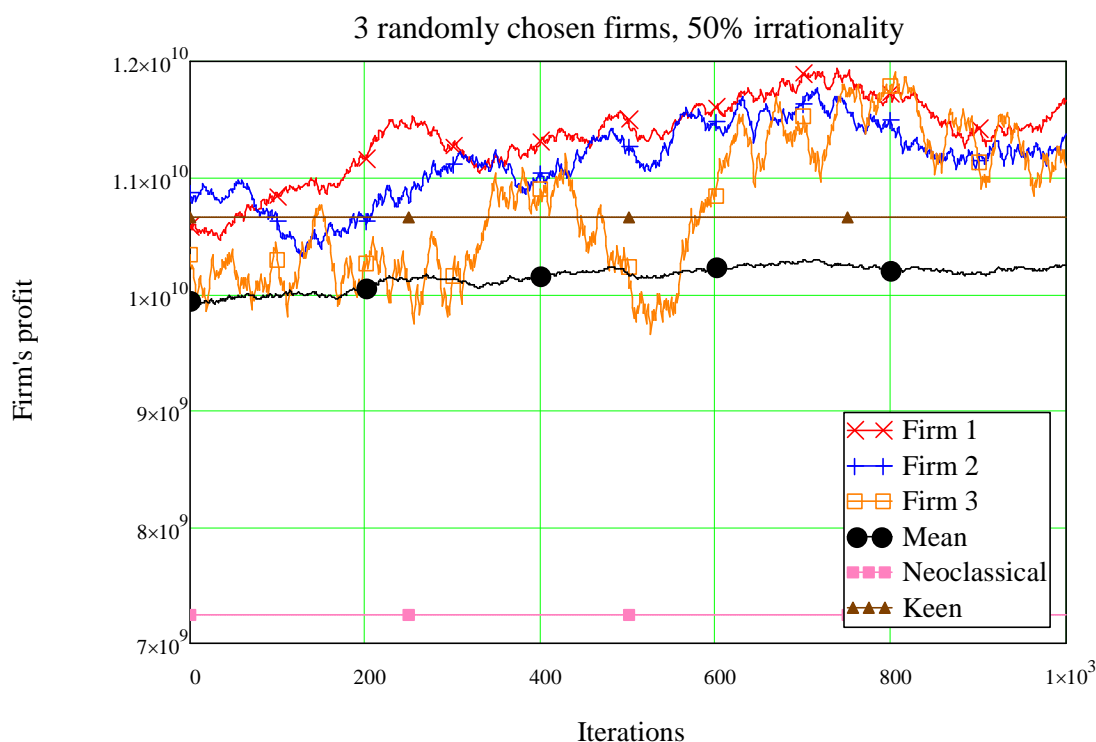


Figure 14: Firm profits with 50% irrationality

A degree of irrational behavior thus saves the neoclassical preferred outcome of price equal to marginal cost—though with some collateral damage, since it is now clearly neither profit-maximizing, nor rational. The question remains, what might help ensure this level of irrationality? Cournot-Nash game theory appears to provide an answer in strategic interactions between firms—though this answer is only unequivocal at a very superficial level of analysis. Though Cournot-Nash mathematical analysis is not subject to the flaws that infest the Marshallian argument, Keen and Standish 2006 shows that the Cournot outcome is dynamically unstable.

9. The empirical reality of competition

A plethora of empirical studies have established that at least 89 per cent of output—and perhaps as much as 95 per cent—is produced under conditions of constant or falling marginal cost, and rising economies of scale.¹⁵ Given such circumstances, there is no profit-maximizing level of output for the individual firm: so long as the sale price exceeds average costs, the firm will profit from additional sales. The key presumption of the neoclassical model—that there is a profit-maximizing level of sales—is thus not fulfilled in the real world.

The most recent such survey was carried out by Alan Blinder and a team of PhD students in 1998. Blinder's results are also arguably the most authoritative, given the scale of his study, and Blinder's prestige as an economist.

Blinder et al. surveyed a representative weighted sample of US non-agricultural corporations with annual sales of more than US\$10 million; a 61% response rate resulted in a study of 200 corporations whose combined output represented 7.6% of the USA's GDP. The interviews were face to face, with Blinder and a team of Economics PhD students conducting the interviews; the interviewees were top executives of the firms, with 25% being the President or CEO, and 45% a Vice President.

Blinder summarized the results in the following way:

“First, about 85 percent of all the goods and services in the U.S. nonfarm business sector are sold to "regular customers" with whom sellers have an ongoing relationship ... And about 70 percent of sales are business to business rather than from businesses to consumers...

Second, and related, contractual rigidities ... are extremely common ... about one-quarter of output is sold under contracts that fix nominal prices for a nontrivial period of time. And it appears that discounts from contract prices are rare. Roughly another 60 percent of output is covered by Okun-style implicit contracts which slow down price adjustments.

Third, firms typically report fixed costs that are quite high relative to variable costs. And they rarely report the upward-sloping marginal cost curves that are ubiquitous in economic theory. Indeed, downward-sloping marginal cost curves are more common... If these answers are to be believed ... then [a good deal of microeconomic theory] is called into question... For example, price cannot approximate marginal cost in a competitive market if fixed costs are very high.” (p. 302)

The key final point about falling marginal cost deserves elaboration. Given that, as they discovered, “marginal cost is not a natural mental construct for most executives.” they

translated marginal cost into ‘variable costs of producing additional units,’ and posed the following question:

B7(a). Some companies find that their variable costs per unit are roughly constant when production rises. Others incur either higher or lower variable costs of producing additional units when they raise production.

How would you characterize the behavior of your own variable costs of producing additional units as production rises? (Blinder 1998: 102)

The survey team collated the responses into five groups, as summarized in Table 1:

Structure of Marginal Costs	Percentage of firms
Decreasing	32.6
Decreasing with discrete jumps	7.9
Constant	40
Constant with discrete jumps	7.9

Increasing	11.1
------------	------

Table 1: Marginal cost structure of American corporations (Blinder et al. 1998: 102-103)

Blinder et al. pithily observed that:

“The overwhelmingly bad news here (for economic theory) is that, apparently, only 11 percent of GDP is produced under conditions of rising marginal cost.” (102)

The overall results of Blinder's survey are summarized in Table 2. Given the empirically common circumstances detailed here, the pre-requisites for being able to identify a profit-maximizing level of output do not exist for at least 89 per cent of US firms.¹⁶ Instead, for these firms, the only profit-maximizing strategy is to sell as much as they can—and at the expense, where possible, of competitors' sales.

Summary of Selected Factual Results Price Policy	
Median number of price changes in a year	1.4
Mean lag before adjusting price months following	
Demand Increase	2.9
Demand Decrease	2.9
Cost Increase	2.8
Cost Decrease	3.3
Percent of firms which	
Report annual price reviews	45
Change prices all at once	74
Change prices in small steps	16
Have nontrivial costs of adjusting prices	43
of which related primarily to	
the frequency of price changes	69
the size of price changes	14
Sales	
Estimated percent of GDP sold under contracts	
which fix prices	28
Percent of firms which report implicit contracts	65
Percent of sales which are made to	
Consumers	21
Businesses	70
Other (principally government)	9

Regular customers	85
Percent of firms whose sales are	
Relatively sensitive to the state of the economy	43
Relatively Insensitive to the state of the economy	39
Costs	
Percent of firms which can estimate costs at least moderately well	87
Mean percentage of costs which are fixed	44
Percentage of firms for which marginal costs are	
Increasing	11
Constant	48
Decreasing	41

Table 2: Summary of Blinder et al.'s empirical findings

The only practical way that this can be done is via product differentiation, and that indeed is the obvious form that real competition actually takes. Innovation and heterogeneity are the true hallmarks of competition, yet these concepts are effectively excluded by the neoclassical model.

A model of how this actual form of competition works would be extremely useful to economic theory—and perhaps even to economic policy, if we could scientifically identify those industry structures that truly promote innovation. The continued teaching of the neoclassical model, and the continued development of a research tradition in which rising marginal cost plays a key role, are a hindrance to developing an adequate model of real world competition.

Our closing observation on this theory is perhaps the most important. A theory is more than a scholastic exercise: a good theory is also an attempt to understand reality, and, where possible, to alter it for the better. There are, therefore, few things more dangerous than an applied bad theory. Unfortunately, neoclassical competition theory is applied throughout the world, in the guise of policies intended to promote competition.

10. The anti-capitalist nature of neoclassical competition policy

The neoclassical vision of competition has been enshrined in competition policies adopted by governments and applied to key industries such as telecommunications, power, sanitation, and water supply. The major practical implications of accepted theory are that more firms equates to increased competition, increased competition means higher output at lower prices, and market price should ideally be equal to marginal cost.

Since the theory is flawed, these implications are at best unproven, and at worst false. There are now numerous instances around the world where competition policies have resulted in deleterious outcomes; a special issue of *Utilities Policy* in 2004 details several of these for the USA (and Australia). Loube, for example, examined the US Telecom Act

of 1996, and found that “this policy has actually raised prices for residential customers” (Trebing & Miller 2004: 106).

Proponents of competition policy normally ascribe such outcomes to poor implementation of policy, poor regulatory oversight, or unanticipated circumstances. However, if the theory is flawed as we argue, then these outcomes are not accidents, but the systemic results of imposing a false theory on actual markets. Some predictable negative consequences are rising costs due to reduced economies of scale, underinvestment caused by imposed prices that lie below average cost, and reduced rates of innovation in related industries, caused by the inadequate “competitive” provision of infrastructure.

That these policies were imposed in a well-meaning attempt to improve social welfare cannot detract from the fact that, if the theory guiding these policies was false, then the policies are likely to cause more harm than good. Real world markets would function far better if competition policy, as it stands, were abolished.

11. Conclusion

A careful examination of the neoclassical theory of competition thus finds that it has little, if any, true content.

The Marshallian argument, which forms the backbone of neoclassical pedagogy, is strictly false in its belief that a downward-sloping market demand curve is compatible with horizontal individual firm demand curves. Once this error is corrected, the model's major conclusion, that competitive industries are better than concentrated ones, is overturned. Given identical demand and cost conditions, competitive industries will produce the same output as monopolies, and sell at the same price—and there are good grounds for expecting that monopolies would have lower costs (see Appendix One).

Given this intrinsic barrenness of the theory, its empirical irrelevance is even more important. Neoclassical economists have ignored a multitude of empirical papers that show that marginal cost does not rise, that firms do not compete on price, and so on, on the basis of Friedman's belief that asking businessmen what they do is not “a relevant test of the associated hypothesis.” But if the “associated hypothesis” is in fact false, or irrelevant, then “asking businessmen what they do” is at least a good place from which to derive stylized facts that a relevant hypothesis would have to explain. It is high time that economists abandoned what superficially appears to be “high theory”, and got their hands dirty with real empirical research into actual firms and actual competition.

Here the picture that emerges from even a cursory examination of the data is very different to neoclassical belief. Table 3 shows the aggregate distribution of firm sizes in the USA in 2002: large firms make up well under 0.3 per cent of the total number of firms, but are responsible for over 60 per cent of sales.

Industry		2002		
		Total	0-499	500+
Total	Firms	5,697,759	5,680,914	16,845
	Estab.	7,200,770	6,172,809	1,027,961
	Emp.	112,400,654	56,366,292	56,034,362

Ann. pay.(\$000)	3,943,179,606	1,777,049,574	2,166,130,032
Receipts(\$000)	22,062,528,196	8,558,731,333	13,503,796,863

Table 3: US firm size data (US Office of Small Business Advocacy)

At the same time, small firms are not negligible: all industries are characterized by a wide distribution of firm sizes, from sole trader through to large conglomerates (see Table 4). Perhaps the real story of competition is the survival of such diversity.

	Manufacturing				
	Firms	Estab.	Emp.	Ann. pay.(\$000)	Receipts(\$000)
Total	297,873	344,341	14,393,609	580,356,005	3,937,164,576
0 *	21,731	21,761	0	2,231,805	15,166,970
1-4	97,197	97,232	219,951	5,823,048	27,659,982
5-9	55,597	55,702	372,245	10,533,204	44,184,220
10-19	46,851	47,200	639,036	19,888,764	80,892,263
0-19	221,376	221,895	1,231,232	38,476,821	167,903,435
20-99	58,198	62,443	2,375,691	82,257,351	346,024,892
100-499	14,124	23,727	2,488,018	91,152,085	460,526,128
0-499	293,698	308,065	6,094,941	211,886,257	974,454,455
500+	4,175	36,276	8,298,668	368,469,748	2,962,710,121

Table 4: Distribution of firm sizes in manufacturing (US SBA)

In the light of both its theoretical weaknesses and its irrelevance to the empirical data, Sraffa's advice in 1930 about what to do with Marshall's theory bear repeating today:

the theory cannot be interpreted in a way which makes it logically self-consistent and, at the same time, reconciles it with the facts it sets out to explain. Mr. Robertson's remedy is to discard mathematics, and he suggests that my remedy is to discard the facts; perhaps I ought to have explained that, in the circumstances, I think it is Marshall's theory that should be discarded. (Sraffa 1930: 93)

The Marshallian theory of competition is a hindrance to understanding real markets and real competition, and it should be abandoned.

12. *Appendices*

Appendix One: Conditions for comparability of cost structures

Economists blithely draw diagrams like Figure 15 below to compare monopoly with perfect competition. As shown above, the basis of the comparison is false: given Marshallian assumptions, an industry with many "perfectly competitive" firms will produce the same amount as a monopoly facing identical demand and cost conditions—and both industry structures will lead to a "deadweight loss". However, in general, small competitive firms would have different cost conditions to a single firm—not only because of economies of scale spread result in lower per unit fixed costs, but also because of the impact of economies of scale on marginal costs.

The Inefficiency of Monopoly...




Figure 15: Mankiw's monopoly versus perfect competition comparison

Rospot (1993) gives a good illustration of this latter point in relation to gas utilities. One of the fixed costs of gas supply is the pipe; one of the variable costs is the compression needed to move the gas along the pipe. A larger diameter pipe allows a larger volume of gas to be passed with lower compression losses, so that the larger scale of output results in lower marginal costs:

Simply stated, the necessary first investment in infrastructure is the construction of the pipeline itself. Thereafter, additional units of throughput can be economically added through the use of horsepower to compress the gas up to a certain point where the losses associated with the compression make the installation of additional pipe more economical than the use of additional horsepower of compression. The loss of energy is, of course, a function of, among other things, the diameter of the pipe. Thus, at the outset, the selection of pipe diameter is a critical ingredient in determining the economics of future expansions of the installed pipe: the larger the diameter, the more efficient are the future additions of capacity and hence the lower the marginal costs of future units of output (Rospot 1993: 288).

Thus a single large supplier is likely to have lower costs—in which case, the marginal cost curve for the monopoly should be drawn *below* that for the competitive industry. Given the same demand curve and the same profit-maximizing behavior, a monopoly is thus likely to produce a higher output than a competitive industry, and at a lower cost.

The cost examples in this paper were artificially constructed to ensure that the assumption of identical costs embodied in Figure 15 were fulfilled—something that we doubt has been done by neoclassical authors in comparable papers. The cost functions were:

$$\begin{aligned} \text{Monopoly: } MC(Q) &= C + D \cdot Q + E \cdot Q^2 \\ \text{Competitive: } mc(q, n) &= C + D \cdot n \cdot q + E \cdot n^2 \cdot q^2 \end{aligned} \quad (1.10)$$

Obviously, it is very arbitrary to have the number of firms in an industry as an argument in the marginal cost function of a single firm—and also highly unlikely. Yet without that “heroic” assumption, the aggregate of marginal costs curves for a competitive industry will *necessarily* differ from the marginal cost curve for a monopoly. If a monopoly has greater access to economies of scale than smaller competitive firms, as in Rosput's example of gas transmission, then on conventional profit-maximizing grounds, a monopoly would produce a higher output for a lower price.

It is also easily shown that the neoclassical pedagogic assumption that the same marginal cost curve can be drawn for a competitive industry and a monopoly is true in only two circumstances: either the monopoly simply changes the ownership of plants in the industry—so that there is no technical difference between one industry structure and the other—or both industry structures face identical *constant* marginal costs.¹⁷

Marginal cost is the inverse of marginal product, which in turn is the derivative of total product. The condition of identical marginal costs—that is, that the marginal cost curve for a monopoly is identically equal to the sum of the marginal cost curves of an industry with many competitive firms, for all relevant levels of output—therefore requires that the total products of two different industry structures differ only by a constant. This constant can be set to zero, since output is zero with zero variable inputs.

Consider a competitive industry with n firms, each employing x workers, and a monopoly with m plants, each employing y workers, where $n > m$. Graphically this condition can be shown as in Figure 16.




Figure 16: Production functions required for identical marginal cost curves

Using f for the production function of the competitive firms, and g for the production function of the monopoly, the equality of total products condition can be put in the following form:

$$n \cdot f(x) = m \cdot g(y) \tag{1.11}$$

Substitute $y = \frac{n \cdot x}{m}$ into (1) and differentiate both sides of (1.11) by n :

$$f(x) = \frac{x}{m} \cdot g'\left(\frac{n \cdot x}{m}\right) \tag{1.12}$$

This gives us a second expression for f . Equating these two definitions yields:

$$\frac{g\left(\frac{n \cdot x}{m}\right)}{n} = \frac{x}{m} \cdot g'\left(\frac{n \cdot x}{m}\right)$$

or

$$\frac{g'\left(\frac{n \cdot x}{m}\right)}{g\left(\frac{n \cdot x}{m}\right)} = \frac{m}{n \cdot x} \tag{1.13}$$

The substitution of $y = \frac{n \cdot x}{m}$ yields an expression involving the differential of the log of g :

$$\frac{g'(y)}{g(y)} = \frac{1}{y} \quad (1.14)$$

Integrating both sides yields:

$$\ln(g(y)) = \ln(y) + c \quad (1.15)$$

Thus g is a constant returns production function:

$$g(y) = C \cdot y \quad (1.16)$$

It follows that f is the *same* constant returns production function:

$$f(x) = \frac{m}{n} \cdot C \cdot \frac{n \cdot x}{m} \quad (1.17)$$

With both f and g being identical constant returns production functions, the marginal products and hence the marginal costs of the competitive industry and monopoly are constant and identical. The general rule, therefore, is that welfare comparisons of perfect competition and monopoly are only definitive when the competitive firms and the monopoly operate under conditions of constant identical marginal cost.

13. References

- Blinder, A.S., Canetti, E., Lebow, D., & Rudd, J., (1998). *Asking About Prices: a New Approach to Understanding Price Stickiness*, Russell Sage Foundation, New York.
- Eiteman, W.J. & Guthrie, G.E., (1952). "The shape of the average cost curve", *American Economic Review* 42: 832-838.
- Craig Freedman (1998). "No ends to means: George Stigler's profit motive", *Journal of Post Keynesian Economics* 20; 621-648.
- Milton Friedman (1953). "The Methodology of Positive Economics", in *Essays in Positive Economics*, University of Chicago Press, Chicago: 3-43.
- Haines, W.W., (1948). "Capacity production and the least cost point", *American Economic Review* 38: 617-624.
- Steve Keen (1993a). "The misinterpretation of Marx's theory of value", *Journal of the History of Economic Thought*, 15 (2), Fall, 282-300.
- Steve Keen (1993b). "Use-value, exchange-value, and the demise of Marx's labor theory of value", *Journal of the History of Economic Thought*, 15 (1), Spring, 107-121.
- Steve Keen (2001). *Debunking Economics*, Pluto Press & Zed Books, Sydney & London.
- (2004a). "Why economics must abandon its theory of the firm", in Salzano, M., & Kirman, A. (eds.), *Economics: Complex Windows*, Springer, New York, pp. 65-88.
- (2004b). "Deregulator: Judgment Day for Microeconomics", *Utilities Policy*, 12: 109 – 125.
- (2004c). "Improbable, Incorrect or Impossible: The persuasive but flawed mathematics of microeconomics", in Fullbrook, E. (ed.), *Student's Guide to What's Wrong with Economics*, Routledge, London, pp. 209-222.

- Steve Keen and Russell Standish (2005). "Irrationality in the neoclassical definition of rationality", *American Journal of Applied Sciences* (Sp.Issue): 61-68.
- Steve Keen and Russell Standish (2006). "Profit Maximization, Industry Structure, and Competition: A critique of neoclassical theory", *Physica A* 370: 81-85.
- Lee, F., (1998). *Post Keynesian Price Theory*, Cambridge University Press, Cambridge.
- A. Mas-Colell, M.D. Whinston, J.R. Green (1995). *Microeconomics*, Oxford University Press, New York.
- Piero Sraffa (1930). "Increasing Returns And The Representative Firm A Symposium", *Economic Journal* 40 pp.79-116.
- George J. Stigler (1957), Perfect competition, historically contemplated. *Journal of Political Economy* 65: 1-17.
- Harry M. Trebing and Edythe S. Miller (2004). "Introduction", *Utilities Policy* 12: 105-108.
- Leon Walras (1899 [1954]). *Elements of Pure Economics*, 4th edition, George Allen & Unwin, London.

¹ This is the subject of a subsequent paper. The early proofs in this paper were also developed with the assistance of John Legge (La Trobe University)

² A referee for the *Economic Journal* wrote that "we always consider the perfect competition case as a polar case which represents an extreme scenario, and is largely a benchmark. I would prefer to see the equation: $(AR - MR)/AR = 1/(nE)$, so that for E at a normal value of say 2, and n at 1000, then the divergence of AR and MR is 1/20¹ of 1%. Then price equals MR seems a pretty good approximation!"

³ A referee for the *Journal of Economic Education* commented that "Stigler's many attempts to save neoclassical theory have always caused more problems than they have solved. His version of the chain rule is contrary to the partial equilibrium method and thus is irrelevant".

⁴ This proof was first developed by John Legge, of La Trobe University.

⁵ We are using a symbolic mathematics program both to reduce the need for some tedious manual calculations, and because on several occasions, neoclassical economists have disputed the results of manual calculations—by in effect distorting the definition of a derivative.

⁶ "furphy" is a delightful Australian word meaning "an irrelevant or minor issue raised to specifically divert attention away from the real issue". It deserves wider currency—especially amongst economists!

⁷ Though not necessarily identical, since $n \cdot q_i$ only equals Q if $q_i = \frac{Q}{n} \forall i$. This impact of dispersal in firm size may explain some of the simulation results shown later.

⁸ We use standard undergraduate terms here because the analysis we are challenging is, up to this point, that served up to undergraduates. We address game theoretic concepts later.

⁹ Equation (1.9) can be put in another form which partly addresses this criticism, and also emphasizes the error in the conventional formula. The profit-maximizing level of output is not to equate firm-level marginal revenue and marginal cost, but to make the gap between them a fraction of the gap between price and marginal cost: $MR(q_i) - MC(q_i) = \frac{n-1}{n}(P - MC)$. The fraction tends to 1 as $n \rightarrow \infty$, so the more "competitive" an industry is, the easier it is to apply this formula.

¹⁰ In fact, we argue later that the assumption that there is some profit-maximizing level for a firm is a furphy. The profit-maximizing strategy for actual firms is simply sell as much as possible, at the expense where possible of your competitors and other possible avenues for consumers' discretionary expenditure.

¹¹ The behavior modeled was deliberately made as simple as possible, to avoid the rejoinder that the results were the product of our algorithm rather than raw profit-motivated behavior. It could only have been simpler by having each firm vary its output by one unit at each time step—a modification which, as it happens, results in a much slower but absolute convergence to the Keen equilibrium.

¹² In effect, F is a matrix where the j^{th} and i^{th} column contains the vector of outputs by an i firms industry at the j^{th} iteration.

¹³ A referee for the *Economic Journal* commented that “if firms act the same way, they will all get higher profits if and only if they reduce outputs. Then the algorithm will continue to lead them to the monopoly outcome since there is no chance any firm can realize the true impact of its own output change. Thus the result is not surprising.”

¹⁴ The function call `runif(firms, -i/100, -i/100 + 1)` generates a vector of numbers between $-i/100$ and $1 - i/100$; when $i=0$, all these numbers will be positive and thus not affect the value of the `sign()` function; when $i>0$, $i\%$ of these numbers will be negative and thus the sign of the `sign()` function will be reversed. The firms that have this randomly assigned negative number against their output change will increase output at the next step if profit rose when the decreased output on the previous step (and vice versa). This is instrumentally irrational behavior.

¹⁵ See Lee 1998 for a comprehensive survey of the 20th century studies.

¹⁶ We say at least because all previous surveys have reported a lower proportion of products that are produced under conditions of diminishing marginal productivity—typically 5 per cent of output (Eiteman & Guthrie 1952).

¹⁷ This argument was first published in Keen 2004a.