9 A pluralist approach to microeconomics

Steve Keen

A great strength of traditional economics is the absence of a well-developed, coherent alternative. The pressure to teach *something* often results in orthodox microeconomics ruling the roost. However, political economists should not be afraid to teach approaches which, in apparent contrast to the logically complete traditional economics of the firm, are inchoate and do not answer every question. A new approach is never born complete but evolves, and the process of teaching an alternative from an incomplete starting point can lead to its development over time.

However, an essential first is to demonstrate to students that the ostensibly well-developed and coherent traditional model is in fact an empty shell. That is difficult in a principles course, since the increase in business majors and the relegation of economics to a service role has dumbed down the content so much that critiquing it is problematic: it is hard to critique something that is itself so nebulous.

The potential for a critique arises at the intermediate level, where the mathematical treatment is first encountered. Many political economists eschew mathematics, often because it is seen as part of why traditional economics is so flawed. Ironically, however, it is precisely when traditional economics is presented mathematically that it is most vulnerable – especially at the level of intermediate microeconomics, where the foundations are still essentially Marshallian – because the mathematics itself is fallacious.

I begin my teaching of heterodox microeconomics by recapping traditional microeconomics, and then demonstrating that the following two key aspects of the theory are mathematically false.

1 Under the assumptions of the traditional model, a competitive market populated by profit-maximizing firms will produce a higher output than a monopoly, and at a lower price (Keen 2004; Keen and Standish 2006).
2 The market demand curve derived from a set of utility-maximizing consumers is necessarily downward sloping (Gorman 1953; Shafer and Sonnenschein 1982).

Once these two assertions are demonstrated to be fallacies, the task of convincing students that a different approach should be considered – even if it is incomplete – is much easier.

Testing Marshall

I begin my course with a computer simulation that demonstrates the falsity of the first proposition, and effectively turns Friedman's methodological defense of orthodoxy on its head (Friedman 1953). Friedman argued that while expert billiard players did not know "the complicated mathematical formulas that would give the optimum directions of travel … unless in some way or other they were capable of reaching essentially the same result, they would not in fact be expert billiard players" (1953: 21). By analogy, he argued that the same could be said of firms: while they did not do calculus to set their output levels, unless they behaved

> as if … they knew the relevant cost and demand functions, calculated marginal cost and marginal revenue … and pushed each line of action to the point at which the relevant marginal cost and marginal revenue were equal … it seems unlikely that they would remain in business for long. Let the apparent immediate determinant of business behavior be anything at all – habitual reaction, random chance, or whatnot. Whenever this determinant happens to lead to behavior consistent with rational and informed maximization of returns, the business will prosper … whenever it does not, the business will tend to lose resources.
>
> (1953: 21–22)

I put Friedman to the test in class, using a multi-agent model of a market. This model uses standard market demand and aggregate marginal cost curves, with equations and parameter values as shown in Equation (9.1).

$$P(Q) = a - b \cdot Q$$

$$MC(Q) = c + d \cdot Q$$

$$MR(Q) = a - 2 \cdot b \cdot Q$$

$$TC(Q) = c \cdot Q + \frac{1}{2} \cdot d \cdot Q^2 + k$$

$$\text{where } a = 800;\ b = 10^{-8};\ c = 100;\ d = 10^{-8};\ k = 10^6 \tag{9.1}$$

These parameter values generate a model market with realistic output levels so that a simulated comparison can be made between a single monopoly producer and a competitive industry with, say, 10,000 firms.[1] Neoclassical theory then makes the following predictions for the output levels of a monopoly and competitive industry respectively:

$$\text{Monopoly } M_Q = \frac{a - c}{2b + d} = 2.333 \cdot 10^{10}$$

$$\text{Perfect Competition } PC_Q = \frac{a - c}{b + d} = 3.5 \cdot 10^{10} \tag{9.2}$$

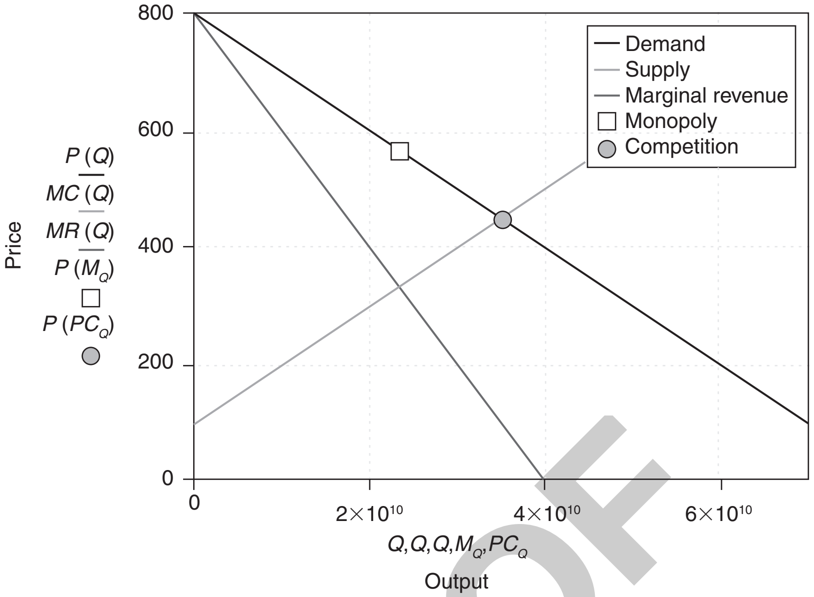

Figure 9.1 Predictions of the Marshallian model.

The model and its predictions are shown in Figure 9.1.

As Friedman notes, actual firms do not do calculus, but follow other procedures which, *if the orthodox model is correct*, must nonetheless correspond to them behaving as if they were setting marginal cost equal to marginal revenue. My model tests this by populating this artificial market with agents who follow the simplest possible rule of thumb for profit maximization: choose an output level, and then change it by a fixed amount (either positive or negative). If profit increases, keep moving in the same direction; if profit decreases, move in the opposite direction by the same amount. I then run the model with a single firm, and also 10,000 firms, and check the results. The results[2] of two typical runs are shown in Figure 9.2.

The theory's prediction for the monopoly level of output is correct, but the prediction for the competitive industry is clearly wrong: rather than producing where supply equals demand, the competitive industry produces much the same level as the monopoly. "Oh dear, something has gone terribly wrong": these instrumental profit-maximizers *don't* do what neoclassical theory predicts! The individual firms all follow very different strategies (see Appendix A), which are extremely complex despite the simple nature of the behavioral algorithm (see Figure 9.3).

The firms also achieve much higher profits from their simple rule of thumb than orthodox theory predicts (see Figure 9.4). They are clearly better at making profits than orthodoxy is at predicting the profit-maximizing output level. This simulation thus sets the scene for a comprehensive demolition of the Marshallian

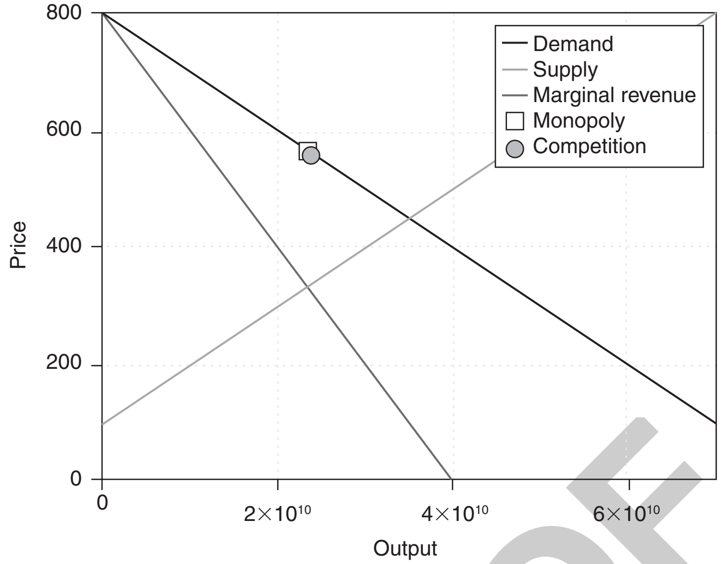

Figure 9.2 Simulation results.

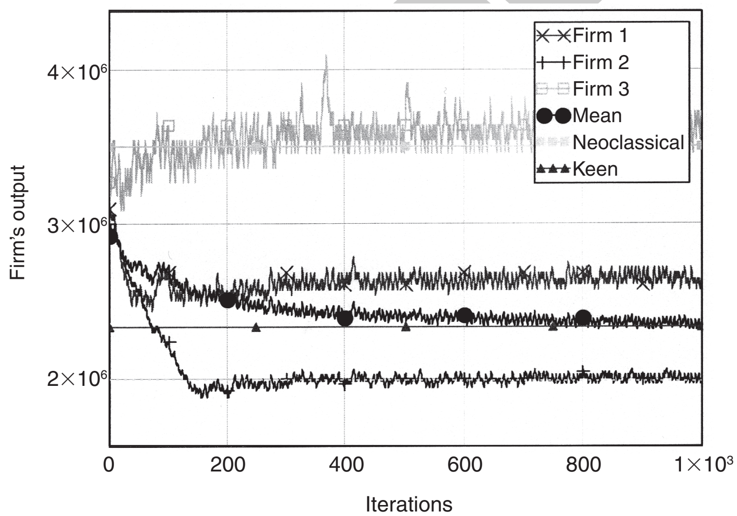

Figure 9.3 Convergence of individual outputs (three randomly chosen firms and average outcome).

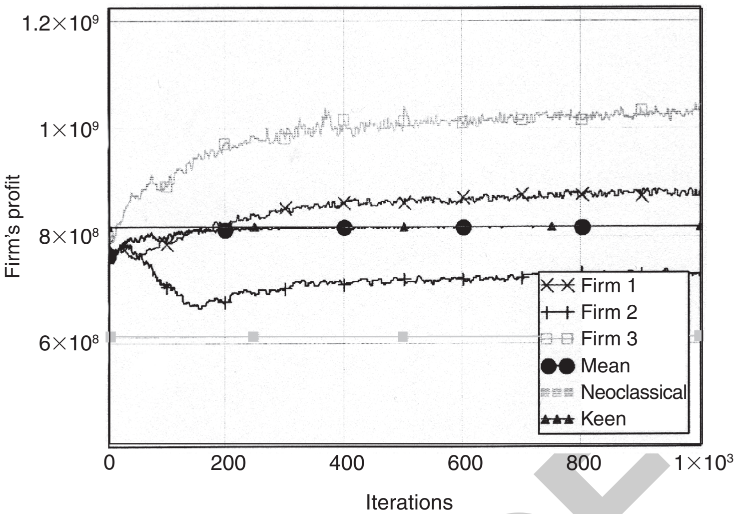

Figure 9.4 Much higher profits result from the firms' "rule of thumb" (three randomly chosen firms and average outcome).

model. The first step in this process is proving that a key proposition of the neoclassical model, that competitive firms face a horizontal demand curve, is mathematically false under the Marshallian assumption of atomism.

Refuting Marshall

This result was first proven in 1957 by, of all people, George Stigler (1957: footnote 31), as shown in Figure 9.5.

$$\frac{dp}{dq_i} = \frac{dp}{dQ}$$

Stigler's logic simply applied the assumption of atomism which characterizes the Marshallian model of competition[3] – that competitive firms neither know of, nor react strategically to, the output decisions of other firms. Given that assumption, if the i^{th} firm changes its output by an amount dq_i, other firms in the industry don't react – and therefore industry output Q changes by the same amount, so that $\frac{dQ}{dq_i} \cdot 1$. Given this result, the conclusion that the slope of the demand curve perceived by the competitive firm $\frac{dp}{dq_i}$ is precisely the same as the slope of the market demand curve $\frac{dp}{dQ}$ is derived by simply applying the chain rule:

$$\frac{dp}{dq_i} = \frac{dp}{dQ} \cdot \frac{dQ}{dq_i} = \frac{dp}{dQ} \tag{9.3}$$

This in turn means that the demand curve perceived by the individual firm is *not* horizontal, and that marginal revenue for the competitive firm is less than price:

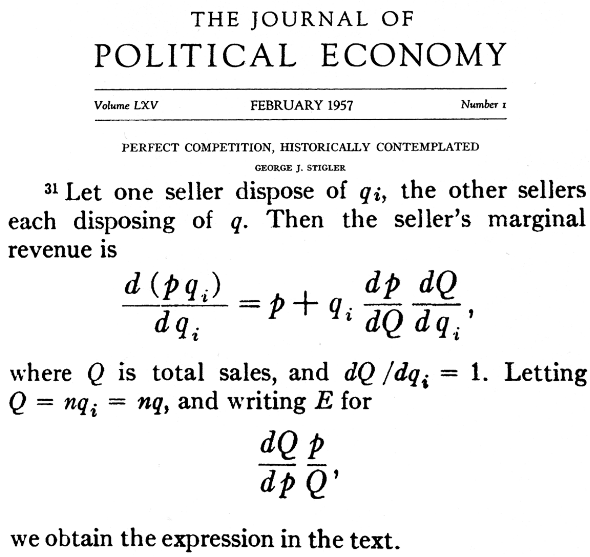
THE JOURNAL OF
POLITICAL ECONOMY

Volume LXV FEBRUARY 1957 Number 1

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 9.5 Stigler's proof.

$$\begin{aligned} MR_i &= \frac{d}{dq_i}(P \cdot q_i) \\ &= P \cdot \frac{d}{dq_i}q_i + q_i \cdot \frac{d}{dq_i}P \\ &= P + q_i \cdot \frac{d}{dQ}P < P \end{aligned} \quad (9.4)$$

Though the mathematics of this result is straightforward, the fallacy of the horizontal demand curve is so strongly ingrained[4] that I find I have to provide a multi-pronged attack on the commonly held defenses of this fallacy. Multiple counters to defenses of it in the light of this result. The three most common defenses are:[5]

1 The equation $\left(\frac{dp}{dq_i} = 0\right)$ is just an assumption.
2 The omniscient consumer argument, that if a firm charges above the market price, it will have no customers, while if it charges below the market price, it will face the entire industry demand curve (see, for example, Varian 2006: 6, Figure 22.1).
3 That competitive firms behave as if they face a horizontal demand curve, or that they are too small to perceive the negative slope of the demand curve they face.[6]

The first proposition appears to be an application of Friedman's dictum that a theory cannot be tested by the "realism" of its "assumptions" (Friedman 1953: 23), but in fact it is a mathematical fallacy. It asserts that it is valid to have a model in which mathematically incompatible assumptions play an essential role. Assuming a negatively sloped market demand curve $\frac{d}{dQ}P < 0$, and atomism (so that $\frac{\partial}{\partial q_i}q_j = 0$), then it follows that $\frac{d}{dq_i}P = \frac{d}{dQ}P < 0$, as Stigler showed.

The second contradicts the assumption of price-taking behavior, which is also an essential aspect of the model of competitive behavior: competitive firms do not set price, but produce a quantity and then accept whatever price the market demand curve throws back at them. Once a single firm has changed its output, then all firms will receive the new market price, and *there is no seller charging a lower price to whom the consumers can turn*.

The third argument is a possibility, but only if firms behave *irrationally*. *If* the demand curve for the market is negatively sloped, *and* atomism applies, *then* the demand curve for the individual firm is negatively sloped: to believe otherwise is to behave irrationally[7] (see Figure 9.6). The too-small-to-perceive slope argument is also contradicted by the computer simulation shown above: even with 10,000 firms in the artificial market, the aggregate result contradicts the outcome that would apply if this defense were valid.

Once students have accepted the mathematical truth that $\frac{d}{dq_i}P = \frac{d}{dQ}P$ under the assumption of atomism, we proceed to the coup de grace for the Marshallian model: the neoclassical mantra that profits are maximized by equating marginal cost to marginal revenue is false in a multi-firm industry. The easiest proof for intermediate micro students[8] is the following: assume that all competitive firms follow the advice of neoclassical theory and set their marginal revenue equal to their marginal cost. Then for an n-firm industry, the sum of this across all firms will also be zero:

$$\sum_{i=1}^{n}(MR_i(q_i) - MC_i(q_i)) = 0 \qquad (9.5)$$

This can be expanded to the following, using the crucial result that $\frac{d}{dq_i}P = \frac{d}{dQ}P$:

$$\sum_{i=1}^{n}\left(P(Q) + q_i \cdot \frac{d}{dQ}P(Q)\right) - \sum_{i=1}^{n}MC_i(q_i) = 0 \qquad (9.6)$$

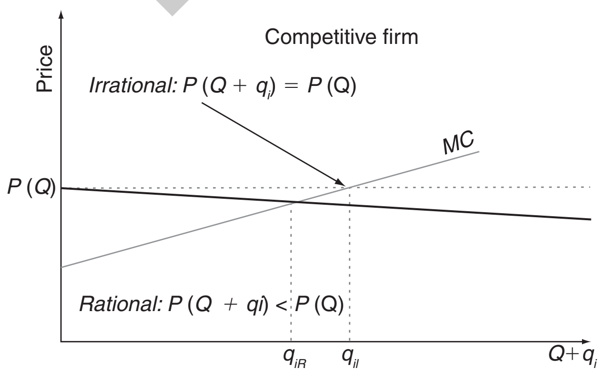

Figure 9.6 The belief that the firm faces a horizontal demand curve.

Expanding the summation over n firms from equation (9.6) yields n copies of P from the first term, $Q \cdot \frac{d}{dQ}P$ from the second, and n copies of marginal cost ($MC(Q)$) from the third[9] so that:

$$n \cdot P(Q) + Q \cdot \frac{d}{dQ}P(Q) - n \cdot MC(Q) = 0; \text{ or } (n-1) \cdot P(Q) + MR(Q) - n \cdot MC(Q) = 0 \tag{9.7}$$

It is then possible to rearrange equation (0.7) to yield this expression in terms of industry-level marginal revenue, marginal cost, and price:

$$MR(Q) - MC(Q) = -(n-1)(P(Q) - MC(Q)) \tag{9.8}$$

This result demonstrates the aggregation fallacy in the neoclassical so-called profit-maximizing formula: if each firm sets its output so that its marginal revenue equals marginal cost at the level of the individual firm, market output will *exceed* the point at which marginal revenue equals marginal cost. As a result, some of the output produced will be produced at a loss – and therefore each individual firm is producing part of its output at a loss if it produces where marginal revenue equals marginal cost. The actual profit-maximizing rule in terms of marginal revenue and marginal cost can be derived by equating equations (9.5) and (9.8):

$$\sum_{i=1}^{n}(MR_i(q_i) - MC_i(q_i)) = (n-1) \cdot P(Q) + MR(Q) - n \cdot MC(Q) \tag{9.9}$$

and then rearranging terms to leave market-level MR(Q) and MC(Q) on one side:

$$\sum_{i=1}^{n}(MR_i(q_i) - MC_i(q_i)) - (n-1) \cdot (P(Q) - MC(Q)) = MR(Q) - MC(Q) \tag{9.10}$$

then bring terms inside the summation and equate market-level marginal revenue and marginal cost to find the aggregate profit maximum:

$$\sum_{i=1}^{n}\left(MR_i(q_i) - MC_i(q_i) - \frac{n-1}{n} \cdot (P(Q) - MC(Q))\right) = 0; \text{ so that}$$

$$\text{For profit maximization set } MR_i(q_i) - MC_i(q_i) = \frac{n-1}{n} \cdot (P(Q) - MC(Q)) \tag{9.11}$$

The actual profit-maximizing rule – the one the instrumental profit maximizers in the multi-agent simulation were clearly following – is thus not to equate marginal cost and marginal revenue, but to make the gap between them equal to $(n-1)/n$ times the gap between price and marginal cost.

The final step in establishing the hollowness of the Marshallian model of competition is to demonstrate that, if a competitive industry produces the same amount as a monopoly when their cost structures happen to coincide[10] then on Marshallian grounds a monopoly should be preferred to a competitive industry if its costs are lower, since it will produce a larger amount at a lower price. In the real world, economies of scale normally mean that a monopoly has lower marginal costs than the smaller firms of hypothetical competitive industry, further strengthening the neoclassical case in favor of monopolies over competitive industries!

So much for the neoclassical model of supply. Turning to the model of demand, we find that it is equally flawed.

The shape of the market demand curve

The derivation of an individual demand curve from a set of indifference curves and a budget constraint is straightforward. However, the process of summing individual demand curves to derive a market demand curve is a non-trivial problem because, in the traditional model of a market economy, consumer incomes are determined by prices and quantities set in markets. Changing relative prices therefore changes incomes – something that is ignored when an individual's demand curve is derived, but which can't be ignored when aggregating to derive the market demand curve in a single market.

Over half a century ago, Gorman proved that the only condition under which a market demand curve necessarily had the same characteristics as an individual demand curve is "that an extra unit of purchasing power should be spent in the same way no matter to whom it is given" (Gorman 1953: 64) – in other words, that the distribution *and scale* of income have no effect on consumption. This in turn requires (a) that all Engels curves are straight lines (homothetic preferences); and (b) that all consumers have parallel Engels curves. Without these restrictions, then *a market demand curve can have any shape at all*.[11] This result – now known as the Sonnenschein–Mantel–Debreu (SMD) conditions after their rediscovery by these researchers in the 1970s – is clearly and emphatically articulated in the authoritative *Handbook of Mathematical Economics*:

> First, when preferences are homothetic and the distribution of income (value of wealth) is independent of prices, then the market demand function (market excess demand function) has all the properties of a consumer demand function.... Second, with general (in particular non-homothetic) preferences, even if the distribution of income is fixed, market demand functions need not satisfy in any way the classical restrictions which characterize consumer demand functions.... The utility hypothesis tells us nothing about market demand unless it is augmented by additional requirements.
>
> (Shafer and Sonnenschein 1982: 671–672)

In contrast, the treatment of this same issue in Varian's *Intermediate Microeconomics* borders on mendacity. In his discussion of individual demand he spends several pages discussing homothetic preferences before concluding that they aren't very realistic (Varian 2006: 102). Later, in his chapter on market demand, he notes the dilemma that the aggregate demand (for a market) will generally depend on prices and the *distribution* of incomes, but continues:

> However, it is sometimes convenient to think of the aggregate demand as the demand of some representative consumer who has an income that is just the sum of all individual incomes. The conditions under which this can be done are rather restrictive, and a complete discussion of this issue is beyond the scope of this book.
>
> (Varian 2006: 267)

Students are thus left with the impression that realistic individual Engels curves are compatible with well-behaved market demand curves – an impression intensified by Varian's chapter summary that begins: "The market demand curve is simply [*sic!*] the sum of the individual demand curves" (Varian 2006: 281).

In reality, the SMD conditions are the transformation problem of neoclassical economics: the two conditions can only strictly apply in a one-consumer *and one-commodity* world, since in a multi-agent world changing prices will change the distribution of income, while in a multi-commodity world increasing income will alter relative demand, which in turn will change the distribution of income.

An intellectually honest response to these results is,

> If we are to progress further we may well be forced to theorise in terms of groups who have collectively coherent behaviour. Thus demand and expenditure functions if they are to be set against reality must be defined at some reasonably high level of aggregation. The idea that we should start at the level of the isolated individual is one which we may well have to abandon.
>
> (Kirman 1989: 138)

The SMD conditions thus validate the focus of the classical economists on class-based analysis: while it is nonsensical to aggregate all consumers into a representative agent and all products into a representative commodity, there is some validity in treating different classes as having coherent tastes, and consuming uniform commodities. Distribution of income within a class can then be ignored – but distribution *between* classes cannot. Nor can the distribution of income be reduced to a market process, because the dilemma of market demand curves having any shape at all undermines the proposition that the return to a factor of production is its marginal product.

From an empty shell to emergent properties

It would be possible to continue with other flaws in traditional microeconomics[12] but I prefer to use the SMD conditions to segue into a crucial, but hard to understand, insight from complexity theory: the concept of emergent properties – that a complex system will have properties that can't be understood simply by understanding the isolated properties of the entities that compose it. The SMD conditions show that an economy consisting of perfectly well-behaved neoclassical agents will *not* behave like a scaled-up individual consumer at the market level, because the relations between consumers – the distribution of income – dominate the isolated behavior of each individual at the aggregate level of a market. The fact that a market demand curve derived from aggregating downward-sloping individual demand curves can have any shape at all is thus a classic emergent property.

This draws a line in the sand between micro and macro – the phenomenon of emergent properties in a complex system means that there is a limit to reductionism, whereas the neoclassical research program is essentially reductionist. Neoclassical economics effectively demolished Keynesian macroeconomics in the 1960–1970s, arguing that it did not have good microfoundations – with the explicit proposition that an economy populated by neoclassical-defined agents could not demonstrate the macro-phenomena of involuntary unemployment, a key tenet in the Keynesian perspective. In fact, an economy populated by neoclassically defined agents can't even generate the essential neoclassical parable of a downward-sloping market demand curve. It can thus be said that neoclassical microeconomics doesn't have good microfoundations either!

The SMD conditions enable the instructor to preface a pluralist approach to microeconomics with the caveat that there is a legitimate divide between microeconomics and macroeconomics: macroeconomics cannot be reduced to additive microeconomics.

A political economy alternative

The ultimate reason why a pluralist approach to microeconomics should supplant a monist one is that realism should be the guiding principle of economic analysis – and traditional neoclassical microeconomics is both internally flawed and unrealistic. The starting point of pluralist microeconomics therefore should be the facts: the actual data on industry structure, and the behavior of firms and consumers. That alone will distinguish a pluralist course from a neoclassical one – by way of illustration, there is not one single table of empirical data in Varian (2006).

However, I prefer to take an interlude, after critiquing neoclassicism, with a market simulation known as Starpower. Not only does it put the many functions of a market and self-interested exchange in context, it also enlivens a topic that neoclassical pedagogy has made mind-numbingly dull, and forces students to engage directly and personally with each other, very early on during a semester.

Starpower

This is a multi-person trading game that works best with between fifteen and twenty participants, but can work with as few as twelve and as many as thirty. I introduce it after a free-ranging discussion of the merits and drawbacks of the market system, where comments by students normally provide a fertile basis for debriefing after the game – and a reasonable measure of the extent to which it affects their opinions. I have played this game over 500 times, and every time it has caused a dramatic shift in initial perceptions of a market economy.

I describe it as a trading game. Its objective is to amass 1200 points, and it is over once three players have reached that level. Each player takes five poker chips from one of three boxes – which commence with identical distributions of five colors of chips – and then calculate their initial point score on the basis of Table 9.1.

The value of different combinations of chips produces opportunities for trade: for example, while two green chips are worth only 20 points, three are worth 60. A player with two green chips will gain from trading a red or white in return for an extra green. Trading must be on a one-for-one basis, and players are required to shake hands while th-y trade. They also can't reveal their hand to anyone else, and can only trade by saying what they will give in return for what they want – a red for a green, for example. If they can't reach agreement, then the traders are forced to remain holding hands until the end of the trading session – which lasts about two minutes.

After the first round, players are ranked by score, divided into three groups – Squares, Circles, and Triangles – and required to wear corresponding badges. The chips are then collected from each group and returned to a box *which then belongs to that group for the remainder of the game* – which of course simulates the inheritance of wealth (and of poverty). Two further two rounds ensue, and after each round students are moved between groups if their scores warrant it – and of course, they take their chips with them. Before round 4, it is announced that the Squares have been doing so well that they can make the rules for the next round.

As you can imagine, pandemonium can ensue: the top group can make any rule changes they like – so long as they reach a consensus – and normally they make rules that favour them (as well as often removing technicalities like having to shake hands to trade). The game then continues, ending normally with three of the Squares winning (though in the 1970s it would often end in revolution!), after which the students are debriefed.

Obvious questions arise – such as whether the game was fair, whether it was realistic, and so on. Important major points involve the role of inheritance and chance, the relatively minor role of trading in enhancing wealth, and the three tendencies that invoke inequality – inheritance, mobility between groups, and the trading table itself that necessarily encourages the top group to accumulate

Table 9.1 Starpower scoring table

Colours	*Number of chips*				
	1	*2*	*3*	*4*	*5*
Yellow	60	120	180	240	300
Blue	40	70	130	170	220
Red	30	50	100	160	180
White	20	40	90	120	160
Green	10	20	60	80	130

high-value chips and the lowest group to accumulate low-value ones. As well as provoking discussion about how good a model of capitalism Starpower itself is, it also helps emphasize the point that models, even heterodox ones, are models and not the real thing. The closest we will get to that comes from surveys and empirical data.

Just the facts, ma'am

There is fortunately a wealth of empirical data on firms available from the U.S. Census Bureau (2008) (www.census.gov/csd/susb/susb.htm) and the U.S. Small Business Administration Office of Advocacy (2008). The latest available aggregate dataset is shown in Table 9.2.

The first two columns in the secondary table at the bottom of Table 9.2 are graphed on a log-log scale in Figure 9.7: U.S. firm size follows a "scale free" power law distribution, which emphasizes the irrelevance of the neoclassical model of the firm: there is no ideal firm size. Instead, the distribution of firm sizes in the real world follows what physicists have dubbed a Power Law. Graphically, this results in a straight line plot between the number of firms of a given size and that size when both are plotted in logs.[13] Intuitively, this means there is no average, representative, or ideal firm size – the distribution of firm sizes is instead scale free. The actual process of competition has resulted in a distribution from many very small firms to many very large, so that just as there is no representative-sized animal in biology, there is no representative-sized firm in economics.

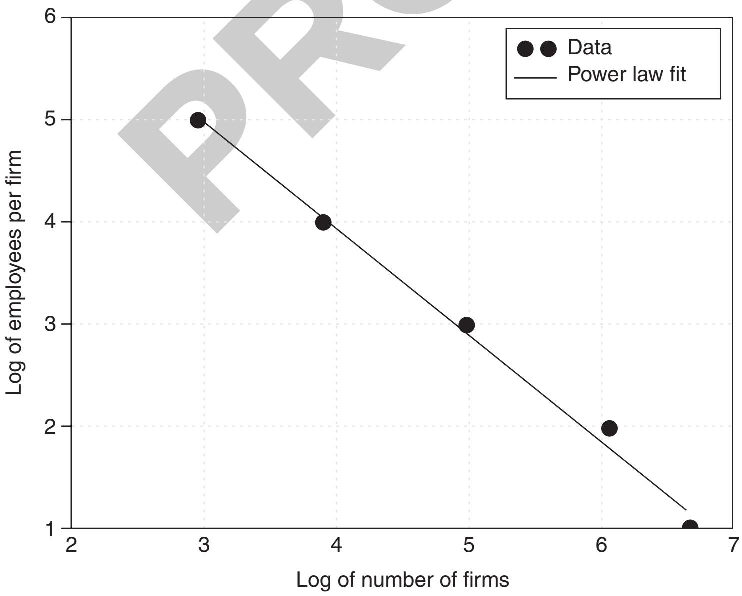

Figure 9.7 U.S. firm size follows a "scale free" power law distribution (2005).

Table 9.2 The size distribution of U.S. firms

Employer firms, establishments, employment, and annual payroll small firm size classes, 2005

Employment size of firm	*Firms*	*Establishments*	*Employment*	*Annual payroll ($1,000)*
Total	5,983,546	7,499,702	116,317,003	4,482,722,481
0–3	823,832	824,952	0	42,182,002
1–4	2,854,047	2,859,095	5,936,859	177,827,102
5–9	1,050,062	1,062,907	6,898,859	206,178,084
10–14	415,989	432,470	4,865,539	153,325,562
15–19	213,957	229,727	3,588,315	116,091,356
20–24	131,514	147,060	2,870,060	94,111,977
25–29	88,097	101,840	2,365,072	78,099,071
30–34	63,260	76,225	2,016,475	67,807,561
35–39	47,373	50,241	1,746,960	59,433,250
40–44	36,656	48,154	1,535,517	52,703,860
45–49	29,143	39,773	1,366,993	47,040,730
50–74	84.607	130,095	5,095,569	178,105,960
75–99	40,247	75,994	3,447,703	123,150,994
100–149	38,694	93,959	4,673,931	169,007,646
150–199	18,538	61,697	3,189,340	115,639,275
200–299	17,383	82,949	4,208,878	153,071,046
300–399	7,999	52,447	2,756,388	103,080,535
400–499	4,671	40,947	2,082,503	75,725,730
500–749	5,823	67,664	3,539,488	135,650,216
750–999	2,878	43,464	2,478,859	95,138,017
1,000–1,499	2,845	56,614	3,456,833	139,104,676
1,500–2,499	2,314	75,406	4,435,321	185,189,876
2,500–4,999	1,787	111,752	6,199,781	276,630,183
5,000–9,999	918	123,808	6,438,639	297,593,815
10,000+	912	600,462	31,123,497	1,340,823,957
Max employees	*Firms*	*Establishments*	*Employment*	*Payroll*
10	4,727,941	4,746,954	12,835,342	426,187,188
100	1,159,843	1,341,579	28,898,203	969,870,321
1,000	95,986	443,127	22,929,387	847,322,465
10,000	7,864	367,580	20,530,574	898,518,550
100,000	912	600,462	31,123,497	1,340,623,957

These data confirm Marshall's assertion: "the Mecca of the economist lies in economic biology rather than in economic dynamics" (Marshall 1920: 19), because this kind of distribution manifests itself in systems subject to evolutionary competition. Of course, Marshall did not develop this apt analogy – that was done by Schumpeter, to whom I turn after considering two further pieces of empirical research which confirm that neoclassical micro is a dead-end.

The law of constant marginal product

It is not commonly appreciated that Friedman's methodology paper was intended to derail empirically researched actual firm behavior, but evidence abounds throughout:

> The lengthy discussion on marginal analysis in the *American Economic Review* some years ago … neglect[s] what seems to me clearly the main issue – the conformity to experience of the implications of the marginal analysis – and concentrate[s] on the largely irrelevant question whether businessmen do or do not in fact reach their decisions by consulting … marginal cost and marginal revenue.
>
> (Friedman 1953: 15)

> 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)

> The evidence cited to support this assertion is generally taken either from the answers given by businessmen to questions about the factors affecting their decisions – a procedure for testing economic theories that is about on a par with testing theories of longevity by asking octogenarians how they account for their long life.
>
> (Friedman 1953: 31)

In one sense, Friedman's critique is reasonable: what businessmen *say* they do and what the market forces them to do may be very different.[14] Just as asked octogenarians to account for their longevity will give spurious reasons, but nonetheless reliable data on, amongst other things, whether they drink a bottle of scotch a day, asking businessmen about their businesses yields important data – *including how many face rising marginal cost*. These data, to cite Alan Blinder, yield overwhelmingly bad news for economic theory in that apparently only 11 percent of GDP is produced under conditions of rising marginal cost (Blinder *et al.* 1998: 102).

Blinder's survey was merely the last in a long line of empirical work that contradicted an essential *structural* assumption in the traditional model: if firms do not *in fact* face rising marginal cost, then the model of perfect competition can't function. The best survey of this long, ignored tradition of work is in Lee (1998), but Blinder's survey is the most recent, and has impeccable professional standing and empirical methods.[15] The key empirical findings are summarized in Table 9.3, and it describes a world that is very different than the traditional model.

Table 9.3 Blinder's summary of his empirical results

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 price all at once	74
Change prices in small steps	16
Have nontrivial costs of adjusting prices of which related primarily to:	
The frequency of price changes	14
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	
Percentage of firms for which marginal costs are:	
Increasing	11
Constant	48
Decreasing	41

Source: Blinder *et al.* (1998, p. 106).

The first is that most firms are price-setters, and prices are normally set for extended periods (and are not generally subject to discounting); price-taking is an interesting phenomenon in agricultural and some commodity markets, but not the modus operandi. Second, sales to other businesses dominate demand: commodities are produced, as Sraffa (1960) emphasized, by commodities in an input-output system, and the net product sold to final consumers and government is only a fraction of total product (a phenomenon which helps explain the overall stability and infrequency of changes to prices). Third, sales are not anonymous but generally repeat sales to a network of customers – on average, 85 percent of sales are to existing customers.

Finally, if there is any law of product in the real world, it is the law of constant marginal product, in contrast to the law of diminishing marginal product of

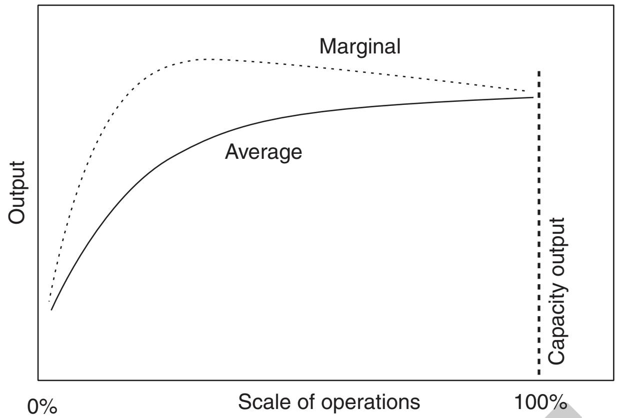

Figure 9.8 Eiteman's representation of marginal product.

orthodox fantasy (Varian 2006: 329). The fantasy arises from the parable of one factor being held constant by a producer while the other is varied.[16] The reality results from a business being designed by engineers

> so as to cause the variable factor to be used most efficiently when the plant is operated close to capacity. Under such conditions an average variable cost curve declines steadily until the point of capacity output is reached. A marginal curve derived from such an average cost curve lies below the average curve at all scales of operation short of peak production.
>
> (Eiteman 1947: 913)

This now almost ancient literature on firm costs that Friedman actively dissuaded economists from considering remains the best basis on which to teach the actual cost structure and decision-making processes of firms, and Fred Lee's *Post Keynesian Price Theory* (1998) gives an excellent survey.

Dimensionality, habit, and rationality

The developing field of behavioral economics provides a good foundation for discussing actual consumer behavior and its departure from the traditional vision of rationality, but I prefer to commence any discussion of consumer behavior with the results of an experiment that confirmed the irrationality of the traditional vision of irrationality – Sippel's test of Samuelson's model of revealed preference (Sippel 1997). This experiment was in fact an unsuccessful attempt to apply the theory, but the very careful manner in which it was done, and its failure examined by Sippel, makes it possible to reverse the economic definition of rationality – from an emphasis upon considering all alternatives, to limiting choice in a manner that makes decision-making in finite time possible.

Table 9.4 Goods in Sippel's experiment

Goods	*Description*	*Range*
Videoclips	Watching videoclips with rock and pop music	30–60 min
Computer game	Playing Super Blast (in Exp1) or Pinball (in Exp2)	27.5–60 min
Magazines	Reading a selection of German newspapers and magazines	30–60 min
Coca-Cola	Cold soft drink	400–2,000 g
Orange juice	Cold drink	750–2,000 g
Coffee	Prepared when demanded	600–2,000 g
Haribo	Popular German brand of candy, licorice, etc.	600–2,000 g
Snacks	Pretzels, peanuts, etc.	600–2,000 g

Table 9.5 Violations of axioms of revealed preference in Sippel's experiment

	Consistent subjects (%)	*Inconsistent subjects (%)*	*With … violations*						
			1–2	*3–4*	*5–6*	*7–8*	*9–10*	*11–20*	*>20*
Exp1									
SARP	1 (8.3)	11 (91.7)	7	3	–	–	–	–	1
GARP	7 (58.3)	5 (41.7)	3	1	–	–	–	1	–
Exp2									
SARP	8 (26.7)	22 (73.3)	7	4	–	1	43	3	3
GARP	11 (36.7)	19 (63.3)	8	1	2	3	1	1	3

Sippel's experiment replicated the standard panoply of orthodox consumer choice theory in a set of consumption choices his subjects were given with varying incomes and prices. The subjects were required to make choices between eight different commodities – an apparently restricted range of options – and were actively encouraged to work out the combinations that maximized their utility, by both the subject's participation fee, and the fact that they were required to consume one of their chosen bundles at the end of the experiment.

The results of the experiment constituted a clear refutation of the neoclassical model of consumer behavior. Overall, more than 75 percent of subjects violated SARP – the strong axiom of revealed preference, the formal definition of a utility maximizer, whereby if a consumer prefers bundle A to B and bundle B to C, he/she will never choose C when A is also affordable – and over 50 percent violated the weaker GARP (generalized axiom of revealed preference).

Though Sippel's examination of this experimental contradiction is exemplary, he did not provide an interpretation of why the model failed: Why do consumers, in a well-designed experiment, fail to behave rationally? What is known as the curse of dimensionality provides a simple explanation that is a good starting point for introducing behavioral economics: the consumers were overwhelmed by the range of choice available, even in this simple situation.

Sippel's consumers were allowed to choose any quantity of the nominated goods, but even if we discretize the choices made to just four options for each commodity (so that we group all choices made of video clips into 0, 30, 31–45 and 46–60 minutes), each consumer was confronted with four possible quantities of each of eight commodities – which results in $4^8 = 65{,}536$ different bundles of commodities to compare with each other. The human brain simply isn't designed to store and rank so many options – let alone as many as consumers confront every day when they enter a supermarket, where 10,000+ items are on display, and the number of distinct bundles blows out to inconceivable numbers.[17]

Instead, true rationality in practice is not considering every option, as neoclassical theory emphasizes, but reducing the bewildering complexity of options available to enable decisions to be made in finite time. Here all the behaviors that neoclassical theory effectively ignores – culture, convention, habit – become truly rational, because they make decision-making possible.

Elements of a political economy microeconomics

Though many non-traditional schools of thought can contribute to a heterodox microeconomics, my personal preference is to begin with Schumpeter's *Theory of Economic Development* (1934). I present this innovation-focused perspective prior to an exposition of the price formation theories of post-Keynesian or Sraffian economists because there is a danger that a desire to provide a replacement for every aspect of the traditional panoply still lets it set the agenda: orthodoxy has a theory of price formation, political economists need one too. Non-traditional theories still over-emphasize price formation of homogeneous products in isolated markets. Yet a political economy theory of price formation that abstracts from product diversity and innovation may be as flawed as the neoclassical one it attempts to replace.

While Schumpeter focuses on explaining cycles, his explanation of why and how firms compete remains ground-breaking, because it ascribes an evolutionary and far from equilibrium perspective upon firm behavior. Schumpeter's analysis of competition begins, not with the process of price formation, but with the means by which a firm distinguishes itself from competitors, and make a profit via innovation:

> (1) The introduction of a new good (2) The introduction of a new method of production (3) The opening of a new market (4) The conquest of a new source of supply of raw materials or half-manufactured goods (5) The carrying out of the new organization of any industry.
>
> (Schumpeter 1934: 66)

Schumpeter's modern descendant is Porter's *Competitive Advantage of Nations*, which provides an application of Schumpeter's framework to explain why some countries have developed competitive advantages in some industries – Italy in fast cars, for example. There is much material of interest for political economists, but for systemic analysis of microeconomics, his conclusion of two sources of competitive advantage is vital:

> It is difficult, though not impossible, to be both lower-cost and differentiated relative to competitors. … Any successful strategy, however, must pay close attention to both types of advantage while maintaining a clear commitment to superiority on one.
>
> (Porter 1998: 38)

After discussing case studies from Porter, I give an exposition of one of the most accessible and compelling models of evolutionary competition available – Paul Ormerod's model of competition in a newly deregulated industry (Ormerod *et al.* 2002).[18] The model simulates competition in both price and quality, and considers what happens to a once-monopolized industry opened to competition.

Traditional economics measures the degree of competition based on the number of firms and the dispersal of market share. Ormerod instead defines the degree of competition on the basis of the degree to which price *and quality* improve over time. Traditional economics predicts – spuriously, as shown above – that a high degree of market concentration will be correlated with a high price, but Ormerod's model finds no correlation (see Figure 9.9).

An important aspect of Ormerod's model is that, with a spectrum of firms competing on both price and quality, there is no single price for the hypothetical commodity in this simulated market. This fits the real-world phenomenon whereby there is no such thing as the price of, for example, a car. Instead, there are a multitude of products that all fit the generic classification of a car, but which have vastly differing qualitative features and widely dispersed prices. This

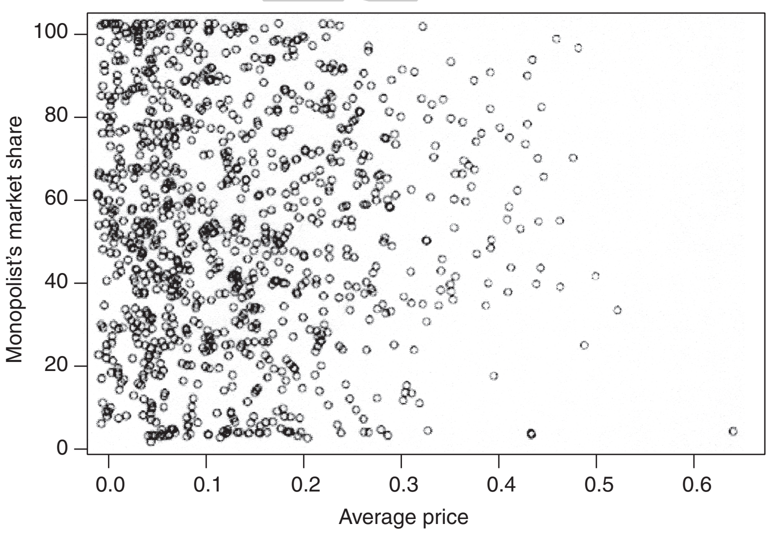

Figure 9.9 No correlation between monopolist's share of output and average price.

was first treated analytically in Farjoun and Machover's (1983), using concepts from statistical physics long before the modern school of econophysics developed. Though its primary objective was to dispute the concept of a uniform rate of profit, this now freely available work underscores that the emphasis on a price of a homogenous commodity is misplaced.

That, nonetheless, is the primary manner in which Kalecki's degree of monopoly price-setting model developed, as well as Sraffa's model of price-setting in an input-output framework. Both deserve coverage in a political economy microeconomics course, and there is an extensive literature on both. As well as using two of Kalecki's original expositions (Kalecki 1940, 1942) and Lee's masterful survey (Lee 1998), I juxtapose a modern exposition of Kalecki (Kriesler 1988), a Sraffian critique of markup pricing (Steedman 1998), and my argument, from a dynamic modeling perspective, that Kaleckian markup pricing and Sraffian input-output pricing are compatible in a dynamic, non-equilibrium context (Keen 1998).

This achieves three results: it covers the major political economy theories of price formation; introduces dynamics; and shows students that economics remains open to debate. Orthodox pedagogy has done an enormous disservice to education by pretending that economics is a done deal, with no outstanding areas of disagreement. I find that students respond positively to the realization that economics is a contested discipline.

Game theory

One manifestation of the dumbing down of economics tuition is that conventional intermediate microeconomics texts treat Cournot–Nash game theoretic analysis as an advanced topic, but only at an introductory level (Varian 2006: chs. 27–29). Nonetheless, game theory is the refuge of choice when orthodox economists are confronted with my critique of Marshallian economics – somehow it seems excusable to teach bad mathematics if it reaches the same result as a sound but more complicated analysis. Game theory will also be confronted by students who move on to postgraduate economics courses – so its applicability as an analysis of microeconomic behavior must be considered in a course on heterodox microeconomics.

I commence by acknowledging that Cournot–Nash analysis does not rely upon the fallacy that afflicts Marshallian analysis: if firms behave strategically as outlined by Cournot, then as the number of firms in an industry increases, price will converge to marginal cost – not because firms are profit-maximizing, as Marshallian analysis erroneously asserts, but because strategic interactions with other firms force them to produce a greater than profit-maximizing quantity.

However there are at least two major problems with Cournot–Nash analysis – one of which is well known. Although the defect strategy is a Nash equilibrium in a single-shot Prisoners' Dilemma game, the cooperate strategy is dominant in repeated games – and real-world competition. If the analysis is to be modeled as a strategic game, it is clearly better modeled by repeated games than a single shot. Cournot–Nash analysis, though a fertile ground for academic papers, is thus an unsatisfactory pedagogic device from the traditional perspective; this probably explains its reluctance to abandon the tired (and fallacious) Marshallian argument.

The second major flaw in the game theoretic analysis of competition is that, while the defect strategy is a global Nash equilibrium *if* firms have perfect knowledge of each other's possible strategies, it can be shown to be locally unstable; and while the cooperate strategy is not a Nash equilibrium, it is locally stable if firms lack perfect knowledge of their rivals' strategies. I illustrate this using the same numerical example as above for a duopoly.

The first step in the illustration shows that the defect strategy actually amounts to firms following the orthodox quantity-setting rule of equating marginal cost to marginal revenue, while the cooperate strategy results from firms being profit-maximizers: orthodoxy here is, in effect, criticizing firms for being self-interested! Using the algebraic example on p. 000, each firm in a duopoly will produce b $a - c\backslash 4 \cdot b + 2 \cdot d$ b units of output if both firms profit-maximize, and b $a - c\backslash 3 \cdot b + 2 \cdot d$ b if both set marginal revenue equal to marginal cost. These output levels provide the diagonal elements in the Prisoners' Dilemma table; the off-diagonal amounts occur when one firm follows an equilibrium profit-maximizing strategy while the other sets $MR = MC$: the firm following the neoclassical strategy will produce a larger amount while the output of the other firm will fall. The full pattern of outputs is shown in Table 9.6.

The profit numbers derived from this table in Table 9.7 ostensibly tell a convincing argument in favour of the Cournot strategy as a Nash equilibrium. Taking the Keen strategy pair as a reference point, each firm will gain (and the other lose) if it changes to the Cournot strategy *while the other firm maintains the Keen strategy*. Therefore, the Cournot strategy pair is a Nash equilibrium.

But what if one firm doesn't know what the other firm's costs or strategies might be, and its strategy instead amounts to varying output to see what happens? And what if the firm reacts to the impact of a change in the other firm's strategy on its profits? Then a very different picture emerges – because for a pair of strategies to be maintained, it must be true that both firms benefit from that pair.

Table 9.8 shows what happens if both firms are at the Keen equilibrium, and each firm experiments with changing its output by +/–1 unit. For a pair of

Table 9.6 Quantity strategy combinations for duopoly

Quantities Firm	*Firm strategies*	*Firm outputs*	*1 Cournot*	*Firm outputs*	*1 Keen*
2	Cournot	1	$\frac{a-c}{3 \cdot b - 2 \cdot d}$	1	$\frac{a \cdot b + 2 \cdot a \cdot d - b \cdot c - 2 \cdot c \cdot d}{5 \cdot b^2 + 10 \cdot b \cdot d + 4 \cdot d^2}$
2	Cournot	2	$\frac{a-c}{3 \cdot b - 2 \cdot d}$	1	$\frac{2 \cdot (a \cdot (b+d) - c \cdot (b-d))}{5 \cdot b^2 + 10 \cdot b \cdot d + 4 \cdot d^2}$
3	Keen	1	$\frac{2 \cdot (a \cdot (b-d) - c \cdot (b-d))}{5 \cdot b^2 + 10 \cdot b \cdot d + 4 \cdot d^2}$	1	$\frac{a-c}{4 \cdot b - 2 \cdot d}$
2	Keen	2	$\frac{a \cdot b + 2 \cdot a \cdot d - b \cdot c - 2 \cdot c \cdot d}{5 \cdot b^2 + 10 \cdot b \cdot d + 4 \cdot d^2}$	2	$\frac{a-c}{4 \cdot b - 2 \cdot d}$

Table 9.7 Profit outcomes from quantity strategies

Relative profits Firm	*Firm strategies*	*Firm outputs*	*1 Cournot*	*Firm outputs*	*1 Keen*
2	Cournot	1	$-5.7 \cdot 10^{10}$	1	$-1.8 \cdot 10^{11}$
2	Cournot	2	$-5.7 \cdot 10^{10}$	1	$-1.3 \cdot 10^{11}$
3	Keen	1	$-1.3 \cdot 10^{11}$	0	
2	Keen	2	$-1.8 \cdot 10^{11}$	0	

Table 9.8 Profit outcomes from varying output at Keen equilibrium

Output changes		*Profit results for Firm 1*		
Firm 2 ⇓	⇒ Firm 1	–1	0	–1
–1		218.7	$-1.1 \cdot 10^{-7}$	–218.8
0		218.7	0	–218.7
–1		218.7	$-1.1 \cdot 10^{-7}$	–218.8
		Profits results for Firm 2		
Firm 2 ⇓	⇒ Firm 1	–1	0	–1
–1		218.7	218.7	218.7
0		$-1.1 \cdot 10^{-7}$	0	$-1.1 \cdot 10^{-7}$
–1		–218.7	–218.7	–218.8

strategies to be maintained, a positive change in profit must occur in the corresponding cells for *both* firms – otherwise the firm that saw its profit fall because of the strategy pair will change its strategy. It shows that there is no pair of strategies which is self-reinforcing: if Firm 1 chooses a strategy that initially will cause its profit to rise, Firm 2 will adopt a strategy that turns Firm 1's originally successful strategy into a losing one. Therefore the Keen equilibrium is locally *meta-stable* – any deviation from it by one firm will cause responses by the other firm that push its competitor back to the Keen equilibrium.

Table 9.9 shows the corresponding situation for the Cournot equilibrium. Here there *are* stable strategy pairs: if both firms reduce their output by one unit, then both will gain profits. A second try of the strategy by both firms (reduce output by two units from the Cournot level) will add even more profit, and so on. The Cournot equilibrium is thus locally *meta-unstable* in the direction of reductions in output by both firms.

Thus profit-maximizing behavior will destabilize the Cournot equilibrium while reinforcing the Keen one. The game theoretic defence of the proposition that competition will force firms to produce where marginal revenue equals marginal cost is thus at best fragile. Even with traditional assumptions on the nature of costs and demand, profit-maximizing firms will tend to produce where their marginal revenue greatly exceeds marginal cost, and market price will exceed marginal cost.

Table 9.9 Profit outcomes from varying output at Cournot equilibrium

Output changes		*Profit results for Firm 1*		
Firm 2 ⇓	⇒ Firm 1	–1	0	–1
–1		218.7	$-1.1 \cdot 10^{-7}$	–218.8
0		218.7	0	–218.7
–1		218.7	$-1.1 \cdot 10^{-7}$	–218.8
		Profits results for Firm 2		
Firm 2 ⇓	⇒ Firm 1	–1	0	–1
–1		218.7	218.7	218.7
0		$-1.1 \cdot 10^{-7}$	0	$-1.1 \cdot 10^{-7}$
–1		–218.7	–218.7	–218.8

Conclusion

This is a demanding course, but one I find students respond well to.[19] At its end, students have a very deep and critical understanding of neoclassical microeconomics, and the beginning of a vision of what a heterodox alternative might be. Hopefully, some of them will be encouraged to help us develop such a richly needed alternative.

Appendix A: market simulation program

This program is written in Mathcad and uses arrays rather than agents, but has the same effect. It could easily be implemented in any number of programming

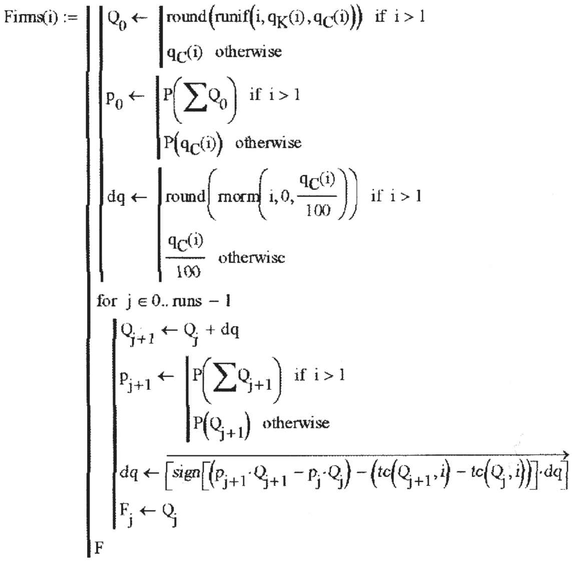

```
Firms(i) := | Q_0 ← | round(runif(i, q_K(i), q_C(i)))  if i > 1
            |       | q_C(i)  otherwise
            | P_0 ← | P(ΣQ_0)  if i > 1
            |       | P(q_C(i))  otherwise
            | dq ← | round(rnorm(i, 0, q_C(i)/100))  if i > 1
            |      | q_C(i)/100  otherwise
            | for j ∈ 0.. runs − 1
            |   | Q_{j+1} ← Q_j + dq
            |   | P_{j+1} ← | P(ΣQ_{j+1})  if i > 1
            |   |           | P(Q_{j+1})  otherwise
            |   | dq ← →[sign[(P_{j+1}·Q_{j+1} − P_j·Q_j) − (tc(Q_{j+1}, i) − tc(Q_j, i))]·dq]
            |   | F_j ← Q_j
            | F
```

Figure 9A.1 Mathcad implementation of a multi-agent simulation.

environments – from NetLogo to C# – but I prefer Mathcad because its code is so compact and readable.

Line by line:

1. The program assigns each of the *i* firms in the simulation a starting amount which is uniformly randomly distributed between my equilibrium prediction $q_K = \frac{a-c}{n \cdot (2 \cdot b + d)}$ and the neoclassical prediction $q_C = \frac{a-c}{bn \cdot (n+1) - n \cdot d}$.[20] This is a vector operation, so for the 10,000 firm simulation, 10,000 different initial amounts are set; the subscript 0 in Q_0 refers to the first time step in the 1000 iterations of the model.
2. This covers the case when a monopoly is simulated.
3. The initial market price is determined from the sum of the initial outputs of all firms – this is a classic price-taker simulation where producers simply determine an output level and then receive the price set by the market demand curve.
4. This covers the case of a monopoly.
5. Each firm is assigned an amount *dq* to change its output by each time step, following a Normal distribution with a mean of zero and a standard deviation equivalent to 1 percent of the neoclassical prediction of the equilibrium output for a firm.
6. This covers the case of a monopoly.
7. A loop is set up to run over *runs* instances (in the simulations shown here *runs* = 1000).
8. A new output level for each firm is set by adding its change amount *dq* to its initial output level Q_0 (this is a vector operation, so with 10,000 firms, 10,000 different *dq*s are added to 10,000 different initial amounts Q_0).
9. A new price P_{j+1} is established based on the new aggregate industry output level ΣQ_{j+1}.
10. This covers the case of a monopoly.
11. Each firm calculates the change in profit between the i^{th} and $i + 1^{th}$ iteration. If profit rose, then the firm changes output in the same direction; if profit fell, then the firm reverses direction.
12. The array *F* then stores the results of the output of each firm at the j^{th} iteration.
13. The program returns *F*, which is a matrix where the rows contain the output of each firm and each column is a time step in the simulation.

Appendix B: true profit-maximizing behavior

If market demand and the cost function of the firm can be expressed mathematically, then the output level that maximizes the firm's profits π_i can be objectively defined. *Whether or not a given market structure – or a given type of strategic interaction between firms – actually results in the profit-maximizing level being the equilibrium level is irrelevant to the question of what the profit-maximizing level actually is*.

Orthodox pedagogy asserts that this maximum is given by the quantity at which the firm's marginal revenue equals its marginal cost:

$$\pi_{iMax}(Marshall) : MR_i(q_i) = MC_i(q_i) \tag{9.11}$$

Given the definition of marginal revenue and the substitution that $\frac{d}{dq_i}P = \frac{d}{dQ}P$, this expands to:

$$P(Q) + q_i \cdot \frac{d}{dQ}P = MC_i(q_i) \tag{9.12}$$

However, the profit-maximizing output level for the i^{th} firm is a function not merely of its output, but also of the output of all other firms in the industry – regardless of whether or not the i^{th} firm can influence their behavior, or knows what that behavior is. The true profit maximum is therefore given by the zero, not of the *partial* differential of the i^{th} firm's profits π_i with respect to its output q_i, but by the *total differential* of its profits with respect to industry output Q: not by the value of q_i for which $\frac{\partial}{\partial q_i}(\pi_i) = 0$ – which economists normally erroneously write as $\frac{d}{dq_i}(\pi_i) = 0$ – but by the value of Q for which $\frac{d}{dQ}(\pi_i) = 0$. Though the individual competitive firm can't ensure that the market produces this amount, it can work out what its own output level should be, given a specified market inverse demand function $P(Q)$ and firm cost function $TC_i(q_i)$. We start by expanding $\frac{d}{dQ}(\pi_i) = 0$ in terms of P, Q, q_i and TC_i:

$$\frac{d}{dQ}\pi_i = \frac{d}{dQ}(P(Q)q_i - TC_i(q_i)) = 0 \tag{9.13}$$

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

$$\frac{d}{dQ}(P(Q)q_i - TC_i(q_i)) = \sum_{j=1}^{n}\left\{\left(\frac{\partial}{\partial q_j}(P(Q) \cdot q_i - TC_i(q_i))\right) \cdot \frac{d}{dQ}q_i\right\} \tag{9.14}$$

In the Marshallian case, atomism lets us set $\frac{d}{dQ}q_j = 1 \ \forall\ j$. Expanding the *RHS* of (9.15) yields:

$$\sum_{j=1}^{n}\left(P(Q) \cdot \frac{\partial}{\partial q_j}q_i + q_i \cdot \frac{\partial}{\partial q_j}P(Q) - \frac{\partial}{\partial q_j}TC_i(q_i)\right) = 0 \tag{9.15}$$

Under the Marshallian assumption of atomism, the first term in the summation in (9.16), $P(Q) \cdot \frac{\partial}{\partial q_j}q_i$, is zero where $j \neq i$, and $P(Q)$ where $j = i$. The second term is equal to $q_i \cdot \frac{\partial}{\partial q_j}P(Q)\forall j$, and $\frac{\partial}{\partial q_j}P(Q) = \frac{d}{dQ}P$, so that this yields n copies of $q_i \cdot \frac{d}{dQ}P$; the third term $\frac{\partial}{\partial q_j}TC_i(q_i)$ is zero where $j \neq i$, and equal to marginal cost $MC_i(q_i)$ where $j = i$. Equation (9.16) thus reduces to

$$P(Q) + n \cdot q_i \cdot \frac{d}{dQ}P = MC_i(q_i) \tag{9.16}$$

This is the true profit-maximization formula, and it coincides with the neoclassical formula only in the case of a monopoly, when $n = 1$. It is easily shown that the rule in (9.17), which I call the Keen formula, results in a substantially higher profit than the standard Marshallian formula. This formula, which is more accurate for individual firms than the representative firm derivation given in the body of this chapter, may explain the variation in individual firm behavior displayed in the multi-agent simulation.

Appendix C: differing cost structures

The standard orthodox diagram comparing a monopoly to a competitive industry blithely assumes that exactly the same line can be drawn to represent the marginal cost curve for the monopoly and the aggregate marginal cost curve (a.k.a. the supply curve) for the competitive industry. In fact, the identity of these two cost curves can be shown to occur only when either (a) marginal costs are identical and constant or (b) a quirk applies so that the sum of the marginal cost curves of the competitive firms happen to overlap with the marginal cost curve for the monopoly (this requires that the number of firms in the industry is an argument into the marginal cost function of the firm, which is of course bizarre).

If we treat labor as the variable input, then marginal cost is the wage rate w times $\frac{dL_i}{dq_i}$, where L_i is the labor input of the i^{th} firm. $\frac{dL_i}{dq_i}$ is the inverse of marginal product for the firm $\frac{dL_i}{dq_i}$, so the identity of marginal costs for many competitive firms and a monopoly also requires the identity of marginal products – and this lets us transfer the problem from the realm of costs to output. For marginal products to be identical, total products can only differ by a constant (since marginal product is the derivative of total product). If output with zero labor input is zero, then this constant of integration is also zero. So the identity of marginal cost functions requires that the output of the monopoly with its labor input is identical to the output of the competitive industry with its labor input at all scales of output. If we consider an n-firm competitive industry where each firm employs x workers and has a production function $q_i = f(x)$, and a monopoly with m plants each employing y workers with production function $q_j = g(y)$, then for the respective marginal cost curves to be identical, the following condition must apply:

$$n \cdot f(x) = m \cdot g(y)$$

$$\text{where } y = \frac{n \cdot x}{m} \qquad (9.17)$$

It is then easily shown using a straightforward application of Euler's theorem that this is only possible if marginal costs are identical and constant (Keen 2004: 121). Another happenstance possibility – that the number of firms in an industry is an argument in an individual firm's marginal cost function, so that in the aggregate marginal costs are identical and rising for all industry scales – was used in the multi-agent program above. Thus while marginal cost at the industry was as defined above, marginal cost for an individual firm in an n-firm industry was defined as $MC_i(q_i) = c + d \cdot n \cdot q$. Aggregation then ensures that $MC(Q)$ is independent of the number of firms in the industry. Of course, in real life, it is highly likely that an industry's marginal costs will be lower when it is dominated by large (non-competitive) firms than when it is dominated by smaller competitive ones. Rosput gives an excellent example of this with respect to the natural gas industry (Rosput 1993).

Notes

1 The identical fixed cost figure for vastly differing industry sizes can be replaced by one that depends on the scale of output, with no effect on the simulation results.
2 The program is shown in Appendix A.
3 In contrast to the Cournot–Nash game theoretic model, which I discuss later.
4 Even in the minds of political economists!
5 A fourth offered by Stigler himself in his paper (Stigler 1957: 8) is that marginal revenue for the i^{th} firm converges to market price as the number of firms in the industry increases. This is true, and understood by a minority of neoclassical economists – but it is also irrelevant, as shown below (p. 000).
6 A fifth defense, that if price exceeds marginal cost, other firms will enter from other industries, is also fallacious since all other industries will likewise have price greater than marginal cost, if the firms in them are profit maximizers.
7 Using a modified version of this program, it can be shown that the neoclassical result applies if about 25–50 percent of firms behave *irrationally* – by increasing output when this decreases profits, and vice versa (Keen and Standish 2005).
8 Although this is difficult, I find that intermediate micro students can understand it. Orthodox economists might object to this proof, given that it requires firms to know the aggregate industry output. More advanced proofs without this ostensible flaw are given in Keen and Standish (2006, 2008). One is shown in Appendix B.
9 The marginal cost of aggregate industry output Q is the same as the marginal cost for the i^{th} firm in that industry of producing q_i.
10 Appendix C shows that this is the exception rather than the rule.
11 Strictly speaking, the market inverse demand curve $D(P)$ can have any shape that can be described by any polynomial equation. Orthodox economists prefer that, like an individual Hicks-compensated demand curve, it necessarily slopes down – implying restrictions on the coefficients of the polynomial so that $D(P + \Delta P) < D(P)$.
12 See Keen (2001) for a non-mathematical overview of these flaws.
13 The number of firms of a given size (measured here in terms of employees per firm) is a function of the number of firms raised to a negative power. See also Axtell (2001).
14 As I have demonstrated above, neoclassical theory's predictions are erroneous, even when the (simulated) empirical data conform to the theory's assumptions.
15 Downward and Lee (2001) review this work from a post-Keynesian perspective.
16 For why this is a parable, see Keen 2001: ch. 6.
17 Even if we classify commodities into 100 bundles and discretize decisions to either buy or not buy a single item of each, 2^{100} = 1,267,650,600,228,229,401,496,703,205,376 distinct bundles result! A neoclassical shopper would need a brain the size of a galaxy to store its preference map.
18 Many other attempts have been made to effectualize Schumpeter's vision – see for example, Andersen (www.business.aau.dk/evolution/esa/) – but I find this model the most tractable and easiest to reproduce.
19 This can also be taught with additional content on macroeconomics and finance, under the guise of managerial economics.
20 See Appendix C for an explanation of these predictions.

References

Axtell, R.L. (2001) "Size Distribution of U.S. Firm Sizes," *Science*, 293: 1818–1820.

Blatt, J.M. (1983) *Dynamic Economic Systems*, Armonk, NY: M.E. Sharpe.

Blinder, A.S., Canetti, E., Lebow, D., and Rudd, J. (1998) *Asking About Prices: A New Approach to Understanding Price Stickiness*, New York: Russell Sage Foundation.

Downward, P. and Lee, F. (2001) "Post Keynesian Pricing Theory Reconfirmed? A critical Review of Asking About Prices," *Journal of Post Keynesian Economics*, 23: 465–483.

Eiteman, W.J. (1947) "Factors Determining the Location of the Least Cost Point," *American Economic Review*, 37: 910–918.

Farjoun, E. and Machover, M. (1983) "Laws of Chaos: A Probabilistic Approach to Political Economy. Verso: Farjoun and Machover," http://iwright.googlepages.com/lawsofchaos. See also www.probabilisticpoliticaleconomy.net/.

Friedman, M. (1953) *Essays in Positive Economics*, Chicago: University of Chicago Press.

Gorman, W.M. (1953). "Community Preference Fields," *Econometrica*, 21: 63–80.

Kalecki, M. (1940) "The Supply Curve of an Industry Under Imperfect Competition," *Review of Economic Studies*, 7: 91–112.

—— (1942) "Mr Whitman on the Concept of Degree of Monopoly – A Comment," E*conomic Journal*, 52: 121–127.

Keen, S. (1998) "Answers (and Questions) for Sraffians (and Kaleckians)," *Review of Political Economy*, 10: 73–87.

—— (2001) *Debunking Economics: The Naked Emperor of the Social Sciences*, London: Zed Books.

—— (2004) "Deregulator: Judgment Day for Microeconomics," *Utilities Policy*, 12: 109–125.——

Keen, S. and Standish, R. (2005) "Irrationality in the Neoclassical Definition of Rationality," *American Journal of Applied Sciences* (special issue): 61–68.

—— (2006) "Profit Maximization, Industry Structure, and Competition: A Critique of Neoclassical Theory," *Physica A*, 370: 81–85.

—— (2008 forthcoming) "Debunking the Theory of the Firm – A Chronology," *American Review of Political Economy.*

Kirman, A. (1989) "The Intrinsic Limits of Modern Economic Theory: The Emperor Has No Clothes," *Economic Journal*, 99: s126–s139.

Kriesler, P. (1988) "Kalecki's Pricing Theory Revisited," *Journal of Post Keynesian Economics*, 11: 108–130.

Lee, F. (1998) *Post Keynesian Price Theory*, Cambridge: Cambridge University Press.

Mankiw, G. (2001) *Principles of Microeconomics*, Orlando, Fla.: Harcourt College Publishers.

Marshall, A. (1920) *Principles of Economics* (8th edn) London: Macmillan.

Ormerod, P., Rosewell, B., and Smith, L. (2002) "An Agent-Based Model of the Evolution of Market Structure and Competition," www.paulormerod.com/documents/paper23Sept2002.pdf.

Porter, M.E. (1998) *The Competitive Advantage of Nations*, New York: Free Press.

Rosput, P.G. (1993) "The Limits to Deregulation of Entry and Expansion of the US Gas Pipeline Industry," *Utilities Policy*, 3: 287–294.

Schumpeter, J.A. (1934) *Theory of Economic Development*, New Brunswick, NJ: Transaction Publishers.

Shafer, W. and Sonnenschein, H. (1982) "Market Demand and Excess Demand Functions," in Arrow, K.J., and Intriligator, M.D. (eds.), *Handbook of Mathematical Economics* Vol. 2, Amsterdam: North Holland Press.

Sippel, R. (1997) "Experiment on the Pure Theory of Consumer's Behaviour," *Economic Journal*, 107: 1431–1444.

Sraffa, P. (1960) *The Production of Commodities by Means of Commodities*, Cambridge: Cambridge University Press.

Steedman, I. (1992) "Questions for Kaleckians," *Review of Political Economy*, 4: 125–151.

Stigler, G. (1957) "Perfect Competition, Historically Contemplated," *Journal of Political Economy*, 65: 1–17.

U.S. Census Bureau (2008) www.census.gov/csd/susb/susb.htm.

U.S. Small Business Administration Office of Advocacy (2008) www.sba.gov/advo/research/data.html.

Varian, H. R. (2006) *Intermediate Microeconomics: A Modern Approach* (7th edn) New York: W.W. Norton.