Toward an Understanding of Keen and Standish's Theory of the Firm: A Comment

David Rosnick
Center for Economic and Policy Research, Washington, DC, USA

Abstract

In a series of papers, Steve Keen and Russell Standish criticize the textbook approaches to firm behavior under conditions of perfect and imperfect competition. These papers misstate the assumptions underlying the models and err in mathematics. The critiques do not follow through on their theoretical arguments, and do not explain what drives the results of their computer simulations. Consequently, their contributions confuse rather than clarify understanding of firm behavior.

1. Introduction: Approaching the work of Keen and Standish

Our launching point for this primer will be the article "Rationality in the Theory of the Firm" [see reference 8] – a paper appearing in this issue of *World Economic Review* which was placed into Open Peer Discussion for comment in January, 2014. However, that paper builds directly upon other papers its authors have written previously on subject of supply theory and will necessarily rely to some extent on those previous works. This review, then, summarizes and critiques that research.

Keen and Standish's papers on the subject of supply [see references 3, 4, 6, 8] pursue three distinct lines of thought. First, they argue that the textbook model of perfect competition[1] is "strictly false" [see ref. 3] in assuming the demand function has "dual [contradictory] properties" [see ref. 4] and thus the model contains a "fundamental flaw" [see ref. 8]. The second thread questions the standard Cournot-Nash oligopoly result as deficient in concluding that firms fail to find the collusive level of output. Third, simulating out an infinitely-repeated Cournot-Nash game, they argue that competitive firms all pursuing the authors' suggested strategy will find the collusive level of industry supply.

Each of these threads contain serious flaws, and this paper will address each thread in turn. However, it should be noted that each thread is in fact distinct. Each thread addresses a different model of competition (perfect, Cournot-Nash, and infinitely-repeated Cournot-Nash) and there is no particular reason to believe that the results of one thread supports another.

2. A conflicting definition

It is not entirely surprising that the critiques of Keen and Standish often seem to conflate these different models of competition. Rather, it appears to stem from non-standard definitions. According to Standish, perfect competitors "are defined as agents with no market power, which I took as being a constraint that firms must produce at marginal cost."[2] Of course, perfectly competitive firms are not constrained to produce at marginal cost; they merely have incentive to produce at that level because they have no power to change the price they will receive for their goods.

Standish continues:

[1] Nebulously and variously described in the papers as "the standard economic theory of competitive markets", "Marshallian" theory or model or analysis, and "Marshallian derivation of the model of 'perfect competition.'" In private correspondence, Keen confirmed that "modern textbook model of perfect competition" was a safe interpretation on the part of any reader when reading "the theory of atomisitic competition" or "Marshall's pure case" or "the standard Marshallian theory of the firm" or "standard Marshallian analysis."
[2] See comment submitted 19 March 2014 at 1:07 AM.

> "We are modelling 'price takers', however, that is the agent has choice over level of production, and the market delivers the clearing price. Price takers are not constrained to having 'no market power' – indeed they will always have some power, albeit diluted by the number of firms in the market place."

This is not what economists generally mean when they say a perfectly competitive firm is a price-taker. Economists generally mean that individual price-takers have no influence over the current price. Consequently, the firms described by Keen and Standish have the same market power as Cournot oligopolists; we should expect a competition among such firms to resemble not a perfectly competitive industry but a Cournot oligopoly.

Their unusual definition helps explain some of the extraordinary claims laid out in Standish and Keen's critique.

3. A relatively generic framework for a theory of firm production

Suppose, as above, that firms face an industry inverse demand curve given by $P(Q^d)$and firm i has total costs $\mathrm{TC}_i(q_i)$. Suppose that for any given period of production, firms accept some common unit price p for their current production which they may or may not know at the time they choose their individual levels of output. Then total revenues of firm i are

$$\mathrm{TR} = pq_i$$

and therefore in the given period firm i receives profits

$$\pi_i = pq_i - \mathrm{TC}_i(q_i) \tag{3.1}$$

Note that we have not yet specified how p is determined. Nor have we made any assumption regarding how the firm selects its level of production. This framework will support all of the discussion which follows.

4. Thread 1: On the logical consistency of textbook perfect competition

The opening paragraph of Standish's most recent work [8] summarizes the authors' position (references adjusted to this document):

> "Keen [2, Ch. 4] pointed out a fundamental aw with the standard Marshallian theory of the firm, whereby the market demand function $P(Q)$ (price of a good given total market production Q) is assumed to be a decreasing function of Q (i.e. $P'(Q) < 0$), yet at the same time, for a large number of firms, each individual firm's production qi has no effect on market price, i.e. $\partial P/\partial q_i = 0$. Yet it is easy to see from elementary calculus, that these two conditions cannot be true simultaneously, as first noted by Stigler [9, footnote 31]."

There are at least three errors here. First, the market (inverse) demand function is not one of "total market production" – that is, quantity supplied. Second, $\partial P/\partial q_i = 0$ (though true) is an incorrect rendering of the price-taking assumption. Third, Stigler's argument does not directly address perfect competition; rather, Stigler is exploring Cournot's argument that imperfect competition appears increasingly perfect as the number of firms grow large.

4.1 Inverse demand is removed from the decision facing perfectly competitive firms

The first error is obvious, but points to further problems with the critique. Inverse demand is a function not of quantity supplied, but quantity demanded. Of course, it is legitimate to evaluate inverse demand at quantity supplied, but interpretation of the result requires care. In textbook models of perfect competition, price determines quantity demanded, so if inverse demand is described by $P(Q^d)$, then $P(Q^s)$ is seen in Figure 1.

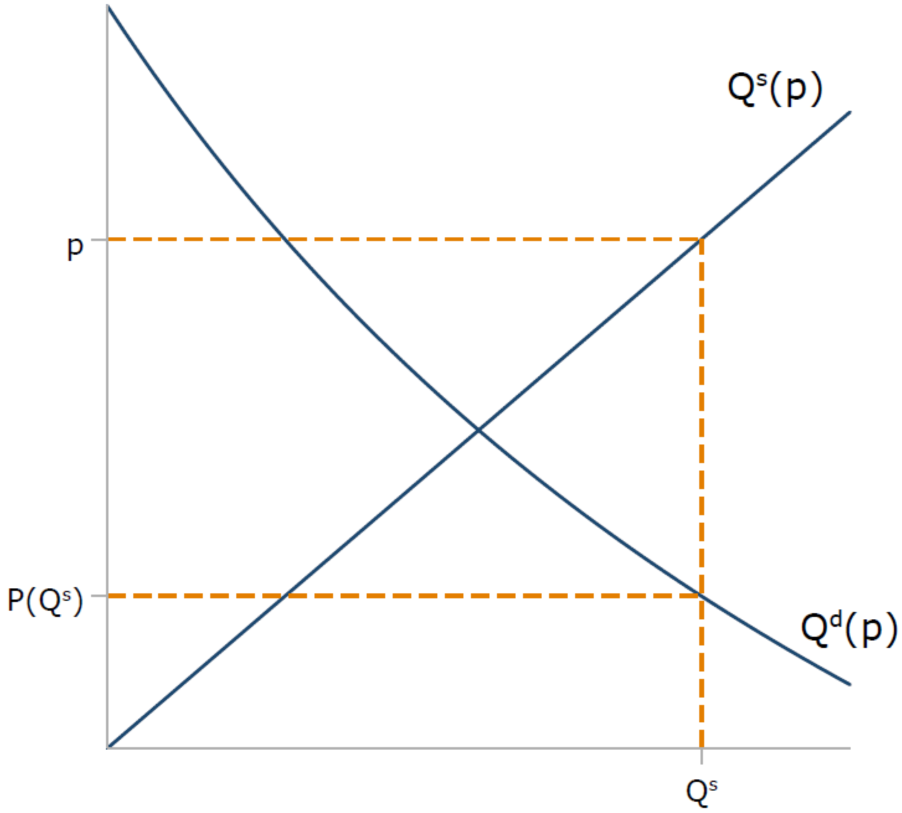

Figure 1: Determination of $P(Q^s)$

Note, as in Figure 1, that $P(Q^s)$ is not in general equal to the market price p except when p happens to be such that the market clears. While $P(Q^s)$ may indeed fall as Q^s rises, this is not the same thing as assuming P to be a decreasing function of Q^s. Rather, $P(Q^s)$ is the price which would have resulted in a quantity Q^s demanded regardless of the actual quantity supplied.

The critique fails to recognize that perfect competition (as opposed to Cournot oligopoly) allows for a market price p such that $Q^d \neq Q^s$ and hence the possibility of a failure to clear the market. Indeed, Standish argues that the two must be identical.[3] Yet the possibility of a non-clearing market price p is critical to textbook analyses of binding minimum wage and rent-control laws; the assumption that market prices always clear the market is not fundamental to perfect competition.

Even if p is assumed to clear the market, the relationship between Q^s and p is not so clear. A positive demand shock increases both quantity supplied and the market-clearing price.

Indeed, in [4] above Figure 9 the authors declare "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" (emphasis in original.) This is plainly false. According to the textbook model of perfect competition, by the time firms begin to make their production decisions the market price is already fixed and so may not fall. Rather, an increase in output results in additional excess supply. It's absolutely rational for the price-taking firm to believe that it will not depress its supply price by increasing output because it operates in a universe where the period's price is already set.

By contrast, this same logic need not apply to models of *imperfect* competition. In Cournot oligopoly, for example, firms do have market power and the market is *assumed* to clear *ex-post* in the

[3] See comment submitted 18 March 2014 at 2:34 AM.

manner described by Standish and Keen. Cournot firms make their output decisions prior to receiving a supply price, and only after the decisions are revealed is the market-clearing supply price $p^C = P(Q^s)$ known to the firm. By insisting that firms have the power to move market prices by varying their individual outputs, we see that the "price takers" of Standish are actually Cournot oligopolists – imperfect competitors, not perfect.

4.2 $\partial P/\partial q_i$ has no obvious connection to price-taking

As to the second error, again note that P is defined as inverse demand - a function of quantity demanded, and not of any firm's quantity supplied. By definition then, $\partial P/\partial q_i = 0$. It is not clear why this conflicts with $P' < 0$.

If, however, when Standish and Keen write "each individual firm's production q_i has no effect on market price" they mean not the market-clearing price but the supply price then this ought to be written either $\partial p/\partial q_i = 0$. Textbook price-taking requires this to be true: each period's p is determined prior to any firm's actual decision. It is not possible for the firm decision to influence directly or indirectly market price p even though p may be determined endogenously via the supply schedule.

Still, Standish's non-standard definition of "price-taking" helps clarify what the authors intent. Let us start by defining a function Q^{s*} so that quantity supplied is given by $Q^s = Q^{s*}(q_1, q_2, \ldots, q_n) \equiv \sum_j q_j$ and that $p^*(q_1, q_2, \ldots, q_n) \equiv P\left(\sum_j q_j\right) = P(Q^s)$. Now, $\partial p^*/\partial q_i = P'(Q^s) < 0$ so long as demand is downward-sloping. This is just as we discussed above – firm choices may alter the *ex-post* market-clearing price.

Likewise, under the non-standard definition, an individual consumer may alter the ex-post market-clearing price. Suppose a consumer may walk into a store where apples are advertised at $1.27 per pound, happy to purchase two pounds at that price. According to the logic laid out in the critique, the consumer may be able to purchase only one pound for only $1.26. This contrasts with the observed practice of discounting bulk purchases. In any case, this describes imperfect and not perfect competition.

It is also worth noting that nothing here depends on "a large number of firms" as the authors suggest. These relations are just as true for a perfectly competitive monopoly. Of course, perfect competition might be a poor choice of model with which to analyze the monopoly. But the results are no less consistent.[4]

4.3 Stigler is correct, but does not undermine price-taking

Finally, we come to Stigler. In his article [9, p. 8], Stigler argues (emphasis added, footnote marker in original):

> It is intuitively plausible that with infinite numbers all monopoly power (and indeterminacy) will vanish, and Edgeworth essentially postulates rather than proves this. But a simple demonstration, in the case of sellers of equal size, would amount only to showing that
>
> $$\text{Marginal revenue} = \text{Price} + \frac{\text{Price}}{\text{Number of sellers} \times \text{Market elasticity}}$$
>
> and that this last term goes to zero as the number of sellers increases indefinitely.[31]
> This was implicitly *Cournot's* argument.

Stigler's "demonstration" reflects the Cournot Theorem – that as the number of imperfect competitors in a Cournot oligopoly becomes increasingly large, the industry behavior increasingly approaches that of perfect competition. Marginal revenues approach market price, and firm production is such that marginal costs approach market price.

[4] A monopoly need not have any pricing power even if textbook models of monopoly generally assume they do. Perhaps the government fixes the price on behalf of the firm, yet allowing it to produce any volume of output it chooses. The government may even agree in advance to purchase any and all of the firm's supply while banning all other sales. In other words, a market may be structured so that the firm faces a perfectly flat demand curve and $P'(Q^d) = 0$. Price-taking is not incompatible with monopoly.

If it is true that as the number of competitors increase firms market power tends to zero, then it is reasonable to imagine a competition in which firms have zero market power. We may then as well ask directly how firms behave in the absence of such market power. We may find such perfect competition does a poor job of representing a real-world market, but we are free to make the assumption.

The important point is that Stigler is working with a model of imperfect competition; his precise construction depends upon assumptions in the Cournot oligopoly model. By contrast, market price is determined prior to the decision of any price-taking firm (under the generally-accepted definition) so the perfectly competitive marginal revenue must be exactly equal to market price.

5. Thread 2: On the correctness of the Cournot-Nash result

Suppose firms compete in the fashion of a Cournot oligopoly. That is, they individually produce whatever they choose, then the supply price adjusts to clear the market. The authors state [8]:

> Mathematically, global equilibrium will occur when all partial derivatives $\partial\pi_i/\,\partial q_j$ vanish. However, this situation can never pertain, as $\forall_i \neq j$, $\partial\pi_i/\,\partial q_j = q_i P' < 0$, except possibly for the trivial solution $Q = 0$.

The critique here confuses equilibrium and extremum. Competitive profit maximization does not mean that firms are in equilibrium if they arrive at a global profit maximum. Reaching global profit maximum for all firms is impossible even with collusion as each firm's maximum is achieved when all competition vanishes and it monopolizes the market. Such conditions are mutually incompatible across firms. Rather, competitive equilibrium requires only that each firm is satisfied with its own production decision given the production levels of its competitors. So satisfied, no firm has incentive to change, and the production levels are therefore stable. The authors continue (emphasis added, reference adjusted):

> The key concept of the *rational agent*, or *homo economicus* is that the agent chooses from an array of actions so as to maximise some utility function. In the context of the theory of the firm, the utility functions are given by π_i in eq (5.1), and the choices are the production values chosen by the individual firms.
>
> Intrinsic to the notion of rationality is the property of determinism. Given a single best course of action that maximises utility, the agent must choose that action. Only where two equally good courses of action occur, might the agents behave stochastically. This deterministic behaviour of the agents is the key to understanding the stability of the Keen equilibrium, and the instability of the Cournot equilibrium, *which is the outcome of traditional Marshallian analysis*.

It is perhaps better to write that utility functions describe the choices agents make and that if agents aim to maximize profits, then the modeler's choice of profit function as utility will most accurately describe the choices agents make. This aside, the last sentence underlines the problem of the previous chapter. No known model of perfect competition has ever brought forth the Cournot equilibrium. It may be argued that the traditional analysis of the Cournot oligopoly has limiting behavior which reflects the outcomes of perfect competition, but there is no analysis which starts with perfect competition and arrives at the Cournot equilibrium.

Again, this suggests that Standish's non-standard definition of price-taking had led the authors to confuse textbook models of perfect and imperfect competition.[5] Putting that aside, we clarify. Given *the*

[5] Seen also in the blurring of models in [6] around equation 16, where the authors describe a "Cournot" result by taking the limit as the number of firms goes to infinity. While the industry production they write is indeed the limit of production in a Cournot oligopoly, it is *not* clearly the level of output produced by any number of firms in perfect competition. If the supply price is indeed $p = c$, then

expected actions of its competitors and a single best course of action for the *ith firm* that maximizes π_i, the ith firm must choose that action. One takeaway from the Cournot analysis is that there is a consistent set of outputs which so satisfy all the firms.

Consider the Cournot duopoly. Firm 1 is satisfied with its choice of operating at the Cournot level if its competitor has also chosen to operate at its own Cournot level. Not only that, but so long as firm 1 operates at its Cournot level, it has structured the incentives of its competitor so that firm 2 will naturally choose to operate at its own Cournot level. Firm 1's belief in firm 2's choice becomes a prophecy fulfilled. The fact that $\partial\pi_1/\partial q_2 < 0$ becomes irrelevant to firm 1 at Cournot-Nash because firm 1 believes (sensibly) that from the Cournot level of output, $dq_2 = 0$.

Standish and Keen reject this and "propose the condition that all firm's profits are maximized with respect to total industry output $d\pi_i/dQ = 0$." Now, let us write more carefully,

$$\pi_i(q_i, Q^s) = P(Q^s)q_i - TC(q_i) \tag{5.1}$$

So that

$$\frac{d\pi_i(q_i, Q^s)}{dQ^s} = \frac{dP(Q^s)}{dQ^d}\frac{dQ^d}{dQ^s} + P(Q^s)\frac{dq_i}{dQ^s} - \frac{d\mathrm{TC}(q_i)}{dq_i}\frac{dq_i}{dQ^s}$$

Assuming that post production the price will adjust to always clear the market ($Q^d = Q^s$) and noting that

$$\frac{dQ^s}{dq_i} = \sum_j \frac{\partial Q^s}{\partial q_j}\frac{dq_j}{dq_i} = \sum_j \frac{dq_j}{dq_i}$$

And the authors' condition reads

$$\left(\sum_j \frac{dq_j}{dq_i}\right)\frac{dP(Q^s)}{dQ^d}q_i + P(Q^s) - \mathrm{MC}(q_i) = 0 \tag{5.2}$$

This appears to confuse firm behavior with comparative statics. As a description of firm behavior, it is simply a restatement of the Cournot equilibrium. Because all firms reveal their outputs simultaneously, the rest of industry cannot tailor its output in response to the whim of any single firm. Thus, each firm must necessarily believe that $dQ^s/dq_i = 1$, leading to the standard Cournot result.

This is not the same as saying that the firm must believe dQ^s/dq_i must be 1 in a dynamic setting. As a matter of comparative statics, for example, we might ask how the equilibrium changes in response to a demand shock. For example, if we assume linear (inverse) demand $P(Q^d) = \alpha - \beta Q^d$ and N firms with zero production costs. Then the Cournot-Nash level of production for each firm is $= \alpha/3\beta$. If demand is greater by $d\alpha$, then all firms produce at a level $d\alpha/3\beta$ greater. In the face of a demand shock, then, $dQ^s/dq_i = N$.

In other words, the firm may know that if it has incentive to produce at a higher level, then so do its competitors; this is different from believing that if it then produces at a lower level, then so will its competitors. Continuing, (edited for clarity, reference adjusted):

> When setting up a game, it is important to circumscribe what information the agents have access to. Clearly, if the agents know what the total market production Q^s will be in the next cycle, as well as their marginal cost MC, the rational value of q_i can be found by setting the partial derivative of π_i to zero:

---

firm profits are exactly zero no matter how much the firm produces. *In the limit* constant marginal costs imply indeterminate supplies. This is why perfectly competitive modeling usually require *increasing* marginal costs, rather than constant.

$$\frac{\partial \pi_i\,(q_1, q_2, \dots, q_n)}{\partial q_i} = P + q_i\,\frac{\partial \mathrm{P}}{\partial q_i} - \mathrm{MC}(q_i\,) = 0 \tag{5.3}$$

> Indeed, the Marshallian theory further assumes that in the limit as the number of firms n tends to infinity, the term $\partial P/\,\partial q_i \to 0$, to arrive at the ultimate result that price will tend to the marginal cost (assuming a unique marginal cost exists over all firms) [5, p. 322]. This assumption is strictly false, as shown by [9]. Instead, $\partial P/\,\partial q_i = dP/dQ$, which is independent of the number firms in the economy.

Once again, this discussion leaves much unpacking do be done. Of course agents need not "know" what total production will be. The firm will maximize based on its belief respecting the choices of all other firms. Taking into account the expected production coming from the rest of industry, the firm may infer the price it will receive as a function of its own decision. The formula should read

$$\frac{\partial \pi_i\,(q_1, q_2, \dots, q_n)}{\partial q_i} = P + q_i\,P'\left(\sum_{j=1}^{n} q_j\right) - MC(q_i\,) = 0 \tag{5.4}$$

Neither "Marshallian theory" nor any known textbook derivation of the Cournot result "assumes" that $\partial P/\,\partial q_i$ approaches zero. Rather, it is zero by definition as seen in Section 4.2 and the reference to Mas-Colell [5, p. 322] does not appear to support the claim. Likewise, Stigler [9, p. 8] nowhere shows this "assumption" or result to be false, but rather outlines proof of the result as we saw in Section 4.3.

Finally, the expression $dP = dQ$ is *not* in general independent of the number of firms in the oligopoly unless the demand curve is linear. As the number of Cournot competitors increases, so does industry production and therefore equilibrium moves down the demand curve.

There is no dispute as to whether or not profits *would be higher* at, say, the collusive result. Objectively, profits would be higher at that level than Cournot-Nash, and firms would be better off producing at that level. The neoclassical argument is that collusion is rational, but firms *competing* for the greatest profits will not forgo opportunities to increase their individual profits and so will over-produce (relative to the collusive level) even if that would result in lower profits for the industry on the whole. That is, competition hurts profits.

This is the orthodox argument for anti-trust actions; it is in the public interest to make sure firms compete rather than collude. Standish and Keen's generalization to asymmetry cannot address this deficiency in their analysis.

While the authors' critique of Cournot-Nash is up to this point is awed, they proceed to argue that the assumptions underlying the model are unrealistic, and begin to offer their own alternative model for competition. While Standish and Keen are free to offer an alternative model of production – perhaps one they consider more realistic – this can in no way prove illogical any textbook model. Nor can it prove "the instability of the Cournot equilibrium."

6. Thread 3: An alternative model of firm behaviour

Imagine you are trying to climb a smooth hill with neither ruts nor stones to stumble over. The peak is in front of you and there is a valley behind. You take a step forward and for some reason not immediately obvious, your GPS records that your elevation fell. So you take a step back, and even more confusing, you find that your elevation fell again. You keep changing course: forward, back, forward, back and all the while, you race toward the valley despite your blindingly sensible hill-climbing technique.

Why are you failing? Did you forget to look around and note that you are actually dancing back and forth in a wagon racing downhill with no brakes? You are failing because you falsely believe you are in control of whether you go up or down the hill when reality puts you at the mercy of your environment.

Such is the case with the firms of Standish and Keen. They literally misunderstand the consequences of their own decisions, and are too "rational" to question their approach. Standish and Keen introduce their alternative oligopoly model as follows (reference adjusted):

> [I]t is completely unrealistic for the firms to be able to predict market production (and hence price). Firms cannot know whether their competitors will act completely rationally, and details such as the marginal cost curve for each firm, and even the total number of players is unlikely to be known. So equation (5.3) cannot be correct. Instead, firms can really only assume that the price tomorrow will most likely be similar today, and that the best they can do is incrementally adjust their output to "grope for" the optimal production value. So in our model, firms have a choice between increasing production or decreasing it. If the previous round's production change caused a rise in profits, the rational thing to do is to repeat the action. If, on the other hand, it leads to a decrease in profit, the opposite action should be taken. At equilibrium, one would expect the production to be continuously increased and decreased in a cycle with no net movement.

The discussion starts out just fine. Standish and Keen argue that Cournot assumptions are unrealistic in that firms may not have quite so much information as the model allows. This is a completely valid critique of Cournot oligopoly. It disputes the applicability of the Cournot model – in contrast to the first thread, where the authors dispute the internal consistency of the textbook model of perfect competition. It also stands in contrast to the second thread, where the authors dispute the result of competition under Cournot assumptions. Here, the authors ask firms to play a different game. Just as someone playing backgammon behaves differently than someone playing chess, so firms playing a Standish-Keen game may behave differently than firms in a Cournot oligopoly in turn behaving differently than firms in perfect competition.

Then the authors reasonably – given the change in information given to firms – offer up an alternative approach to firm decision-making. Where the authors start to fall down is when they call their algorithm for output adjustment "rational." The hill-climber dancing in the wagon may impress with such commitment to an erroneous idea of how the world works, but "rational" may not describe the hill-climber very well. As we will see, the behavior of firms of Standish and Keen is better described as erratic rather than reasonable.

6.1 A note on "profit maximization"

It is important to note that Standish and Keen suggest an iterative procedure for reaching equilibrium. In each iteration, firms play a Cournot game. This is significant because the authors outline a function for profit for a single period only. However, as Anglin [1] correctly points out, "proposing a longer horizon would also add other dimensions to the optimization problem that KS do not consider." Among these other dimensions, if the firms are expected to play a game over multiple periods, the objective function ought be defined differently.

If we call $\pi_{i,t}$ the profits for firm i in period t, then what exactly is the firm aiming to maximize? The authors might have chosen any of the following examples – among many others – as legitimate "profits" for firm i to maximize:

$$\pi_{i,t} \qquad \sum_{j=1}^{n} \pi_{i,j} \qquad \sum_{j=j_0}^{j_0+n-1} \pi_{i,j}$$

$$\sum_{j=1}^{n} \rho^{j-1}\, \pi_{i,j} \qquad \pi_{i,1000} \qquad \sum_{j=1}^{\infty} \rho^{j-1}\, \pi_{i,j}$$

Obviously, firm behavior depends on the choice of objective function. In particular, it is a well-established result of game theory that firms playing an infinitely-repeated Cournot-Nash game may produce a rich variety of competitive equilibria.

Of course, the variety of equilibria possible in the infinitely-repeated Cournot-Nash game does not disprove the textbook analysis of the simple Cournot oligopoly. They are different games, and so the different outcomes are not in conflict. Standish and Keen address this concern arguing in [7]:

> This can be interpreted as firms anticipating what their competitors might do, although we tend to regard it as describing reactions to competitors in a "time-free" model, so the variation is not conjectural but reactionary.

In other words, they would have each firm offering some kind of tentative level of output and reacting to the revealed plans of other firms before everyone adjusts output, and so on until the firms reach some sort of equilibrium. Indeed, this is often how Cournot production is initially presented, with the understanding that in Cournot-Nash, each firm privately simulates out this adjustment process to learn its optimal production level. However, this is totally inadequate for Standish-Keen competition as they presume firms cannot know the price they will receive except through the market. In other words, they must receive actual prices in the market, and lock in actual profits as they progress.

Implicitly, Standish-Keen firms discount to zero all profits made in any finite time, leaving $\Pi_i = \pi_{i,\infty}$. But of course, tomorrow never comes. Thus, the authors view equation (5.1) as a description of a particular condition when firms are in equilibrium - not to be confused with an equilibrium condition. However, as we will see, this belies the authors' description of their simulations. Though the firms of Standish and Keen's simulations do not perform calculus to find the solution, they do quite actively solve equation (5.2). That is, it is no mere emergent outcome, but does in fact represent the firms' "behavioral rule."

Absent a clear objective function, though, the trade-offs between current and future profits are not specified. Still, the analyses of Standish and Keen are deficient in other important ways.

6.2 A wagon, not a hill

Though they say firms "grope," the authors imagine firms as hill-climbers. Firms increase or decrease output, rather than move forward or back, and they observe whether their profits – rather than elevation – rose or fell.

In the extreme case of Cournot-Nash equilibrium, all firms have optimized profits with respect to their own output and so each firm's profits are by construction insensitive to small changes in its own output. There, firm profits rise and fall based exclusively upon the choices of its competitors. Yet Standish and Keen would have the firms act as though they believed the opposite – that their own choice of output is the sole cause of any change in their own profits.

If a firm's wagon is climbing, then a step forward will coincide with an increase in elevation. However, so will a step backward. This is seen in Figure 2 where a symmetric 1,000-firm oligopoly as in [8, Sec. 6] starts in Cournot-Nash equilibrium, but all firms decide to cut output. If one firm had chosen instead to expand output, its profits would have been the tiniest bit larger. Both outcomes are shown, along with the iso-profit lines for each case.

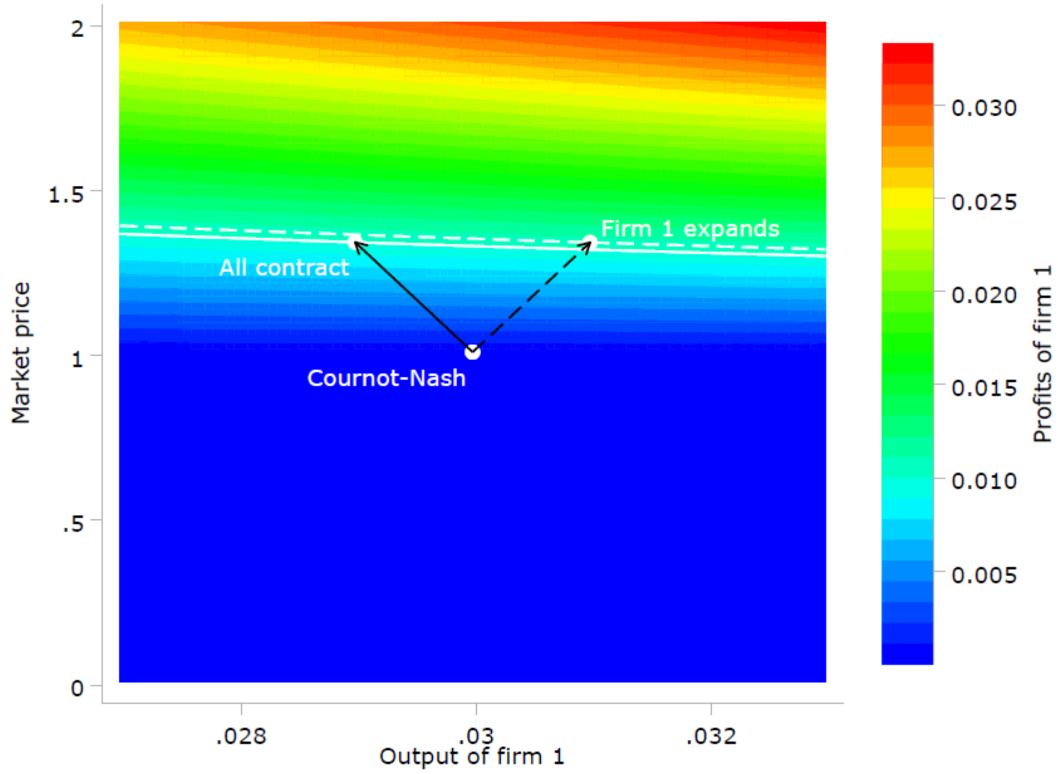

Figure 2: Oligopoly firm 1's profits rise no matter what choice it makes.

The possibilities of firm 1's expansion or contraction is entirely irrelevant to the direction its profits move. Rather, the profits of firm 1 rise because all other firms contracted (though profits rise even more if firm 1 expands.) Still, if the firm is ignorant of its alternatives in the vicinity of Cournot-Nash, the problem is far worse nearing the collusive level of production as seen in Figure 3.

Again it is true that by continuing to contract, the profits of all firms increase. However, the increase is marginal – almost nonexistent. Any firm would see a much larger increase in profits by expanding.

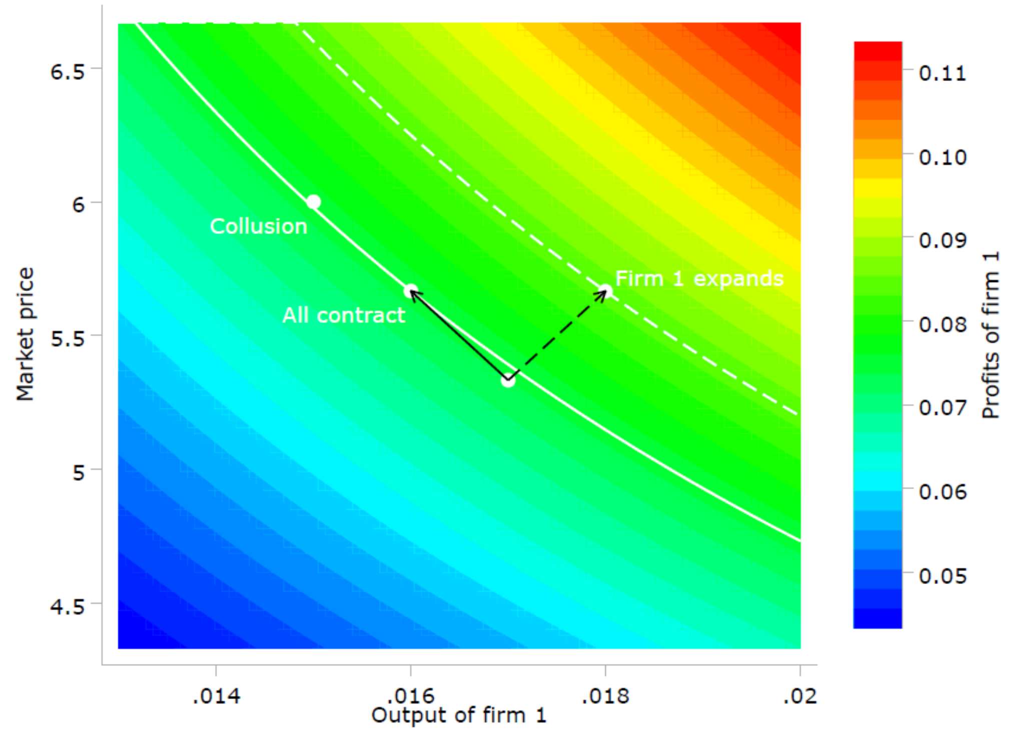

Figure 3: Oligopoly firm 1's profits rise either way, but rise much faster when expanding.

Of course, if all firms expanded, each would suffer losses, however marginal. The firms are technically in a Prisoner's Dilemma.

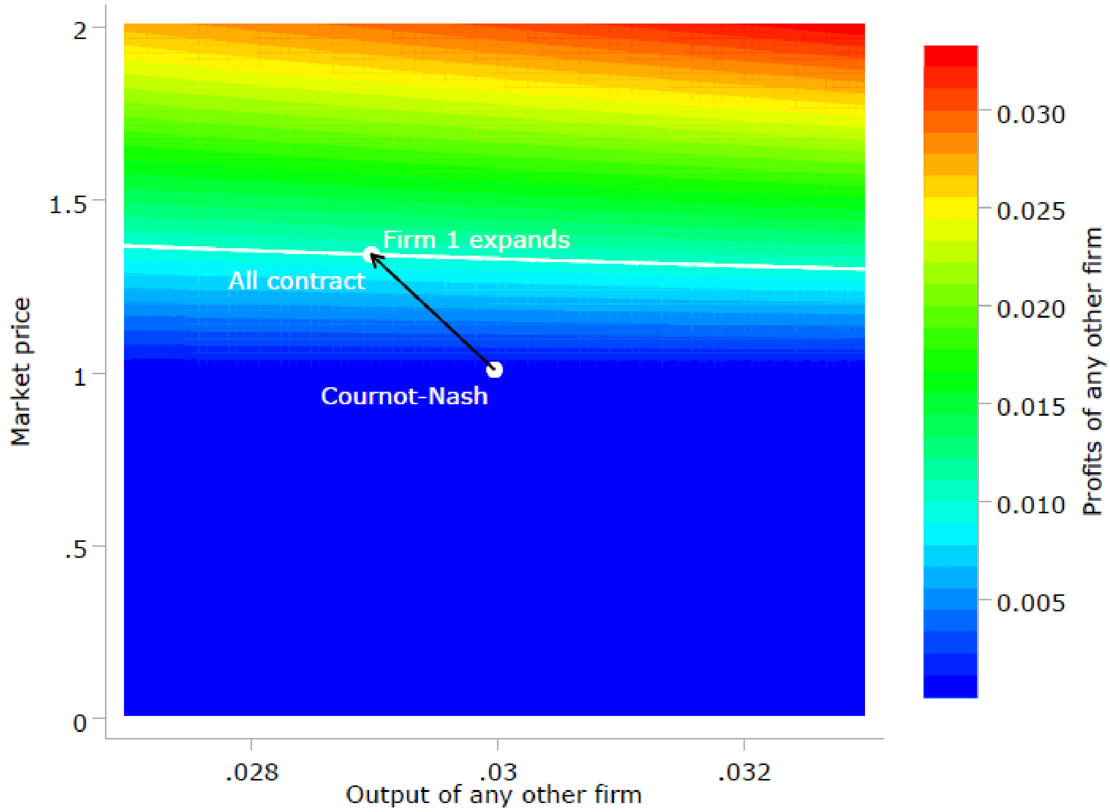

Figure 4: Other firms in a large oligopoly hardly perceive any difference of a single firm's expansion.

Note that in such a large oligopoly, expansion on the part of a single firm from Cournot-Nash does not appreciably change the market price and so the contracting firms can hardly even notice if one firm expands, as seen in Figure 4.

We will investigate further the actual behavior of these "groping" firms in Section 7. For now, it suffices to recognize that such firms regularly fail to move in the direction which maximizes profits and inquire as to why they so fail.

The terrain for every firm is in fact fixed. The reason why competitive firms do not magically head off toward the peak is that the terrain is different for every firm, even if the firms are identical. In general, $\pi_1(q_1, q_2) \neq \pi_2(q_1, q_2)$. Further, no firm would see the collusive outcome as the top of the hill - but this is precisely *because* as in the above figures the peak of the mountain is where they monopolize the market and everyone else goes away. The peak will never be reached because firms disagree as to where the peak exists, and no firm has control over how the other firms operate.

This is exactly the significance of the Cournot equilibrium. Looking at its own map, every firm is there indifferent to marginal changes in its own output and so no firm has incentive to change. Thus, the Cournot levels are stable. This is what defines Cournot production as an equilibrium. On the other hand, at the point of collusion every firm looks at its map and sees greater profits with increased output and so every firm has incentive to expand. This instability in the face of competition for greater profits is what distinguishes the collusive outcome from a competitive equilibrium.

Nor may it be considered rational for firms to behave so that the vector of firm outputs moves in the direction of steepest ascent and so tend toward the collusive result. Rather, each firm reckons the steepest ascent of *its* profits to be in a different direction. There is no consistent direction for the vector of firm outputs to move. Some other explanation is required.

6.3 A note on "irrationality"

The authors note that "30% irrationality...is sufficient to ensure competitive pricing." Irrationality here being "the probability that an agent makes the opposite decision to the rational one" [6, Sec. 3.1]. Of course, by "rational" the authors mean that a firm follows their hill-climbing algorithm. That is, it is "irrational" for the firm to believe that it will have higher profits by stepping in the opposite direction, even if we know objectively that the firm is correct in its belief.

The reason that some measure of "irrationality" yields "competitive pricing" is that such firms are not consistently fooled by the wagon effect. What the authors call "irrationality" would be better described as "the firm occasionally testing its assumption that its own production decision was responsible for the direction of its change in profits."

Such firms take a more experimental approach rather than follow blindly the advice of Standish and Keen. Suppose that the firm changes course with nonzero probability $1/[1+\exp(\Delta\pi/T)]$. That way, when profits are rising rapidly, the firm is almost certain to continue, flips a coin when profits are unchanged, and when profits fall rapidly the firm is almost certain to reverse course.

The results are to some degree dependent upon the choice of T. An overly large T approaches a random walk, while a decidedly small T approaches the asserted firm behaviour of Keen and Standish. Between, firms are less fooled by the wagon and yet do not spend an inordinate amount of time experimenting.

Far from "irrational", such activity directly fills in an important blind spot in the "rational" approach of Standish and Keen.

6.4 Standish-Keen dynamics

The authors argue that their firms' behavior implies "that the Cournot [duopoly] equilibrium is unstable in the direction of both firms decreasing production." But the authors say nothing whatsoever about the stability of the Cournot model of imperfect competition. Simply, they argue that if firms are given different information and forced to follow a particular strategy, that they will not operate like firms which are not so restricted.

The authors implicitly argue that if firms follow their given strategy, that they all get higher profits than if the firms produce at Cournot levels. Thus, they argue, their strategy is, *prima facie*, more rational. It is certainly more rational in the sense that collusion is more rational than ruinous competition. Neoclassical oligopoly theory agrees that the rational strategy is to collude. Rather, the question is whether such an equilibrium is *competitive*. We have seen above that oligopolists following the production strategy of Standish and Keen often will be misled into passing on opportunities for higher profits.

Consider, then, the simple case of an oligopoly of N identical firms with zero costs facing linear demand such that

$$P(Q^d) = \alpha - \beta Q^d$$

According to the authors, the industry will supply $Q^s = \alpha/2\beta$, resulting in a price $p = \alpha/2$.
Each firm then receives a profit of $\pi_i = \alpha^2/4N\beta$.

Suppose one firm recognizes that movements in its short-run profits are driven by industry and not itself. It then pursues the following strategy: Select a new output, hold there until the industry settles into a new *equilibrium*, and then evaluate its decision based on the new equilibrium rather than the transient response of the other firms. To avoid confusion with any equilibrium of strategies, let us call this post-transient response the "medium-run" levels of output and profits.

So long as firm i holds its output at q_i, this effectively removes that amount from demand as seen by the rest of industry. That is, the remaining $N-1$ firms face residual demand

$$P(Q^r) = (\alpha - \beta q_i) - \beta Q^r$$

and therefore the $N-1$ firms supply $\alpha/2\beta + q_i/2$,resulting in a price $p = \alpha/2 - \beta q_i/2$. Thus, firm i received medium-run profits of

$$\pi_i(q_i) = \frac{\alpha - \beta q_i}{2} q_i$$

If all other firms follow Standish and Keen, then firm i maximizes its long-run profits by selecting the level of output which maximizes medium-run profits. Specifically, the firm produces N times its Standish-Keen amount ($q_i = \alpha/2\beta$) and accepts a price half that of Standish-Keen ($p = \alpha/4$) for profits of $\pi_i = \alpha^2/8\beta -$ that is, $N/2$ times greater.

Thus, firm i is no worse off in the case of duopoly, and has greater long-run profit incentive the greater the number of competitors. So each firm has incentive to change strategy pursuit of greater long-run profits. There are two ways of preventing such a competitive firm from seeking these greater profits. Either such profit-seeking behavior is assumed away or the firms must be allowed to agree among themselves not to engage in the behaviour – that is, to collude.

7. Understanding the simulated dynamics

As we have seen in Section 6 the proposed “Keen equilibrium” is not in fact a competitive equilibrium as all firms must agree not to pursue an alternative strategy that exposes opportunities for greater profit. Standish and Keen attempt to demonstrate that their claims are still valid by running computer simulations of firms following the suggested strategies. One important feature of the strategy is that the firm is assign a fixed but not necessarily identical step size. This leads to an unequal division of the market even if the firms are otherwise identical. As it turns out, the market share of each firm playing the author’s strategy is directly related to its assigned step size. Finally, we see how this result is connected to the wagon effect. Consider again our greatly stripped-down example of linear demand and zero production costs, so that

$$\pi_i(q_{1,}q_2, \ldots, q_N) = \left[\alpha - \beta \sum_{j=1}^{N} q_j\right] q_i$$

Now Standish and Keen’s firms do not know the specific shape of the demand curve, so let us weaken again our example by saying that they do know that demand is linear and so firm i does believe that

$$\pi_i(q_{1,}q_2, \ldots, q_N) = \left[\alpha_i - \beta_i \sum_{j=1}^{N} q_j\right] q_i$$

Suppose then, that having received a price p for its product firm i expands production from q_i to $q_i + \delta_i$ and its profits fall, so that it believes that its profits would have been no worse had it instead contracted. That is,

$$(p + \beta_i\delta_i)(q_i - \delta_i) \geq (p - \beta_i\delta_i)(q_i + \delta_i)$$

or

$$\beta_i\delta_i \geq p \tag{7.1}$$

That is, from “Keen equilibrium” when the firm thinks it over-expanded from its “equilibrium” share $\ell_i = q_i/Q^s$, it believes that

$$\beta_i \geq \frac{p}{q_i} = \frac{p}{\ell_i Q^s} = \frac{\alpha/2}{\ell_i\alpha/2\beta} = \frac{\beta}{\ell_i}$$

Likewise, when the firm thinks it over-contracted, it believes that $\beta_i \leq \beta/\ell_i$. Thus, every firm systematically overestimates the elasticity of demand.

Recall from the previous section, however, that a firm facing competitors that follow Standish and Keen and which *believes* that the slope of the inverse demand curve is $-\beta_i$ would then *believe* that it maximizes long-run profits by producing at the level $q_i - \alpha/2\beta_i$. So in "Keen equilibrium" such a firm would *think* it is maximizing long-run profits by producing at $q_i = \ell_i\alpha/2\beta -$ exactly where the authors argue the equilibrium exists.

In other words, a firm following Standish-Keen gets a false idea of how elastic demand is, and so thinks it would do exactly as well pursuing the strategy described above. Unfortunately, the notion is predicated on the false estimate of demand elasticity; a firm which actually varied its output in the manner described in the previous section would find itself with higher profits – in the short term as well as the long. What remains is to more carefully derive the actual "equilibrium" outputs.

7.1 A simplified derivation

A full derivation of market shares requires a deal more space than this note allows. However, suppose that firm i has constant marginal costs c_i and constant step size δ_i. Suppose also that the industry faces linear demand such that if the output of firm i in period t is given by $q_{i,t}$ then market price in the same period is

$$p_t = \alpha - \beta Q_t = \alpha - \beta \sum_j q_{j,t}$$

Profits in the next period are given by

$$\begin{aligned}\pi_{i,1} = (p_1 - c_i)q_{i,1} &= (p_1 - c_i)q_{i,0} + (p_1 - c_i)\big(q_{i,1} - q_{i,0}\big) \\ &= (p_0 - c_i)q_{i,0} + (p_1 - p_0)q_{i,0} + (p_1 - c_i)(q_{i,1} - q_{i,0})\end{aligned}$$

Which is to say that a firm switches between expansion and contraction when

$$\Delta\pi_{i,t} = (-\beta\Delta Q_t)q_{i,t} + (p_{t+1} - c_i)\Delta q_{i,t} \leq 0$$

That is, (5.2) describes not merely the proposed equilibrium, but the general behavior of firms following the strategy of Keen and Standish. To speed things along, let us stipulate that in "equilibrium" expansion and contraction of all firms are synchronized so that

$$\frac{\Delta q_{i,t}}{\Delta Q_t} = \frac{\delta_i}{\sum_j \delta_j}$$

When contracting, firms switch to expansion when

$$\left(\beta \sum_j \delta_j\right) q_{i,t} \leq (p_{t+1} - c_i)\delta_i$$

and when expanding, firms switch to contraction when

$$\left(\beta \sum_j \delta_j\right) q_{i,t} \geq (p_{t+1} - c_i)\delta_i$$

Thus, synchronized firms move toward

$$\frac{\sum_j \delta_j}{\delta_i}(-\beta)q_{i,t} + p_{t+1} - c_i = 0 \tag{7.2}$$

which is to say they behave almost identically to Cournot oligopolists, except that they perceive the slope of the inverse demand curve to be not $-\beta$ but rather

$$\frac{-\beta}{\delta_i/\sum_j \delta_j}$$

When (7.2) applies to all firms in the industry,

$$\left(\beta\sum_i \delta_i\right)Q_t = \left(\beta\sum_i \delta_i\right)\sum_i q_{i,t} = \left(\alpha - \beta Q_t \mp \beta\sum_i \delta_i\right)\sum_i \delta_i - \sum_i \delta_i c_i$$

so

$$2\beta Q_t = \alpha - \frac{\sum_i \delta_i c_i}{\sum_i \delta_i} \mp \beta\sum_i \delta_i$$

and

$$2\beta Q_{t+1} = 2\beta\left(Q_t \mp \sum_i \delta_i\right) = \alpha - \frac{\sum_i \delta_i c_i}{\sum_i \delta_i} \mp \beta\sum_i \delta_i$$

so that

$$\frac{Q_t + Q_{t+1}}{2} = \frac{\alpha - c_\delta}{2\beta}$$

where c_δ is the δ -weighted average marginal cost across firms. Thus, the industry operates on average as if a monopoly with marginal cost c_δ. Call this industry average Q_δ. In addition,

$$\beta q_{i,t} + \beta q_{1,t+1} = 2\beta q_{i,t} \mp \beta\delta_i = (p_{t+1} - 2c_i)\frac{\delta_i}{\sum_j \delta_j} - \beta\delta_i$$

$$= \left(\alpha + c_\delta + \beta\sum_j \delta_j - 2c_i\right)\frac{\delta_i}{\sum_j \delta_j} - \beta\delta_i = [\alpha - c_\delta - 2(c_i - c_\delta)]\frac{\delta_i}{\sum_j \delta_j}$$

so

$$\frac{q_{i,t} + q_{i,t+1}}{Q_t + Q_{t+1}} = \frac{2\beta Q_\delta - 2(c_i - c_\delta)}{2\beta Q_\delta}\frac{\delta_i}{\sum_j \delta_j} = \left[1 - \frac{c_i - c_\delta}{\beta Q_\delta}\right]\frac{\delta_i}{\sum_j \delta_j}$$

Note also that if marginal costs are equal across firms (or, if n any case, $c_i = c_\delta$) then

$$\frac{q_{i,t} + q_{i,t+1}}{Q_t + Q_{t+1}} = \frac{\delta_i}{\sum_j \delta_j}$$

That is, the assignment of δ to the firm determines the firm's market share and therefore firm profits. These results are confirmed via simulation, and are available in HTML/Javascript at Test of "Keen Equilibrium".

7.2 Discussion

As we have seen, the firms of Standish and Keen actually do behave as if following equation (5.1), even if this is an illusion brought on by the strategy assigned to (not chosen by) the firm. As we have just seen, even if the firm followed broadly the assigned strategy, it would benefit the firm to choose a larger step size. Thus, the proposed equilibrium depends on the firm convincing itself in the non-existence of profit opportunities which do in fact exist. If all other firms follow the strategy, the firm need only expand to increase its profits.
Such results depend on the firm's near-total ignorance. Even if the firm believes it is in a Keen equilibrium and by any means happens to discover its market share, then it would know β and would know it would profit by expanding output.

As suggested earlier, the firm could profit by recognizing that its short-run profits are often driven by its competitors' movements more so than its own. If, as in Section 6.4, firm i picks an output q_i and holds it there, then rest of industry will see

$$p(Q^r) = (\alpha - \beta q_i) - \beta Q^r$$

where Q^r is supply from of all other firms. From the perspective of all other firms it is as though demand had fallen by q_i, but the number of firms in the oligopoly had also fallen by one. Thus, the rest of industry reacts to firm i's choice of output by supplying

$$Q^r = \frac{(\alpha - \beta q_i) - c_\delta^r}{2\beta} = \frac{\alpha - c_\delta^r}{2\beta} - \frac{q_i}{2}$$

where c_δ^r is the δ-weigthed average firm marginal cost among all other firms. Thus, industry supplies

$$Q^s = \frac{\alpha - c_\delta^r}{2\beta} + \frac{q_i}{2}$$

leaving a market price of

$$p = \frac{\alpha + c_\delta^r}{2} - \frac{\beta}{2} q_i$$

and therefore firm i received equilibrium profits of

$$\pi_i = \frac{1}{2}[(\alpha - c_i) - (c_i - c_\delta^r) - \beta q_i] q_i$$

As firm I experiments with variations in q_i, it finds that is maximizes profits by producing where

$$2\beta q_i = (\alpha - c_i) - (c_i - c_\delta^r)$$

Noting that

$$c_\delta^r \left(\delta_i - \sum_j \delta_j \right) = \delta_i c_i - c_\delta \sum_j \delta_j$$

So

$$(c_\delta^r - c_i) \left(\delta_i - \sum_j \delta_j \right) = c_i - c_\delta \sum_j \delta_j$$

We find that the best long-run profits seen by such a competitor is achieved at

$$q_i = \frac{1}{2\beta}\left(\alpha - c_i - \frac{c_i - c_\delta}{1 - \delta_i / \sum_j \delta_j}\right)$$

and so operates near to the monopoly level.

A more detailed analysis is provided in the Technical Appendix, but the bottom line is that so long as all firms are following the strategy of Standish and Keen, then few (if any) firms fully pursue opportunity for profit. It does not matter that such competition would be ruinous; only by choosing not to pursue such profits is the authors' result obtained. The result is not a *competitive* equilibrium.

8. Conclusions

The critique by Keen and Standish on the theory of the firm is deficient on every front. Their critique of perfect competition depends on an unaccepted definition of price-taking that confuses textbook models of perfect and imperfect competition. Their critique of Cournot-Nash relies upon firms neither understanding the tradeoff it faces nor how to deal with that lack of understanding. As a result, the firms of Keen and Standish superficially seek out greater profits but forgo profit opportunities by failing to act in their objective best interests. In contrast to the authors contention, firms forced to grope in order to maximize profits are capable of working around the informational limitations Keen and Standish impose in modifying the Cournot game. Standish and Keen may not explicitly model firms' communications with one another, but the firms must somehow come to agreement not to pursue individual strategies which yield greater profits. Either the firms agree among themselves or Standish and Keen simply forbid the strategies. The “Keen equilibrium” is therefore no more competitive than any collusive oligopoly.

References

[1] Paul Anglin. On the proper behavior of atoms: A comment on a critique. *Physica A*, 387:277-280, 2008.

[2] Stephen L. Keen. *Debunking Economics: The Naked Emperor of the Social Sciences*. Zed Books, 2001.

[3] Stephen L. Keen and Russell K. Standish. Profit maximization, industry structure and competition: A critique of neoclassical theory. *Physica A*, 370:81-85, 2006.

[4] Stephen L. Keen and Russell K. Standish. Debunking the theory of the firm - a chronology. *real-world economics review,* (53):56-94, 2010.

[5] Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, 1995.

[6] Russell K. Standish and Stephen L. Keen. Emergent effective collusion in an economy of perfectly rational competitors. In Proceedings of the 7th Asia-Pacific Conference on Complex Systems, December 2004.

[7] Russell K. Standish and Stephen L. Keen. Comment on “On the proper behavior of atoms” by Paul Anglin, September 2013.

[8] Russell K. Standish and Stephen L. Keen. Rationality in the theory of the firm, *World Economic Review*, (5):101-106, 2015.

[9] George J. Stigler. Perfect competition, historically contemplated. *The Journal of Political Economy*, 65(1):1-17, February 1957.