Homo Oecon (2016) 33:135–156
DOI 10.1007/s41412-016-0008-1

RESEARCH PAPER

Industrial Organization and Experimental Economics: How to Learn from Laboratory Experiments

Max Albert[1] · **Andreas Hildenbrand**[2]

Received: 29 June 2015 / Accepted: 15 June 2016 / Published online: 29 June 2016

Abstract Experimental economists frequently claim that they can contribute to industrial organization (IO) by observing individual behavior in, for example, Cournot or Stackelberg games. In these experiments, they regularly falsify the hypothesis of profit maximization, which is, by and large, retained in applied IO. However, what experimental economists test in the laboratory is, at best, the theory of single-person firms. In contrast, empirical IO studies quite large organizations on the basis of field data, while theoretical IO is concerned with internal organization and its links to market behavior. From a modern theory-testing perspective, many experiments should therefore be considered as irrelevant to modern IO; the focus of experimental IO must change if experimental economists want to contribute to IO.

Keywords Experimental economics · Industrial organization · Oligopoly theory · Theory of the firm

JEL Classification B4 · C7 · C9 · D2 · D4

✉ Andreas Hildenbrand
andreas.hildenbrand@agrar.uni-giessen.de

Max Albert
max.albert@wirtschaft.uni-giessen.de

[1] Behavioral and Institutional Economics, Justus Liebig University Giessen, Licher Strasse 66, 35394 Giessen, Germany

[2] Farm and Agribusiness Management, Justus Liebig University Giessen, Senckenbergstrasse 3, 35390 Giessen, Germany

1 Introduction

Industrial organization (IO) was one of the areas in economics where experimental methods were applied first (see Plott 1982, pp. 1485–1486; Roth 1995b, pp. 13–19). Nevertheless, modern IO developed quite independently of experimental economics (see Holt 1995, pp. 352–355). Specialized IO journals typically did not publish experimental work. This is not surprising. Early experiments in IO were concerned with individual behavior in experimental markets (see Plott 1982, pp. 1486–1507) or oligopoly games (see Plott 1982, pp. 1507–1519). In contrast, IO is primarily concerned with the structure and behavior of firms (see, e.g., Tirole 1988, p. 3; Scherer and Ross 1990, p. 1). Especially in oligopoly theory, firms are usually assumed to be quite large (see, e.g., Shapiro 1989, p. 330; Scherer and Ross 1990, p. 199). Of course, single-person firms—like Hotelling's (1929) ice cream vendors—also exist and are legitimate subjects of inquiry. However, as a matter of fact, they play a rather peripheral role in IO (see, e.g., Schmalensee 1989; Sutton 2007; Aghion and Holden 2011).

In the early days of experimental economics, oligopoly games certainly made an important point. The neoclassical theory of the firm does not say anything about the internal structure of a firm. It just assumes profit maximization on the basis of some production function, mostly combined with the assumption of rational expectations, and, in strategic settings, game-theoretic equilibrium assumptions. It does not distinguish between multi-person or single-person firms. For this reason, the neoclassical theory of the firm can be tested in the laboratory. If individuals acting as single-person firms are shown to be unable to maximize profits in an oligopoly game, it follows that a theory claiming that all kinds of firms maximize profits is false. However, it does not follow that multi-person firms are also unable to maximize profits. Moreover, even more importantly, it does certainly not follow that the quirks of individual behavior observed in the laboratory tell us anything about the behavior of multi-person firms in the field.

Let us distinguish, with Bardsley et al. (2010, pp. 196–197), between experimental IO and applied IO. Experimental IO mostly considers individual behavior in experimental markets or oligopoly games, while applied IO is mostly concerned with the behavior of multi-person firms in non-experimental markets. Once it has been accepted that the neoclassical theory as a theory of all kinds of firms is false, it becomes at least unclear what applied industrial economists can learn from individual behavior in experimental markets or oligopoly games. Whatever individuals may do in such markets or games, the multi-person firms that stand in the center of interest in applied IO may behave quite differently.

The difference in subject matters of experimental IO and applied IO may very well explain their separate development. However, the explosive development of experimental economics has changed the situation. Applied industrial economists increasingly begin to take note of experimental IO. Unfortunately, this does not mean that the relation between experimental IO and applied IO has been clarified.

Experimental industrial economists emphasize the relevance of their research to applied IO. In their introduction to a recent special issue of the *International Journal of*

Industrial Organization devoted to experimental IO, Normann and Ruffle (2011, p. 1) point out the advantages of experimental methods in IO, which "allow us to address questions that are otherwise outside the purview of empirical research due to the unobservability of the underlying cost and demand parameters or the outcome variable, namely, the firm's behavior." They even claim special advantages for experimental IO in relation to experimental economics in general. Experiments in IO are supposed to be "largely immune to many of the criticisms launched at laboratory methods in the social sciences" and, allegedly, framing is "ordinarily not an issue" because "experimental market results are robust to the context in which the decision is imbedded". Moreover, "while experiments on individual choice and social preferences have been shown to be sensitive to moral considerations and the degree of anonymity of laboratory subjects' decisions, these considerations are typically orthogonal to any experiment in industrial organization."

We will come back to the issues of framing and "moral considerations" (social preferences), arguing that, in fact, these issues belong into the center of interest in experimental IO.

A much more important issue is conspicuous by its absence. Normann and Ruffle (2011) do not even mention the problem of how individual behavior in experiments relates to the behavior of multi-person firms. This problem is most relevant in the context of the contributions to the special issue in question because there are eleven experiments on individual behavior not mentioning this aspect. In these eleven experiments, individuals act in the experimental markets but their behavior is interpreted as the behavior of firms. A nice example is provided by Lange et al. (2011). They compare the behavior of individuals in their experiment with data on firm behavior—small firms, to be sure, but not single-person firms—in the field. Only three experiments explore the behavior of groups, but they also leave it open if it is appropriate to consider these groups as firms: Barreda-Tarrazona et al. (2011) consider seller collectives and Ahn et al. (2011) consider rent seeking by groups. In both cases, there is no organizational structure or decision rule. Gillet et al. (2011) contribute the only experiment where there are, at least, groups with decision rules.

This is rather typical for the literature on experimental IO. Many experiments that are supposedly concerned with the behavior of firms actually consider the behavior of individuals. Such contributions can only be considered as important to applied IO if we can learn something about the behavior of multi-person firms in, for instance, oligopolistic markets by looking at the behavior of individuals in oligopoly games. But how can this be possible?

Answering this question involves two problems: (1) learning from experiments about non-experimental settings and (2) learning from observations of individual behavior in certain situations about organizational behavior in the same kind of situations.[1] Both problems seem to have a similar structure. We try

[1] The first problem is often referred to as "the problem of external validity". External validity is one of the quality criteria for judging experiments that have been proposed originally in psychology (see Campbell 1957; Campbell and Stanley 1963). For a criticism from the theory-testing perspective see, most notably, Gadenne (1976) and further sources discussed by Shadish et al. (Shadish et al. 2002, ch. 14). In experimental economics, the theory-testing view had for a long time been prevalent (see, e.g., Plott 1982); only recently, external validity considerations have become popular (see Schram 2005).

to learn from cases we can easily observe about difficult-to-observe cases we are really interested in. Stating the problem in this way suggests that, if learning is possible in the first case (which experimentalists typically assume), it should also be possible in the second case. This view is supported by the inductivist idea that learning proceeds by generalizing and/or analogical reasoning, which could be applied in both cases.

However, a deductivist, or theory-testing, perspective, which we defend in this paper, leads to a very different conclusion.[2] Roughly speaking, deductivism assumes that learning from experience proceeds by testing theories and rejecting those that fail while accepting those that pass, where rejection or acceptance decisions are, of course, provisional and may change with new empirical evidence. From a deductivist perspective, it is a matter of the theory under consideration whether we can learn something about one situation by looking at another. In order to clarify the role of experiments for applied IO, then, we have to discuss the theories involved and their relation to the experiments conducted. This will lead us to the conclusion that, in the case of IO theories, we can indeed learn from experiments about non-experimental settings; however, learning from observing individual behavior in certain settings about organizational behavior in these settings is impossible.

In this paper, we argue that the deductivist view, together with what we have learned so far from economic experiments, has rather strong implications for the design of experiments that are relevant to applied IO. Our point is not that experimental IO is irrelevant, or that the hypothesis of profit maximization must be maintained in applied IO at all costs. We argue that the focus in experimental IO should be on different classes of experiments. We do not claim that such experiments do not exist; indeed, we have mentioned one and will mention some more. However, there seems to be no common understanding in IO which kinds of experiments are relevant to applied IO and which kinds are irrelevant (or no longer relevant). Obviously, such a situation can lead to a waste of resources. There is no simple remedy, however, because views about the relevance of experiments are closely connected with entrenched methodological positions.

The paper is organized as follows. In Sect. 2, we review oligopoly games and oligopoly experiments. We will use them as our main examples because they are well-known. Section 3 considers the methodological question of what the theories of applied IO are about and how they are related to the purported experimental tests. We argue that the oligopoly experiments reviewed in Sect. 2 are not relevant to modern applied IO. In Sect. 4, we present an alternative approach to testing applied IO. Section 5 concludes.

[2] Deductivism means, of course, modern hypothetico-deductivism, where theories are taken to be conjectures and, therefore, in need of tests. Hypothetico-deductivism is opposed to inductivism, see Musgrave (2011). External validity concerns have inspired a "new inductivism" in economics; see Guala (2005), Bardsley et al. (2010), and Gadenne (2013) for a thorough methodological criticism.

2 Oligopoly Theory in the Laboratory

2.1 Oligopoly Games

In many areas of applied IO, economists consider the output decisions of a small number of competing firms under the assumption of market-clearing prices. The defining characteristics of an oligopoly are that a firm's market behavior does not only affect its own profit but also affects the profits of the other firms in the market, and that all firms are aware of this interdependence. Moreover, as already explained, industrial economists are mostly interested in quite large or, more precisely, multi-person firms, which, therefore, possess non-trivial organizational structures.

The two benchmark models of oligopoly theory, to be found in every IO textbook, are the Cournot and the Stackelberg model of duopolistic quantity competition with homogeneous products (see, e.g., Tirole 1988; Scherer and Ross 1990). These are probably the models that are most often implemented in the laboratory.

Let us shortly review both models. There are two firms, *A* and *B*, which produce homogeneous products, meaning that, in the eyes of the consumers, each unit of output is as good as any other, no matter who produced it. In the Cournot model, firms decide simultaneously (that is, not knowing the other firm's decision) how much to produce. In the Stackelberg model, firms decide sequentially: one firm moves first (leads) and the other firm, knowing the leader's decision, moves second (follows); thus, the follower reacts to the leader's decision. In both models, the total output is then sold at a uniform market-clearing price, that is, at a price just low enough so that consumers want to buy all units of output. Firms' profits are equal to their revenues minus their production costs.

The theoretical predictions of neoclassical theory for the situations described above result, if we take market clearing for granted, from two independent hypotheses: first, firms maximize profits on the basis of their expectations, and, second, firms have rational expectations or beliefs. In the Cournot model, equilibrium means Nash equilibrium: each firm correctly anticipates the other firm's choice and plays a profit-maximizing answer. In the Stackelberg model, equilibrium means subgame-perfect Nash equilibrium: the second-moving firm maximizes its profit given the first moving firm's choice, and the first mover correctly anticipates the second mover's reaction to all possible choices and maximizes profit in the light of these anticipations. In the most popular specification (linear demand and linear cost functions), unique Cournot and Stackelberg equilibria exist, and equilibrium outputs of both firms are positive if the costs of producing the first unit of output are not too high.

This description specifies the five elements of a game: (1) the players, (2) players' action possibilities, that is, the order of moves and players' possibilities for action when they move, (3) the probability distributions over exogenous events, (4) the players' beliefs, and (5) the players' motivations (in traditional game theory described by utilities and the assumption of utility maximization) and cognitive abilities (unlimited in traditional game theory).[3] Elements (1, 2, 3) are often referred to as the "game form". Specifically, (1) players are the two firms, (2) action possibilities are possible quantities, (3) exogenous events are absent in the basic versions of the models considered here, (4) players' beliefs are rational, and (5) players' maximize their own profits.

2.2 Experimental Games

There are important differences between the description of a game in game theory and the description of an experimental game. An experimental game is defined by the experimental design. The relation between the experimental game and the description of a game in terms of the five elements above is entirely hypothetical. Some of the hypotheses are those the experimenter wants to test. These are often hypotheses about the motivations and cognitive abilities of the participants. Other hypotheses are auxiliary hypotheses that the experimenter hopes to be unproblematic. These are often hypotheses about the effects of design elements on participants' beliefs.

An important element of many experiments is that payoffs are unknown. Experimental games specify material (mostly monetary) payoffs, but participants may care for other things besides their own earnings. This has already been mentioned under the heading of "moral considerations" or social preferences. It is, therefore, not known which game participants play in the context of a given experimental design. There is some hypothesis the experimenter wants to test with the help of the experiment, the target hypothesis. The design is chosen such that it is possible to test the target hypothesis. Let us assume that the target hypothesis is concerned with (5) players' motivations and cognitive abilities. The experimenter may then choose a game form (1, 2, 3) as a basis for the experimental design, a game form that is suitable to test the target hypothesis if (4) participants believe that they act within this game form. The experimental design must be specified such that the experimenter can be reasonably sure that participants actually have these beliefs. Moreover, predictions are often equilibrium predictions, which presuppose that participants have rational expectations about other players' actions. If it cannot be assumed that participants have rational expectations from the start, the experimental design must allow participants to learn about other participants' behavior.

In general, the experimenter must know enough about the five elements of a game in order to draw conclusions about the target hypothesis. Although the presentation of the experimenter's arguments is mostly informal, the central part is a deductive argument from several premises to a prediction which is then borne out by the data generated in the experiment or not. The premises consist of the target hypothesis, the description of the experimental design, and several auxiliary hypotheses that connect the design elements with the five elements of the game. Of course, the target hypothesis may be derived from a general theory of behavior like

[3] The five elements are a generalization from Fudenberg and Tirole (1991, pp. 77–82). Our version also covers behavioral extensions of game theory.

the *homo oeconomicus* (HO) model or some other theory, which may be the ultimate target of the test.

Although this situation can be quite messy in any specific case of experimentation, it is straightforward enough from a methodological point of view. The fact that we need auxiliary hypotheses in a test of some target hypothesis means that we are, here as in every other field, confronted with the so-called Duhem-Quine problem (see, e.g., Cross, 1998 and, in the experimental context, Bardsley et al. 2010, ch. 3). If the deduction of the prediction from the premises is correct—which it is often not, as when critics discover implicit auxiliaries—and the prediction fails, at least one of the premises must be false. From a logical point of view, however, the failure can be blamed on any of the premises. This ambiguity would allow the scientific community to follow a strategy of blaming predictive failures on the auxiliaries and counting predictive successes as corroborations of the target hypothesis, thus making experimentation a pointless exercise that never can lead to the falsification of a target hypothesis even if it is false.

Duhem, who is usually credited with seeing the problem first, did not think of it as unsolvable (see Gadenne 1998, pp. 119–121). The developments in experimental economics demonstrate how the solution looks like in principle. Economists consider many hypotheses (including many auxiliaries) and many different experiments involving different subsets of the hypotheses. Given some predictive failures and some predictive successes, the scientific community decides which assignments of truth values to hypotheses provide the best explanation for the observed failures and successes (see Albert 2010). Of course, at some points in time, the picture may be unclear; several assignments of truth values may compete. Moreover, new hypotheses change the situation, and new experimental results further constrain the possible assignments of truth values. Critics of an existing consensus may combine several accepted hypotheses to derive a new prediction that is then shown to be false. However, although there is no guarantee of progress, it seems that the Duhem-Quine problem is not so severe that progress becomes impossible.

In this paper, we assume that the Duhem-Quine problem can be dealt with along these lines. We are not concerned with this problem in general but with the specifics of oligopoly experiments. Thus, we assume, for the sake of this discussion, that the many auxiliary hypotheses oligopoly experiments share with other experiments are unproblematic. Instead, we focus on the specific target hypotheses and auxiliary hypotheses that come into play when oligopoly theory, viewed as a part of applied IO, is brought into the laboratory.

2.3 Experimental Oligopoly Games

Oligopoly games have been played in the laboratory for about 50 years. Hoggatt (1959) and Sauermann and Selten (1959) published the first relevant experiments, followed by Siegel and Fouraker (1960) and Fouraker and Siegel (1963). However, most of these experiments were of an exploratory nature. In contrast, recent experimental studies, beginning with Holt (1985), are explicitly designed as tests of hypotheses.[4]

Holt (1985) tested whether individuals who are just told that they are "sellers" behave according to a symmetric Nash equilibrium prediction in a Cournot game.[5] He found that average quantities per round were close to the Nash equilibrium quantities. Similar results have been found in more recent experimental studies (see, e.g., Huck et al. 2004).

The first to implement the Stackelberg game were Huck et al. (2001). They tested whether individuals who were told to represent "firms" (IO framing) behave according to the unique (asymmetric) subgame-perfect Nash equilibrium prediction in a Stackelberg game. They found that leaders mostly chose a lower quantity than predicted, and followers typically chose a higher quantity. As the number of rounds increased, the variance remained approximately constant. Similar results have also been found in other experimental studies (see, e.g., Huck and Wallace 2002).[6]

Müller and Tan (2013) also implemented a Stackelberg game with IO framing. However, firms were not only represented by individuals, but also by three-member groups. Compared to individuals' quantity choices, groups' choices were not found to be more in line with the assumption of profit maximization. The same answer was given by Raab and Schipper (2009) on the basis of a Cournot game with IO framing. Firms were also represented by individuals and by three-member groups. Neither Müller and Tan nor Raab and Schipper made use of institutional theories of the firm. Groups were organized according to ad-hoc theories of the firm at best.

3 Experimental Evidence and Theories of the Firm

As already explained, applied IO is mostly concerned with multi-persons firms that, in many cases, are quite large organizations. The basic hypothesis in oligopoly theory, as it is employed in applied IO, is that these firms maximize profits. Do we have experimental evidence that speaks against profit maximization by such organizations?

3.1 Learning in the Laboratory About Theories

Oligopoly theory is part of neoclassical economics. In principle, neoclassical economics is based on methodological individualism, that is, the requirement that all social phenomena should be explained in terms of the actions of

[4] For an early survey, see Cyert and Lave (1965). For a recent survey of the early Cournot oligopoly experiments, see Bosch-Domènech and Vriend (2008). For an overview of the early and recent Cournot oligopoly experiments, see Requate and Waichman (2011, p. 39). For the credit to Holt (1985), see, e.g., Plott (1989).

[5] Holt's design leads to further, asymmetric Nash equilibria.

[6] In the recent past, many more such experiments were run. Often, the standard duopoly games were extended by a pre-play stage in order to endogenize the sequence of play. For example, see Huck et al. (2002), Fonseca et al. (2005, 2006), and Müller (2006). For a survey, see Hildenbrand (2010).

individuals.[7] In practice, the neoclassical theory of the firm, of which oligopoly theory is a special case, does not conform to this requirement. It assumes that firms maximize profits on the basis of a given technology but ignores the question of how individuals within the firm coordinate their actions in order to maximize profit (see, e.g., Nadiri 1982). It is a kind of reduced-form theory that does not spell out the conditions under which profit maximization is to be expected. Because of this, it is hard to test.

Can we bring firms into the laboratory? How, exactly, does a firm differ from a single person or some multi-person group? The neoclassical theory of the firm provides no answers. It only says that a firm maximizes profits on the basis of a production technology. The theory does not say how to recognize a firm.

The simplest (and, presumably, the original) interpretation of the neoclassical theory of the firm is that it provides a theory of all firms, including single-person firms. This interpretation is not consistent with methodological individualism. Moreover, interpreted in this way, the theory can be considered as falsified by the oligopoly experiments discussed above which show that individuals are not able to maximize profits even in quite simple experimental markets. If we wish to uphold some version of the neoclassical theory, its domain must be restricted in an appropriate way, excluding at least single-person firms.[8]

One way to do this in a systematic way is to turn to extensions of neoclassical economics that make the theory of the firm (more) consistent with methodological individualism.[9] This extension is institutional economics (see, e.g., Furubotn and Richter 2005) or, more precisely, the institutional economics of the firm (see, e.g., Milgrom and Roberts 1992; Aghion and Holden 2011).

The neoclassical version of institutional economics retains the HO model and combines it with game theory. Under the HO model, we understand, roughly, the hypothesis that people are perfectly rational, purely egoistic, and materialistic. The HO model together with game theory is a very general theory (T) that can be applied to derive predictions for many specific circumstances or environments (see Fig. 1).

On the basis of T, we can consider individuals or groups that are endowed with a technology and predict their decision behavior in market-like environments. If the central part of T, the HO model, were well-tested and well-corroborated, the situation in applied IO would be relatively straightforward. Like meteorologists, who can rely on well-corroborated physical theories, applied industrial economists

[7] See, e.g., Kincaid (1998). For a review and a discussion of different meanings of the term "methodological individualism", see Hodgson (2007).

[8] One referee suggested the analogy of a map. Representative-agent models, and corresponding experiments, are similar to maps without all the details, good for driving but not for an orienteering competition. Our criticism could then be considered as a request for more detailed maps (maybe because we have come as far as we can go by car). This is a nice picture but we would not like to emphasize it too much because the analogy between theories and maps (and even more between experiments and maps) may not go far.

[9] Perfect consistency with methodological individualism is not a reasonable aim given the complexity of large organizations. Theories connecting organizational structure with market behavior offer a reasonable compromise between representative-agent approaches on the one hand and (infeasible) reductions of economics to psychology on the other hand. This compromise is, for instance, exemplified by Ostrom (1990) approach to governance of a commons.

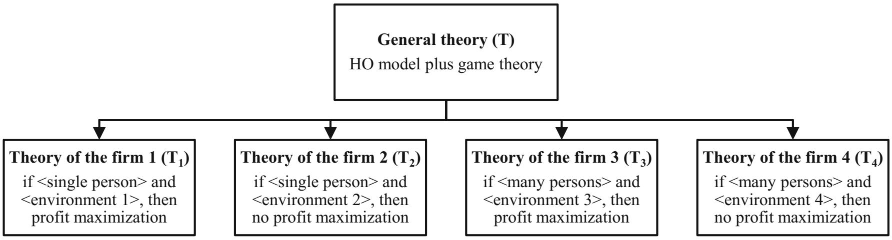

Fig. 1 Deductive tree. The different theories of the firm follow deductively from T by applying the model to specific institutional environments where a single person or several persons take choices that result in profit-maximizing behaviour or not. In such a deductive tree, falsifications go up but not down: if T_1 is falsified, T is also falsified, but T_2 is not. See Albert (1996, pp. 464–466) for a general discussion

would have to find out about the relevant initial conditions—in this case, about the environments in which people make choices—and derive the consequences. As meteorology shows, this may be a difficult task even when the basic laws are known. Two problems are involved: the empirical problem of improving the descriptive part of complex models, and the theoretical problem of finding the implications of the basic laws when applied to such a description. The predictive performance of the models must be evaluated and turned into a program of further improving the descriptive parts. In meteorology, laboratory experiments play a relatively small role. Instead, simulations become very important because it is impossible to derive predictions from these complex models by analytic methods.

The situation in economics is, unfortunately, very different. The HO model must be considered as falsified, not least due to experimental work. The logical and methodological consequences of this fact are often overlooked.

Consider, for instance, the Stackelberg game for two single-person firms discussed above. Let us call the environment of the two firms “environment 1”. The general theory T, applied to environment 1, implies that both firms maximize their monetary payoffs. This is a very simple theory of the firm, called T_1. This theory has been falsified experimentally by Huck et al. (2001). It follows that T is falsified, because T must be false if T_1 is false.

The falsification fits in with many other experiments where participants do not maximize their own monetary payoffs. The simplest ones are dictator and ultimatum experiments (see, e.g., Camerer, 2003, ch. 2). Many of these experiments differ not much from oligopoly games except for the IO framing, that is, the fact that participants in oligopoly experiments are told that they are sellers or represent firms.[10] Consider, for instance, a bilateral monopoly of two single-person firms trading one unit of some good. Let one firm, called “proposer”, be a unilateral price setter; the other firm is called “responder”. The situation is equivalent to an ultimatum game. The general theory T, applied to this ultimatum-game environment, implies that the responder maximizes his monetary payoff. This is, again, a

[10] Thus, the claim of Normann and Ruffle (2011, p. 1) that framing plays no role in IO experiments is rather dubious: it seems that framing is the main difference between many IO experiments and other experiments.

very simple theory of the firm, which has been falsified experimentally (see Hoffman et al. 1994). Moreover, since the responder needs to form no expectations, the failure of T is usually blamed on the HO model. The same holds for the behavior of the second mover in the Stackelberg game by Huck et al. (2001).

However, even if the HO model must be considered as falsified, it does not follow that any of the other theories of the firm that can be derived from it can be considered as falsified. Consider a theory like T_3 that considers multi-person firms in "environment 3", which is also characterized by Stackelberg competition between firms but specifies, in addition, that the agents of each firm have entered a verifiable and binding contract that forces them to pay a lot of money to some third party if they collectively deviate from profit-maximizing behavior. With this specification, T_3 follows from T. While T is falsified, T_3 is not falsified, and it is quite plausible that people who failed to maximize profits in environment 1 maximize profits in environment 3. Thus, we have no reason at all to reject T_3 just because we falsified T.

Even if the HO model is false in general, it may very well be true, then, that profit maximizations fails for single persons in environment 2 and holds for groups in some environment 3, just as the HO model predicts (see Fig. 1). A falsification of the HO model implies that all the different theories of the firm, which before could have been viewed as deductive consequences of the general theory T, become isolated, so to speak: if the HO model is true, all its deductive consequences must be true; but if the HO model is false, any proper subset of its consequences can be true or false.

In fact, given that the neoclassical theory T is falsified, its critical potential is also lost. Consider the theory T^* that all multi-person firms, no matter how they are organized, maximize profits. As long as it was rational to believe in T (if it ever was), it was rational to reject T^* since T implies that, in some environments, groups fail to maximize profits because of, for instance, organizational problems like team production problems (see also T_4 in Fig. 1). Once T is falsified, T^* can again be taken seriously. Of course, we may argue that we know enough about team production problems from other experiments in order to reject T^*. However, we do not know enough from experiments in order to reject a hypothesis like T^{**}, namely, that the management of a large firm is able and willing to choose an internal structure that solves organizational problems and leads to, or at least approximates, profit maximization.

Thus, the deductivist methodology does not force applied industrial economists to give up their basic hypothesis because of the experiments we have discussed so far. Of course, the hypothesis of profit maximization started out as a general hypothesis about all firms, including single-person firms. (Never mind that, in this general form, the hypothesis would be inconsistent with T.) The duopoly experiments show that this general theory is false. From this perspective, retreating to a hypothesis like T**, which predicts that only large firms maximize profits, is a defensive move of restricting the domain of a hypothesis so that it no longer covers those situations in which it had been falsified.

While such domain restrictions cannot just be ruled out—the resulting theory may, after all, be true—, the present instance is problematic because of the presumption, well expressed by Normann and Ruffle (2011), that the hypothesis of profit maximization by large firms is difficult to test in the field. Restricting the basic hypothesis of applied oligopoly theory such that it becomes untestable would turn this part of IO from science into ideology.

3.2 Learning from Observations of Individual Behavior

Let us apply the deductive tree to the problem of learning from individual behavior about large organizations' behavior in similar settings. From an inductivist point of view, one must generalize from some observed cases to non-observed cases that are different on some and similar on other dimensions. Inductivists reconstruct this as analogical reasoning, a special kind of inductive inference (cf., e.g., Guala 2005, pp. 194–199). Deductivists accept analogical reasoning but reconstruct it as deductive reasoning with an implicit premise (see Musgrave 2011). Inductivists think that analogical reasoning can lend empirical support to a theory. Deductivists say that this depends on the status of the implicit premise: unless the premise is already well-confirmed, the theory must be tested before we can say that it is supported by the evidence.[11]

The deductivist reconstruction of generalizations makes explicit the missing premise in the argument from experimental observations to general theories. In the case of IO, we would need a hypothesis that bridges the gap between the behavior of individuals and multi-person firms. Thus, the claim that the duopoly experiments are relevant to applied IO might be considered as assuming, implicitly, that individuals and multi-person firms always show the same behavior in the same situation. This hypothesis is a very strong version of the representative-agent hypothesis. We therefore call it the strong representative-agent (SRA) hypothesis.

Under the SRA hypothesis, all kinds of experiments involving individual behavior would become directly relevant to applied IO. For instance, deviations from profit maximization in the single-person experiments would not only falsify the profit-maximization hypothesis of multi-person firms but also indicate how such firms would deviate from profit maximization. We could, then, develop a theory of single-person behavior in Stackelberg games (where the hypothesis of profit maximization fails) and test it in further single-person experiments, all the time relying on the fact that our theorizing is relevant to applied IO because of the SRA hypothesis.[12]

However, in order to make use of the SRA hypothesis in this way, the SRA hypothesis itself should be well-confirmed. At least, we should be able to show that

[11] Inductivists do not accept all kinds of analogical reasoning. They try to state conditions under which experiments are externally valid; cf., e.g., Guala's (2005) "eliminative inductivism". Jones (2011) rightly argues that Guala's methodology condemns experimental economics to irrelevance. This should at least motivate experimentalists to consider the deductivist alternative before they buy into eliminative inductivism.

[12] While theories explaining the behavior of individuals in experiments would give rise, via the SRA hypothesis, to theories about the behavior of multi-person firms, it is rather doubtful that these theories would count as explanations. Since the SRA hypothesis does not explain why firms behave like individuals, it would at best support a kind of as-if theorizing about firms as defended by Friedman (1953). This criticism applies also to the heuristic uses of the SRA hypothesis mentioned below.

the assumption that the SRA hypothesis is true explains the results of many different tests (see the discussion of the Duhem-Quine problem above). However, this is not at all the case.

From the perspective of neoclassical theory T, the SRA hypothesis is not convincing since, as we have already shown, it implies that the SRA hypothesis is false: depending on the organizational structure, the behavior of multi-person firms may differ from individual behavior. This can be shown by a simple example.

Consider the following four-player game. Each player $i = 1, 2, 3, 4$ chooses a nonnegative variable L_i called effort. Players act simultaneously. Efforts determine outputs $x_A = f_A(L_1, L_2)$, $x_B = f_B(L_3, L_4)$ according to neoclassical production functions f_A, f_B. The output price is determined according to an inverse demand function: $p(x_A + x_B)$. Players receive fixed shares of the revenue they generate with their team partner. Players' material payoffs are equal to their revenue share minus their costs, which are increasing in effort. Thus, each two-player team forms a firm that competes in a standard Cournot duopoly against another firm. Each player bears the costs of his own effort but shares the revenue with his colleague, implying that each firm faces an internal team production problem.[13]

In this environment, the neoclassical HO assumption that players rationally maximize their own material payoffs does not lead to profit maximization by firms. The team production problem induces free-riding and leads to firm outputs that are lower than those of profit-maximizing firms.

From the perspective of behavioral economics, the SRA hypothesis is also not convincing. Individuals exhibit social preferences in many experiments like dictator or ultimatum games (see, e.g., Roth 1995a). It seems to be highly implausible that firms behave in this way. Individuals with social preferences will, when acting as members of groups, take their group members into account, not only outsiders. In fact, prisoner's dilemma experiments, which can be thought of as simplified Cournot experiments, frequently show in-group favoritism, that is, participants interacting with an out-group member are less cooperative than participants interacting with an in-group member (see, e.g., Yamagishi and Kiyonari 2000; Yamagishi and Mifune 2009). Hence, it seems that members of groups are more likely to cooperate with their colleagues, which, as we have shown, can easily lead to more competitive behavior of the group against outsiders.

Actually, as far as the experimental evidence goes, more speaks in favor of the idea that groups maximize profits than for the idea that groups behave like individuals. For example, Bornstein et al. (2008) experimentally examine a market for a homogeneous good in which two producers compete in prices. Producers act simultaneously. Each producer is represented by either a participant or a group of participants consisting of two or three members. Sessions consisted of one hundred

[13] Here is a simple specification with explicit solutions. Production functions are $f_A(L_1, L_2) = L_1^{0.5} L_2^{0.5}$ and $f_B(L_3, L_4) = L_3^{0.5} L_4^{0.5}$. The inverse demand function is $p = a - b(x_A + x_B)$. The material payoff of players $i = 1, 2$ is $\pi_i = p x_A / 2 - c L_i$ and the material payoff of players $i = 3, 4$ is $\pi_i = p x_B / 2 - c L_i$, where $0 < c < a$. With free-riding, the unique Nash equilibrium in which both firms serve the market is $L_i = \frac{a - 4c}{3b}, i = 1, 2, 3, 4$. Internal cooperation leads to the standard Cournot-Nash equilibrium with $L_i = \frac{a - c}{3b}, i = 1, 2, 3, 4$. With internal cooperation and collusion between the firms, the solution is $L_i = \frac{2a - 4c}{8b}, i = 1, 2, 3, 4$. If $a = 10c$, free-riding leads to collusion.

rounds. The number of rounds was not revealed to the participants. Producers and groups were matched randomly at the beginning of the sessions. The matching remained constant. Thus, participants played repeated Bertrand duopoly games. Bornstein et al. found that single agents behaved differently from groups: prices were significantly higher in competition between individuals than in competition between groups, that is, groups came closer to profit-maximizing behavior.

Other results from experiments on group decision making also suggest relevant behavioral differences between individuals and groups. In a series of prisoner's dilemma experiments, group behavior is found to be more competitive and less cooperative than individual behavior (see Schopler and Insko 1992). In ultimatum experiments, similar behavioral patterns are observed. Groups in the role of proposers offer less than individuals, and groups in the role of responders are typically willing to accept less than individuals (see Robert and Carnevale 1997; Bornstein and Yaniv 1998). That is, in both prisoner's dilemma and ultimatum experiments, the (subgame-perfect) Nash equilibrium solution is a better predictor of group behavior than individual behavior. This is also true for dictator, gift exchange, trust, and centipede experiments (for reviews, see Bornstein 2008; Kugler et al. 2012).

However, Raab and Schipper (2009) do not find significant behavioral differences between individuals and groups in a Cournot triopoly experiment. The same is true for Müller and Tan's (2013) Stackelberg experiment on individual and group decision making. Thus, it seems that there are conditions where groups and individuals behave alike in oligopoly games, but other conditions where this is not the case.

In view of the theoretical arguments and the experimental results, then, we are highly skeptical about the SRA hypothesis that single-person firms and multi-person firms show the same behavior. Moreover, we see no tenable alternative to the SRA hypothesis that would allow us to consider the Cournot and Stackelberg experiments on individual decision making as relevant to applied IO. Some other experimental approach to IO, then, is needed—an approach that avoids the implicit reliance on the dubious hypothesis that single-person and multi-person firms show the same behavior. Answering the opening question, we do not have experimental evidence that speaks against profit maximization by all kinds of organizations.

While the SRA hypothesis must be rejected in the light of the available evidence, it can still be used in a heuristic argument (that is, in order to derive new theories). However, it cannot be used to argue that the experimental evidence lends empirical support to some applied IO theory. When it comes to empirical support, then, we are back to testing.

We agree with Normann and Ruffle (2011) that it would be a great advantage if one could rely on experimental testing in IO because it is difficult to test applied IO theories using field data. Testing a theory in the laboratory requires that the experimental design falls within the domain of the theory; then, and only then, the experiment is a test of the theory. To put it more concisely, testing a theory by an experiment requires (1) that the theory (possibly extended by auxiliary hypotheses) together with a description of the experimental design implies a prediction, and (2) that the experimental design makes it possible to check this prediction by observation.[14] Since applied IO theories are concerned with the market behavior of large organizations, they imply no predictions about the behavior of individuals in experimental markets or oligopoly games. If experimental economists want to contribute to a field concerned with organizational behavior, they must test theories of organizational behavior.

3.3 Organizational Structure

As we have seen, many experiments falsify the HO model. Some of these experiments can even be viewed as falsifications of simple theories of single-person firms, theories that predict profit maximization. However, it does not follow from these experiments that *the neoclassical theory of the firm as a theory of organizational behavior* must be considered as falsified.[15] In fact, in many environments, other general theories of behavior contradicting the HO model might predict profit maximization by groups, even in situations where the HO model would not predict profit maximization.

We have already considered a simple four-player game of two firms with an internal team-production problem where the assumption that players rationally maximize their own material payoffs does not lead to profit maximization by the firms. However, let us assume that the HO model is false and that players have social preferences. It may, then, be possible that the two-player teams solve their internal cooperation problems. This would lead to profit maximization by teams and, consequently, to the standard Cournot-Nash equilibrium.[16]

Thus, depending on the motivational assumptions, a theory of the firm that takes the firm's internal organization into account may predict profit maximization or not. Specifically, demonstrating that participants in laboratory experiments are able to overcome cooperation problems, like the problem of team production, falsifies the HO model but not the hypothesis that multi-person firms maximize profits. Quite to the contrary: in some environments (like the four-player game considered above), solving internal cooperation problems is a precondition of profit maximization. Of course, finding the solution to the combined problems of internal cooperation and profit maximization is difficult for real people, who are boundedly rational at best. In addition to the motivational aspects, we have to take cognitive limitations into

[14] Bardsley et al.'s (2010, ch. 2) more complicated account boils down to the account given in the text.

[15] Moreover, it is well known that deviations from profit maximization may lead to behavior that, from the outside (i.e., without direct access to a firm's technology), still looks like profit maximization. This can easily be seen from the example above, where an observer who just knows the inverse demand function would be unable to say whether the quantity choices of a firm are affected by internal free-riding or not (see also the numerical specification in note 13, where free riding just looks like profit maximization with twice the costs of effort). For some purposes, the hypothesis of apparent profit maximization may suffice. Obviously, this hypothesis can be defended even if the HO model has been falsified.

[16] Obviously, the team production problem, if unsolved, reduces efforts in comparison to the standard model of profit maximization. A reduction in outputs may improve the situation for all players if it brings them closer to the collusive solution. Depending on the specification, the equilibrium outcome achieved by four free-riding homines oeconomici may be higher than, lower than, or equal to the collusive optimum. Situations where players profit from solving their team production problems are possible.

account. Because difficult optimization problems may be solved better by a cooperating team (see Davis 1992 for a review of selected examples on group decision making), the consideration of cognitive aspects creates further possibilities for theories that predict profit maximization by teams but not by single persons.

To summarize, organizational behavior depends on three elements, namely, agents' motivations, agents' cognitive abilities, and the rules of the game. It is not impossible, and, given what we know about human behavior, not even implausible that well-designed organizations maximize profits—or at least come close to profit maximization—while single persons do not.

Our point is not that we have strong evidence for profit maximization by large firms. Rather, we make two points, one logical and one empirical. The logical point is that the experimental falsification of the HO model does not imply a falsification of certain applied economic theories like the theory of the profit-maximizing multi-person firm or the theory of international trade. The empirical point is that, as far as we know, it is not inconceivable that organizations are better at profit maximization than individuals, implying that we cannot dismiss the logical possibility as empirically irrelevant.

In consequence, experimental research on single-person decision making in market-like environments or oligopoly games cannot test a theory that assumes profit maximization of multi-person firms. Whether firms maximize profits or not depends on how they solve their internal problems of cooperation. Of course, nothing in the argument presented here is new or surprising. What is surprising is the fact that experimental IO, by and large, seems to ignore this problem. Currently, the vast majority of contributions to experimental IO fall into just one of two categories: investigations of market behavior or oligopoly-game behavior of individuals or single-person firms; or investigations of the internal organization of multi-person firms. Largely missing are experiments connecting the two,[17] and a clear idea of what should be tested.

What, then, can and should experimental economists test if they wish to contribute to a theory of multi-person firms? If it is admitted that experimental economics has falsified the HO model, and with it the assumption that individuals maximize profits, the domain of the neoclassical theory of the firm has to be restricted. Only in a restricted form, the theory can be upheld. Restricting it to multi-person firms is one reasonable possibility. And this restricted theory cannot be tested in experiments where firms are represented by single individuals.

The alternative is to consider theories of the firm that focus on the link between organizational structure and market behavior. This link takes center stage in current IO (see, e.g., Holmstrom and Tirole 1989), which incorporates the early insights

[17] See Sauermann and Selten (1959) and subsequent work by Selten (cf. Engel 2010, pp. 449–450) for early exceptions, and Kirstein and Kirstein (2007) and Raab and Schipper (2009) for recent ones. See also Engel (2010) for a useful but quite terse survey of the empirical literature on the behavioral differences between individuals and corporate actors (or, as we say, groups). Engel refers to many results where groups behave more in line with the HO model than individuals, both cognitively and motivationally. However, there are also many results that point in the opposite direction. Importantly, and as emphasized by Engel, the internal organization of groups matters. However, most of the literature surveyed by Engel bears no immediate relation to applied IO, although almost all of it may have some relevance.

from Berle and Means' (1932, pp. 352–357) "new concept of the corporation", later developments like Alchian and Demsetz's (1972) contractual view of the firm, and Williamson's (1975) transaction cost approach.

If we wish to retain the assumption that firms maximize profits, it must be assumed that organizational structure transforms individual utility maximization (or individual boundedly rational behavior) within the firm into profit maximization in the market. However, the number of experimental studies that examine the relation between organizational structure and oligopolistic market behavior is currently small (cf. Kirstein and Kirstein 2007, pp. 3–5). Nevertheless, there is a clear-cut role for experimental economics in such a context, namely, to test hypotheses on the link between organizational structure and profit maximization (or other kinds of firm behavior if the assumption of profit maximization by multi-person firms cannot be retained).

The advantage of experimental economics is that it can test basic hypotheses in this area. One obvious question concerns the relation between cooperation in a group and collusion in the market. For instance, does internal cooperation come at the cost of cooperation with other groups? Other questions concern the relation between delegation and hierarchy within the firm and behavior on markets (see Güth et al. (2011, 2012), Fershtman and Judd (1987), or Fershtman (1985) on theoretical oligopoly markets, and Berninghaus et al. (2009), Kirstein and Kirstein (2007), or Fehr and Falk (1999) on experimental labor markets). In all these contexts, agents' social preferences will most likely affect internal cooperation and, in this way, behavior on markets. For this reason, agents' "moral considerations" are not orthogonal to the problems of interest in IO, as claimed by Normann and Ruffle (2011, p. 1); rather, they are in the center of interest, as in other areas of experimental economics.

4 Example of How to Proceed

A way to bring empirical methods, and experimental methods specifically, to bear on the theory of multi-person firms can, in our view, be achieved if we follow the principle of methodological individualism and turn to approaches that link the market behavior of multi-person firms, or teams, to their internal structure. As already discussed, such approaches yield hypotheses that can be tested, and are tested, with experimental methods. Subsequently, we outline one experimental approach along these lines, although there are, of course, other possible approaches.

For example, an experimental approach to IO could be based on Alchian and Demsetz' (1972) contractual view of the firm. For Alchian and Demsetz (p. 783), a firm is a contractual structure "with (1) joint input production, (2) several input owners, (3) one party who is common to all the contracts of the joint inputs, (3a) who has rights to renegotiate any input's contract independently of contracts with other input owners, (3b) who holds the residual claim, and (3c) who has the right to sell his central contractual residual status" (the numbers are modified). Moreover, (4) individuals within the firm maximize (expected) utility on the basis of a utility function that is increasing in income and leisure and has no further arguments. And, last but not least, (5) the firm is organized in a way that individual utility maximization within the firm is transformed into profit maximization on the market (see also Holmstrom, 1982).

Items (1) to (3) put restrictions on the internal structure of the firm. Item (4) is a consequence of the HO assumption; it follows from the HO assumption if leisure is more satisfying than work throughout and if there is no possibility of consumption within the firm. Item (5) requires further specifications of the organizational structure within the limits given by items (1) to (3); these specifications depend on the relevant theory of individual behavior, here given in (4).

One might be tempted to think that this theory of the firm is doomed from the start, because its central behavioral assumption follows from the HO model, which we consider as falsified. However, just to repeat this elementary point, it does not follow from the falsity of the HO model that people are never rational, egoistic, and materialistic. They may still behave like this under specific circumstances, like those in a firm organized according to the principles considered by Alchian and Demsetz (1972): certain environments may make people behave in line with the HO model, even if the HO model is false in general.[18]

As long as we are explicit concerning the kind of environment that supposedly triggers HO behavior, this restricted HO model is testable in principle. Moreover, if the relevant kind of environment is described in universal terms, that is, as long as we refer to an organizational structure and not just to "the environment of a big 21st century US corporation", it is quite likely testable by experimental methods.[19]

It is, of course, not necessary to stick to (restricted versions of) the HO model. As we have shown by the simple model used in the second section, a theory of team production where free riding is prevented not by material incentives but by something like group solidarity or informal sanctions may also predict profit maximization. Moreover, we need not restrict theories of the firm to those that imply profit maximization in the market. There are other theories of the firm, like Leibenstein's (1966) theory of X-inefficiency, which can also be derived from theories of human behavior together with assumptions on the organizational structure.[20] Hence, we do not intend to defend (restricted versions of) the HO model or the profit maximization hypothesis. These hypotheses serve as examples. We show that it is possible to take them seriously without either retreating to a sterile ideology (by declaring them untestable) or by relying on an experimental research program that requires the acceptance of highly dubious, untested, or even falsified auxiliary hypotheses.

A general strategy for devising experimental designs that are relevant to applied IO would be to focus on hypotheses about how individual motivations are influenced by organizational structure. Take, for instance, the case of contracts. Even if the HO model is false in general, and even if individuals do not maximize profits in oligopoly games despite being told that they represent firms, it may still be

[18] According to Schotter (1981, p. 11) and Hodgson (1988, p. 10), it is characteristic of social institutions that they trigger specific motivations and expectations and, consequently, specific behaviours.

[19] There are, of course, exceptions. If the relevant environment is one of very large teams with a lot of face-to-face interactions (thus ruling out internet experiments), it may become too expensive to test the theory.

[20] For a survey of the new institutional theories of the firm, see Furubotn and Richter (2005, pp. 361–469).

true that they turn out to be profit maximizers in oligopoly situations if they act within an incentive-compatible contract.

This strategy opens up a broad range of experiments that are relevant to applied IO. On the one hand, it is possible to analyze existing firms and build hypotheses concerning the link between organizational structure and individual behavior that translates into organizational behavior. If these hypotheses are general und stated in universal terms, it is quite likely that they can be tested in the laboratory. The results of these tests can then feed back into the process of theory formation. It is in no way necessary to stick to restricted HO models and profit maximization; we have just used these (after all, still prominent) hypotheses as an example.

Of course, such a research strategy may or may not be successful. At least, success is conceivable. That is more than can be said of the alternatives that we have analyzed in this paper.

5 Conclusion

Oligopoly games have been played in the laboratory for about 50 years. Unfortunately, as we have argued at length, this does not mean that oligopoly theory, as it is used in modern IO, has always been put to the test. In many experiments, the role of firms is taken by single persons. From the theory-testing or deductivist perspective, this is problematic. Either the experimental designs are not within the domain of the theory—because the theory is about multi-person firms while the instructions just tell single persons that they represent firms. Or the conclusions rest, implicitly, on a highly dubious auxiliary hypothesis, namely, that multi-person firms and single persons show the same behavior. This problem arises with all theories in which multi-person firms are treated as corporate decision makers. For that reason, much of IO is affected. Since Amazon, IBM, or what have you cannot be brought into the laboratory, there are no experimental tests of these theories.

However, this does not mean that experimental methods are irrelevant to applied IO. Methodological individualism holds that theories about corporate decision making are themselves in need of explanation. Specifically, the assumption that multi-person firms maximize profits, which is the basic hypothesis in oligopoly theory and other areas of applied IO, is itself in need of explanation. Such an explanation is especially important in view of the fact that the theory is difficult to test with data on firm behavior on real markets. Current IO theory tries to explain the behavior of firms by linking their organizational structure with their market behavior. According to methodological individualism, the link is provided by theories of individual behavior or, more precisely, institutional theories of the firm. It is, as we have argued, exactly these theories that can, and should, be tested by experimental economists.

Acknowledgments For detailed comments, we are grateful to Siegfried Berninghaus, Björn Kuchinke, Manfred Stadler, and two anonymous referees. For further comments and suggestions, we thank Hartmut Kliemt and participants at the 2011 Karlsruhe Seminar on Behavioral Economics, the 2011 ESA European Conference, the 2011/12 MAGKS Workshop, the 2011/12 Tübingen Economic Workshop, the 2011/12 Hohenheimer Oberseminar, and the 2013 Session of the Economic Theory Group ("Theoretischer Ausschuss") of the Verein für Socialpolitik.

References

Aghion, P., & Holden, R. (2011). Incomplete contracts and the theory of the firm: What have we learned over the past 25 years? *Journal of Economic Perspectives, 25*, 181–197.

Ahn, T. K., Isaac, R. M., & Salmon, T. C. (2011). Rent seeking in groups. *International Journal of Industrial Organization, 29*, 116–125.

Albert, M. (1996). "Unrealistische Annahmen" und empirische Prüfung. *Zeitschrift für Wirtschafts- und Sozialwissenschaften, 116*, 451–486.

Albert, M. (2010). Critical rationalism and scientific competition. *Analyse & Kritik, 32*, 247–266.

Alchian, A. A., & Demsetz, H. (1972). Production, information costs, and economic organization. *American Economic Review, 62*, 777–795.

Bardsley, N., Cubitt, R., Loomes, G., Moffatt, P., Starmer, C., & Sugden, R. (2010). *Experimental economics: Rethinking the rules*. Princeton: Princeton University Press.

Barreda-Tarrazona, I., García-Gallego, A., Georgantzís, N., Andaluz-Funcia, J., & Gil-Sanz, A. (2011). An experiment on spatial competition with endogenous pricing. *International Journal of Industrial Organization, 29*, 74–84.

Berle, A. A., & Means, G. C. (1932). *The modern corporation and private property*. New York: Macmillan.

Berninghaus, S.K., Güth, W., Hoppe, C., & Paul, C. (2009). *Duopolistic Hiring and Sales Competition: A Theoretical and Experimental Analysis*. Working Paper, Karlsruhe Institute of Technology, Department of Economics and Business Engineering.

Bornstein, G. (2008). A Classification of Games by Player Type. In A. Biel, D. Eek, T. Gärling, & M. Gustafsson (Eds.), *New issues and paradigms in research on social dilemmas* (pp. 27–42). New York: Springer.

Bornstein, G., Kugler, T., Budescu, D. V., & Selten, R. (2008). Repeated price competition between individuals and between teams. *Journal of Economic Behavior & Organization, 66*, 808–821.

Bornstein, G., & Yaniv, I. (1998). Individual and group behavior in the ultimatum game: Are groups more "rational" players? *Experimental Economics, 1*, 101–108.

Bosch-Domènech, A., & Vriend, N. J. (2008). The classical experiments on cournot oligopoly. In C. R. Plott & V. L. Smith (Eds.), *Handbook of experimental economics results* (pp. 146–152). Amsterdam: North-Holland.

Camerer, C. F. (2003). *Behavioral game theory: Experiments in strategic interaction*. New York: Russell Sage Foundation.

Campbell, D. T. (1957). Factors relevant to the validity of experiments in social settings. *Psychological Bulletin, 54*, 297–312.

Campbell, D. T., & Stanley, J. C. (1963). *Experimental and quasi-experimental designs for research*. Boston: Houghton Mifflin.

Cross, R. (1998). The Duhem-Quine thesis. In J. B. Davis, D. W. Hands, & U. Mäki (Eds.), *The handbook of economic methodology* (pp. 107–110). Cheltenham: Edward Elgar.

Cyert, R. M., & Lave, L. B. (1965). 'Collusion. *Conflit et Science Économique', Économie Appliquée, 18*, 385–406.

Davis, J. H. (1992). Some compelling intuitions about group consensus decisions, theoretical and empirical research, and interpersonal aggregation phenomena: Selected examples 1950–1990. *Organizational Behavior and Human Decision Processes, 52*, 3–38.

Engel, C. (2010). The behaviour of corporate actors: How much can we learn from the experimental literature. *Journal of Institutional Economics, 6*, 445–475.

Fehr, E., & Falk, A. (1999). Wage rigidity in a competitive incomplete contract market. *Journal of Political Economy, 107*, 106–134.

Fershtman, C. (1985). Managerial incentives as a strategic variable in duopolistic environment. *International Journal of Industrial Organization, 3*, 245–253.

Fershtman, C., & Judd, K. L. (1987). Equilibrium incentives in oligopoly. *American Economic Review, 77*, 927–940.

Fonseca, M. A., Huck, S., & Normann, H.-T. (2005). Playing cournot although they shouldn't: Endogenous timing in experimental duopolies with asymmetric cost. *Economic Theory, 25*, 669–677.

Fonseca, M. A., Müller, W., & Normann, H.-T. (2006). Endogenous timing in duopoly: Experimental evidence. *International Journal of Game Theory, 34*, 443–456.

Fouraker, L. E., & Siegel, S. (1963). *Bargaining behavior*. New York: McGraw-Hill.

Friedman, M. (1953). The methodology of positive economics. In M. Friedman (Ed.), *Essays in positive economics* (pp. 3–43). Chicago: University of Chicago Press.

Fudenberg, D., & Tirole, J. (1991). *Game theory*. Cambridge: MIT.

Furubotn, E. G., & Richter, R. (2005). *Institutions and economic theory* (2nd ed.). Ann Arbor: University of Michigan Press.

Gadenne, V. (1976). *Die Gültigkeit psychologischer Untersuchungen*. Stuttgart: Kohlhammer.

Gadenne, V. (1998). Spielarten des Duhem-Quine-Problems. *Logos, 5*, 117–148.

Gadenne, V. (2013). External validity and the new inductivism in experimental economics. *Rationality, Markets and Morals, 4*, 1–19.

Gillet, J., Schram, A., & Sonnemans, J. (2011). Cartel formation and pricing: The effect of managerial decision-making rules. *International Journal of Industrial Organization, 29*, 126–133.

Guala, F. (2005). *The methodology of experimental economics*. Cambridge: Cambridge University Press.

Güth, W., Pull, K., & Stadler, M. (2011). Intrafirm conflicts and interfirm competition. *Homo Oeconomicus, 28*, 367–378.

Güth, W., Pull, K., & Stadler, M. (2012). Strategic delegation in price competition. *Theoretical Economics Letters, 2*, 355–360.

Hildenbrand, A. (2010). *Cournot or stackelberg competition? A survey of experimental research*. MPRA Paper 26845, University of Munich, Munich Personal RePEc Archive.

Hodgson, G. M. (1988). *Economics and institutions: A manifesto for a modern institutional economics*. Cambridge: Polity.

Hodgson, G. M. (2007). Meanings of methodological individualism. *Journal of Economic Methodology, 14*, 211–226.

Hoffman, E., McCabe, K., Shachat, K., & Smith, V. (1994). Preferences, property rights, and anonymity in bargaining games. *Games and Economic Behavior, 7*, 346–380.

Hoggatt, A. C. (1959). An experimental business game. *Behavioral Science, 4*, 192–203.

Holmstrom, B. R. (1982). Moral hazard in teams. *Bell Journal of Economics, 13*, 324–340.

Holmstrom, B. R., & Tirole, J. (1989). The theory of the firm. In R. Schmalensee & R. D. Willig (Eds.), *Handbook of industrial organization* (pp. 61–133). Amsterdam: North-Holland.

Holt, C. A. (1985). An experimental test of the consistent-conjectures hypothesis. *American Economic Review, 75*, 314–325.

Holt, C. A. (1995). Industrial organization: A survey of laboratory research. In J. H. Kagel & A. E. Roth (Eds.), *The handbook of experimental economics* (pp. 349–443). Princeton: Princeton University Press.

Hotelling, H. (1929). Stability in competition. *Economic Journal, 39*, 41–57.

Huck, S., Müller, W., & Normann, H.-T. (2001). Stackelberg beats Cournot: On collusion and efficiency in experimental markets. *Economic Journal, 111*, 749–765.

Huck, S., Müller, W., & Normann, H.-T. (2002). To commit or not to commit: Endogenous timing in experimental duopoly markets. *Games and Economic Behavior, 38*, 240–264.

Huck, S., Normann, H.-T., & Oechssler, J. (2004). Two are few and four are many: Number effects in experimental oligopolies. *Journal of Economic Behavior & Organization, 53*, 435–446.

Huck, S., & Wallace, B. (2002). Reciprocal strategies and aspiration levels in a cournot-stackelberg experiment. *Economics Bulletin, 3*, 1–7.

Jones, M. K. (2011). External validity and libraries of phenomena: A critique of guala's methodology of experimental economics. *Economics and Philosophy, 27*, 247–271.

Kincaid, H. (1998). Methodological individualism/atomism. In J. B. Davis, D. W. Hands, & U. Mäki (Eds.), *The handbook of economic methodology* (pp. 294–300). Cheltenham: Edward Elgar.

Kirstein, R., & Kirstein, A. (2007). *Inefficient Intra-Firm Incentives Can Stabilize Cartels in Cournot Oligopolies*. Working Paper 04, Otto von Guericke University Magdeburg, Faculty of Economics and Management.

Kugler, T., Kausel, E. E., & Kocher, M. G. (2012). Are groups more rational than individuals? Property rights, and anonymity in bargaining games. *WIREs Cognitive Science, 3*, 471–482.

Lange, A., List, J. A., & Price, M. K. (2011). Auctions with resale when private values are uncertain: Evidence from the lab and field. *International Journal of Industrial Organization, 29*, 54–64.

Leibenstein, H. (1966). Allocative efficiency and X-efficiency. *American Economic Review, 56*, 392–415.

Milgrom, P., & Roberts, J. (1992). *Economics, organization and management*. Englewood Cliffs: Prentice-Hall.

Müller, W. (2006). Allowing for two production periods in the Cournot duopoly: Experimental evidence. *Journal of Economic Behavior & Organization, 60*, 100–111.
Müller, W., & Tan, F. (2013). Who acts more like a game theorist? Group and individual play in a sequential market game and the effect of the time horizon. *Games and Economic Behavior, 82*, 658–674.
Musgrave, A. (2011). Popper and hypothetico-deductivism. In D. M. Gabbay, S. Hartmann, & J. Woods (Eds.), *Handbook of the history of logic* (Vol. 10, pp. 205–234). Amsterdam: North-Holland.
Nadiri, M. I. (1982). Producers theory. In K. J. Arrow & M. D. Intriligator (Eds.), *Handbook of mathematical economics* (pp. 431–490). Amsterdam: North-Holland.
Normann, H. T., & Ruffle, B. (2011). Introduction to the special issue on experiments in industrial organization. *International Journal of Industrial Organization, 29*, 1–3.
Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge: Cambridge University Press.
Plott, C. R. (1982). Industrial organization theory and experimental economics. *Journal of Economic Literature, 20*, 1485–1572.
Plott, C. R. (1989). An updated review of industrial organization: Applications of experimental methods. In R. Schmalensee & R. D. Willig (Eds.), *Handbook of industrial organization* (pp. 1109–1176). Amsterdam: North-Holland.
Raab, P., & Schipper, B. C. (2009). Cournot competition between teams: An experimental study. *Journal of Economic Behavior & Organization, 72*, 691–702.
Requate, T., & Waichman, I. (2011). A profit table or a profit calculator? A note on the design of cournot oligopoly experiments. *Experimental Economics, 14*, 36–46.
Robert, C., & Carnevale, P. J. (1997). Group choice in ultimatum bargaining. *Organizational Behavior and Human Decision Processes, 72*, 256–279.
Roth, A. E. (1995a). Bargaining experiments. In J. H. Kagel & A. E. Roth (Eds.), *The handbook of experimental economics* (pp. 253–348). Princeton: Princeton University Press.
Roth, A. E. (1995b). Introduction to experimental economics. In J. H. Kagel & A. E. Roth (Eds.), *The handbook of experimental economics* (pp. 3–109). Princeton: Princeton University Press.
Sauermann, H., & Selten, R. (1959). Ein Oligopolexperiment. *Zeitschrift für die gesamte Staatswissenschaft, 115*, 427–471.
Scherer, F. M., & Ross, D. (1990). *Industrial market structure and economic performance* (3rd ed.). Boston: Houghton Mifflin.
Schmalensee, R. (1989). Inter-industry studies of structure and performance. In R. Schmalensee & R. D. Willig (Eds.), *Handbook of industrial organization* (pp. 951–1009). Amsterdam: North-Holland.
Schopler, J., & Insko, C. A. (1992). The discontinuity effect in interpersonal and intergroup relations: Generality and mediation. *European Review of Social Psychology, 3*, 121–151.
Schotter, A. (1981). *The economic theory of social institutions*. Cambridge: Cambridge University Press.
Schram, A. (2005). Artificiality: The tension between internal and external validity in economic experiments. *Journal of Economic Methodology, 12*, 225–237.
Shadish, W. R., Cook, T. D., & Campbell, D. T. (2002). *Experimental and quasi-experimental designs for generalized causal inference*. Boston: Houghton Mifflin.
Shapiro, C. (1989). Theories of oligopoly behavior. In R. Schmalensee & R. D. Willig (Eds.), *Handbook of industrial organization* (pp. 329–414). Amsterdam: North-Holland.
Siegel, S., & Fouraker, L. E. (1960). *Bargaining and group decision making: Experiments in bilateral monopoly*. New York: McGraw-Hill.
Sutton, J. (2007). Market structure: Theory and evidence. In M. Armstrong & R. Porter (Eds.), *Handbook of industrial organization* (pp. 2301–2368). Amsterdam: North-Holland.
Tirole, J. (1988). *The theory of industrial organization*. Cambridge: MIT.
Williamson, O. E. (1975). *Markets and hierarchies: Analysis and antitrust implications*. New York: Free.
Yamagishi, T., & Kiyonari, T. (2000). The group as the container of generalized reciprocity. *Social Psychology Quarterly, 63*, 116–132.
Yamagishi, T., & Mifune, N. (2009). Social exchange and solidarity: In-group love or out-group hate? *Evolution and Human Behavior, 30*, 229–237.