

<https://helda.helsinki.fi>

Two Strands of Field Experiments in Economics : A Historical-Methodological Analysis

Nagatsu, Michiru

2020-01

Nagatsu , M & Favereau , J 2020 , ' Two Strands of Field Experiments in Economics : A
Historical-Methodological Analysis ' , Philosophy of the Social Sciences , vol. 50 , no. 1 , pp.
45-77 . <https://doi.org/10.1177/0048393119890393>

<http://hdl.handle.net/10138/313942>

<https://doi.org/10.1177/0048393119890393>

acceptedVersion

Downloaded from Helda, University of Helsinki institutional repository.

This is an electronic reprint of the original article.

This reprint may differ from the original in pagination and typographic detail.

Please cite the original version.

Philosophy of the Social Sciences

Two strands of field experiments in economics: A historical- methodological analysis

Journal:	<i>Philosophy of the Social Sciences</i>
Manuscript ID	POSS-19-0059.R1
Manuscript Type:	Original Manuscript
Keywords:	experimental economics, field experiments, randomized experiments, internal and external validity, exhibits, history of economics, Development Economics, methodology of economics
Note: The following files were submitted by the author for peer review, but cannot be converted to PDF. You must view these files (e.g. movies) online.	
Nagatsu_Favereau_Final_LaTeX2e class file for SAGE Publications.zip	

SCHOLARONE™
Manuscripts

Two strands of field experiments in economics: A historical-methodological analysis

August 30, 2019

Abstract

While the history and methodology of laboratory experiments in economics have been extensively studied by philosophers, those of field experiments have not attracted much attention until recently. What is the historical context in which field experiments have been advocated? And what are the methodological rationales for conducting experiments in the field as opposed to in the lab? This paper addresses these questions by combining historical and methodological perspectives. In terms of history, we show that the movement toward field experiments in economics has two distinct roots. One is the general orientation of medical and social sciences to evidence-based policy evaluation, which gave rise to *randomized* field experiments in economics (e.g., behavioral public policy, poverty alleviation policy, etc.) The other is an awareness of several methodological limitations of lab experiments in economics, which required practitioners to get out of the lab and into the field. In these senses, the movement is a consequence of influences from both outside and inside economics: the general evidence-based trend in policy science and an internal methodological development of experimental economics. In terms of methodology, we show that these two roots resulted in two somewhat different notions of “external validity” as methodological rationales of field experiment. Finally, we suggest that analysis of experiments as *exhibits* highlight a methodological strategy in which both strands methodologically complement each other.

Keywords: Experimental economics, Field Experiments, Randomized Experiments, Methodology of Economics, History of Economics, Internal and External Validity, Exhibits

JEL codes: B210, B410, C930, C900

1 Introduction

The experimental turn in economics in the last 50 years has invited substantial historical and methodological analyses (Guala, 2005; Bardsley et al., 2010; Heukelom, 2014; Svorenčík and Maas, 2015).¹ Many of these analyses focus on describing and justifying how economics, a discipline once regarded as non-experimental, has established a canon of laboratory experimental methodology despite skepticism from both within and outside the profession.

In contrast, the use of *field* experiments, an increasingly popular trend in economics, has just started to be reflected on methodologically and historically. While the existing discussions on this topic by practitioners (Harrison and List, 2004; Levitt and List, 2009; DellaVigna, 2009; Dufflo, 2006; Charness et al., 2013) suggest the involvement of heterogeneous traditions such as behavioral and experimental economics, development economics, education economics and political economy, much philosophical attention has been narrowly focused on the methodological problems of randomization. This state of the art makes it difficult to understand the significance of field experiments in economics and their historical and methodological complexity.

The aim of this paper is to offer a historically-informed methodological analysis of field experiments in economics. We argue that once a basic but crucial distinction between two historical strands of field experiments in economics is made, their respective methodological advantages and weaknesses, and the ways in which they complement each other, will become clear. The first strand, which we call the

¹This turn is part of the broader applied or empirical turn that economics have taken place in the last 50 years. For an analysis of these broader turns see ?. In this paper we focus on the experimental turn of such broader turns.

1
2
3
4
5 lab-in-the-field experiments (LFEs) strand, extends the laboratory experimental eco-
6 nomics to address the *artificiality* of traditional lab experiments. The second strand,
7 which we call the randomized field experiments (RFEs) strand, applies randomized
8 controlled trials, widely and increasingly-used in the social sciences on economic
9 problems, to provide ‘evidence’ of ‘what works’ for policy. In this strand, the key
10 methodological challenge concerns the *generalizability* of results from RFEs for pol-
11 icy. Although both artificiality and generalizability concerns are often framed as the
12 problem of the external validity of extrapolation from experimental findings, this is
13 misleading since their underlying methodological concerns are distinct. After mak-
14 ing this distinction clear, we propose that the exhibit analysis (Sugden, 2005, 2008;
15 Bardsley et al., 2010) highlights a useful methodological strategy to exploit each
16 strands’ methodological complementarity.
17
18
19
20
21
22
23
24
25
26
27
28

29 The paper proceeds as follows: Section 2 characterizes the two historical strands
30 of field experiments in economics. Section 3 outlines the corresponding method-
31 ological rationales of each strands, and how they relate to the notion of external
32 validity discussed in the methodological literature. Section 4 offers a strategy to
33 exploit methodological complementarity of both strands of field experiments with an
34 illustrative case. Section 5 concludes.
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52

2 Two historical origins of field experiments in economics

In this section, we show that field experiments in economics have two distinct origins. One stems directly from the advancement of lab experiments in economics at the turn of the century (2.1). This strand gave rise to the lab-in-the-field experiments (LFEs) as a way to overcome the limitations of laboratory experiments: their artificiality relative to more naturalistic real-world settings. The second strand, the randomized field experiments (RFEs), originates from social field experiments in the second half of the 20th century, which have been later adopted to evaluate smaller-scale economic policy interventions (2.2).

Our basic distinction is not purely historical but motivated by the methodological notion of *controlled variation*. Controlled variation, or simply control, is a hallmark of experimental methods, which enables experimenters to establish a causal link between the putative cause and its putative effect in an experiment (Guala, 2005, 67). While the LFEs have their roots in the efforts of experimental economists to *directly* control variation, RFEs originate from the attempt to *indirectly* control variation by random assignment of subjects into treatment(s) and a control group. As we will see, these distinct strategies to achieve control in experiments give rise to distinct traditions and methodological concerns.

2.1 The first strand: Field experiments as an extension of lab economic experiments

As Guala (2008a) points out, experimental economics has its origins in many disciplines such as mathematics, philosophy, psychology, sociology and political science, in addition to economics, which makes it difficult for us to tell its history comprehensively. While there are different ways to tell this history (e.g. Roth, 1995), it is useful for our purposes to focus on the methodological notion of *control*. The first historical strand we identify originates in the attempt to increase control in the economics laboratory. The driver of this strand was to tighten control in order to secure causal inferences (2.1.1). At the same time, however, the lab creates an artificial space to secure control of variables of interest. This motivates a call to relax such control in order to decrease the artificiality of the lab. This historical strand, the lab-in-the-field experiments (LFEs), is defined by this move from tightening to relaxing control. However, LFEs stay close to the idea of direct control, as the way this relaxation takes place is systematic (2.1.2).

2.1.1 Tightening the control: The establishment of experimental economics in the lab (1950-2000)

According to the orthodox account (Roth, 1995), three historical origins can be identified in laboratory experimental economics. The first can be traced to the psychologist Thurstone's experiment in Chicago, which tried to measure individuals' indifference curves in the 1930s. This line has continued as experiments on individual judgement and decision-making under labels such as mathematical psychology and

1
2
3
4
5 behavioral decision research, forming the foundations for behavioral economics (see
6 Heukelom, 2014). The second line starts with the experiments at RAND by the
7 mathematicians Dresher and Flood in 1950, now known as the Prisoner's Dilemma
8 game. The third line originates from the Harvard economist Chamberlin's classroom
9 experiments to test market equilibrium predictions in the 1940s, leading to Vernon
10 Smith's seminal contributions to market experiments.
11
12
13
14
15
16

17 A common element of experimental economics from these lines is its incremental
18 and systematic effort to control for economically important variables. For example,
19 Vernon Smith's (1976) Induced Value Theory proposes sufficient conditions for a
20 microeconomic market experiment to achieve control for subjects' preferences over
21 outcomes of their actions by using a medium of reward that (the experimenter knows)
22 the subjects value in specific ways (Guala (2005, 232-233); Camerer (2003, 35)). Al-
23 though psychologists' decision-theoretic experiments traditionally used hypothetical
24 payments, the use of task-related incentives has also become more common through
25 empirical and methodological debates with economists (see Hertwig and Ortmann,
26 2001, for the methodological differences between experimental economics and ex-
27 perimental psychology). As a result, requiring task-related monetary incentives has
28 become part of the convention in experimental economics to control not only for
29 preferences, but also for the features of the objects of choice, such as private-public
30 goods and risky lotteries (Bardsley et al., 2010, 248-249). As this example of incen-
31 tives shows, control has been a central interest of experimental economics, and the
32 laboratory has been regarded as conducive to increasing control.
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48

49 From this methodological vantage point, many of the early experimental studies
50
51
52

1
2
3
4
5 look more ‘field-like’ in the sense that they do not control for many variables. For
6
7 example, Ledyard (1995, 124) criticizes a lack of control in one of the earliest public
8
9 goods experiments by Bohm in 1972, in which subjects’ willingness to pay, and belief
10
11 about the group size, were unmeasured and uncontrolled. Another example is Cham-
12
13 berlin’s (1948) market experiment, which allowed his students to move around freely
14
15 in the classroom and negotiate bilaterally with one another without any predefined
16
17 rules. These ‘field-like’ features have been gradually eliminated from experimental
18
19 economics practices.
20

21 22 **2.1.2 From the lab to the field: Relaxing control in a controlled way** 23 24 **(2000-)** 25

26
27 Since the late 1990s and 2000s, there has been a movement to conduct economic
28
29 experiments “in the field” as opposed to the lab (Harrison and List, 2004). The early
30
31 work of John List, one of the central figures in this movement, and collaborators (List
32
33 and Shogren, 1998; List and Lucking-Reiley, 2000) aimed at studying methodological
34
35 questions, such as whether or not the calibration of hypothetical and actual bidding is
36
37 specific to goods and contexts. A variety of goods is needed to answer such questions,
38
39 so they conducted a field auction experiment with actual sports card (baseball and
40
41 football) collectors at a sports card show in Denver, CO in December 1995. As
42
43 evident in this early methodological interest, the field experiments movement is a
44
45 natural extension of lab economic experiments to address questions originating in
46
47 the lab.
48

49 In addition to an inability to address such methodological questions, some of
50
51

1
2
3
4
5 the limitations of lab experiments have been recognized by the practitioners. One
6
7 worry has been that an effort to standardize and control everything tightly in the
8
9 lab might lead to a result that is not relevant to or informative about real world
10
11 policy-making. Even worse, such artificial experiments may result in the unintended
12
13 loss of control over relevant variables. For example, the unnatural or artificial lab
14
15 environment might make the subjects confused about what the task is about, which
16
17 means that the experimenter loses control over the subjects' framing of or beliefs
18
19 about the game (Harrison and List, 2004). Therefore, the key slogan of experimental
20
21 economists in conducting field experiments is that one should "lose control in a
22
23 controlled way" (Bolton and Ockenfels, 2012, 671). Levitt and List (2009, 2) offer a
24
25 useful characterization:

26
27
28 This [...] movement approaches field experiments by taking the tight
29
30 controls of the lab to the field. In doing so, the analyst bridges labora-
31
32 tory and naturally-occurring data by systematically relaxing the controls
33
34 inherent in a laboratory experiment.

35
36
37 The key phrase here is "systematically relaxing the controls". Field experimen-
38
39 talists' responses to the problem of the artificial laboratory and the difficulties of
40
41 interpreting and extrapolating the results to outside the laboratory wall have two
42
43 components, namely systematicity and controls. We discuss them in turn.

44
45
46 Systematicity refers to how different experiments are designed, which allows the
47
48 adding of one or more extra components that represent the natural settings in which
49
50 economic decisions are made. Table 1 shows this systematic character of field ex-
51
52 periments by comparing the influential taxonomy by Harrison and List (2004) and

a more recent alternative by Charness et al. (2013). While details of terminology differ, both taxonomies are defined in terms of the systematic ways in which the modifications are made to the lab experiments (see Table 1). The table makes explicit an important methodological characteristic of field experiments in this strand, namely the systematic way in which experimental variables are varied.

Taxonomies		Variables			
Charness et al.	Harrison&List	subject	referent	location/situation	awareness
Conventional lab exp.		students	imposed	artefactual	yes
2*Extra-lab exp.	Artefactual field exp.	nonstandard	imposed	artefactual	yes
	Framed field exp.	nonstandard	field	artefactual	yes
2*Field exp.	n/a	nonstandard	field	natural	yes
	Natural field exp.	nonstandard	field	natural	no

Table 1: Taxonomies of field experiments by Charness et al. (2013) and Harrison&List (2004)

For example, if you observe a fair offer in the Ultimatum game in the conventional lab experiment, but observe a less fair offer in the experiment with non-student subjects, you can infer that this difference is likely to be due to characteristics of subject pools (Henrich, 2004). Similarly, if you observe differences in how people choose in the Diner's dilemma game in both artefactual and framed field experiments, then you can infer that the difference is due to the nature of the task or referents, e.g., a dish in a restaurant versus an abstract experimental currency on a computer screen (Gneezy et al., 2004). In contrast, if you see no difference in the results between these experiments, then you become more confident that those differences do not matter.² The main strength of this approach is that you can attribute a cause of difference in observations to one difference maker. That is, you can make a stronger inference about what causes what by changing one variable of interest while

²For more examples, see Gerber and Green (2012, 10, footnote 20).

1
2
3
4
5 holding the others constant, which Guala (2005, 67) calls *controlled variation* in the
6 context of laboratory experiments. Although controlled variation in this context is
7 far from perfect, since there are many other potentially relevant variables that are
8 different across experiments, the idea behind comparing findings across these types
9 of experiments is clearly analogous to controlled variation in the laboratory.
10
11
12
13
14

15 A famous example of this type of comparison is a labor markets experiment by
16 Gneezy and List (2006), which is a (natural) field version of Fehr et al. (1993). The
17 latter experiment was designed to test the theory of “gift exchange” originally pro-
18 posed by Akerlof (1982), who argued that in labor markets employers offer wages
19 higher than their reservation utility and employees in turn reciprocate with extra
20 effort, which is distinct from the prediction of sub-game perfect equilibrium. Gneezy
21 and List (2006) demonstrate substantial but short-lived reciprocal behavior of em-
22 ployees in natural one-time task job markets.
23
24
25
26
27
28
29
30

31 The relaxation of controls seems to contradict the idea of controlled variation, but
32 this is not the case. Note that the notion of control is used in an ambiguous way, as in
33 the phrase “losing control in a controlled way”. The latter control refers to controlled
34 variation, which is considered to define experimental science (e.g. Camerer, 2003). In
35 contrast, the former control refers to artificial features of the lab experiments, which
36 are undesirable either (i) because they lead to unintended loss of control (e.g. List,
37 2007; Bardsley, 2008) or (ii) because they make interpretation difficult by creating an
38 artificial environment that does not correspond to naturally occurring economic set-
39 tings in any straightforward way. In sum, although practitioners disagree about the
40 exact taxonomy of experiments as well as the extent to which the field experiments
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

1
2
3
4
5 should be tied strictly to the lab experiments (e.g. Bardsley et al., 2010, 241), there
6
7 is a common understanding of the importance of controlled variation and concern
8
9 about the artificiality of lab experiments.
10

11 12 13 **2.2 The Second Historical Strand: Field Experiments as a** 14 15 **Tool for Evaluating Policies** 16 17

18 Field experiments in economics are also directly inspired by large social scale public
19 policy evaluation. Such inspiration defines the second strands of field experiments
20 in economics. During the first part of the 20th century, field experiments were ad-
21 vocated as a useful tool to evaluate large scale public policy. Economists, more
22 specifically labor economists, implemented several field experiments to assess such
23 large scale policies (2.2.1). More recently, field experiments have been extensively
24 conducted in economics in the last decade, in particular within development eco-
25 nomics, to assess the impact of aid development programs (2.2.2.). Such evaluations
26 are local rather than large-scale. However, they both emphasize the statistical rigor
27 that the random allocation design of the field experiments allows. The random as-
28 signment is a particular and indirect method of achieving controlled variation, the
29 emphasis on which defines the second strand of field experiments³, in contrast to the
30 first strand, in which the notion of control is not tied to particular methods such as
31 random assignments of treatments.
32
33
34
35
36
37
38
39
40
41
42
43
44
45

46 ³For a discussion about the control or balance over variables in randomized experiments, see
47 Heckman (1992), Heckman et al. (1998), Heckman and Smith (1995), Heckman et al. (1997).
48
49
50
51
52

2.2.1 Social field experiments in economics (1960-2000)

As stated above, many field experiments in economics are rooted in the history of social field experiments. The latter aim to assess public policies, primarily on a large scale (Levitt and List, 2009). Greenberg and Shroder (2004), who have provided the largest work on reporting all the social experiments that have been conducted so far, define a social field experiment as an experiment that aims to (1) assess the impact of (2) a political intervention by (3) collecting data on such political intervention. In order to properly collect data, the participants of the experiment are (4) randomly assigned into at least two groups. In one of these two groups the participants receive the political intervention that is being assessed and the other group do not. Such a design allows the comparison of the two groups and thus allows the impact of the intervention to be identified. The random assignment also provides another advantage: it removes the selection bias that confounds results. The evaluation of public policies is the focus of such social experiments. Indeed, the primary objective of these experiments is to “speak to policy-makers” ((Greenberg and Shroder, 2004, 4). Greenberg and Shroder’s (2004) definition exclude many experiments conducted in economics which do not satisfy (2) and (3).

Ferber and Hirsch (1982) define field social experiments in a very similar manner to Greenberg and Shroder (2004) but clearly relate them to the domain of economics, suggesting that a social field experiment evaluates the impact of both economic and social political intervention. A social field experiment is defined by four characteristics in the definition, namely, (1) public funding; (2) a rigorous statistical design; (3) targeting human subjects; (4) assessment of the effect of a political intervention

1
2
3
4
5 on social and economic variables.⁴
6

7 Common characteristics in these two definitions are rigorous statistical design
8 and the evaluation of political intervention, which enable proper assessment of the
9 impact of public policy. Indeed, Ferber and Hirsch (1982) underline the fact that
10 although other methodologies were available to evaluate public policy during the
11 second half of the 20th century, they were not as reliable as field experiments, and
12 especially in economics, because of the many unobservable variables. Thus, social
13 field experiments appear to be a very powerful tool for economics because of their
14 potential to control for unobservable variables. The major difference between the two
15 definitions is that Ferber and Hirsch (1982), but not Greenberg and Shroder (2004),
16 insist on public funding, highlighting the fact that such experiments are extremely
17 costly because of their large scale interventions.
18

19 During the second half of the 20th century, labor economics constituted a major
20 domain in economics where field experiments were implemented. During the Reagan
21 administration in 1986, a vast job program was also evaluated through a field experi-
22 ment, which was also randomized: the Job Training Partnership Act Title II-A. This
23 program offered financial support and training to people facing important barriers
24 on the job market (such as low qualifications). This experiment is the paradigm case
25 used by Heckman et al. (1997) to question the power of social evaluations in contrast
26 to structural econometrics. Since then, more than 235 social field experiments eval-
27 uating public policy programs have been conducted over the second half of the 20th
28

29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47 ⁴“A publicly funded study that incorporates a rigorous statistical design and whose experimental
48 aspects are applied over a period of time to one or more segments of a human population, *with the*
49 *aim of evaluating the aggregate economic and social effects of the experimental treatments.*” (Ferber
50 and Hirsch, 1982, 7, our emphasis)
51

1
2
3
4
5 century (Greenberg and Shroder, 2004) offer an overview of all these experiments).
6
7

8 **2.2.2 Randomized field experiments for evidence-based policy (2000-)** 9

10
11 During the turn of the century, field experiments began to be advocated not only
12 as a useful tool for policy evaluation but as a methodologically rigorous approach
13 to produce evidence for policy in general. As noted above, Greenberg and Shroder
14 (2004) and Ferber and Hirsch (1982) defined one of the key characteristics of social
15 field experiments as a random assignment of the experimental subjects into at least
16 two groups. Such assignment enables control for both observable and unobservable
17 variables, thereby offering a reliable evaluation of the impact of policies. During
18 the 1990s this reliable evaluation was highlighted, following the trend first initiated
19 within the movement of evidence-based medicine.
20
21
22
23
24
25
26
27
28

29 The random assignment (one of the main characteristics defining a randomized
30 field experiment) allows for reliable results by removing key statistical biases. Specif-
31 ically, an experiment needs to control for both observable and unobservable con-
32 founders to offer reliable results. Even carefully designed laboratory experiments
33 (see 2.1) cannot control every variations, for the simple reason that the experimenters
34 cannot foresee every possible variables that might affect the results. While the LFEs
35 strand tries to achieve controlled variation through the laboratory methods such as
36 Induced Value Theory, the RFEs strand adopts an indirect control through random-
37 ization as the central method. Random assignment of the participants to at least
38 two groups removes or balances selection bias.⁵ Biases concern the possibility that
39
40
41
42
43
44
45
46
47
48

49 ⁵See again here the discussion of Heckman (1992) for example.
50
51
52

1
2
3
4
5 other factors than the one the experiment is testing (e.g. the treatment) would ex-
6 plain the experimental results (e.g. the unmeasured differences between the groups).
7
8 In other words, randomization allows the experimenter to isolate only the effect of
9 the treatment. Thus randomized field experiments possess a strong internal and
10 statistical conclusions validity. Such strong internal validity gives RFEs' results the
11 status of evidence,⁶ leading, as we have seen in the previous section, to the so-called
12 evidence-based movement.
13
14
15
16
17
18

19 In 1992, a seminal paper was published by a group of thirty researchers who
20 labelled themselves the evidence-based medicine working group (Guyatt et al., 1992).
21 The paper argued that advances in medicine have been substantial thanks to the
22 recent massive use of randomized controlled trials (also using random assignment
23 as the key aspect of the experimental design); however, such scientific advancement
24 was not used in daily medical practice. The claim of the working group was that the
25 practice of medicine should be based on evidence provided by randomized controlled
26 trials. Randomized experiments are thus used and justified during this period for
27 two concomitant reasons: (1) the fact that they produce evidence, (2) the fact that
28 this evidence can directly help decisions-makers (here physicians).
29
30
31
32
33
34
35
36
37
38

39 At the beginning of the 21st century, the social sciences followed the same trend,
40 using social field experiments to base policy decisions (instead of medical ones) on
41 evidence. The two key characteristics defined by both Greenberg and Shroder (2004)
42 and Ferber and Hirsch (1982), the random assignment and the policy evaluation aim,
43 are still present, but the emphasis has shifted from the latter to the former. Tony
44
45
46
47
48

49 ⁶for a discussion about such status see Cartwright and Duflo REF
50
51
52

1
2
3
4
5 Blair was one of the first to popularize the idea of and the term “evidence-based
6 policy” in 2000. The United States and Australia also played a major role in the de-
7 velopment of evidence-based policy. As this trend shows, the evidence-based policy
8 movement is prominent in the Anglophone part of the world, whereas its introduc-
9 tion was later and smaller in other parts of the world, for example in France, even
10 though some randomized experiments were and are still being conducted. RFEs
11 are extensively used in development economics in particular, since the creation of
12 the Jameel Abdul Latif Poverty Action Lab (J-PAL) at the Massachusetts Institute
13 of Technology (MIT) in 2003 by Abhijit Banerjee, Esther Duflo and Sendhil Mul-
14 lainathan. J-PAL, one of the main centers of evidence-based policy movement, aims
15 to systematically promote the implementation of RFEs in order to provide evidence
16 to policy-makers of developing countries about the efficacy of development programs.
17
18

19 On the one hand, field experiments were introduced in economics to overcome
20 the limitations of lab experiments, by offering them more “natural” settings along
21 different dimensions than in the lab. On the other hand, field experiments entered
22 economics as a tool to assess policy interventions and then as the main tool to produce
23 evidence to guide policy-makers. The history of field experiments in economics, we
24 suggest, becomes clear once these two strands are clearly distinguished. These two
25 historical strands also create two methodological rationales for distinct types of field
26 experiments, lab-in-the-field experiments (LFEs) and randomized field experiments
27 (RFEs). These two methodological rationales originate in two distinct approaches
28 to controlled variation, a direct one for the LFEs and an indirect one for the RFEs.
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

3 Two faces of external validity

The principal methodological concern for experiments in general is that of their validity. Commonly, the validity of experiments is defined as twofold: internal and external. Internal validity refers to the validity of an experiment's result within the frame of that experiment. Internal validity thus refers to the level of control one experimental frame allows; the more control, the stronger the internal validity. In contrast, external validity relates to the validity of the experiment's result outside the frame of the experiment, and it has been discussed more extensively by philosophers and practitioners alike.

Our aim in this section is to illuminate this methodological debate concerning the external validity of field experiments in economics by drawing on the distinction between the two field experiment strands we have defined in the last section. Specifically, we argue that the two strands are concerned with two distinct types of external validity issues: what is at stake concerning the external validity for the RFE strand is the possibility to *generalize* the experiment's result to another setting, population, context, etc. This challenge comes from the indirect way in which RFEs achieve control. We call this the issue of *generalizability* (3.1); in contrast, the LFEs strand, with its strong connection to the lab and thus a high degree of direct control over variables, creates an experimental frame that can be characterized as *artificial* relative to the naturally-occurring decision-making contexts which the experiment intends to inform. This concerns the issue of *artificiality*. Artificiality in this sense diminishes the external validity of the results of lab experiments, but for distinct reasons (3.2).

3.1 Validity concepts in experimental social science

As we have discussed above, the notion of validity is commonly used as a methodological criterion to evaluate experiments. Internal and external validity concepts were first developed during the 19th century in psychology as the “accuracy of a [psychological] test” (Heukelom, 2011). During the 1950s they came to mean the accuracy of the experiments conducted in psychology. In the middle of the 1950s the American Psychological Association (APA) aimed to specify the methodological standard for psychology and produced a textbook with methodological recommendations for doing research in psychology. Donald Campbell (1957) opposed those recommendations, which led to several debates between Campbell and the APA that were concluded by a settled definition of internal and external validity in psychology.⁷ The resulting definitions were immediately adopted within the RFE strand. The historical and methodological continuity between psychology and the RFE strand might explain this immediate conceptual adoption. Building on the seminal definition, Shadish et al. (2002) distinguish four types of validity: statistical conclusions validity, internal validity, construct validity, and external validity.

1. Statistical conclusions validity refers to the existence of a statistical variation and its degree; it asks “whether the presumed cause and effect covary and how strongly they covary” (Shadish et al., 2002, 42).
2. Internal validity shows that the observed statistical variation is a causal relation

⁷“The discussions of validity in the Technical Recommendations and by Campbell between the early 1950s and early 1960s defined the experimental vocabulary in the psychology for subsequent decades, up into the twenty first century.” (Heukelom, 2009, 6)

1
2
3
4
5 and not a correlation and that this statistical variation is due only to the
6 treatment that is being assessed through the experiment.⁸
7
8
9

10 3. Construct validity concerns the validity of measurement, in particular the la-
11 beling of the operation of an experiment. In order to study a phenomenon it
12 is necessary to link its operationalization and the label; for example, to study
13 intelligence, it is necessary to operationalize the notion using a certain psycho-
14 logical test. Such a test of intelligence is said to be constructively valid if it
15 really measures the intelligence that we are interested in, and not something
16 else: “construct validity involves making inferences from the sampling partic-
17 ular of a study to the higher-order constructs they represent” (Shadish et al.,
18 2002, 65).
19
20
21
22
23
24
25
26
27

28 4. Finally, external validity is about the generalization of a causal inference from
29 the context of the experiment to another one, and-or from the population of the
30 experiment to another one, and-or from the treatment tested in the experiment
31 to another one.
32
33
34
35
36

37 These four types of validity are deeply linked to each other, in particular (1) and
38 (2), and (3) and (4). (1) Statistical conclusion validity and (2) internal validity are
39 both concerned with the statistical variation between the treatment that is tested
40 through the experiment and the experimental result. They concern two different
41
42
43
44

45 ⁸ “[W]e use the term internal validity to refer to inferences about whether observed covariation
46 between A and B reflects causal relationship from A to B in the form in which the variables were
47 manipulated or measured. To support such an inference, the researcher must show that A preceded
48 B in time, that A covaries with B (already covered under statistical conclusions validity) and that
49 no other explanations for the relation are plausible.” (Shadish et al., 2002, 53)
50
51
52

1
2
3
4
5 steps of one process: the former is about the validity of the inference that a statistical
6 variation exists, and the latter concerns the validity of the inference that a variation
7 is a causal one. Internal validity thus presupposes statistical conclusion validity.
8
9 Similarly, (3) construct validity and (4) external validity are deeply linked. While the
10 former concerns the validity of the degree of generalization from particular research
11 operation *to* higher-order constructs, the latter concerns the validity of the degree of
12 replication *across* different constructs and populations. To use a spatial metaphor,
13 the former regards *vertical* extrapolation, the latter *horizontal*. The term “external
14 validity” in the literature of experimental economics does not distinguish between
15 the two, probably due to the general disciplinary distance between economics and
16 psychology.⁹ But the distinction is relevant to understanding the methodological
17 challenges for both strands of field experiments, as we shall see below.
18
19
20
21
22
23
24
25
26
27
28
29
30

31 **3.2 External validity as the problem of generalizability of** 32 **randomized field experiments** 33 34 35

36 As we saw in the previous section, while the internal validity of RFEs seems to be
37 strong, their external validity has been extensively criticized from several perspec-
38 tives, includign the philosophy of science (Cartwright, 2007, 2009, 2010; Cartwright
39 and Hardie, 2012; Teira, 2013; Teira and Reiss, 2013; Davis, 2013), development eco-
40 nomics (Deaton, 2010; Ravallion, 2009; Rodrik, 2008; Acemoglu, 2010; Basu, 2014;
41 Barrett and Carter, 2010), experimental economics (Harrison, 2011, 2013), labor
42
43
44
45
46
47

48 ⁹In contrasting the two disciplines, Ross (2019) points out that “[t]he language of ‘constructs’ is
49 foreign to economists; referential terms in their models are taken to be directly isomorphic to real
50 objects and processes.”
51

1
2
3
4
5 economics and econometrics (Heckman, 1992; Heckman and Smith, 1995; Heckman
6 et al., 1997, 1998; Leamer, 2010). Even though these criticisms target different
7 methodological, practical and technical aspects of RFEs, they all somehow concen-
8 trate on the fact that RFEs produce a result termed “black-box test of ‘what works’”
9 (Deaton, 2010, 451). In other words, RFEs’ results would not illuminate the mech-
10 anisms behind such results, making their possible evidential use difficult or even
11 impossible. Indeed, both observables and unobservables that could offer explana-
12 tions about *why* a treatment is working are masked by the random allocation of the
13 treatments between the groups. For instance, an experimental result highlighting the
14 efficacy of deworming in a specific region in Kenya (see Miguel and Kremer, 2004)
15 to reduce students’ absenteeism might turn out otherwise in India. Indeed, in India
16 students might not suffer from worms but rather from anemia (Bobonis et al., 2006),
17 therefore iron would be effective while deworming would be not. Not knowing the
18 mechanisms behind the experimental result (e.g., efficacy of deworming) does not
19 tell us if such results will hold up in another context. Therefore, it appears difficult
20 to *generalize* from the RFE results, and thus to produce policy recommendations be-
21 yond particular experimental set-ups. In order, to favor such generalizability, J-PAL
22 researchers, such as Abhijit Banerjee and Esther Duflo, propose to concentrate on
23 “structured speculative judgements” (Banerjee et al., 2017), or to combine machine
24 learning and RFEs Chernozhukov et al. (2018). However, such propositions are still
25 at an early stage of development.
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

3.3 External validity as the problem of artificiality of lab experiments

As Heukelom (2011) points out, the experimental economists who established the lab tradition did not refer to the term “external validity” until the 1990s. Heukelom (2011) suggests that experimental economists’ reluctance to adopt the methodological distinction between internal and external validity is due to their fear that it would have pushed experimenters in economics to distinguish between an “inside” and an “outside” world.¹⁰ As Heukelom (2011) points out, a series of methodological analyses overcame this reluctance and made it a key methodological concept for laboratory experiments in economics (Guala, 1999, 2003, 2005; Guala and Mittone, 2005). Indeed, Guala explicitly defines internal and external validity in terms of artificiality, highlighting the distinction between an “inside” and an “outside” world.

Experimental economists have acknowledged and analyzed this gap between the lab and the naturally occurring economic decision-making set-up as *artificiality*. In particular, Bardsley et al. (2010) define four types of artificiality: the artificiality of (i) isolation, (ii) omission, (iii) contamination, and (iv) alteration. (Bardsley et al., 2010, 215). The development of the lab-in-the-field experiments (LFEs) strand can be seen as a response to these artificiality concerns.

Regarding the artificiality of isolation and omission, this is an inevitable consequence of experimental control to study one putative cause at a time, in contrast

¹⁰“The reason that [Vernon] Smith and other experimental economists in the 1970’s such as Plott were reluctant to adopt validity as understood by the psychologists was that the psychologists’ way of validity risked creating a division between an inside world of the laboratory and an outside ‘real’ world of the economy and its actors” (Heukelom, 2011, 20)

1
2
3
4
5 to the natural environment in which multiple causes are at work at any given time.
6
7 As we have seen, the solution proposed by LFEs is to study a putative cause while
8
9 reducing isolation and omission in a controlled way. For example, lab phenomena
10
11 such as gift-exchange or the endowment effect were first studied in the conventional
12
13 labs, and then in the LFEs by varying referent goods, locations and situations closer
14
15 to the field counterparts we are interested in.
16

17 The artificiality of contamination refers to the possibility that the participants'
18
19 very awareness of being studied will introduce new causes, such as a wish to confirm
20
21 to or disrupt the perceived goal of the experimenters.¹¹ What Harrison and List
22
23 (2004) call natural field experiments address this problem by making participants
24
25 unaware of the fact that they are being studied.
26

27 While these types of artificiality are methodologically well-addressed, at least in
28
29 principle, the less discussed artificiality of alteration may pose a bigger challenge
30
31 to the LFEs as it refers to the gap between the lab and the target phenomenon of
32
33 interest, which cannot be filled in an incremental way. In other words, going to the
34
35 field with standardized experiments does not solve this problem, but already assumes
36
37 that the alteration does not happen. The alteration refers to the possibility that
38
39 the experimental setup changes the nature of the phenomenon that the experiment
40
41 is designed to study. For example, do behavioral experiments to study risk, social
42
43 and time preferences represent naturally occurring situations in which people make
44
45 relevant risky, interpersonal, and intertemporal choices? In other words, do risk, so-
46
47 cial and time preferences measured in the experiments measure the same preferences
48

49 ¹¹See Jimenez-Buedo and Guala (2016) for a detailed discussion of contamination as reactivity.
50
51
52

1
2
3
4
5 that operate in the wild? We will discuss this challenge in detail with examples in
6
7 the next section.

8
9 The external validity problem thus construed as the artificiality of alteration is,
10
11 we argue, distinct from the external validity generalizability problem that the RFEs
12
13 strand faces. To use Cook et al.'s (2002) distinction that we introduced earlier, arti-
14
15 ficiality by alteration concerns (3) construct validity while generalizability concerns
16
17 (4) external validity. Although both concerns the validity of extrapolative infer-
18
19 ences, the grounds for extrapolation are different. Construct validity is grounded in
20
21 the assumption that the psychological or economic construct of beliefs or preferences,
22
23 operationalized by a particular elicitation procedure of an experiment and concep-
24
25 tualized or identified by a standard theoretical model, represents what it intends to
26
27 represent. External validity, in contrast, is grounded in a more specific assumption
28
29 about the comparison between the experimental model and the target phenomenon.
30
31 Although in this paper we use the term external validity in a broader sense to re-
32
33 fer to both generalizability and artificiality problems, thus departing from Shadish
34
35 et al. (2002), the distinction is important for the methodology of field experiments
36
37 in economics.

38
39 In this section, we have argued that the RFEs and LFEs strands face two distinct
40
41 challenges, generalizability and artificiality, which are both discussed under the term
42
43 external validity in the literature in experimental economics. We suspect that this is
44
45 part of the reason why philosophers contest the very usefulness of the notion of ex-
46
47 ternal validity (Jiménez-Buedo, 2011; Reiss, 2018). Our proposal has been to make a
48
49 methodologically relevant historical distinction explicit, before giving up on external
50
51

1
2
3
4
5 validity on conceptual grounds. In the next section, we propose a methodological
6 strategy to address both concerns, while making the two strands complement each
7 other in contemporary practices of field experiments in economics.
8
9
10

11 12 13 **4 Integrating the two strands** 14 15

16
17 Despite their distinct historical origins, the two strands of field experiments are
18 gradually becoming mixed in practice, resulting in hybrid field experiments (Viceisza,
19 2016). What are the affordances of such hybrid experiments? In particular, how do
20 they address the two distinct methodological issues defined in the previous section?
21 We argue that there is some methodological complementarity between RFEs and
22 LFEs, in particular the latter informing the data-generating process of the former.
23 However, although LFEs become less artificial by being embedded in RFEs, the
24 former's construct validity remains a problem.
25
26
27
28
29
30
31

32
33 To provide a methodological strategy in which both strands play complementary
34 roles, we draw on the notion of *exhibits* introduced by Robert Sugden (4.1). We then
35 show how this strategy combines both strands to address respective external validity
36 concerns and inform both theory and policy (4.2). Finally, we illustrate our point by
37 examining a recent field experimental study (4.3).
38
39
40
41
42
43

44 **4.1 Three functions of economic experiments as exhibits** 45 46

47 Roth (1995) distinguishes three functions of economic experiments, namely *informing*
48 *theory*, *searching for facts*, and *informing policy*. In terms of this three-part catego-
49
50
51

1
2
3
4
5 rization of economic experiments, the two strands of field experimental traditions—
6 the ‘lab-in-the-field’ experiments and randomized field experiments—are strongly
7 associated with the first and the last functions, respectively. But we propose that
8 a key to understanding the different external validity concerns of both strands in a
9 coherent framework and realizing their complementary roles is to pay more attention
10 to the second, fact-searching function. For this purpose, we draw on the notion of
11 *exhibits* developed by Sugden (2005, 2008) and his collaborators (Bardsley et al.,
12 2010).¹² An exhibit is “a replicable experimental design that reliably produces some
13 interesting result” (Bardsley et al., 2010, 156). Examples include famous anoma-
14 lies in behavioral experiments, such as Allais and Ellsberg paradoxes in judgements
15 and decision-making under uncertainty. There are three things to highlight in this
16 definition of an exhibit.

17
18
19 First, an exhibit as an experiment produces a phenomenon, which is distinct from
20 what occurs naturally in the world.¹³ In other words, the produced phenomenon is
21 an experimental artefact. Daniel Kahneman calls such an experimental invention as
22 an act of ‘bottling phenomena’, crediting the social psychologist Lee Ross for coining
23 the expression (Andersson and Holm, 2002, 45).

24
25
26 Second, an exhibit has to produce a phenomenon in a reliable and replicable man-
27 ner. That is, the experimental design is expected to produce the same phenomenon,
28 within a range of design variables. Replicability makes the bottled phenomenon dis-

29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

¹²Several authors have developed related notions such as paradigmatic experiments (Guala, 2008b; Guala and Mittone, 2010) building on exhibits. For our purposes, however, we mainly use Sugden’s original term.

¹³In the following literature some define an exhibit as a phenomenon. We follow Sugden’s original definition and distinguish the exhibit as an experiment and the phenomenon it produces.

1
2
3
4
5 tinct from a 'mere artifact', which is thought to be caused by some idiosyncratic
6 features of the experiment and therefore unworthy of further investigation. In con-
7 trast, the reliability of an exhibit to produce the same phenomenon motivates exper-
8 imenters to vary the original design systematically to investigate its regularity and
9 mechanisms.
10
11
12
13
14

15 Third, the resulting phenomenon is characterized as 'surprising' or 'interesting'
16 to the relevant research community. Often the surprise comes from the fact that
17 an exhibit constitutes an anomaly to the standard theory, as in the case of Allais
18 paradox violating the independence axiom of expected utility theory, or the fair offer
19 in the ultimatum game contradicting the sub-game perfect equilibrium prediction of
20 minimum offer. However, a phenomenon can be interesting or surprising without
21 clearly contradicting the prediction. For example, in finitely repeated public goods
22 games with the voluntary contribution mechanism, contributions typically start off
23 around 50% in the initial round, and then steadily decline toward the last round
24 (Ledyard, 1995). This pattern seems to confirm the free-riding hypothesis (subjects
25 are learning how to play rationally), or disconfirm it (the contribution usually does
26 not reach zero even in the final round), depending on what one sees in the phe-
27 nomenon. Either way, the phenomenon has interested a generation of experimental
28 economists, who have conducted numerous experimental variations and offered differ-
29 ent explanations. Thus, although surprise or interest sounds subjective, it is defined
30 against some shared or contested background knowledge in the research community.
31 A phenomenon may also simply be 'unexplained by' a received theory, in the sense
32 that the received theory does not take into account its alleged causal mechanism,
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52

1
2
3
4
5 in which case the phenomenon neither confirms nor disconfirm the theory. The fact
6
7 that the initial contributions in public goods games are typically around 50% is an
8
9 example of such phenomena.

10
11 What are exhibits good for, then? Sugden proposes that, in addition to challeng-
12
13 ing existing theories and helping in the construction of new ones (i.e., theory-testing
14
15 and building), exhibits serve at least three distinct functions in empirical investiga-
16
17 tions.

18
19 First, they serve as a “public arena in which empirical insights can be presented
20
21 and discussed, structured so that the more promising insights tend to have more
22
23 prominence” (Sugden, 2005, 298). This social-epistemic function has two mecha-
24
25 nisms. First, an exhibit provides an inter-subjective reference phenomenon that
26
27 other scientists can verify or challenge, and that the subsequent research can refer to
28
29 (Bardsley et al., 2010, 163), which we call coordination. Second, Sugden (2005) also
30
31 conjectures that an exhibit may attract many rival explanations for the phenomenon
32
33 when multiple causal mechanisms are at work *in the same direction* to create the
34
35 phenomenon (attention). In other words, while there are many ways to coordinate
36
37 research efforts, exhibits may tend to facilitate coordination on causally important
38
39 phenomena. On this account, a prominent exhibit will help produce and sharpen
40
41 alternative explanations and enrich our causal understanding of that relevant phe-
42
43 nomenon, without necessarily resulting in a replacement of the standard theory by
44
45 some unified theory that wins the competition. Still, this process facilitates empirical
46
47 progress in two ways. On the one hand, an exhibit’s ability to produce a robust regu-
48
49 larity can be examined through a range of changes in experimental design (robustness
50
51
52
53
54
55
56
57
58
59
60

1
2
3
4
5 check). On the other hand, an exhibit can be further broken down to finer-grained
6 exhibits that demonstrate regularities (decomposition) (Bardsley et al., 2010).
7
8

9 A second function of an exhibit is to serve as a *construct*, or an experimental
10 measure or representation of some feature of real-world behavior (Sugden, 2008, see
11 also (Guala, 2008b; Guala and Mittone, 2010)). Shadish et al.'s (2002) discussion
12 of construct validity (see Section 3.1 above) is useful in explicating this function.
13 When an exhibit produces a valid construct, the latter represents some higher-order,
14 more *abstract* disposition of a given population. Sugden's example is 'trust' in the
15 trust game.¹⁴ In the trust game, a varying level of 'trust' is measured as the ratio
16 of the endowment transferred from the Truster to the Trustee. The level of 'trust'
17 is expected to vary across populations, capturing their different levels of trust. In
18 other words, valid constructs permit *vertical* extrapolation to the more abstract
19 target phenomena, rather than *horizontal* extrapolation to other concrete situations.
20 How do we increase the validity of a construct? Guala (2008b) emphasizes the
21 importance of the *standardization* of experimental design, or the establishment of
22 what he calls 'paradigmatic experiments'. Clearly, however, standardization as such
23 is not sufficient because assuming so would imply that each standardized operation
24 has its own construct, which undermines the attempt to let the procedure measure
25 a higher-order construct. As we discussed, this worry, the artificiality of alteration,
26 cannot be addressed by merely moving the standardized experiments to the field

27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45 ¹⁴The trust game (Berg et al., 1995) is a two-person sequential game in which the first mover
46 (Truster) will decide the proportion of transfer of the initial endowment, which will be multiplied by
47 a certain factor; the second mover (Trustee) then decides the proportion to return. Although trust
48 and reciprocal cash transfer can benefit both players, since the Trustee can maximize his income
49 by not returning anything, and since the Truster knows this, the sub-game perfect equilibrium is
50 no transfer, or no trust.
51

1
2
3
4
5 (LFEs). What is additionally needed is the more or less *convergence* of the measures
6 from these experiments, as well as between these measures and other indicators of
7 the construct coming from various non-experimental data on a given population,
8 such as survey data about trust in the government, prevalence of trustful behavior
9 in the field, etc. In this sense, LFEs are not a sufficient but necessary step toward
10 building a valid construct.
11
12
13
14
15
16

17 Finally, exhibits serve as analogous surrogates for ‘real-world’ phenomena to be
18 compared, in particular behavior in non-experimental settings. This means that by
19 studying the causal mechanism of the exhibit we aim to learn about that of the anal-
20 ogous real-world behavior in question. Sugden’s example is a lab coordination game,
21 to which some observable real-world regularity, namely the tendency for trade deals
22 to be struck at round numbers (Schelling, 1960, 67), is compared. To reason that
23 the lab experiment and the real world regularity share some relevant mechanism is a
24 type of inductive reasoning called analogical reasoning (see Guala, 2010; Steel, 2010).
25 Although Sugden is very cautious about identifying the mechanism of an exhibit with
26 that of a target real-world phenomenon, he does not exclude the possibility of us-
27 ing exhibits as “explanatory devices in their own right” (Sugden, 2005, 298) in this
28 way. Since analogical reasoning is fallible, several methodological approaches have
29 been proposed to make it successful, the most prominent one being comparative
30 process-tracing (Steel, 2008).¹⁵ In what follows, we highlight complementary roles of
31 LFEs and RFEs, focusing on which functions each type of experiments serve.
32
33
34
35
36
37
38
39
40
41
42
43
44
45

46
47 ¹⁵See Khosrowi (2019) as a critical evaluation of comparative process-tracing as a research strat-
48 egy in econometrics.
49
50
51
52
53
54
55
56
57
58
59
60

4.2 Combining RFEs and LFEs as exhibits

How can RFEs and lab-in-the-field experiments (LFEs) be construed and combined as exhibits? In this subsection we present a schematic way to do so, which will be illustrated with an actual experimental study in the following sub-section.

Table 2 summarizes the three functions of exhibits we discussed above in column, applied to RFEs and LFEs in row.

Functions of exhibits	Coordinate (Attention)	Construct (Abstraction)	Compare (Analogue)
Randomized field experiments (RFEs)	①	-	③
Lab-in-the-field experiments (LFEs)	-	②	-

Table 2: Functions of field experiments as exhibits: Circled numbers indicate the step-wise hybrid use of RFEs and LFEs that we advocate.

First, we discuss whether and how RFEs perform the three functions as exhibits, in comparison to lab experiments. Then, we argue that, when combined, RFEs and lab experiments, in particular LFEs, complement each other in terms of these functions.

First of all, RFEs, just like lab experiments, fulfill the social-epistemic function as research coordination devices (① in Table 2). That is, some RFEs can attract research efforts by their prominence by virtue of their surprising character (against some theoretical background), replicability and accommodation of multiple causal explanations. For example, the exhibit that the poor do not invest enough in cheap preventive health products, whereas they should if they follow rational choice models (Dupas, 2014) illustrates such a surprise. When such behavior is observed not only

1
2
3
4
5 for bednets, but also for other products such as deworming pills and chlorine, the
6 exhibit is more replicable, suggesting some regularity in health-related behavior in
7 general. And when the exhibit invites multiple explanations (present-bias, social
8 pressures from family and relatives, credit constraints, etc.), its prominence increases.
9 Another example is an experiment that shows that behavioral intervention that the
10 received theory considers to be of no relevance turns out to have a significant impact
11 on behavior, which also invites more experiments and explanations.
12
13
14
15
16
17
18

19 Concerning the second function of exhibits, namely to serve as a construct, RFEs
20 are not suited for this purpose. This is because they are not as standardized as lab
21 experiments, which is a necessary feature for them to function as portable measure-
22 ment devices of particular constructs (Guala, 2008b). RFEs are primarily designed
23 to measure the impact of interventions, not the level of trust or altruism in a given
24 population. Lab-in-the-field experiments, in contrast, are intended to function this
25 way, as their standardized experimental designs are implementable in different lo-
26 cations with different populations ((2) in Table 2). When a LFE measures a valid
27 construct, it can be used to tell us whether and to what extent a given population
28 is heterogeneous in terms of the level of the construct it measures. As we will see
29 below, this information facilitates the use of RFEs as surrogates for counterfactual
30 scenarios of policy interventions.
31
32
33
34
35
36
37
38
39
40
41
42

43 Regarding the third function of serving as an analogue to be directly compared to
44 the real world target system, RFEs are intended to serve as such surrogates, as they
45 are intended to demonstrate what works in the real policy contexts ((3) in Table 2).
46 However, as we have seen in the last section the main concern with RFEs is their
47
48
49
50
51
52

1
2
3
4
5 generalizability. In particular, it is difficult to soundly infer that what happens in
6 an RFE is going to happen in an analogous counterfactual policy scenario, in which
7 a similar intervention in another sample, time, scale, or location is implemented.
8
9
10
11 This use of RFEs as surrogates is more demanding than the case Sugden considers,
12 in which a lab experiment is compared to some observable real-world regularity.
13 Unlike this use of a lab experiment as an explanatory device, a RFE is intended
14 to function as a simulation device that tells us what would happen if a particular
15 policy intervention would be implemented in a particular setting. This demands
16 more careful and detailed confirmation of the relevant analogy between the RFE and
17 its implementation in the intended context. In the next sub-section, we will illustrate
18 how the use of LFEs as measures of constructs will help such analogical reasoning.
19
20
21
22
23
24
25
26
27
28

29 **4.3 Illustrations**

30
31 Our example comes from studies on health-related purchasing/saving behavior in
32 developing countries. As underlined above, RFEs have suggested that the poor
33 do not invest in preventive health products although they are relatively affordable;
34 instead, they often end up spending a substantial amount of money on curative
35 products, while incurring the cost of diseases (see Dupas, 2011, for a review). RFEs
36 on bednets against malaria illustrates this more in detail. First, the demand for
37 bed-nets drops by 80% as soon as they cost some money (Cohen and Dupas, 2010,
38 20). Second, individuals who benefited from bed-nets for free in an initial period
39 are more willing to pay for bed-nets (Dupas, 2014). Third, the poor do not invest
40 in bed-nets, whereas they are intensively used when given out for free, which Dupas
41
42
43
44
45
46
47
48
49
50
51
52

(2009, 230) calls a “puzzling fact”.

We propose that this puzzling fact can be construed as an exhibit in the first sense. Indeed, it highlights a surprising fact contradictory to standard rational choice theory. It would be rational for individuals to invest in cheap and effective preventive health care (such as bed-nets) rather than to spend more on expensive curative health products. Such a fact is not only surprising but also regularly observed in many RFEs (Dupas, 2011). The fact that similar RFEs show similar patterns with other preventive health products such as deworming pills and chlorine tablets to purify water in other regions point to the regularity of the exhibit. By producing a surprising experimental phenomenon, these RFEs help coordinate research attention and efforts.

However, RFEs alone are unable to explain such results as they are like a black-box (e.g. Deaton, 2010). What explains the under-investment in products with apparent low cost and high expected returns? Mahajan and Tarozzi (2012) provide an excellent case in which this question was addressed in a large-scale RFE that evaluated alternative micro financial tools to provide insecticide treated nets (ITNs), and how they impact health and socio-economic outcomes of potential users in India.¹⁶

¹⁶The study is part of the broader project published as Tarozzi et al. (2014). We are focusing on this study for the following reasons: first, its topic, the ITNs use, is the same as Dupas’s RFEs that gave rise to the more general phenomenon that we have been discussing, namely the poor’s under-investment in preventive health products; second, we consider the methodological standard of the study to be high, even compared to other published studies. Although it is still a working paper, it has already been cited 55 times, and its larger study, the RFE part, (Tarozzi et al., 2014) has been published in *American Economic Review* and cited 162 times (both on Google Scholar, checked on 27th August 2019), suggesting that the study is also well-received in the research community. In general, in economics many influential and important texts are published in respected working paper series, followed by publications in journals a few years later. See Viceisza (2016) for a more systematic review of more examples of hybrid experiments, many of which are working papers as of 2016.

1
2
3
4
5 The motivation of the study was to find an effective micro-credit for the poor to
6 buy ITNs as an alternative to free distribution. Given that health organizations and
7 local governments also have budget constraints, free distribution of preventive health
8 products is not a sustainable policy option. Moreover, free distribution will target
9 the most credit-constrained households more effectively if other households share or
10 fully pay the cost of the products using existing affordable micro-credit schemes. In
11 order to devise such effective schemes, however, we need to understand why some
12 households choose particular micro loan scheme while others choose another, and yet
13 others do not purchase ITNs at all.
14
15
16
17
18
19
20
21
22

23 Mahajan and Tarozzi (2012) addressed this question by conducting a LFE inside
24 the RFE. In particular, the lab-in-the-field part of the experiment focused on what
25 explains the take-up rates of different micro-loan schemes to purchase ITNs. In
26 the RFE they gave the households two alternative micro-loan offers through a local
27 micro lending association, one with the service of future re-treatments of the net,
28 and the other without. The former contract has a feature of a commitment device,
29 binding the buyer to the payment of the future re-treatments, but with lower cost
30 than when the buyer later decides to pay for re-treatments. The households also had
31 the option to not purchase any net. So within this treatment group, the observed
32 choices were (i) purchase of a net with a commitment loan; (ii) purchase of a net
33 without a commitment loan; and (iii) no purchase.¹⁷ The question that the RFE
34 alone cannot answer is: what explains take-up rates and different choices between
35 (i) -(iii)?
36
37
38
39
40
41
42
43
44
45
46
47
48

49 ¹⁷The households could buy up to two ITNs.
50
51
52

1
2
3
4
5 To answer this question, the lab-like field experiment elicited these households'
6 time preferences, slightly modifying the standard lab version of the intertemporal
7 "preference reversal" experiment.¹⁸ The experimenters call this elicitation process
8 a 'baseline survey,' but the technique comes from lab experiments, and also the
9 elicitation of preferences are incentivized. Therefore in our exposition it is more
10 appropriate to call it a lab-in-the-field experiment (LFE).
11
12

13
14
15
16
17 In this experiment, the households were asked to choose between 12 binary inter-
18 temporal payments, e.g., 10 rupee to be paid one month later vs. 15 rupee four
19 months later; 10 rupee four months later vs. 15 rupee seven months later.¹⁹ This
20 design enabled the researchers to (i) identify time-inconsistent choice patterns, such
21 as choosing 10 rupee in the first set while choosing 15 rupee in the second set, and (ii)
22 associate this choice data to the choice data from the RFE concerning the purchase of
23 ITNs. In addition, the experimenters also elicited households' beliefs about the risk of
24 becoming infected with malaria within a year when sleeping regularly under a treated
25 net, an untreated net, and no net in cases where one is a child, a pregnant woman,
26 or an adult. Based on these data and the literature on intertemporal decision-
27 making, the researchers hypothesized three types of households: (1) time-consistent
28 agents; (2) time-inconsistent agents who are aware of their own future present-bias
29 ("sophisticated" time-inconsistent agents); and (3) time-inconsistent agents who are
30 unaware of their own future present-bias ("naive" time-inconsistent agents).
31
32
33
34
35
36
37
38
39
40
41
42
43
44

45 Using these types, the experimenters identified a full structural model of the data-

46
47 ¹⁸The experimenters modified some details based on contextual considerations. See Mahajan and
48 Tarozzi (2012, footnote 14, page 7)

49 ¹⁹The participants were told that one of the 12 chosen rewards would be paid by the local micro-
50 lending association.
51
52

1
2
3
4
5 generating process underlying the take-up choices in the RFE, and then estimated its
6 parameters using standard goodness of fit methods (maximum likelihood estimation
7 with the finite mixture approach). An important point is that not only was the data-
8 generating process of the RFEs have been fully modelled, but also the model was
9 constructed and estimated based on the LFE, which revealed the heterogeneity of
10 agent types and other characteristics. In other words, the LFE helped solve the so-
11 called identification problem by providing data about the subjects' risk preferences,
12 cost preferences, and beliefs regarding risks of malaria infection. Analysing these
13 data from the LFE, together with the choice data from the RFE in which realistic
14 services (nets with different loan instruments) are offered, enabled the researchers to
15 study which types of subjects chose which micro-credit loans.
16
17
18
19
20
21
22
23
24
25
26

27 One of the main findings of Mahajan and Tarozzi (2012) is that the mapping
28 between the subjects' types (40% of the subjects are consistent, 50% "naive" hy-
29 perbolic, and 10% "sophisticated" hyperbolic) and the field purchase patterns is not
30 straightforward. In particular, they found that the product with commitment (i.e.,
31 the loan with future re-treatment of the net included) was more popular among
32 wealthier (and "naive") households than among the "sophisticated" ones who, ac-
33 cording to the received view, should be the main target of such a product. They
34 also found that the belief in ITN efficacy and wealth predict the choice of the com-
35 mitment device better than being "sophisticated" does. This result challenges the
36 prominent behavioral economic idea that the poor's under-investment in preventive
37 health products is caused by hyperbolic time discounting.
38
39
40
41
42
43
44
45
46
47
48

49 As this example illustrate, a hybrid design combining LFEs and RFEs can ex-
50
51
52

1
2
3
4
5 exploit their complementarities as exhibits. First, a series of prominent RFEs have
6 exhibited the poor's under-investment in preventive health products as a prominent
7 phenomenon to be investigated further ((1) in Table 2). Second, an RFE has been
8 designed as an analogue of real-world interventions, in particular different micro
9 credit schemes, *together with* an LFE that measured households' relevant constructs
10 such as risk and cost preferences, as well as beliefs regarding the efficacy of bednets
11 ((2) in Table 2). Contingent on the validity of these constructs, these data then help
12 identify the data-generating processes underlying the RFE, in the form of a struc-
13 tural model, which is needed for explicit and quantitative evaluations of analogous
14 counterfactual interventions ((3) in Table 2).
15
16
17
18
19
20
21
22
23
24

25 Our proposed strategy is a schematic sketch of how research can proceed, so the
26 temporal order of the actual processes can vary. In particular, as we noted earlier, the
27 validity of experimental constructs cannot be verified simply by combining it with
28 RFEs, but instead it has to be gradually developed as the convergence of different
29 experimental and non-experimental measures of a given construct in the long run.
30 Further, the generalizability of the results from a RFE faces more challenges as
31 the intended context of application moves further away from the original RFE, for
32 example to different populations or to a larger scale, which cannot be fully addressed
33 by structural modelling. Despite all these challenges, we suggest that coupling RFEs
34 and LFEs is a concrete and promising step forward to strategically exploit low-
35 hanging methodological advantages of the two strands of field experiments.
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

5 Conclusion

In this paper, we argued that field experiments in economics come from two distinct historical strands, the first as an extension of laboratory experiments, which we labelled as lab-in-the-field experiments (LFEs), and the second strand defined by the implementation of randomized-control field experiments (RFEs), originating from social sciences more broadly. While the LFEs strand developed to decrease the artificiality that tight laboratory control can create, the RFEs strand tries to achieve control indirectly for both observable and unobservable variables. Thus the two strands face two distinct external validity issues: artificiality for the LFE strand and generalizability for the RFE strand. While both are serious challenges that cannot be fully addressed in this paper, we proposed a promising methodological strategy to couple RFEs and LFEs, building on Sugden's functional analysis of experiments as exhibits for empirical investigations. Although still in an early stage, several hybrid experiments in the same spirit have begun to appear (see Viceisza, 2016, Table 1). Our historical and methodological analysis clarifies why such hybridization is happening—because these strands have been distinct traditions in field experiments in economics—and why it is a good idea—because it can exploit methodological complementarities of both strands.

References

Acemoglu, D. (2010). Theory, general equilibrium, and political economy in development economics. *Journal of Economic Perspectives*, 24(3):17–32.

- 1
2
3
4
5 Akerlof, G. A. (1982). Labor Contracts as Partial Gift Exchange*. *The Quarterly*
6 *Journal of Economics*, 97(4):543–569.
7
8
9
10 Andersson, F. N. and Holm, H., editors (2002). *Experimental Economics: Financial*
11 *Markets, Auctions, and Decision Making: Interviews and Contributions from the*
12 *20th Arne Ryde Symposium*. Springer Science & Business Media, New York.
13
14
15
16
17 Banerjee, A., Chassang, S., and Snowberg, E. (2017). Decision theoretic approaches
18 to experiment design and external validity. In Banerjee, A. V. and Duflo, E.,
19 editors, *Handbook of Field Experiments*, volume 1 of *Handbook of Economic Field*
20 *Experiments*, chapter 4, pages 141 – 174. North-Holland.
21
22
23
24
25
26 Bardsley, N. (2008). Dictator game giving: altruism or artefact? *Experimental*
27 *Economics*, 11:122–133.
28
29
30
31 Bardsley, N., Cubitt, R., Loomes, G., Moffat, P., Starmer, C., and Sugden, R. (2010).
32 *Experimental economics: Rethinking the rules*. Princeton University Press.
33
34
35
36 Barrett, C. B. and Carter, M. R. (2010). The power and pitfalls of experiments
37 in development economics: Some non-random reflections. *Applied Economic Per-*
38 *spectives and Policy*, 32(4):515–548.
39
40
41
42 Basu, K. (2014). Randomisation, causality and the role of reasoned intuition. *Oxford*
43 *Development Studies*, 42(4):455–472.
44
45
46
47 Berg, J., Dickhaut, J., and McCabe, K. (1995). Trust, reciprocity, and social history.
48 *Games and Economic Behavior*, 10(1):122 – 142.
49
50
51
52

- 1
2
3
4
5 Bobonis, G. J., Miguel, E., and Puri-Sharma, C. (2006). Anemia and school partic-
6 ipation. *Journal of Human resources*, 41(4):692–721.
7
8
9
10 Bolton, G. E. and Ockenfels, A. (2012). Behavioral economic engineering. *Journal*
11 *of Economic Psychology*, 33(3):665–676.
12
13
14 Camerer, C. F. (2003). *Behavioral Game Theory*. Princeton University Press, Prince-
15 ton, NJ.
16
17
18
19 Campbell, D. T. (1957). Factors relevant to the validity of experiments in social
20 settings. *Psychological bulletin*, 54(4):297.
21
22
23
24 Cartwright, N. (2007). Are rcts the gold standard? *BioSocieties*, 2(1):11–20.
25
26
27 Cartwright, N. (2009). What is this thing called "efficacy"? In Mantzavinos, C.,
28 editor, *Philosophy of the Social Sciences: Philosophical Theory and Scientific Prac-*
29 *tice*, chapter 7, pages 185–206. Cambridge University Press, Cambridge, England.
30
31
32
33
34 Cartwright, N. (2010). What are randomised controlled trials good for? *Philosophical*
35 *Studies*, 147(1):59.
36
37
38
39 Cartwright, N. and Hardie, J. (2012). *Evidence-based policy: A practical guide to*
40 *doing it better*. Oxford University Press.
41
42
43
44 Charness, G., Gneezy, U., and Kuhn, M. A. (2013). Experimental methods: Extra-
45 laboratory experiments—extending the reach of experimental economics. *Journal*
46 *of Economic Behavior & Organization*, 91(0):93–100.
47
48
49
50
51
52

- 1
2
3
4
5 Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey,
6 W., and Robins, J. (2018). Double/debiased machine learning for treatment and
7 structural parameters. *The Econometrics Journal*, 21(1):C1–C68.
8
9
10
11
12 Cohen, J. and Dupas, P. (2010). Free distribution or cost-sharing? evidence from a
13 randomized malaria prevention experiment. *The Quarterly Journal of Economics*,
14 125(1):1–45.
15
16
17
18 Davis, J. (2013). Economics imperialism under the impact of psychology: The case
19 of behavioral development economics. *Economia*, 3(1):119–138.
20
21
22
23 Deaton, A. (2010). Instruments, randomization, and learning about development.
24 *Journal of Economic Literature*, 48(2):424–55.
25
26
27
28 DellaVigna, S. (2009). Psychology and economics: Evidence from the field. *Journal*
29 *of Economic Literature*, 47(2):315–72.
30
31
32
33 Duflo, E. (2006). *Field Experiments in Development Economics*, volume 2 of *Econo-*
34 *metric Society Monographs*, pages 322–348. Cambridge University Press.
35
36
37
38 Dupas, P. (2009). What matters (and what does not) in households' decision to
39 invest in malaria prevention? *American Economic Review*, 99(2):224–230.
40
41
42
43 Dupas, P. (2011). Health behavior in developing countries. *Annual Review of Eco-*
44 *nomics*, 3(1):425–449.
45
46
47
48 Dupas, P. (2014). Short-run subsidies and long-run adoption of new health products:
49 Evidence from a field experiment. *Econometrica*, 82(1):197–228.
50
51
52

- 1
2
3
4
5 Fehr, E., Kirchsteiger, G., and Riedl, A. (1993). Does Fairness Prevent Market
6 Clearing? An Experimental Investigation*. *The Quarterly Journal of Economics*,
7 108(2):437–459.
8
9
10
11
12 Ferber, R. and Hirsch, W. Z. (1982). *Social experimentation and economic policy*.
13 Cambridge University Press, Cambridge, England.
14
15
16
17 Gerber, A. S. and Green, D. P. (2012). *Field experiments: Design, analysis, and*
18 *interpretation*. WW Norton.
19
20
21
22 Gneezy, U., Haruvy, E., and Yafe, H. (2004). The inefficiency of splitting the bill*.
23 *The Economic Journal*, 114(495):265–280.
24
25
26
27 Gneezy, U. and List, J. A. (2006). Putting behavioral economics to work: Testing for
28 gift exchange in labor markets using field experiments. *Econometrica*, 74(5):1365–
29 1384.
30
31
32
33 Greenberg, D. H. and Shroder, M. (2004). *The digest of social experiments*. The
34 Urban Insitute Press, Washington, D.C., 3rd edition.
35
36
37
38 Guala, F. (1999). The problem of external validity (or “parallelism”) in experimental
39 economics. *Social science information*, 38(4):555–573.
40
41
42
43 Guala, F. (2003). Experimental localism and external validity. *Philosophy of Science*,
44 70:1195–1205.
45
46
47
48 Guala, F. (2005). *The Methodology of Experimental Economics*. Cambridge Univer-
49 sity Press, Cambridge, England.
50
51
52

- 1
2
3
4
5 Guala, F. (2008a). *Experimental Economics, History of*, pages 1958–1962. Palgrave
6 Macmillan UK, London.
7
8
9
10 Guala, F. (2008b). Paradigmatic experiments: The ultimatum game from testing to
11 measurement device. *Philosophy of Science*, 75:658–669.
12
13
14 Guala, F. (2010). Extrapolation, analogy, and comparative process tracing. *Philos-*
15 *ophy of Science*, 77(5):pp. 1070–1082.
16
17
18
19 Guala, F. and Mittone, L. (2005). Experiments in economics: External validity and
20 the robustness of phenomena. *Journal of Economic Methodology*, 12(4):495–515.
21
22
23
24 Guala, F. and Mittone, L. (2010). Paradigmatic experiments: The dictator game.
25 *The Journal of Socio-Economics*, 39(5):578–584.
26
27
28
29 Guyatt, G., Cairns, J., Churchill, D., Cook, D., Haynes, B., Hirsh, J., Irvine, J.,
30 Levine, M., Levine, M., Nishikawa, J., Sackett, D., Brill-Edwards, P., Gerstein, H.,
31 Gibson, J., Jaeschke, R., Kerigan, A., Neville, A., Panju, A., Detsky, A., Enkin,
32 M., Frid, P., Gerrity, M., Laupacis, A., Lawrence, V., Menard, J., Moyer, V.,
33 Mulrow, C., Links, P., Oxman, A., Sinclair, J., and Tugwell, P. (1992). Evidence-
34 Based Medicine: A New Approach to Teaching the Practice of Medicine. *JAMA*,
35 268(17):2420–2425.
36
37
38
39
40
41
42
43
44 Harrison, G. W. (2011). Randomisation and its discontents. *Journal of African*
45 *Economies*, 20(4):626–652.
46
47
48
49 Harrison, G. W. (2013). Field experiments and methodological intolerance. *Journal*
50 *of Economic Methodology*, 20(2):103–117.
51
52

- 1
2
3
4
5 Harrison, G. W. and List, J. A. (2004). Field experiments. *Journal of Economic*
6
7 *Literature*, 42(4):1009–1055.
8
9
- 10 Heckman, J., Ichimura, H., Smith, J., and Todd, P. (1998). Characterizing selection
11
12 bias using experimental data. *Econometrica*, 66(5):1017–1098.
13
14
- 15 Heckman, J. J. (1992). Randomization and social policy evaluation. In Manski, C. F.
16
17 and Garfinkel, I., editors, *Evaluating Welfare and Training Programs*, chapter 5,
18
19 pages 201–230. Harvard University Press.
20
- 21 Heckman, J. J., Smith, J., and Clements, N. (1997). Making the most out of pro-
22
23 gramme evaluations and social experiments: Accounting for heterogeneity in pro-
24
25 gramme impacts. *The Review of Economic Studies*, 64(4):487–535.
26
27
- 28 Heckman, J. J. and Smith, J. A. (1995). Assessing the case for social experiments.
29
30 *Journal of Economic Perspectives*, 9(2):85–110.
31
32
- 33 Henrich, J. P. (2004). *Foundations of human sociality: economic experiments and*
34
35 *ethnographic evidence from fifteen small-scale societies*. Oxford University Press,
36
37 Oxford.
38
39
- 40 Hertwig, R. and Ortmann, A. (2001). Experimental practices in economics: A
41
42 methodological challenge for psychologists? *Behavioral and Brain Sciences*,
43
44 24:383–+.
45
46
- 47 Heukelom, F. (2009). *Kahneman and Tversky and the Making of Behavioral Eco-*
48
49 *nomics*. PhD thesis, University of Amsterdam.
50
51
52

- 1
2
3
4
5 Heukelom, F. (2011). What to conclude from psychological experiments: The con-
6 trasting cases of experimental and behavioral economics. *History of Political Econ-*
7 *omy*, 43(4):649–681.
8
9
10
11 Heukelom, F. (2014). *Behavioral economics: a history*. Cambridge University Press,
12 Cambridge.
13
14
15
16 Jiménez-Buedo, M. (2011). Conceptual tools for assessing experiments: some well-
17 entrenched confusions regarding the internal/external validity distinction. *Journal*
18 *of Economic Methodology*, 18(3):271–282.
19
20
21
22 Jiménez-Buedo, M. and Guala, F. (2016). Artificiality, reactivity, and demand effects
23 in experimental economics. *Philosophy of the Social Sciences*, 46(1):3–23.
24
25
26
27 Khosrowi, D. (2019). Extrapolation of causal effects – hopes, assumptions, and the
28 extrapolator’s circle. *Journal of Economic Methodology*, 26(1):45–58.
29
30
31
32 Leamer, E. E. (2010). Tantalus on the road to asymptopia. *Journal of Economic*
33 *Perspectives*, 24(2):31–46.
34
35
36
37 Ledyard, J. (1995). Public goods: a survey of experimental research. In Kagel, J. H.
38 and Roth, A. E., editors, *The handbook of experimental economics*, pages 111–194.
39 Princeton University Press.
40
41
42
43 Levitt, S. D. and List, J. A. (2009). Field experiments in economics: The past, the
44 present, and the future. *European Economic Review*, 53(1):1 – 18.
45
46
47
48 List, J. A. (2007). On the interpretation of giving in dictator games. *Journal of*
49 *Political Economy*, 115(3):482–493.
50
51
52

- 1
2
3
4
5 List, J. A. and Lucking-Reiley, D. (2000). Demand reduction in multiunit auc-
6 tions: Evidence from a sportscard field experiment. *American Economic Review*,
7 90(4):961–972.
8
9
10
11
12 List, J. A. and Shogren, J. F. (1998). Calibration of the difference between actual
13 and hypothetical valuations in a field experiment. *Journal of Economic Behavior*
14 *& Organization*, 37(2):193 – 205.
15
16
17
18
19 Mahajan, A. and Tarozzi, A. (2012). Time inconsistency, expectations and tech-
20 nology adoption: The case of insecticide treated nets. WPS 019, University of
21 Californica Berkeley.
22
23
24
25
26 Miguel, E. and Kremer, M. (2004). Worms: Identifying impacts on education and
27 health in the presence of treatment externalities. *Econometrica*, 72(1):159–217.
28
29
30
31 Ravallion, M. (2009). Should the randomistas rule? *The Economists' Voice*, 6(2).
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
- Reiss, J. (2018). Against external validity. *Synthese*.
- Rodrik, D. (2008). The new development economics: We shall experiment, but
how shall we learn? *HKS Working Paper No. RWP08-055*. Revised, Octo-
ber 2008. There is a lot more convergence between macro- and micro-development
economists than meets the eye.
- Ross, D. (2019). Empiricism, sciences, and engineering: cognitive science as a zone
of integration. *Cognitive Processing*, pages <https://doi.org/10.1007/s10339-019-00916-z>.

- 1
2
3
4
5 Roth, A. E. (1995). Introduction to experimental economics. In Kagel, J. H. and
6
7 Roth, A. E., editors, *Handbook of Experimental Economics*, chapter 1. Princeton
8
9 University Press.
10
11
12 Schelling, T. C. (1960). *The strategy of conflict*. Harvard University Press.
13
14
15 Shadish, W., Cook, T. D., and Campbell, D. T. (2002). *Experimental and quasi-*
16
17 *experimental designs for generalized causal inference*. Houghton Mifflin Company,
18
19 Boston, New York.
20
21
22 Smith, V. L. (1976). Experimental economics: Induced value theory. *The American*
23
24 *Economic Review*, 66(2):274–279.
25
26
27 Steel, D. (2008). *Across the boundaries: extrapolation in biology and social science*.
28
29 Oxford University Press, Oxford.
30
31
32 Steel, D. (2010). A new approach to argument by analogy: Extrapolation and chain
33
34 graphs. *Philosophy of Science*, 77(5):pp. 1058–1069.
35
36
37 Sugden, R. (2005). Experiments as exhibits and experiments as tests. *Journal of*
38
39 *Economic Methodology*, 12(2):291–302.
40
41
42 Sugden, R. (2008). The changing relationship between theory and experiment in
43
44 economics. *Philosophy of Science*, 75(5):pp. 621–632.
45
46
47 Svorenčík, A. and Maas, H. (2015). *The making of experimental economics: Witness*
48
49 *seminar on the emergence of a field*. Springer.
50
51
52

1
2
3
4
5 Tarozzi, A., Mahajan, A., Blackburn, B., Kopf, D., Krishnan, L., and Yoong, J.
6 (2014). Micro-loans, insecticide-treated bednets, and malaria: Evidence from
7 a randomized controlled trial in orissa, india. *American Economic Review*,
8 104(7):1909–41.
9
10
11
12

13
14 Teira, D. (2013). Blinding and the non-interference assumption in medical and social
15 trials. *Philosophy of the Social Sciences*, 43(3):358–372.
16
17

18
19 Teira, D. and Reiss, J. (2013). *Causality, Impartiality and Evidence-Based Policy*,
20 pages 207–224. Springer Netherlands, Dordrecht.
21
22

23
24 Viceisza, A. C. G. (2016). Creating a lab in the field: Economics experiments for
25 policymaking. *Journal of Economic Surveys*, 30(5):835–854.
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52